



Pacific Northwest
NATIONAL LABORATORY

*Proudly Operated by **Battelle** Since 1965*

Transactive Control of Commercial Building HVAC Systems

December 2016

CD Corbin
A Makhmalbaf
S Huang
VV Mendon
M Zhao

S Somasundaram
G Liu
H Ngo
S Katipamula

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 Battelle Memorial Institute, 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, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <<http://www.ntis.gov/about/form.aspx>>
Online ordering: <http://www.ntis.gov>



This document was printed on recycled paper.

(8/2010)

Transactive Control of Commercial Building HVAC Systems

CD Corbin	S Somasundaram
A Makhmalbaf	G Liu
S Huang	H Ngo
VV Mendon	S Katipamula
M Zhao	

December 2016

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Abstract

This document details the development and testing of market-based transactive controls for commercial building heating, ventilation, and air-conditioning systems. These controls are intended to serve the purposes of reducing electricity use through conservation, reducing peak building electric demand, and providing demand flexibility to assist with power system operations. This report is the summary of the first year of work conducted under Phase 1 of the Clean Energy and Transactive Campus Project. The methods and techniques described here were first investigated in simulation, and then subsequently deployed to a physical testbed on the Pacific Northwest National Laboratory campus for validation. In this report, we describe the models and control algorithms we have developed, testing of the control algorithms in simulation, and algorithm deployment in a physical testbed. Results from physical experiments support previous simulation findings and provide insights for further improvement.

Summary

This report is a summary of the first year of work conducted under Phase 1 of the Clean Energy and Transactive Campus (CETC) project. In this report, we detail the development, deployment, and testing of a market-based transactive control application for commercial building heating, ventilation, and air-conditioning (HVAC) systems. Specifically, we discuss the application of this control application to a variable-air-volume (VAV) system common to many commercial buildings.

Market-based control, an example of transactive control, is a distributed control strategy. In this document, these terms are used interchangeably. In a market-based control system, a virtual market enables transactions between HVAC components for the exchange of “commodities,” such as electricity power or cooling/heating energy. Each component is represented by an “agent” that is self-interested and tends to maximize its own benefit. Agents submit bids for commodities based on the benefit they receive. The market receives bids from all agents, and determines the clearing price of the commodity. Each agent then adjusts its consumption based on the cleared price. Market-based control is distributed and scalable, making it suitable for large-scale applications.

Markets may be defined within a building according to commodity (e.g., hot water), physical relationship (e.g., all VAV boxes connected to an air handler), or some combination thereof. In this work, we have taken this concept and created a market in a commercial building HVAC system that allows zones to bid for cooling energy with the air handler and chiller, which then bids for electricity from the electric market to generate the necessary amount of cooling. The purpose of this system is to expose the building’s inherent electric demand flexibility, and thus allow integration of building operation with power system operation. The structure of our market is bi-level—both cooled air and electricity are commodities. VAV agents, representing the thermal zones needing cool air for conditioning, purchase the cool air from the AirMarket. This market has a single supplier, the AHUChiller agent, which in turn purchases the electricity it requires to generate the cool air from the ElectricityMarket. The ElectricityMeter supplies electricity to the ElectricityMarket. In addition, we extend the market-based control for VAV terminals to heating operation and describe the market structure for a VAV system that has hot water reheating capability.

The control system is composed of a set of models, each representing separate conditioned areas, equipment, and markets. Models are control-oriented models—all of which are inverse models—and are therefore relatively simple compared to those used in detailed energy simulation. The developed models include 1) a zone model to predict the HVAC energy demand based on outdoor dry-bulb temperature and other zone parameters, 2) an air handler model used to estimate fan power and cooling load given real-time measurements from the building automation system, 3) a simple chiller model that estimates the electric demand of the district chilled water plant required to serve the cooling load calculated by the air handling unit, and 4) a set of rooftop unit models for the future deployment of transactive controls to commercial buildings that have one or more zones conditioned by packaged rooftop heat pumps. For each model, we describe the mathematical formulation and the method by which we tune the models to predict the performance of the physical system.

Simulations of transactive market controls were performed using a new co-simulation capability developed for this project. This new co-simulation capability, built upon the VOLTTRON™ platform, allows testing and validation of the developed algorithms against our EnergyPlus building simulation model within the target deployment environment. A VOLTTRON agent was developed to manage communication between the VOLTTRON message bus and the EnergyPlus simulation engine. We have validated the VOLTTRON-based market application against previous simulation cases using this capability.

The first simulation case used a fixed price to test market operation independent of changes in price from one clearing interval to the next. Depending on the price used, this case resulted in demand reduction and energy savings. In the second demand-limiting case, we imposed a demand limit to the total (fan plus chiller) electric demand at each market-clearing interval. The third dynamic price case tested the ability of the market and building to react to price signals that changed over time. Table S.1 summarizes the impacts of the three incentive signals on three metrics of interest: building peak demand, total energy consumption, and zone temperature deviation. Figure S.1 shows simulated electricity demand from the fixed price case compared to the baseline.

Table S.1. Results of transactive market simulation cases.

Case		Peak Load (%)			Energy (%)	Temperature (°C)
		Mean	Min	Max		
Fixed Price	\$55	5.6	-0.6	11.4	4.7	-0.07
	\$60	1.1	-4.2	6.9	1.0	0.30
	\$65	-4.2	-9.3	1.4	-3.1	0.61
	\$70	-8.5	-12.3	-5.1	-7.1	0.85
Demand Limit		-6.2	-6.3	-6.1	-1.5	0.17
Dynamic Price		5.8	-4.2	14.2	3.1	-0.03

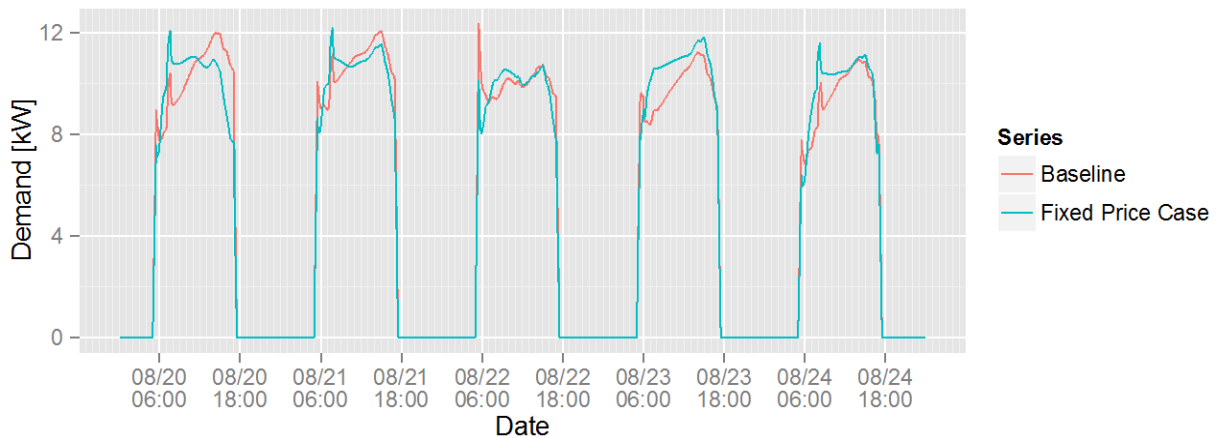


Figure S.1. Comparison of \$65/MW fixed price simulation to baseline.

These market-based controls have been deployed in a small commercial building on the Pacific Northwest National Laboratory (PNNL) campus to test the market-based control on a physical system. Each agent was configured to receive measured real-time values from the building automation system (BAS). These values set the state of the models so that calculations—and thus market bids—reflected current conditions. VAV agents controlled the cooling set points in occupied, standby, and unoccupied modes strictly within limits established with the building manager. Two types of experiments were performed, corresponding to the flat pricing and demand-limiting cases explored in simulation. A number of issues were encountered during testing that invalidated early results.

We evaluated the performance of the market-based controls during the physical tests using two methods: 1) a simple, naïve approach that compares experimental results from performance measured during a similar “baseline” day, and 2) a statistical approach that attempts to model building performance based on a set of relevant predictors selected by their association with the underlying physical processes that determine building energy use. Results from these experiments generally supported previous findings from simulation studies, and provided insights for further improvement. Predicted baseline and

experimental results are shown in Figure S.2. In contrast to the direct comparison method, the regression method shows reduced demand and energy during the June 29 experiment across both fan and chiller.




