A DISTRIBUTED INTELLIGENT AUTOMATED DEMAND RESPONSE BUILDING MANAGEMENT SYSTEM

Final Report

March 31, 2013

1 DOE Award Number & Name of Recipient

1.1 DOE Award Number:

DE-EE0003847

1.2 Recipient:

Regents of the University of California, (UC Berkeley campus), with Lawrence Berkeley National Lab and Siemens Corp as sub-awardees.

2 Project Title and Principal Investigators

2.1 Project Title:

A DISTRIBUTED INTELLIGENT AUTOMATED DEMAND RESPONSE BUILDING MANAGEMENT SYSTEM

2.2 Principal Investigators:

Professor David Auslander (PI)
Department of Mechanical Engineering, U.C. Berkeley

Professor David Culler (co-PI)
Department of Electrical Engineering, U.C. Berkeley

Professor Paul K. Wright (co-PI)
Director, Center for Information Technology Research in the Interest of Society (CITRIS)
and the Department of Mechanical Engineering, U.C. Berkeley

Dr. Yan Lu (PI for Siemens Corp.)
Siemens Corporate Research Inc

Ms. Mary Ann Piette (PI for LBNL)
Lawrence Berkeley National Laboratory

2.3 Key personnel:

Siemens Corporate Research Inc (SCR): Thomas Gruenewald, Prasad Mukka, Siyuan Zhou, Jianmin Zhu, Ling Shen, and Dr. Mohammad Al Faruque
Lawrence Berkeley National Laboratory (LBNL): Sila Kiliccote
CIEE: Karl Brown, systems consultant; Dr. Therese Peffer, project manager
Graduate Student Researchers, UC Berkeley:
Daniel Arnold, Tyler Jones, Michael Sankur, and Jason Trager, Mechanical Engineering
Stephen Dawson-Haggerty, Andrew Krioukov, Jorge Ortiz, Jay Taneja, Electrical Engineering Computer Science
Graduate Student Researcher, Lawrence Berkeley National Laboratory & Architecture, UC Berkeley:
Rongxin Yin, Joyce Kim
Bin Chen, Architecture
Research staff: Steven Yen, Nathan Murthy

Acknowledgement

This material is based upon work supported by the Department of Energy under Award Number DE-EE0003847

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

3 Contents

3.1 Table of Contents

3.2 Table of Figures

3.3 Acronyms

ADR: Automated Demand Response

BAS: Building Automation System

CLSC: Central Load-Shed Coordinator

CITRIS: Center for Information Technology Research in the Interest of Society

DR: Demand Response

DRAS: Demand Response Automated Server

HVAC: Heating, Ventilation, and Air Conditioning

LBNL: Lawrence Berkeley National Laboratory

OpenADR: open Automated Demand Response

SCR: Siemens Corporate Research

SEB: Smart Energy Box

sMAP: simple Measurement and Actuation Profile

UC: University of California

4 Executive Summary

The goal of the 2.5 year Distributed Intelligent Automated Demand Response (DIADR) project was to reduce peak electricity load of Sutardja Dai Hall at UC Berkeley by 30% while maintaining a healthy, comfortable, and productive environment for the occupants. We sought to bring together both central and distributed control to provide “deep” demand response[1] at the appliance level of the building as well as typical lighting and HVAC applications. This project brought together Siemens Corporate Research and Siemens Building Technology (the building has a Siemens Apogee Building Automation System (BAS)), Lawrence Berkeley National Laboratory (leveraging their Open Automated Demand Response (openADR), Auto-Demand Response, and building modeling expertise), and UC Berkeley (related demand response research including distributed wireless control, and grid-to-building gateway development).

Sutardja Dai Hall houses the Center for Information Technology Research in the Interest of Society (CITRIS), which fosters collaboration among industry and faculty and students of four UC campuses (Berkeley, Davis, Merced, and Santa Cruz). The 141,000 square foot building, occupied in 2009, includes typical office spaces and a nanofabrication laboratory. Heating is provided by a district heating system (steam from campus as a byproduct of the campus cogeneration plant); cooling is provided by one of two chillers: a more typical electric centrifugal compressor chiller designed for the cool months (Nov-March) and a steam absorption chiller for use in the warm months (April-October). Lighting in the open office areas is provided by direct-indirect luminaries with Building Management System-based scheduling for open areas, and occupancy sensors for private office areas. For the purposes of this project, we focused on the office portion of the building. Annual energy consumption is approximately 8053 MWh; the office portion is estimated as 1924 MWh. The maximum peak load during the study period was 1175 kW.

Several new tools facilitated this work, such as the Smart Energy Box, the distributed load controller or Energy Information Gateway, the web-based DR controller (dubbed the Central Load-Shed Coordinator or CLSC), and the Demand Response Capacity Assessment & Operation Assistance Tool (DRCAOT). In addition, an innovative data aggregator called sMAP (simple Measurement and Actuation Profile) allowed data from different sources collected in a compact form and facilitated detailed analysis of the building systems operation. A smart phone application (RAP or Rapid Audit Protocol) facilitated an inventory of the building’s plug loads. Carbon dioxide sensors located in conference rooms and classrooms allowed demand controlled ventilation.

The extensive submetering and nimble access to this data provided great insight into the details of the building operation as well as quick diagnostics and analyses of tests. For example, students discovered a short-cycling chiller, a stuck damper, and a leaking cooling coil in the first field tests. For our final field tests, we were able to see how each zone was affected by the DR strategies (e.g., the offices on the 7th floor grew very warm quickly) and fine-tune the strategies accordingly.

[1] Note that the University does not currently participate in a demand response program, but the Operational Excellence Energy Management program does provide a structure for bringing savings back to the departments.

Most of the reduction in peak demand came from the HVAC system by increasing supply air and zone temperatures, and reducing ventilation rates. Other reduction involved dimming and turning off lights. The web interface for the distributed load controller or gateway allowed the prioritization of curtailable appliances (such as lamps, laptops, fans, heaters, printers) and visualization of energy consumption.

We found challenges in achieving this goal in the office portion of the building, given an overcooled building not fully commissioned, grossly oversized chillers, continually increasing demand, and difficulties isolating the office portion of the building from the nanofabrication laboratory[2]. Nonetheless, we reduced peak electricity by 14-24% and cooling by 25-78 tons using the absorption chiller, the deeper reduction for hotter weather conditions. We estimate that had the centrifugal chiller been used on the hot test day, we would have achieved a 30% reduction from peak electrical load. We also applied many of the demand response strategies (such as reducing minimum ventilation rates and expanding the zone temperature range) towards daily energy efficient strategies for overall energy savings; we estimate that this savings would be $44k annually were it not confounded by increased load in other areas.

5 Introduction

This project brought together Siemens Corporate Research (since the building under study had a Siemens Apogee Building Automation System (BAS)), Lawrence Berkeley National Laboratory (leveraging their OpenADR, Auto-DR, and building modeling expertise), and UC Berkeley (related demand response research including distributed wireless control, and grid to building gateway development).

This 2.5-year project proposed to develop a Distributed Intelligent Automated Demand Response (DIADR) system for Sutardja Dai Hall, a relatively new building at UC Berkeley. This system responds to a demand response signal by reducing the peak load of the office portion of the building by 30% while maintaining a healthy, comfortable, and productive environment for the occupants.

5.1 Background

In the last ten years, much research has focused on fully automated demand response, or Auto-DR, where a building management system automatically initiates energy saving strategies upon receipt of a signal (Piette et al. 2006). The OpenADR communication specification developed by LBNL is now used worldwide, has been adopted by many vendors, used in utility programs, and implemented in hundreds of commercial buildings (Kiliccote et al. 2010). Typical peak reductions are approximately 10%, with some buildings achieving as much as 20% (Piette et al. 2006). The vast majority of these reductions has been in HVAC and lighting systems. However, in recent years, there has been increasing interest in controlling distributed loads, especially in the residential sector (Peffer 2009), and very recently the commercial sector (LeGrand 2012), (Sator 2008).

We endeavored to bring together both central and distributed control to provide "deep" demand response at the appliance level of the building as well as typical lighting and HVAC applications. In doing

[2] Many key loads—such as the chillers and associated chilled water and condenser water pumps—were shared between the nano fab lab and the office portion of the building.

so, we strove to achieve greater savings—30%—than is typically achieved in demand response programs.

5.2 Infrastructure

Sutardja Dai Hall houses the Center for Information Technology Research in the Interest of Society (CITRIS), which includes four of the UC campuses (Berkeley, Davis, Merced, and Santa Cruz). The 141,000 square foot building, occupied in 2009, includes offices, a few classrooms, café, and a nanofabrication laboratory. For the purposes of this project, we did not include the loads of the nanofabrication laboratory (about 15,000 square feet). The building has two separate chillers: an absorption chiller meant to use steam from the campus during the summer months and a typical compressor chiller. DR strategies were considered for both chillers. Even though demand response events typically occur in the summer months in California when air conditioning contributes to peak loads, we wanted to consider strategies for the compressor chiller since that is the more common chiller found in office buildings.

The main architecture of the system included the building (Sutardja Dai Hall), a Demand Response Automated Server, a central controller (Smart Energy Box or web-based control), and distributed load control gateways with a user interface.

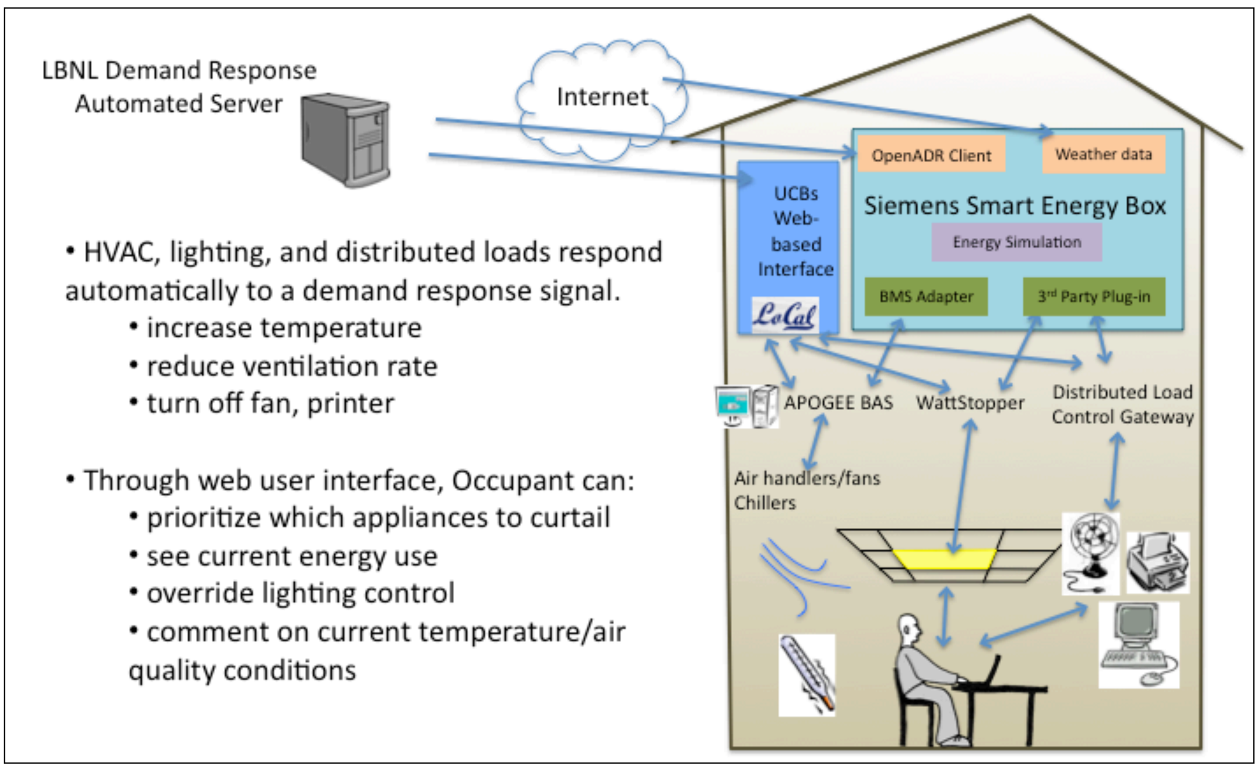

Figure 1: Schematic of system architecture showing parallel controller: the Smart Energy Box and sMAP-based controller.

We ended up exploring two parallel means for interfacing with the controls: the Smart Energy Box (Siemens) and an open web-based interface that directly acted upon the BACnet points through a secure server with its own supervisory control (UCB). A distributed load control gateway was developed for the project, based on the residential energy gateway developed at UC Berkeley (see (Arnold 2011). These gateways use Raritan metered and switched plugstrips as well as wireless ACme receptacle meter and relays previously developed at UC Berkeley. The gateway also had a user interface through which a building occupant could add appliances, prioritize which appliance to curtail and override appliance and lighting curtailment during events.

Simulation models developed included an extensive EnergyPlus model and predictive load modeling for developing a baseline.

A previous project had added 27 revenue grade DEM 2000 power submeters to the building; this involved most subpanels including individual subpanels on each floor for lighting and receptacle power. However, in order to isolate the loads of the nanofabrication laboratory, a couple of electrical subpanels required submetering and flowmeters were required to determine the portion of chilled water for the office versus the lab. An audit of the plugloads of the building was conducted using an innovative smart phone application and StreamFS (data management software) developed for this project.

UC Berkeley, LBNL, and Siemens Corporate Research all developed demand response algorithms for the HVAC system. SCR included a thorough study of the weather patterns of Berkeley, and an adaptive DR strategy tool that would automatically select the best strategy. Basic scenarios for the air-side of the HVAC system included: increase the supply air temperature, increase the zone temperatures (Global air temperature), reduce ventilation rate, and reduce static pressure. For the lighting system, we decided to try a couple of strategies; reduce the lighting by one step or reduce the lighting to 33% and allow overrides. A web-based Personalized Lighting Control system provided a simple interface to occupants in open plan offices to control the lighting for their zone and control the level of lighting.

The test plan consisted of first isolating the office load from the rest of the building. Then we identify peak loads, and the components of these loads for both warm and hot days, and with either chiller. While not required by the project, we endeavored to improve the energy efficiency of the building before we started our DR strategies.

For testing, diagnostics, and analyzing results, we used the simple measurement and actuation profile (sMAP) developed on campus as a data source-agnostic aggregator (Dawson-Haggerty, Krioukov, and Culler 2011).

The next section (6) of the report describes task-by-task project accomplishments in comparison with official project goals. Section 7 outlines the major accomplishments and significant results of the project.

6 Comparison of Project Accomplishments and Project Goals

The DIADR project consisted of two main phases: Research and Development, and Implementation and Validation. The following section describes the goals of each task for each phase and how the actual accomplishment compared to those goals.

6.1 Research and Development

6.1.1 Task 2: System architecture

Task 2 required the outlining of functional requirements and the development of the system architecture.

The building selected was Sutardja Dai Hall, also known as the CITRIS building, at the UC Berkeley campus in Berkeley, CA (described further in Appendix A and B). The building management system for the heating, ventilation, and air conditioning system was a Siemens Apogee Building Automation System.

To provide the Distributed Intelligent Automated Demand Response (DIADR) system, we created two new parallel controllers with an external Automated Demand Response (ADR) Server and interface with the existing building controller or Building Automation System (BAS) for the building HVAC systems. This interface communicated with the building lighting system via the existing WattStopper control system and distributed loads through a gateway. For this interface we began using Siemens' Smart Energy Box (outlined later in this document), but by the end of the project UC Berkeley had developed a second parallel DR controller.

The Siemens Smart Energy Box includes an external connection with an OpenADR client and access to weather data. A built-in energy simulation allows for optimization of DR strategies. The BMS adapter provides communication to legacy building equipment and can be customized. Third party plug-in allows addition of new algorithms and custom interfaces (such as communication with gateway).

For initial field testing, UC Berkeley graduate students developed a simple Measurement and Actuation Profile (sMAP)-based controller (Appendix C); other students added DR algorithms (Appendix D) and an optimizer (Appendix E) to create a parallel DR control architecture called the Central Load-Shed Coordinator (CLSC). sMAP[3] (**http://www.cs.berkeley.edu/~stevedh/smap2/intro.html**) allows access to building information through Internet web protocols; time-based data are streamed to a database with a web portal for easy non-proprietary access. Thus the DR controller needs no special interfaces and can be customized for any building.

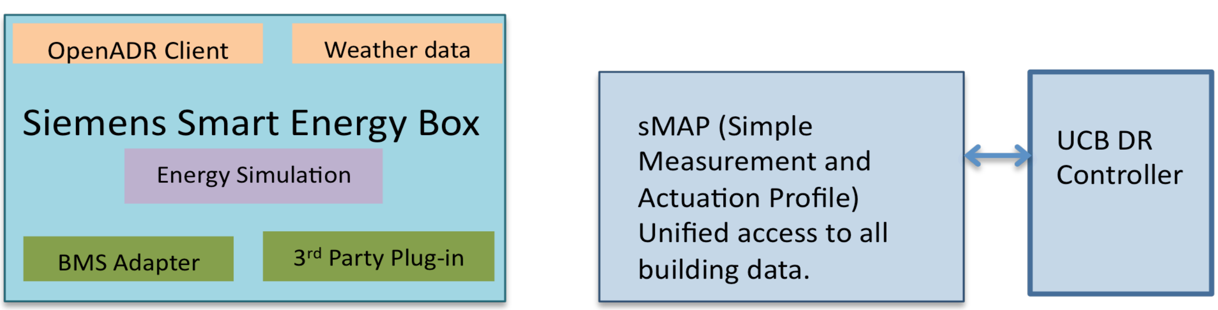

Figure 2: Parallel DIADR controllers.

[3] sMAP development was supported under two NSF grants (LoCal and ActionWebs).

6.1.1.1 Siemens Smart Energy Box

Smart Energy Box (SEB) is an existing integration platform developed by Siemens which allows communication among building subsystems, building occupants and building operators in order to form coordinated load management and demand control. For this project, SCR adopted the existing SEB architecture and customized it as the central DR controller for the DIADR system.

The core of SEB are a runtime system and several basic function components developed earlier through Siemens internal funded effort in order to perform HVAC demand response control, including a BACnet adaptor, an OpenADR client and a weather service adaptor. Within the scope of DIADR project, Siemens has extended SEB's capabilities to bring lighting and plug load control into the load shaping programs. Additional SEB function components were developed for this purpose, including an energy simulation engine based on EnergyPlus, an adaptive and intelligent DR manager and a distributed DR agent integration interface based on JADE (Java-based Agent Development Environment). To make the DR controller more user-friendly, SCR also developed a windows-based user interface to configure and monitor SEB. This interface allows user configure Building Information such as Building energy model, control points, control point mapping to energy model, strategies, monitor DR event progress, monitor strategy execution and cancel/suspend/resume scheduled strategies. The details of SEB design and adaptive demand response control strategy can be found later in this document.

The extended SEB hosted in Siemens SIMATIC Industrial PC has been installed in the CITRIS Building. Several tests have been conducted and the function of SEB as a central DR controller was validated.

6.1.2 Task 3: Develop a Service-Oriented Architecture

Task 3 required the development of a service-oriented architecture. Since UC Berkeley had already been developing a residential energy gateway (Dan Arnold, http://mechatronics.berkeley.edu/gateway.htm), we decided to build upon this architecture for the purposes of this project. Michael Sankur of UC Berkeley worked with Dan to develop the commercial version of this gateway, referred to as the Energy Information Gateway or EIG; researchers at SCR in turn worked to improve the performance of the gateway.

6.1.2.1 UCB Energy Information Gateway

The description of the UC Berkeley gateway may be found in Appendix F.

6.1.2.2 Siemens' amendments/improvements

Siemens has extended the capabilities of UCB Gateway software in the following aspects.

- A multi-agent platform is enabled by deploying Java Agent Development (JADE) framework. Each Gateway acts as an agent compliant with Foundation for Intelligent Physical Agents (FIPA) (http://www.fipa.org/), and can communicate with SEB with Agent Communication Language (ACL) messages (http://www.fipa.org/repository/aclspecs.html).
- A sequence of behaviors is implemented in Gateway to execute complete Adaptive Demand Response Strategy with SEB. As a result, when a DR event is scheduled, both the centralized loads, (HVAC and Central lighting), and the Distributed loads contribute together towards the DR power reduction goal. The details of Adaptive Demand Response Strategy can be found in Section 6.3.2.
- An optimization algorithm based on Particle Swarm Optimization (PSO) is implemented in the Gateway core. User preferences, real-time power consumptions and environment information are the key factors to this optimization to form appliance-level utility functions. Therefore, the power reduction can be reached in a smart, considerate and optimal way.

- SCR developed several drivers for Gateway, including a printer driver to control printers over USB cable or on the network, a laptop battery driver to monitor laptop battery status, and a more resilient Raritan driver for fast on/off commanding and power consumption monitoring.
- Integration of Zigbee Sensors was developed and implemented in Gateway, including an occupancy sensor, a light sensor, a temperature sensor, a humidity sensor. The Gateway can use the sensors to collect important environment information for appliance control.
- A Web-based User Interface (Web UI), based on JavaScript and AJAX, is provided for users to configure Gateway and monitor appliances. With a successful login to the Web UI, User can monitor power usage of each appliance, add/delete/edit DR preferences on each appliance, monitor DR event and environment information from the sensor box.
- An XML-based configuration system is added to provide persistency to the Gateway. No matter if Gateway runtime is stopped normally or unexpectedly, the user login information, Gateway and appliances configuration can be restored with the XML configuration.
- The communication between SEB and Gateways is facilitated by a ManagerWrapper component, which manages the translation and message forwarding between SEB and Gateways.

More details about Siemens Gateway features and architecture can be found in the section on Gateway Architecture, 6.1.2.6.

6.1.3 Task 4: Open ADR integration

The goal of Task 4 was to integrate OpenADR into the DR controller. SCR configured Smart Energy Box's Open ADR Client module to enable the communication between CITRIS DIADR system and Demand Response Sever (DRAS) hosted initially by Akuakom and later by AutoGrid. The communication is compliant with Open ADR protocol version 1.0 and was fully tested.

6.1.4 Task 5: DR algorithm development

The goal of Task 5 was to develop demand response algorithms for the heating, ventilation, and air conditioning equipment, lighting, and plugloads. LBNL took on the additional task of developing an extensive EnergyPlus model; the details of this are described in Appendix G. In order to develop strategies for the plugloads, we needed to understand what exactly was in the building. Jason Trager and Jorge Ortiz developed the Rapid Audit Protocol using a smart phone app to log and categorize all appliances in the office portion of the building (Appendix H).

Typical demand response strategies often include open loop strategies as well as closed loop strategies. For example, to reduce load, one could shut down the air handling units directly—an open loop strategy—or reduce ventilation to the zones—a closed loop strategy. Because we were interested in maintaining control for the comfort and productivity of the occupants, we developed and tested only closed loop DR strategies.

As we developed demand response strategies, we were also interested in improving the energy efficiency of the building's everyday operational performance. As an example, see Table 1 below.

Table 1: How Building automation systems support both energy efficiency and demand response (Goldman et al. 2010).

Function: Energy Efficiency	Function: Energy Efficiency and Demand Response	Function: Demand Response
Lighting Control Features of BAS		
▪ Centralized on/off controls, timers	▪ Central dimming ▪ Bi-level/zonal switching	▪ Demand limiting ▪ Lighting sweep ▪ Overrides
HVAC Control Features of BAS		
▪ Optimal start ▪ Variable speed drive control ▪ Demand-controlled ventilation ▪ Chilled water temperature control ▪ Condensing temperature control ▪ Cooling tower/evaporative condenser fan control	▪ Global zone reset ▪ Duct static pressure reduction	▪ Equipment lockout ▪ Pre-cooling ▪ Thermal energy storage ▪ Cooling reduction ▪ Fan, pump, or chiller quantity reduction

Source: Based on Kiliccote and Piette, 2005.

While we developed and simulated strategies for the water- and air-sides of the HVAC system, it was only feasible to field test the air side strategies. Late in the project, we developed demand control ventilation strategies using carbon dioxide sensors (see Appendix Q). We developed and successfully tested lighting strategies. We also developed plug-load strategies, but only tested these in a lab environment (e.g., one office). The simulation of laptop battery optimization turned into a paper, available in Appendix O. We began to develop a multi-gateway simulation to understand the effects of using plug load control throughout the building, but due to personnel changes late in the project, did not complete this.

6.1.4.1 Load baseline development

In order to compute peak load savings, we needed to develop a baseline model to predict the demand load. The initial baseline work was conducted by Tyler Jones and summarized in (Jones and Auslander, 2010), found on the project website. Jason Trager took over the baseline work when Tyler left. Ultimately, however, there were too many changes to the building performance throughout the project to allow proper training of the predictor tools; we ended up using day before data for each demand response event day as the baseline data.

6.1.4.2 UCB DR Controller Optimizer

Jason Trager developed an optimizer for the DR controller. When scheduling a demand response event, we want to assure that a certain amount of power will be reduced during every time point in the day. We designed an autonomous arrangement of the order of strategies for saving power during a DR event. In order to do this, we classify all loads for DR in terms of how much power they save over time, and how "inconvenient" they are for the occupants during that time. Details of the optimizer may be found in Appendix E.

6.1.4.3 Demand Response Capacity Assessment & Operation Assistance Tool (DRCAOT)

SCR developed a Demand Response Capacity Assessment & Operation Assistance Tool (DRCAOT) during the project execution to better assess the load reduction capacity of the CITRIS building in response to the demand response requests from the utility provider, with various levels of awareness of the weather condition at the planned demand response day. In addition, DRCAOT also allows user to design control strategies for the building HVAC system, and is able to identify the optimal strategy based on simulation evaluation though a strategy library. The strategy library is a collection of possible actions from which to use.

6.1.5 Task 6: Local control testing in lab

The goal of Task 6 was to run the DIADR controller in a small-scale lab environment.

6.1.5.1 UCB's test lab

Initially, the test lab was a private office, room 464, in Sutardja Dai Hall; for the last six months of the project, this lab was moved to room 456. The equipment installed includes desktop and laptop computers, a small refrigerator, a laser printer, small fan, and small heater, along with UPS (Uninterruptible Power Supply) devices. In addition, the lab has a combination of ACme plug-load electricity meters (developed at UC Berkeley) and Raritan Dominion PX8 metered and switched power strips for monitoring the power consumption of appliances. This instrumentation reports its data to sMAP, a physical data store. Room 456 had a lighting upgrade during the project: the ballasts for the overhead lighting were replaced with continuous dimming ballasts, and the lighting circuits connected to an Ethernet network; however, we were unable to implement this into the DR controller. We did however successfully control the other appliances with the DR controller and gateway.

6.1.5.2 Siemens' test lab

Before the onsite deployment, SCR performed intensive lab testing of DIADR system, including both SEB and Distributed Gateway. The picture below shows the lab setup in SCR's Princeton office. The table lists all the equipment for the lab testing.

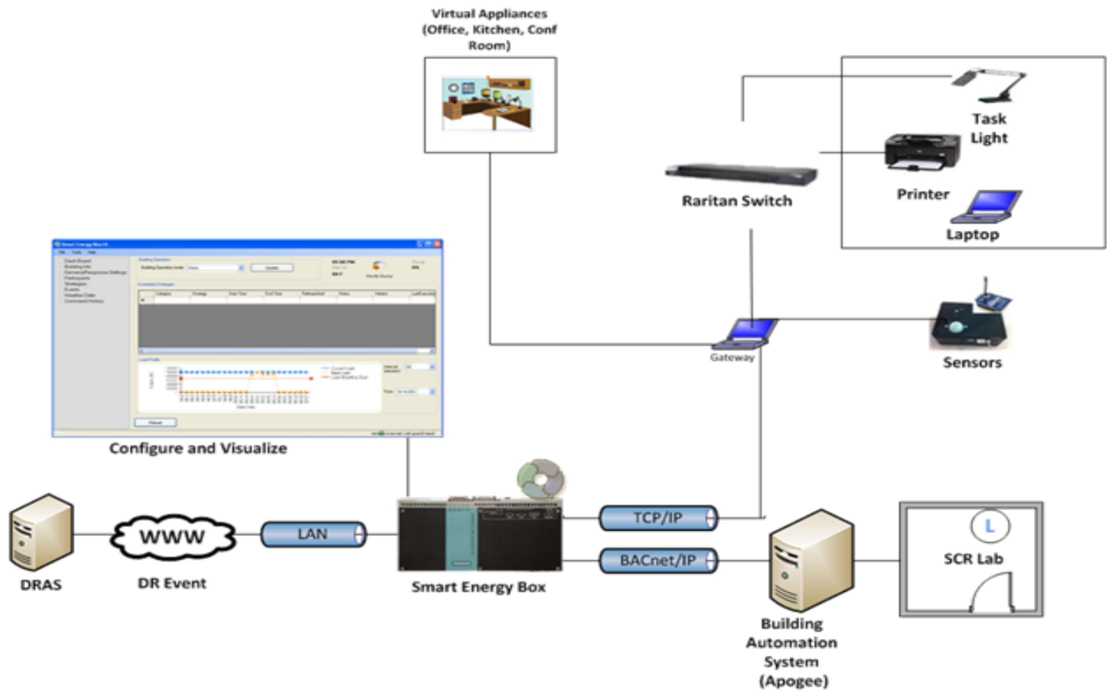

Figure 3 DIADR setup at Siemens lab.

The following table shows list of lab equipment and the purpose of each.

Equipment	Purpose
Siemens Apogee Building management system	Similar to CITRIS building management system
SEB Machine	Runs the Smart Energy Box (SEB) UI, the ManagerWrapper and the SEB runtime
DRAS Machine	Demand Response Automated Server (DRAS) that hosts in-house developed demand response server based on Open ADR protocol version 1.0
Laptop	The laptop is used to run multiple instances of Gateways. It uses Ethernet connection to negotiate with the SEB machine and allows the Gateways to receive instructions from the SEB. It also connects with the Raritan and Printer through Ethernet, and the Sensor box through Zigbee wireless communication.
Raritan	The Raritan is an Ethernet device which allows the Gateways to off/on and monitor energy consumption of plugged in appliances. In our lab, the Gateway can monitor and manage the task light and laptop power supply through this device.
Xbee Chip/Sensor box	The sensor box monitors the light, humidity, temperature and occupancy of a room. Gateway accesses this information over Zigbee.
Printer	The Printer is connected to the laptop hosting Gateways through Ethernet and is configured as a network printer on the laptop. Gateway can stop/resume jobs on the printer to contribute to load shedding targets requested from SEB.

6.2 Implementation and Validation

6.2.1 Task 8: Integration

The goal of Task 8 was to integrate the DIADR controller with the various systems of the building (e.g., HVAC, lighting, and plugloads). The Smart Energy Box provided integration to the HVAC system through the Apogee system, lighting through the WattStopper-BACnet interface, and the plugloads through the gateway. As mentioned in Task 2, the sMAP interface provided integration for the UCB DR controller to communicate with the HVAC system and lighting system; the gateway provided the integration from the DR controller and the plugloads.

Section 6.1.1.1 describes the integration of the DR controller with the HVAC system and lighting system via sMAP, led by Andrew Krioukov. The Building Management System of Sutardja Dai Hall is a Siemens Apogee system, which is accessible to sMAP over a BACnet/IP address. The WattStopper lighting control system required some integration work by WattStopper, but eventually was also made accessible in a similar way. This integration allowed both control of systems and visualization and monitoring of the data.

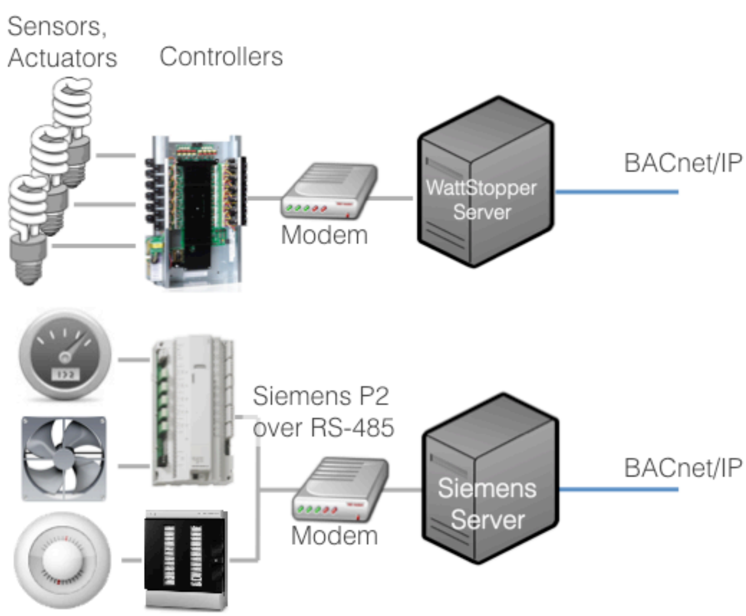

Figure 4: Lighting and HVAC control integration through sMAP web interface (Krioukov).

Andrew Krioukov developed a web-based interface to enable the field testing. The controllers send updates to the BMS server; the sMAP Gateway polls data from BMS.

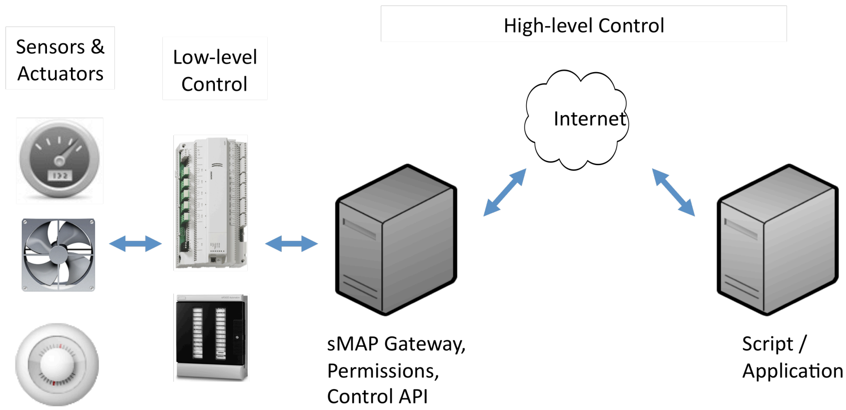

Figure 5: The field tests required entering the prepared scripts via the Internet to the sMAP gateway to control the HVAC and lighting systems (Krioukov).

6.2.2 Task 9: Enhanced Scale Testing

The goal of task 9 was to conduct testing of the DIADR system at the building scale. Both the UC Berkeley/LBNL team and the Siemens Corporate Research team conducted full-scale tests in the office portion of the building.

6.2.2.1 UC Berkeley and LBNL field tests

The UC team began testing the HVAC system in September 2011 by raising the zone temperature setpoints and increasing the supply air temperature. We discovered that the centrifugal chiller was short cycling, an economizer damper was stuck, and the cooling coil valve for one of the air handling units for the office was leaking. Ventilation tests continued in November 2011; the minimum ventilation rate was reduced by 30%, 50%, and 70%, and carbon dioxide sensors were deployed to ascertain the effect on occupants. Even a large reduction in ventilation did not increase CO2 levels above 700 ppm. However, the secure doors to the nanofabrication lab on the fifth floor were not closing properly due to the difference in pressure. In late January and February 2012, Andrew Krioukov continued the ventilation tests, turning off systems at night, and fine-tuning the ventilation so that the doors were closing properly.

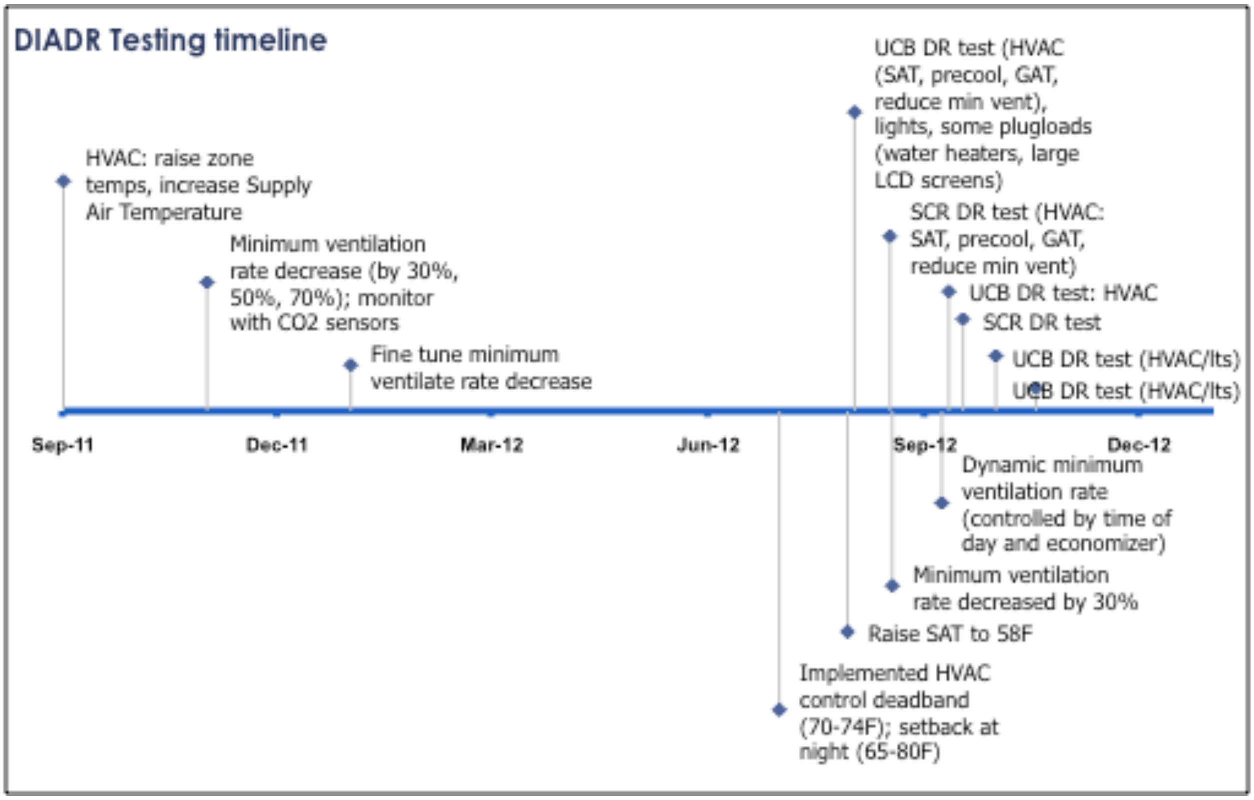

Figure 6: Testing events in Sutardja Dai Hall.

To understand the effect of the various tests on occupants, we worked with the Center for the Built Environment at UC Berkeley to administer a series of occupant surveys. A long survey on detailed aspects of the indoor environment was administered in August 2011; a series of short surveys to determine the immediate "right now" effects of tests on thermal comfort, lighting, and air quality were conducted during the tests in September and November 2011, and for the final DR event on Oct 17-18. We determined that in general the building was overcooled and overventilated.

In July 2012, we began to implement a series of energy efficient measures to bring the building to a more efficient benchmark; these included developing a deadband for zone temperatures (from the single 70F heating/cooling setpoint to heating below 70F and cooling above 74F), raising the supply air temperature from 56F to 58F, and reducing the minimum ventilation rate to 70% of original values.

In August 2012, we began the first of several integrated test DR events, curtailing both the HVAC and lighting loads, and asking for occupant participation in reducing plugloads; these tests took place August 2 (HVAC, lights, and plugloads), Sept 11 (HVAC only test with DR controller), Oct 1 (HVAC and lights controlled with DR controller), and Oct 18, 2012 (HVAC and lights controlled with DR controller, Occupant survey). The next section has the results of these tests.

6.2.2.2 Siemens Field Tests

SCR conducted two field tests to validate SEB implementing the central load control optimization algorithm (the Optimal Strategy Pool algorithm (OSP)), on August 17 and September 17 respectively. We successfully finished optimal strategy selection, issued control directives from SEB to the BACnet server, and executed the directives of controlling the global temperature setpoint, supply air temperature setpoint and minimum air volume set points. According to the recorded data by sMAP system, the building responded to our strategies as expected. The test design and result will be detailed in the later section.

6.2.3 Task 10: Commercialization plan

The goal of Task 10 was to outline the steps of commercialization any feasible product of the DIADR project. Based on Siemens work under this project, three potential products were identified by Siemens for commercialization, including Smart Energy Box (SEB), Distributed Gateway, and Demand Response Capacity Assessment & Operation Assistance Tool (DRCAOT). SBT developed commercialization plans; the Task 10 commercialization plan report may be found on the project website and in the attached Appendix I.

6.2.4 Task 11: Demonstration (September 18, 2012)

On September 18, 2012, UC Berkeley, LBNL and Siemens team members made an oral presentation of the DIADR project to Department of Energy officials Alan Schroeder and George Hernandez, and others at the all day i4Energy research symposium. In the morning session, Dan Arnold described the gateway. At the lunch break, we provided a brief live demonstration of the gateway optimally controlling distributed plugloads and limited lighting control in the laboratory, room 456. Jason Trager (UCB) described and demonstrated the energy audit and database association with an Android smart phone application, Prasad Mukka and Siyuan Zhou (SCR) demonstrated the DRAS-SEB-gateway, and Michael Sankur (UCB) presented the UCB DR controller and HVAC/lighting and typical office plugloads.

In the auditorium downstairs, David Auslander (UCB) began the three-hour presentation by providing an introduction and scope of project. Mary Ann Piette (LBNL) described OpenADR. Prasad Mukka outlined the functions of the Smart Energy Box. This included communication between the SEB and DR Server, Central Load Control from SEB, Central Lighting Control from distributed gateway, and Distributed load Control Gateway. Michael Sankur and Siyuan Zhou described the gateway. Michael then described the DR controller. Jason Trager talked about the Building-wide Energy Audit & Association. David Culler (UCB) described the work of Stephen Dawson-Haggerty and Andrew Krioukov by introducing sMAP and presenting the control integration. After a brief break, Tyler Jones (UCB) outlined the process of establishing a baseline. Then Rongxin Yin (LBNL) described developing the EnergyPlus model as well as DR algorithm development and simulations. Jason Trager then discussed the UCB DR controller optimization followed by Prasad Mukka and Siyuan Zhou describing the SEB and gateway optimization. Then Thomas Gruenewald (SCR) outlined the DR Capacity Assessment & Operation Assistance Tool. In the final section, Prasad presented SCR's test results and Rongxin presented those of UCB. Pornsak Songkakul of SBT presented the Commercialization plan and cost analysis. David Culler provided the wrap-up discussion.

All powerpoint presentations may be found at **http://i4energy.org/index.php/projects/affiliate-projects/6-sutardja-dai-hall**; the youTube videos of the presentations may be found at: http://www.youtube.com/playlist?list=PLYTiwx6hV33t-y_gt8IZKiHm0W8WYffG_&.

7 Major Activities, Significant Results, Major Findings

7.1 Architecture of System

7.1.1 Overall scheme: parallel approach

Sutardja Dai Hall is the new headquarters for the Center for Information Technology Research in the Interest of Society (CITRIS), which fosters collaboration among industry and faculty and students of four UC campuses. The 141,000 square foot building, occupied in 2009, includes both private and open plan office space, a few classrooms, light laboratories, café, auditorium, data center, and a nanofabrication laboratory; it also houses the Main Distribution Center (MDC) for the northeast quadrant of campus. For the purposes of this project, we did not include the loads of the energy-intensive nanofabrication laboratory, the data center, nor the MDC.

The building has a fuel-flex system, which allows either the 600 ton Trane centrifugal compressor chiller or the 600 ton absorption chiller (powered with campus steam) to provide chilled water; these are controlled through the Siemens Apogee Building Automation System. The absorption chiller was designed to use steam from April through October when steam on the UC Berkeley campus is not in high demand for heating. A centrifugal compressor chiller with hot gas bypass was designed to run from November through March. Steam also provides reheat for 130 variable air volume boxes throughout the office portion of the building. DR strategies were considered for both chillers. Even though demand response events typically occur in the summer months in California when air conditioning contributes to peak loads, we wanted to consider strategies for the centrifugal chiller since that is a more typical chiller. More detail about the building and monitoring systems may be found in Appendices A and B.

The building has a Siemens Apogee Building Automation System (BAS). The WattStopper lighting system in the open plan offices (found on floors 4-7) has tri-level stepped dimming capability (three 25w lamps with two ballasts) and is on a timed schedule. The private offices have Lutron wall switches with stepped dimming and an occupancy sensor; no central control of these lights is available.

We explored two parallel means for interfacing with the controls: the Smart Energy Box (SEB) and an open web-based interface that directly acted upon the BACnet points through a secure server with its own supervisory control and optimizer. Siemens developed SEB with the intent of commercialization. The web-based system was initially used for control during field testing. Following the spirit of open-source horizontal integration espoused by Professor David Culler and the LoCal group at UC Berkeley, it used the Simple Measurement and Actuation Profile (sMAP) developed on campus as a data source-agnostic aggregator (Dawson-Haggerty, Krioukov, and Culler 2011). The plan was to use and compare both systems.

As Sutardja Dai Hall was intended to be a "living laboratory", CITRIS invested over $200k for submetering and infrastructure. This includes 27 revenue grade DEM 2000 power submeters on most subpanels (including a submeter each for lighting and receptacle power on each floor) for a total of 7000 sensor points in the BAS. However, in order to isolate the loads of the office from the nanofabrication laboratory, two additional electrical subpanels required submetering (installed Sept 2012); in addition, flowmeters were required to determine the portion of chilled water for the office versus the lab (installed Oct 2012).

7.1.1.1 UCB sMAP control interface

UC Berkeley made available the data from the Apogee system through a Simple Measurement and Actuation Profile (sMAP) which provides a RESTful web services integration of the submetering, as well as the California Independent System Operator (CAISO) feeds, weather feeds, and many other physical information sources that are important for developing the modeling and intelligent control needed to energy efficient, grid-responsive operation. The LoCal team designed and implemented a data acquisition architecture and wrote a BACnet-to-sMAP converter that makes 400+ HVAC-related streams available.[4] This has enabled all of the BACnet data points from the building to be monitored continuously and made accessible via an open interface. StreamFS (**http://www.eecs.berkeley.edu/~jortiz/**) is the data collection system, metadata manager, and data stream processing system (Ortiz 2012); the BACnet-to-sMAP interface is integrated into the data collection system. This data availability supports the development of algorithms for curtailing load in DR events, the generation of better baseline data, and even a richer understanding of the building processes in their non-DR state. The WattStopper lighting control system was integrated with the Siemens Apogee system through a BACnet interface. The LoCal graduate students enabled all the BACnet data points for the WattStopper (lighting-related) data. The pieces of the sMAP system are designed to separate concerns and allow users to, for instance, run their own web front-end while using hosted infrastructure for storing the actual data and metadata. See Appendix C for more details.

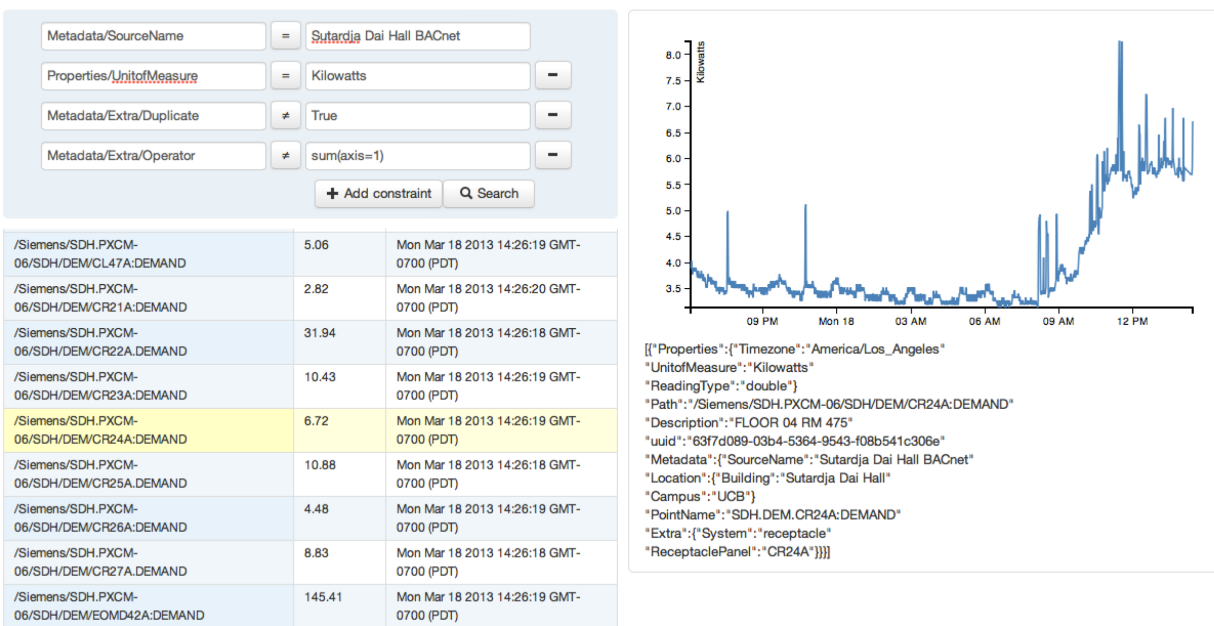

Figure 7: New sMAP interface developed for the project (http://new.openbms.org/status/).

7.1.1.2 UCB's DR controller: the Central Load-Shed Coordinator

The Central Load-Shed Coordinator (CLSC) is designed as a program that facilitates coordinated control over building lighting, HVAC, and plug loads. As shown in Figure 6, the CLSC has the ability to poll DR servers for event information, and can display pertinent building information to a building manager via a user interface. The CLSC is also an energy related information aggregator, quite similar to an Energy

[4] Extensive studies were performed on utilizing conventional BACnet building management systems for deep analytics. When businesses moved to a 24x7 operational model, we found that care must be exercised not to over-tax these legacy systems. For example, the Siemens Apogee Insight system in Sutardja Dai Hall actually comprises over 6,000 points, but sampling too many of them introduces failures.

Information Gateway (EIG). An EIG communicates with, and gathers data from, connected plug loads in its domain, whereas the CLSC gathers data on lighting and HVAC systems, as well as pertinent plug load information from EIGs. The CLSC communicates with, and controls, lighting and HVAC systems. The CLSC also communicates with EIGs to allow the participation of their plug loads in the load-shed process.

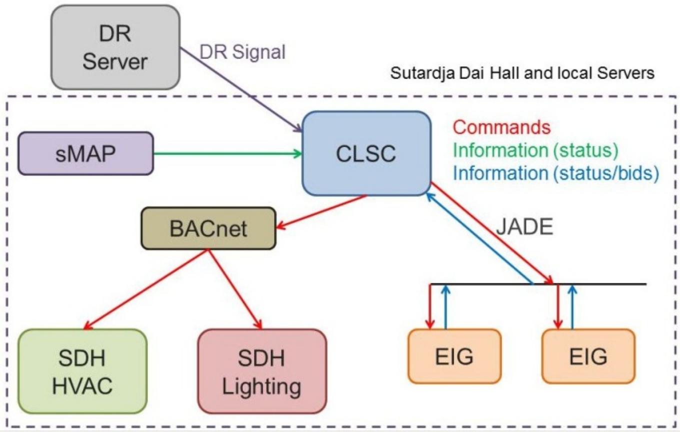

Figure 8: Deployment of CLSC within SDH as per DIADR project task

In order to generalize actuation resources across the three systems, a measure of power and occupant inconvenience is needed so actuation of different resources can be compared. Details of this controller are discussed in Appendix D.

7.1.2 Siemens Smart Energy Box architecture

7.1.2.1 Smart Energy Box

Smart Energy Box (SEB) is an integration software platform developed for advanced energy management and comfort control for commercial buildings. The enhanced SEB through the effort under this project intelligently integrates with existing building management systems to provide optimal demand response results by using real-time weather forecast and energy simulation as well as runtime adaptive negotiations with distributed load controllers for load reduction. SEB allows controlling both central load and distributed plug load seamlessly. SEB adapted market based adaptive approach to negotiate and achieve load shedding targets either set by utility provider or set by the building operator for certain period of time with minimal or no comfort loss for the occupants.

Smart Energy Box has been designed as component based system, where each component is responsible for one or more responsibilities. The components can be plug-and-play; it has many components as explained in the following sections. The following picture shows the overall system architecture with SEB acting as central DR controller and interacting with Demand Response Server, Distributed Control Coordinator, Distributed Control gateway, and user interface.

The details of the runtime environment may be found in Appendix J.