Figure S.2. Comparison of measured and model-predicted chiller and fan power, June 29.

Experimental results show promise for the application of market-based controls in commercial building HVAC systems. Savings were modest compared to results from previous simulations, but we believe they can be improved. Some differences between experiment and simulation were expected due to a mismatch between the EnergyPlus model and the physical building, while many differences may likely be attributed to imperfect modeling of building dynamics. In simulation, decisions made by VAV agents were acted upon immediately and control was perfect; in a physical experiment, control decisions were delayed, and control was imperfect.

This work forms the foundation for a larger deployment of control methods in multiple buildings on the PNNL campus. The larger deployment is intended to demonstrate the use of transactive controls in commercial buildings for constructing transactive campuses capable of energy-efficient operations that enable greater penetration of renewable resources and a more robust and reliable electric grid. During year two of this CETC project, the transactive market will be deployed in additional buildings on the PNNL campus. Models will be extended to heating operation, and new models will be developed to enable transactions in common HVAC system topologies.

Acknowledgments

We acknowledge Robert Lutes for his many hours spent configuring the platform and troubleshooting integration issues; Ron Underhill for bridging organizations and driving the development of process and documentation for experiments; Jereme Haack, Craig Allwardt, Kyle Monson, and the entire VOLTTRON development team for their assistance with troubleshooting, additional feature development to support the project, and guidance on agent development; Tom Haller, Curt Nichols, Steve Hultman, Sanjay Sanan, Shan Below, and many others in Facilities and Operations, without whose assistance with integration and testing would have made the physical experiments impossible; Yannan Sun for guidance on our baselining methodology; and He Hao for his thorough review and helpful comments.

Acronyms and Abbreviations

°C	degree(s) Celsius
°F	degree(s) Fahrenheit
ACM	Ascent Control Module
AHRI	Air-Conditioning, Heating, Refrigeration Institute
AHU	air handling unit
ANN	artificial neural network
ANSI	American National Standards Institute
BACnet	Building Automation and Control Network
BAS	building automation system
bhp	brake horsepower
CAISO	California Independent System Operator
CETC	Clean Energy and Transactive Campus project
CFM	cubic feet per minute
COP	coefficient of performance
CRBM	Continuous Restricted Boltzmann Machine
ft ²	square foot(feet)
EER	energy efficiency ratio
HP	heat pump
HMM	Hidden Markov Model
HVAC	heating, ventilating, and air-conditioning
IP	Internet Protocol
KKT	Karush-Kuhn-Tucker
kW	kilowatt(s)
LBNL	Lawrence Berkeley National Laboratory
m ³ /s	cubic meter(s) per second
MW	megawatt(s)
NAE	Network Automation Engine
NUC	Next Unit of Computing
PC	Personal Computer
PNNL	Pacific Northwest National Laboratory
R ²	coefficient of determination
RTU	rooftop unit
SVM	support vector machine
SWM	social welfare maximization
UML	Unified Modeling Language
VAV	variable air volume

Contents

Abstract	iii
Summary	v
Acknowledgments.....	ix
Acronyms and Abbreviations	xi
1.0 Introduction	1.1
1.1 Objective	1.1
1.2 Outcomes.....	1.1
1.3 Structure of Document	1.2
2.0 Transactive Market Model.....	2.1
2.1 Background and Significance.....	2.1
2.2 Market-Based Control	2.1
2.3 Market Structure for Commercial Building VAV Systems	2.2
2.3.1 General Bidding and Clearing Process.....	2.3
2.3.2 VAV Bidding and Clearing.....	2.4
2.4 Hot Water Markets for Heating Operation.....	2.5
2.4.1 Market Structure.....	2.5
2.4.2 Problem Statement	2.6
2.4.3 Market-Based Control for VAV Terminal	2.8
2.4.4 Discussion	2.12
3.0 Thermal and HVAC Component Models.....	3.1
3.1 Zone Model	3.1
3.1.1 Model Formulation.....	3.1
3.1.2 Parameter Identification	3.4
3.1.3 Future Work	3.5
3.2 Air Handler Model	3.6
3.2.1 Model Formulation.....	3.6
3.2.2 Parameter Identification	3.6
3.2.3 Model Validation.....	3.6
3.3 Chiller Model	3.7
3.3.1 Model Formulation.....	3.7
3.3.2 Parameter Identification	3.8
3.3.3 Model Limitations	3.9
3.4 Rooftop Unit Models.....	3.9
3.4.1 Introduction.....	3.9
3.4.2 BUILDING6 Rooftop Unit Characteristics.....	3.9
3.4.3 Model Formulation.....	3.10

3.4.4	Parameter Identification Using Catalog Performance	3.14
4.0	Market Simulation and Cross Validation	4.1
4.1	Building Model	4.1
4.2	VOLTTRON Co-Simulation.....	4.2
4.2.1	EnergyPlus Agent.....	4.2
4.2.2	Market and HVAC Agents.....	4.2
4.3	Case Studies	4.3
4.3.1	Fixed Price.....	4.4
4.3.2	Demand Limit	4.5
4.3.3	Dynamic Price	4.6
4.3.4	Summary and Discussion of Results	4.7
5.0	Physical Experiments	5.1
5.1	Experiment Description.....	5.1
5.2	Deployment and Integration.....	5.2
5.2.1	Model Tuning.....	5.2
5.2.2	Agent Configuration.....	5.3
5.2.3	Challenges and Resolution	5.5
5.3	Analysis Method	5.6
5.3.1	Method 1	5.7
5.3.2	Method 2	5.10
5.3.3	Limitations and Assumptions.....	5.11
5.4	Results and Analysis	5.12
5.4.1	Flat Price Experiments	5.12
5.4.2	Demand-Limiting Experiments.....	5.16
6.0	Conclusions and Next Steps	6.1
7.0	References	7.1
	Appendix A – Supplement to Alternative Market Structures	A.1
	Appendix B Supplement to Rooftop Unit Model	B.1

Figures

S.1	Comparison of \$65/MW fixed price simulation to baseline	vi
S.2	Comparison of measured and model-predicted chiller and fan power, June 29	vii
2.1	UML object diagram of the two-commodity transactive market implemented in this project.....	2.3
2.2	UML object diagram of a commercial building market system in which natural gas and electricity are both considered.....	2.6
2.3	Single Maximum control logic for a VAV terminal	2.7
2.4	Dual Maximum control logic for a VAV terminal.....	2.7
2.5	The aggregation of VAV demand curves.....	2.10
2.6	The conversion between power and electricity demand curves	2.11
2.7	The problem caused by inappropriately selected m	2.12
3.1.	Comparison of the actual and predicted zone temperature and cooling rate for a sample thermal zone within the building	3.2
3.2	A residual plot of the predicted cooling rate for the initial model	3.2
3.3	A residual plot of the predicted cooling rate for the revised model	3.3
3.4.	A comparison of the predictive performance of the initial and revised model with the actual measured performance for a sample thermal zone within the building	3.4
3.5.	Comparison of measured and model-predicted fan power during 1 day of the validation data set.....	3.7
3.6.	Scatter plot of total chiller plant cooling load versus plant electric demand.....	3.8
3.7	Fitted line plots showing estimated and actual cooling capacity for Models B and C.....	3.15
3.8	Fitted line plots showing estimated and actual heating capacity for Models A and B.....	3.16
3.9	Fitted line plots showing estimated and actual heating electric power for Models A and B	3.16
4.1	Floor plan of the modeled building	4.2
4.2	Comparison of \$65/MW fixed price simulation to baseline	4.5
4.3	Comparison of demand limit simulation to baseline.....	4.6
4.4	Comparison of dynamic price simulation to baseline	4.7
5.1	Direct comparison of chiller power, fan power, and outdoor air temperature for similar days.....	5.9
5.2	Comparison of measured and model-predicted chiller and fan power, June 29	5.13
5.3	Comparison of measured and model-predicted chiller and fan power, August 18	5.14
5.4	Comparison of measured and model-predicted chiller and fan power, September 1	5.14
5.5	Comparison of measured and model-predicted chiller and fan power, September 8.....	5.15
5.6	Comparison of measured and model-predicted chiller and fan power, September 13.....	5.15
5.7	Comparison of measured and model-predicted chiller and fan power, September 14.....	5.16
5.8	Comparison of measured and model-predicted chiller and fan power, August 25	5.17
5.9	Comparison of measured and model-predicted chiller and fan power, August 31	5.17

5.10	Comparison of measured and model-predicted chiller and fan power, September 2.....	5.18
5.11	Comparison of measured and model-predicted chiller and fan power, September 15.....	5.18
B.1	Fitted line plot showing the estimated vs. actual cooling capacity for Model E.....	B.4
B.2	Fitted line plot showing estimated vs. actual cooling electric power for Model E.....	B.5
B.3	Fitted line plots showing estimated vs. actual heating capacity for Models E and F.....	B.5
B.4	Fitted line plots showing estimated vs. actual heating electric power for Models E and F.....	B.6

Tables

S.1	Results of transactive market simulation cases	vi
3.1	Parameters for a sample thermal zone within the building	3.5
3.2	BUILDING6 heating and cooling systems	3.10
3.3	BUILDING6 indoor and outdoor fan specifications	3.10
3.4	Data and modeling methods classification	3.10
3.5	Data points required to characterize a rooftop heat pump using the inverse method	3.11
3.6	Fan motor heat calculation for BUILDING6 RTUs	3.11
3.7	Fan motor power calculated at different external static pressures	3.12
3.8	BUILDING6 cooling capacity coefficients and R^2 s	3.15
3.9	BUILDING6 heating capacity coefficients and associated R^2 s	3.15
3.10	BUILDING6 heating electric power consumption coefficients and associated R^2 s	3.16
4.1	Summary of HVAC equipment characteristics and operating parameters	4.1
4.2	Simulation parameters used in transactive market simulations	4.4
4.3	Minimum and maximum zone airflow configuration for VAV agents in simulation experiments	4.4
4.4	Results of transactive market simulation cases	4.8
5.1	Summary of experiments performed on the physical testbed	5.1
5.2	Chiller model parameter for physical experiments	5.2
5.3	Fan model parameters for physical experiments	5.3
5.4	Example of zone model parameters for a single zone used in physical experiments	5.3
5.5	Input points configured for the AHUChiller agent	5.3
5.6	Input points configured for each VAV agent	5.3
5.7	Temperature limits for VAV agents	5.4
5.8	Control points configured for each VAV agent	5.4
5.9	Minimum and maximum zone airflow configuration for VAV agents in physical experiments	5.5
5.10	Data points required for analysis of physical tests	5.7
5.11	Difference in energy and demand between June 28 and June 29	5.9
5.12	Difference in energy and demand for flat price experiments	5.13
5.13	Difference in energy and demand for demand limit experiments	5.17
B.1	BUILDING10 heating and cooling systems	B.1
B.2	BUILDING10 indoor and outdoor fan specifications	B.1
B.3	A sample of general data provided by the manufacturer	B.2
B.4	Sample of manufacturer's performance data	B.3
B.5	BUILDING10 cooling capacity coefficients and associated R^2 s	B.4
B.6	BUILDING10 cooling electric power coefficients and associated R^2 s	B.4
B.7	BUILDING10 heating capacity coefficients and associated R^2 s	B.5

B.8 BUILDING10 heating electric power consumption coefficients and associated R^2 s B.6

1.0 Introduction

This document describes the development and testing of market-based transactive control of commercial building heating, ventilation, and air-conditioning (HVAC) systems to serve the purposes of reducing electricity use through conservation, reducing peak building electric demand, and exposing demand flexibility to assist with electric system operations. This report is a summary of the first year of work conducted under Phase 1 of the Clean Energy and Transactive Campus Project (CETC). The methods and techniques described here were first investigated in simulation, and then later deployed in a physical testbed on the Pacific Northwest National Laboratory (PNNL) campus for validation. In this report, we describe the models and control algorithms we have developed, testing of the control algorithms in simulation, and subsequent testing of the algorithms in the physical testbed.

Algorithms were deployed as agents on the VOLTTRON platform. This work forms the foundation for a larger deployment of control methods in multiple buildings on the PNNL campus—and on project partner campuses abroad—to demonstrate the use of transactive controls in commercial buildings for constructing transactive campuses capable of energy-efficient operations that enable greater penetration of renewable resources and a more robust and reliable electric grid.

1.1 Objective

This work represents the first phase of a multi-year project to design, test, and deploy transactive controls on the PNNL campus. The goal of this work is to lay a foundation for further development and deployment of transactive controls in the small commercial buildings thereon. The algorithms, models, and software agents that are the product of this work will serve as a template for the future phases that follow. In addition to these specific products, we are guided three high-level objectives aimed at supporting our future efforts:

- Develop a scalable agent-based software application for market-based transactive building control methods.
- Demonstrate and validate the developed methods and software in a physical testbed.
- Produce deployment-ready software and documentation that allows others to adopt transactive building controls.

This report deals primarily with the theory and methodology behind market-based controls for commercial buildings, while the companion document (Corbin 2016) describes the software agents and their configuration in more detail.

1.2 Outcomes

The outcomes of the first phase of the CETC project are as follows:

- Implementation of a market-based transactive control algorithms based on prior PNNL research (Hao et al. 2016) in VOLTTRON. This has enabled us to deploy our algorithms to a secure, scalable platform capable of communicating with a building automation system (BAS).
- Development of a co-simulation capability that couples VOLTTRON to EnergyPlus. This has allowed us to validate the newly implemented VOLTTRON-based transactive control agents in simulation against a high-fidelity building energy model.

- Development of model-tuning methods that are easily automated and encapsulation of these methods in a VOLTTRON agent. Automated model tuning enables scaling of market-based controls to multiple buildings and is central to widespread adoption.
- Testing of the VOLTTRON agents and the market-based controls in physical testbed. While additional testing is required, the lessons learned through the testbed process informs future experiments and modifications to our algorithms.

1.3 Structure of Document

This document is divided into four main sections:

- Section 2.0, Transactive Market Model, covers the basics of the market-based method for commercial building HVAC cooling control applied in this project. This section describes the market structure, bidding and clearing process, and an extension of the method to heating control.
- Section 3.0, Thermal and HVAC Component Models, describes the models used in the market-based method to simulate the energy behavior and performance of HVAC components. We present the mathematical formulation of these models, their validation, and simple methods for tuning model parameters.
- Section 4.0, Market Simulation and Cross Validation, describes a co-simulation capability developed in VOLTTRON for the testing of the market-based system. Within this environment, we perform three case studies using a high-fidelity model of the physical building in which the market-based system will be deployed.
- Section 5.0, Physical Experiments, describes the physical testing performed on a small commercial building on the PNNL campus. We begin with a summary of the tests performed and a brief description of the integration and deployment of the market-based system in the building's BAS. Next, we describe our analysis methodology and conclude with a discussion of the experimental results.

The document concludes with Section 6.0, Conclusions and Next Steps, in which we briefly summarize the results of the work, lessons learned, and potential next steps.

2.0 Transactive Market Model

In this section, we describe the basic operating principles of transactive markets and their application to the control of commercial building HVAC systems. We introduce the concept of market-based control, then discuss the market structure applied to the variable-air-volume (VAV) system investigated in the CETC project, and describe the bidding and clearing process that occurs for such a system. We conclude the section with an extension of the developed method to the heating operation of commercial buildings.

2.1 Background and Significance

In the United States, the building sector accounted for the largest portion of primary energy consumption in 2010 (DOE 2011). Building energy use is expected to increase by around 31% from 2010 to 2030 (EIA 2015, DOE 2011). Employing the use of appropriate control systems can achieve highly energy-efficient buildings. Although conventional building control systems are highly capable and able to realize desired operation at relatively high levels of efficiency, they do not guarantee optimal building operation in terms of energy performance or ability to coordinate with the electric power grid (Katipamula et al. 2006).