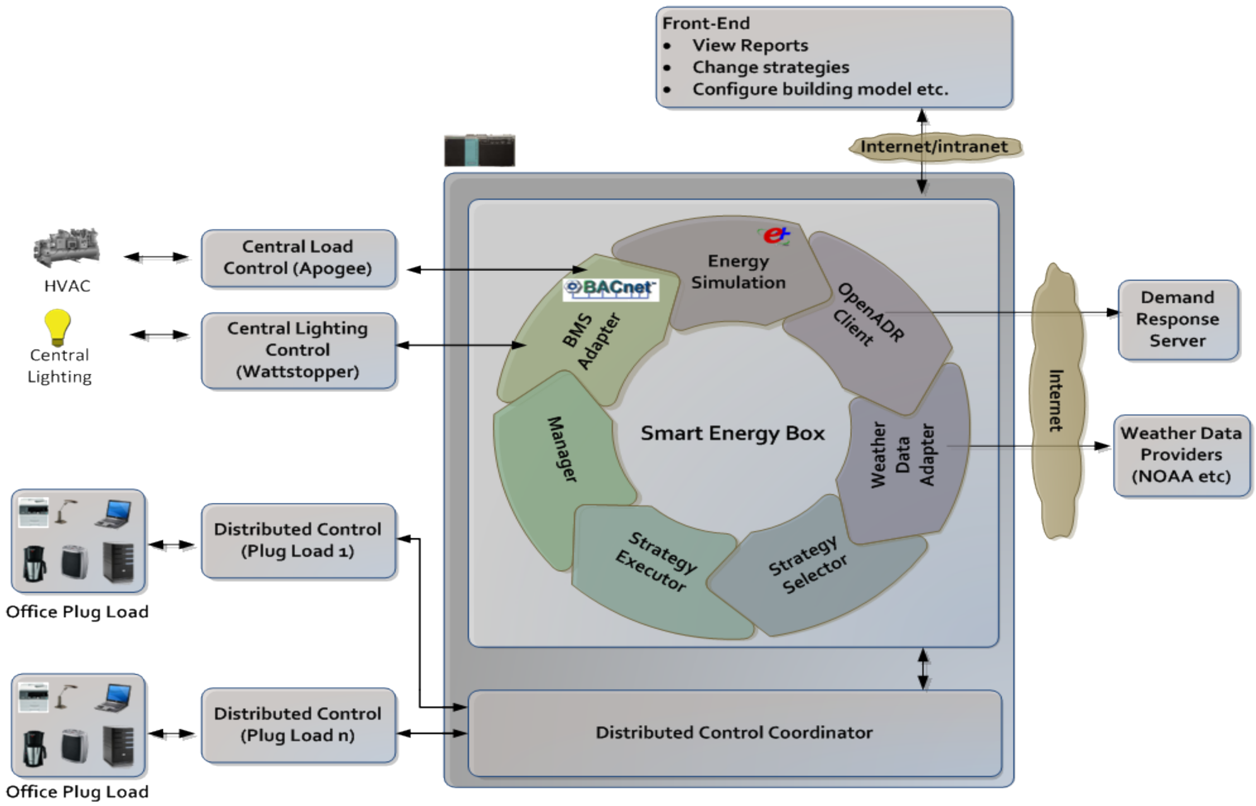

Figure 9: Smart Energy Box System Architecture

7.1.2.1.1 Central HVAC Control

Smart Energy Box has built-in central Control manager. This manager evaluates all available strategies with current weather and brings the most optimal one for the defined demand period. There are 13 optimal strategies defined for Berkeley's weather.

The CITRIS Building HVAC is being controlled by Siemens Apogee Building management system. SEB uses its BACnet interface to perform communication with the building equipment. During demand control, SEB takes the control of allowed building equipment and relinquishes at the end of the demand period to its normal state.

7.1.2.1.2 Central Lighting Control

Central Lighting Manager (CLM) was introduced as runtime component for making Lighting control to be participated in the market based negotiation during demand response event. The office portion of CITRIS building from Levels 2-7 participated in the demand response. For each level multiple strategies were designed based on Siemens adaptive approach. CLM provides a means for all levels to participate and negotiate with SEB Central manager to perform load shedding based on the adaptive market based approach.

The central lighting system was controlled with relays from WattStopper. Each relay control is assigned with one unique BACnet ID so that it can be commanded from SEB. A relay is responsible to control a zone of lighting on the floor, by lamp ballast. There were multiple relays configured to control certain area of lighting. For example, if two relays controlling an area of lighting, then the lights can be controlled at 4 levels through the combinations of the two relay operation ([off, off], [off, on] (one lamp ballast—low light), [on, off] (two lamp ballast—med light), [on, on] (all three lamps—high level of light)). Turning On/off the relay turns on/off lighting it controls.

We have adapted Siemens adaptive agent based negotiation approach to control lighting. So we designed one Agent per floor of lighting for floors 2 to 7 of CITRIS building. Based on the relays available, we have designed multiple lighting strategies per floor. A lighting strategy is defined as a set of status of all relays in this floor with a cost function and a load it can shed.

During demand reduction period each agent proposes its cost function to negotiate with SEB. A cost function is defined with multiple points, where each point represents the relationship between its load reduction and the cost when a specific strategy is chosen. For example, if strategy 1 for floor 1 defines the status of all relays to be “On”, then power reduction is 0, and the cost of reduction is also 0. The number of points in a cost function equals to the number of strategies for the agent. To form a cost function, we need to consider the power reduction when a relay is turned “Off” while all others are “On”, which is equivalent to the power consumption when each relay is turned “On” by its own. The power consumption information can come from appliance energy survey.

The previously conducted appliance energy survey on CITRIS building has power consumption information for each Energy Plus zone. However, the Energy Plus zones do not have one-to-one relationship with the relays. In order to get power consumption of each relay, a mapping between the Energy Plus zones and the relays was calculated. The area mapping and the cost functions for each floor were defined in Lighting_Strategy_Points.xlsx, which is attached.

7.1.2.1.3 User Interface

A Windows based user interface was designed and developed to configure and monitor Smart Energy Box by the facility manager. This user interface accesses SEB through web service hosted by SEB runtime. It has the following major features

- Dash board that shows comprehensive view of the system
- Monitor and update Building information
- Configure Demand response related information
- Monitor Weather data, Energy Simulation results
- Create Energy Strategies and administer them
- Import/Export BACnet configuration
- Monitor Participants, building load shedding goal

7.1.2.1.3.1 Dashboard

User interface was enhanced by introducing a dashboard, where a facility manager can get a glance of whole system. And also he can see the strategies that were selected by the system and override for any reason. That means selected strategies can be cancelled, suspended, or resumed.

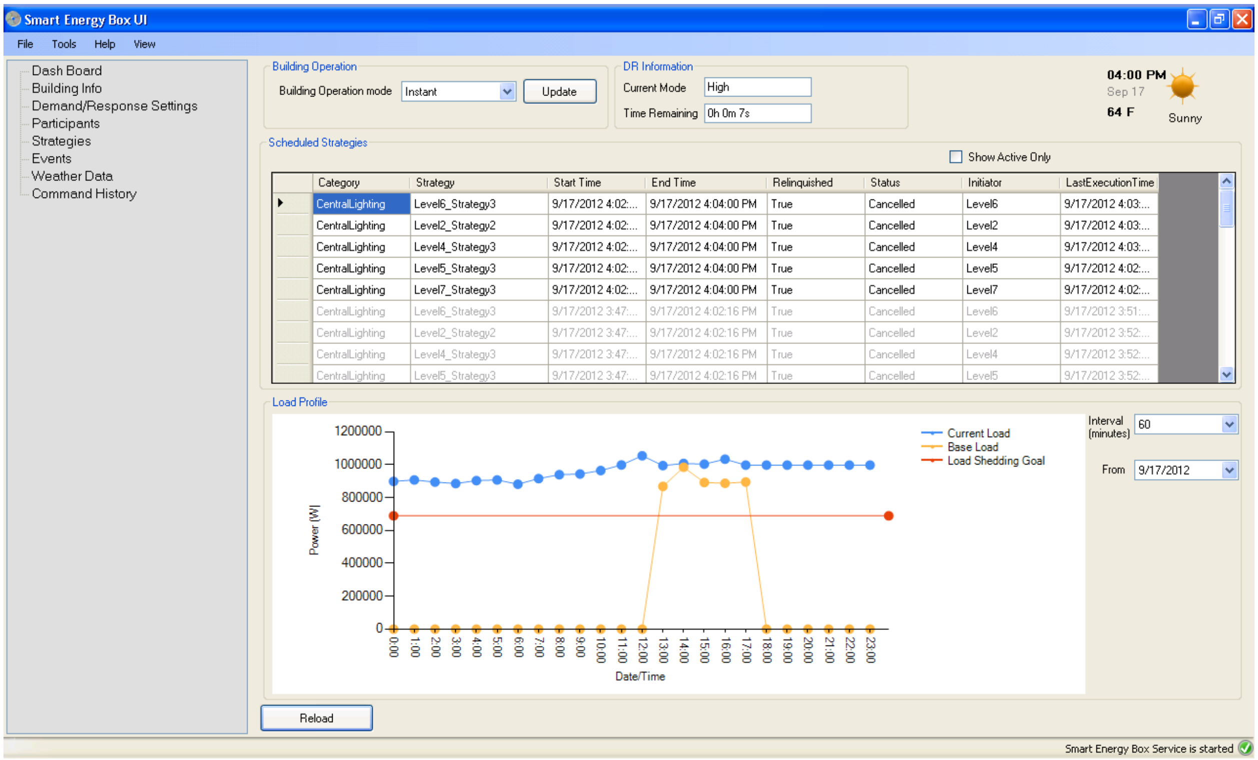

Figure 10: User Interface Dashboard

7.1.2.1.4 Importing lighting Strategies

An XL based strategy import was introduced to import Participants and Strategies. This creates Participants and strategies in one stretch.

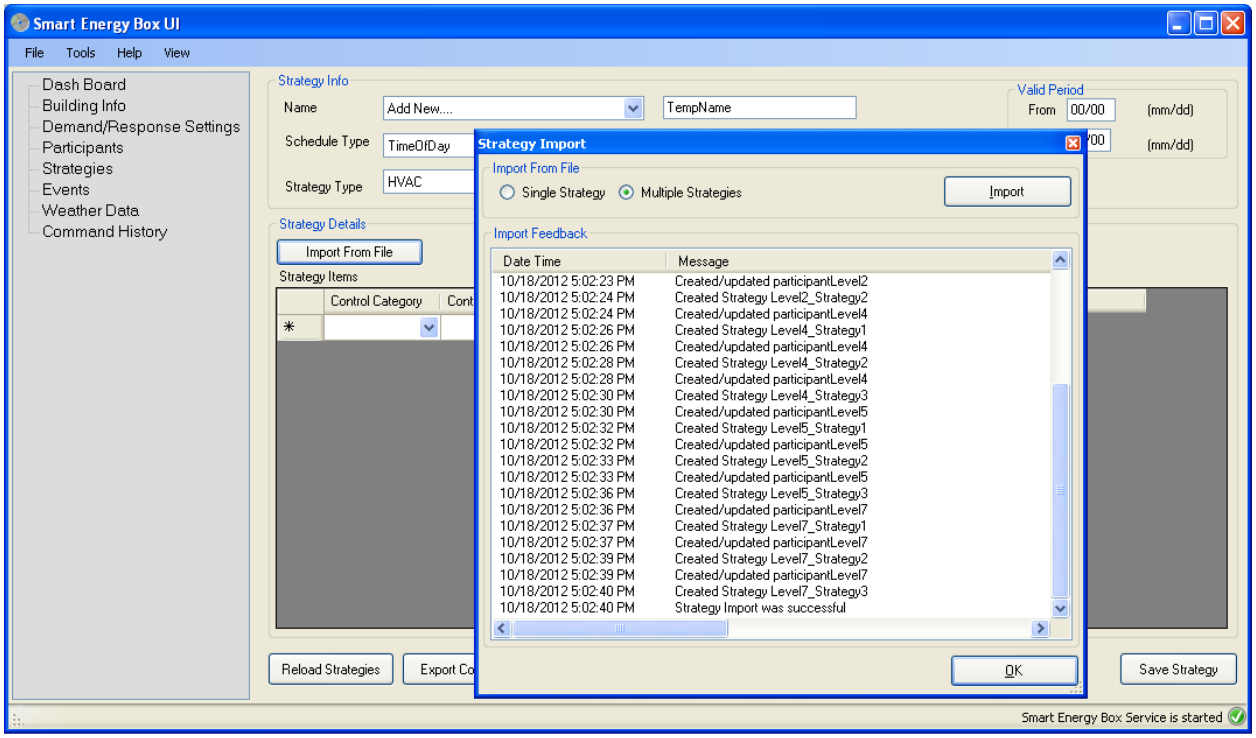

Figure 11: User Interface Importing of Strategies

7.1.2.2 Events

This screen shows DR event information and also the current strategy that has been selected.

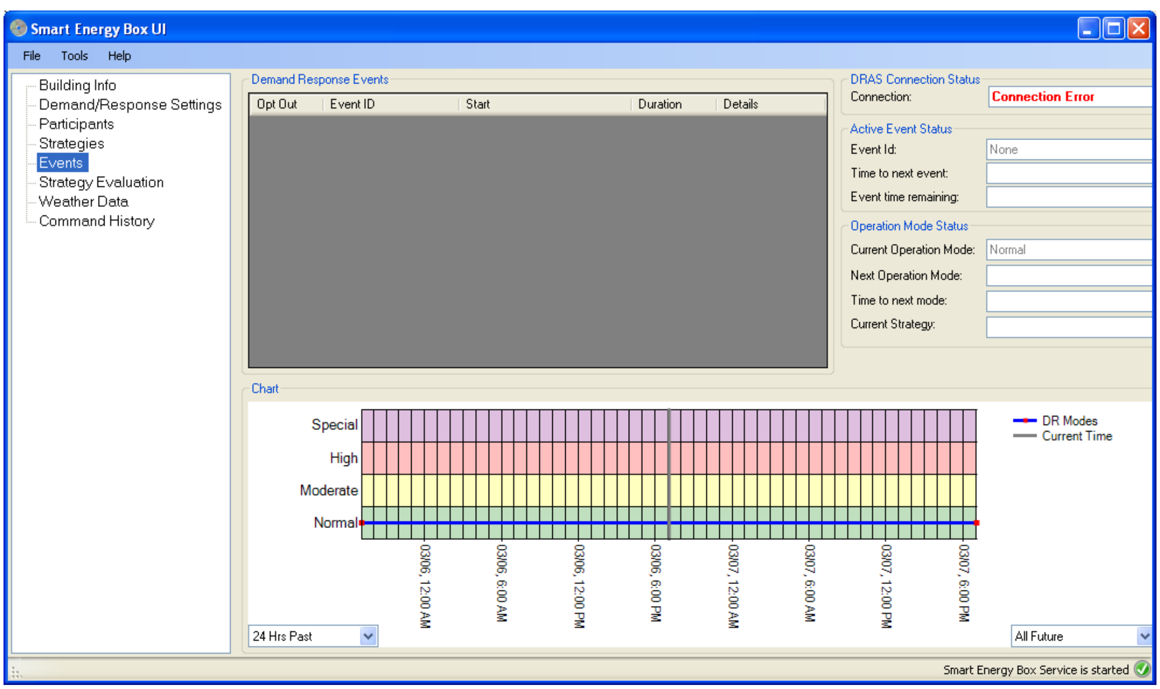

Figure 12: User Interface DR Event information

7.1.2.3 Weather data

This screen shows real time weather data that has been imported from NOAA weather service by the SEB.

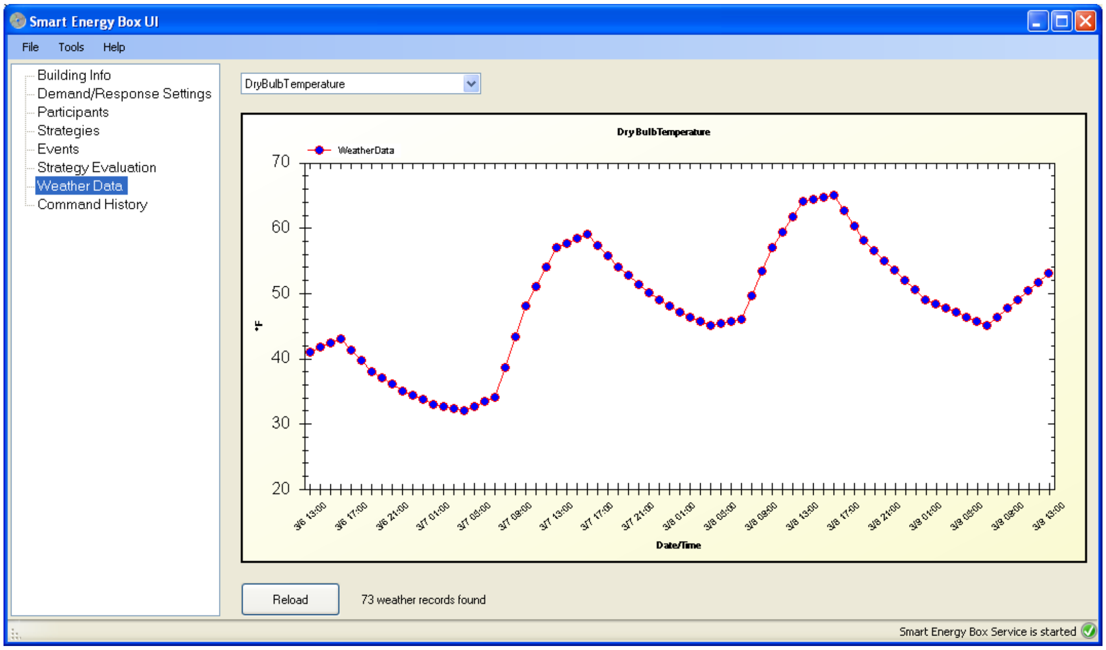

Figure 13: User Interface Weather data View

7.1.2.4 Point List

Building control points can be configured using this screen. This provides a file based import, edit or export mechanism.

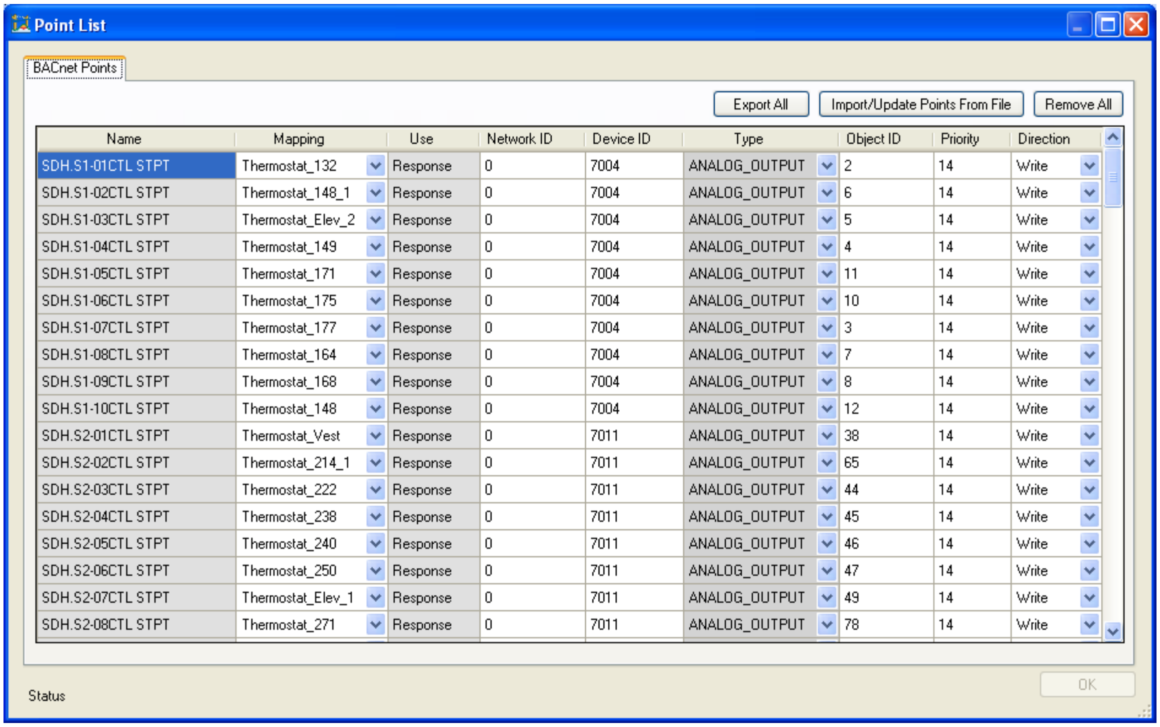

Name	Mapping	Use	Network ID	Device ID	Type	Object ID	Priority	Direction
SDH.S1-01CTL STPT	Thermostat_132	Response	0	7004	ANALOG_OUTPUT	2	14	Write
SDH.S1-02CTL STPT	Thermostat_148_1	Response	0	7004	ANALOG_OUTPUT	6	14	Write
SDH.S1-03CTL STPT	Thermostat_Elev_2	Response	0	7004	ANALOG_OUTPUT	5	14	Write
SDH.S1-04CTL STPT	Thermostat_149	Response	0	7004	ANALOG_OUTPUT	4	14	Write
SDH.S1-05CTL STPT	Thermostat_171	Response	0	7004	ANALOG_OUTPUT	11	14	Write
SDH.S1-06CTL STPT	Thermostat_175	Response	0	7004	ANALOG_OUTPUT	10	14	Write
SDH.S1-07CTL STPT	Thermostat_177	Response	0	7004	ANALOG_OUTPUT	3	14	Write
SDH.S1-08CTL STPT	Thermostat_164	Response	0	7004	ANALOG_OUTPUT	7	14	Write
SDH.S1-09CTL STPT	Thermostat_168	Response	0	7004	ANALOG_OUTPUT	8	14	Write
SDH.S1-10CTL STPT	Thermostat_148	Response	0	7004	ANALOG_OUTPUT	12	14	Write
SDH.S2-01CTL STPT	Thermostat_Vest	Response	0	7011	ANALOG_OUTPUT	38	14	Write
SDH.S2-02CTL STPT	Thermostat_214_1	Response	0	7011	ANALOG_OUTPUT	65	14	Write
SDH.S2-03CTL STPT	Thermostat_222	Response	0	7011	ANALOG_OUTPUT	44	14	Write
SDH.S2-04CTL STPT	Thermostat_238	Response	0	7011	ANALOG_OUTPUT	45	14	Write
SDH.S2-05CTL STPT	Thermostat_240	Response	0	7011	ANALOG_OUTPUT	46	14	Write
SDH.S2-06CTL STPT	Thermostat_250	Response	0	7011	ANALOG_OUTPUT	47	14	Write
SDH.S2-07CTL STPT	Thermostat_Elev_1	Response	0	7011	ANALOG_OUTPUT	49	14	Write
SDH.S2-08CTL STPT	Thermostat_271	Response	0	7011	ANALOG_OUTPUT	78	14	Write

Figure 14: User Interface Building Control Points Import/Export/View

7.1.2.5 Strategy Evaluation

Strategies that were evaluated by Smart Energy Box can be viewed by facility manger using the following screen. It explains the basis for selection of a particular strategy

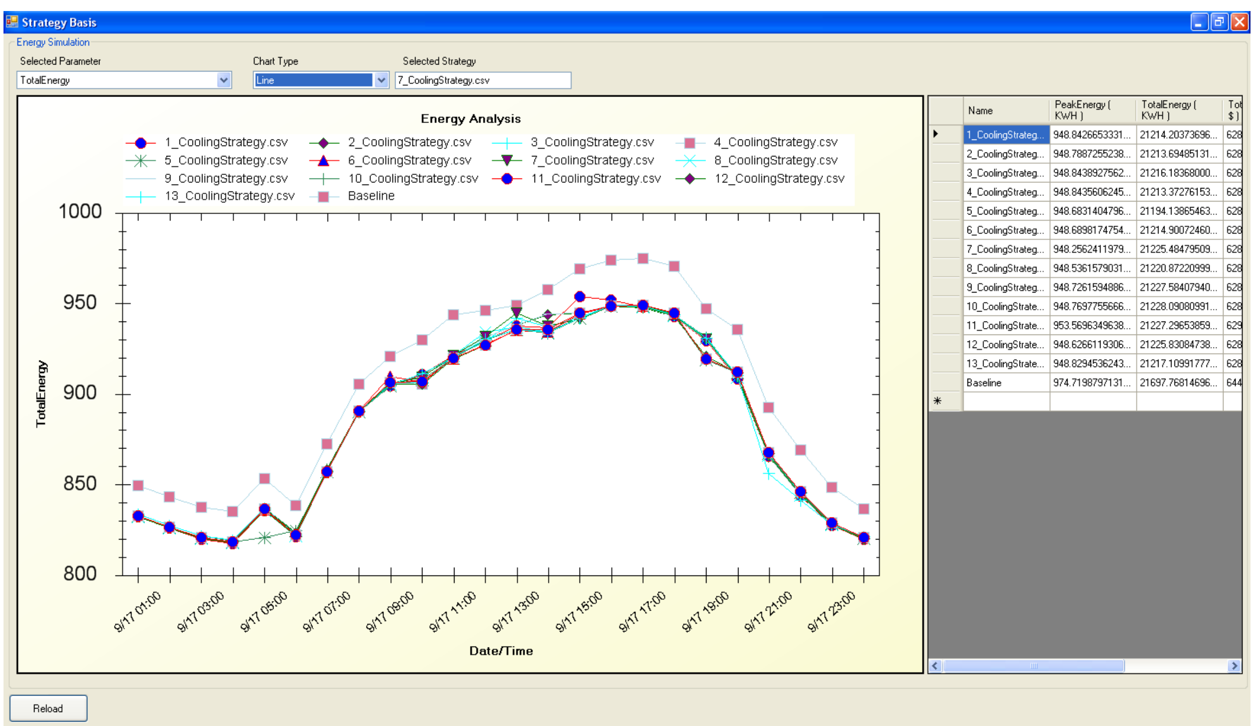

Name	PeakEnergy (KWH)	TotalEnergy (KWH)	Tot $)
1_CoolingStrateg...	948.8426653331...	21214.20373696...	628
2_CoolingStrateg...	948.7887255238...	21213.69485131...	628
3_CoolingStrateg...	948.8438927562...	21216.18368000...	628
4_CoolingStrateg...	948.8435606245...	21213.37276153...	628
5_CoolingStrateg...	948.6831404796...	21194.13865463...	628
6_CoolingStrateg...	948.6898174754...	21214.90072460...	628
7_CoolingStrateg...	948.2562411979...	21225.48479509...	628
8_CoolingStrateg...	948.5361579031...	21220.87220999...	628
9_CoolingStrateg...	948.7261594886...	21227.58407940...	628
10_CoolingStrate...	948.7697755666...	21228.09080991...	628
11_CoolingStrate...	953.5696349638...	21227.29653859...	629
12_CoolingStrate...	948.6266119306...	21225.83084738...	628
13_CoolingStrate...	948.8294536243...	21217.10991777...	628
Baseline	974.7198797131...	21697.76814696...	644

Figure 15: User Interface HVAC Strategy Evaluation

7.1.2.6 Gateway architecture

Gateways allow a fully integrated demand response system that not only are able to control central systems (HVAC, lighting, etc.), but also are able to control distributed plug-in loads in various zones within a building such as office rooms, kitchen, lab, auditoriums, conference rooms etc. An independent gateway can be configured to manage non-centrally controlled equipment such as a task lights, laptops printers, kitchen equipment etc. A Gateway can be configured with sensor box that can sense humidity, temperature, occupancy, and illumination to make an intelligent decision of load reduction within its zone.

Figure below shows the architecture of Gateway; Appendix K provides the detailed description of the Gateway Components.

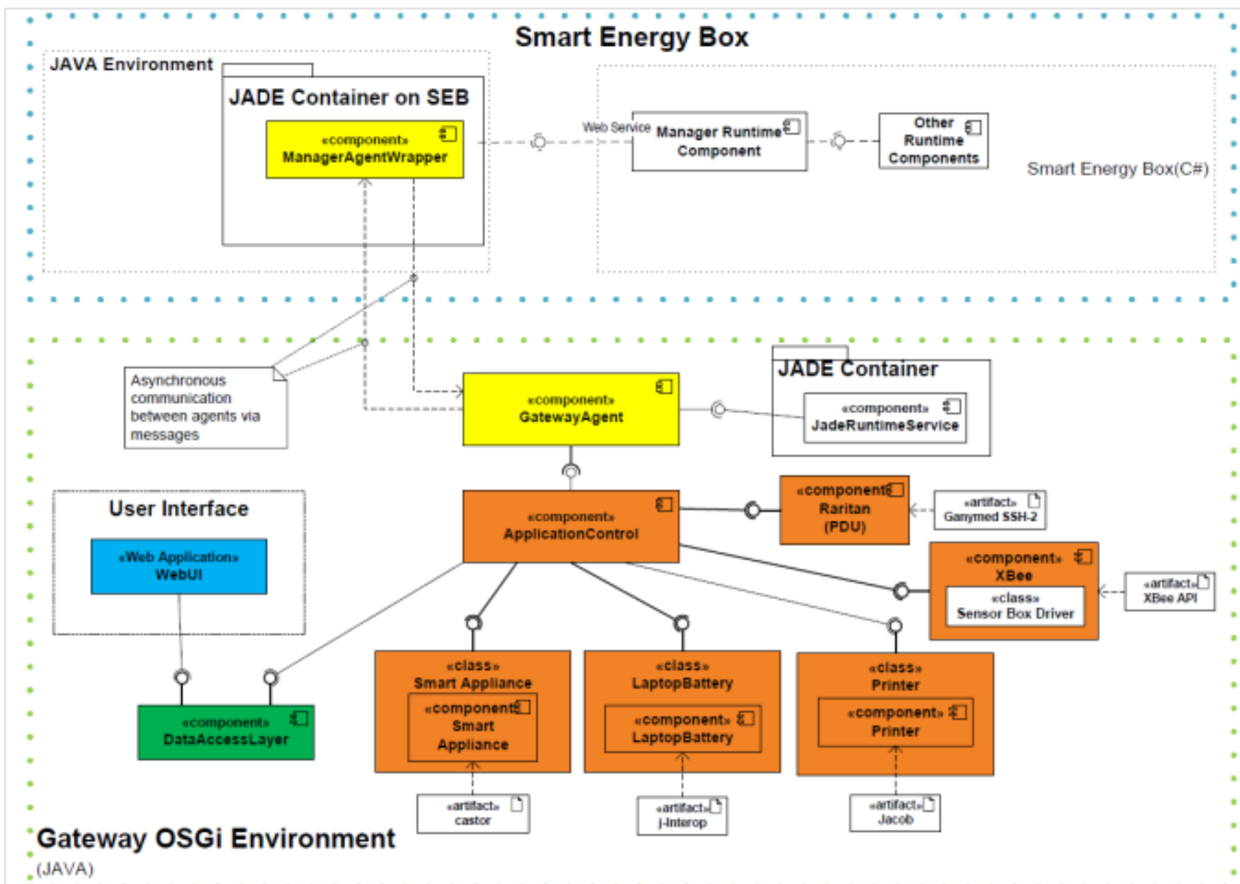

Figure 16: Component diagram of the Gateway

7.1.2.7 Integration with BM and, DRAS

Siemens developed the Smart Energy Box (SEB) as a building to grid connection. It can receive a signal from the power grid (DRAS) and automatically moderate the actions on building automation systems (BMS Adapter) for HVAC, lighting load management and on local Gateway Controllers for plug load management as shown on the Figure below. The following figure shows the details of the communication channels among DRAS server, SEB, BMS systems and local Gateway Controllers.

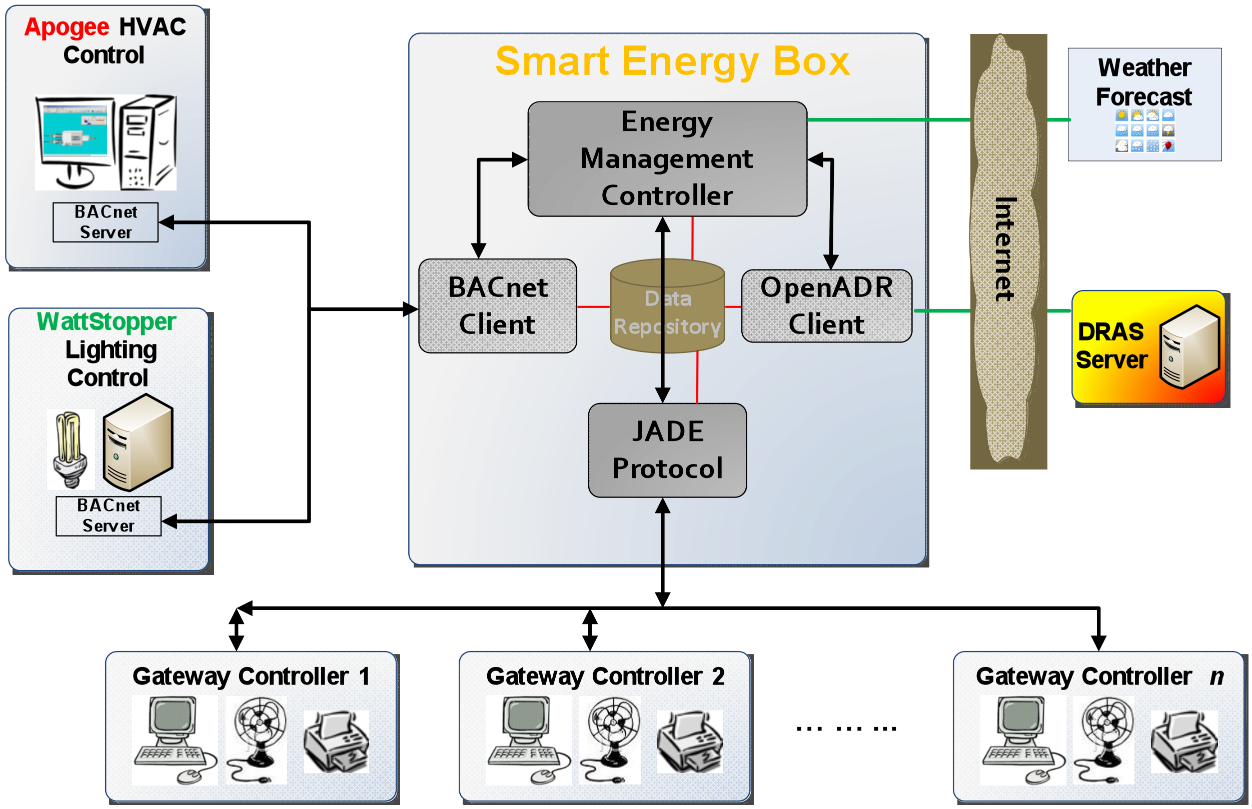

Figure 17: Communication channels of Gateway.

As explained earlier, the OpenADR Client (version OpenADR 1.0) implemented within the SEB brings the DR signal from the DRAS server. Upon receiving the DR event, if there is enough time to perform pre-cooling, then the controller simulates all available strategies using real-time weather forecast and selects best one. The selected strategy will be applied on building on start of the DR day. If there was not enough time to perform pre-cool then the static strategies that were configured for mode of the DR event (Moderate or High) will be selected to apply during DR period. BACnet Adapter that was available within the SEB performs the real write /read operation of the control set points.

However as explained in Adaptive approach, it may be possible the goal will not be able to met during DR period due to load prediction error, abnormalities in electric equipment usage or occupancy fluctuations. In this case SEB contracts the remaining load to instant load controls such as central lighting, distributed plug load control.

To perform plug load control for demand response, the distributed Gateway is integrated into DIADR system with Java based Agent Development Environment (JADE) framework. A JADE coordinator within SEB bridges SEB and gateway controllers to perform interactive negotiations, as shown in the figure above.

7.1.2.8 Optimization

Three-stage optimization was adopted to perform the demand response control using Smart Energy Box and distributed Gateways.

- Stage I: Off-line optimization and generate OSP (Optimal Strategy Pool)
- Stage II: On-line simulation evaluation on OSP
- Stage III: Instant Load Control (Central Lighting and Distributed Plug Load control)

7.1.2.8.1 Stage I: Off-line optimization and generate OSP (Optimal Strategy Pool)

Siemens have used offline optimization mechanism to generate optimal strategy pool. The result showed that the Optimal Strategy Pool (OSP) optimization scheme has superior advantages over Exhaustive Search (ES), Genetic Algorithm (GA) and Pattern Based Selection (PBS), in terms of both high accuracy and low computation requirement for on-line implementation. The simulation-based experiment showed that an optimal HVAC control strategy is able to reduce CITRIS peak load by as much as 109 kW (~18%) in a typical hot day of Berkeley, CA. The OSP optimization scheme will be detailed in the later section, together with the result of its field tests done with SEB. The compact HVAC control strategies are done off-line by the Optimal Strategy Pool (OSP) algorithm, which is described in Appendix L.

7.1.2.8.2 Weather Clustering

The historical August weather data of Berkeley, CA for the years between 2002 and 2010 has been collected. For each August day, the hourly dry bulb temperature and its simulated baseline peak load are included in the feature space and subject to dimension reduction by principal component analysis (PCA). And then K-means clustering algorithm is applied. In this study, at least 19 clusters are required to ensure the variance in each cluster is lower than a pre-determined threshold. The centroid weather profile of each cluster is then obtained by taking average over all member profiles. All 19 centroid August weather profiles are depicted in the following figure.

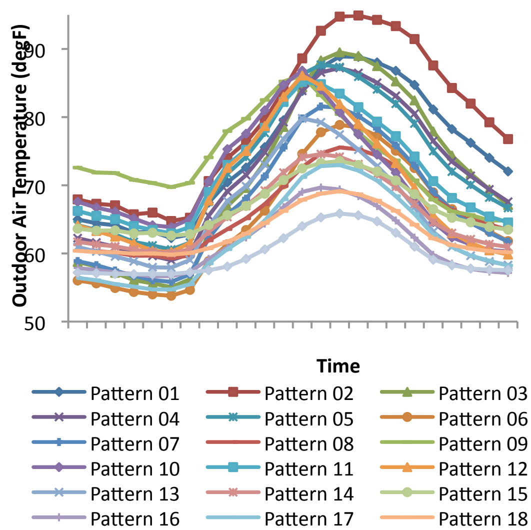

Figure 18: Berkeley historical weather patterns--daily outdoor air temperature.

7.1.2.8.3 DR control strategy design

Global temperature setpoint adjustment (GTA), supply air temperature setpoint (SAT), supply fan pressure setting (SFP) and minimum ventilation (MinVent) are considered in DR control. The controls are detailed as the following.

GTA: The cooling set points of all zones are subject to change throughout the day. Pre-cooling and exponential set-up strategy is applied. As depicted in the following Figure, between 0:00 and T1, the cooling setpoint is set at the current baseline value, which is 72°F; between T1 and T2, the cooling setpoint is set at 70°F (pre-cooling); between T2 and T3, the cooling setpoint is set up exponentially to 78°F (exponential set-up); and between T3 and 24:00, the cooling setpoint is set back to 72°F. All zones are using the same GTA strategy. To reduce the size of solution space, only the three time points (i.e., T1, T2 and T3) are considered as decision variables. The setpoint values at T1, T2 and T3 are fixed at 72°F, 70°F and 78°F, respectively. Furthermore, time points can only be integer hours within the following ranges: 5≤ T1 ≤9, T1<T2 ≤14, and 17 ≤ T3 ≤19.

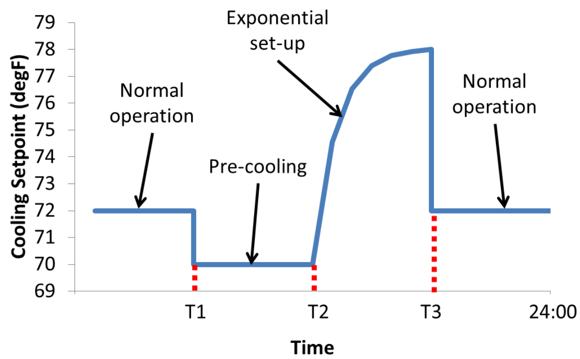

Figure 19: Setpoint strategy

SAT: There are two AHUs dedicated for office spaces. They are controlled by the same SAT setpoint, whose current value is 56°F; and they share the same supply air duct. SAT setpoint values between 51 and 60°F were explored, with interval of 1°F. We assume that SAT setpoint only changes at the beginning of the DR day, to simplify the problem formulation.

SFP: The two supply fans (SF-2A and SF-2B) in the building are variable volume fans. The operation speed is controlled by a proportional-integral-derivative controller (PID controller) to maintain the fan pressure at a fixed setpoint, which is currently 1350 Pa. SFP setpoint values between 1150 Pa and 1350 Pa were explored, with intervals of 50 Pa. Again, to reduce the complexity of the problem, we assume that SFP setpoint only changes at the beginning of the DR day.

MinVent: The current minimum air flow rate settings for most of zones are found to be higher than the required levels defined by the standard (ASHRAE 62.1-2010). New minimum air flow requirements for all zones are calculated based on the area, occupancy density and functionality (ASHRAE 62.1-2010), and are implemented in a retrofitting model. The retrofitting model is the same as the base model, except for the adjusted minimum air flow requirements. The total minimum supply air volume is reduced by 28%, in the retrofitting model.

A DR strategy is defined by five decision variables, which are GTA(T1), GTA(T2), GTA(T3), SAT and SFP. The total number of strategies is 5250. All strategies are tested on the retrofitting model, rather than on the base model. Therefore, MinVent applies to all tested strategies, by default, even though it is not mentioned in the strategy description.

7.1.2.8.4 Stage II: On-line simulation evaluation on OSP

SEB is able to make decisions, for a DR event, on the control of building central loads, i.e., the HVAC system and central lighting, based on the forecast of weather. The decision making module and the on-line control optimization algorithm is essential to such functionality. Upon receiving of DR event, SEB simulates all strategies against real-time weather forecast and selects the optimal strategy.

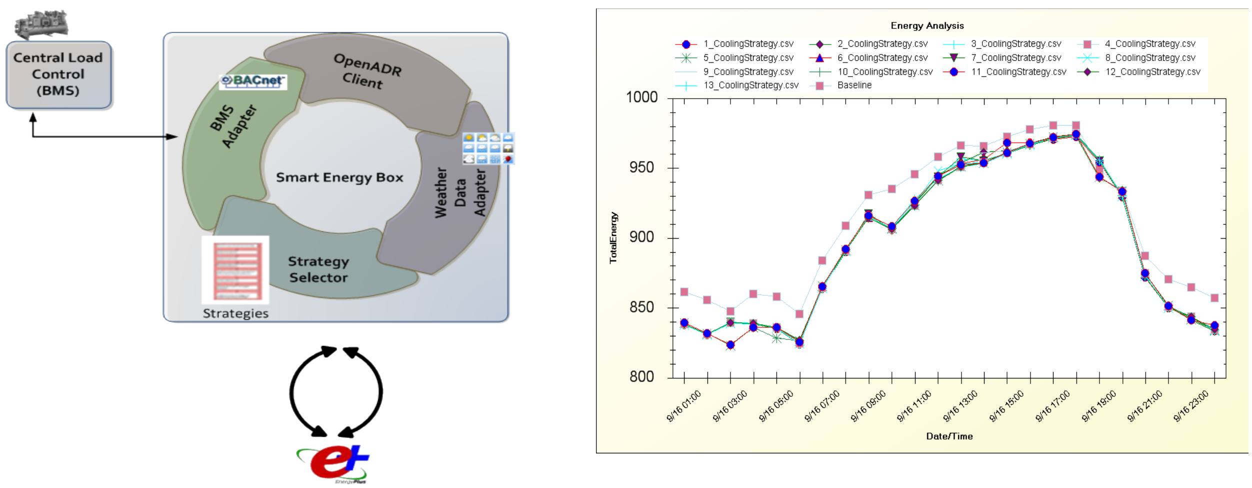

Figure 20: Online Strategy simulation by SEB

7.1.2.8.5 Stage III: Instant Load Control (Central Lighting and Distributed Plug Load control)

As explained in the Adaptive Framework, during DR period, the load reduction goal that still needs to be achieved will be contracted to instant load control agents through Market based negotiation mechanism. This negotiation happens every 15 minutes; however SEB continuously monitors the status of load reduction almost every minute and assesses if any additional load reduction needs to be performed next 15 minutes.

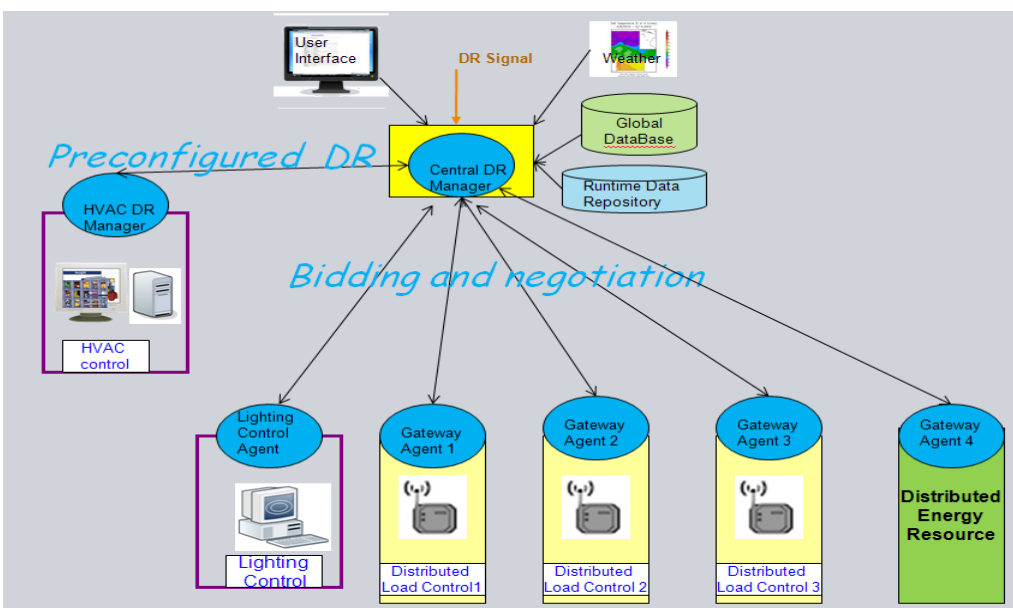

Figure 21: Stage III, Market based Negotiation with agents

7.1.3 Characterization of building operation

The average electrical demand of Sutardja Dai Hall in Academic Year 2011-2012 (July-June) was approximately 894 kW when the building used the steam-driven absorption chiller and 964 kW when the building used the electricity-driven centrifugal chiller. The figure below shows the load from the two main substations, MSA and MSB, beginning with the first sMAP feeds in May 2011 through December 2012. Over the course of this project, many factors affected the energy consumption: which chiller[5] was running, the gradual installation and use of tools in the nanofabrication lab, the addition of laboratories and other rooms to the first floor, and energy efficiency measures, such as the addition of Variable Frequency Drives (VFDs) to the chilled and condenser water pumps and a dynamic ventilation regime (minimum ventilation rates dynamically changed based on economizer and assumed occupancy to maintain 15 cfm of outside air). In addition, daily factors, such as outdoor air and solar loads and occupancy played a role.

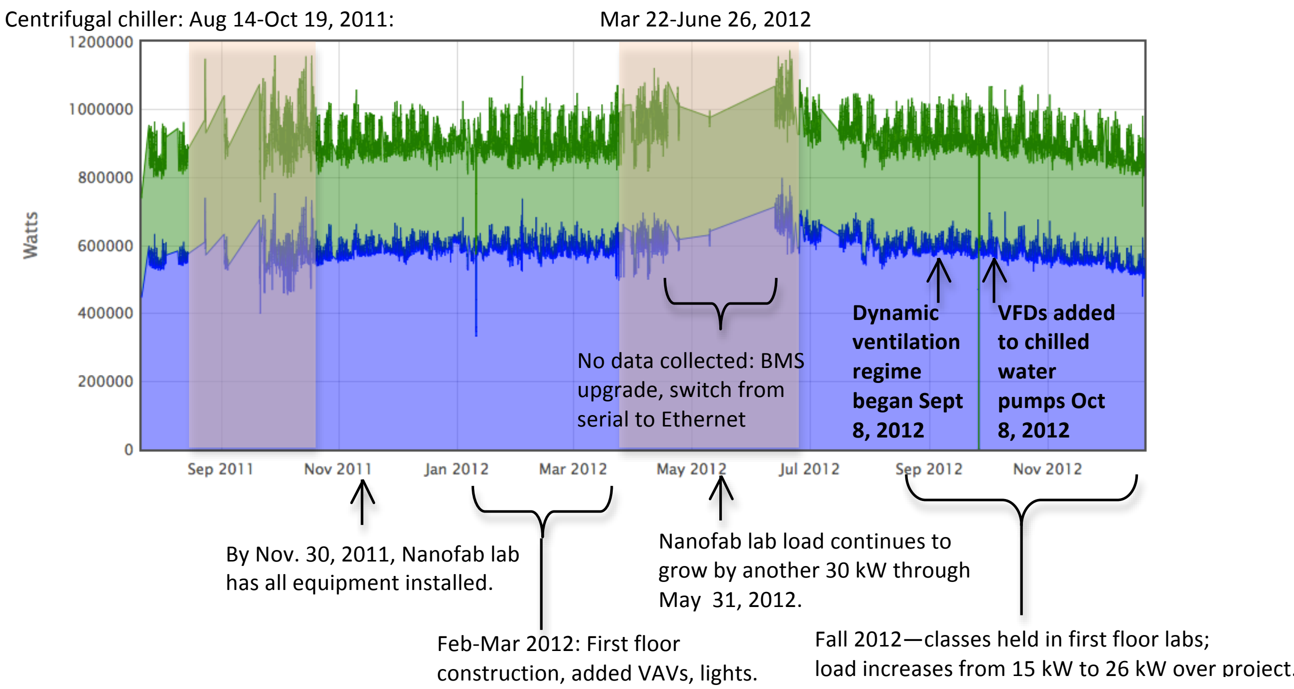

Figure 22: Whole building electrical load of Sutardja Dai Hall from May 22, 2011 to Dec 31, 2012 (MSA upper/green, MSB lower/blue).

7.1.3.1 Discussion of components

Steve Yen took all the data streams in sMAP, calculated 15-minute average data, and developed the following figures (see Appendix P). Of the annual 8053 megawatt-hour energy consumption (July 2011-June 2012), about 32% is attributable to the nanofabrication lab process tools; this increased to 35% by the end of the project. About 16% is consumed by a combination of the fab lab air handling units and the absorption chiller (a load which is split or shared between the lab and the office); this increased to 18% by the end of the project (primarily because the facilities manager decided to run the absorption

[5] The absorption chiller broke down in mid-August 2011, and thus the centrifugal chiller ran in August-October 2011, but we discovered this chiller was short-cycling. As soon as the absorption chiller was fixed, the building was switched to this chiller while the problem with the centrifugal chiller was addressed.

chiller instead of the centrifugal chiller. About 12% is for emergency lighting (a load split between the fab lab and the office) and exhaust fans for the fab lab.

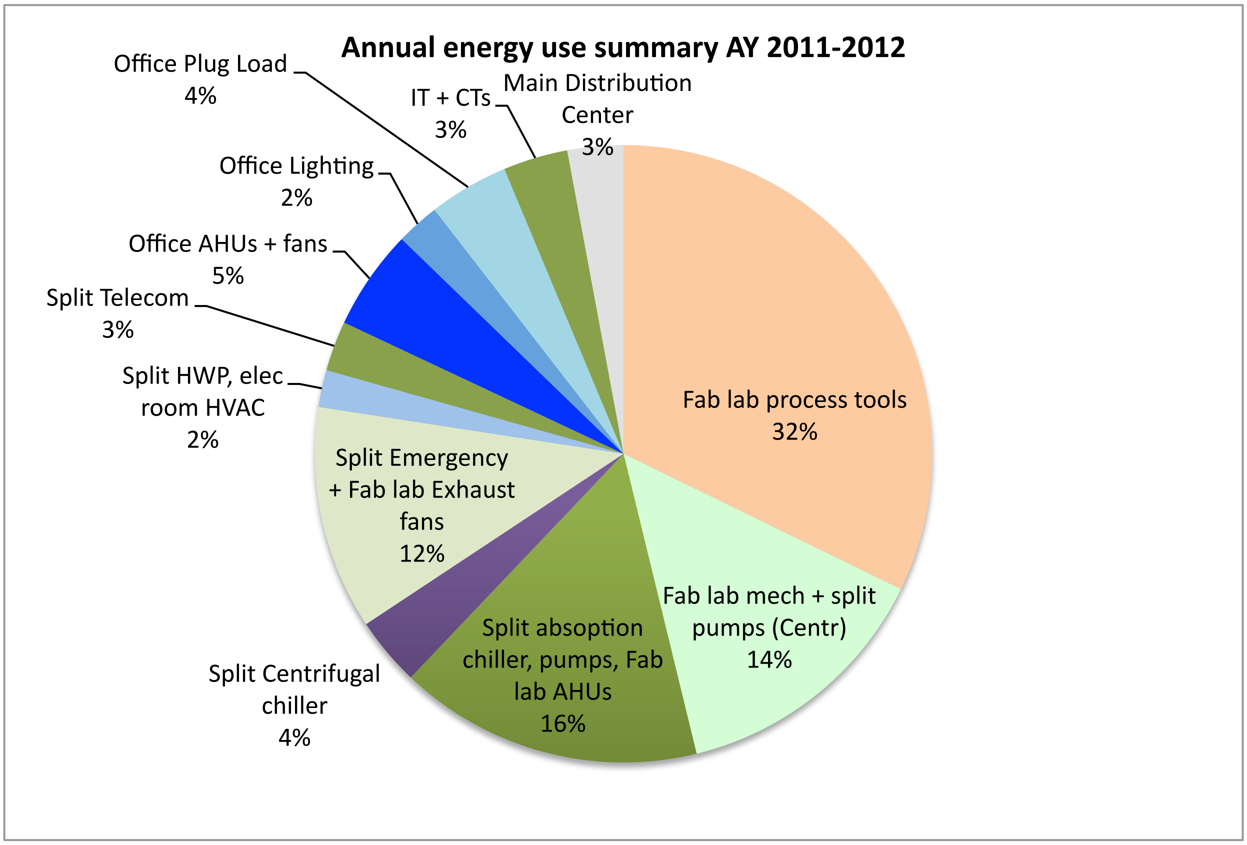

(IT = data center, CTs = cooling towers, HWP = hot water pumps, AHUs= Air Handling Units)

Figure 23: Total building annual energy use by submeter data for July 2011-June 2012; some submeters are split between office and nano fab lab.

The office portion of the building was isolated to the best of our ability; the methodology is described briefly in Appendix B and in more detail in Appendix P. The office portion of the building consumed approximately 1924 mWh per year—about one quarter of the whole building energy consumption—from July 2011-June2012. The office portion is about 81,000 assignable square feet; if one can assume typical office spaces are approximately 0.6 ASF/GSF, then this corresponds to 14 kWh/gross square feet. For comparison: 11 kWh/gsf UC/CSU is the 1999 average UC/CSU benchmark for classroom and office buildings in the Berkeley Climate (personal communication with Karl Brown 2013). As another point of reference, 14 kWh/gsf is the average benchmark for classroom and office buildings in the hot summer inland Merced California climate, and 8.5 kWh/gsf is the well documented performance of the energy efficient Classroom and Office Building at UC Merced[6] (Brown 2010).

The following figure shows the results of the plug load audit—how many of each appliance were found in the building, and the approximate power draw. The aggregate plug load, not accounting for concurrence (diversity), is 68 kW. For comparison, the average demand of the office portion of the building (using the absorption chiller) can be estimated at one-quarter of 894 kW or 224 kW. We note

[6] These numbers are as-measured at the building, downstream of campus distribution and transformation losses.

here that there are 26 electric tankless water heaters for lavatories in kitchens and bathrooms throughout the office portion of the building (as well as an electric tank water on the first floor), each at 6000 watts. These are included in the lighting submeters.