New building control strategies are of interest in order to improve energy performance and achieve better integration with electric grid operations. These new building control strategies should have the following characteristics:

- They should not break existing control loops to simplify integration and minimize possible conflicts.
- They should be readily extensible such that limited efforts are required to adjust the control strategies if there are changes to the controlled systems, or if they are applied to different building types.
- They should enable coordination of energy use at various scales—from campus to city to region—and allow buildings to assist with the stable, efficient, and cost-effective operation of the electric grid.

Market-based control is an approach that can potentially realize these objectives.

2.2 Market-Based Control

Market-based control, an example of transactive control, is a distributed control strategy. In market-based control, a virtual market is developed to encourage the competition among controlled devices for “commodities,” such as electricity power or cooling/heating energy. Each controlled device is represented by an “agent.” Each agent is self-interested and tends to maximize its benefit. The agents submit bids for the commodities based on how much benefit they receive from a certain amount of the commodity. The market clears after receiving bids from each agent and determines the clearing price of the commodity. Each agent then adjusts its consumption based on the price.

Market-based control has several advantages that make it very suitable for large-scale application. These advantages include the following:

- Information exchange occurs only between the agents and the market. Each agent does not need to communicate with other agents. This feature makes the market-based control easy to extend. For example, if new agents are introduced into the market, no additional modification is required, in contrast to central control schemes. In addition, because the only information exchanged is the bid, the privacy of each agent is respected.

- Market abstractions provide an intuitive high-level view of the optimal allocation problem. For example, the commodity and bid price are familiar concepts to most people. This helps to establish trust and understanding in the system. This feature makes the market-based control more attractive and acceptable to building owners and operators compared to many “black-box” optimization approaches proposed in the literature.

Because of the advantages mentioned above, market-based control is getting more attention in the building research arena. For example, Li et al. (2016) developed a market-based control for a system in which a population of residential air conditioners was engaged to reduce peak electricity power. Each residential air conditioner adjusted its thermostat according to the output of the market-based control. Huberman and Clearwater (1995) designed a market-based controller for a VAV system in order to fairly distribute supply airflow among the thermal zones in a commercial building. We have previously proposed a market-based control for a commercial building VAV system, which can be used for peak load shaving, load shifting, and building energy conservation (Hao et al. 2016).

2.3 Market Structure for Commercial Building VAV Systems

Power markets are used extensively today to negotiate the delivery price and quantity of electricity in the bulk electricity generation and transmission system. The result is a dynamic flux of prices and quantities that change on a minute-to-minute basis. As consumers of electricity, we are insulated from much of the volatility of electricity prices, but this may soon change as power markets extend into the distribution system to address opportunities introduced by an increased adoption of distributed energy resources (DERs). By doing so, the flexibility of these DERs may be leveraged to address power system constraints and incentivize more efficient generation and distribution of electricity.

A power market may be organized into multiple layers consisting of regional and local markets arranged in a hierarchy: market participants bid into their local market, which then bids into the regional market above. Each level of the market provides a layer of abstraction that enables scalability and security. We can extend this multi-level market analogy to commercial building systems. That is, we can define markets within a building according to commodity (e.g., hot water), physical relationship (e.g., all VAV boxes connected to an air handler), or some combination thereof.

In this work, we have applied this concept to the control of a commercial building VAV system in order to leverage the inherent flexibility provided by the building, and to better balance cooling demands within the building, while responding to external signals indicative of power system objectives. The proposed market structure is shown in Figure 2.1,

The bi-level structure of our market consists of both cooled air and electricity as commodities. VAV agents, representing the thermal zones needing cool air for conditioning, purchase the cool air from the AirMarket. This market has a single supplier, the AHUChiller agent, which in turn purchases the electricity it requires to generate the cool air from the ElectricityMarket. The ElectricityMeter supplies electricity to the ElectricityMarket.

Although only three VAV agents are shown in Figure 2.1, there may be any number of VAVs associated with a single AHUChiller, as well as any number of AHUChillers within the building, allowing for AHUChillers to “compete” for the electricity resource. In addition, other agents may be introduced to buy electricity within the building electricity market, or if storage or distributed generation is present, to sell electricity to the electricity market. The flexibility of the structure and, importantly, the bidding and clearing process allows for any number of variations.

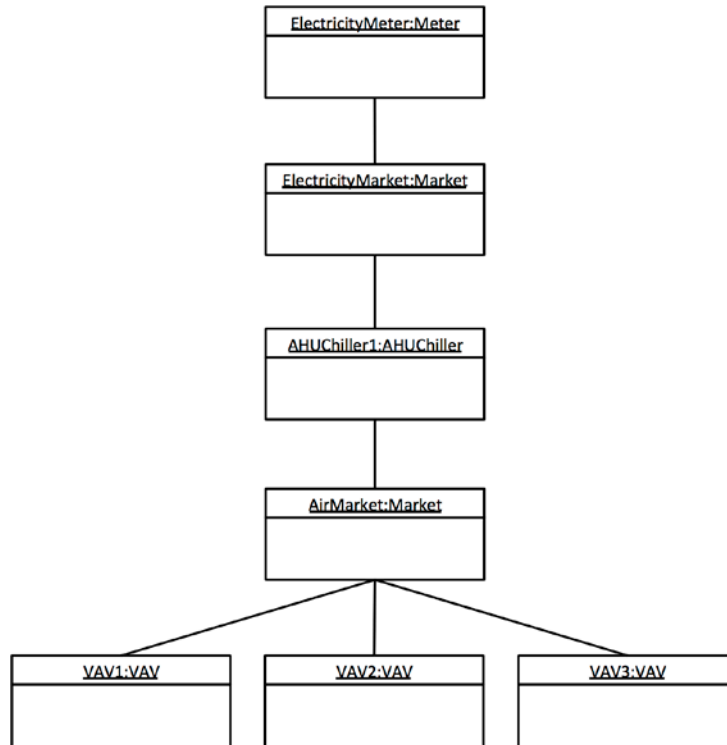


Figure 2.1. Unified Modeling Language (UML) object diagram of the two-commodity transactive market implemented in this project.

In the simulations that follow in Section 4.0, the ElectricityMeter is responsible for bidding the electricity supply curve. In alternative structures, the meter may instead be responsible for bidding the building electricity demand curve into a market one layer above. For example, it may bid into a campus, micro-grid, or distribution market where price is determined, or, an aggregate electricity demand curve at this layer may be bid into yet another layer at the regional level, or higher still.

2.3.1 General Bidding and Clearing Process

The market we have designed is a variation of a double-blind auction, in which each market participant bids to buy or sell a commodity for a given price. In contrast to other common implementations, participants do not bid single price-quantity pairs. Instead, they bid a price-quantity curve, or “flexibility curve” into their respective markets. Market participants may be both buyers in one market and sellers in another; the AHUChiller agent is an example of a market participant.

Settling of the market is a “single shot” process that begins with bidding that progresses from the bottom up and concludes with a clearing of the markets from the top down. We refer to this as “single shot” because there is no iteration required to find the clearing price or quantity at any level of the market structure. The steps of the process are as follows:

1. Each VAV generates an air demand curve and bids its curve into the air market.
2. The AirMarket aggregates the bids of the zones into an aggregate air demand curve and communicates this to the AHUChiller.¹

¹ Because the air handling unit (AHU) receives the aggregate air demand curve, the proposed market cannot be considered a pure double-blind auction. Alternatively, the AHU may instead bid an electricity demand curve based

3. The AHUChiller computes the electricity and price to serve each point on the aggregate curve.²
4. The AHUChiller submits the resulting aggregate electricity demand curve to the electricity market.
5. The ElectricityMeter submits an electricity supply curve to the electricity market.
6. The ElectricityMarket determines the intersection of electricity supply and demand curves and broadcasts the cleared electric price and quantity.
7. The AHUChiller receives the cleared electric price and quantity, and calculates the cleared air quantity.
8. The AHUChiller bids the cleared quantity into the air market as an air supply curve³.
9. The AirMarket determines the intersection of the air supply curve and the aggregate air demand curve and broadcasts the cleared air price and quantity.
10. Each VAV receives the cleared air price and adjusts its thermostat to modify its cooling demand.

Once the market has cleared, the process begins again for the next market interval, and new bids are submitted based on the updated states of the agents.