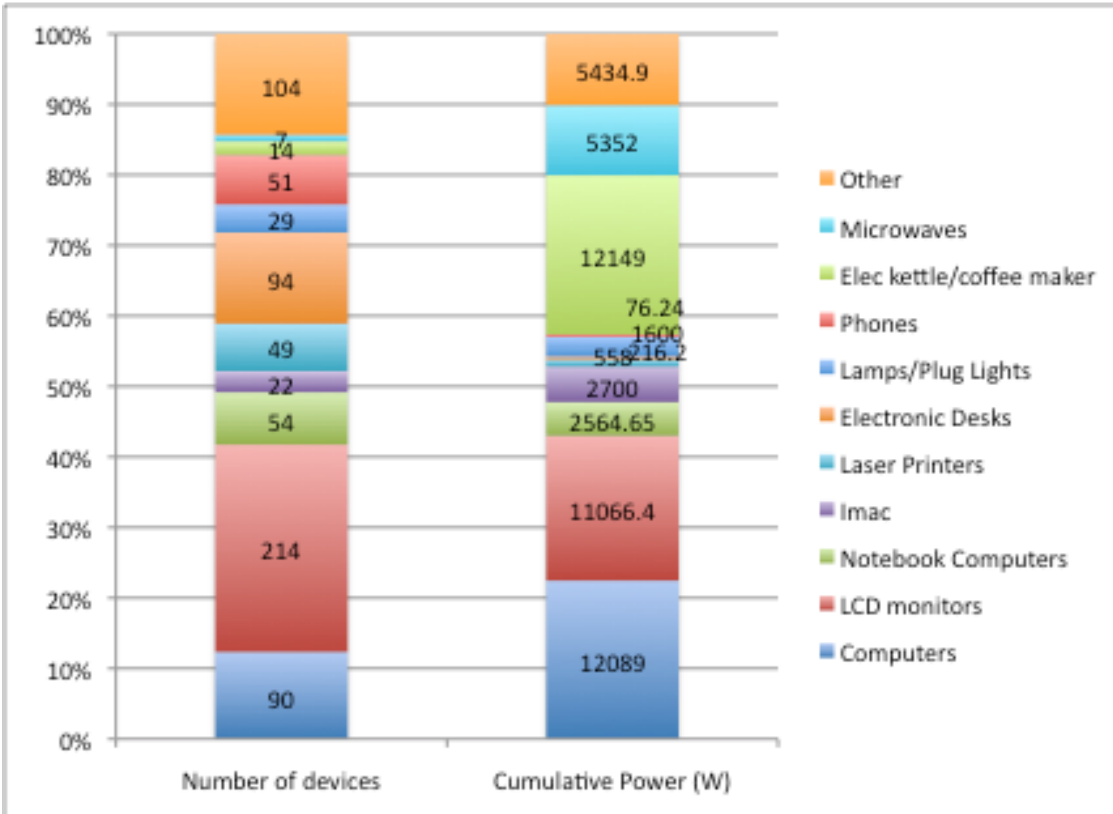

Figure 24: Appliances found in SDH in 2011 audit.

7.1.3.2 Discussion of peak days

We initially analyzed the submeter components of the peak days from the first baseline year (July 2011-June 2012), seen in the figure below. We chose peak days when school was in session as well as out of session, and using the absorption chiller as well as the centrifugal chiller.

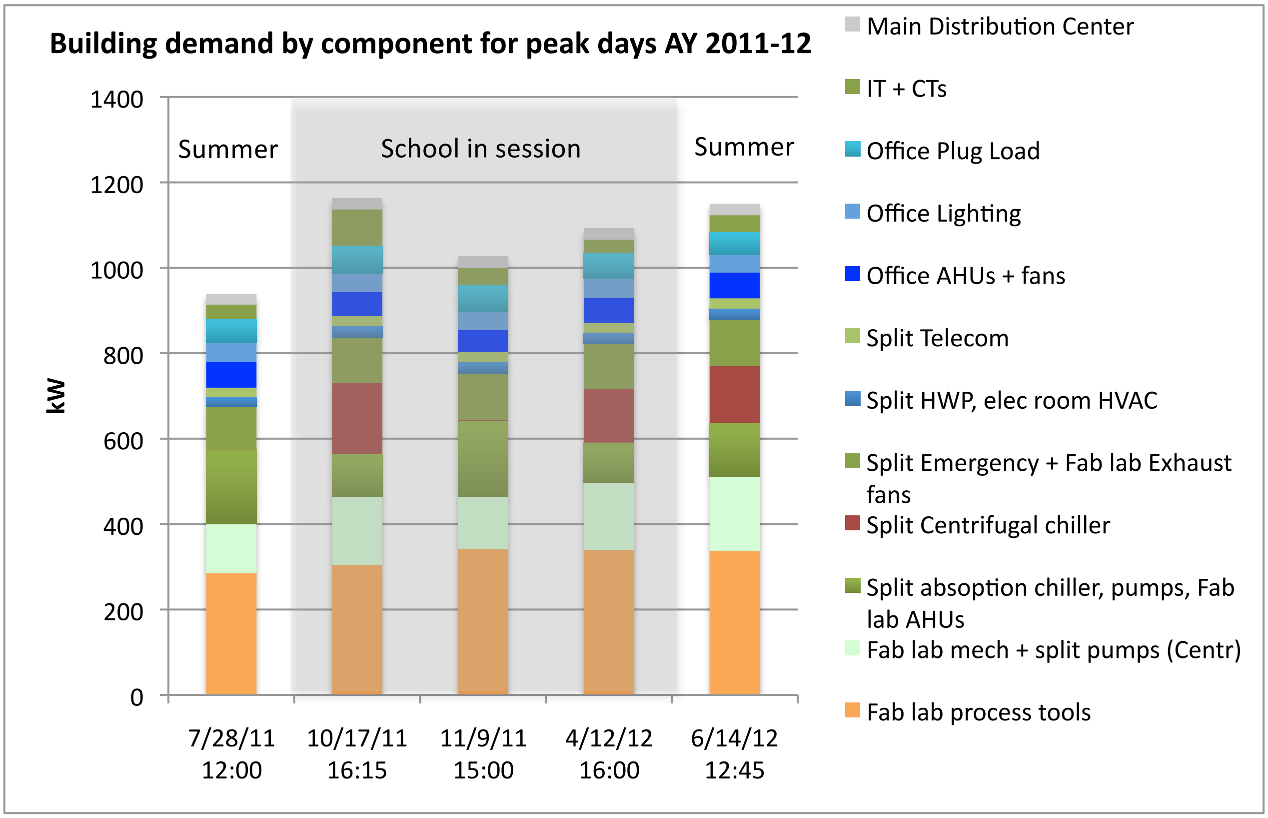

Figure 25: Whole building peak demand days in the Academic Year 2011-2012 (begins in July, ends in June).

The highest electrical demand peak recorded over the course of the project occurred on a warm day in June 2012 using the centrifugal chiller: 1176 kW. The highest peak recorded using the absorption chiller was 1085 kW on a very warm day, Oct 17, 2012.

We estimated the office portion of load for these days, using the formula described in Appendix B. We note here that initially we had estimated the office portion of the building used approximately 25% of the chilled water load on average (the figure used for the estimation below). However, after a subsequent analysis, we believe that on hot days the office portion may use up to 40% of the chilled water load. We estimated that the office portion of the load on peak days was approximately 260-303 kW.

7.2 Software tools

7.2.1 EnergyPlus model

Rongxin Yin of Lawrence Berkeley National Laboratory developed the EnergyPlus model; the process of developing and using this model is outlined in a paper referenced in Appendix G.

7.2.2 Demand Response Capacity Assessment & Operation Assistance Tool (DRCAOT)

Demand Response Capacity Assessment & Operation Assessment Tool (DRCAOT) is a tool designed to assist building facility managers to better assess the load reduction capacity of the building in response to the demand response requests from the utility, with various levels of awareness of the weather condition at the planned demand response day. DRCAOT is also allowing the users to design control strategies for the building HVAC system, and is able to identify the optimal strategy based on simulation evaluation though a strategy library. The corresponding energy savings and/or cost reductions will be quantified with respect to different weather conditions, demand response rate structures and HVAC control strategies. By consulting this tool, cost effective outcomes can be expected, when the building owners signing up demand response program with service providers; energy efficient operation could also be achieved in non-DR periods.

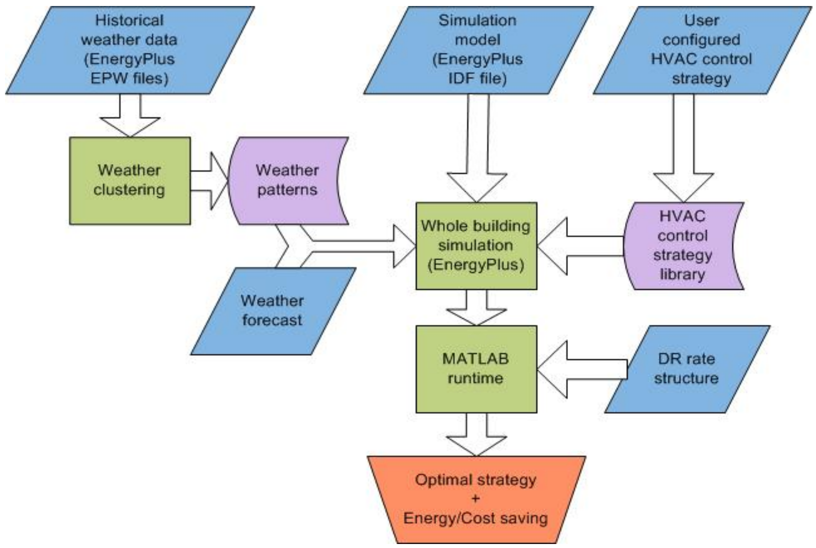

Figure 26: DRCAOT architecture design

DRCAOT has three main function features, including weather clustering, energy simulation and operation strategy configuration.

7.2.2.1 Weather clustering

Given the location information, the local historical weather data can be obtained from the Internet. This data is stored in multiple files with EnergyPlus weather input format (.epw). The hourly temperature profile of each day in history will be analyzed, and categorized into a number of patterns. The representative profile of each pattern will be stored for the later use in energy simulation.

When the weather forecast information is available, the forecast weather can be directly fed to the simulation for the day-ahead decision-making. However, in case that such information is not available or DR capacity for an extended period of time is of interest, the weather patterns can be good representatives for different weather conditions.

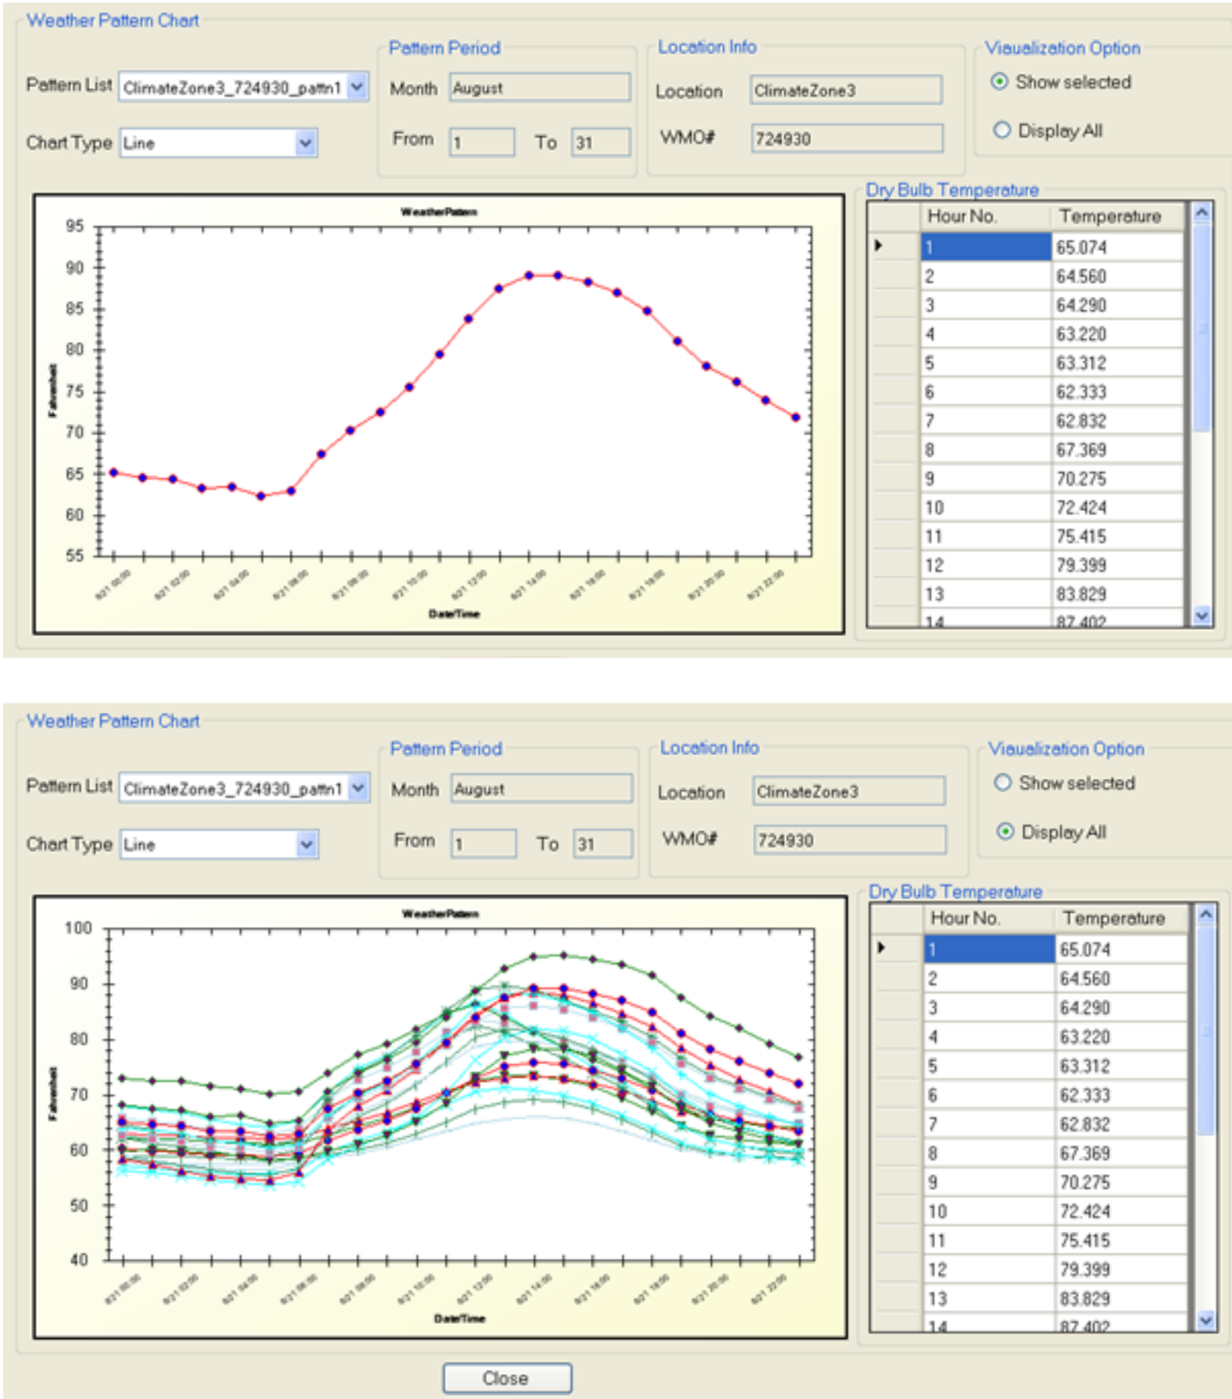

Figure 27: Weather pattern illustration.

7.2.2.2 Whole building energy simulation model

DRCAOT incorporates the state-of-the-art whole building simulation engine, EnergyPlus, to perform simulation based control optimization. The heat balance, HVAC performance and loads are calculated in the background. The system state and energy usage are updated sub-hourly (configurable between once per hour and once per minute), thus DR programs with various time resolution requirements can be accommodated easily. EnergyPlus simulation model used in the simulation can be very simple, or detailed with respect to the building envelop, system configuration and operations, providing the users with flexibility to balance between simplicity and prediction accuracy.

Figure 28: DRCAOT energy simulation configuration and result presentation

7.2.2.3 Strategy configuration approaches

Several controllable points are available for configuring control strategies, including zone temperature, supply air temperature, etc. DRCAOT provides an initial library of HVAC control strategies, which includes a set of pre-cooling strategies. It also allows the users to create and customize those strategies, by selecting the controllable points and specifying point values for each hour of the day. Strategies can also be created in batch when the points are customized in ranges.

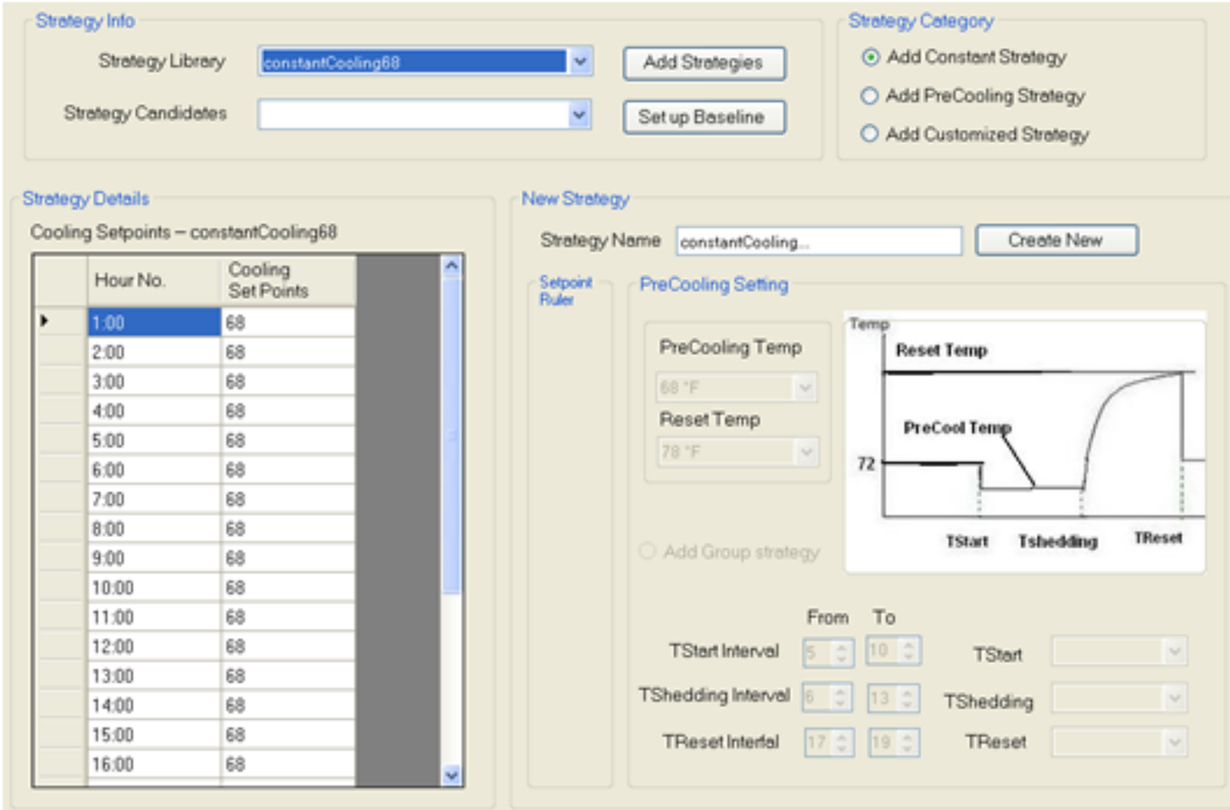

Figure 29: HVAC control strategy customization

7.2.2.4 MATLAB runtime

MATLAB runtime component is communicating with EnergyPlus simulation engine through the External Interface functionality. MATLAB coordinates all simulation evaluations and collects the energy consumption data during the simulation. Data processing and optimization are conducted by MATLAB. Advanced optimization techniques can be applied if necessary. In the end, the best strategy among all the candidates, together with its simulated consumption profile, peak load reduction, energy reduction and cost savings, will be graphically presented to the user.

DRCAOT is an advanced managing tool that supports the facility manager with simulation proved optimal control strategy, in DR decision-making. The clustering technique provides as accurate weather input as possible for the building energy simulation, when the forest weather does not exist. The EnergyPlus simulation engine, coupled with MATLAB runtime components, is able to estimate the sub-hourly energy usage by different control strategies. Thus, the best control strategy can be obtained, and the capacity of peak load reduction, energy and/or cost savings under different DR scenarios can be assessed.

7.3 Test plan

The ultimate goal was to automatically reduce peak electrical load from Sutardja Dai Hall based on an external signal while maintaining a reasonable environment for the occupants. From knowledge of California's weather patterns, common sense, and simulations, we knew that we would get the most curtailment during coastal California "heat storms" — one-to-five day periods of much hotter than typical daytime temperatures coincident with offshore wind patterns. While Berkeley rarely has the temperatures in excess of 90 deg F commonly experienced further inland, the difference in maximum temperature between heat storm days and typical days can be as great on the coast as it is inland. So we anticipated heat storm days with maximum temperatures in the 80s (deg F) would produce the demand response opportunities sought by the project. As luck would have it, we actually experienced an extremely unusual hot 2-day period with maximum temperatures above 94 F, during which we achieved the most dramatic demand response results.

We knew that the office peak load would be approximately 234-303 kW and thus our goal was to shed 70-101 kW.

The basic components of the test plan are shown below. First we isolated the office load from the rest of the building. We identified peak loads, and the components of these loads for both warm and hotter days, and with either chiller. While not explicitly required by the project, as we developed demand response scenarios, we saw an opportunity to develop long-term solutions with respect to energy efficiency in conjunction with the peak energy curtailing strategies[7]. Finally, in determining the DR strategies, we explored strategies that could be used with day-ahead notification as well as hour ahead notification. We envision different "depths" of demand response for different periods of time.

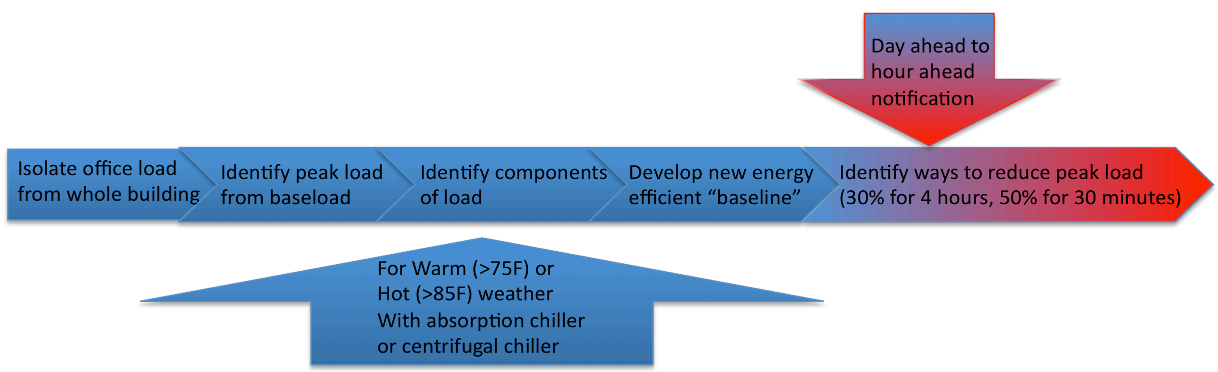

Figure 30: Basic components of the test plan.

By nature, demand response projects assume some discomfort and inconvenience on the part of the occupants. Nevertheless we wanted to monitor the occupants' comfort with thermal conditions, lighting, and air quality during the tests. The Center for the Built Environment (CBE) at UC Berkeley issued several online surveys to assess the occupants' satisfaction.

[7] In actual practice, development of demand response capability usually identifies opportunities for reducing energy use that can be implemented continuously for long-term energy efficiency. This efficiency is captured for its valuable GHG emission and cost reduction, creating a new baseline from which demand response can be implemented.

The facilities manager gave the final approval for any test, and needed advanced notice. The web interface required the proper syntax of commands to be prepared and eventually developed within the DR controller. The UC Berkeley/LBNL teams had the luxury of proximity. We conducted many smaller scale tests, especially to test the DR controller and confirm the matching of BACnet command with intended relay; some tests could be conducted initially at a smaller scale (e.g., only the fourth floor) or degree (e.g., change the zone temperature by 1F) or even at night (e.g., turn on and off lighting to check correct labeling of relays). Finally, we planned full-scale tests with both HVAC and lighting systems during typical peak periods (e.g., 2-6p). We thought we would decrease complaints from occupants if we notified them in advance, and asked for their participation.

7.3.1 DR algorithms and energy efficiency

The demand response algorithms included automatic control of the air-side HVAC system and the lighting system; a separate series of tests used the gateway to control plug loads automatically. During some tests we turned off some loads manually. We did not test strategies on the water-side of the HVAC system (e.g., reducing the chilled water temperature or increasing the condenser water temperature), although these were explored in simulation. This was precluded by operational requirements of the nanofabrication facility served by the same chiller. We also did not change the static air pressure.

	Energy Efficiency	*Demand Response*
HVAC		
Increase supply air temperature	Increase to 58°F	Increase to 60°F
Zone temperature	Install deadband of 70°F-74°F during day and to 65°F (heating) and 80°F (cooling) at night (10p-5a)	Increased to 78F on all floors except 7th floor (76F), allow “cool blast” (immediate short-term occupant-control)
Reduce ventilation rate	Reduce minimum ventilation rate by 30-70%; reduce by 70-85% at night. Use economizer more effectively. Use demand controlled ventilation in conference rooms.	Reduce minimum ventilation rate by 70-85% for short periods with air monitoring.
Increase condenser water temperature		(needs empirical testing—thought to be best for hot days)
*VFDs on chilled/ condenser water pumps	(automatically controlled)	(not pursued)
Lighting		
Daylighting		Turn off lights in daylit zones.
Open plan office dimmers	Timers on open plan office space, especially at night; web-based interface to dim lights	Dim light to 33-66% and allow user overrides
Corridors/core areas		Turn off all lighting except emergency

*Required new equipment

Table 2: Comparison of demand response and energy efficient strategies.

Details of the enhanced scale test plan may be found in Appendix M.

7.3.2 Adaptive Demand Response Framework

Siemens Corporate Research implemented an adaptive demand response framework that integrated central DR control (day ahead planning) with distributed DR control (real run-time) seamlessly.

With our centralized DR strategy on HVAC system, Siemens implemented to use building thermal mass as storage to shift the building load in order to reduce peak demand. Many pre-cooling strategies are designed with a schedule of global zone temperature settings, supply air temperature reset, minimum VAR Airflow setting during the day before DR event. The predicted load reduction is based on weather forecast.

However, the fluctuating nature of the electrical usage (e.g., subject to error in weather forecast and people's changing activities) increases the difficulty to achieve the load-shedding goal by simply applying pre-calculated control to centralized systems like the HVAC system and the central lighting system. The distributed loads can also contribute to compensate for the electrical usage variation. The adaptive control procedure is proposed to combine the contribution from both central load and distributed load reduction, which is controlling HVAC system with pre-calculated set points and compensating for the difference between the load-shedding goal and real-time power consumption with controlling plug-in loads and part of the central lighting according to occupancy status. The figure below shows an example scenario. Line 1 shows the baseline power consumption of a hypothetical building. Line 2 shows the predicted power consumption by the HVAC system based on the day-ahead simulation result to provide the fixed set points. The prediction is calculated based on weather forecasting and occupancy if available. But the HVAC power consumption may be different from the prediction on a particular day due to a change in weather; for example, the day might be much hotter than predicted and therefore the HVAC system would be more active more of the time which consumes more power and energy. The real-time power consumption by an HVAC is shown using Line 4. Therefore, the total power consumption of the building (shown by Line 5) may go beyond the demand response target (in this case, reducing peak load by 30% shown by Line 3). The proposed adaptive demand response scheme will be effective in such a scenario by controlling the distributed loads, such as plug loads and lighting loads; we assumed that the HVAC set points determined in a day ahead simulation will not be changed. The adaptive demand response will control the distributed loads so that the total power consumption is below the 30% peak load reduction target (see the highlighted part in the figure).

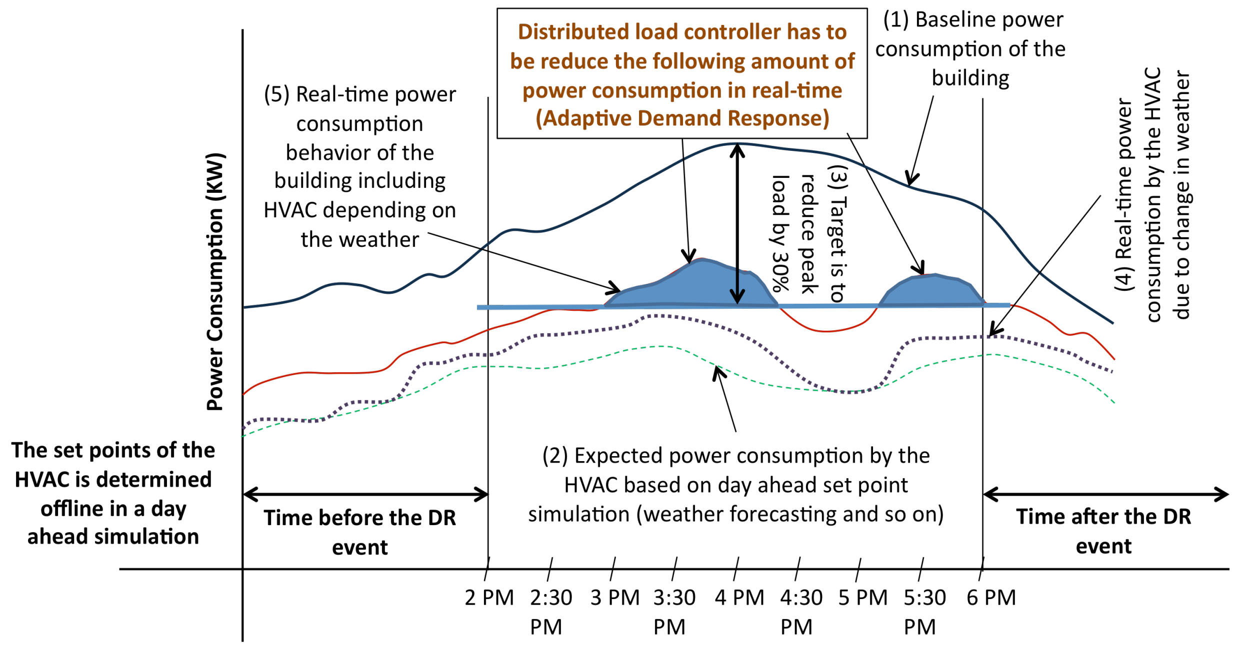

Figure 31: Objective of the Adaptive Demand Response based on distributed load control

The advantage of this framework is to utilize the nature of plug loads—their power consumption is closely related to people's changing activity and can be controlled in real time. Therefore, the plug loads are used to compensate for the gap between real-time power consumption and demand response target.

7.3.2.1 Control structure and sequence

In the proposed Adaptive Demand Response Framework, the entire controllable area is defined; for example a building has one central controller that will receive the Demand Response event signal from the utility. The controllable area is divided into multiple control zones. There is one distributed controller in each control zone. The plug loads in one control zone, for example an office room, are controlled by the distributed controller. All distributed controllers, central lighting system and the HVAC system communicate with the central controller. [Note that other frameworks may have distributed autonomous control]. The hierarchy is illustrated in the figure below.

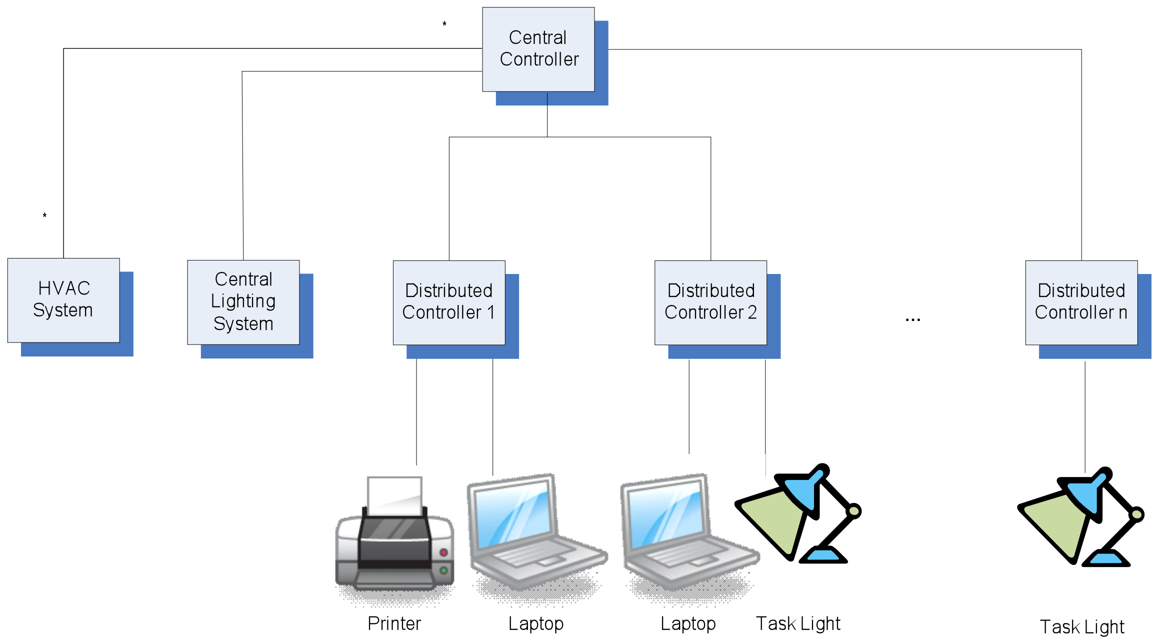

Figure 32: Control Hierarchy

The complete control procedure includes the following steps in the day before a DR event, and on the day of the DR event.

One day before DR event:
User preferences, such as temperature, lighting level, humidity, etc. and activity schedule information are gathered at each distributed controller. The distributed controllers send all information to the central controller. The central controller calculates the temperature set points of the HVAC system using an optimization algorithm. (The specific optimization algorithm is beyond the scope of this discussion).

On the day of DR event:
Step 1: Central control manages the HVAC system according to the temperature set points.

Step 2: Before the start of the DR event, the central controller reads from the sub meter and compares the real-time power consumption and required power consumption due to the DR Event.

Step 3: If the real-time power consumption is smaller than or equal to the required power consumption, the central controller will apply no control. Otherwise, the central controller controls Central Lighting directly and controls plug-in loads through the Distributed controller using the following steps:

- The central controller initiates a bid request for each distributed controller.
- Each Distributed controller will report Occupancy Sensor Information and plug load appliance information when receiving the bid request.

- Central controller will decide the power allocation and control commands for Central lighting and each Distributed controller. (The algorithm used by the central controller to allocate power and generate control commands to the distributed controllers is outside the scope of this discussion).

- For lighting, the central controller can determine the control based on the Occupancy Sensor Information, user lighting level preference and activity schedule. If the occupancy information shows nobody is in the control zone, the central controller will turn off the lights. Otherwise, the control command is created according to the user lighting level preference and activity schedule.
- At the start of the DR Event, the Central controller will apply control to the lighting and give the command to the Distributed controller to apply control to the plug load appliances.

Step 2 and Step 3 will repeat, for example every 15 minutes, during the DR Event. The interval can be chosen based on the fluctuation of the electrical usage.

Step 4: At the end of DR Event, the central lighting and plug loads will resume to the state before the DR event.

The figure below illustrates the sequence diagram of the procedure.

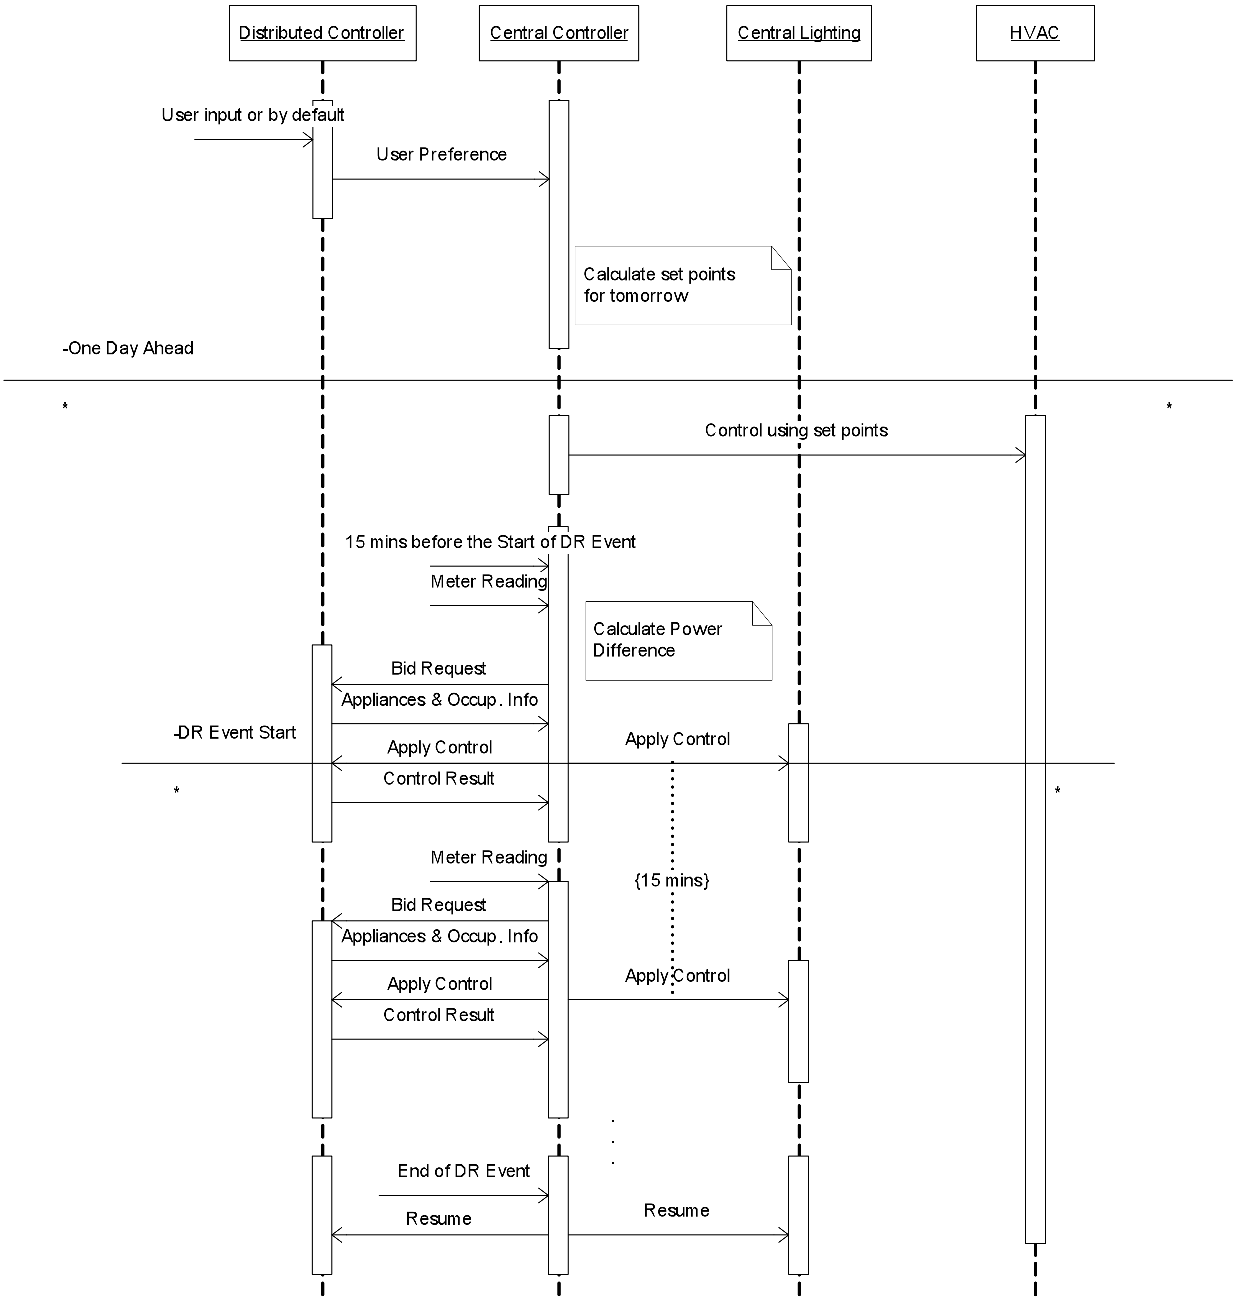

Figure 33: Sequence Diagram of the Control Procedure

7.4 Results

7.4.1 Energy Information Gateway Testing

At UC Berkeley, the gateway controlled both actual appliances (desk lamp, fan, heater, laptop, and Uninterruptable Power Supply with a desktop) using both wired (Raritan plugstrip) and wireless (ACme) monitoring receptacle controls as well as virtual smart appliances.

The following screenshots describe a sequence of events describing how the gateway responds to a demand response (DR) event

Figure 34 shows the gateway web user interface (UI), with the gateway running before an event signal is received. The table at the upper right shows that a Raritan plugstrip (SPS) with eight outlets is connected to the gateway. The SPS has five loads including: an uninterruptible power supply (UPS) on outlet 1, a fan on outlet 4, a desk lamp on outlet 5, a desk lamp with a compact fluorescent light on outlet 6, and a laptop on outlet 8. The bottom left shows how the web UI displays detailed information for resources. Here, a sensing resource for the number 1 outlet of the SPS is selected. The table shows the power is on ("state" is true), and the UPS is drawing 167 W. Here the total load is roughly 285 W.

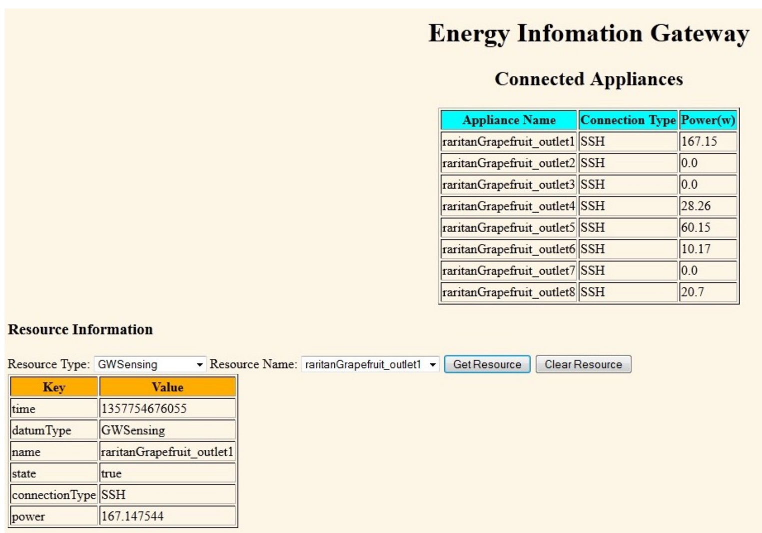

Energy Infomation Gateway

Connected Appliances

Appliance Name	Connection Type	Power(w)
raritanGrapefruit_outlet1	SSH	167.15
raritanGrapefruit_outlet2	SSH	0.0
raritanGrapefruit_outlet3	SSH	0.0
raritanGrapefruit_outlet4	SSH	28.26
raritanGrapefruit_outlet5	SSH	60.15
raritanGrapefruit_outlet6	SSH	10.17
raritanGrapefruit_outlet7	SSH	0.0
raritanGrapefruit_outlet8	SSH	20.7

Resource Information

Resource Type: GWSensing Resource Name: raritanGrapefruit_outlet1 Get Resource Clear Resource

Key	Value
time	1357754676055
datumType	GWSensing
name	raritanGrapefruit_outlet1
state	true
connectionType	SSH
power	167.147544

Figure 34: Web UI before event signal received

Figure 35 displays the bottom half of the web UI. Here, a DR event is being created to simulate a signal that might be sent from a DR server or other entity. An event is scheduled to start at 10:08 a.m. and to end at 10:12 a.m., with a goal of shedding 220 W. A binary actuator control scheme is selected for the Raritan SPS connected. This is a controller that selects one of two binary states (on or off) for the SPS outlets. Control is applied dependent on user pre-defined appliance ranking. It can be seen in the upper right; an event is scheduled and is inactive, meaning its start time has not passed.

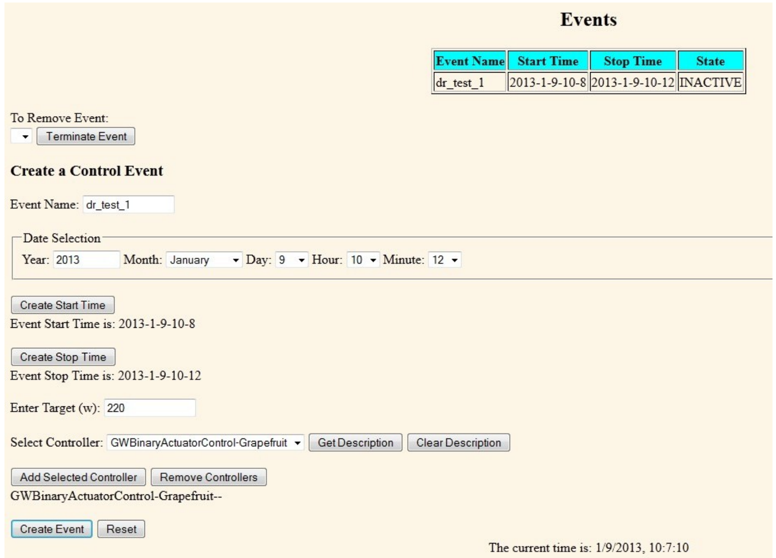

Figure 35: Creating a DR event to simulate receiving one from a server.

Figure 36 depicts the web UI during the DR event. It can be seen, at the bottom right, that the event is active. The appliance table shows that control of outlets 1 and 5 has been applied. The aforementioned binary controller switched the outlets to the off position. These outlets were controlled first as they were ranked low in importance. As the applied control resulted in the load reduction goal being met, the gateway returns to a passive state and no more control is applied. Here the total load is roughly 53 W.

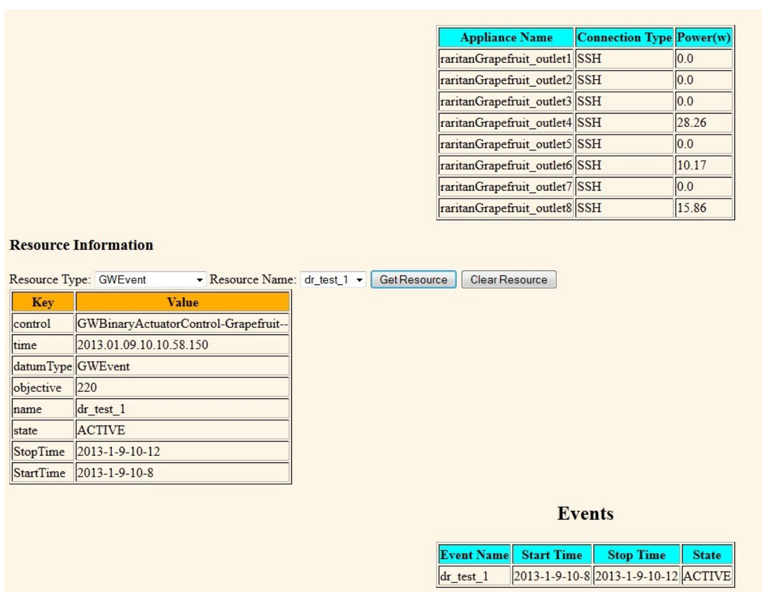

Figure 36: Web UI after control is applied during event

Figure 37 shows the UI after completion of the DR event. The event resource is shown in detail in the table to the middle left. Here, the event is shown to be over. The appliance table shows that the gateway has restored the outlets it previously controlled to their pre-event states, in the on position. It should be noted that the UPS, on outlet 1, is now drawing twice the power than before. This is due to the nature of the UPS battery charging algorithm.

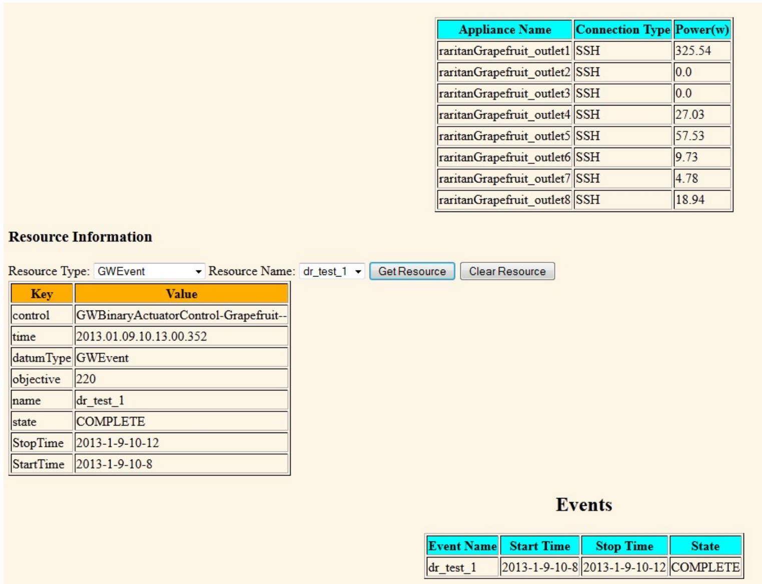

Figure 37: Web UI post event

7.4.2 SEB testing (Siemens)

A field test was conducted at CITRIS building, on August 17 and September 17, 2012. The HVAC system was controlled by SEB empowered with OSP optimization scheme. The conducted test included:

1. The control on CITRIS building management system was granted for the period from 9:00 AM and 3:00 PM. As a result, the GTA pre-cooling phase could not begin until 9:00 AM, and GTA reset could not reach 78 °F;
2. All water-side system control points and supply air static pressure were not controllable during the test. Therefore, we were unable to implement full optimal control strategy, which should have included both air- and water-side controls. Consequently, our HVAC control strategy only had GTA, SAT and Minimum Ventilation elements. And the OSP only contained 13 candidates;
3. Weather forecast of the test day was obtained. And the optimal control strategy was identified by simulation evaluation of each candidate strategy in OSP;
4. The optimal strategy was programmed in SmartEnergyBox;
5. Strategy was executed according to the schedule, from 9:00 AM to 3:00 PM.

7.4.2.1 Weather forecast

The following figure shows the weather forecast (Dry Bulb Temperature) for September 16~20. The test day was a mild day, with highest temperature about 65 °F.

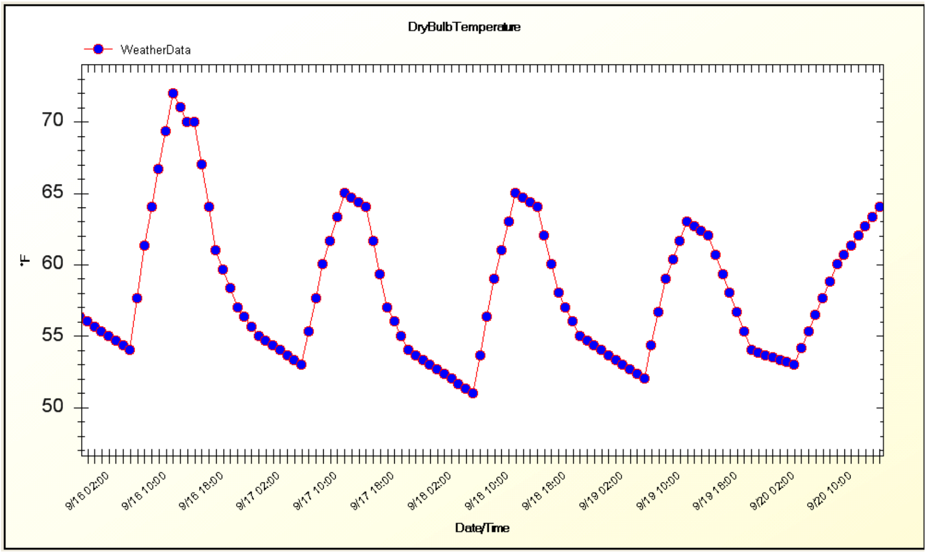

Figure 38: Weather forecast for September 16-19, 2012.

7.4.2.2 HVAC control optimization

13 candidate strategies in OSP were evaluated by the simulation model integrated in SEB. The simulated building HVAC energy consumption was plotted in the following figure. Strategy 7 was selected as the optimal strategy. Strategy 7 is characterized as GTA shown in the following figure, together with SAT=60 °F and Minimum Ventilation.