2.3.2 VAV Bidding and Clearing

VAV agents bid a demand curve defined by two points: one that represents the minimum desired cooling rate, which is associated with an increased cooling set point and maximum bid price; and one that represents the maximum desired cooling rate, which is associated with a decreased cooling set point and minimum bid price. To avoid excessive temperature deviations and minimize comfort impacts, minimum and maximum cooling set points are bounded by upper and lower limits, effectively setting the maximum and minimum cooling demand bid. In addition, cooling delivered to a zone is bounded by the physical constraints on airflow. Therefore, minimum and maximum airflow rates configured for the VAV terminal box are used to calculate the absolute boundaries for the cooling demand curve submitted by each VAV agent.

Once the market clears and the VAV agents receive the cleared price, each agent responds by adjusting its cooling demand to match the cleared price. It does so by finding the point on the cooling demand curve it previously bid into the air market that corresponds to the clearing price. Next, it calculates the temperature set point that will result in the found cooling demand. High thermostat set points are associated with reduced cooling demand, and low thermostat set points are associated with increased cooling demand. Therefore, high clearing prices will tend to drive thermostat set points up and low clearing prices will have the opposite effect. A description of the demand and temperature set point calculations is provided in Section 3.1.

on its own flexibility, thus satisfying double-blind requirements. Our method provides a tighter coupling between demand estimated at the zone level, and that bid at the electricity market level, which, in theory, guarantees that the quantity cleared at the electricity market level can be achieved at the zone level. The aggregation procedure is depicted in Figure 2.5.

² In our implementation, air and electricity prices are equivalent; the price of the maximum air demand in the aggregate air demand curve is submitted as the price of the maximum electricity demand in the electricity demand curve. However, the proposed process allows for these prices to be decoupled.

³ The air supply curve is effectively a vertical line spanning the minimum and maximum prices bid by the zones.

2.4 Hot Water Markets for Heating Operation

In this section, we extend the market-based control for VAV terminals to heating operation. We first describe the market structure for a VAV system with hot water reheat. We then introduce the problem encountered when two commodities are introduced, resulting in a decision between operating modes. We conclude with set of options for solving this problem, followed by a discussion of potential shortcomings of the proposed solutions.

2.4.1 Market Structure

The market structure explored in this section is an extension of that depicted in Figure 2.1. In the extended structure, shown in Figure 2.2, we have included an additional market and agents that represent the heating components of the VAV system. These include the HWMarket, which brokers the exchange of hot water; a Boiler agent, which supplies the hot water; a GasMarket agent, from which the Boiler agent purchases natural gas; and a GasMeter agent, which sells natural gas and represents the interface to the utility. As in the previously discussed structure, there may, in practice, be more than one supplier of hot water and more than one purchaser of natural gas, depending on the specific system configuration.

We note that there is a connection or dependency between the quantities cleared in the hot water, electricity, and cool air markets by way of the AHUChiller and VAV agents. In this example, we are only considering the case in which the air handling unit (AHU) consumes electricity and does not use hot water to pre-heat supply air. This simplifying assumption avoids some of the complexity associated with settling markets when they constrain one another's operation. When such connections exist, the problem may not be separable (and may not satisfy other necessary conditions), and an optimal result cannot be guaranteed using a market mechanism. The issue of separability and the relationship between optimization and markets is discussed briefly in Appendix A.1.

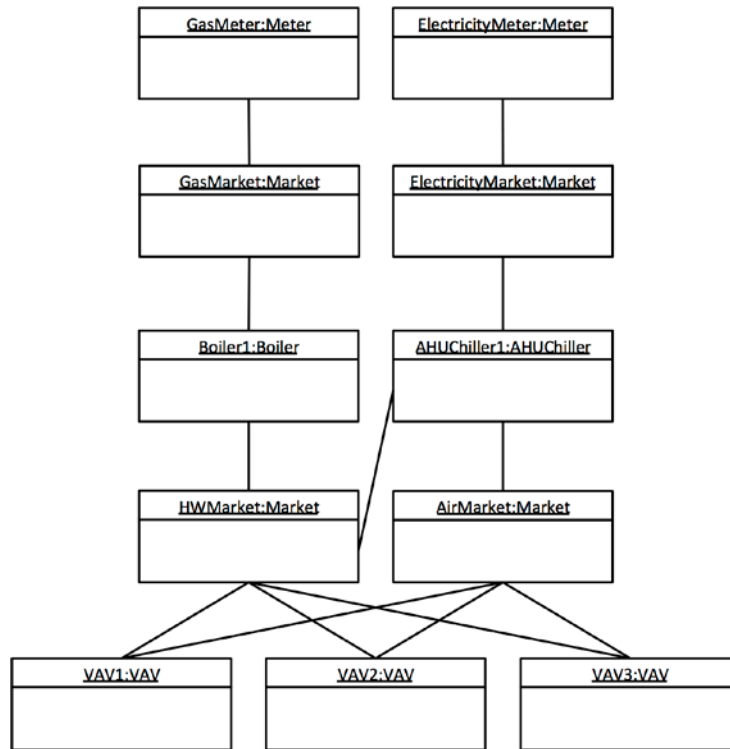


Figure 2.2. UML object diagram of a commercial building market system in which natural gas and electricity are both considered.

2.4.2 Problem Statement

The controllers in the VAV terminals regulate the position of the damper and reheat valve to maintain a zone temperature according to a set point. There are two commonly used control strategies for the VAV controllers: Single Maximum and Dual Maximum control. In Single Maximum control, the logic for the local controller of the VAV terminal is as follows, which is also depicted graphically in Figure 2.3:

- When the zone temperature is larger than the cooling set point of the HVAC zone, $T_{set,cooling}$ (cooling mode), the reheat coil valve is closed. The damper position is adjusted so that the supply air flow rate is proportional to the deviation of the zone temperature from $T_{set,cooling}$.
- When the zone temperature is between the heating set point, $T_{set,heating}$ and $T_{set,cooling}$ (i.e., the dead band), the reheat coil valve is closed while the damper position is adjusted to maintain a minimum air flow rate.
- When the zone temperature is less than $T_{set,heating}$ (heating mode), the damper position is adjusted to keep the air flow rate as a constant value and the position of the reheat coil valve is proportional to the deviation of the zone temperature from $T_{set,heating}$.

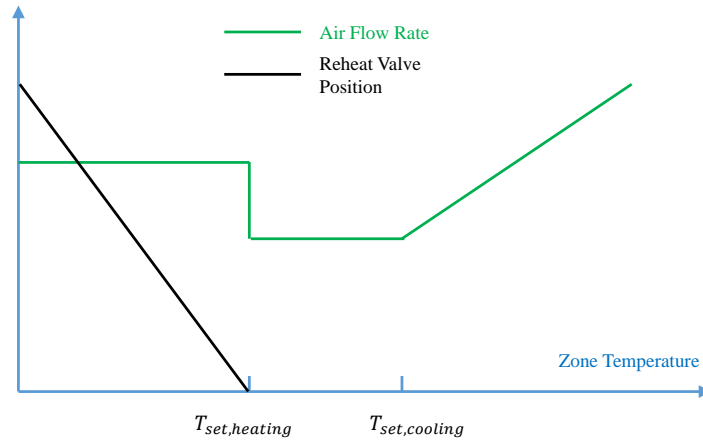


Figure 2.3. Single Maximum control logic for a VAV terminal.

In Dual Maximum control, the logic for the VAV terminal follows a slightly modified sequence shown in Figure 2.4:

- When the VAV terminal is in cooling mode, or when the zone temperature is in the dead band, the control logic for the reheat coil and the damper is the same as that in Single Maximum control.
- When the VAV terminal is in heating mode, the position of the reheat coil valve is adjusted to be proportional to the deviation of the zone temperature from $T_{set,heating}$. If the valve of the reheat coil is fully open, but the zone temperature is still less than $T_{set,heating}$, the damper position is adjusted so that the supply air flow rate is proportional to the deviation of the zone temperature from $T_{set,heating}$.

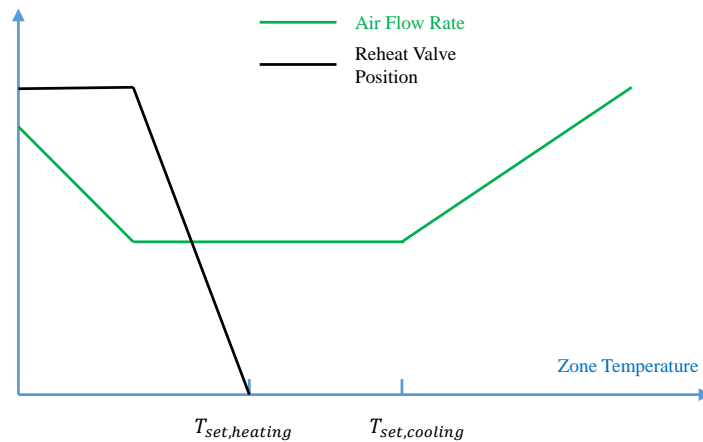


Figure 2.4. Dual Maximum control logic for a VAV terminal.

We can see that by modulating $T_{set,heating}$ and $T_{set,cooling}$ for each VAV terminal, we are able to adjust their heating/cooling demand. Adjusting the heating/cooling demand can serve the following purposes:

- Limit the peak electricity or natural gas demand in order to increase the efficiency on the energy supply side. At the same time, we may still deliver the total amount of heating or cooling needed by the zone served by the VAV terminal.
- Enable VAV terminals to respond to changing prices for electricity and natural gas. Specifically, VAV terminals may be controlled such that additional cooling or heating energy is delivered when

prices are low, and less cooling or heating is delivered when prices are high. By responding to the changing price, we may reduce the cooling and heating cost and increase the efficiency on the energy supply side as well.

- Reduce the electricity or natural gas demand. By better allocating the cooling and heating energy among VAV terminals, we can reduce the total cooling and heating demand. For example, if only one HVAC zone (critical zone) is requesting more cooling while other HVAC zones are overcooled, instead of increasing the total cooling energy supply, we can assign more cooling energy to the critical zone.

2.4.3 Market-Based Control for VAV Terminal

Several market-based control strategies for VAV terminals have been proposed in the literature (Hao et al. 2016; Huberman and Clearwater 1995). In those strategies, VAV terminals were assumed to be in cooling mode. In this section, we propose a market-based control strategy for VAV terminals in heating/reheating mode:

1. Determine the mode, either heating or cooling, into which the VAV terminal will be placed.
 - a. Consider the several strategies that can be used to make such a decision:

- Prediction-based strategy

A prediction of the HVAC load for the next market interval is generated. If the predicted load market is negative (cooling), the VAV terminal will be placed in cooling mode, otherwise, it will be placed in heating mode.

- Budget-based strategy

Each VAV terminal receives an amount of virtual money, m , for bidding. Based on m , we estimate the corresponding maximum cooling energy $\dot{Q}_{cool,max}$ and the maximum heating energy $\dot{Q}_{heat,max}$ that we can afford:

$$\dot{Q}_{cool,max} = \frac{m}{p_{cool,pre}} \quad (2.1)$$

$$\dot{Q}_{heat,max} = \frac{m}{p_{heat,pre}} \quad (2.2)$$

where $p_{cool,pre}$ and $p_{heat,pre}$ are the price of the cooling and heating energy at the previous market interval. If $\dot{Q}_{cool,max} \geq \dot{Q}_{heat,max}$, the VAV terminal will move to/stay in the cooling mode, otherwise, it will be placed in heating mode. The value of m can be determined based on the cooling design load of the VAV terminal:

$$m \propto \dot{Q}_{cool,design} \quad (2.3)$$

- Cost-based strategy

Based on a user-defined range of acceptable zone temperature, $[T_{low,lim}, T_{high,lim}]$, we can determine the maximum cooling and heating energy as follows:

$$\dot{Q}_{cool,max} = f_{zone}(T_{low,lim}) \quad (2.4)$$

$$\dot{Q}_{heat,max} = f_{zone}(T_{high,lim}) \quad (2.5)$$

If $\dot{Q}_{cool,max}p_{cool,pre} \leq \dot{Q}_{heat,max}p_{heat,pre}$, the VAV terminal will be placed in cooling mode, otherwise, it will be placed in heating mode.

2. Develop the following functions to connect heating set points to heating energy:

$$T_{set,heating} = f_{heating}(\dot{Q}_{heating}) \quad (2.6)$$

where $\dot{Q}_{heating}$ is the heating energy delivered by the supply air. Similarly, we can develop the following function to connect the cooling set point to cooling energy:

$$T_{set,cooling} = f_{cooling}(\dot{Q}_{cooling}) \quad (2.7)$$

where $\dot{Q}_{cooling}$ is the cooling energy delivered by the supply air.

3. Each VAV terminal develops the following function to connect price to cooling energy:

$$\dot{Q}_{cool} = f_{cool,demand}(p_{cool}) \quad (2.8)$$

where p_{cool} is the price of the desired cooling energy and $f_{cool,demand}$ is a non-increasing function. Then we can determine p_{cool} according to one of two strategies described below.

- Budget-based strategy

Using the previously determined budget, m , we can define $p_{cool,min}$ as follows:

$$p_{cool,min} = \frac{m - p_{hot,water}\dot{Q}_{hot,water,max}}{\dot{Q}_{cool,max}} \quad (2.9)$$

$$p_{cool,max} = \frac{m - p_{hot,water}\dot{Q}_{hot,water,min}}{\dot{Q}_{cool,min}} \quad (2.10)$$

where $p_{hot,water}$ is the price of hot water and $\dot{Q}_{hot,water,min}$ and $\dot{Q}_{hot,water,max}$ are the minimum and maximum heating energy available from the hot water, respectively.

- Historical data-based strategy

Alternatively, the price may be a function of historical prices:

$$p_{cool,min} = \bar{p}_{cool,pre} - k \quad (2.11)$$

$$p_{cool,max} = \bar{p}_{cool,pre} + k \quad (2.12)$$

where k is the user-defined parameter within a given range $[k_{min}, k_{max}]$, and $\bar{p}_{cool,pre}$ is the average cooling price over a preceding interval, e.g., 24 hours.

4. Aggregate the received cooling demand curves by summing the bid quantities as shown in Figure 2.5. Note that this is the same process mentioned in Section 2.3.1.




Figure 2.5. The aggregation of VAV demand curves.

- a. After the aggregated cooling load demand curve is created, we convert it into an electricity demand curve (Figure 2.6). In this curve, the maximum and minimum electricity demand is determined according to the following functions:

$$E_{max,cool} = f_{energy,load}(\dot{Q}_{cool,ag,max}) \quad (2.13)$$

$$E_{min,cool} = f_{energy,load}(\dot{Q}_{cool,ag,min}) \quad (2.14)$$

where $\dot{Q}_{cool,ag,max}$ and $\dot{Q}_{cool,ag,min}$ are the maximum and minimum aggregated cooling energy, respectively. The transform $f_{energy,load}$ is specific to the system modeled.




Figure 2.6. The conversion between power and electricity demand curves.

- b. Obtain the price for electricity by

$$p_{ele,min} = p_a - \sigma \quad (2.15)$$

$$p_{ele,max} = p_a + \sigma \quad (2.16)$$

where p_a is the average electricity price over a past period and σ is the standard deviation. Based on Equations (2.13), (2.14), (2.15), and (2.16), we may generate a demand curve to be bid into the local electricity market.

5. After the clear price for electricity is obtained, we can find the price of cooling power according to the one-to-one relationship between the electricity and cooling power prices. Each VAV terminal can determine the cooling power it will use according to its demand curve.
 - a. Based on cooling power, we can calculate the zone set point according to Equations (2.6) and (2.7).
 - b. Finally, based on the cleared cooling power, we can calculate the clearing price for heating energy as follows:

$$p_{heat} = \frac{p_{cool}\dot{Q}_{cool,tot} + p_{hot,water}\dot{Q}_{hot,water,tot}}{\dot{Q}_{heat,tot}} \quad (2.17)$$

where $\dot{Q}_{cool,tot}$ is the cooling energy cleared in the cool air market, $\dot{Q}_{hot,water,tot}$ is the heating energy cleared in the hot water market, and $\dot{Q}_{heat,tot}$ is the sum of net ($\dot{Q}_{cool,tot} + \dot{Q}_{hot,water,tot}$) energy cleared in the market. The effective heating price, p_{heat} , is used in subsequent bidding in the next market interval.