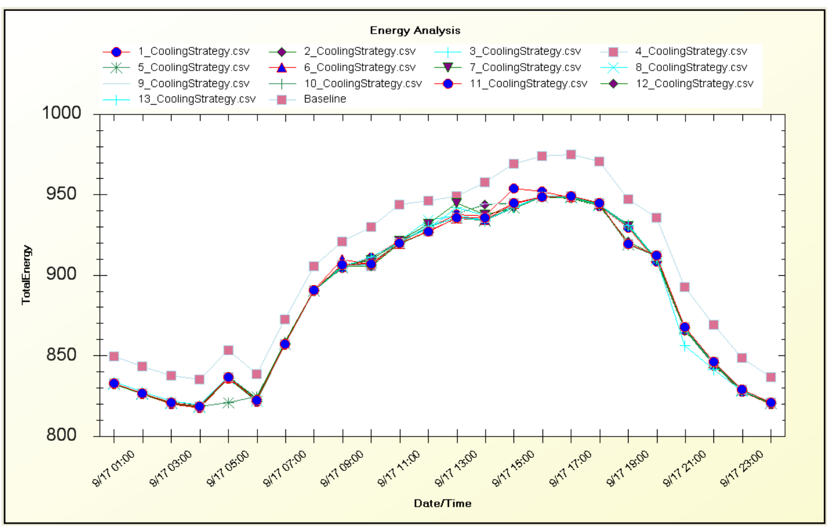

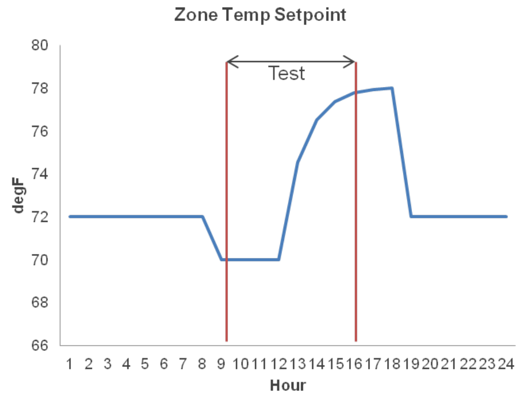

Figure 39: Setpoing control strategy.

Again, only the control for the period between 9:00 AM and 3:00 PM was tested.

7.4.2.3 Measured response

The response of the system was measured, and the data was accessed from UCB sMAP2.0 service. The following figure shows the supply and return air temperature in air handling unit (AH2B). SAT was successfully set at 60°F. And return air temperature, which, to some extent, represents the average zone air temperature across the controlled zones, was responding to our GTA settings.

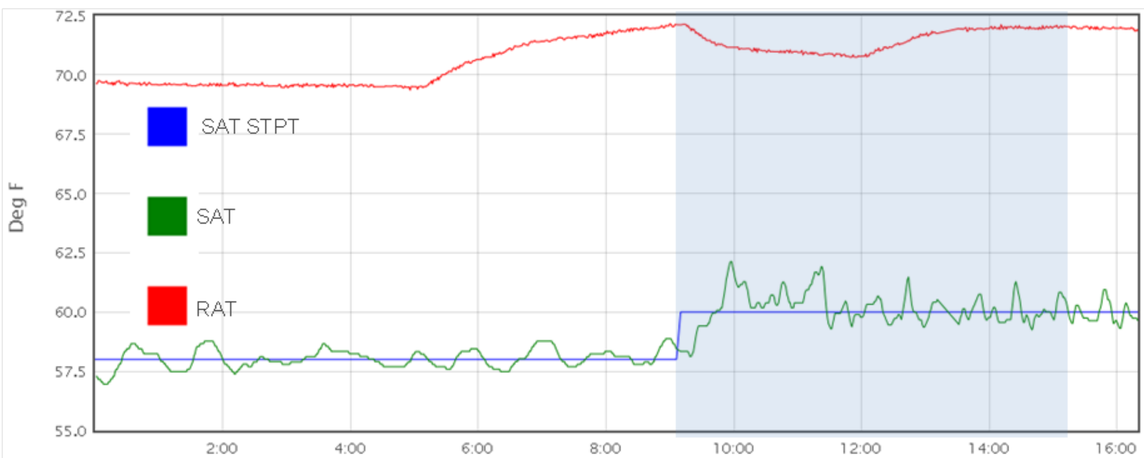

Figure 40: Temperature of air system.

The following figure shows the supply air volume flow rate of both air handlers. It is clear that during pre-cooling phase (9:00 to 12:00), more conditioned air was demanded, and less demands during the reset phase (12:00 to 15:00).

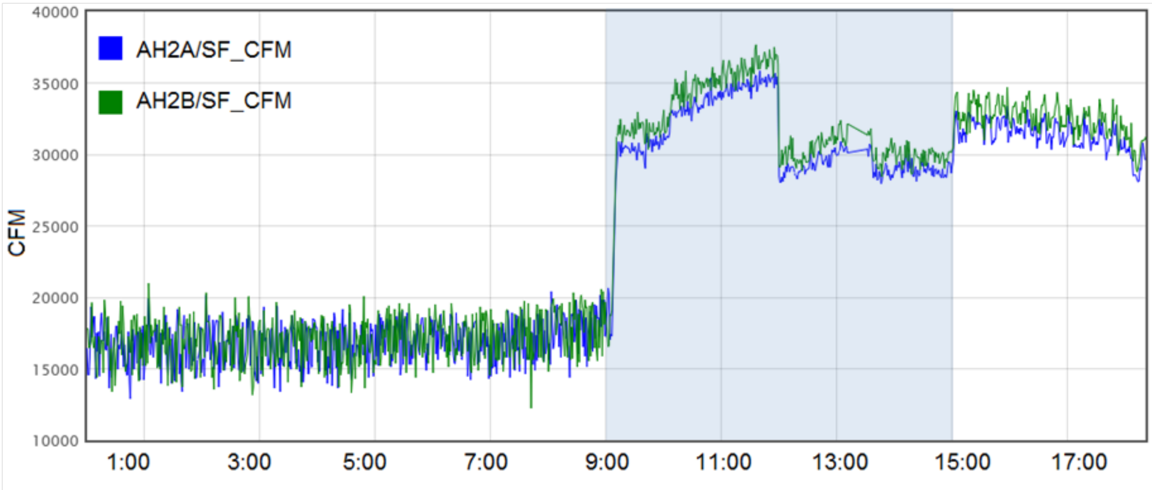

Figure 41: Supply air volume flow rate (shadow indicates the test period)

Zone temperature and ventilation rate were also monitored (the following figure). It shows that the zone temperature setpoint followed the GTA strategy that we designed.

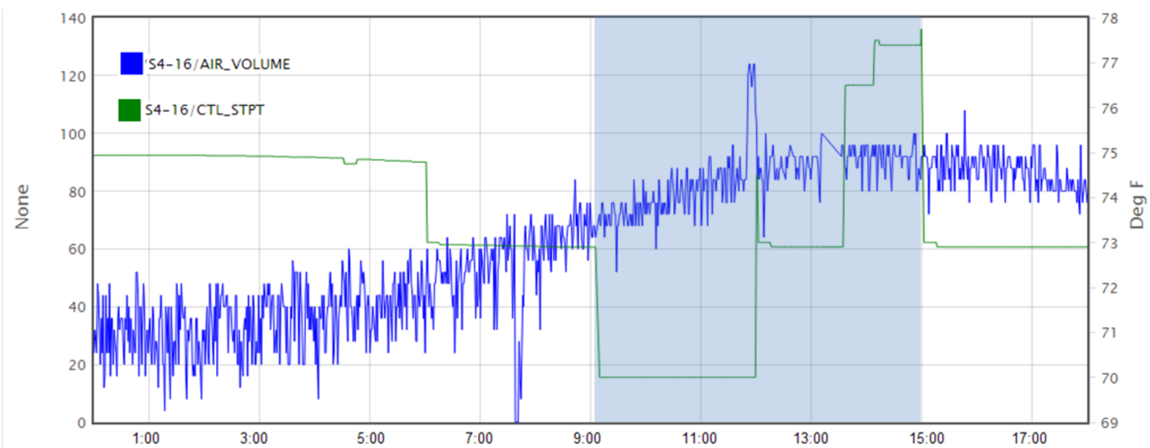

Figure 42: Control setpoint and air volume for one sample VAV zone on the fourth floor.

From the metered consumption of HVAC system (the following figure), we estimated about 5KW peak load reduction was achieved by our strategy.

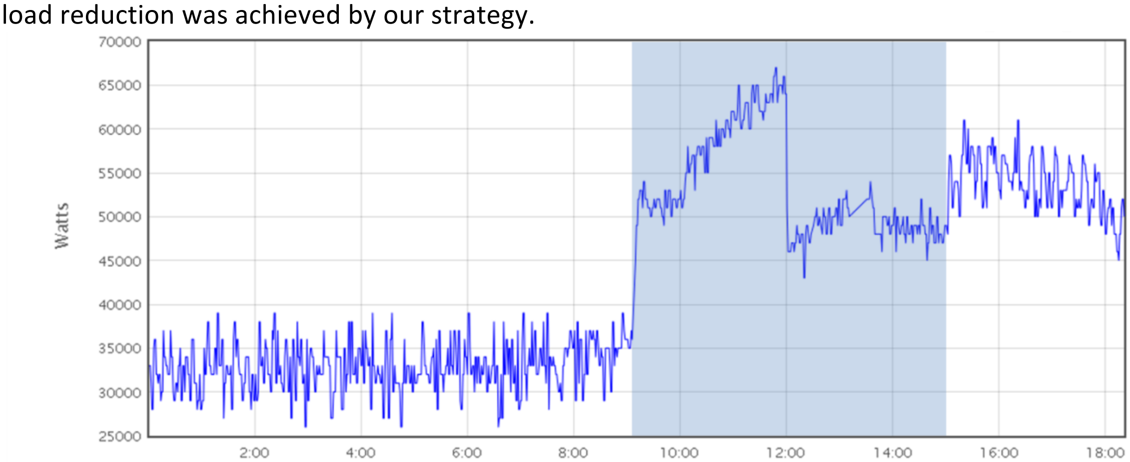

Figure 43: HVAC consumption during the test.

7.4.2.4 Conclusion

The following conclusion can be drawn:

1. Our two-stage optimization scheme (OSP optimization) successfully generated an optimal strategy on-line within relatively short time;
2. The strategy was programmed in SmartEnergyBox, and was successfully executed during the test;

Partially implementation of the strategy achieved certain level peak load reduction. However, due to some interruptions during test, and relatively cool outdoor condition, the peak load reduction might not seem very significant.

7.4.3 DR controller testing (UCB)

7.4.3.1 Load Reduction Testing

The Central Load-Shed Coordinator (CLSC) has undergone several tests to determine its efficacy in load-shedding for a hypothetical DR event. During tests, the CLSC controlled both lighting and HVAC in SDH, and networked with a handful of Energy Information Gateways (EIGs) for plug load control. The CLSC is run on a desktop computer in SDH. Tests involve several EIGs within an office, connected to a variety of plug loads. The small number of EIGs and connected plug loads means that plug loads are a negligible portion of potential load-shed, but these tests have proven the ability of the CLSC to coordinate multiple EIGs with building wide lighting and HVAC. Several custom strategies for each resource are designed by the authors prior to each test, representing a range of load-shedding options.

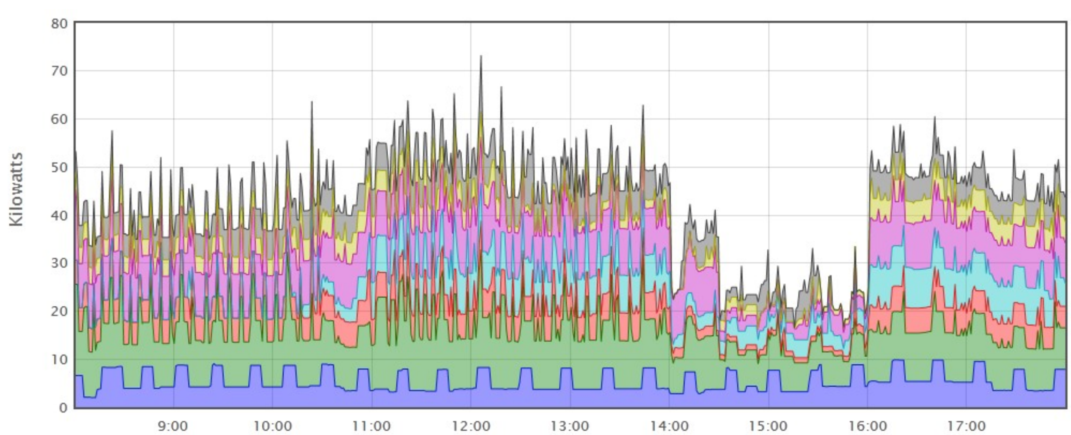

Figure 44: Aggregate lighting and local tankless water heater power for each floor in SDH, stacked vertically from 1 to 7.

Figure 44 displays the total lighting and local tankless water heater power for each floor in SDH[8]. During normal midday operation, the total building lighting power is between 45 and 55 kW with the exception

---

Local tankless electric water heaters for each restroom are on the same circuit as lighting. The spikes in this illustration are the water heating use. The spikes are suppressed and not as prominent in 15-minute average demand analysis typically used for energy billing and demand response savings accounting.[8]

It can be seen, at 2 p.m., immediately after the start of the DR event, the total lighting load drops significantly, only to rise shortly thereafter. This immediate rise is surmised to be in part from occupants turning their local lights back on. The CLSC, through querying sMAP, learned of lighting zones that were on after the initial command was issued. It then reissued as appropriate commands to switch the lighting zones off. After the initial rise, the total lighting power drops again and settles at roughly 25 kW. This is an approximately 50% drop in total lighting power relative to before the DR event.

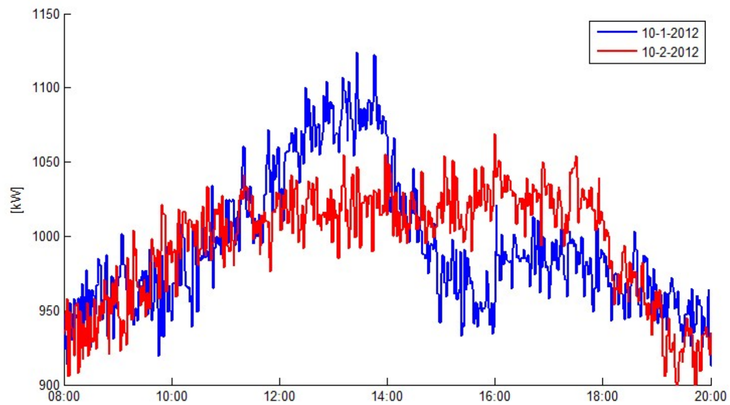

Figure 45: Total SDH building electrical load for October 1 (red) and October 2 (blue).

Figure 45 shows the total building power from the test day of October 1 (blue), and the next day October (red), which is used as a baseline (normal operation) due to the very similar weather. The second strategy used in this test was an HVAC precooling, wherein the entire building would be set to cool to 70°F from 10 a.m. to 2 p.m., at which time the building would be set to cooling mode with a set point of 76°F until 6 p.m. Essentially this strategy cools the building when energy is cheap, and only cools when necessary when electricity is expensive. This time-shift of building load is easily seen in Figure 45, where the total load for the test day is higher than normal until the start of the event and lower until the end of the event. When both the lighting and HVAC strategy are activated the maximum shed is roughly 70 kW around 3:30 p.m. When only the HVAC strategy is active the load-shed is roughly 50 kW. These promising results demonstrate the effectiveness of the CLSC as a load-shed coordinator between the three resources of lighting, HVAC and plug loads. This illustrates ability to generalize actuation of these systems, and choose how to meet a load-shed while minimizing a measure of the negative effects of doing so.

7.4.3.2 Carbon Dioxide Concentration Control Algorithm and Testing

An important parameter for which the CLSC utilizes closed-loop control is CO_2 concentrations within private offices and conference rooms. A caveat of DR is that reducing building ventilation rates to save power reduces air turnover and thus CO_2 concentrations may be significantly higher than normal, especially in enclosed spaces. Maintaining a predictable CO2 level is often difficult if not impossible with reduced ventilation levels often used in HVAC based DR strategies. Therefore, the CLSC employs a threshold-based controller with hysteresis to prevent unacceptable levels. This control logic is based in the HM and supersedes control of HVAC when necessary. When the CO_2 concentration reaches 700 ppm, below ASHRAE allowable maximum, the ventilation rate is set to 100% of the maximum flow for the corresponding area. When the CO_2 concentration reaches 550 ppm, the ventilation rate is set to its previous value. (See Appendix Q for background on ASHRAE standards and a similar strategy applied to conference rooms throughout the building).

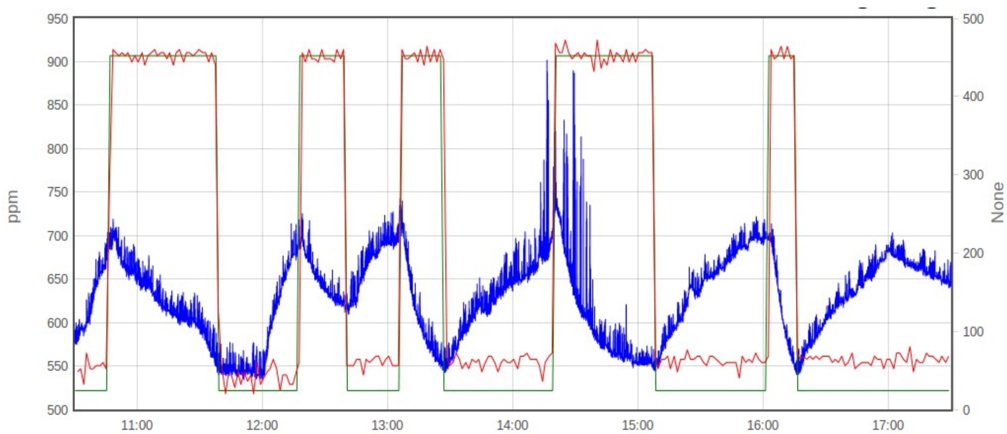

Figure 46: Results from running CO_2 control algorithm on November 13, 2012 with both one and two people in a private office.

As shown by Figure 46 the control strategy for CO_2 is successful in preventing the concentration from reaching the ASHRAE allowable maximum. In this test, the upper threshold was set at 700 ppm and the lower threshold at 550 ppm, due to the lengthy time for the concentration to reach the lower threshold even when using the maximum ventilation rate of 452 CFM. The lower ventilation rate was set in the HVAC controller as 25 CFM, however the real rate hovers above 50 CFM. The maximum high ventilation rate was chosen for fastest air turnover such that many cycles could be completed. In a real deployment of such a system, the ventilation rate may be adaptive to account for the number of people inside the room (if known) and for the current concentration.

7.4.4 UC Berkeley/LBNL Field tests

As mentioned in the previous section, we implemented several demand response event tests in 2012; we report here on the August 2, October 1, and October 18 events. For all three events, the building was using the absorption chiller. From the simulations, we expected to see most savings on the hottest day—the goal of the DIADR project was to curtail load by 30% on hot days[9]; we expected less curtailment on cooler days. By chance, these three events captured a cool day, a warm day, and a very hot day. The data are from sMAP but processed in Excel to provide 15-minute average data.

7.4.4.1 August 2 event

August 2 was the first DR test to curtail HVAC, lighting and plugloads at the same time using the web-based interface. The zone temperature of the VAVs in the office portion of the building was set to cool to 70°F from 11:30a to 2:30p to precool the zones. The supply air temperature was raised to 60°F from 58°F, zone temperatures set to 76°F, and minimum ventilation rate reduced by 70% from 2:30 to 6p. To minimize the more tangible discomfort to the occupants, we decided to limit the lighting curtailment to 2-4p. The lighting levels in the open plan areas were dimmed to the lowest level (one 30 watt fluorescent lamp in a two ballast-three lamp fixture), all possible corridor and lobby lighting were turned

[9] Per section 7.3, Berkeley temperatures in the 80s (deg F) during California heat storms are equivalent to the originally envisioned test temperature range (>95 F) for inland locations.

off, and all lighting in daylit zones (e.g., east and west zones of the open plan office space) were turned off. An email went out to all occupants to notify them of the event and to request that they reduce load from 2-4p (see figure 47 below). They were allowed to request a cool "blast" from the temperature web-based control (thisroom.is), as well as higher light levels from another web-based control. Five of the six large LCD monitors (four drawing 180 watts, but two drawing 600 watts) in the lobbies were manually turned off as were many of the (26) 6kW tankless electric water heaters (which are submetered with the lighting).

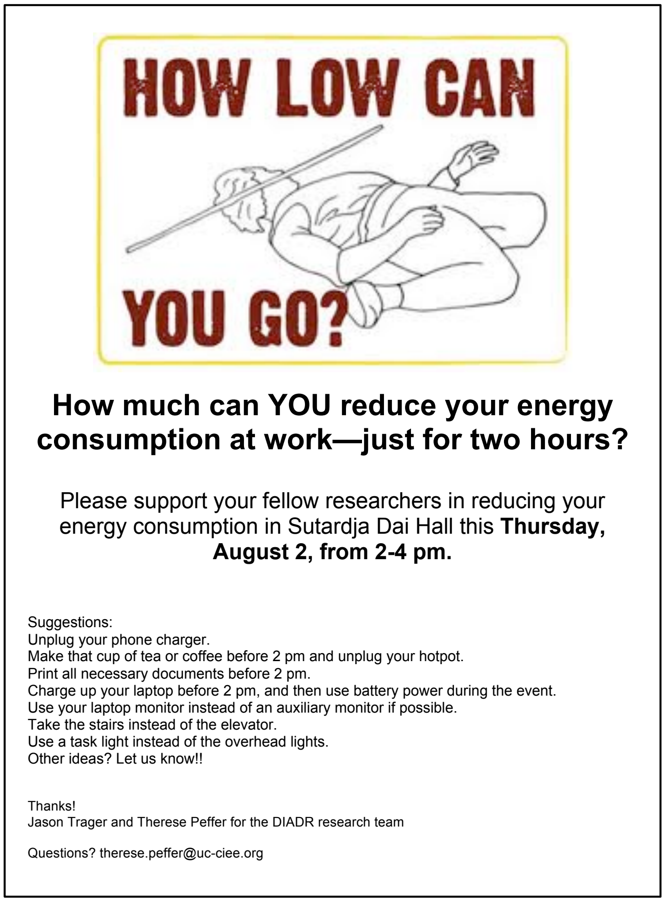

Figure 47: Flyer emailed to occupants and posted at most entrances to each floor.

Because there have been so many changes to the building over time, we decided the best "baseline" would be the day before (Aug 1) to compare the actual load with the predicted load in order to determine the curtailment. August 2 was classified as a cool day, with peak temperature just above 70°F. Thus the load shed from cooling[10] was not large, but still effective (an estimated equivalent of 5.7 kW in chiller cooling[11]). All figures below show 15-minute average data.

[10] While we have flow meters to determine the chilled water flow to the office, we did not have the appropriate temperature sensors in place to determine the portion of the cooling load for the office. An initial calculation by Rongxin Yin using the temperature sensor display at the meters indicated about 25% of the total chilled water flow went to the office portion of the building. However, a recent calculation by Bin Chen indicated that the flow may be as much as 40% during warm weather. Bin Chen calculated the office cooling load in BTU/h using the supply air temperature (SAT) and mixed air

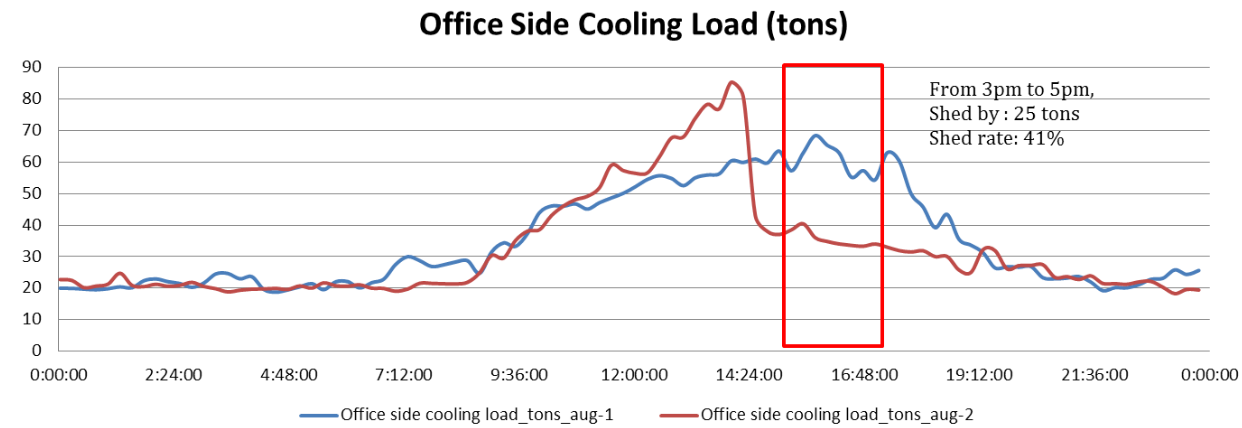

Figure 48: Aug 2 estimated cooling load shed compared to baseline.

The curtailed load from the air handling units was substantial—nearly 27 kW. The precooling caused an increase in the AHU load before the curtailment; while we still achieved the goal of reduced peak, we'd like to improve this performance.

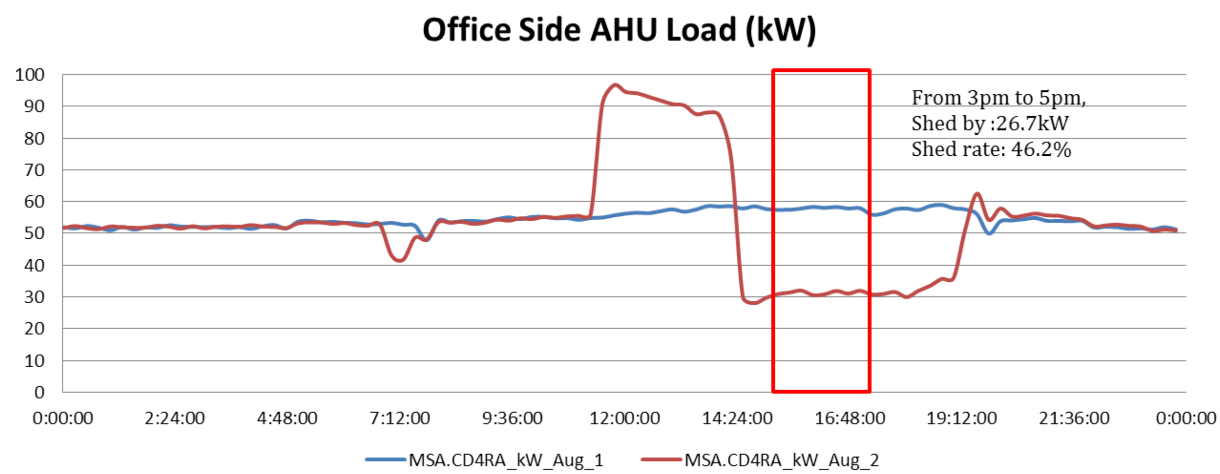

Figure 49: Aug 2 office AHU load shed compared to baseline.

The lighting load provided a curtailment of 17 kW, as shown below. We note that in early August, the building is in summer schedule; the lighting load seems to peak between 11a-1p, and gradually decline.

---

temperature (MAT) for the two air handling units as follows: Office Cooling load (Btu/h) = sensible load + latent load => H total = H sensible + H latent = 1.08 * Air Volume (CFM) * (delta temp F = MAT - SAT) + 0.68 * Air Volume (CFM) * (humidity ratio difference) (We assume the humidity ratio difference is 0).

[11] Bin Chen calculated the total Cooling Load assuming 40% of the chilled water flow was going to the office and using a COP of 0.75 for the absorption chiller as follows
Total cooling load = Office Side Cooling load / 40% --> Abs chiller load = total cooling load/ COP

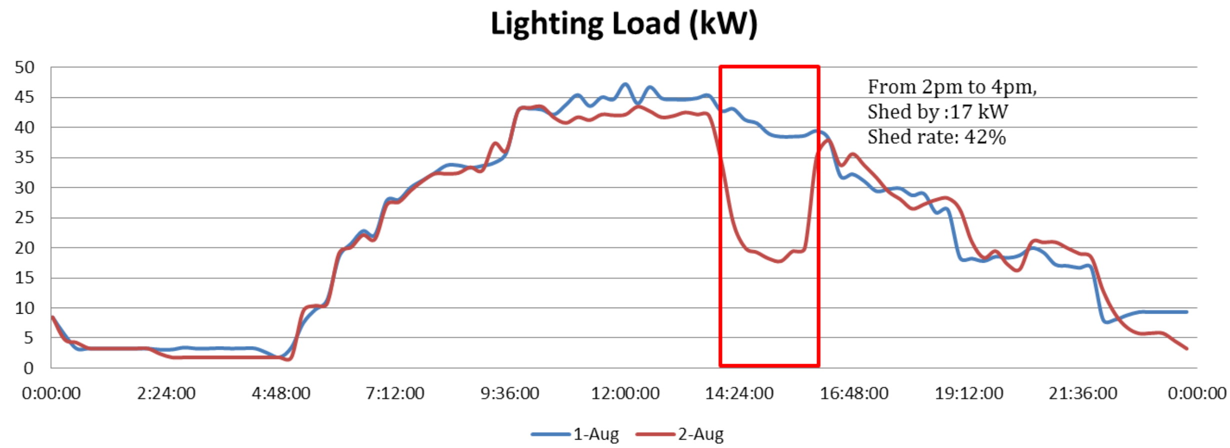

Figure 50: Aug 2 lighting load shed compared to baseline.

The curtailment of the load due to the appliances (also called receptacle or plugloads) was minor—only about 3 kW. We see a similar pattern as the lighting load—peaking at 11a-1p. The café on the second floor, one of the largest receptacle loads, closes around 3p in the summer.

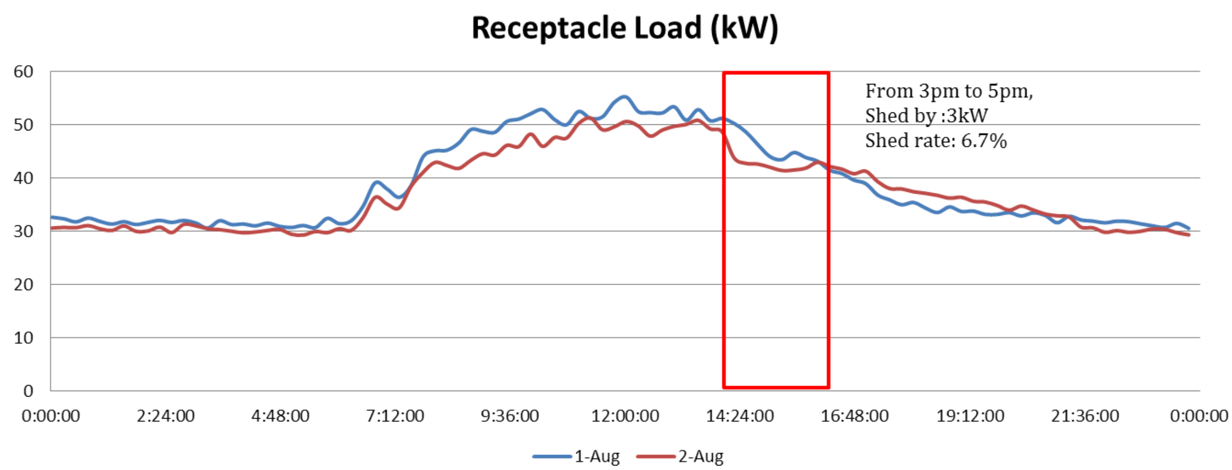

Figure 51: Aug 2 receptacle or plug load shed compared to baseline.

The load shed from the absorption chiller and chilled water and condenser water pumps was minor—perhaps 0.5 kW.

In summary, we estimated at least 47 kW load was shed as a result of the DR event, as well as slight reduction in steam consumption from the reduction in the absorption chiller load. Our estimate of the baseline load for the office for this day is approximately 258 kW. Thus the curtailment or shed is about 18%. We note that while this day was not hot in Berkeley, it was declared a DR event day by the local utility, Pacific Gas & Electric.

7.4.4.2 October 1 event

On October 1, we used the DR controller to automatically start the precool event as well as control the HVAC and lighting systems. The strategies were similar to the Aug 2 test, except we raised the zone temperatures to 78°F (instead of 76°F) and we were able to turn off lighting in a few more of the core areas (corridors and lobbies). The same notice went out to occupants, and the LCD monitors were turned off as well as many of the water heaters.

This test day was fortuitous in that the weather was unusually hot for Berkeley—94°F—for two weekdays in a row; we were able to use Oct 2 as the baseline day. Because of the hot weather, we hoped to see much more load shed from the cooling system. The figure below shows that indeed, the cooling load curtailment was more than three times the curtailment in cool weather (the previous event)—78.5 tons (equivalent to 33.5 kW of electricity).

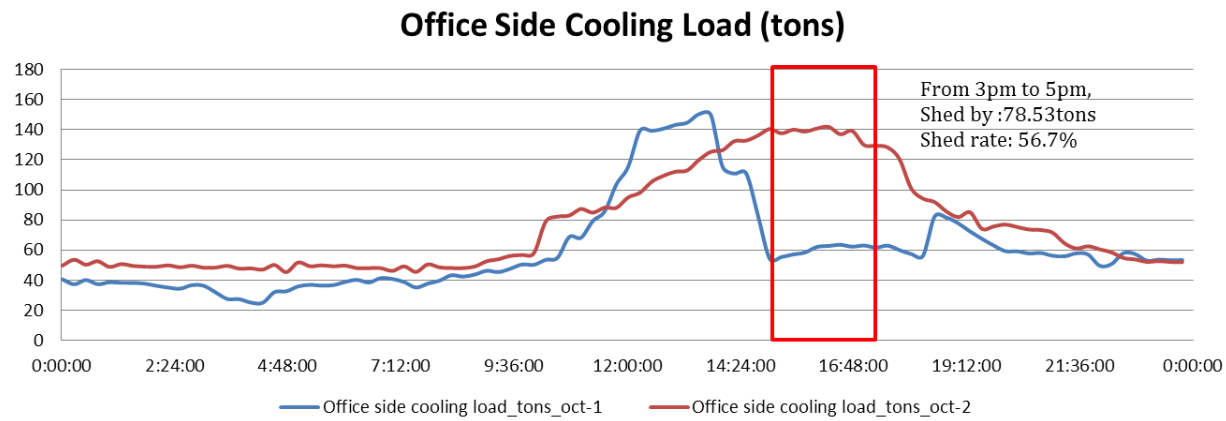

Figure 52: Oct 1 estimated office cooling load shed compared to baseline.

The reduction in the load from the air handling units was greater too—34 kW (versus 27 kW on the cool day). See figure below.

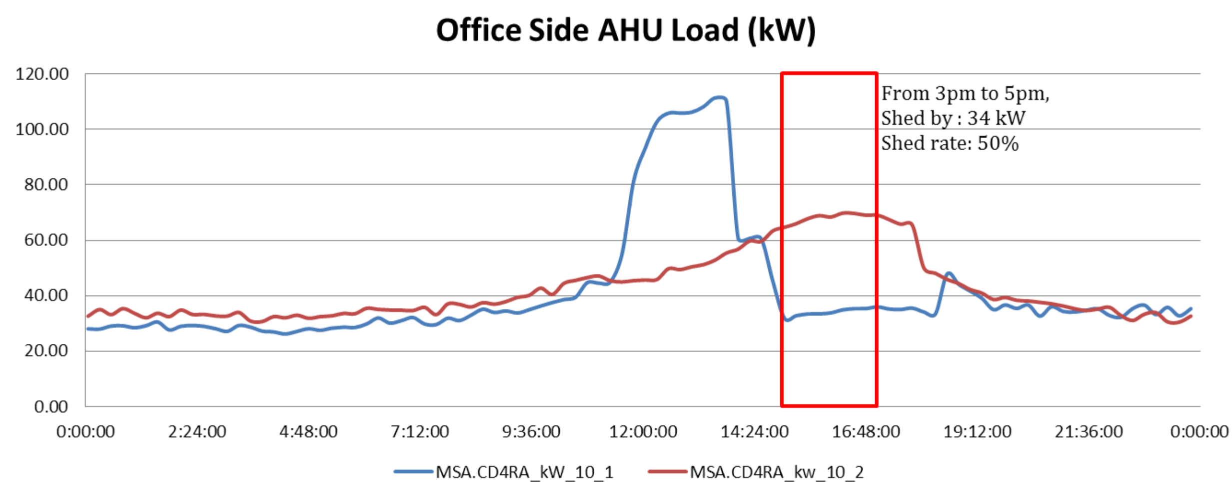

Figure 53: Oct 1 office AHU load shed compared to baseline.

The lighting load curtailment increased slightly to 20 kW from 17 kW; the shed is calculated over the two hour period. While we were able to turn off a few more lights, much of this increase is due to an increase in the baseload, since school was in session, and lighting levels were pretty steady during the DR event period. We discovered that some of the lighting commands did not go through the first time—perhaps a limitation of how many BACnet commands can be processed at the same time. The DR controller was designed to automatically resend the commands every 15 minutes; we noticed that the lighting was not curtailed until 2:30p. If we consider the period from 3-4p, the load shed was about 25 kW.

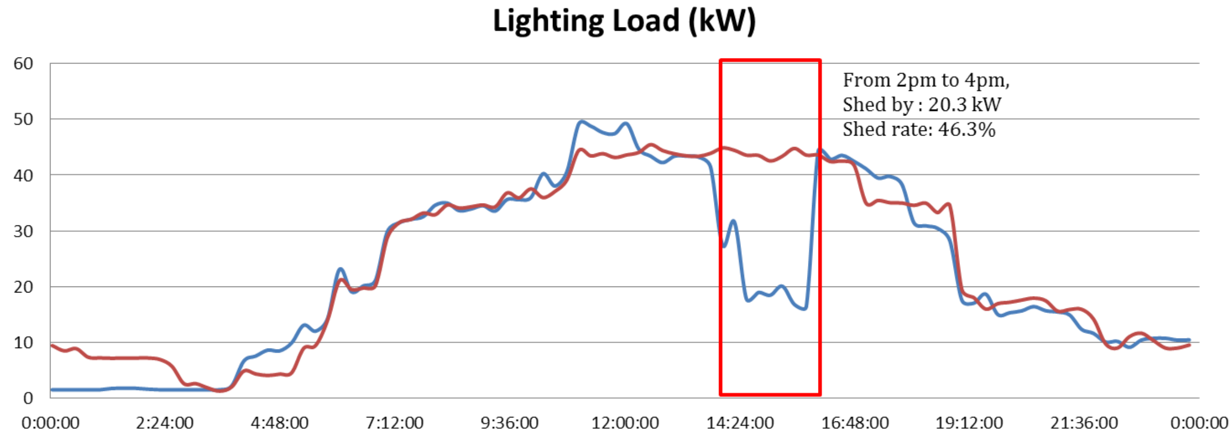

Figure 54: Oct 1 office lighting load shed compared to baseline.

A new submeter installed in September allowed us to see the cooling tower load more directly for the first time. The system has two cooling towers that are staged, so the second one only runs as needed (e.g., on hot days). Indeed, we were able to see an average shed of 7.3 kW (see figure below); peak shed was nearly 11 kW.

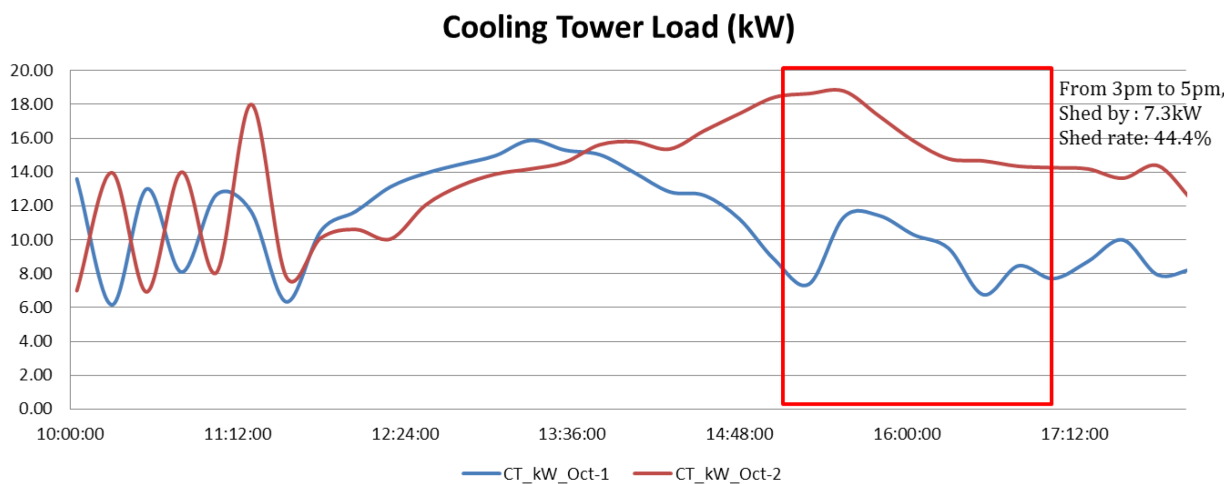

Figure 55: Oct 1 cooling tower load shed compared to baseline.

The load shed from the absorption chiller and associated pumps was approximately 5 kW.

The previous calculations of the contribution of each component to the load shed show the average of the load shed over a certain period of time. However, in just viewing the period from 3-4p, we saw the greatest load shed of 75 kW. Our best estimate of the total peak electrical office load for the baseline is about 300 kW. Thus the shed is approximately 25% from the peak electrical load.

7.4.4.3 October 18 event

On October 18, we again used the DR controller to automatically start the precool event as well as control the HVAC and lighting systems. The strategies were similar to the Oct 1 test, except we kept the zone temperatures on the 7th floor to 76°F instead of 78°F; in the last test we found that the temperatures on the 7th floor rose to 78°F rather quickly, most likely due to the windows on the east, south and west facades of this floor. We did not turn off the LCD monitors nor turn off the water heaters. We also did not override the lighting controls as in previous event days. We used October 17 as the baseline day; both days were categorized as warm with outdoor temperatures in the mid 80s.

As expected, the cooling load shed was between that of the cool and hot DR days, with a reduction of 47 tons of cooling (equivalent to 20 kW).

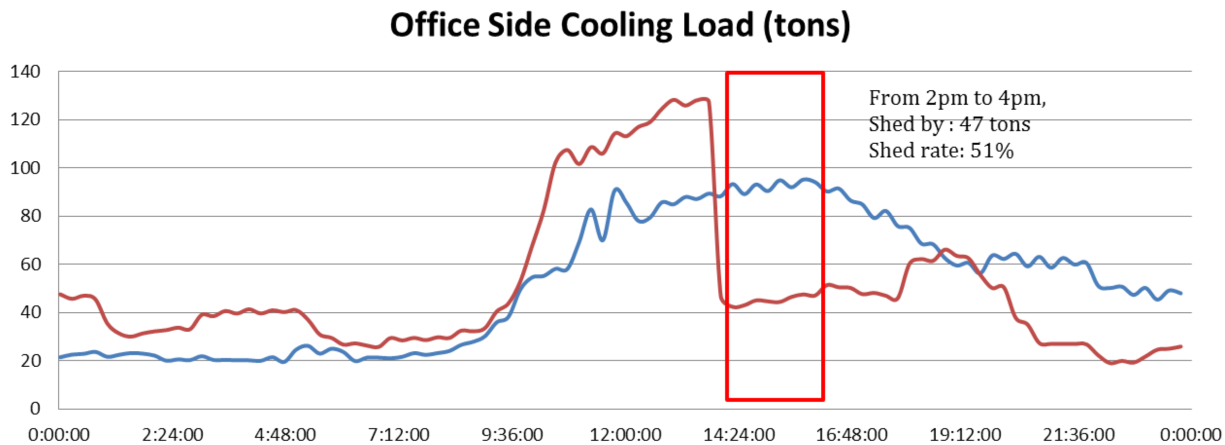

Figure 56: Oct 18 estimated office cooling load compared to baseline.

The curtailment from the air handling units was 21 kW, less than expected. While the total consumption during the event is about the same as on August 2 at approximately 31 kW, we note that the baseline on Oct 17 is lower—approximately 53 kW compared to 58 kW baseline on Aug 2).

After some investigation, we discovered this is due to an experiment run by Andrew Krioukov beginning on Sept 8; this test was supposed to be suspended for the Oct 17 baseline, but apparently was not. In this test, minimum ventilation rates were substantially reduced (5-30% of original) and controlled dynamically based on expected occupancy schedule and zone temperature. On October 1-2, apparently the increase in zone temperatures driven by the hot weather increased the minimum ventilation rate, but not so on October 17-18.

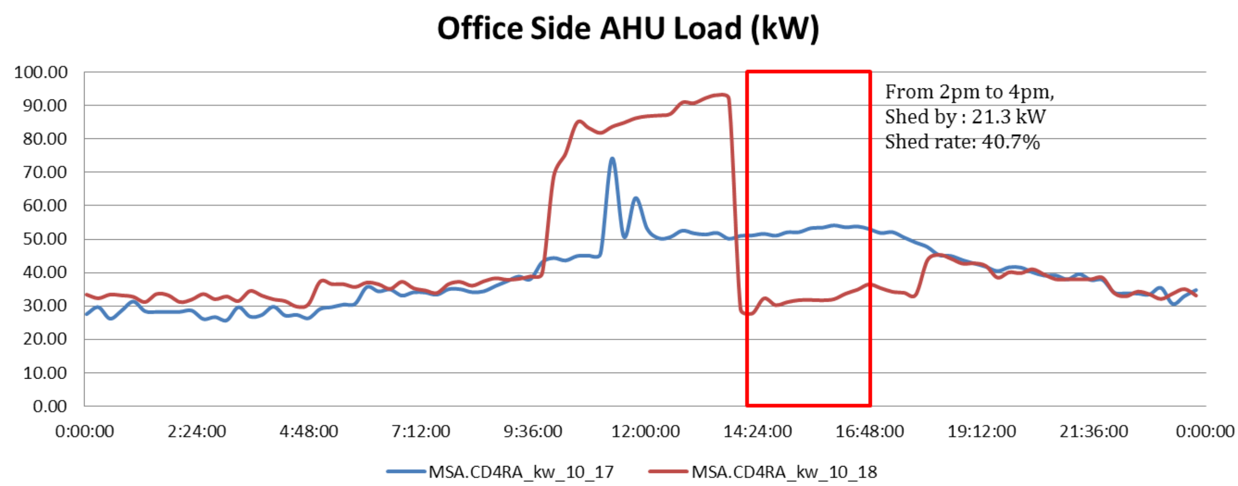

Figure 57: Oct 18 AHU load shed compared to baseline.

The lighting load curtailment was less than expected at 15.3 kW. In both previous DR events, we achieved a reduced lighting load of around 20 kW; however, for this event the lowest load is approximately 27 kW. The water heaters were not turned off for this event, but we would expect the short duration of the 6 kW water heaters would not be seen in the 15-minute average data. We think that the occupants overrode the light switches more, perhaps a bit fatigued by the DR events.

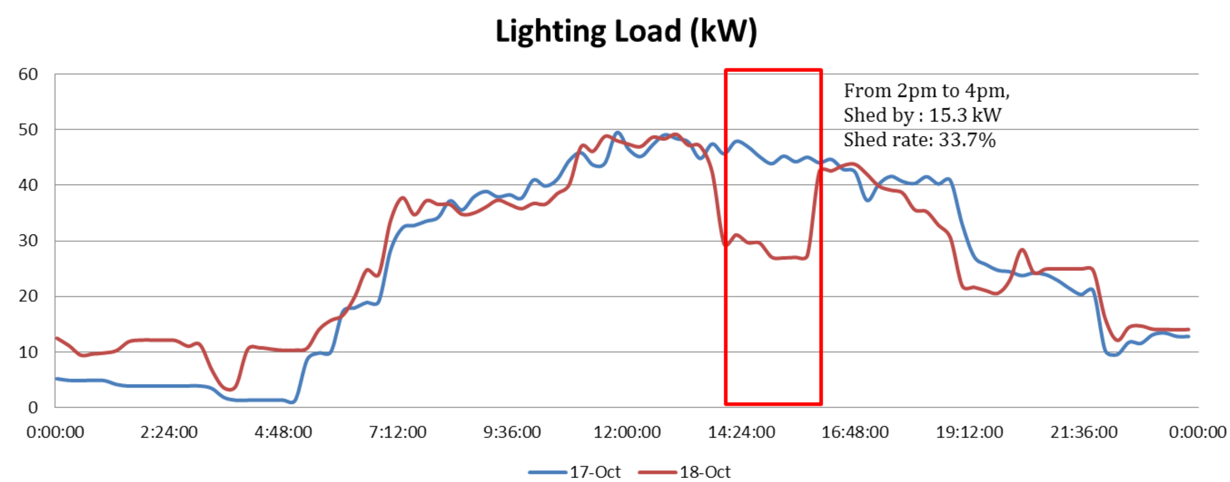

Figure 58: Oct 18 lighting load shed compared to baseline.

The cooling tower load was not curtailed by much, as shown below. This was not unexpected, since the weather was not hot.

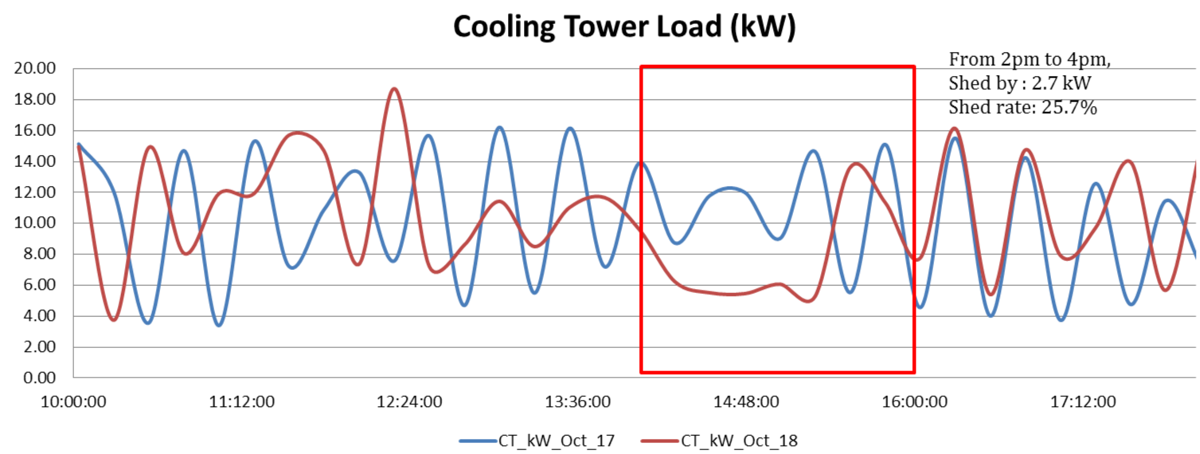

Figure 59: Oct 18 estimated cooling tower load shed compared to baseline.

The load shed from the absorption chiller and pumps hit a high of 35 kW for about one hour. We have not analyzed this, but note that the VFDs were installed prior to this event.

The total load curtailment was approximately 39.3 kW; if we compare the AHU curtailment against the Aug 1 baseline, this increases to 45 kW. This provides a 14-16% average peak load reduction from the baseline electrical peak load of 273 kW. If we include the peak curtailment including the extreme shed from the chiller, this increases to 80 kW, for a 29% peak load shed.

A summary graphic is provided below:

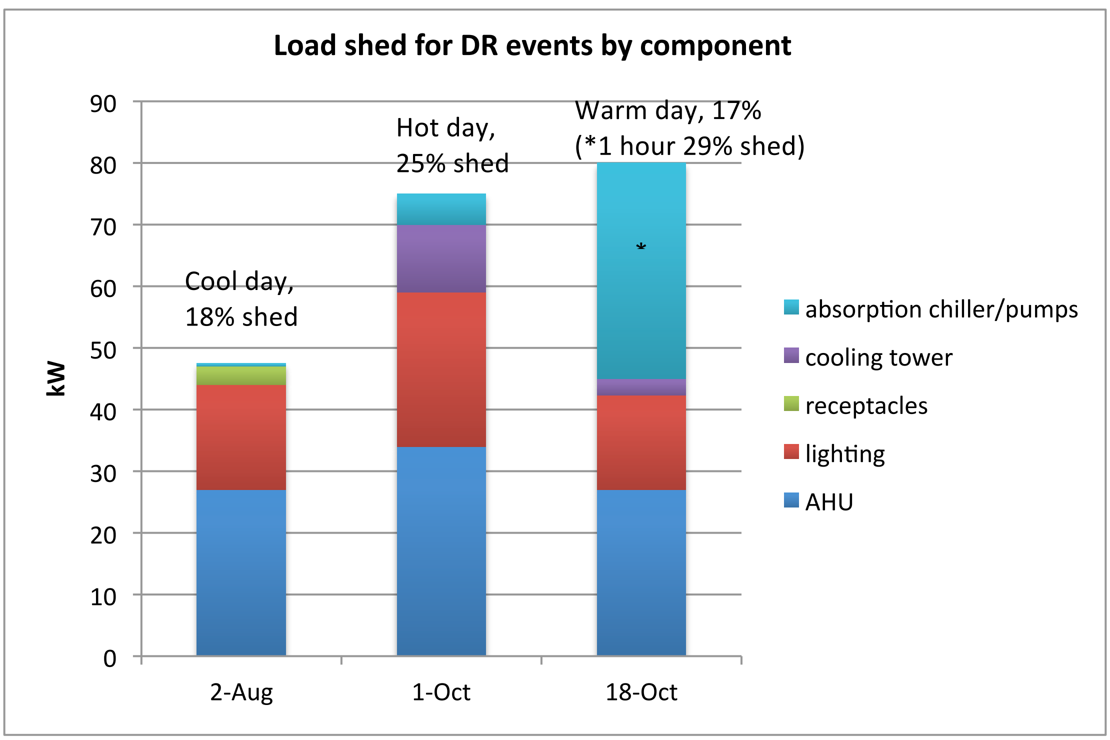

Figure 60: Summary of load shed for DR event days.

7.4.4.4 Discussion

We achieved 14-24% average peak load reduction during our DR events using the absorption chiller, which uses steam for cooling and thus consumes far less electricity than the centrifugal chiller. This scenario is realistic for this building, that is, the absorption chiller was designed to run in the warm months of the year when steam demand is low on campus, which is when most demand response events would occur. However, we were curious what might have been the savings for an office building with a more typical centrifugal chiller. Bin Chen ran this analysis (described in Appendix N) and developed the graphs in Figures 61 and 62.

If the centrifugal chiller had been used on the hot day, Oct 1, 2012, the reduced cooling load would have been an estimated 116 kW. Added to the other office load curtailment, we would have achieved a 30% electrical peak load reduction. In the same way, if the centrifugal chiller had been used on the warm day, Oct 18, the load shed would have been approximately 66 kW, resulting in a total shed of approximately 26%.