2.4.4 Discussion

One key parameter in the budget-based strategy is the quantity m . We propose to determine m according to the cooling design load of the VAV terminal. However, this may potentially cause problems in the real implementations. For example, it is possible that $\dot{Q}_{cool,max}$ and $\dot{Q}_{cool,min}$ of one VAV terminal (named the VAV Terminal 1) are much larger than those of other VAV terminals when all of the terminals are in cooling mode. Then $p_{cool,min}$ and $p_{cool,max}$ should be much less than those of other VAV terminals according to Equations (2.9) and (2.10). In that case, as shown in Figure 2.7, it is likely that VAV Terminal 1 will always be assigned $\dot{Q}_{cool,min}$. In the future study, we will explore ways to avoid this problem by better selecting m .




Figure 2.7. The problem caused by inappropriately selected m .

In Equation (2.17), we propose a heating price calculation based on the cleared quantities and prices of the cool air and hot water markets by way of $\dot{Q}_{heat,tot}$. However, this quantity may approach zero, in which case the heating price approaches infinity, significantly affecting the bid for heating by the agent in subsequent market-clearing intervals. To solve this problem, we may instead define p_{heat} as follows:

$$p_{heat} = \frac{|p_{cool}\dot{Q}_{cool,tot}| + |p_{hot,water}\dot{Q}_{hot,water,tot}|}{|\dot{Q}_{cool,tot}| + |\dot{Q}_{hot,water,tot}|} \quad (2.18)$$

We note that in either definition, we have not yet performed simulations to test the effect of the effective heating price on market behavior.

3.0 Thermal and HVAC Component Models

This section describes the models used by the agents in the transactive system to estimate the quantities bid into their respective markets. These models are control-oriented models, the majority of which are inverse models, and therefore relatively simple compared to those used in detailed energy simulation. For each model, we describe the mathematical formulation and the method by which we tune the models to predict the performance of the physical system.

3.1 Zone Model

As described earlier, the market-clearing process begins with the bidding of individual zone-level demand curves by the VAV agents, which are aggregated by the AHUChiller agent into a system-level demand curve. In order to generate such zone-level demand curves for each zone, we use a simple zone model to predict the HVAC energy demand based on outdoor dry-bulb temperature and other zone parameters. Once the market clears, an analogous model is used to find the new cooling set point temperature. To ensure an equitable market clearing and appropriate control response, it is important that these models have accurate predictive capabilities. Moreover, as the market system is extended in future work to include forward market settlements, the predictive ability of the models must be accurate over long time frames, from several hours to ideally an entire day.

Lumped parameter models have long been used for predicting the performance of a building for applications where the building level response that is both sufficiently accurate and computationally efficient is of greater interest (Gouda et al. 2002). These models can take a basic lower-order form or increase in complexity to a higher-order combination of resistances and capacitances such as the 3R2C or the 3R4C models (Fraisse et al. 2002). These models are popular because of the simplicity and adequacy they offer in capturing the thermal response of a zone.

3.1.1 Model Formulation

The zone model is initially formulated as a linear regression-based approximation of the zone temperature at a given time-step as a function of outdoor dry-bulb temperature, zone temperature at the previous time-step, and HVAC energy (Hao et al. 2016). This initial form of the model is expressed as follows:

$$T_{t+1}^i = a_1^i T_t^i + a_2^i T_{o,t}^i + a_3^i q_{t+1}^i + a_4^i q_t^i + a_5^i \quad (3.1)$$

or by rearranging the terms as follows:

$$q_{t+1}^i = b_1^i T_t^i + b_2^i T_{o,t}^i + b_3^i T_{t+1}^i + b_4^i q_t^i + b_5^i \quad (3.2)$$

where T_{t+1}^i and q_{t+1}^i represent the zone temperature and cooling rate at the next time-step, and $a_4^i q_t^i + a_5^i$ and $b_4^i q_t^i + b_5^i$ estimate the unobserved external disturbances as measured by the cooling rate at the current time, q_t^i .

The model is trained using measured data for a given thermal zone within the building and validated by comparing the predicted and measured zone temperatures and cooling rate. Figure 3.1 shows a comparison of the measured and predicted zone temperature and cooling rate for a thermal zone within the building using data aggregated at 60-minute intervals. Because the model corrects T_t and q_t at each time-step, the predicted response matches the measured response very closely.




Figure 3.1. Comparison of actual and predicted zone temperature and cooling rate for a sample thermal zone within the building.

The initial zone model is then tested for multi-period prediction by propagating the predicted T_t and q_t through the next time-step, $t + 1$. In the absence of a correcting signal, the predicted response for the cooling rate (q) is observed to quickly diverge from the measured response. A step-wise tuning process is employed to identify and correct deficiencies identified from the residual plots of the initial model. The first deficiency observed is that the model does not capture residual information when the HVAC system is unavailable (scheduled off); Figure 3.2 illustrates this clearly.




Figure 3.2. A residual plot of the predicted cooling rate (q) for the initial model.

We correct this deficiency by introducing the HVAC system availability as a model parameter and modifying the form of the model as follows:

$$T_{t+1}^i = a_1^i T_t^i + a_2^i h_t + a_3^i T_{o,t}^i + a_4^i T_t^i h_t + a_5^i q_t^i h_t + a_6^i q_{t+1}^i h_{t+1} + a_7^i \quad (3.3)$$

and

$$q_{t+1}^i = a_1^i T_{o,t}^i + a_2^i q_t^i + a_3^i h_{t+1}^i + a_4^i T_t^i x h_t + a_5^i T_{t+1}^i x h_{t+1} + a_6^i \quad (3.4)$$

where h_t and h_{t+1} indicate the HVAC system availability at time-step t and $t + 1$, and can have a value of 0 or 1 based on whether the HVAC system was off or on. The availability is specified as an interactive component for the zone temperature (T) and cooling rate (q) terms. This modification allows the model to capture previously untapped information, thereby improving the residuals of the model when the system was unavailable as illustrated in Figure 3.3.




Figure 3.3. A residual plot of the predicted cooling rate (q) for the revised model.

The modification also improves the predictive performance of the model compared to the initial case as illustrated in Figure 3.4.

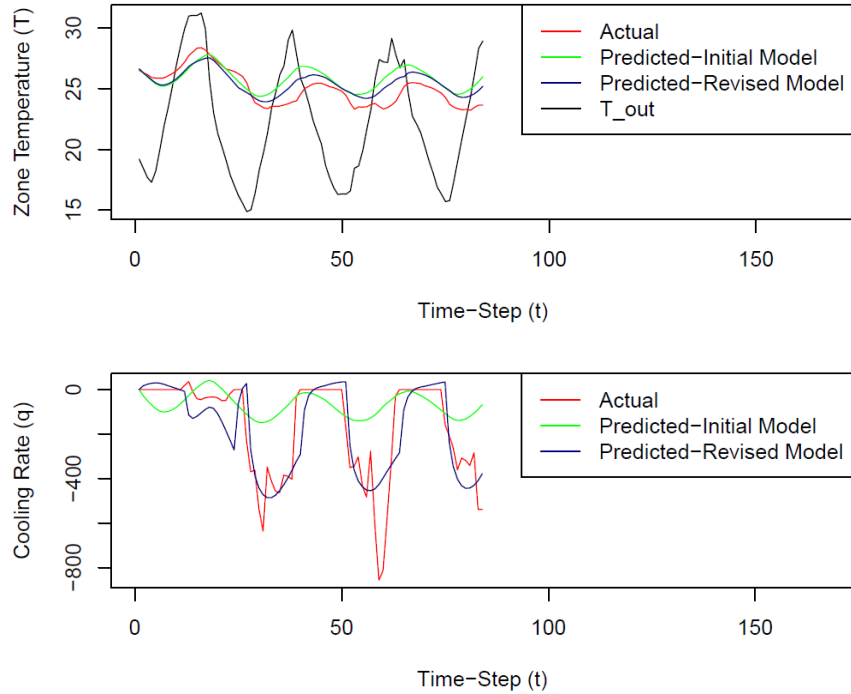


Figure 3.4. A comparison of the predictive performance of the initial and revised model with the actual measured performance for a sample thermal zone within the building.