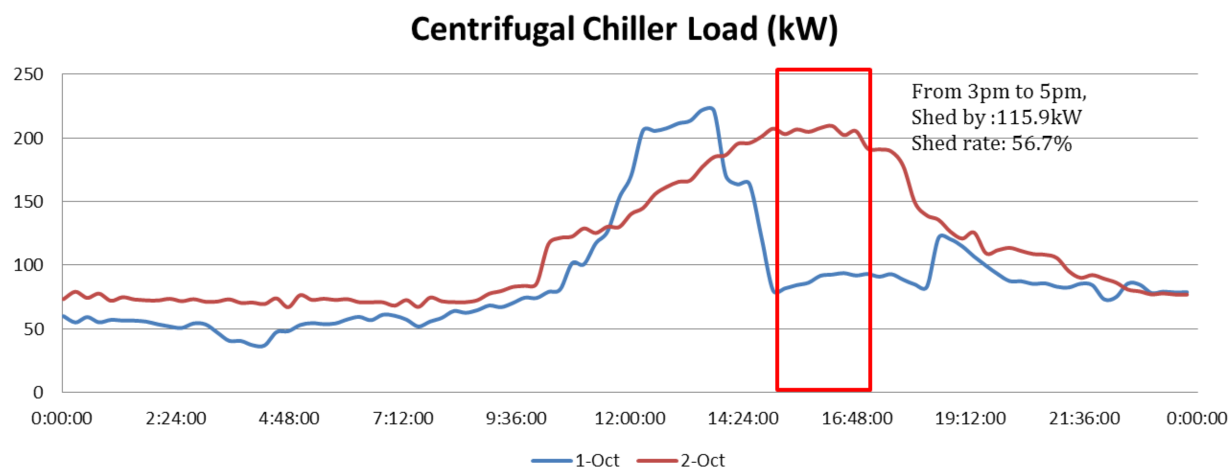

Figure 61: Hypothetical estimated cooling load shed had the centrifugal chiller been used.

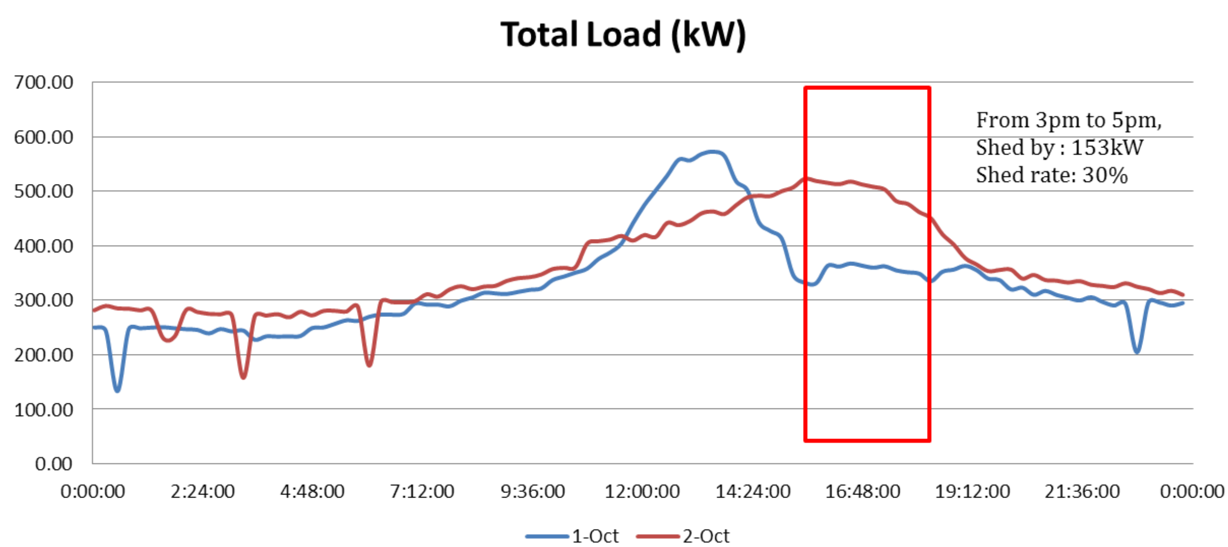

Figure 62: Hypothetical office load shed using the centrifugal chiller on the hot DR day.

7.4.4.5 Energy efficiency measures

Although we have not conducted a detailed analysis of the annual energy savings from the energy efficiency measures, we have calculated some approximate numbers. The load on the air handling units supplying the office portion of the building dropped by 20 kW over the course of the project, mostly due to the dynamic minimum ventilation scheme. The load on the absorption chiller and associated pumps decreased by about 30 kW, probably primarily due to the VFDs installed on the pumps, but some effect from the implementation of the temperature deadband (e.g., control points at 70-74F versus a single control point of 70F), increase of the supply air temperature to 58F from 56F, and the reduced ventilation. This amounts to approximately $44k savings annually. However, over the course of the project, the fab lab process tools load increased by 70 kW, the fab lab mechanical load increased by 15 kW, the addition of offices and labs on the first floor added 16 kW, and the telecom system added about 3 kW—a little over 100 kW total increased load.

There are both additional efficiency opportunities and demand response strategies that were not explored by this project, but will be pursued pending more coordination with the nanofabrication facility in the building (see Table 2 and Section 7.6).

7.5 Commercialization plan/cost analysis

Throughout the project, we utilized many different technologies, such as sMAP, StreamFS, the Rapid Audit Protocol, baselining techniques, the Smart Energy Box, Gateway, and the assessment tool. sMAP will remain an open source protocol and platform; StreamFS is also open source, and is already being used by several companies. The RAP has morphed into part of the Mobile Energy Lens (Ortiz 2012).

SCR focused on three technologies to document a formal commercialization plan: the Smart Energy Box, the Distributed Gateway Controller and the Demand Response Capacity Assessment & operation Assistance Tool (DRCAOT). Siemens Building Technologies followed established internal procedures to develop commercial products and performed commercialization plan for each product concept.

Through market evaluation study, it is confirmed that total spending for ADR deployment is expected to grow substantially. After studying how automated demand response (ADR) can be deployed in a typical commercial office building, Siemens has defined the specific market requirements for Smart Energy Box and Distributed Gateway. Since Siemens Building Technologies has not yet been in the business of selling design software tools, the delivery model for delivering such design software tool to end-user customers could be different than traditional building automation system software.

Siemens Building Technology also conducted technical feasibility and technical evaluation on the three product concepts. It is found that the time and effort of energy modeling is significant which could prohibit the successful launching of SEB and DRCAOT as a commercial product. In practical, the effort for setting up the simulation model has to be substantially reduced to make the deployment of SEB economically feasible. For the building about the size of the Sutardja Dai Hall (140,000 square feet), the time required to set up the simulation model should be reduced to around 40 - 80 man-hours. For Distributed Gateway, the use of Raritan remote control switches, though appropriate for this research project, could be costly in the field. Alternate distributed load control devices should be investigated. The first choice would be to integrate control capability to the local electrical circuit controller without the needs for additional remote switches such as the Raritan or equivalent.

Before concepts such as SEB, Gateway, and DRCAOT, can be further developed into products, the product marketing manager of the demand response family of product has to present a business case to justify the investment (including money, time, and human resource). Further study of fundamental data is required to create the business case include the following data:

- market size and potential product sales,
- development costs,
- sales and marketing expenses,
- other non-monetary factors including required profit margin and internal rate of return

Once the product is justified for the development, Siemens has a rigid product development process to ensure that the product is developed with the highest quality and craftsmanship. After the product completes the field testing phase and all issues have been identified and resolved, the product is released to manufacturing for production. In addition, a field support team for the product will be formed. Support team member will be trained on all aspects of product features and functions.

When the product is released for manufacturing, the marketing team will kick off several product launching activities. Several communication channels will be used to promote and educate users on the benefits and usages of the product. Some of key channels frequently used by Siemens are: Conferences and Tradeshows, Webinar and Web Video and Internal Webinar.

Depend on the complexity of the product, in-person training will be organized to get the installation personnel up-to-speed on the technology as soon as possible. When the product incorporates technology such as on-line simulation model and optimization engine, it could be necessary for certain key personnel to attend the training offered by outside vendors. For training purpose, a live demonstration system will also be set up and used during the rollout training to demonstrate the inner working operation.

A detailed report has been delivered on Commercialization plan separately.

7.6 Next steps

While the DIADR project has ended, the CITRIS building will continue to operate as a testbed for future research projects. For this project we focused on the air-side of the HVAC system; future projects could look at water-side strategies, such as increasing the condenser water temperature on warm-hot days to reduce the cooling tower load or changing the chilled water temperature. We also could review changes to the static pressure.

We will apply what we learned from this project. For example, we discovered that the chiller load decreased 30-60 minutes after the DR HVAC curtailment control measures (e.g., increasing SAT and zone temperatures, reducing minimum ventilation rate). Any DR event would need to take this into account and perhaps employ additional anticipatory control strategies such as restrictions on chilled water flow. With more occupancy sensors, we could perform demand-controlled ventilation in all the private offices, conference rooms, classrooms, and auditorium as described as an energy efficient measure and alter it slightly for demand response events (e.g., not increase minimum air flow in anticipation of meetings).

Toward the end of the project, an ASHRAE Building Energy Quotient (Building EQ) assessment was conducted; certainly measures from that study (such as removing one lamp from many of the corridor sconce fixtures) could continue to improve the energy performance.

We will continue to improve the sMAP interface for evaluating the data from the building. For example, having an interface that can quickly provide 15-minute average energy data is quite useful for comparing to utility-metered data.

7.7 Conclusion

The stated goal of the Distributed Intelligent Automated Demand Response (DIADR) project was to develop a system that could achieve a 30% peak electrical load reduction by going deeper than the typical (often open loop) HVAC and lighting curtailment approaches. With its extensive submetering and nimble data acquisition available from the simple Measurement and Actuation Profile (sMAP), Sutardja Dai Hall provided a good test platform for DIADR. The system received an outside demand response signal, and automatically acted to curtail load from the HVAC and lighting systems. The submetering allowed better diagnosis and faster response to problems. We also demonstrated that given a DR signal, the distributed load controllers (gateway) could curtail receptacle or plug loads; Siemens developed a sophisticated user interface to allow occupants to prioritize (or opt out) appliances to be curtailed. Siemens' Smart Energy Box and the UCB DR controller with sMAP platform provided centralized supervisory control with demand response functionality. Siemens' Smart Energy Box included a real-time demand response assessment tool (DRCOAT) that could select an optimized DR strategy given the weather conditions; a validated EnergyPlus model embedded in the Smart Energy Box enabled this optimization.

While the EnergyPlus model development was not in the original scope of work, LBNL created the model and used the real performance data to refine the model. Another UCB team helped the validation process; by the end of the project the simulations were found to be within 4% of actual consumption values.

Since the building was relatively new, the energy loads of the building continued to increase over the project (100+ kW), confounding our attempts to develop a robust baseline. However, we were able to use adjacent days as baseline. Siemens conducted several DR events using the Smart Energy Box; the SEB correctly chose optimized strategies and implemented these strategies through the Apogee BAS. The UCB/LBNL teams conducted three major DR events, using the UCB DR controller through the sMAP platform. With the building using the absorption chiller, we achieved average load sheds of 14-24% of peak power. As expected, the hotter the weather, the greater the degree of sustained load shed. [The last event saw a one-hour shed of 29%, we are still analyzing the cause]. Energy efficient measures have accounted for a 50 kW reduction in load, which would save about $44k per year. We continue to improve the energy performance of this building, and plan future DR events.

8 References

ANSI/ASHRAE Standard 62.1-2010. (2010). Ventilation for Acceptable Indoor Air Quality. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. Atlanta, GA.

Arnold, Daniel. Residential Energy Gateway 2011. Available from http://mechatronics.berkeley.edu/gateway_research.htm.

(CEC), California Energy Commission. 2005. Implementing California's Loading Order for Electricity Resources, CEC-400-2005-043. Sacramento: California Energy Commission.

Dawson-Haggerty, Stephen, Andrew Krioukov, and David Culler. 2011. Experiences Integrating Building Data with sMAP. UC Berkeley. Available from http://i4energy.org/index.php/projects/affiliate-projects/6-sutardja-dai-hall.

Goldman, C., Reid, M. W., Levy, R., & Silverstein, A. (2010). *Coordination of Energy Efficiency and Demand Response*. Berkeley: Lawrence Berkeley National Laboratory.

Jones, Tyler, and D. M. Auslander. (2012). "Improved Methods to Load Prediction in Commercial Buildings." 2012 ACEEE Summer Study on Energy Efficiency in Buildings 3. Washington, D.C.: American Council for an Energy-Efficient Economy.

Jones, Tyler, and D. M. Auslander. (2010). Commercial Building Load Prediction using Modified Learning from Experience & Recursive Least Squares Methods. Baseline Prediction Calculation report available at **http://i4energy.org/diadr-project-sutardja-dai-hall-0**.

Kiliccote, Sila, M. A. Piette, Johanna Mathieu, and Kristen Parrish. 2010. "Findings from 7 years of Automated Demand Response in Commercial Buildings." In ACEEE 2010 Summer Study on Energy Efficiency in Buildings. Washington, D.C.: American Council for an Energy-Efficient Economy.

Kiliccote, S.**,** and M. A. Piette**,** (2005) **"**Advanced Control Technologies and Strategies Linking Demand Response and Energy Efficiency**",** *5th International Conference on Enhanced Building Operations*, Pittsburgh, PA, 09/02/2005.

Krioukov, Andrew, Stephen Dawson-Haggerty, Linda Lee, Omar Rehmane, and David Culler. 2011. "A Living Laboratory Study in Personalized Automated Lighting Controls." In BuildSys'11. Seattle, WA: ACM.

LeGrand. 2012. Managing Plug load is the Next Challenge for Energy Efficient Buildings. Available from http://www.legrand.us/wiremold-pages/wiremold-white-papers/plugload.aspx#.T0_VP8yaesU.

Ortiz, J., Trager, J., Saldanha, G., Culler, D., & Wright, P. (2012). Live, Continuous Energy Auditing of Miscellaneous Electrical Loads: Methodology and Results. *Proceedings from the 2012 ACEEE Summer Study on Buildings and Energy*.

Peffer, T.E. 2009. California DREAMing: the Design of Residential Demand Responsive Technology with People in Mind, University of California Berkeley, Berkeley, CA.

Piette, Mary Ann, Dave Watson, Naoya Motegi, Sila Kiliccote, and Eric Linkugel. 2006. "Automated Demand Response Strategies and Commissioning Commerical Building Controls." In National Conference on Building Commissioning.

Sator, Spencer. Managing Office Plug Loads 2008. Available from www.cityofpaloalto.org/civica/filebank/blobdload.asp?BlobID...

Trager, Jason, Michael Sankur, Jorge Ortiz, Tyler Jones, Jay Taneja, D. M. Auslander, David Culler, and Paul Wright. 2012. "Rapidly Adaptable Plug-load Simulation for Evaluating Energy Curtailment Strategies." 2012 ACEEE Summer Study on Energy Efficiency in Buildings 7. Washington, D.C.: American Council for an Energy-Efficient Economy.

9 Products of the Project

9.1 A. Publications (list journal name, volume, issue); conference papers; or other public releases of results.

Many published papers are available at http://i4energy.org/index.php/projects/affiliate-projects/6-sutardja-dai-hall.

Dawson-Haggerty, Stephen, Andrew Krioukov, and David Culler. (2011) Experiences Integrating Building Data with sMAP. Berkeley: UC Berkeley.

Jones, Tyler, and David M. Auslander. (2010). Commercial Building Load Prediction using Modified Learning from Experience & Recursive Least Squares Methods. Baseline Prediction Calculation report available at **http://i4energy.org/diadr-project-sutardja-dai-hall-0**.

Jones, T. C., Auslander, D. M., Taneja, J., Trager, J., Sankur, M., & Peffer, T. E. (2012). Improved Methods to Load Prediction in Commercial Buildings. *Proceedings of the 2012 ACEEE Summer Study on Energy Efficiency in Buildings, 3*.

Krioukov, Andrew, Stephen Dawson-Haggerty, Linda Lee, Omar Rehmane, and David Culler. (2011) A Living Laboratory Study in Personalized Automated Lighting Controls. In *BuildSys'11*. Seattle, WA: ACM, 2011.

Murthy, Nathan, Taneja, J., Bojanczyk, K., Auslander, D., and Culler, D. (2012) Energy-Agile Laptops: Demand Response of Mobile Plug Loads Using Sensor/Actuator Networks. In the Third IEEE International Conference on Smart Grid Communications (SmartGridComm'12), November 2012.

Peffer, T. E., Auslander, D. M., Caramagno, D., Culler, D., Jones, T., Krioukov, A., et al. (2012). Deep Demand Response: The Case Study of the CITRIS Building at the University of California-Berkeley. *Proceedings of the 2012 ACEEE Summer Study on Energy Efficiency in Buildings, 3*.

Trager, J., Sankur, M., Jones, T., Auslander, D. M., Wright, P., Ortiz, J., et al. (2012). Rapidly Adaptable Plug-load Simulation for Evaluating Energy Curtailment Strategies. *Proceedings of the 2012 ACEEE Summer Study on Energy Efficiency in Buildings, 7*.

Wei Zhang, Siyuan Zhou and Yan Lu, Distributed Intelligent Load Management and Control System, , 2012 IEEE Power & Energy Society General Meeting, San Diego, USA, July, 2012.

Zhu, J., Shen, L., Yin, R., Lu, Y., A TWO-STAGE SIMULATION-BASED ON-LINE OPTIMIZATION SCHEME FOR HVAC DEMAND RESPONSE, *SimBuild 2012 – the 5th National Conference of IBPSA-USA*, August 1-3, 2012, Madison, USA

9.2 Website with results of this project.

A website for this project may be found at http://i4energy.org/index.php/projects/affiliate-projects/6-sutardja-dai-hall. This site includes final reports for Tasks 2, 3, 4, 5, 6, and 8, all quarterly reports, and the slide decks for the April 27, 2011 and September 18, 2012 demonstrations, as well as other ancillary reports.

9.3 Networks or collaborations fostered.

Besides the collaboration among UC Berkeley, Siemens Corporate Research (Princeton, NJ) and Siemens Building Technology (Buffalo Grove, IL and Hayward, CA), and Lawrence Berkeley National Laboratory, the team established a relationship with Capital Projects at UC Berkeley and Facilities Dynamics (a commissioning firm).

9.4 Technologies/Techniques.

A smart phone (Android) application was developed for the energy audit.

9.4.1 Inventions

1. Tablet Based Distributed Intelligent Load Management, Yan Lu, Siyuan Zhou, Docket Number 2012E21452 US
2. Distributed Intelligent Load Management and Control, Yan Lu, Siyuan Zhou, Wei Zhang, Docket Number 2011E18351 US
3. Adaptive Demand Response Based on Distributed Load Control, Yan Lu, Siyuan Zhou and Mohammad Abdullah Al Faruque, Docket Number 2011E22943 US

9.4.2 Patent

1. Adaptive Demand Response Based on Distributed Load Control, Yan Lu, Siyuan Zhou and Mohammad Abdullah Al Faruque, Application No. 2011P22948US01

9.5 Other products, data or databases, physical collections, audio or video, software or netware, models, educational aid or curricula, instruments or equipment.

EnergyPlus model, Rapid Auditing Protocol, Gateway, Demand Response Capacity & Operation Assistance Tool, sMAP BACnet control interface, Personal Lighting Control

10 Computer modeling

EnergyPlus model is described in Appendix G.

11 Appendices

11.1 Appendix A: Building description

Sutardja Dai Hall (SDH) on the UC Berkeley campus is the headquarters of the Center for Information Technology Research in the Interest of Society (CITRIS). The building has approximately 141,000 gross square feet of space that houses laboratories for collaborative research, open plan and private offices, a 149-seat auditorium, conference rooms on each of seven floors, state-of-the-art classrooms, a data center, walkways that connect with the surrounding buildings, and 12,000 square feet dedicated to the Marvell NanoLab (the most sophisticated academic clean room worldwide). The building also houses the Main Distribution Center (MDC) for the northeast quadrant of campus. See the figure below.

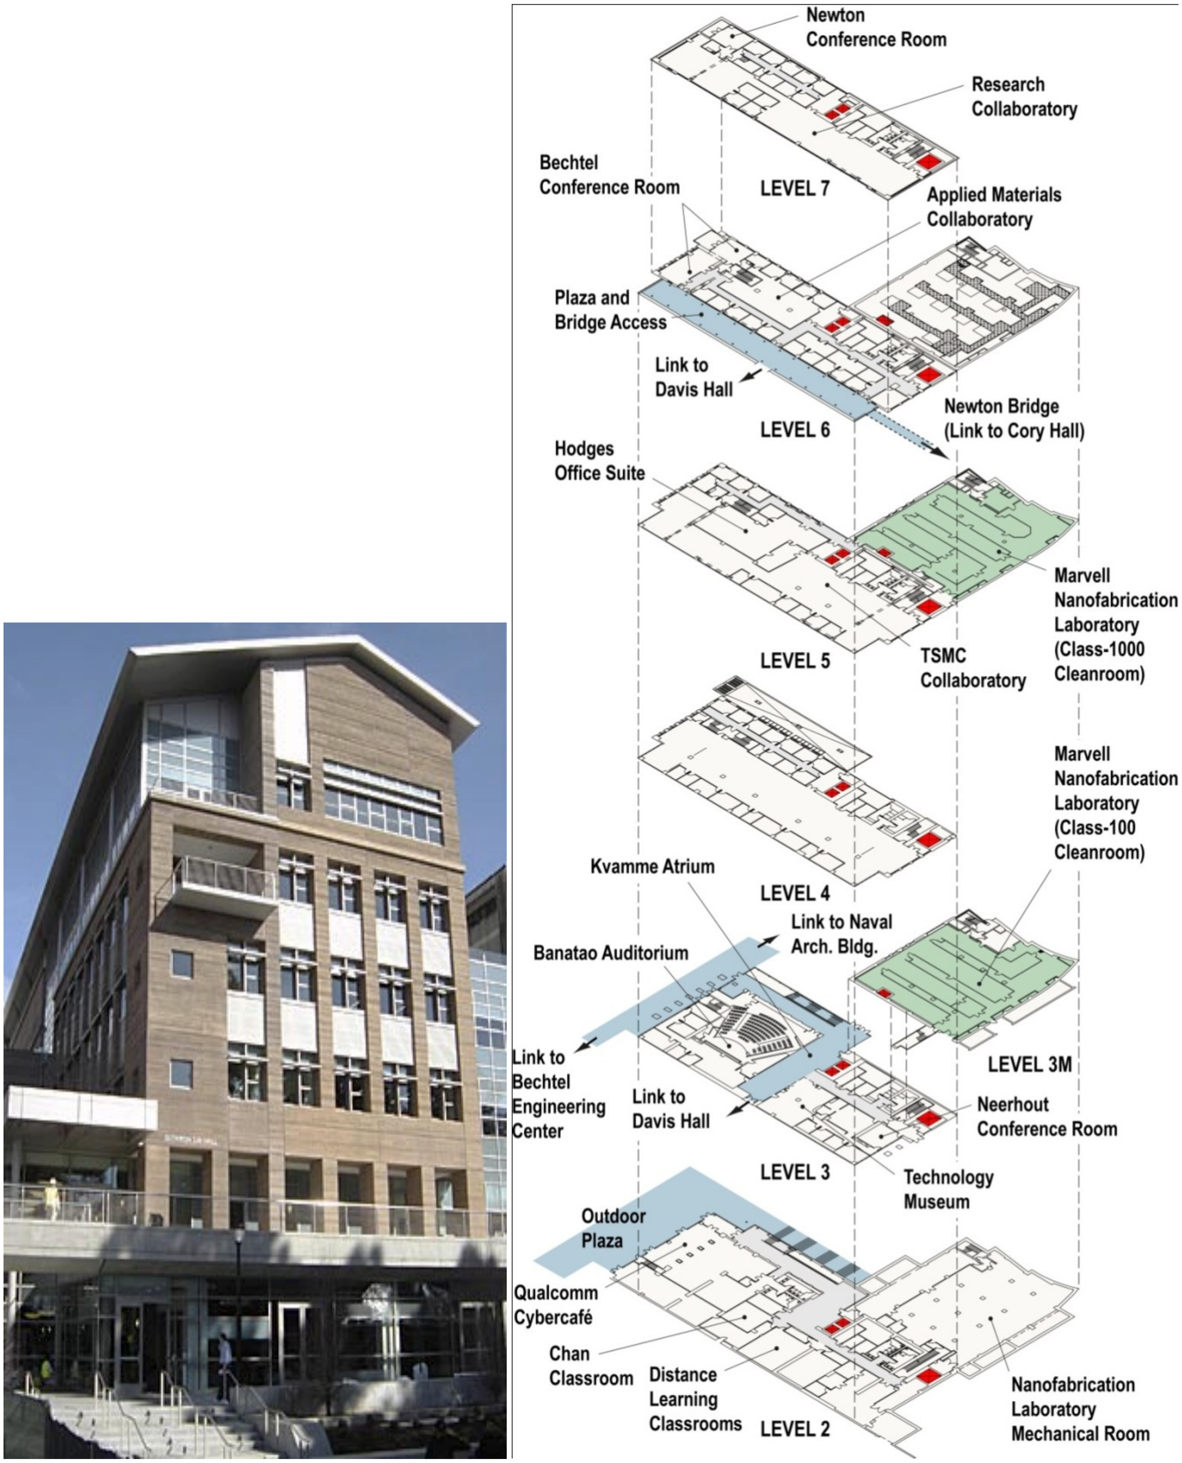

Figure 63: Left: West facade of Sutardja Dai Hall; Right: schematic of building uses.

SDH was constructed using concrete with a high volume of coal fly ash, which releases significantly less CO_2 into the atmosphere than the standard Portland cement. Designed as a hub for scientists and researchers from different disciplines, floors four through seven are for research, mixed with collaborative spaces to promote interaction between the 50 CITRIS faculty affiliates, hundreds of students, and the administrative staff for CITRIS. It was dedicated and occupied in 2009. It is a complex facility providing diverse usage models concurrently, ranging from the NanoLab to a data center, to laboratories, and to office space.

The building has a Siemens Energy Management and Control System (EMCS) and Siemens Apogee Building Automation System (BAS). The WattStopper lighting system in the open plan offices (found on floors 4-7) has tri-level dimming capability and is on a timed schedule. The private offices have Lutron wall switches with dimming and an occupancy sensor.

The building has two 600 ton Trane chillers controlled through the Building Automation System. The absorption chiller was designed to use steam from April through October when steam on the UC Berkeley campus from the 30MW co-generation facility is not in high demand for heating. A centrifugal compressor chiller with hot gas bypass was designed for use from November through March.

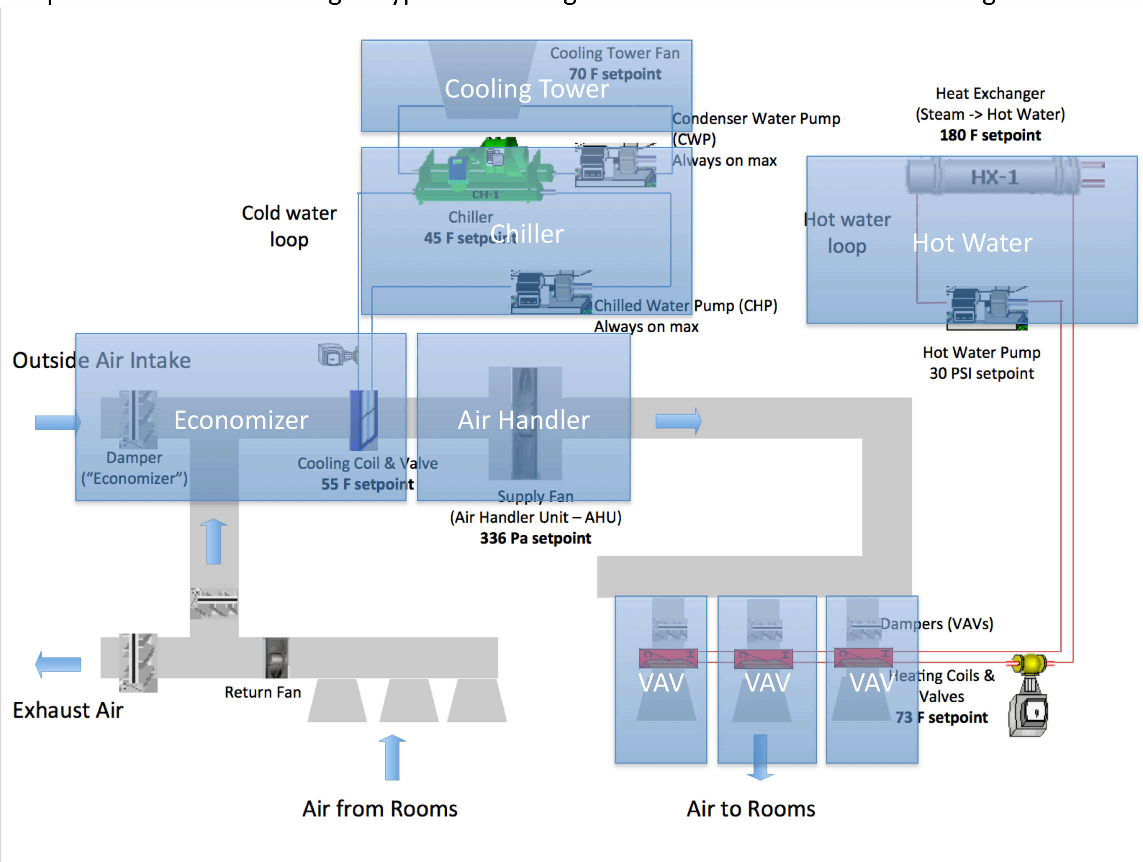

Figure 64: Schematic of the building's conditioning system (Krioukov).

11.2 Appendix B: Instrumentation

11.2.1 Submetering

As Sutardja Dai Hall was intended to be a "living laboratory", CITRIS invested over $200,000 for submetering and infrastructure in 2010. This includes 27 revenue grade DEM 2000 power submeters with modbus connections on most subpanels (including submeters for lighting and receptacle power on each floor) for a total of 7000 sense points in the BAS. These monitor energy (kilowatt-hours), voltage, current, power factor, and peak demand, while also integrating existing BAS data points that currently monitor and trend HVAC control data, providing end users a comprehensive picture of how energy is consumed throughout Sutardja Dai Hall. See the following two figures. However, in order to isolate the loads of the nanofabrication laboratory, two additional electrical subpanels required submetering; in addition, because the office portion of the building and the nanofab share chilled water resources, flowmeters were required to determine the portion of chilled water for the office versus the lab. In early 2012, the panels were upgraded from the modbus RS-485 network to Ethernet to provide faster communications. In October 2012, four variable frequency drives were added to the chilled water and condenser water pumps for each chiller: CHP-2 and CWP-3 for the absorption chiller (CH1) and CHP-1 and CWP-1 for the centrifugal chiller (CH2). These VFDs provide turndown control as well as submetering the power consumption of the device.

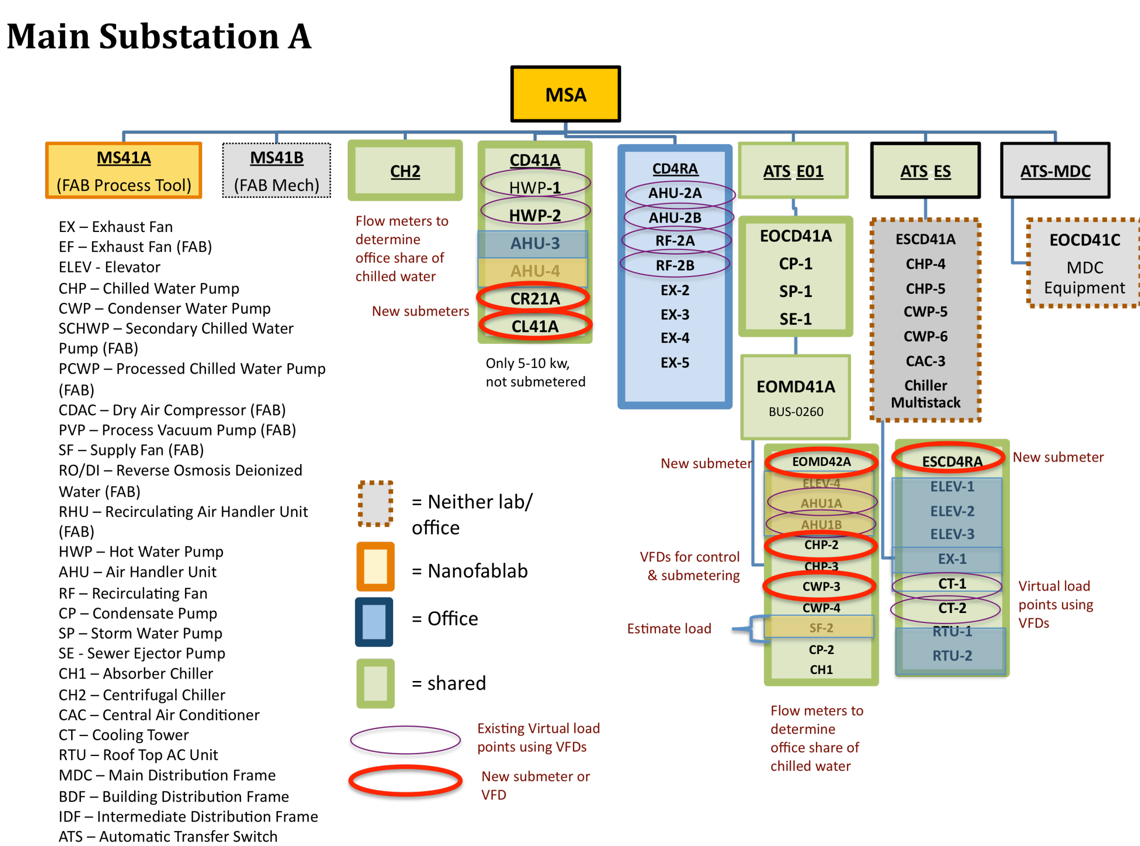

Figure 65: The electrical load diagram of main substation A for Sutardja Dai Hall.

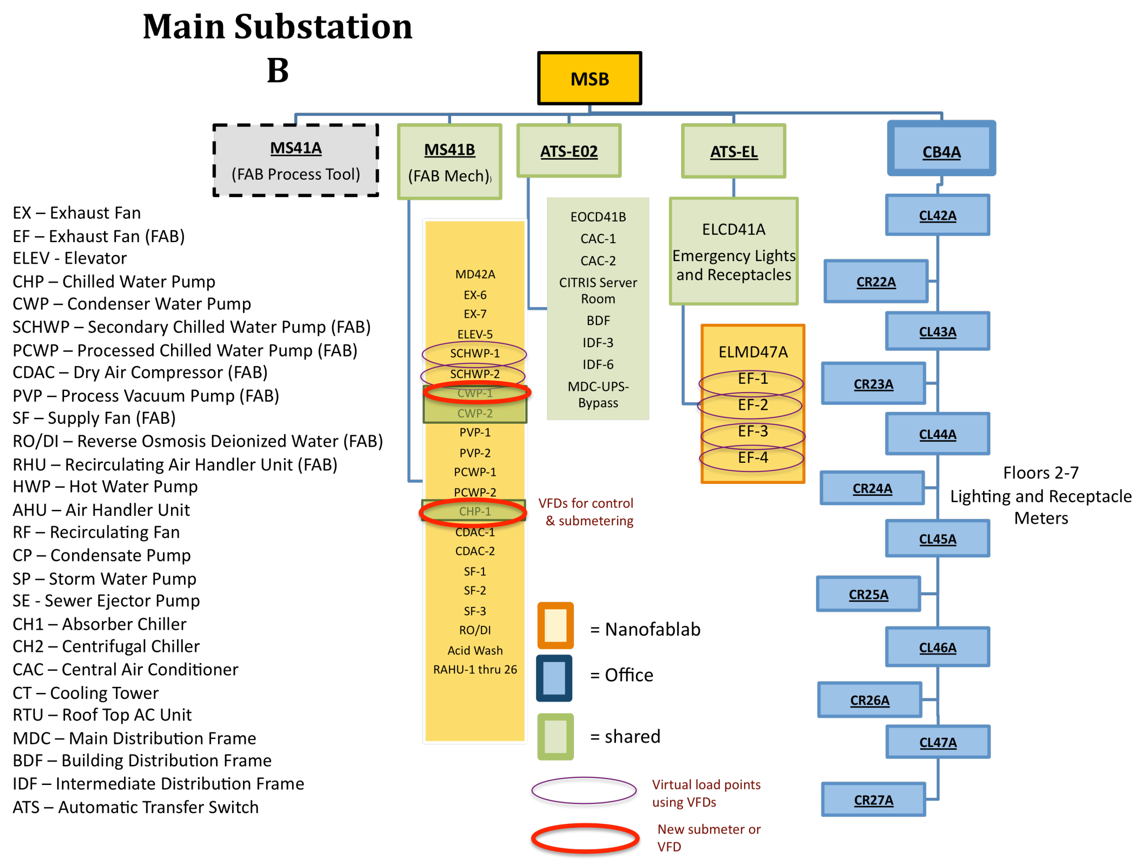

Figure 66: The electrical load diagram for main substation B.

We used the following equation to determine the office portion of the electrical load:
Office Side Electric Load =66%***HWP-1** + 66%***HWP-2** + **AHU-3** (constant of 5.5kw) + 66%*(**ATS-EO1** - **EOMD42A**) + 25%*(**EOMD42A** - (**AHU1A** + **AHU1B**)) + [**ESCD4RA** - 75% * (**CT-1** + **CT-2**)] + **CD4RA** + **CB4A** + 66%* **ATS-EO2** + [66% **(ATS-EL - [EF1+EF2+EF3+EF4])]**

Assumptions:
Domenico Caramagno, the facilities director, suggested 66% for general loads such as the hot water pumps, emergency panel, and general telecommunications.
Rongxin Yin initially calculated the chilled water flow from the chillers to the office at about 25%; this was using point measurements of temperature. When Bin Chen compared the centrifugal chiller performance to the absorption chiller performance on Oct 17 2011 and Oct 17 2012 respectively, he felt that on warm days the percentage was closer to 40%. Once the wireless temperature sensors have been installed, we will be able to calculate and trend the flow changes in hot weather.

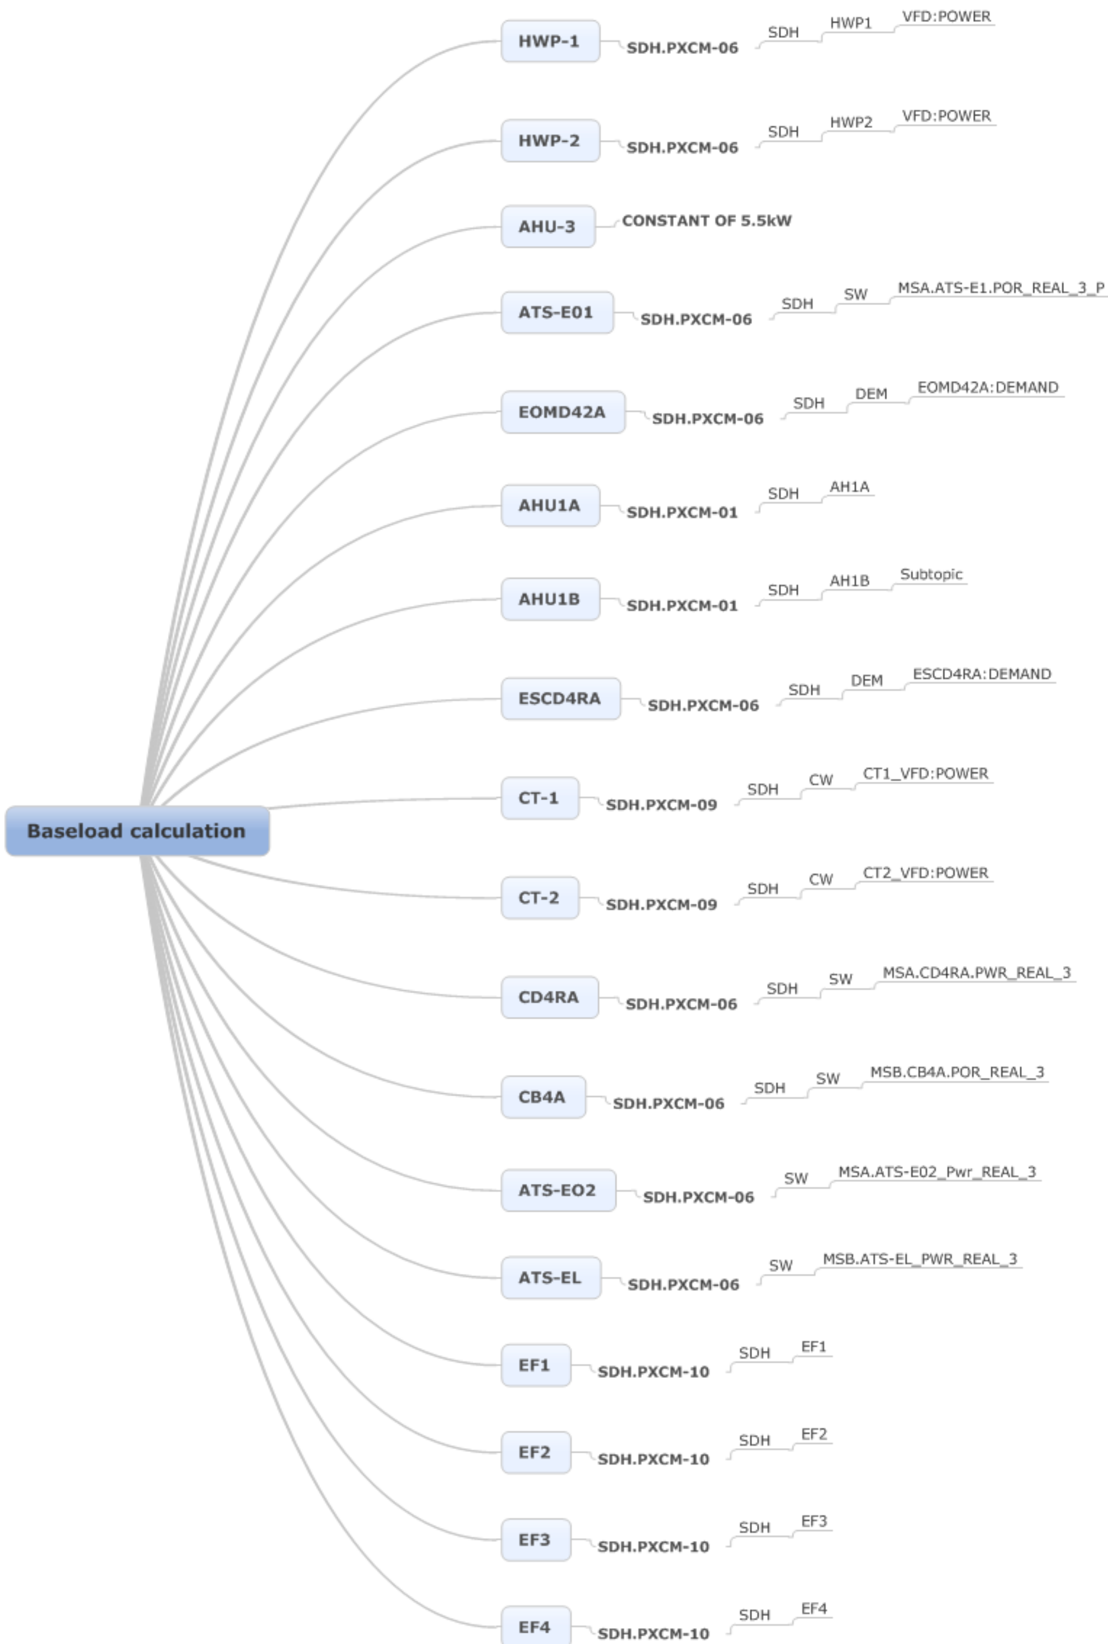
Baseload calculation
HWP-1
SDH.PXCM-06
SDH
HWP1
VFD:POWER
HWP-2
SDH.PXCM-06
SDH
HWP2
VFD:POWER
AHU-3
CONSTANT OF 5.5kW
ATS-E01
SDH.PXCM-06
SDH
SW
MSA.ATS-E1.POR_REAL_3_P
EOMD42A
SDH.PXCM-06
SDH
DEM
EOMD42A:DEMAND
AHU1A
SDH.PXCM-01
SDH
AH1A
AHU1B
SDH.PXCM-01
SDH
AH1B
Subtopic
ESCD4RA
SDH.PXCM-06
SDH
DEM
ESCD4RA:DEMAND
CT-1
SDH.PXCM-09
SDH
CW
CT1_VFD:POWER
CT-2
SDH.PXCM-09
SDH
CW
CT2_VFD:POWER
CD4RA
SDH.PXCM-06
SDH
SW
MSA.CD4RA.PWR_REAL_3
CB4A
SDH.PXCM-06
SDH
SW
MSB.CB4A.POR_REAL_3
ATS-EO2
SDH.PXCM-06
SW
MSA.ATS-E02_Pwr_REAL_3
ATS-EL
SDH.PXCM-06
SW
MSB.ATS-EL_PWR_REAL_3
EF1
SDH.PXCM-10
SDH
EF1
EF2
SDH.PXCM-10
SDH
EF2
EF3
SDH.PXCM-10
SDH
EF3
EF4
SDH.PXCM-10
SDH
EF4

11.3 Appendix C: sMAP (simple Measurement and Actuation Profile)

An enormous amount of physical information, that is, information from and about the world, is available today, especially in buildings, as the cost of communication and instrumentation has fallen. This physical information is essential for building models of building performance and energy consumption, as well as for determining and assessing demand response. However, making use of that information is still challenging. The information is frequently siloed into proprietary systems, available only in batch, and fragmentary and disorganized. sMAP seeks to change this by making available and usable:

- a specification for transmitting physical data and describing its contents,
- a large set of free and open drivers communicating with devices using native protocols and transforming it to the sMAP profile, and
- tools for building, organizing, and querying large repositories of physical data.

The pieces of the sMAP system are designed to separate concerns and allow users to, for instance, run their own web front-end while using hosted infrastructure for storing the actual data and metadata, as shown below.

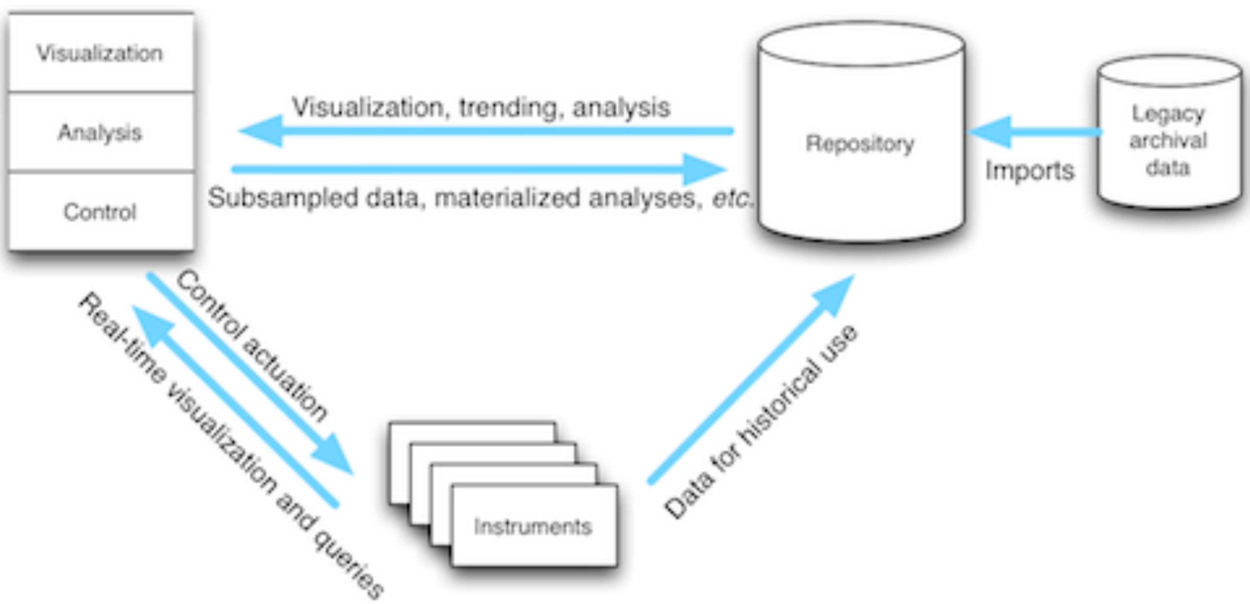

Figure 67: Functional diagram of sMAP system components.

The core object in sMAP is the *Timeseries*, a single progression of (time, value) tuples. Each Timeseries in sMAP is identified by a UUID (universally unique identifier), and can be tagged with metadata; all grouping of time series occurs using these tags. These objects are exchanged between all components in this ecosystem.

The first essential piece of sMAP is a library for writing **instrument drivers**. These drivers connect to existing instrumentation and provide tools for exposing the data over http/sMAP. The library and protocol are designed to support various common scenarios:

- Intermittent connectivity: provide local buffering
- Local metadata: apply tags at the source
- Bulk loading and real-time: support both bulk-loads from existing databases and real-time data from streaming or polling sources in the same framework
- Actuation (using SSL (secure socket layer) encryption)

Information about using the sMAP library is available in pydoc (documentation for Python module code) and can be found at **http://code.google.com/p/smap-data/**.

The **repository** gives the drivers a place for instruments to send their data. It supports:

- Efficient storage and retrieval of time-series data
- Maintenance of metadata using structured key-value pairs
- Metadata querying using the ArdQuery language
- Front-end

Most systems provide some amount of graphic display dashboard and plotting. Out of the box, the **powerdb** project provides plotting and organization of time-series data, built on top of the **ArdApi**. Due to the decoupled nature of the design, this front-end can be run by anyone. The application is designed to give users a large amount of flexibility to organize, display and plot streams using **ArdQuery** to generate tree views of their streams using the **SlicrApi**.

A running instance of sMAP with many live streams is available at **http://new.openbms.org/plot/**

As one example of the utility of this design, the figure below shows a public dashboard built on sMAP for Sutardja Dai Hall. Note that in addition to the usual figures of merit, like energy consumption, it shows how the building is performing relative to real time baselines that are based on an empirical model as a function of time of day, day of week, and outside air temperature. Also, it brings together a variety of different meters. For this particular day, electrical consumption is much higher than expected and steam is much lower because the HVAC system has switched from using the absorption chiller to the electrical chiller. In this timeframe, the office portion of the building is also being controlled through sMAP, via the Apogee BMS, so the dashboard shows a room-by-room accounting of the state of the zones relative to set point. By clicking the star in the upper right, the facilities manager can access deeper information.

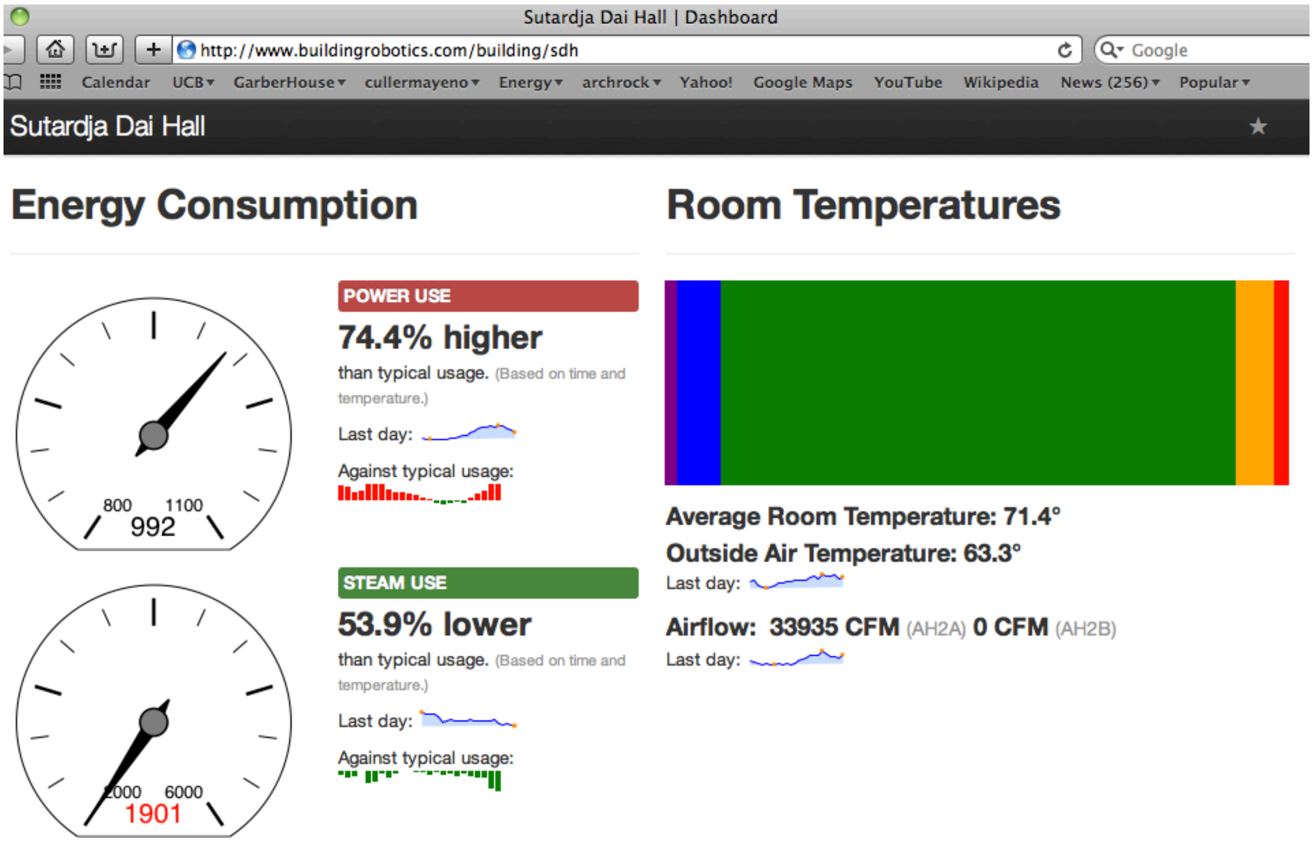

Figure 68: sMAP-based dashboard for Sutardja Dai Hall.

The sMAP infrastructure permits scaling out to many buildings or scaling in to detailed aspects of a particular building. For example, we have built a campus-wide energy portal that comprises over 50 buildings. Below, Figure 63 shows the entry page of the campus portal with a map linking to detailed information on about 50 monitored buildings (yellow in color). Clicking on one of the buildings provides additional detail over various time frames. For example, Figure 64 shows three days of energy consumption for Soda Hall: a hot weekday followed by a cool weekday followed by a cool Saturday. Another example, Figure 71, looks back to a winter day for Sutardja Dai Hall.

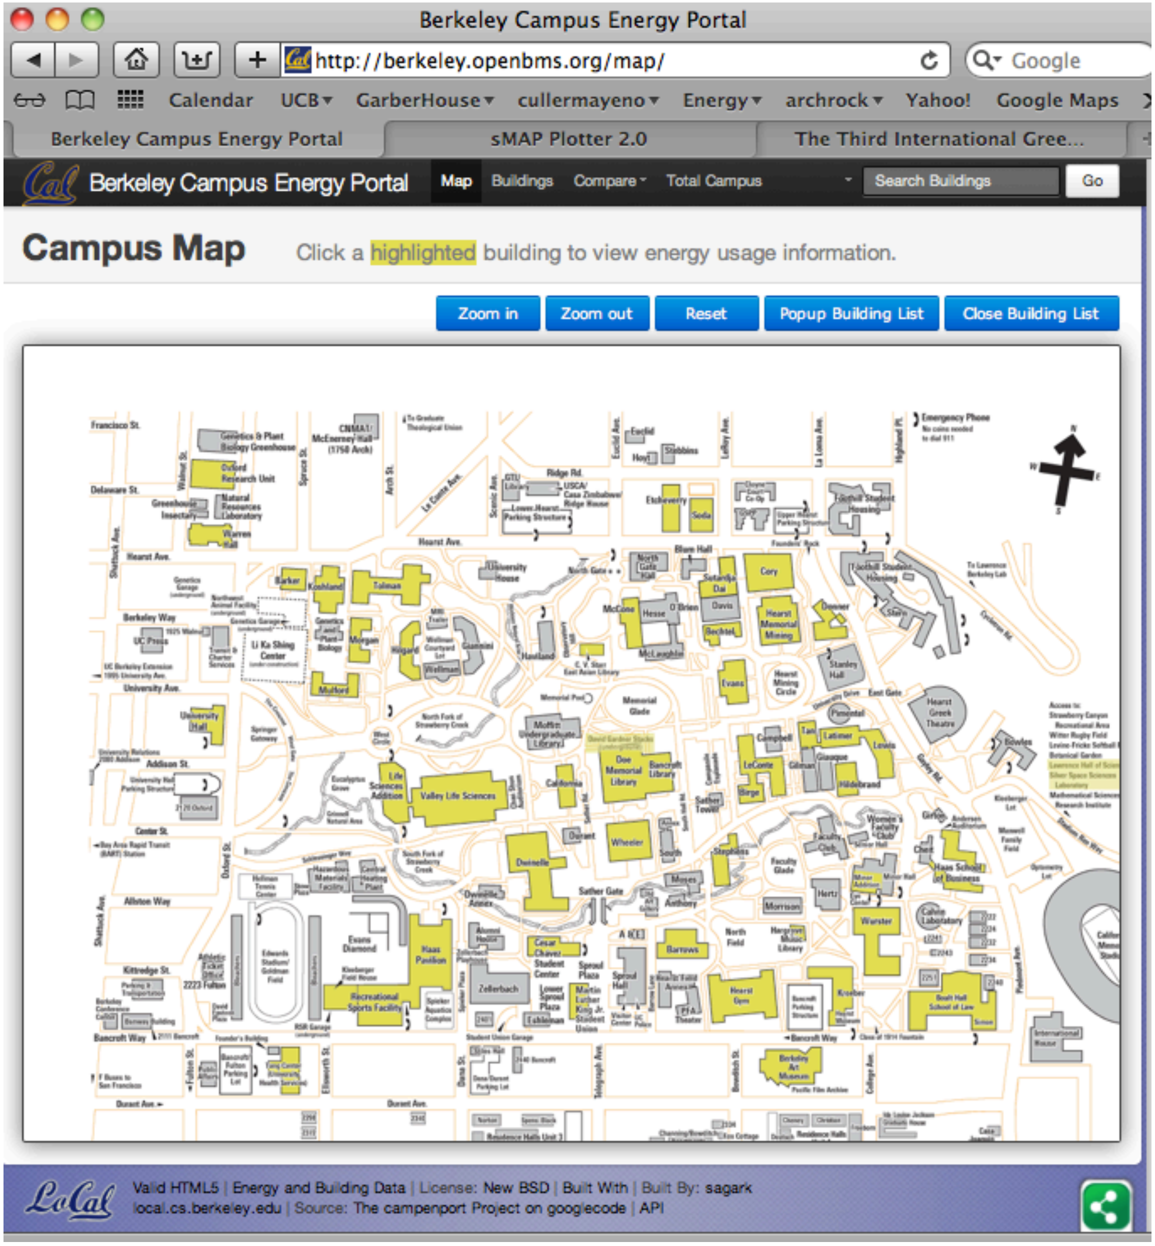

Figure 69: Berkeley sMAP campus portal - map view.

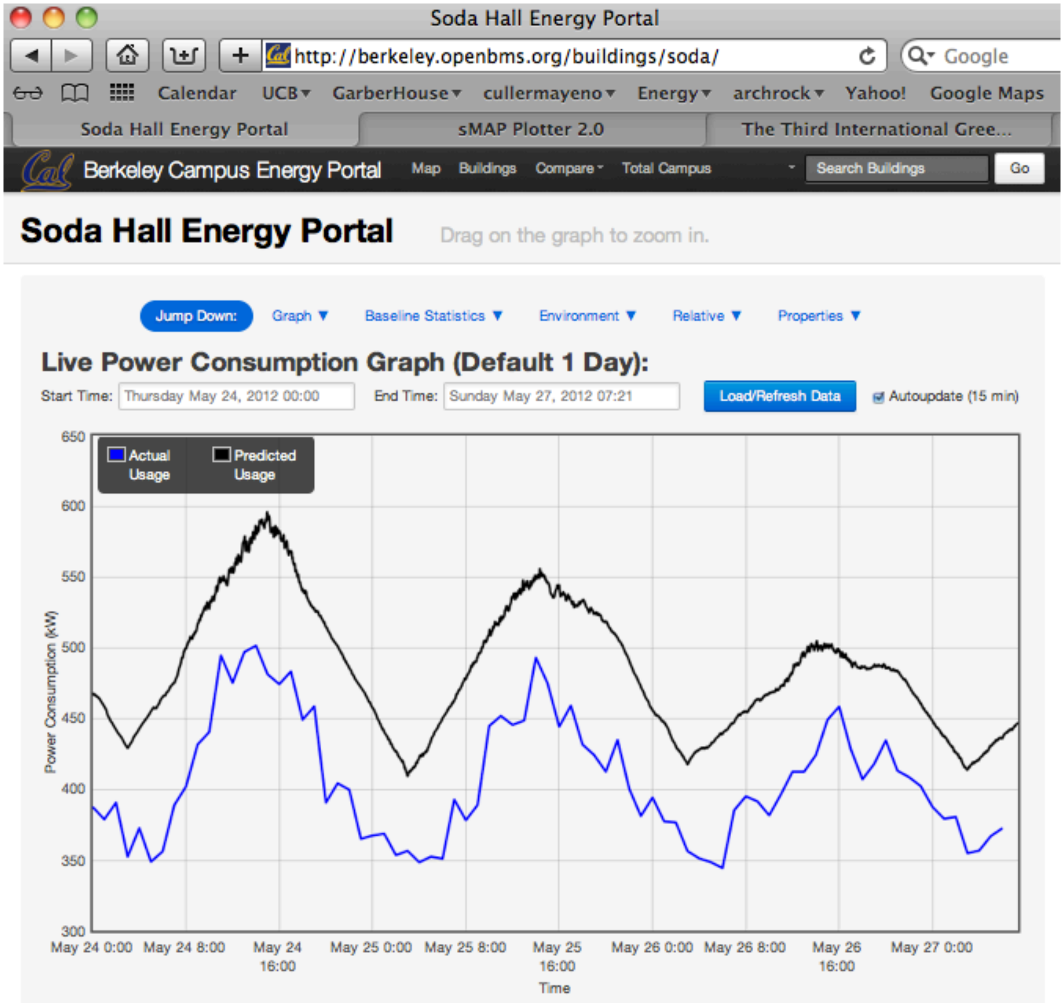

Figure 70: Three days (hot, cool, weekend) of Soda Hall through UCB sMAP portal.

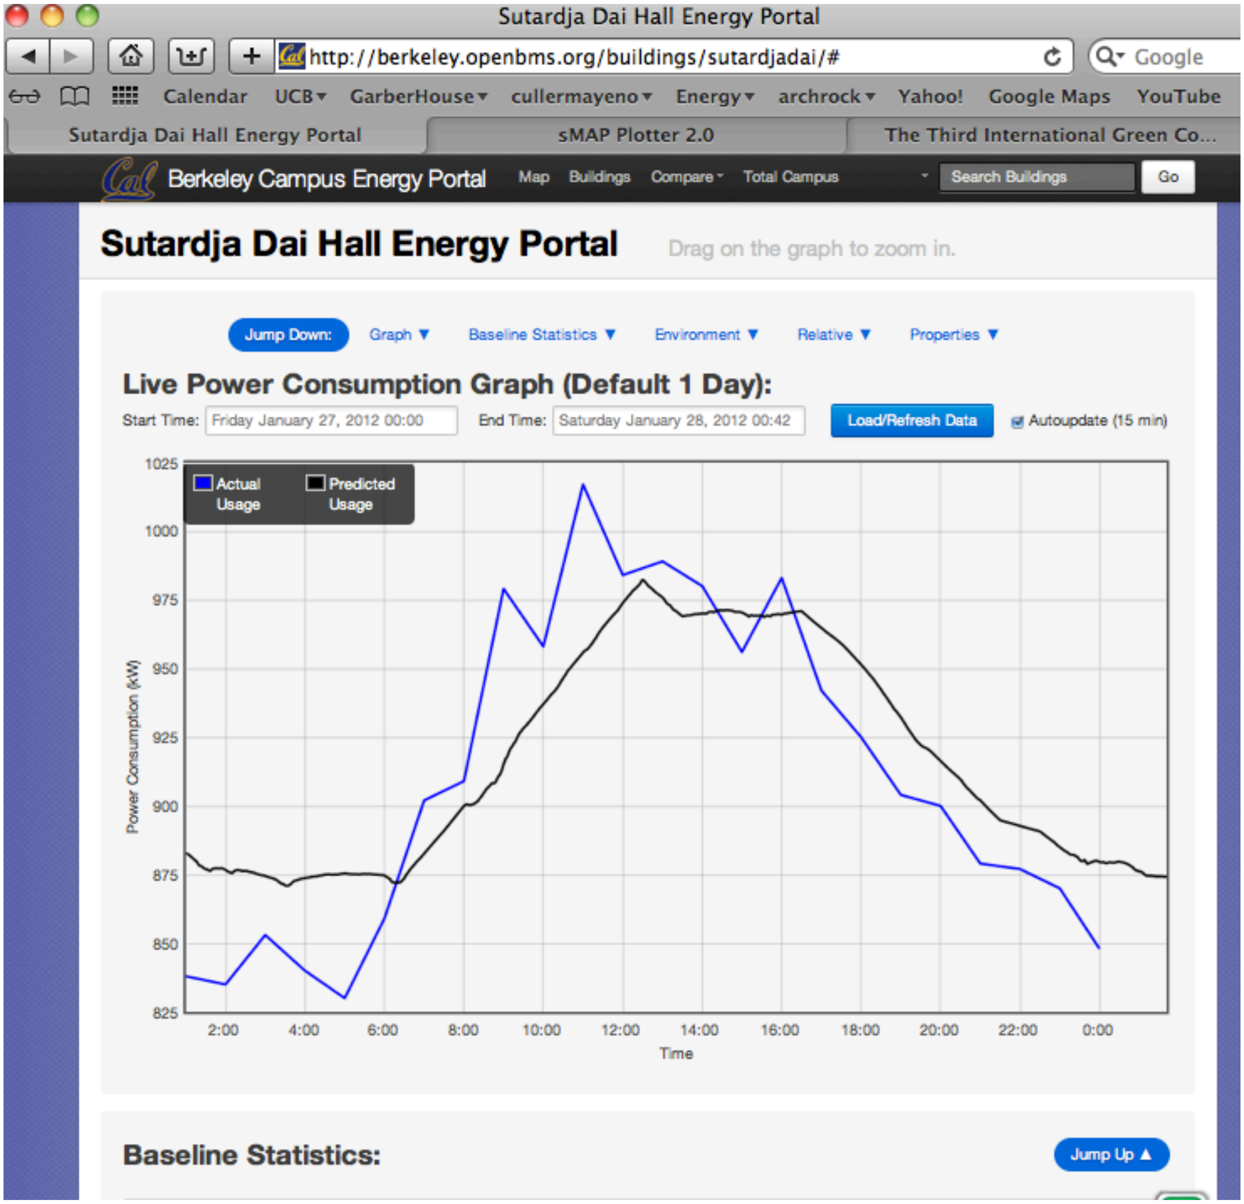

Figure 71: A winter day in Sutardja Dai Hall through sMAP portal.

At the campus scale, Figure 72 shows an sMAP-based aggregation of whole building metering from fifty buildings on the Berkeley campus over a week with weekend days on either side. During weekdays we see a swing from the nighttime load of 12.5 MW to a peak of about 18 MW, whereas on weekends the peak is about 15 MW. This shows both the potential for broad demand response measures in peak management, but also the huge energy efficiency potential of deeper night time setbacks. Also, much of the load is process-oriented, including semiconductor manufacturing, fume hoods in chemical and biological laboratories, refrigerators, and computing equipment. Addressing these process loads requires a very different approach than office environmental conditioning.

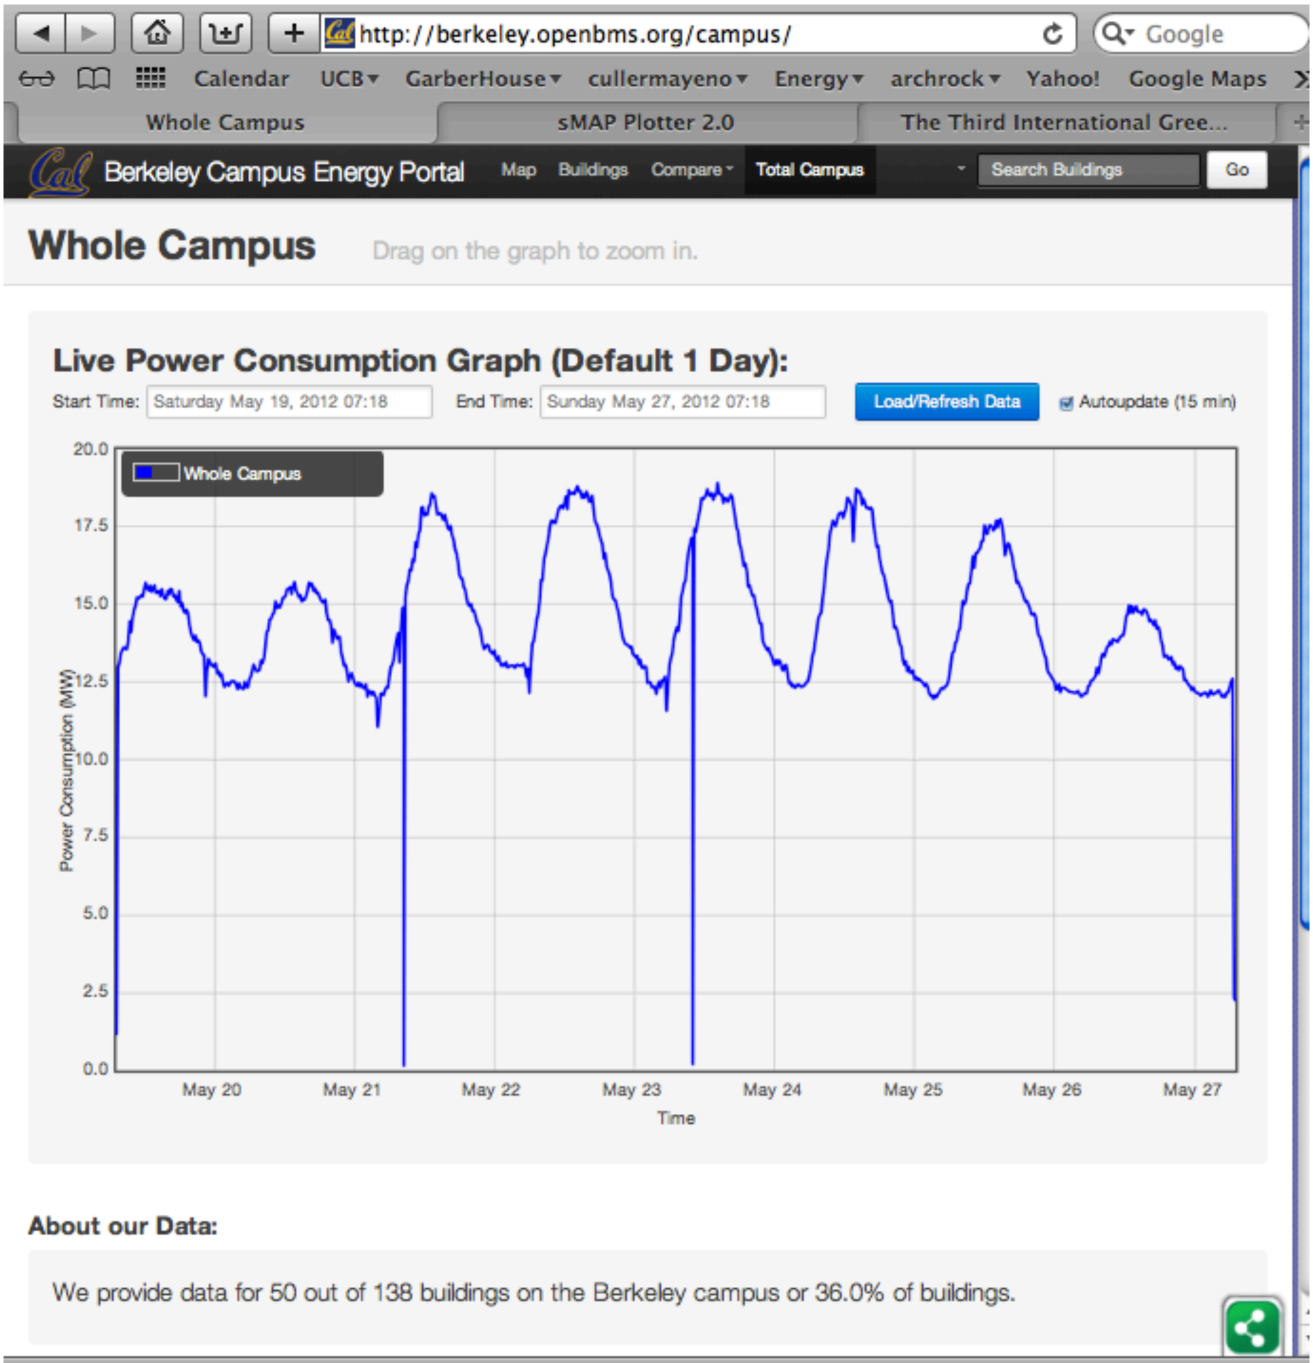

Figure 72: Aggregation of Meters from 50 UCB buildings in Campus sMAP Portal.

Stepping outside the portal to a more detailed data analysis tool, Figure 73 shows an overlay of Sutardja Dai Hall power consumption at 15 second resolution with a whole building analysis of the 114 variable air volume (VAV) damper positions. Using this tool, air handling units, fans, pumps, and chillers can be examined. This detailed analysis has proved essential for calibrating simulation models.

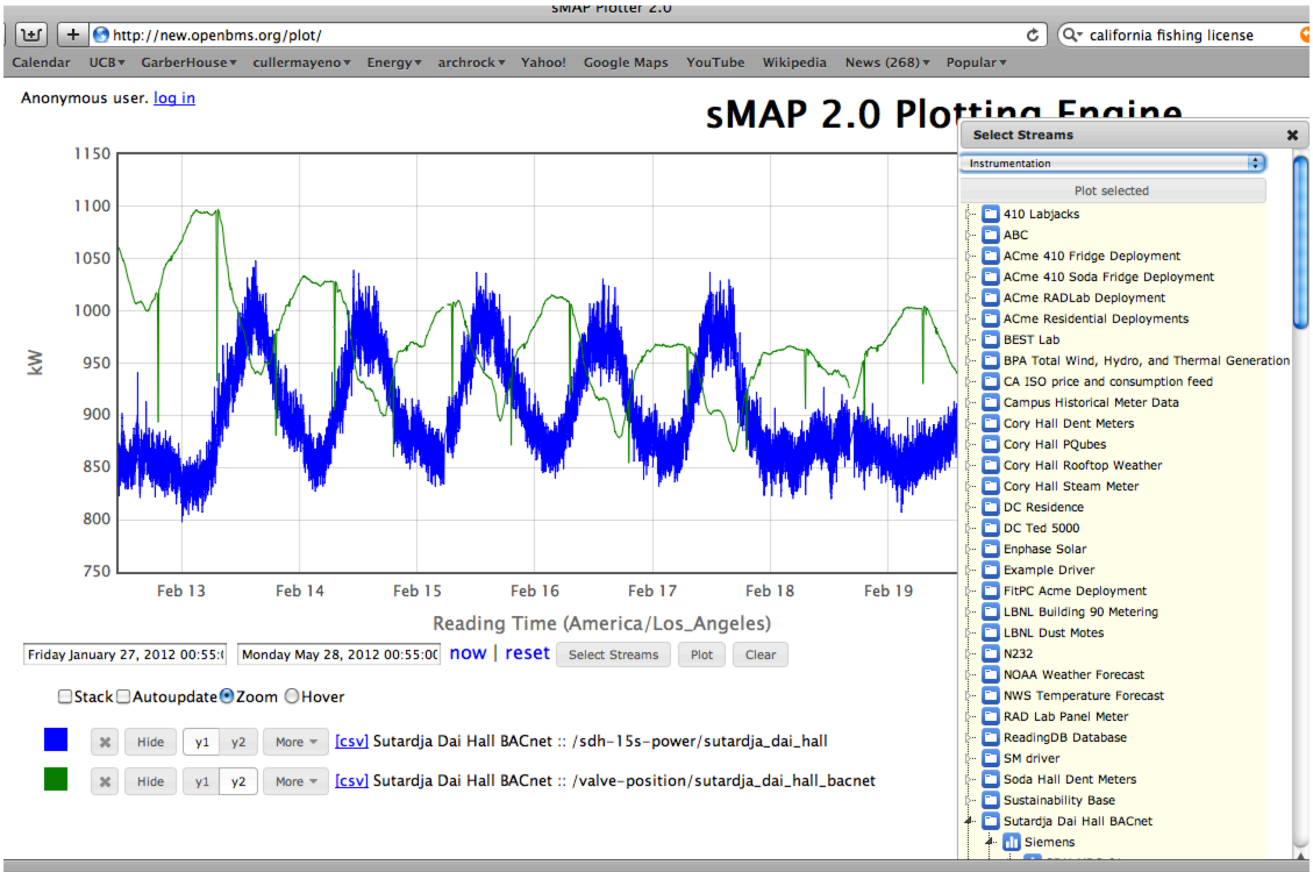

Figure 73: Sutardja Dai whole building power at 15 s resolution and damper position.

Andrew Krioukov also developed a lighting interface for the occupants to control their overhead lighting (see figure below).

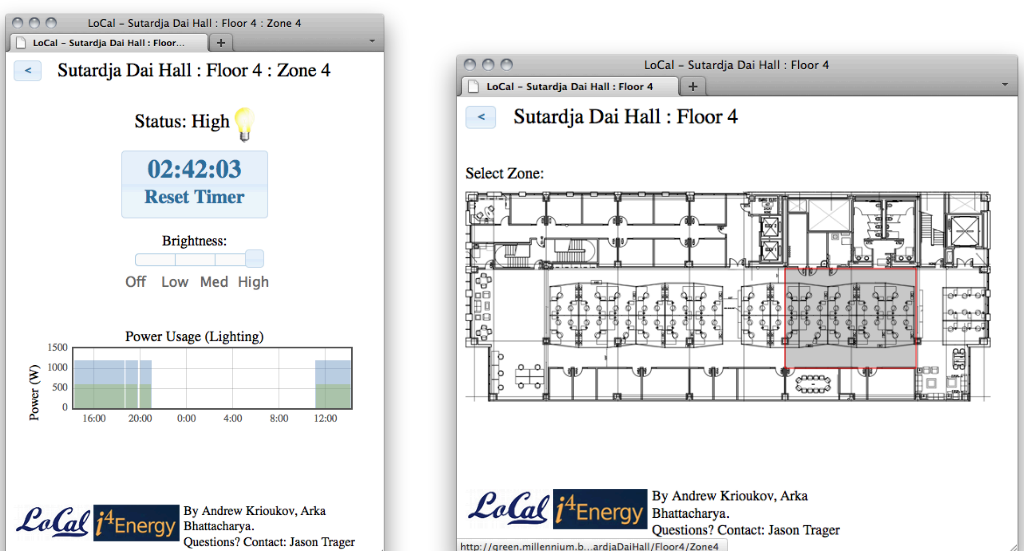

Figure 74: Lighting interface for occupants to choose their lighting zone, brightness level, and time on.

In addition, Andrew Krioukov developed a thermostat interface so occupants could request a cold or hot blast of air through the HVAC.

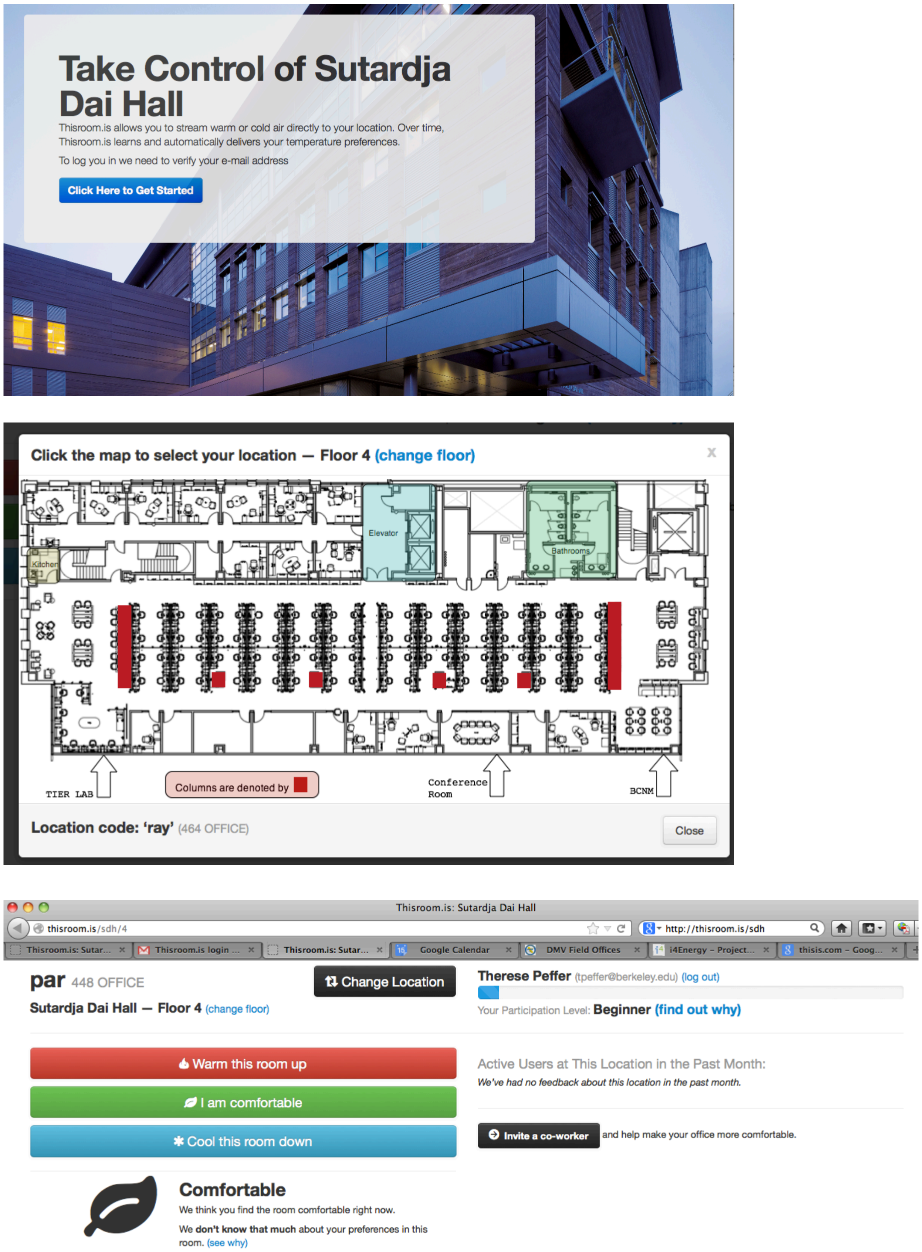

Figure 75: Web-based office temperature control.

11.4 Appendix D: UC Berkeley DR Controller

This section describes the details of the UC Berkeley DR Controller, the Central Load-Shed Coordinator (CLSC), which was designed as a program that facilitates coordinated control over building lighting, HVAC, and plug loads.

First, we define the terminology:
A ***strategy*** refers to what control actions, and at what times, a system (lighting, HVAC or plug loads) will actuate. A strategy has one or more steps (control actions at times relative to the onset), and a time validity window for when it can be used. For example, a lighting strategy may turn off all hallway lights on all floors at its onset and dim all office lights on the fifth floor after 30 minutes.

A ***bid*** is a time profile of power savings with an associated metric of inconvenience (occupant discomfort and/or loss of productivity). Every strategy has an associated bid, which is calculated in the context of the strategy being activated at a certain time. For example, a two hour lighting strategy activated at noon may reduce power by 20 kW for its duration, with a low level of inconvenience, and a three hour HVAC strategy activated at 1 p.m. on a hot day may save 100 kW with a much higher level of inconvenience.

A ***plan*** is a group of strategies and the time(s) they are to be activated. For example, a plan may consist of enacting an hour long lighting strategy at 2 p.m., and a three hour HVAC strategy starting at noon.

Due to the diverse nature of the load-shed resources, and given communication hardware requirements, the CLSC is sectioned into four modules: Core Module (CM), Lighting Module (LM), HVAC Module (HM), and Gateway Module (GM). The CLSC encompasses a two-tier design, where the CM acts as the information coordinator for the LM, HM and GM. The CLSC is designed so that a lower level module may be replaced or modified to accommodate for different communication standards, without the need to redefine intra-module interaction. The process in which the CLSC operates for a single DR event is as follows, which can be repeated for non-overlapping events:
The CM polls for, and receives a DR signal
The CM parses the signal, sends it to each module, and the GM relays the signal to EIGs
The LM, HM and EIGs, with knowledge of a priori defined strategies, calculate associated bids
The CM collects all bids and a plan is selected either by the user or optimization program
The CM schedules strategy execution for lighting and HVAC and relays plan to EIGs
The CM queries each module for pertinent system information

Software Architecture
The two-tier design of the CLSC and the flow of the DR signal, bids, system information and instructions between modules is highlighted in Figure 76. The software architecture of the controller will now be discussed along with the specifics of its interaction with the three load-shed resources.

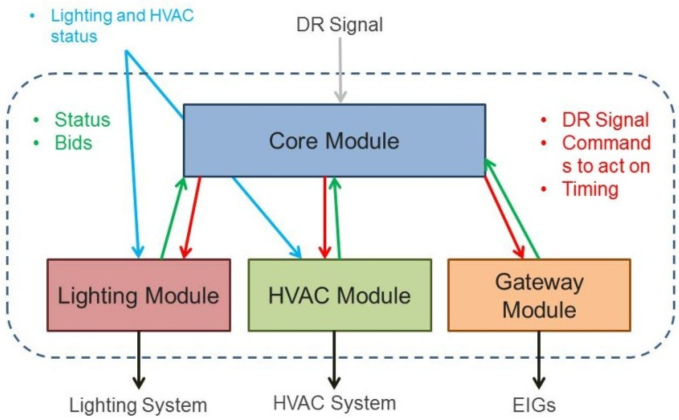

Figure 76: Detailed view of CLSC two-tier architecture.

Core Module

The CLSC core module acts as the architecture administrator. The CM handles DR event timing, and contains scheduling logic for strategy execution and system monitoring. The CM polls openADR compliant DR servers, and when a valid DR signal is received, the CM interprets the signal for its own timing and relays it to the three other modules. Once a plan is chosen, the CM instructs the GM to relay the plan to the EIGs. The CM schedules execution of strategies for lighting and HVAC, instructing the LM and HM to perform a step of a strategy at the appropriate time.

Lighting Module and HVAC Module

The lighting module (LM) is the communication bridge between the CLSC and a building lighting system. The LM has knowledge of user defined lighting strategies. When prompted by the CM calculates associated bids for each strategy, which are returned to the CM when requested. The LM contains methods to send commands to a lighting system for control. It also contains methods that are called by the CM, to query for system information. The HVAC module (HM) parallels the LM in function, but for the HVAC system. The separation of functionality into two modules is to accommodate for differences in lighting and HVAC system communication. Furthermore, lighting and HVAC are likely to have different types of strategies. Lighting is usually actuated with binary control (on/off), and HVAC systems contain multiple control parameters for temperature, flow rates, operation mode, etc. Both the LM and HM utilized BACnet to interface with the building energy management system, and therefore the lighting and HVAC systems, respectively. The LM and HM also poll sMAP for their respective system's information.

Gateway Module

The Gateway Module enables communication between the CLSC and EIGs in SDH. The paradigm in which this architecture is designed, is for the EIGs to aggregate data from, and control, plug loads within their domain. The GM utilizes existing infrastructure, such as WiFi, LAN, or Zigbee, to host a network for CLSC-EIG communication. In this deployment in SDH, the GM uses JADE (Java Agent DEvelopment framework) to create a network over which it can communicate with EIGs. The GM disseminates the DR signal to the EIGs. EIGs read user defined strategies for their connected plug loads and create associated bids. The bids are passed up to the CM through the GM. When a plan is selected, the relevant information is sent to the EIGs through the GM. The GM contains methods to poll EIGs for overall and plug load status.

Optimization Program

As mentioned earlier in this paper, a plan is chosen based on strategies and bids from the three modules. Every strategy has an associated bid, which is calculated from the strategy and models of power saving and inconvenience. The term inconvenience is a measure of the negative effects of actuation of a load-shed resource. Inconvenience can be loss of productivity in man-hours or dollars, or a measure of occupant discomfort, such as the square sum of occupant preferred temperature to actual temperature. Models for power saving and inconvenience are user defined and the architecture is designed for ease of modification and replacement. While a building manager can choose a plan based off the bids, the CLSC is designed to be able to utilize a modular optimization program.

The optimization program takes in bids and associated strategy timing parameters as its input. The program objective is to minimize total inconvenience while meeting a power reduction constraint (a load-shed goal) over a period of time. It returns a course of action (plan) for a user defined cost function and set of constraints. The architecture is also designed for ease of modification of the optimization program parameters. When the optimization program is used, the CLSC is fully automated in the sense that it only requires a user to initiate the program.

User Interface

Though the CLSC is designed as an autonomous architecture, it is still important to provide pertinent information to a building manager and/or occupants. Furthermore, the building manager needs to retain supervisory control over the architecture and its control decisions. For safety reasons, he or she must be able to interrupt CLSC operation, assume control of a system, and return the systems to their defaults, should the need arise. To this end, the CLSC incorporates a simple user interface (UI). The UI displays relevant information such as total building load, DR event parameters and plan/strategy information. It also allows a user to interrupt operation and restore system defaults.

11.5 Appendix E: Optimizer

When scheduling a demand response event, it is important to assure that a certain amount of power will be reduced during every time point in the day. For the DIADR Project, we have arranged for an autonomous arrangement of the order of strategies for saving power during a DR event. In order to do this, we classify all loads for DR in terms of how much power they save over time, and how inconvenient they are during that time.

While we want to look at reducing power use at every time during a day, we also wish to evaluate the cost of that power draw reduction in terms of how it affects comfort and occupant productivity. In order to do so, we have to assign an "Inconvenience cost" (Ip) to every action that we take at every time step. This inconvenience cost over each time step must be paired with the power use reduction at every time step – the savings profile (Sp), and can be imagined to be two time-linked profiles of cost and power use reduction. The two figures below will help visualize this.

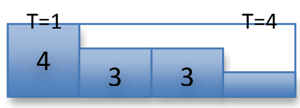

Figure 77: Savings Profile example: each block represents 15 minutes of time. Numbers inside represent savings in kW

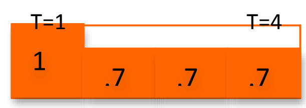

Figure 78: Inconvenience profile example. Each block represents 15 minutes, each number represents an inconvenience score.

After profiles are established, it is important to be able to shift those profiles in order to move the strategy to an appropriate time step that it helps us meet the energy budget that we set at a minimum level of inconvenience. This scheduling vector, which we call X, is divided over k time steps of set K possible time steps. X, with chosen x, would look like this if graphically represented:

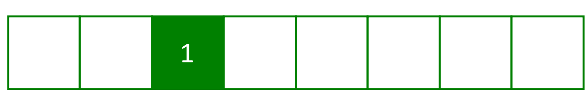

Figure 79: Example scheduling vector X, consistent with scheduling a strategy in the third time step

In order to translate the power and inconvenience profiles, we multiply Sp and Ip by X, and the result is the shifted load. This is represented graphically in the following figure:

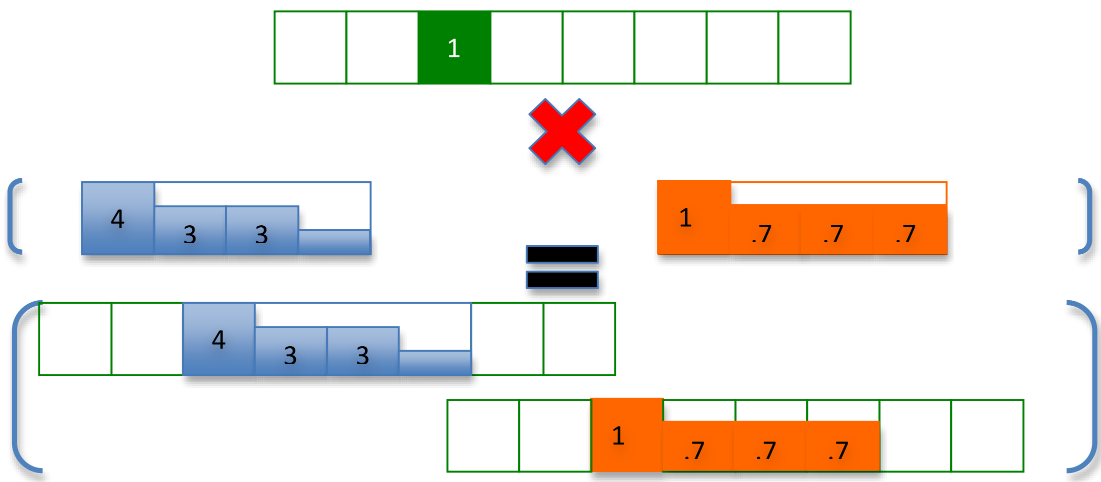

Figure 80: Load shifting procedure.

This graphical representation of the load shifting procedure is done for one "Strategy" – a strategy is defined as a series of actions that reduce power draw for multiple consecutive time steps. While the strategies talked about in this paper are for reducing load on peak, this technique could be used for continuous load management as well.

The Mixed Integer formulation of this procedure is defined below, with the following variables for each k of K strategies, t of T times, and f of Fa schedulable time blocks.

Decision Variable:
x_{kt}: Scheduled time of strategy activation for each strategy k
I_{kt}: Inconvenience Cost for each strategy k, and time t
Sp_{kf}: A Savings profile for each strategy k, denoted into each of f time steps
Ip_{kf}: an Inconvenience profile for each strategy k, denoted into each of f time steps
Ps_{ktf}: The overall power saved for each of k strategies, computed using each strategies local time, f, and global time steps t
I_{ktf}: The overall power saved for each of k strategies, over each strategies local time, f, and all global time steps t
Ps^{DR}: The required power savings for every building.

$$min \sum_{t \epsilon T} \sum_{k \epsilon K} I_{kt}$$

$$\sum_{t=ST_k}^{ET_k - nt_k + 1} x_{kt} \leq 1 \quad \forall \quad k \in K$$

$$Ps_{ktf} = Sp_{kf} x_{k(t-f+1)} \ \forall \quad k \in K, t \in T, f \in F_a : f \leq t$$

$$I_{ktf} = Ip_{kf} x_{k(t-f+1)} \quad \forall \quad k \in K, t \in T, f \in F_a : f \leq t$$

$$\sum_{k \epsilon K} Ps_{kt} \geq Ps_t^{DR} \quad \forall \quad t \in T$$

Results:
The DR planner is effective at lining up strategies in order to save a required amount of power at a minimum inconvenience cost. Below is a graph showing results from an initial test of the optimizer, with the blue line being the required necessary power use reduction at every time step in this simulation, and the green line being the scheduled power use reduction at every time step in the simulation.

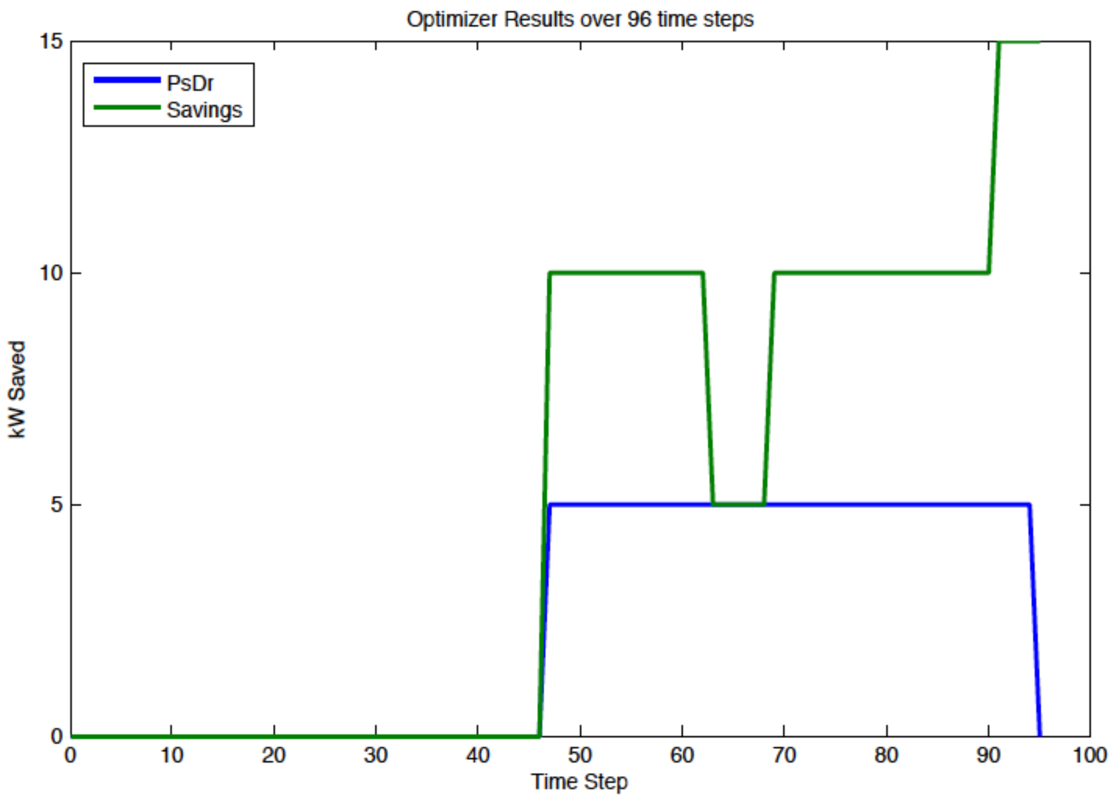

Figure 81: Initial test of the optimizer.

11.6 Appendix F: An Energy Information Gateway for use in Residential and Commercial Environments (Attached IEEE paper)

11.7 Appendix G: EnergyPlus model (Attached)

11.8 Appendix H: Rapid Audit Protocol

While plug loads can use from 10-25% of office load in commercial buildings[12], it is especially important to determine this percentage in the case of demand response, where peak load must be addressed.

Initial data exploration included breakdown of aggregate plug load use in Sutardja Dai Hall. The graphs below show these breakdowns over one day and over the entire month of October 2011.

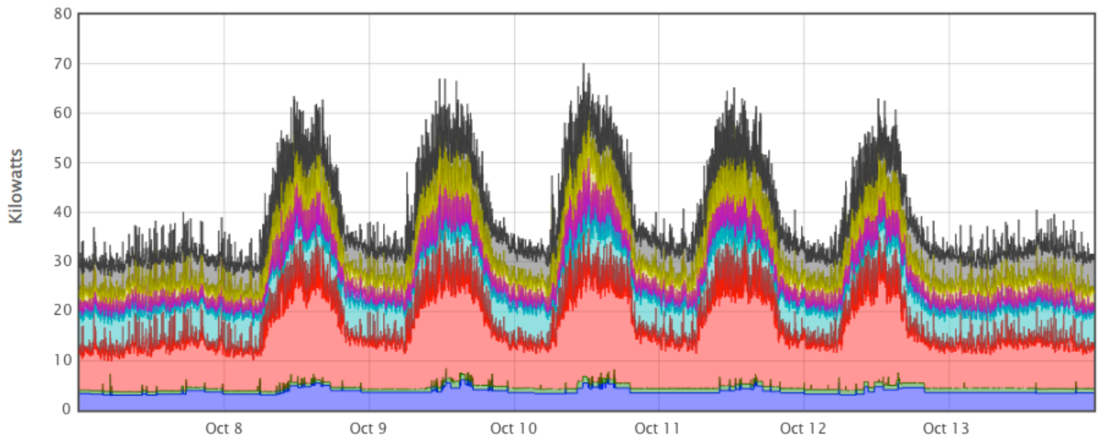

Figure 82: Plugload usage over one week: Every color is a different floor

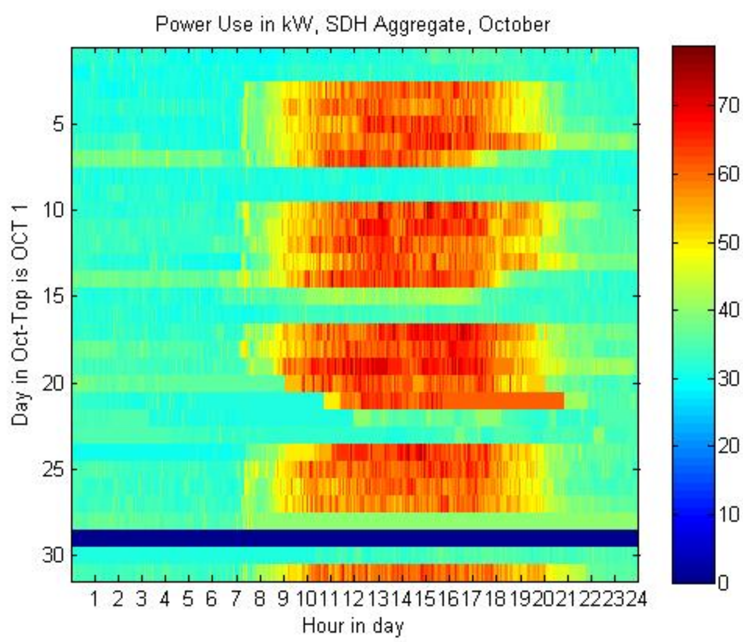

Figure 83: Plug load usage over a whole month in Sutardja Dai Hall. Every Row is a different day, with the color representing kW.

[12] http://www.energy.ca.gov/2006publications/CEC-400-2006-005/CEC-400-2006-005.PDF

This initial exploration revealed that while different floors in the building had varying electricity demands, there exists a pattern in consumption across all floors. In order to determine how to best shift power during peak times, we needed to analyze the component breakdown of this power draw, and for that, the Rapid Audit Protocol (RAP) was designed.

RAP: A design study in classifying plug load devices

In the initial stages of the DIADR project, one question that was asked was "How much do plug loads contribute to the overall energy usage of the building." While this question can be answered in a number of ways, we were interested to see how:

1) The plug loads use power in aggregate
2) Plug loads can be counted quickly and in an individual manner

Toward this end, the Rapid Audit Protocol was developed in order to assess plug load usage in Sutardja Dai Hall. The idea behind this was to make cataloguing devices faster and more efficient than has been done before, as well as enabling ease of attaching sensors to plug load devices. For reference, the procedure for adding devices to a database before RAP existed was as such:

1) Assign each device an ID
2) Assign each meter an ID
3) Record which meter is matched to which device
4) Enter this into some Excel sheet for record keeping
5) Enter the same data into some database for recording the energy used by each device
6) Publish power data to the web

Based on experience, we estimate that this takes approximately six minutes per device, and more if location information about each device is recorded along with power information. This time includes the time taken to move from room to room.

The above procedure, hereafter the static auditing procedure (MAP), has the distinct disadvantage that if the devices move, an entire new audit must be conducted in order to figure out where all the devices are in the building. Clearly this procedure is suboptimal for a world in which mobile applications allow a user to quickly change items in a database or framework. In order to do this, we must transform this static audit into a dynamic one via the use of unique tags. These tags take the form of QR codes, as exemplified below, and effectively make RAP into a cyber-physical system. Every room, device, and meter is tagged with a unique QR code, which links each device to a unique URL.

Figure 84: Example QR code

With this in mind, the RAP was designed as an Android application that allowed each item to be added to the database of possible items via the following method:

1) Before entering a room, scan a QR code assigned to the room

2) Attach QR Code to item
3) Using the RAP app, enter device data
4) Scan QR Code – device is registered

If one wants to attach a meter to an item:

1) Scan QR Code on meter and item, data starts streaming instantly to the internet.

If one wants to move a device:

1) Scan into the new room
2) Re-scan device's QR code in update mode – its position will be updated in the database.

If one wants to check for new items

1) Scan every QR code in a room, compare to existing database.

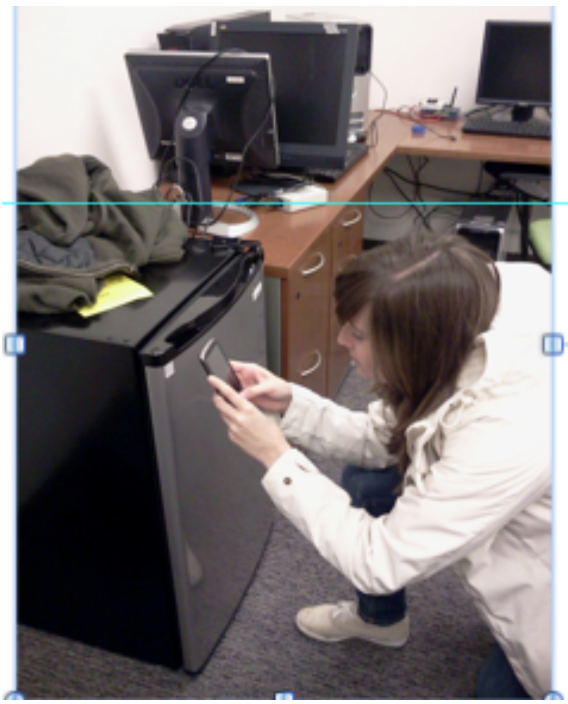

Figure 85: Student using the Android app to capture a QR code on a small refrigerator.

Given that a QR code scan takes less than ten seconds, this technique allows us to do an energy audit of a building at least four times faster than an audit with a clipboard, with fewer mistakes and greater reproducibility. Overall, in our audit, it took 2 minutes per device for 700 devices in a commercial building, *including moving from room to room*.

Audits and data

This information, when combined with an analysis of the total plug load draw in Sutardja Dai Hall, paints a powerful picture of energy usage

AUDIT 1: Complete audit of the building

The first audit was intended to get a breakdown of devices within Sutardja Dai Hall in a complete manner so that estimating future usage would be a trivial task. Given that there were serious time constraints on the auditors, it was considered difficult to maintain a live audit, but one complete audit was possible. Below is a screenshot from the original auditing application, used to create the original database of devices in the building.

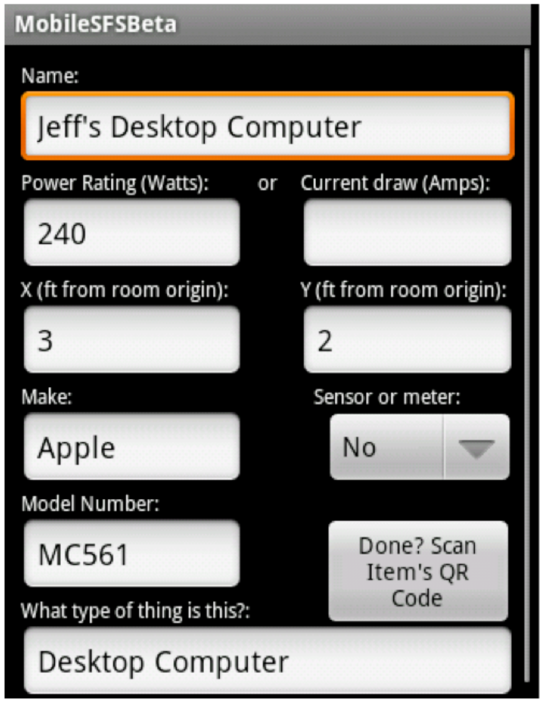

Figure 86: Main screen of RAP V 1.0

This first audit was intended to be largely complete, and also show the possibility of producing a feedback system for users in an office to understand how their energy consumption is in relation to their officemates. Below is an example of this feature of the energy audit. It should be noted that the only thing that was made by hand in this image was the picture of the office from an overhead view – the devices and power draws were extracted from the database. (Streaming data was drawn from StreamFS – a streaming file system for sensors).

Figure 87: Display of the floor plan of test lab in Sutardja Dai Hall, colors over appliances represent real-time power draw.

Attached is the final count from the original audit. The results present in this audit are summarized in the graph below.

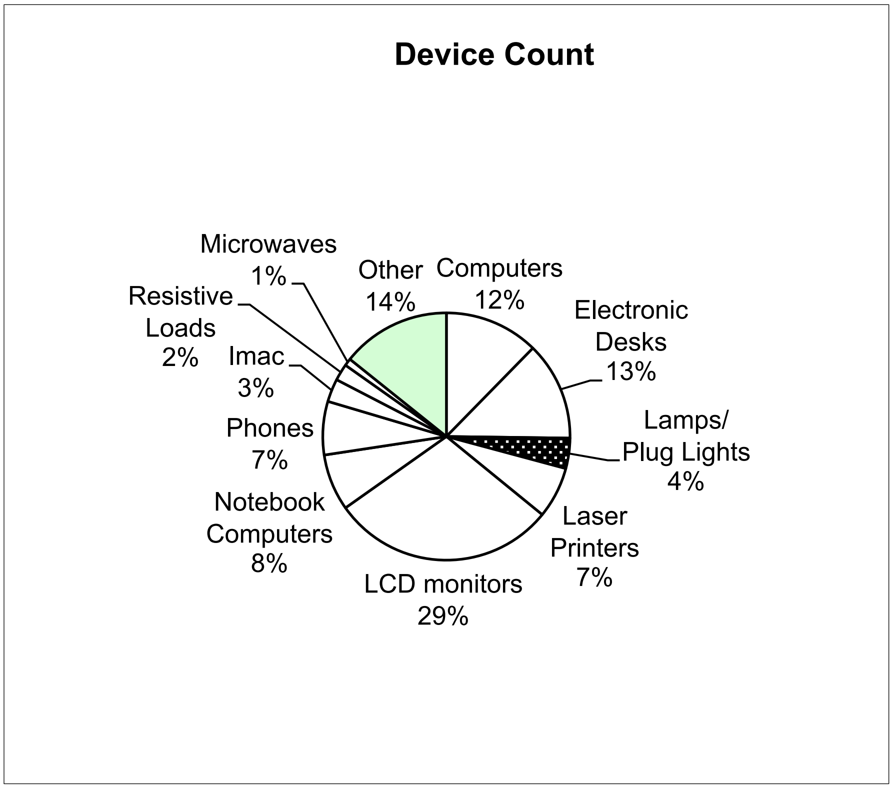

Figure 88: Pie chart of devices in SDH at first count.

AUDIT 2: Re-design for web-based auditing

After the initial audit, the auditing process was streamlined, and a smaller test was done on a design basis for accelerating the process. This audit was streamlined for improved ability to re-count devices and link items to sensors in a streamlined fashion. With this new re-design, it was possible to scan any device with a QR code attached and see its previous power usage up to the present time. This was a novel innovation, and allows users to see their power usage in real time.

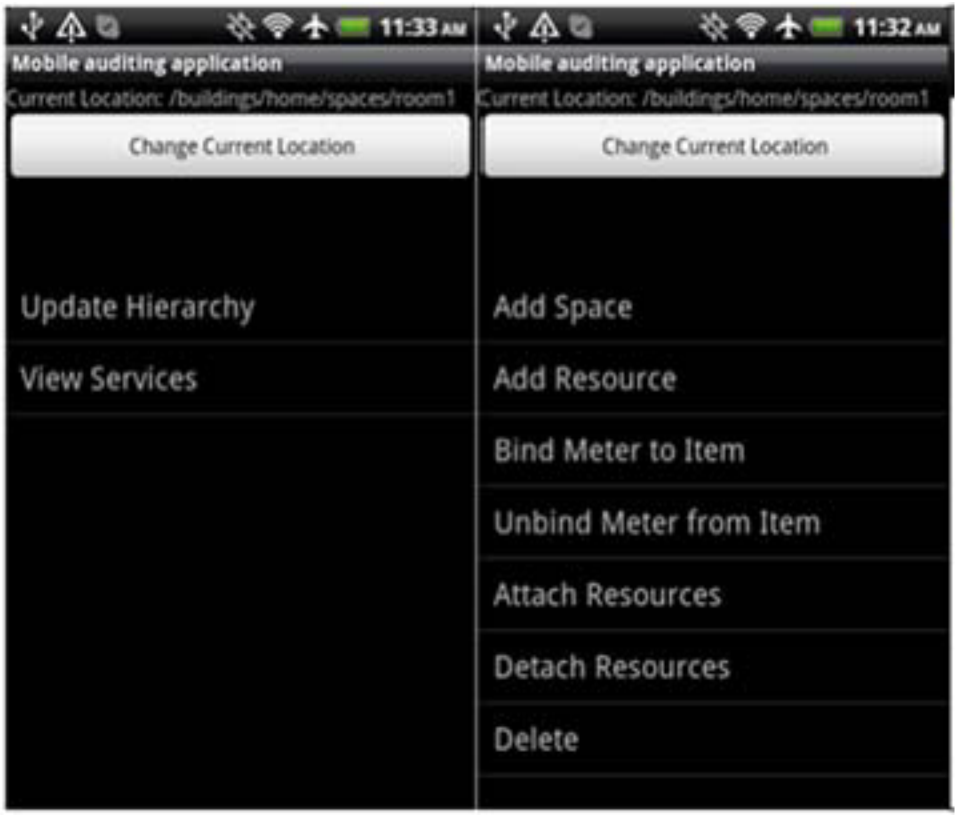

Figure 89: Interface of RAP 2.0

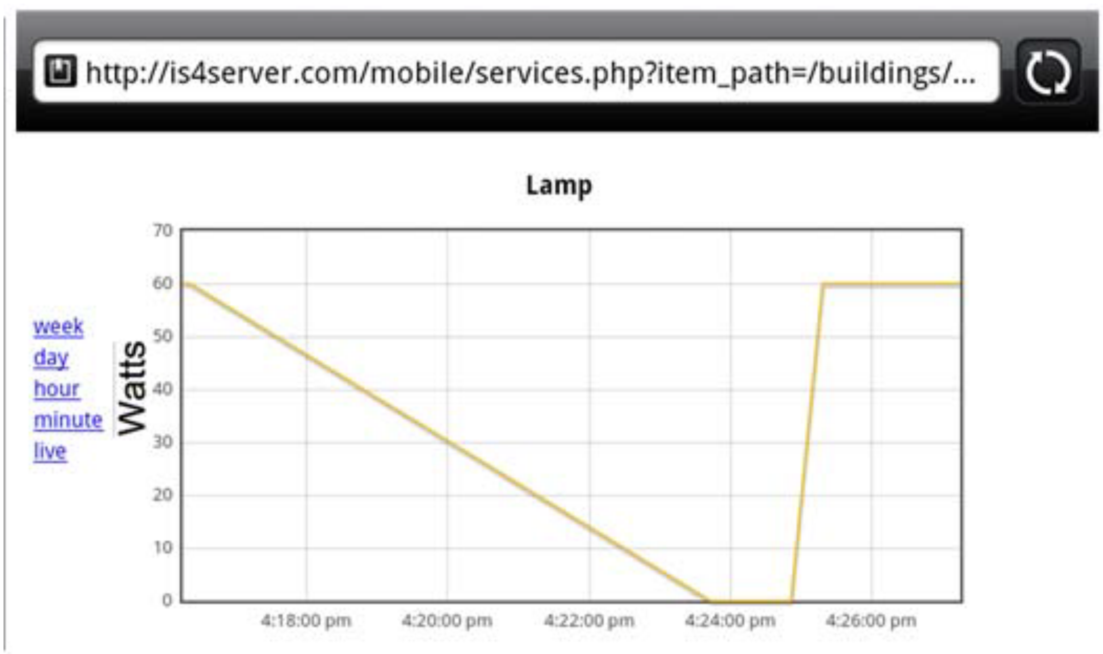

Figure 90: Web based interface for historical power use, RAP 2.0

Recommendations:

Three main kinds of devices exist in Sutardja Dai Hall: Devices which are on all the time (Baseload), devices which are on part of the time but in a persistent pattern (Intermittent load), and devices which are on in a stochastic fashion (Peak loads). As far as our ability to modulate which devices draw power during peak hours, the key recommendations are:

- Baseload: Encourage users to charge laptops before peak hours, thus shifting load, in addition, turn off frivolous loads such as the technology museum in SDH: This can shift approximately 6 kW off of peak
- Intermittent load: Reduce the amount of redundant LCD displays during peak hours – this can shift approximately 5kW off of peak
- Stochastic loads: Encourage users to print before or after peak hours, and to make coffee or tea early in the morning and not during peak. This will prevent unintended peaking events during a demand response event.

The figure below illustrates these kinds of loads in a picture. As we can see, the control strategies implemented were appropriate for DR in this case, and should be noted for other office buildings.

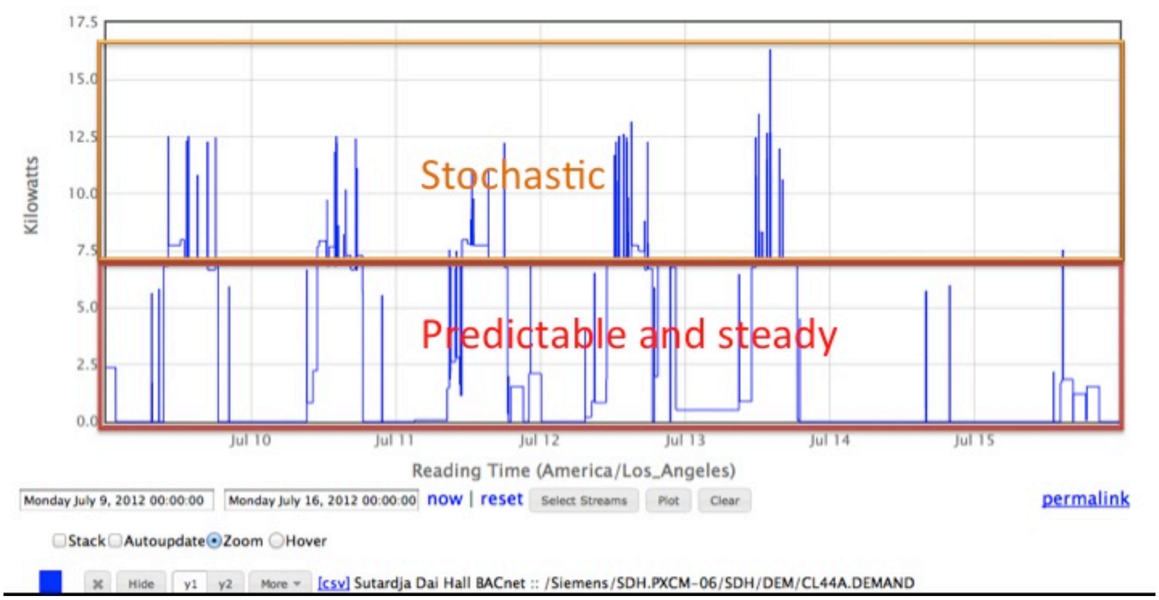

Future directions:

RAP highlighted appropriate control strategies appropriate for the Sutardja Dai Hall, yet could be used in a wider scenario to highlight plug load use in any building. Given that each building will contain a difference plug load device ecosystem, this is a viable, scalable strategy to quickly analyze the present scenario in a building or place of work. It would be interested to see how a more modular application could help RAP become more diverse, potentially opening the door to more kind of plug load audits in a highly customizable fashion.

11.9 Appendix I: Commercialization Plan (Attached)

11.10 Appendix J: Smart Energy Box Runtime Components

11.10.1 SEB Runtime Architecture

The Runtime Components are the functional units of smart energy box which implement the smartness provided by the Box. All Runtime components should derive from the common interface provided by the framework. One Runtime component cannot reference another runtime component directly by design. They must always have an interaction through the data repository. This increases the stability, scalability and reliability of the system.

Runtime component can simply kill the performance of the system if they employ significant amount of processing, so it is highly recommended to run non deterministic behavior in a separate thread and synchronize with repository at a valid point, which should be decided by the runtime component itself. For example the Demand response Client Runtime component polls for the demand response event from the Demand response Server. In order to maximize the Runtime manager turnaround time, we make the polling of demand response event in a separate thread and synchronize with repository whenever there is a DR event retrieved.

Each Runtime component should define the services it provides and its data format to use them by other components. One should refer the documentation to find the services offered by the runtime component. Runtime components can be designed in three ways

- with an embedded business logic implementation
- with business logic running in separate process (an external application) that runs on the same machine
- with business logic in a separate component that runs over the network (internet/intranet)

11.10.2 Runtime Components

11.10.2.1 BACnet Adapter Runtime Component

BACnet Adapter runtime component interacts with the building management system to get or set any centrally controlled devices data through BACnet/IP communication protocol. This component is part of SEB Core.

11.10.2.2 OpenADRClient Runtime Component

OpenADRClient runtime component is responsible for getting DR events from the demand response server and making them available for Smart Energy Box components. Strategy Selector and Energy Simulation Runtime components accesses DR event information and proceed further on identifying and implementation of building energy strategies. This component is part of SEB Core.

11.10.2.3 Weather Data Adapter Runtime Component

WeatherData Adapter runtime component retrieves weather data from publicly available weather data service(s), converts and stores in hourly format. This component is part of SEB Core.

11.10.2.4 Strategy Selector Runtime Component

This component selects the strategy based on the current building strategy mode and DR event. If building is configured to Instant strategy mode, this component simply gets the configured static strategy for the current DR event category (High or Moderate). If building is configured to PeakDayPricing mode, this component involves in the negotiation with the Manager, evaluates all available strategies for the current DR mode and selects the best strategy. Once it receives agreement from the manager then it creates strategy execution request.

11.10.2.5 Strategy Executor Runtime Component

The strategy executer runtime component does what its name states. It is responsible to execute the strategy execution requests. It keeps track of the time and updates/resets building set points according to the strategy schedule. The maximum time that a strategy can last is 24 hours. If selected strategy is applicable only for the DR period then it will be reset after DR event is finished. If strategy is configured for 24 hours then it executes for whole 24 hours. In case of Peak Day Pricing mode, 24 hour strategies are executed, since strategies involve pre event and post event (pre-cooling, pre-heating and post-event operations) configurations.

11.10.2.6 Energy Simulation Component

This component is responsible to perform energy simulation of the supplied strategy and provide results for the users. Users are any other runtime components in the smart energy box. Energy simulation is performed by EnergyPlus with the real-time weather forecast. The runtime component reads all requests for energy simulation during its turn of run by the runtime and passes them in to a separate thread to perform simulation. After performing the simulation, results will be passed back to Energy Simulation Runtime component, which puts in to the repository for the access by the requestors. This component is part of SEB Core.

11.10.2.7 Manager Runtime Component

This is a runtime component that works on a collective intelligence from all participants to control building load. Participants are nothing but HVAC System, Central Lighting, Distributed plug load control systems. This component uses market based adaptive approach to negotiate and decide load shedding goals during DR event.

Participants are dynamically registered with manager to participate in the negotiation. There are two types of participants

1. Participants that receives load shedding target on notification of DR event
2. Participants that receives load shedding target every few minutes during DR event

There are several methods identified to implement collective intelligence as mentioned below, however in DIADR we decided to implement Adaptive Partial-Centralized Load Allocation mechanism.

- Adaptive Centralized Load Allocation
- Adaptive Partial-Centralized Load Allocation

- Non-Cooperative Market-based Load Allocation
- Non-Cooperative with Group Interruption Cooperation Market-based Load Allocation
- Partial-Cooperative Market-based Load Allocation
- Peer-to-Peer Market-based Load Allocation

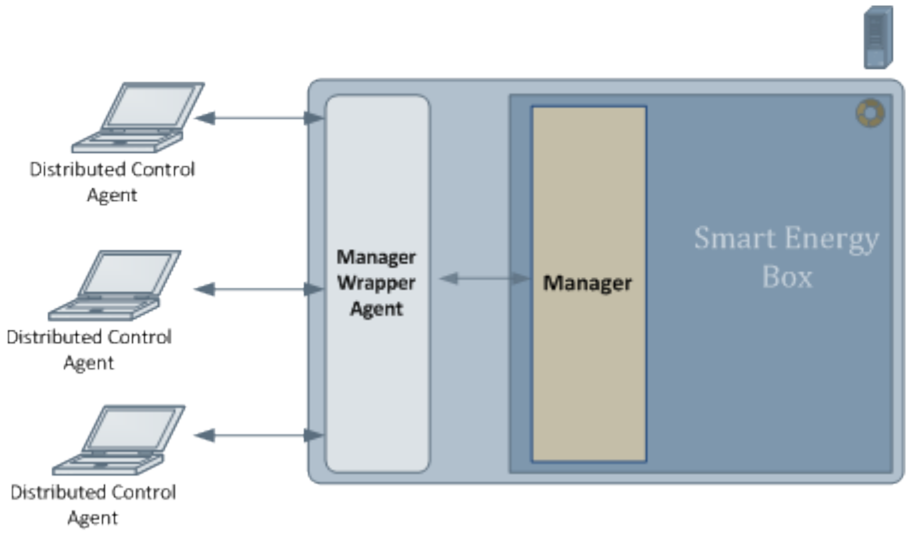

Figure 91: Manager Runtime Component

In Adaptive Partial-Centralized Load Allocation mechanism, Participant is mainly responsible to manage appliances, consolidate cost function and meet load setting goal received from SEB.

11.10.3 Integration with JADE (Agent Based System)

In order to extend SEB capabilities to all office equipment (plug load) in the building, an agent based mechanism (FIPA compliant) called JADE (Java Agent Development Environment) has been adapted. JADE is a FIPA compliant open source agent based platform, which supports variety of communication protocols, standards to implement effective market based solutions among multiple parties in the building such as HVAC, Lighting, Plug Loads, Occupancy etc. using collective intelligence.

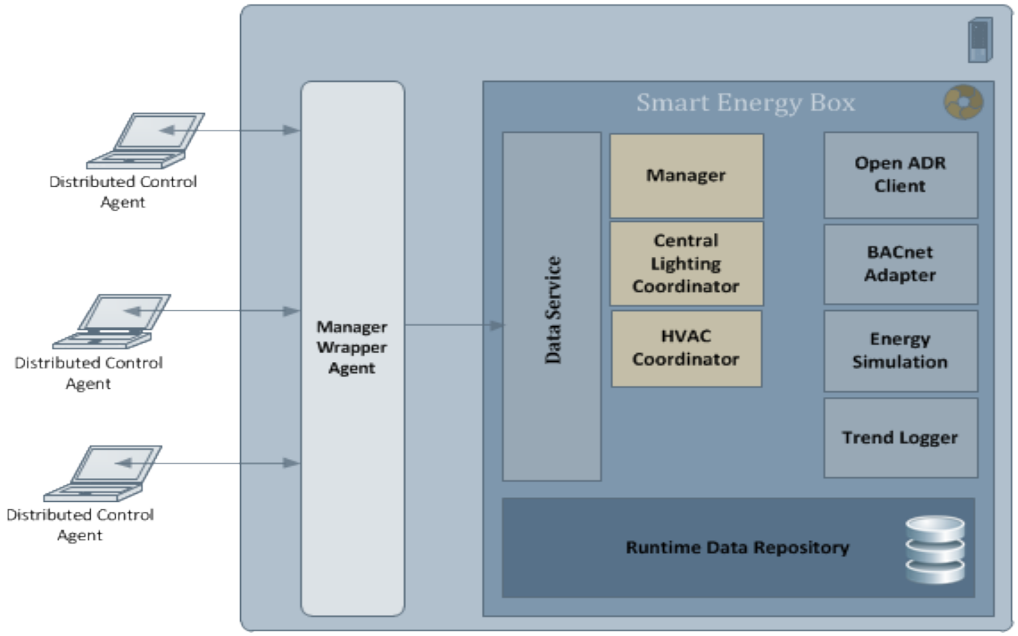

Figure 92: **SEB and JADE Integration**

11.11 Appendix K: SCR Gateway details

11.11.1 Main Components

The Gateway software runs in the OSGi environment. The OSGi Gateway agent consists of three main components:

- DataAccessLayer
  Reads and writes the configuration and other persistency data of the Gateway to an XML-file, creates the appliances objects from the XML-file and provides the newly created appliances to the ApplicationControl

- ApplicationControl
  Provides control over the appliances, Reads sensor data from the sensor box and optimizes the load shedding

- GWAgent
  Communicates with the ManagerWrapper using JADE and Takes action in the load shedding process

11.11.2 Additional Components

Additional components are: Utility, Raritan, XBee, and JADE.

- Utility
  It is comprised of all the necessary interfaces for OSGi services, all of the appliances classes and all the classes for communication with the SEB

- Raritan
  It is used to turn on and off outlets where appliances are connected, reads the power consumption of connected appliances. This kind of device is necessary for devices which cannot perform these action on their own

- XBee
  It is used to communicate between a Gateway and a sensor box. XBee.API, XBee.Lib.Rxtx and XBee.Lib.Log4jConfig belong together
- JADE
  The Gateway Agent is running on the JADE platform and requires the JadeRuntimeService and The OSGi add-on for JADE is needed to have the JadeRuntimeService
- Printer
  It is used to pause/ resume the printing jobs

- LaptopBattery
  It is used to get the battery status of the laptop

- Web UI

Web-based frontend that can be opened from Web browser like IE, Firefox or Chrome (Figure 88 shows a screen shot of the Web UI frontend). It also has a backend serving data request and configuration change from the frontend

The Web User Interface is split into two parts. The left side shows the general Gateway information, sensor data, and DRAS information. The right side shows the appliance information. The appliances status is monitored and its participation to demand response can be configured.

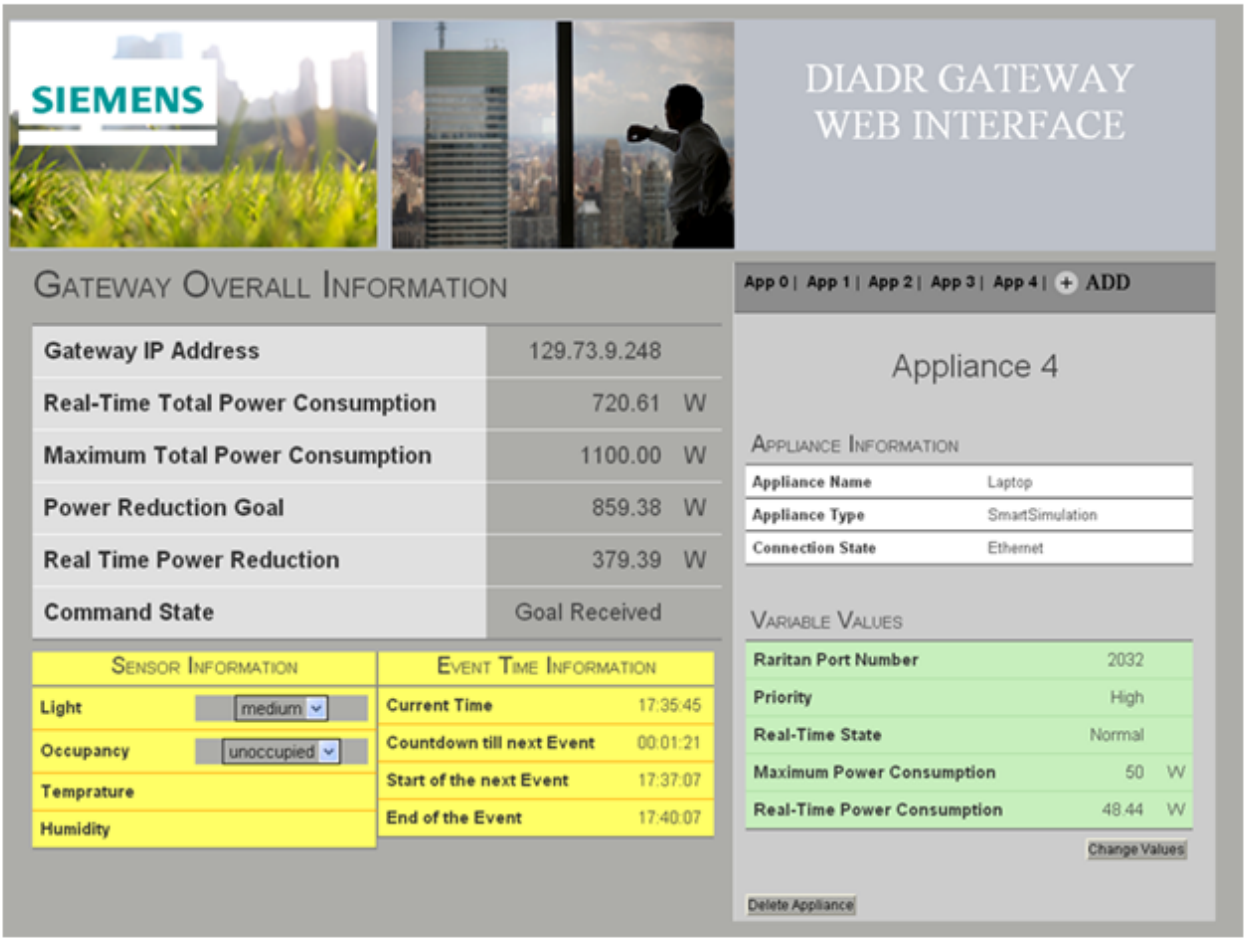

Figure 93: Web UI Front End

The Web User Interface communicates with backend using web service hosted by the Gateway. They exchange data using JSON file.

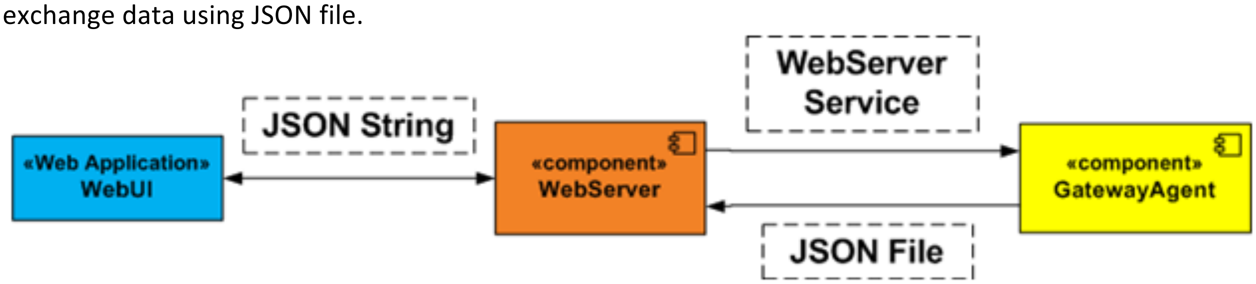

Figure 94: Web User Interface communication

11.11.3 Gateway JADE Agent

JADE is a FIPA compliant open source agent based platform, which is used to bring Gateway plug load control to integrate with Smart Energy Box. Using JADE API Gateway agents (GWAgent) are created with set of behaviors, which can be communicated with other agents in the system. The agent behavior is used to model the Finite State Machine (FSM) in part I: Control Logic of the Gateway. The way JADE communication can be seen in part 2: Communication and Control Sequence.

11.11.3.1 Control Logic of the Gateway

The Gateway uses an adaptive partial centralized algorithm. In this algorithm, a Gateway will follow the following states:

- Reset appliances - The Gateway will start by resetting the appliances so they are in a known state
- Reset values - resets some values to their standard value
- Receive command - receives the bid request demand from the ManagerWrapper
- Read data - reads specific data from its appliances (e.g. power consumption, priority)
- Adjust priorities - calculates the power consumption and adjusts the priorities
- Make decision - decides whether or not to apply control
- Wait - After receiving the power reduction goal it must wait until the event starts
- Apply control - when the event is started it optimizes the power consumption of its appliances and applies control to them
- Adjust priorities - Reads specific parameters of the appliances and the reduction in power is calculated
- Report result - reports to the SEB what actions were taken and what their results were
- Receive command - the Gateway waits for the next command

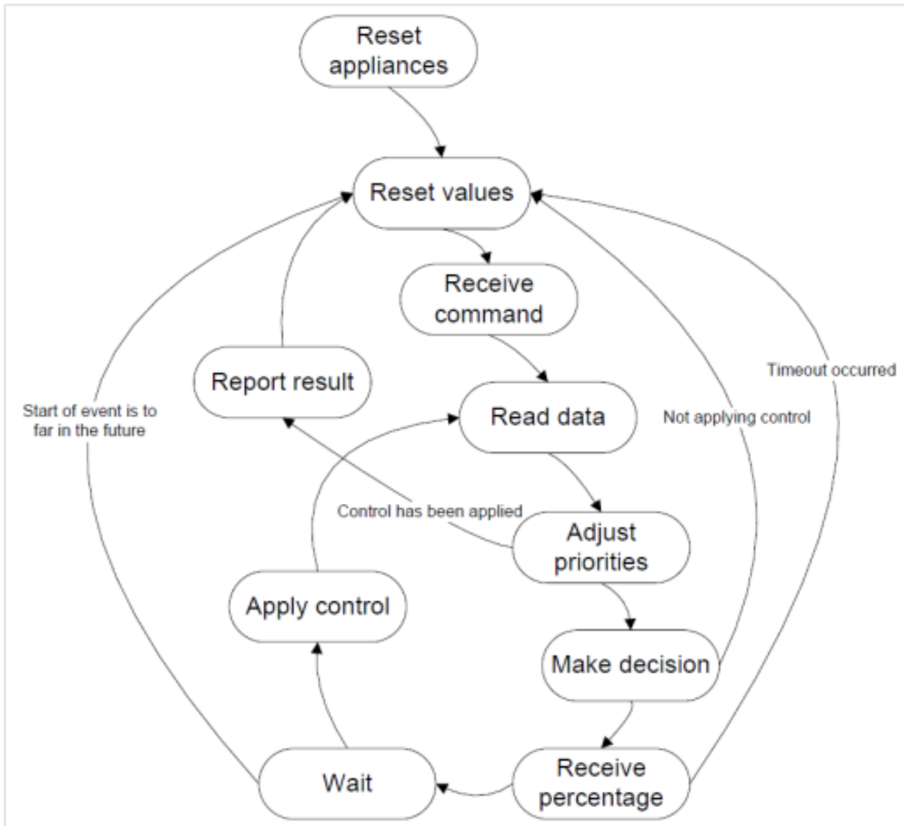

Figure 95: Finite state machine of the behavior of the Gateway Agent

11.11.3.2 Gateway Control Sequence

The communication between the SEB and all Gateways happens through its ManagerWrapper. The ManagerWrapper is a distributed load control coordinator, which is also a JADE agent who knows both SEB and all Gateway agents. It continuously poll for Bid Requests from SEB and passes them to corresponding Gateway agents and brings the responses from Gateway agents back to SEB. The following sequence diagram shows the communication sequence between SEB , manager Wrapper and Gateway Agents.

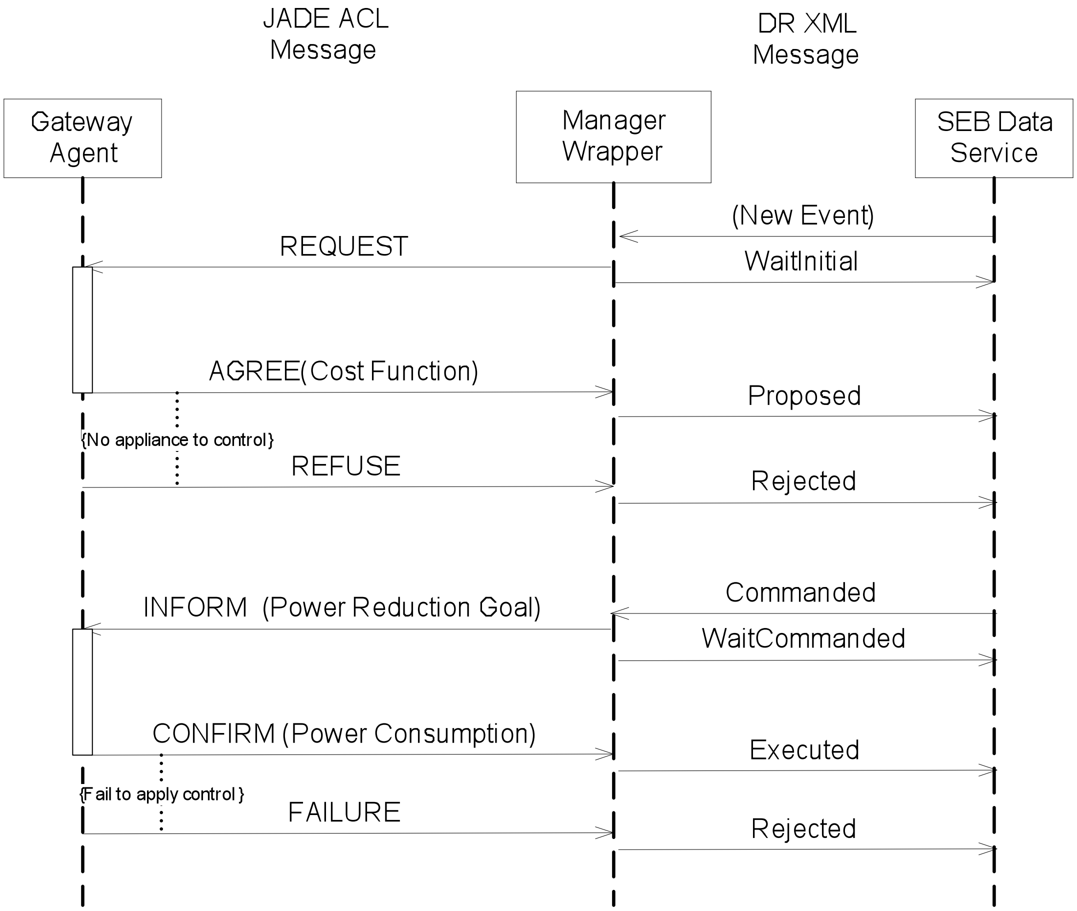

Figure 96: Communication between Gateway Agent, ManagerWrapper and SEB Runtime

11.11.4 Gateway Configuration

In order to run Gateway properly, the two configuration files must be correct, the JADE configuration file and the GWAgent configuration that must be correct.

1. The Jade configuration file– consists of 4 properties as illustrated below.

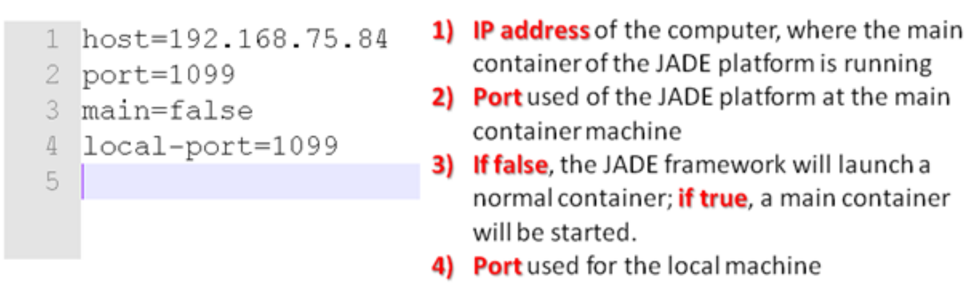

Figure 97: JADE configuration file and parameter definition

2. The GWAgent configuration contains a description of the Gateway and the devices that are under its domain. The required attributes are explained below.

- A unique agent name is required(line 2)
- All Appliances must have these attributes
    - “MaxPower” – describes the maximum power that a device can use.(Line 5)
    - “States” – describes the states a device can be in which must be between 0 and 1. There are three different types on/off, linear states(Ex:0,0.25,0.5,0.75,1) or discrete(Ex: 0,0.3, 1). The state 1 is automatically generated and indicates that the device is operating at max power and a state of 0 meaning the device is off.(Lines 6-8)
    - A unique id(Line 4)
    - A priority has to be set. This value must be greater than or equal to 0. If not set it will default to 0.(Line 4)
- There are more attributes which are device type specific.(see example XML)
- The Raritan and sensor box are also defined in this file. (Line 43 - 52)

```
1.  <?xml version="1.0" encoding="UTF-8" ?>
2.  <GatewayConfiguration name="GatewayAgent1">
3.        <ApplianceList>
4.              <Light id="1" priority="1" outlet="1">
5.                    <MaxPower>60</MaxPower>
6.                    <States>
7.                          <State>1</State>
8.                    </States>
9.              </Light>
10.       <Printer id="2" priority="3" outlet="3">
11.             <Name>HP LaserJet Professional P1102w</Name>
12.             <MaxPower>400</MaxPower>
13.             <States>
14.                   <State>1</State>
15.             </States>
16.       </Printer>
17.       <Laptop id="3" priority="2" outlet="4">
18.             <Domain></Domain>
19.             <Username>Siemens</Username>
20.             <Password>Test!234</Password>
21.             <Hostname>127.0.0.1</Hostname>
22.             <MaxPower>73</MaxPower>
23.             <States>
24.                   <State>1</State>
25.             </States>
26.       </Laptop>
27.       <SmartSimulation id="4" priority="4">
28.             <EthernetConnection write="true" delay="false"
    interrupt="true">
29.                   <ConnectionName>WiFiAppliance1</ConnectionName>
30.                   <Port>2004</Port>
31.                   <Timeout>5000</Timeout>
32.             </EthernetConnection>
33.             <MaxPower>500</MaxPower>
34.             <States>
```

```
                  <State>0.2</State>
                  <State>0.4</State>
                  <State>0.6</State>
                  <State>0.8</State>
                  <State>1</State>
            </States>
      </SmartSimulation>
</ApplianceList>
<Raritan>
      <Hostname>192.168.75.202</Hostname>
      <Username>admin</Username>
      <Password>Siemens</Password>
      <Unitname>Raritan1</Unitname>
</Raritan>
<SensorBox>
      <Devicename>sensorBox</Devicename>
      <Port>COM7</Port>
</SensorBox>
</GatewayConfiguration>
```

11.12 Appendix L: Details of OSP generation

Stage I: For any weather pattern *i* (*i*=1, 2, 3, …, 19), DR strategy *j* (*j* =1, 2, 3, …, 5250) will be evaluated by EnergyPlus simulation. The hourly HVAC energy and hourly PPDs of all zones are calculated by the simulation. A simplified peak day price model (PDP) is applied to calculate the energy cost (*C*).

The 24-hour PPD values of the occupied zones are summed up to generate the "thermal comfort loss" (*U*). After min-max normalization, normalized energy cost ($\bar{C}$) and thermal comfort loss ($\bar{U}$) will time their corresponding weights (w_c and w_u, respectively), and add together to provide the objective value $F_{i,j}$(Equation (1)). Exhaustive Search (ES) and other optimization algorithms can be applied on this stage, to obtain the optimal objective value $F_i^* = \min_j F_{i,j}$.

$$F_{i,j} = w_c \frac{C_{i,j} - C_{min}}{C_{max} - C_{min}} + w_u \frac{U_{i,j} - U_{min}}{U_{max} - U_{min}} \tag{1}$$

$$\text{Min: } F(W,S) = w_c \frac{C(W,S) - C_{min}}{C_{max} - C_{min}} + w_u \frac{U(W,S) - U_{min}}{U_{max} - U_{min}}$$

Stage II: On Stage I, the evaluations of all 5250 DR strategies are done for each of 19 centroid weather profiles, using ES or Genetic Algorithm (GA). On Stage II, for each weather pattern *i*, strategy *j* will be selected if it satisfies Equation (2).

$$F_{i,j} \le \alpha F_i^* \tag{2}$$

where α is the pre-determined threshold, and α=1.1 is used in this study. Denote the total number of selected strategies by N_i. The selected strategies are sorted by ascending objective values. And let j_k be the *k*-th strategy in this rank (*k* = 1, 2, 3, …, N_i). Each of the selected strategies will be assigned with a "likelihood" score (*L*). The likelihood scores are determined by Equation (3) and (4).

$$\sum_{k=1}^{N_i} L_{i,j^k} = 1 \tag{3}$$

$$\frac{L_{i,j^{k+1}}}{L_{i,j^k}} = \beta, k = 1, 2, 3, \dots, N_i - 1 \tag{4}$$

where β is the pre-determined ratio, and β = 0.5 is used in this study. The overall likelihood score of strategy *j* for all weather patterns is given by Equation (5):

$$L_j = \sum_{i=1}^{I} p_i L_{i,j} \tag{5}$$

where, *I* is the total number of weather patterns, and p_i is the probability that the weather of the planning day is of pattern *i*. p_i can be estimated by dividing the number of pattern *i* days in the record with the total number of recorded days.

An "optimal strategy pool" can be created by selecting candidate strategies with large overall likelihood score. As this pool will contain a smaller number of candidate strategies, exhaustive search within the pool can provide the best solution for a given weather condition; and this search can be conducted on-line. The following table shows the OSP created by this algorithm.

The optimal strategy pool

DR strategy ID	Overall likelihood score	GTA (T1)*	GTA (T2)*	GTA (T3)*	SAT*	SFP*
4818	8.2381	9	10	19	60	1150
4755	4.2000	6	7	19	60	1150
4817	2.3963	9	10	18	60	1150
4800	1.8857	8	9	19	60	1150
4728	0.6667	5	6	19	60	1150
4779	0.6095	7	8	19	60	1150
4823	0.5039	9	12	18	60	1150
4820	0.252	9	11	18	60	1150
4826	0.126	9	13	18	60	1150
4754	0.0667	6	7	18	60	1150
4829	0.0315	9	14	18	60	1150
4827	0.0157	9	13	19	60	1150
4824	0.0079	9	12	19	60	1150

11.12.1.1.1 OSP validation by simulation

To validate the OSP algorithm, 13 historical August days of Berkeley, CA were randomly sampled for testing. The result of OSP was compared with the ES optimization and other two on-line optimization algorithms – GA and pattern based strategy selection (PBS). The optimal DR strategies obtained by each algorithm as well as total number of evaluations are shown in the following table.

Validation of the Optimal Strategy Pool

Sample Day (Pattern)	Algorithm				
	ES	PBS	GA	OSP	
1 (pattern 19)	4728	4755	4728	4728	Opt. DR strategy ID
	5250	0	288	13	# of on-line evaluations
2 (pattern 3)	4818	4818	4818	4818	Opt. DR strategy ID
	5250	0	349	13	# of on-line evaluations
3 (pattern 9)	4818	4818	4818	4818	Opt. DR strategy ID
	5250	0	358	13	# of on-line evaluations
4 (pattern 8)	4779	4800	4755	4779	Opt. DR strategy ID
	5250	0	362	13	# of on-line evaluations
5 (pattern 14)	4818	4818	4779	4818	Opt. DR strategy ID
	5250	0	380	13	# of on-line evaluations
6 (pattern 13)	4800	4818	4800	4800	Opt. DR strategy ID
	5250	0	383	13	# of on-line evaluations
7 (pattern 16)	4755	4755	4755	4755	Opt. DR strategy ID
	5250	0	308	13	# of on-line evaluations
8 (pattern 8)	4800	4800	4755	4800	Opt. DR strategy ID
	5250	0	311	13	# of on-line evaluations
9 (pattern 18)	4755	4818	4755	4755	Opt. DR strategy ID
	5250	0	326	13	# of on-line evaluations
10 (pattern 18)	4755	4755	4755	4755	Opt. DR strategy ID
	5250	0	311	13	# of on-line evaluations
11 (pattern 17)	4755	4755	4755	4755	Opt. DR strategy ID
	5250	0	339	13	# of on-line evaluations
12 (pattern 14)	4800	4818	4779	4800	Opt. DR strategy ID
	5250	0	361	13	# of on-line evaluations
13 (pattern 16)	4755	4755	4755	4755	Opt. DR strategy ID
	5250	0	323	13	# of on-line evaluations

By applying PBS, the optimal strategy for sample weather is assumed to be the same for the weather pattern it belongs to. PBS does not require on-line simulation evaluation of DR strategies, therefore, it seems to be a perfect on-line optimization algorithm. However, according to our result, PBS algorithm fails to identify the optimal strategy for Sample Day 1, 4, 6, 9 and 12. GA performs slightly better, as it fails for Sample Day 4, 5, 8 and 12. As a contrast, OSP successfully identifies the optimal DR strategy for all sample days. And furthermore, OSP only requires 13 on-line simulation evaluations, which is 3.5% of evaluations by GA, and 0.2% of evaluations by ES. It takes about 2 minutes to finish one simulation execution (on a personal PC laptop). This infers that ES needs 7.3 days, GA needs 12 hours, but OSP only needs less than 30 minutes to obtain the optimal DR strategy for a given weather profile.

This result shows that the OSP algorithm has superior capability of identifying optimal DR strategies and reducing on-line computation load, simultaneously. And the optimal strategy is able to reduce the HVAC peak load by 18% for a typical hot August day of Berkeley, CA (the following figure).

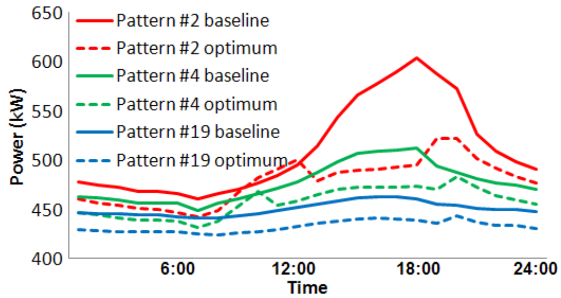

Figure 98: Simulated results of various strategies for different weather patterns.

11.13 Appendix M: Distributed Intelligent Automated Demand Response Enhanced Scale Test Plan

11.13.1 Scope

This document describes the enhanced scale test plan of DIADR project. Enhanced scale testing is to extend building control from Single room to whole building. The testing involves the following cases

1. Communication from Demand Response Server to Smart Energy Box
2. Communication from Smart Energy Box to Building Manager System
3. Identifying HVAC strategy based on real-time weather forecast and energy simulation
4. Controlling Building HVAC system by Smart Energy Box
5. Controlling Building Central Lighting by Smart Energy Box
6. Controlling Building Plug load Control through Local Energy gateway
7. Market based adaptive negotiation during Demand Response Event

11.13.2 Key notions

11.13.2.1 DR signal

DR signal is sent out by Demand Response Automation Server (DRAS) to the buildings at the "event notification time" with the information such as of peak day pricing, start and end of DR period. For this demonstration, we use next day DR signal and instant DR signal

11.13.2.2 DR strategy and baseline strategy for HVAC control

A strategy for HVAC control is a collection of HVAC system setpoint values for certain duration of time. For this demonstration, a strategy includes the "supply air temperature" setpoint, the "supply fan pressure" setpoint, and "zone temperature" setpoint, through 24-hour period of the day. The baseline strategy has fixed values for all the above setpoints for entire day, shown in Table 3.

Table 3: Baseline HVAC control strategy

	Value
Supply air temperature setpoint	56 °F
Supply fan pressure setpoint	1350 Pa
Zone temperature setpoint	70 °F

11.13.2.3 Optimal strategy pool

Optimal strategy pool is a collection of DR strategies for HVAC control designed based on the weather patterns of Berkeley, CA. The performances of these DR strategies are to be evaluated by energy simulation module in SEB with the current weather forecast, and the one with best outcome (maximal peak load reduction and minimal thermal comfort loss) will be identified and applied to the system. For this demonstration, the optimal strategy pool has 13 candidate DR strategies.

11.13.2.4 Base load

The base load is defined as the total building power consumption when the optimal DR strategy for HVAC control is applied. Base load is predicted at 11AM for 1PM to 5PM based on the predicted weather and historical energy consumption of a building.

11.13.2.5 *Peak load reduction target*

For this demonstration, a target of up to 30% peak load reduction is defined. The peak load is defined as the maximum of the projected (base load) total building energy load during DR period. The building energy projection is calculated based on weather and historical energy consumption of the building.

11.13.2.6 *Negotiation cycle*

SEB negotiates with Distributed Load Control (DLC) to perform load shedding during the DR event in standard cycle. Negotiation cycle can be configured from 3 minutes to 15 minutes based on the system response requirement. Typical negotiation cycle is 15 minutes. If negotiation cycle configured as 15 minutes, every 15 minutes starting from DR event SEB negotiates with participating DLC agents, and commands the optimized load shedding targets for DLCs.

11.13.3 Test Layout

The following figure shows the test Layout for enhanced scale testing of DIADR project at UCB. This plan includes Central HVAC, Central Lighting and experimental plug load control.

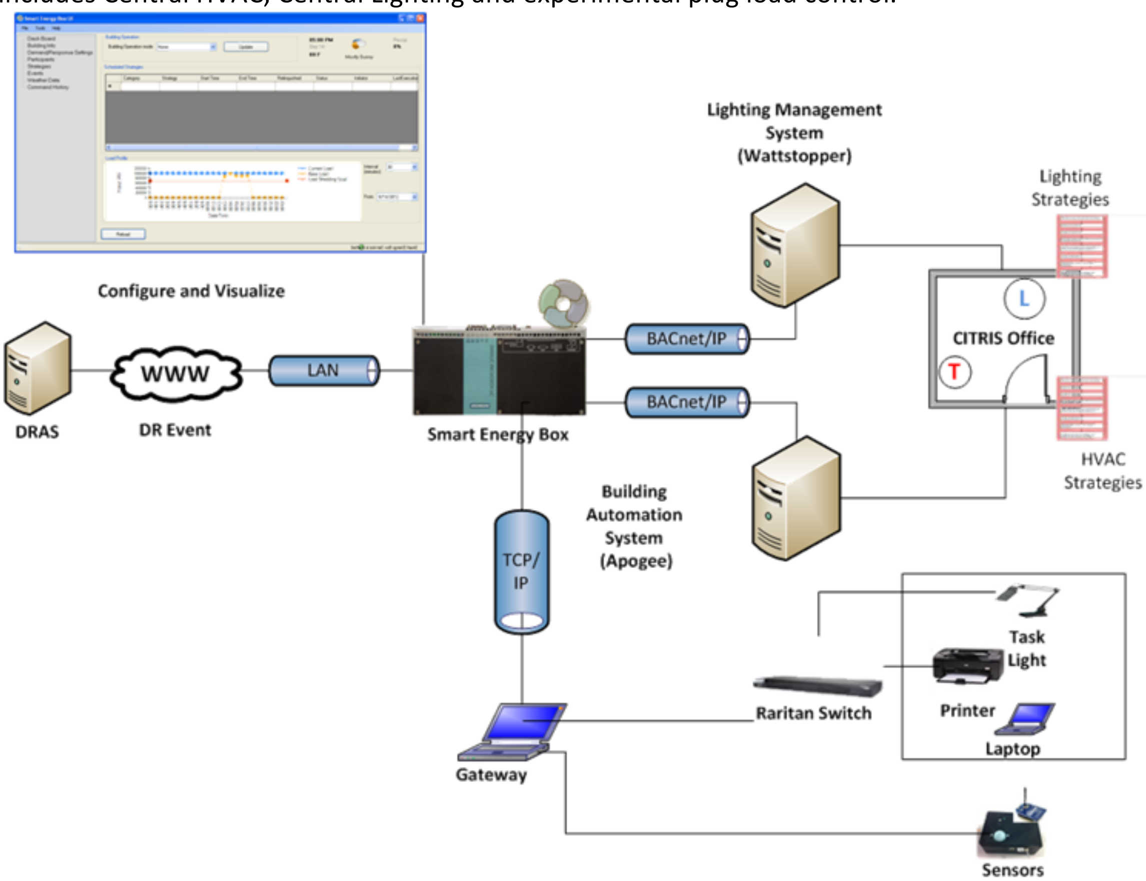

Figure 99: Test layout of various components

11.13.3.1 *HVAC Strategies for Day Ahead DR Event*

All 13 strategies identified for Berkeley weather are attached as a separate appendix.

11.13.3.2 Instant DR mode HVAC Strategies

During instant DR event, there will not be any strategy evaluation performed; however system allows building owner to create fixed set points for the duration of DR event

DR Mode	Room 456 Thermostat Set Point (F)
Moderate	68
High	66

11.13.3.3 Lighting Strategies

Lighting strategies from Level 2-7 are attached as a separate appendix.

11.13.3.4 Peak day pricing event

DR Event property	Definition
Event notification time	9:00 P.M., Prior day of DR Event
DR Mode	High
DR period start	2:00 P.M., Day of DR Event
DR period end	6:00 P.M., Day Of DR Event

11.13.4 Test cases

11.13.4.1 Establishing connection between DRAS and Smart Energy Box

The goal of this test is to test the connection between Demand response server and Smart Energy Box and ensure the DR signal is defined as shown in DR Events or as created by DRAS. DR events issued by DRAS should be able to receive by Smart Energy Box. The following are the sequence of steps that should be executed in order to test the connectivity between DRAS and SEB.

Test Pre-condition: Establish Ethernet connection between SEB and CITRIS Apogee System. Smart Energy Box should be able to access Internet. Configure SEB to be able to receive DR events from DRAS.

Step #	Description	Expected Result
1	Run SEB	SEB runs.
2	Run SEB UI	SEB UI runs.
3	Configure building mode to Instant	Building mode successfully set
4	Issue DR Event using LBNL Auto grid DRAS /Siemens DRAS with notification time as current according to the schedule mentioned in DR Events	• Event should be generated and notified to SEB • SEB should receive and display in the SEB UI Events page

11.13.4.2 Establishing connection between Smart Energy Box and Building Management System

This goal of this test is to test the connectivity between Smart Energy Box and Building management system in response to the DR event. Smart Energy will change thermostat setpoint of Room 464 in response to the DR event.

Test Pre-condition: Establish Ethernet connection between SEB and CITRIS Apogee System. Smart Energy Box should be able to access Internet. Configure SEB to be able to receive DR events from DRAS. Configure Moderate and SEB Strategies to change Room 464 Thermostat setpoint.

Step	Description	Expected Result
1	Run Smart Energy Box and Smart Energy Box UI	Both runs
2	Create and save strategies of type “DR Period” one for moderate and another for High as per data provided in DR Period (Instant mode) Strategies	Strategies should successfully created and saved
3	Configure Instant mode strategies for Moderate and High and save	Strategies should be configured and saved
4	Issue DR Event from Demand Response Server	DR events should be visible in SEB UI
5	When Start time of DR event approached, the corresponding Strategy of the current DR mode should be selected and start applying on the building	• SEB UI should show selected DR strategy and also shows current DR Mode • Building thermostat setpoint should be changed as per the definition of strategy
6	Once current DR event is finished, all set points that were changed set to their normal value	SEB relinquishes all set points
7	Repeat Step 5 To Step 6 until all issued events are finished	

11.13.4.3 Simulation based HVAC control optimization for DR

The DR signal has been received, at 9:00 P.M. of the day before the DR Day. The weather forecast for the DR day is pulled from internet. Simulation based HVAC control optimization starts immediately. All candidate strategies in the optimal strategy pool will be evaluated one by one; the best strategy that saves more peak energy is selected for applying on the building. SEB starts applying strategy on the day of DR event starts from 12 AM.

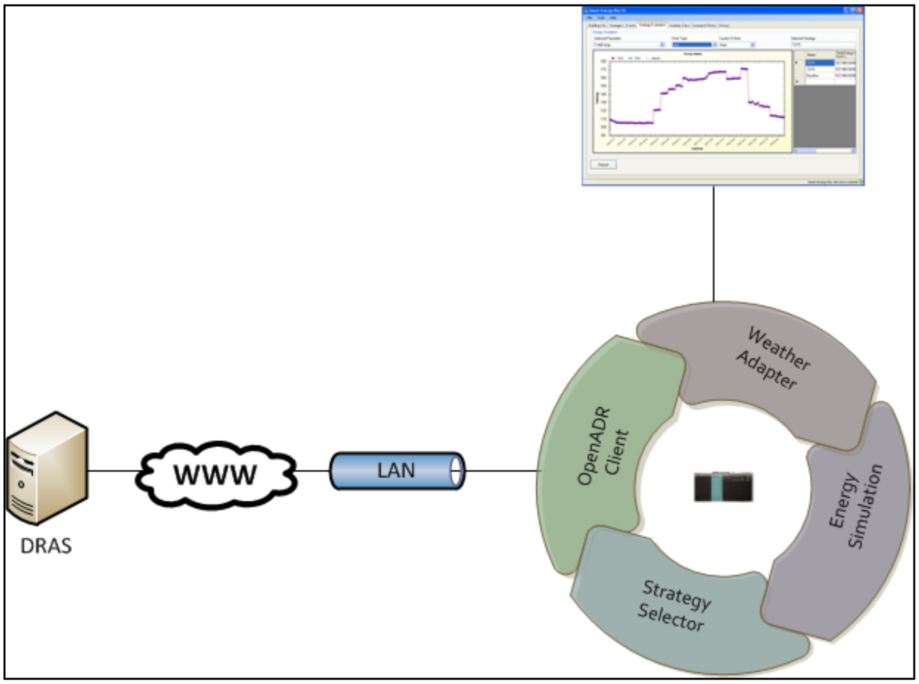

Figure 100: Simulation based DR strategy optimization for HVAC control

Test pre-condition:

- Make sure DRAS, SEB and Building management system are connected

Step #	Description	Expected Result
1.	Start SEB Runtime and SEB user Interface	Both runs successfully
2.	Configure all 13 strategies (Import from file)	All strategies should be saved successfully in to the SEB repository
3.	Configure Strategies applicable for each type of DR mode (High and Moderate)	Configuration should be successfully saved in to repository
4.	Set building mode to PeakDayPricing and save	Building mode should be changed and saved successfully
5.	Issue DR Event for tomorrow as per schedule defined at Peak day pricing event	• Event should be notified to SEB and visible in SEB UI • After 30 minutes (roughly) shows the strategy evaluation results and selected strategy in the SEB UI

11.13.4.4 Execution of the optimal DR strategy for HVAC control

The execution of the optimal DR strategy for HVAC control starts from 0:00 A.M. of the DR day and ends at 11:59 P.M. The change of the setpoint values is applied according to the optimal DR strategy.

Test pre-condition:

- Execute test case 4.3.1

Step #	Description	Expected Result
1.	Selected strategy from test case 4.3.1 will be applied by SEB automatically from 12 AM of the DR day	Set points according to strategy are applied on the building
2.	After 11:59PM of DR day, the strategy should be finished and reset changed building set points to their original values	• Set points should be relinquished to their default values

11.13.4.5 Central Lighting Control

During DR event, if there is still load shedding target to be achieved then Smart Energy Box Central manager will plan load shedding with instant load control Agents such as Central lighting and distributed load control

Step #	Description	Expected Result
1	Configure all Lighting Agents and Strategies before issuing DR event	All Participants should be created and Strategies should be created with Smart Energy Box
2	Central manager negotiates with Lighting Participants every 15 minutes until end of DR event	• During DR event, if load target i.e. 30% is not achieved then Central manager should negotiate with Lighting participants and issue command • Lighting Agents should receive goal • Selected Strategy to be scheduled for next 15 minutes
3	End of DR event	• Set points should be relinquished to their default values

11.13.4.6 Distributed load control

Smart Energy Box starts negotiating with distributed load control 15 minutes before (or as configured) DR period starting time. In each 15-minute negotiation cycle, SEB plans for next 15 minutes only if there was a load reduction goal to be achieved, otherwise it simply skips that cycle. SEB finishes negotiation and commanding goals to gateways within defined interval (15 minutes) and then gateways starts applying control in next cycle.

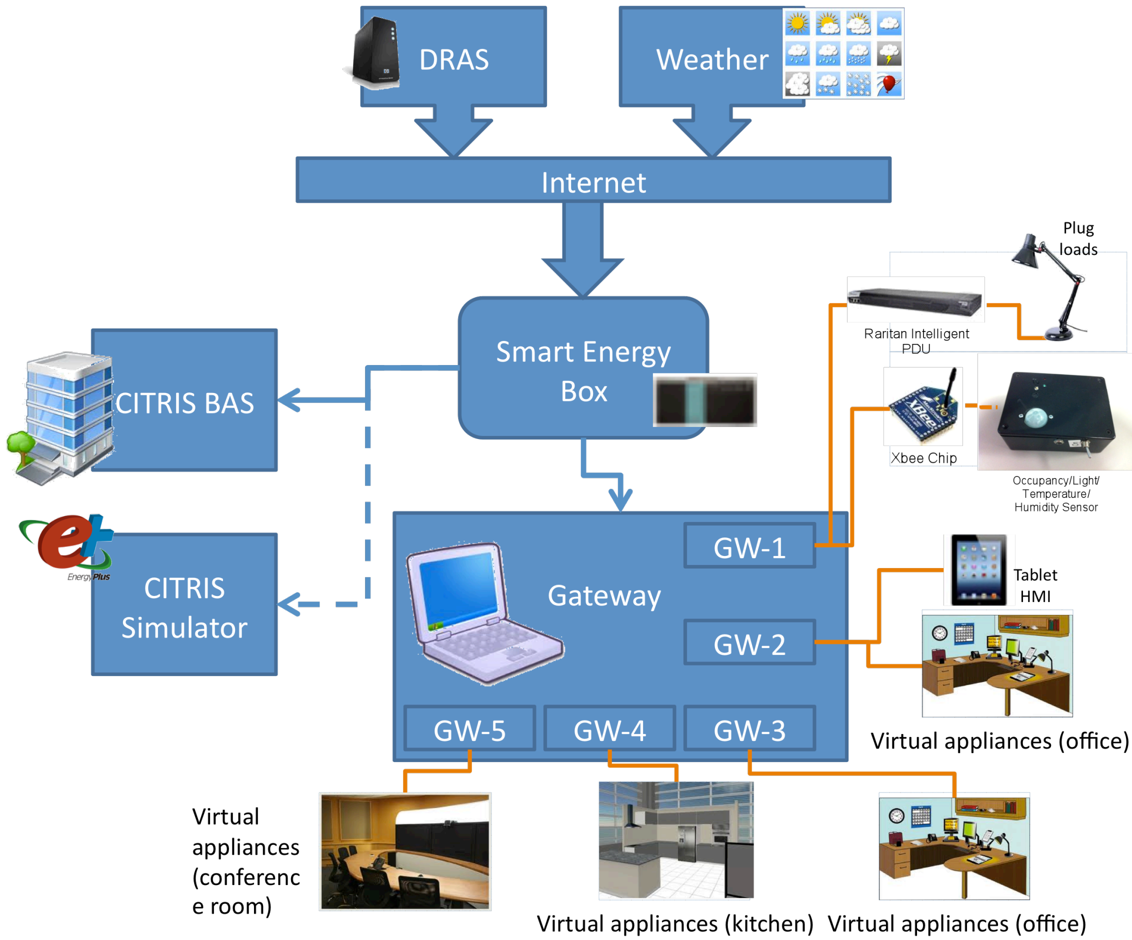

Figure 101: System test layout

11.13.5 Test data

Gateway 1 – Office 1

Device	Max Power	Port	Type	Note
Raritan	-	192.168.1		
Sensor Box	-			With Light and Occupancy Sensors
XBee + usb-serial Cable	-	COM3		
Task Light	60W	R-1	On/Off	Real Appliance
Laptop	60W	R-2	AC/Battery	Real Appliance
Printer	400W	2011	On/Off	Smart Appliance
Total	520W			

Gateway 2 – Office 2

Device	Max Power	Port	Type	Note

HMI Tablet	-	-		
Router	-	-		Router has to be set as DHCP Server for Ethernet enabled Tablets
Task Light	20W	R-4	On/Off	Smart Appliance
Desktop with PSU	200W	2021	On/Off	
Monitor	80W	2022	On/Off	
Humidifier	200W	2023	On/Off	
Printer	500W	2024	Printing/StandBy	
Total	1000W			

Gateway 3 – Office 3

Device	Max Power	Port	Type	Note
Task Light	50W	2031	On/Off	Smart Appliance
Laptop	50W	2032	AC/Battery	
Humidifier	200W	2033	On/Off	
Projector	500W	2034	On/Off	
Small Refrigerator	300W	2035	Dimmable	Smart Appliance, 3 linear levels between 0-300W
Total	1100W			

Gateway 4 – Kitchen

Device	Max Power	Port	Type	Note
Refrigerator I	600W	2041	Dimmable	Smart Appliance, 5 linear levels between 0-600W
Refrigerator II	800W	2042	Dimmable	Smart Appliance, 5 linear levels between 0-800W
Coffee Machine(S)	1200W	2043	On/Off	Smart Appliance
Snack Vending Machine	80W	2044	On/Off	
Beverage Vending Machine	400W	2045	On/Off	
Total	4000W			

Gateway 5 – Conference Room

Device	Max Power	Port	Type	Note
Projector	300W	2051	On/Off	Smart Appliance
Printer	400W	2052	On/Off	
Laptop - 1	50W	2053	AC/Battery	
Laptop - 2	50W	2054	AC/Battery	
Laptop – 3	50W	2055	AC/Battery	

Laptop – 4	50W	2056	AC/Battery	
Total	900W			

<u>Test pre-condition</u>:

- Establish connection between SEB and Gateways
- Establish connections between Gateways and their controlling devices

Step #	Description	Expected Result
1.	Configure test data as mentioned above	Data configuration should be saved and retrieved successfully using gateway web user interface
2.	Run Gateways	• Gateways run successfully • The Gateway Agent GUI runs • The sensors status is shown on GUI • Configured plug load successfully showed and ready to command
3.	Check Gateway SEB connection	• After initialization, the communication between SEB Agent and Gateway Agents is shown on GUI • Gateways are listed in SEB UI
4.	Run test case	Test case should successfully run
5.	DR Negotiation between SEBs and Gateways (if load shedding required)	• SEB UI shows negotiation status • Gateway Agent GUI shows the status of negotiation
6.	Load shedding commanded by SEB	• Gateway applies the control and shows the status in Gateway UI • The changes should reflect in actual meter reading in SEB UI

11.14 Appendix N: Office load calculation, including estimating the centrifugal chiller load

The following graph describes the analysis on the electrical components for the DR events.

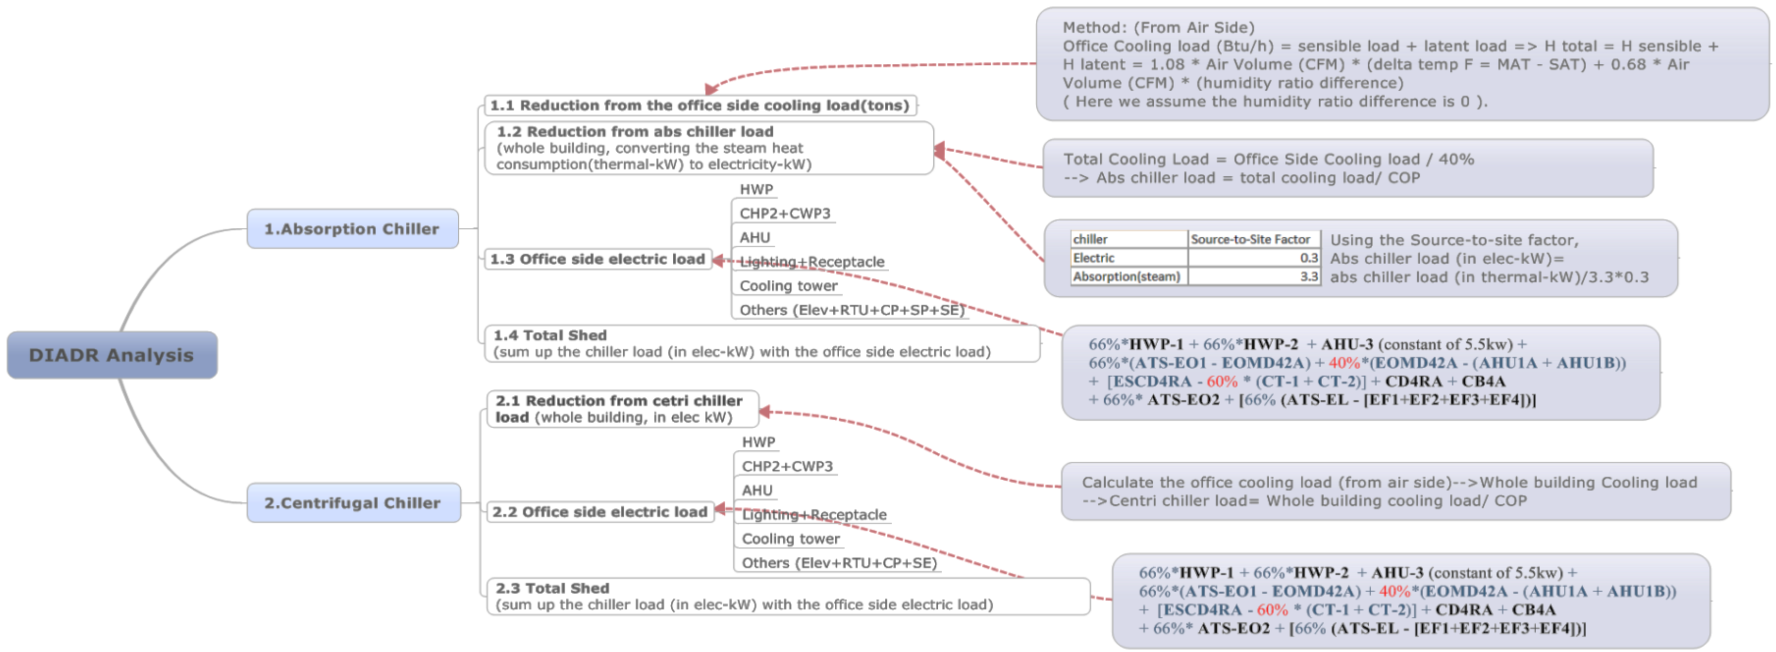

To estimate the electrical consumption had the building been using the centrifugal chiller, the cooling load of the absorption chiller for Oct 17, 2012 was compared to the cooling load for the centrifugal chiller from Oct 17, 2011 (which happened to be about the same temperature). Bin Chen calculated the cooling load using the air side equation (change in temperature for supply air and return air, and air flow), and compared this with the electrical consumption of the centrifugal chiller. He discovered that the chilled water for the office portion of the building is more likely to be 40% rather than 25% previously calculated by Rongxin Yin. He then used the COP of the chiller at 5.96.

11.15 Appendix O: Energy optimization of laptops for demand response (Attached)

11.16 Appendix P: Calculating annual energy consumption and isolation of the office load (Attached)

11.17 Appendix Q: Managing building ventilation using demand-controlled ventilation

The challenge for a modern building operator in selecting ventilation rates and schedules is to achieve energy-efficient operation while ensuring that building denizens receive ample fresh air. Traditionally, these decisions have been made at the commissioning stage, with airflow levels selected to ensure adequate ventilation such that air quality and human bioeffluent levels remain at a comfortable level during full occupancy. This airflow level is called the equilibrium level, and its selection is governed by a variety of factors, including maximum occupancy, usage pattern, air volume, and adherence to state and national building standards.

Looking particularly at commercial buildings, the relevant standard for California nonresidential buildings is the California Code of Regulations (CCR) Title 24, Part 6, Subsection 3, Section 121 (California Buildings Standards Commission, 2010), which draws guidance from a national standard from the American Society for Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) that governs indoor ventilation levels, ASHRAE 62.1 (American Society of Heating, 2010). These standards dictate that mechanically-ventilated spaces must always receive at least 0.15 cubic feet per minute (cfm) of outside air for each square foot of area ventilated. Further, ventilation should be delivered based on occupancy, at a rate of 15 cfm of outside air per occupant of the space. The code defines two possible means by which to detect if the space is occupied: occupant sensors, which generally detect motion and provide only a binary signal, or CO_2 sensors, which reflect the concentration of CO_2 in the indoor air. Use of these sensors to determine ventilation rates is called *demand-controlled ventilation* (DCV).

Though DCV has been around for decades, it is not widely in use. However, emerging technologies, changes in building standards, and more awareness of occupant comfort are increasing deployment of DCV systems. Sensors and the information technology needed to retrieve the data they produce are becoming cost-effective and ubiquitous. Recognizing this and the energy savings potential from reducing excess indoor ventilation, the authors of the Title 24 standard have modified the next iteration of the standard to require that DCV be used in dense settings – specifically, in any room larger than 150 square feet with 40 square feet or less per expected occupant. By itself, this new requirement will drive widespread deployment of DCV systems; coupled with recent studies showing reduced decision-making performance in settings with elevated CO_2 concentration, DCV systems are on the way to becoming standard in most commercial buildings.

In this work, we evaluate the use of DCV in our commercial building testbed. In particular, we use CO_2 sensors as opposed to binary occupancy sensors, as we believe these sensors provide a less-discretized measure of occupancy that allows for increased energy efficiency and improved occupant comfort. We compare our DCV deployment to two other ventilation control strategies. The first is the *baseline* system configuration that reflects the state of the ventilation controls after the commissioning process. The second is an *extreme efficiency* approach that aims to reduce ventilation by as much as safely possible.

Selecting a target CO_2 level

Title 24 of the CCR dictates that DCV systems employing CO_2 sensors should be configured to maintain a CO_2 concentration below 600 parts per million (ppm) beyond outside air concentration (in the absence of a sensor measurement, this is assumed to be 400 ppm). Figure 102, adapted from the ASHRAE 62.1 standard, encapsulates the challenge for building managers: despite dynamic conditions created by the movement of people and a changing environment, maintain a CO_2 concentration in all spaces between

900 and 1175 ppm. Newer research indicates that even this target region may be too high; human decision making performance can show significant reductions at even 1000 ppm CO_2, as compared to 600 ppm (U. Satish, 2012). Given this, for the duration of our study, we chose a target maximum concentration of 800 ppm CO_2.

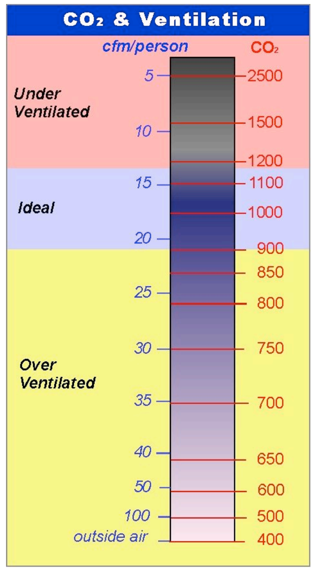

Figure 102: CO_2 concentration, associated ventilation rates, and their acceptability. Adapted from reference material (Airtest Technologies) (American Society of Heating, 2010).

Selecting rooms to monitor and control

To establish a preliminary understanding of building ventilation and air quality, we deployed 8 CO_2 sensors throughout a single floor of office space in our testbed building on campus, as shown in Figure 103. The floor is primarily an open office "cubicle" area, which is surrounded by enclosed offices and conference rooms, and is roughly 10,000 ft^2 total. To capture the variations in behavior among these spaces, six of the CO_2 sensors were deployed in the open office space, one in an office, and one in a conference room. The sensors were all positioned near air return vents and at a height of between three and six feet, according to the guidelines laid out in previous work (W.J. Fisk, 2010).

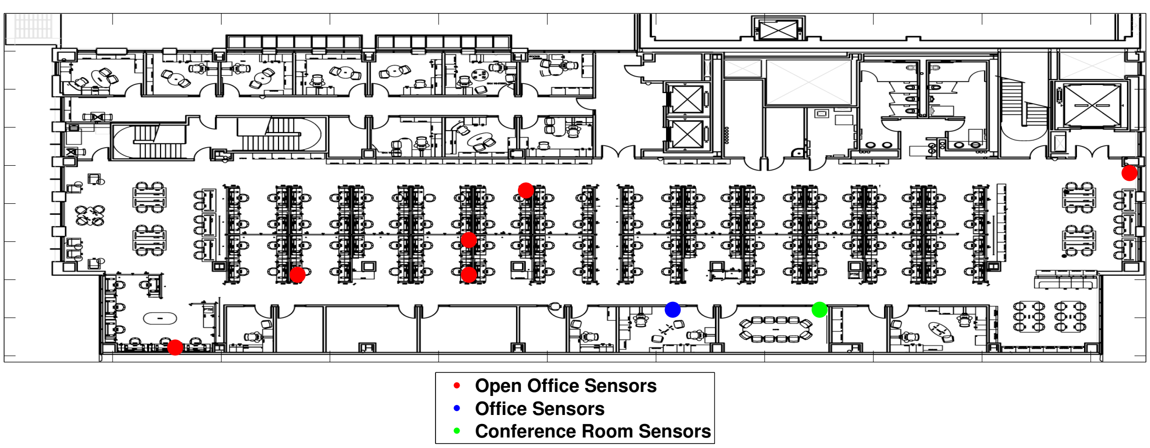

Figure 103: Map of a floor of the testbed building showing CO_2 sensor deployment locations.

Figure 104 shows the CO_2 levels collected across this floor for three typical weekdays. In each area, the CO_2 concentration remains near the concentration of outside air (generally, 400 ppm) during the nighttime hours. As occupants arrive late in the morning, the concentration in all of the areas begins to increase. During the occupied hours, the conference room has intermittent spikes due to meetings, while the other spaces reflect some but minimal occupancy extending into the evening, in line with typical graduate student schedules.

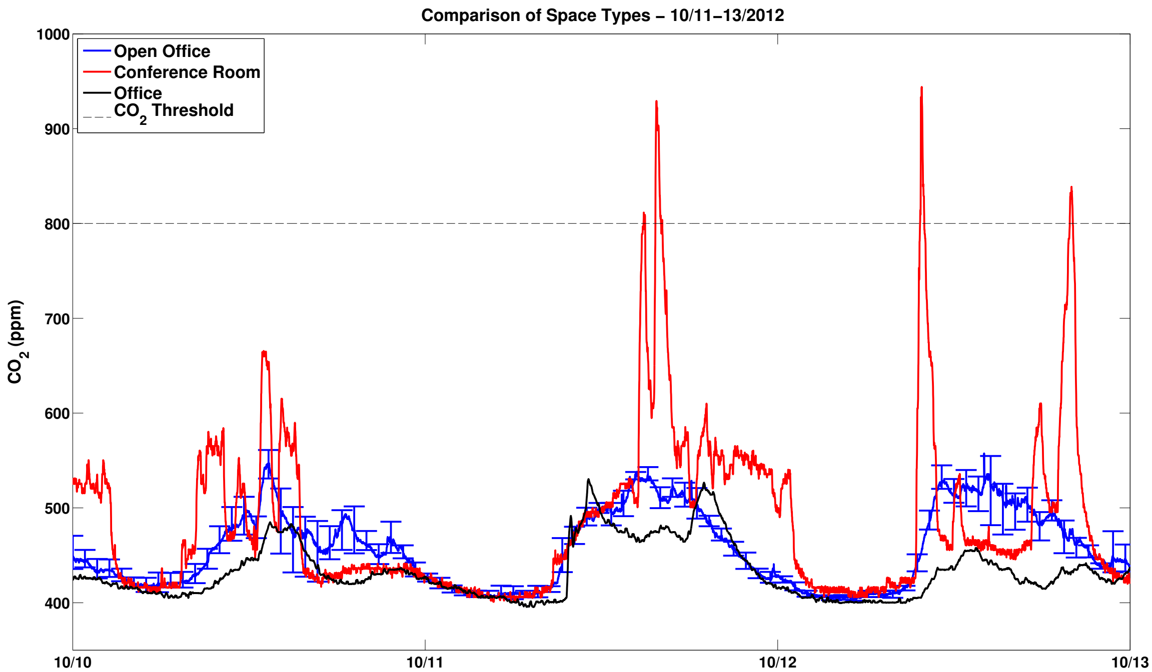

Figure 104: Comparison of CO_2 concentrations for three different types of spaces across our test floor. Error bars on the open office line indicate 10th and 90th percentile measurements.

The conference room approaches our target level of 800 ppm, crossing it four times during the three days. Neither the enclosed office, which seldom has more than 1 or 2 occupants, nor the open office area, which is shared among tens of occupants but is far less dense, approaches the target threshold. In fact, over the month-long duration of our deployment in this area, the open office area never violates the 800 ppm maximum concentration. This is primarily because this area is not densely occupied, has far more baseline ventilation, and comprises a much larger volume of air, all resulting in lower, generally acceptable CO_2 concentration.

The concentration in the enclosed areas, the office and the conference room, generally remains below the threshold, but rises quickly during meetings with multiple people. The breath of the occupants permeates the space, generated at roughly 0.5 L/min and with a concentration of near 5% CO_2 (5000 ppm) for each occupant. As this air mixes with the existing room air, without any increased response from the ventilation system, the air quality in the room quickly degrades. Though the building management system in the room uses a control mechanism to maintain temperature in the room, the system will not introduce increased ventilation until the room exits the configuration *deadband*, a range between the heating and cooling setpoints in which no control actions are taken. For these areas, the heating setpoint is 70 degrees Fahrenheit and cooling setpoint is 74 degrees Fahrenheit. This building, like many others due to conservative minimum ventilation values, is generally overventilated; since supply air is colder than the setpoint, the room temperature tends to stay near or even below the setpoint. Even a large meeting in a small enclosed space can take tens of minutes to generate enough heat in the room to exit the deadband and cause the control system to increase ventilation for the purpose of cooling the room; generally, by this time, the CO_2 concentration in the room has significantly violated the 800 ppm threshold.

These observations lead us to focus our study on enclosed spaces with highly variable occupancy, where air quality is not consistently acceptable and intermittent usage patterns can be exploited for significant energy savings. In our campus building, this description primarily characterizes conference rooms, classrooms, and large presentation rooms.

Baseline building ventilation

With a target room type and CO_2 level selected, we proceed to examine system performance in greater detail and identify metrics to compare in different control scenarios.

Figure 105 shows the interaction of the ventilation system of the conference room, which is a variable-air-volume (VAV) unit configured with minimum and maximum airflow settings, and the CO_2 concentration in the room over a week. The room has an estimated maximum occupancy of 12 people and is an area of 214 ft^2. As explained previously, a PID controller with a deadband maintains the temperature in the room, but seldom needs to cool the room, resulting in air volume that nearly always equals the minimum airflow setting. The average airflow over the week is 222.2 cfm; at 15 cfm/occupant, this estimates that roughly 15 people are continuously in the room. Despite a configuration that aims to ventilate the room as if it is continuously fully occupied, the CO_2 concentration in the room still crosses the acceptable threshold multiple times within this week, for a total of just over 6 hours spent above the threshold (3.6% of the hours). Further, for the grand majority of the time, the room is unoccupied or lightly occupied; it is estimated that the room is occupied less than 10% of the hours of the week. This presents an opportunity not only for increased comfort during occupied hours but also for energy savings during unoccupied hours by employing DCV.

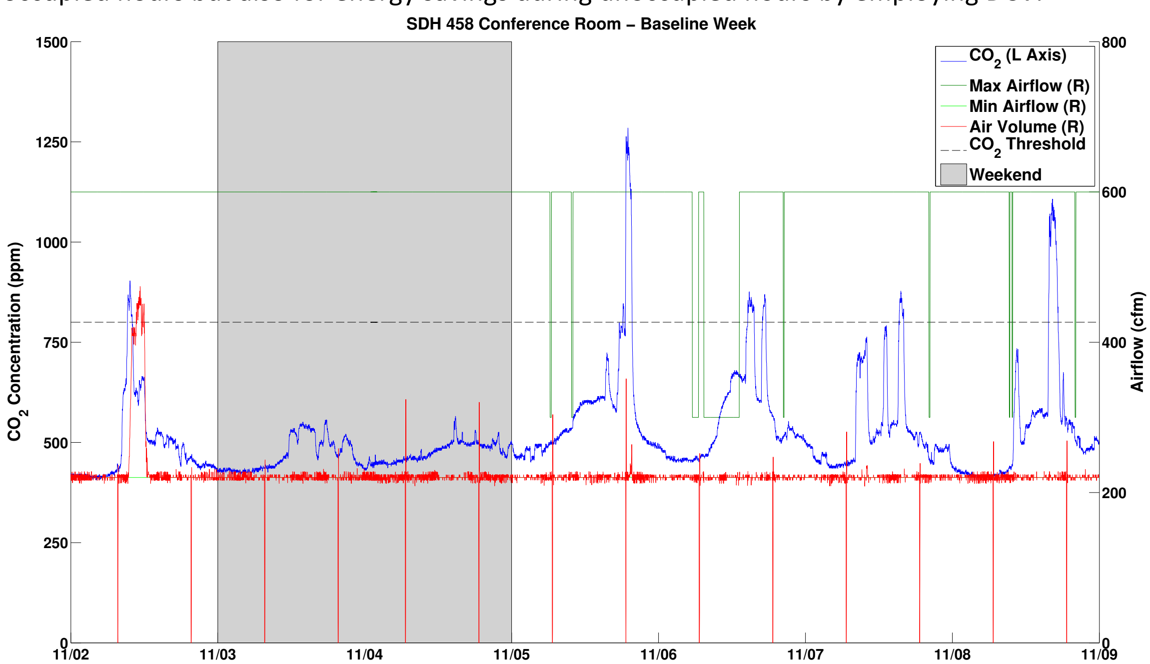

Figure 105: The performance of the baseline ventilation strategy in a conference room over a week. Twice each day, in the morning and in the evening, these VAV units run a self-calibration process that resets their air volume to zero cfm.

Extreme efficiency ventilation

Prior to the deployment of a DCV system, initial ventilation experiments aimed to reduce the ventilation as much as possible while still maintaining sufficient air quality. This effort, which combines an occupancy model, outside air damper control sequence, and significant reductions in default airflow levels, represents an *extreme efficiency* ventilation strategy.

Figure 106 shows a week of operation in the same conference room running an extreme efficiency ventilation system. The average airflow for the room is 79.8 cfm over the week. Though the total airflow to the room, and thus the energy spent on ventilating the room, has significantly decreased (over 64%), we see even more violations of the target maximum CO_2 threshold than we did in the baseline scenario. In fact, there is at least one violation every day, and the total time spent over the maximum threshold is nearly 11 hours (6.5% of the total). Further, the room may exceed the thermal comfort zone range. Each time the airflow deviates from a value below 100 cfm represents the temperature controller responding to a violation of the deadband, meaning that the temperature in the room is 74 degrees Fahrenheit or above; this happens multiple times throughout the week.

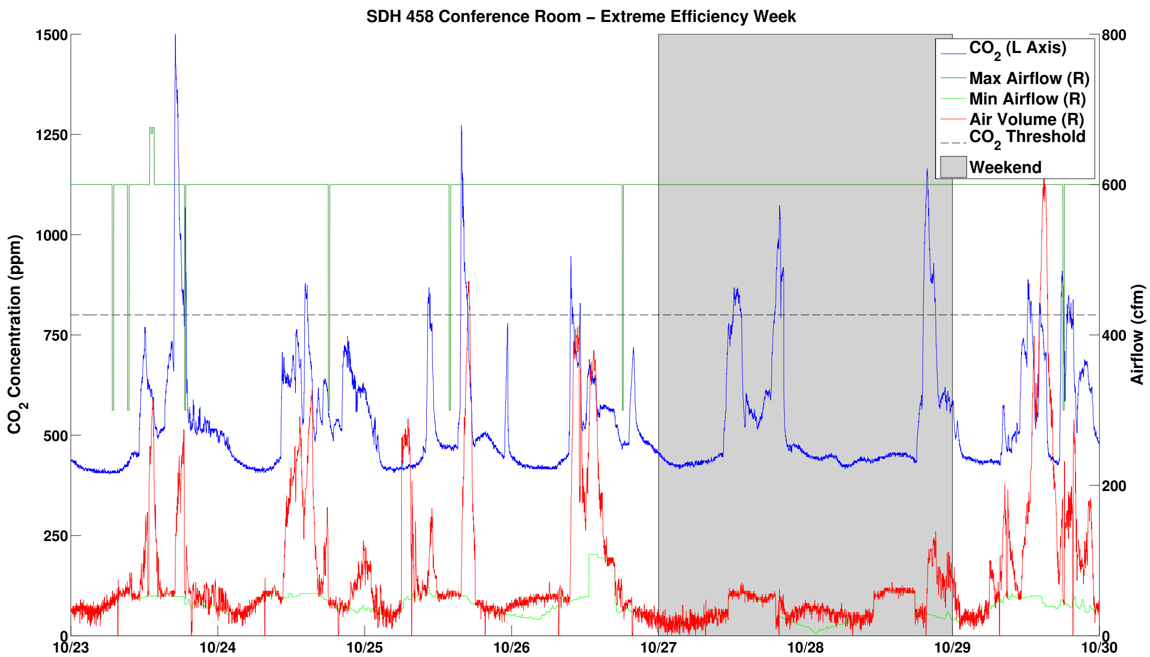

Figure 106: Operation of a conference room ventilation system with an extreme efficiency control strategy.

Demand-controlled ventilation

Since the extreme efficiency ventilation system did not provide adequate airflow to maintain CO_2 at a healthy level, we determined that it was necessary to deploy a ventilation system that could modulate based on occupancy in the room. Initially, we sought to investigate whether a purely software measure of occupancy, leveraging the department room reservation system, would suffice for ventilating the room. In this system, occupants use bConnected, a service in the campus suite that provides a Google Calendar for each of the conference rooms, as seen in Figure 107. By fetching calendar entries, it was possible to discern "scheduled meetings" – when people were expected to be in the room. However, it is not necessarily the case that the room would be occupied when a meeting is scheduled; even more importantly, it is possible that the room would be occupied outside of when a meeting is scheduled – "unscheduled meetings". Looking over a two-week period, we sought to uncover the frequency of unscheduled meetings. During this period, the ventilation system used the extreme efficiency method, and a meeting was determined to be taking place if the threshold CO_2 concentration was surpassed. The results of this investigation showed that over the two-week period, there were 28 total meetings corresponding to violations of the CO_2 threshold. Of these, 15 were during scheduled meetings (out of 26 total scheduled meetings in the period), and 13 were during unscheduled meetings. Half of scheduled meetings did not have enough occupancy to create high CO_2 concentrations even with reduced airflow and a third of all high CO_2 concentration events were during unscheduled meetings. Given this, we believe that these calendar entries can provide useful information about when people are likely going to be in the room, but are not entirely sufficient for providing software-only ventilation control that meets our CO_2 concentration goals.

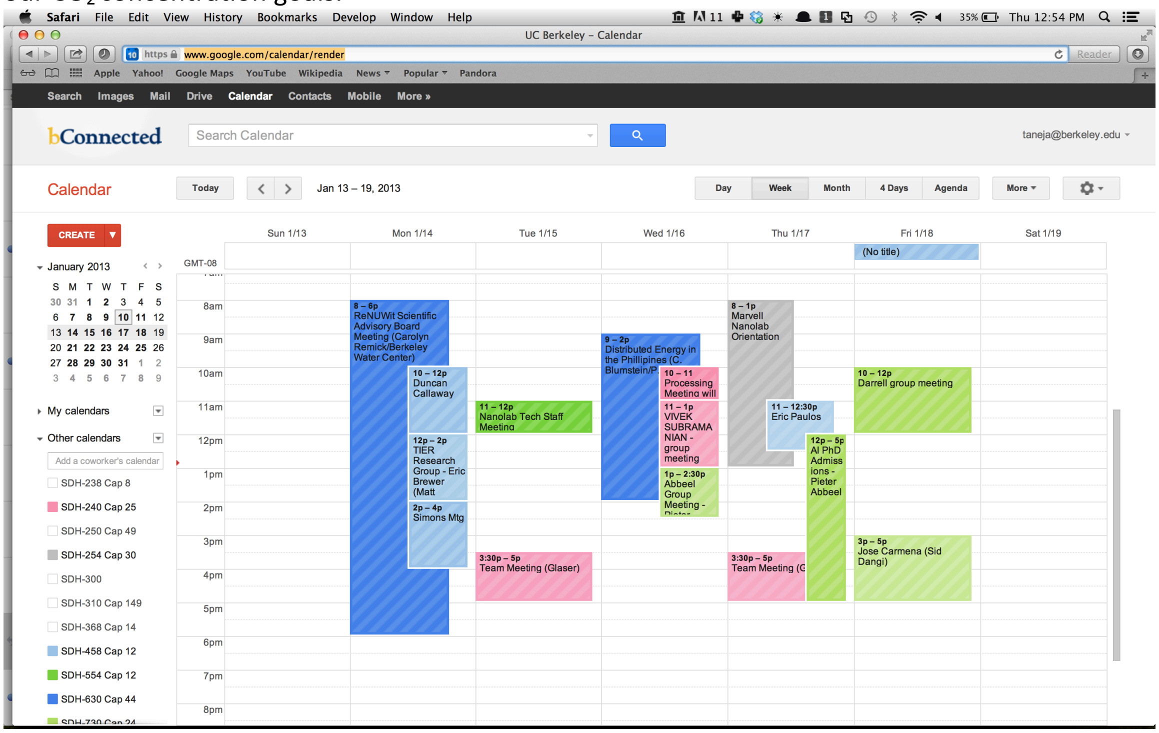

Figure 107: bConnected, UC Berkeley's campus calendar system, showing reservations for the conference rooms over a week.

Given this finding, we designed a demand-controlled ventilation system that utilizes a CO_2 sensor. Our system uses a moving average of CO_2 readings over the previous two minutes; since the CO_2 sensors take a reading every 15 seconds, this averages over enough samples to dampen the effect of outliers, but allows the controller to be agile to somewhat fast changes in the concentration, such as at the beginning of a large meeting. To describe the controller rules, Table 3 introduces some terminology, as well as relevant values for the conference room under study (SDH 458).

Parameter	**SDH 458**
af_{min} : Min airflow (Title 24): 0.15 * ft^2	32.1 cfm
af_{dmax} : Default maximum airflow	600 cfm
af_{dmin} : Default minimum airflow	220 cfm
$CO2_b$: Baseline CO_2 value	425 ppm
$CO2_{max}$: Maximum allowable CO_2	750 ppm
$CO2_{hyst}$: Threshold to reduce airflow	700 ppm

Table 3: Terminology used for DCV controller. Values are provided for the conference room used in this study.

Table 4 shows the rules used by the controller to set the minimum and maximum airflow values under all conditions. Using the calendar entries, the controller ventilates during scheduled meetings by selecting a minimum airflow value that reflects the concentration of CO_2 in the room between a baseline value for the room ($CO2_b$) and a safe maximum ($CO2_{max}$), which is slightly below the target maximum to ensure that the system can respond quickly enough to prevent violations of the target maximum concentration. Additionally, the controller ventilates slightly before and after the scheduled meeting to ensure fresh air for occupants when they arrive and in case the meeting runs over its allotted time. In non-meeting times, as long as the CO_2 concentration is not approaching the target maximum, the minimum airflow remains at the absolute minimum af_{min}, which is based only on square footage and assumes zero occupancy. During these times, the maximum is increased to reflect occupancy, though the wide deadband in the temperature control system generally dictates that air volume will match the minimum airflow setting. In unscheduled meeting times when the CO_2 concentration does approach the target maximum, the system responds by providing full airflow af_{dmin} in order to prevent a violation. In order to reduce cycling between af_{min} and af_{dmin} around the threshold, a measure of hysteresis is added by not reducing the airflow until the CO_2 concentration falls below a level less than $CO2_{max}$, called $CO2_{hyst}$; in our system, this hysteresis level is 700 ppm.

Situation	**Minimum Setting**	**Maximum Setting**
5 mins before/after scheduled meeting	$[af_{min}, af_{dmin}]$ α CO_2 value in $[CO2_b, CO2_{max}]$	af_{dmax}
Other times ($CO_2 < CO2_{max}$)	af_{min}	$[af_{min}, af_{dmin}]$ α CO_2 value in $[CO2_b, CO2_{max}]$
Other times ($CO_2 \geq CO2_{max}$)	af_{dmin} until CO_2 falls below $CO2_{hyst}$	af_{dmax}

Table 4: Rules used for DCV controller.

A week of performance data for the DCV controller is provided in Figure 108. For the grand majority of hours, the airflow closely mimics af_{min} as meetings are not being held and CO_2 concentrations do not approach the target maximum level. During scheduled meetings with significant occupancy, the reactivity of the system maintains the CO_2 concentration at a safe level. During unscheduled meetings, the system responds in time to maintain CO_2 concentration near the target maximum. It may be necessary to lessen the hysteresis threshold to reduce the cycling of the VAV damper.

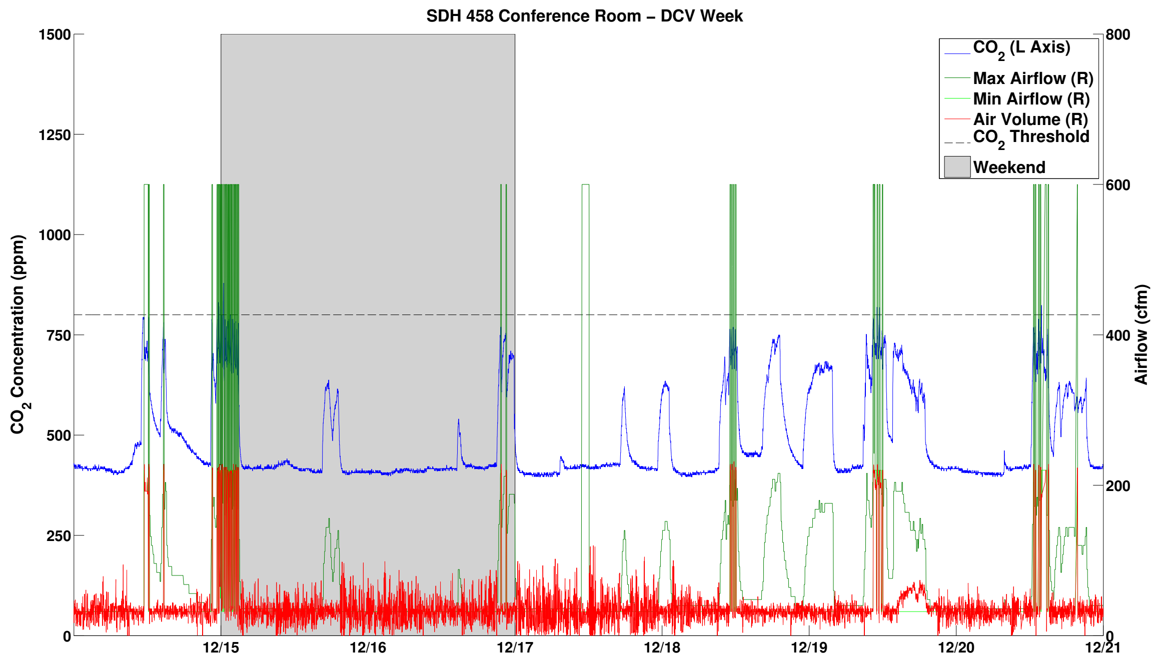

Figure 108: Operation of a conference room ventilation system with a demand-controlled ventilation strategy.

Comparison of results

A table of results comparing the three ventilation strategies is provided in Table 5. Mean ventilation power can be calculated by using a model of supply fan power derived from measurement data, as shown in Figure 109. At each airflow level, we calculate the power required to provide an incremental cfm of airflow; we use this ratio and the total airflow required by that room to calculate its instantaneous ventilation power.

Ventilation Strategy	Scheduled Meetings (Over)	Unscheduled Meetings	Mean Airflow	Mean Ventilation Power	Time > 800 ppm
Baseline	11 (6)	3	222.2 cfm	0.1765 kW	6h3m (3.6%)
Extreme Efficiency	12 (7)	9	79.8 cfm	0.0616 kW	10h57m (6.5%)
DCV	3 (0)	8	40.2 cfm	0.0272 kW	17m (0.2%)

Table 5: Results of a week of operation of three different ventilation strategies. Counts of scheduled meetings are obtained from the department calendar; a measure of events over the 800 ppm threshold is also provided.

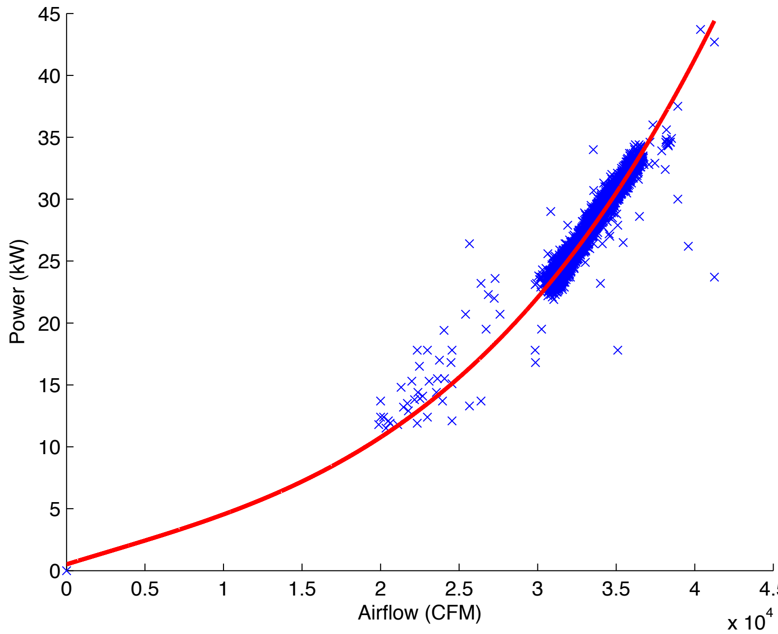

Figure 109: A third-order polynomial model accurately relates supply fan airflow to supply fan power, coinciding with fan affinity laws (American Society of Heating, Refrigerating, and Air Conditioning Engineers, 2012).

Despite the different levels of activities during the three weeks under observation, the scale of the performance differences is significant. By employing DCV, CO_2 concentration is violated a factor of 21 and 38 less time in comparison to the baseline and extreme efficiency systems, respectively, while only using 15% and 44% of the power of those systems. Further, small changes to the DCV rules such as reducing the maximum allowable CO_2 value or adding a derivative term could further improve violation performance with minimal effect on power consumption. In the absence of far-improved localization systems that can provide instantaneous and accurate occupancy estimates, we believe the strength of these results highlights the importance of incorporating CO_2 sensors into ventilation systems in dense settings with variable occupancy such as conference rooms.

Future directions

As we deploy this system throughout our testbed building, we see similar performance in other rooms. As of this writing, there are CO_2 sensors deployed in 7 of the 10 conference and class room settings in the building, with our DCV system running on 6 of them, saving roughly 1.3 kW continuously out of approximately 3.9 kW used for the ventilation systems in these rooms.

Another potential application of a similar idea is *demand-controlled filtration* (DCF) (D. Faulkner, 2007). Our testbed building also has over 15000 ft^2 of Class 100 and Class 1000 cleanroom. In these types of settings, maintaining low particle counts of impurities is critical. As such, besides using VAV systems for injecting fresh air into the space, recirculating air handler units (RAHUs) are used to continuously push air through particle filters. In our building, the total airflow through these RAHUs is around 215000 cfm; this dwarfs the airflow in our 10 conference and class rooms, which is a total of about 5000 cfm. The potential to curtail RAHU operation when the cleanrooms are not in use could save large amounts of power. However, since air volume and airflow in these rooms are large, CO_2 sensors do not provide enough indication of occupancy in these spaces. Instead, basic motion sensors are better, in this instance. Figure 110 shows a few days of operation of a combination CO_2 and binary motion sensor we installed in one bay of the cleanroom. There are substantial periods of no motion in this bay; energy can be saved by turning down the relevant RAHUs during inactive periods.

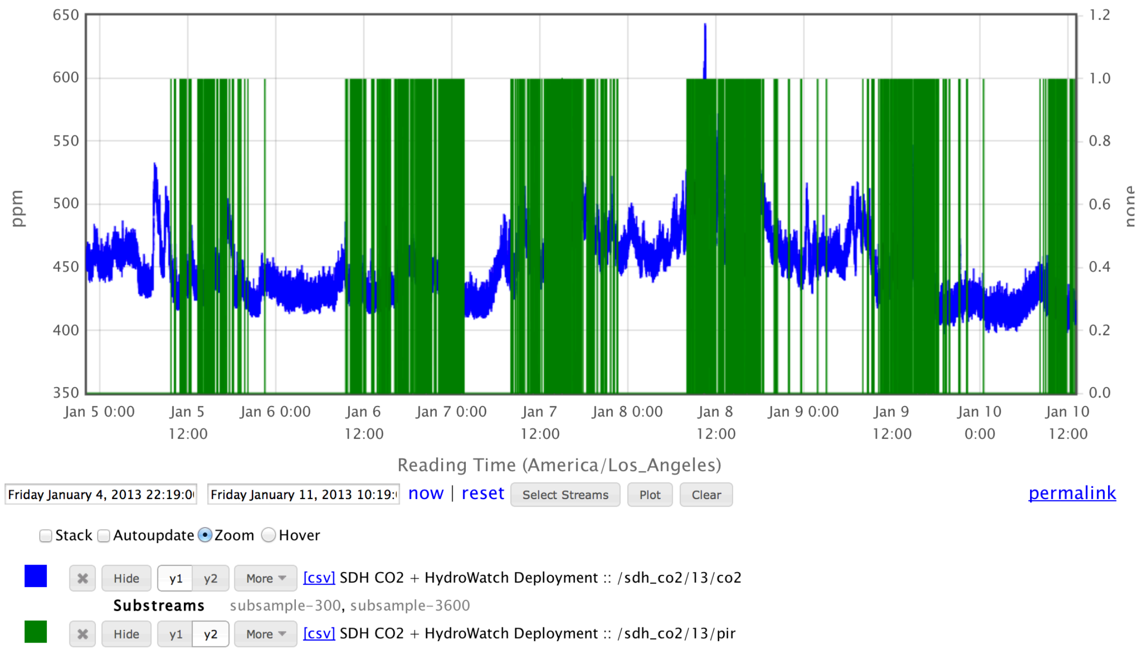

Figure 110: CO_2 (blue) and motion sensor (green) values in one bay of a cleanroom. Identifying periods of inactivity will allow energy savings by turning down the rates of recirculating air handler units (RAHUs).

Another potential application for DCV and DCF systems is as a supply-following load (J. Taneja, 2012). In this scenario, the rate of ventilation would be modulated to make the energy consumption of the supply fan better match the availability of electricity from the grid. This becomes more valuable as non-dispatchable renewable sources such as solar and wind comprise a larger proportion of generation on the electricity grid. It is important to note, though, that the "slack," or capacity to change, in the load is limited in one direction; that is, in nearly all situations, the system is running as efficiently as possible, and energy consumption cannot be reduced any further. However, at the same time, these systems can nearly always increase consumption to better match a surplus of grid electricity. This potential to sink extra electricity could be used in combination with other loads with different characteristics to provide supply-following capacity.

Conclusions

This study examines the appropriate deployment setting for a demand-controlled ventilation system and evaluates its performance in such a setting. In comparison with a baseline system reflecting the results of commissioning and an extreme efficiency system that attempts to reduce airflow as much as safely possible, the system is able to both reduce the time spent in violation of an 800 ppm target CO_2 constraint as well as sharply reduce the energy needed to operate the ventilation system.

Bibliography

Airtest Technologies. (n.d.). CO2 Ventilation Control and California Title 24.

American Society of Heating, R. a. (2010). *Ventilation for Acceptable Indoor Air Quality - Standard 62.1.*

American Society of Heating, Refrigerating, and Air Conditioning Engineers. (2012). *ASHRAE Handbook - HVAC Systems and Equipment.*

California Buildings Standards Commission (2010). California Code of Regulations (CCR), Title 24, Part 6, Section 121, available at http://www.bsc.ca.gov/codes.aspx

D. Faulkner, D. D. (2007). *Demand Controlled Filtration in an Industrial Cleanroom.*

J. Taneja, R. K. (2012). Defining CPS Challenges in a Sustainable Electricity Grid . *International Conference on Cyber-Physical Systems.*

U. Satish, M. M. (2012, September). Is CO2 an Indoor Pollutant? Direct Effects of Low-to-Moderate CO2 Concentrations on Human Decision-Making Performance. *Environmental Health Perspectives* .

W.J. Fisk, D. S. (2010). *CO2 Monitoring for Demand Controlled Ventilation in Commercial Buildings.*