3.1.2 Parameter Identification

The coefficients of the regression models establish the nature of the relationship between the variable terms and the response. It is important to distinguish the function of the coefficients of a regression model from that of a heat balance model, because the regression coefficients establish the correlation between the trend of the associated variable and the trend of the response. By examining each term in isolation and assuming that all other variables remain unchanged from one time-step to the next, the sign of each coefficient can be determined based on whether the expected change in the response is in the same direction as the change in the associated variable (positive coefficient) or opposite in trend (negative coefficient).

When considering the temperature model, we note that the response, i.e., zone temperature at the next time-step, is likely to increase as the zone temperature at the current time-step increases, all other variables being held constant. Thus, the coefficient associated with the T_t^i term, a_1^i , is expected to be positive. We apply the same logic to the other terms leading to an expected negative coefficient for the HVAC system availability term, a_2^i ; a positive coefficient for the outdoor air temperature term, a_3^i ; a positive coefficient for the combined effect of the zone temperature and HVAC system availability at the current time-step, a_4^i ; a positive coefficient for the combined effect of the cooling rate and the HVAC system availability at the current time-step, a_5^i ; and the next time-step, a_6^i . The constant term, a_7^i , defines the fairly constant internal occupancy-related gains within the zone that are not captured by the remaining terms. As the amount of internal gain increases, the zone temperature is expected to increase. Thus, the constant term is expected to be positive.

Assessing the cooling rate model using the same principles, we expect the cooling rate to increase with an increase in the outdoor air temperature, all other variables being held constant. However, because the

cooling rate is expressed as the rate of heat extraction, the sign associated with it is negative. Thus, the cooling rate becomes more negative as the outdoor air temperature increases so a negative coefficient is expected to be associated with the outdoor air temperature term, b_1^i . Using the same logic, a positive coefficient, b_2^i , is expected to be associated with the cooling rate at the current time-step; a negative coefficient, b_3^i , is expected to be associated with the HVAC system availability at the next time-step; a negative coefficient, b_4^i , is expected to be associated with term representing the combined effect of the zone temperature and HVAC system availability at the current time-step; and a positive coefficient, b_5^i , is expected to be associated with the term representing the same effect at the next time-step because if the zone temperature at the next time-step increases, it must be associated with a less negative or increased cooling rate. Similarly, an increase in the internal loads represented by the constant term, b_6^i , is expected to result in an increase in the cooling rate or a more negative cooling rate, thereby resulting in an expected negative coefficient. The coefficients $a_1^i \dots a_7^i$ and $b_1^i \dots b_6^i$ are found through linear regression, where t is the current time-step and $t+1$ is the next time-step.

We have developed a VOLTTRON agent to periodically and automatically perform parameter identification for the zone model. As the model and tuning method are revised, this agent will be updated. More detail about this agent is provided in the *Transactive Control of Commercial Building HVAC Systems: VOLTTRON User Guide* (Corbin 2016).

The coefficients obtained from the model indicate a strong dependence of the zone temperature at a given time-step on the zone temperature at the previous time-step. While this trend is expected, the coefficients associated with the other terms except the HVAC system availability are negligible, indicating a persistence issue. As observed in Figure 3.4, the model in its current form is unable to capture the effects of thermal capacitance. This will be addressed in future modifications of the model.

3.1.3 Future Work

The introduction of lagged temperature and cooling rate as additional model parameters to account for the effects of capacitance is currently being investigated. Based on the results of the investigation, a higher-order resistance-capacitance (RC) model maybe considered for implementation. The potential issue of persistence described above is also under investigation.

Additionally, as seen in Table 3.1, some model coefficients for certain thermal zones are observed to have an opposite trend to what is expected based on the principles of thermodynamics. Constraining the sign of the coefficients during the regression process alleviates the issues related to the sign of the coefficients but the problem of negligible coefficients associated with terms other than the temperature of the previous time-step still persists. Additional investigation is required to identify the source of these issues.

Table 3.1. Parameters for a sample thermal zone within the building.

T_{t+1}^i Coefficient	Value	q_{t+1}^i Coefficient	Value
a_1^i	0.81119	b_1^i	-2.563
a_2^i	1.40796	b_2^i	31.52
a_3^i	0.06018	b_3^i	- 2416
a_4^i	-0.05369	b_4^i	-0.02428
a_5^i	0.00074	b_5^i	86.31
a_6^i	0.00068	b_6^i	65.82
a_7^i	3.5580		

3.2 Air Handler Model

The air handler model is a simple representation of a typical commercial VAV unit. The purpose of the model is to estimate fan power and cooling load given real-time measurements from the BAS. The AHU consists of a variable-speed supply fan and chilled water coil; the return fan, which is approximately one-tenth the size of the supply fan is neglected. Because only cooling operation is considered currently, the hot water coil is not modeled. A more detailed representation of the AHU is not necessary for the transactive control algorithms; other aspects of an AHU normally considered in simulation, i.e., economizer operation, are not modeled.

3.2.1 Model Formulation

Supply fan power is estimated using a linear model based on fan laws, which relate volume flow rate and pressure rise to electric power consumption. Specifically, fan power, P_{fan} , is approximated as a function of volume flow rate, \dot{v} , and static pressure, ρ . This gives rise to the following equation:

$$P_{fan} = a_0 + a_1\dot{v} + a_2\dot{v}^2 + a_3\dot{v}^3 + a_4\rho + a_5\rho^2 \quad (3.5)$$

Cooling coil load, \dot{Q}_{cool} , is modeled simply from first principles using the familiar heat transfer equation:

$$\dot{Q}_{cool} = \dot{m}_{air}C_p(T_{supply} - T_{mixed}) \quad (3.6)$$

where

- \dot{m}_{air} = the mass flow rate of air entering the cooling coil,
- T_{mixed} = the mixed air temperature,
- T_{supply} = the supply air temperatures, and
- C_p = the heat capacity of air¹.

Air mass flow rate is calculated from the volume flow rate; this and all parameters in the preceding equations are available from the BAS.

3.2.2 Parameter Identification

Coefficients of the fan power model are found using least-squares regression. Measured fan power, volume flow rate, and system static pressure were obtained from the BAS. These data were recorded at 1-minute intervals over approximately 2 weeks of normal building operation. The data were split into training and validation sets.

We have developed a VOLTTRON agent to periodically and automatically perform parameter identification for the fan model. More detail about this agent is provided in the user guide (Corbin 2016).

3.2.3 Model Validation

We validate the fan model by comparing power predicted using measured pressure and flow in the validation set to measured power over the same time period. Figure 3.5 below shows a comparison between measured and predicted power for a single day in the validation data set. Model fit is extremely

¹ We assume a fixed value for the heat capacity of dry air, calculated at the nominal supply air temperature.

good and has an adjusted R^2 value of 0.9966 and mean absolute error of 0.04kW, or 4.4% of average power.




Figure 3.5. Comparison of measured and model-predicted fan power during 1 day of the validation data set.

While economizer operation is not modeled explicitly, its operation is reflected in the cooling coil load. This allows economizer operation to be considered in the transactive market, even though it is not under direct control.

3.3 Chiller Model

A simple chiller model estimates the electric demand of the district chilled water plant required to serve the cooling load calculated by the AHU. In practice, this model is embedded inside the AHU model. While trivially simple, this model provides a mechanism for the AHU to estimate the electric demand associated with a given cooling load so that it may be combined with the fan power estimate and bid into the electric market.

3.3.1 Model Formulation

The chiller model is a single parameter model that relates chilled water load to electric demand. The single parameter—coefficient of performance (COP)—can be easily estimated from total chiller plant electric demand, $P_{chiller}$, and total cooling load delivered, \dot{Q}_{cool} , which is estimated by the AHU model. This model takes the following form:

$$P_{chiller} = COP * \dot{Q}_{cool} \tag{3.7}$$

The decision to adopt such a simplified model was not arbitrary; rather it represents recognition that the electric demand required to satisfy the cooling load only needs to be a first-order approximation. In practice, chiller plant performance is characterized by a large number of parameters, including outdoor temperatures, condenser type, supply and return temperatures, chiller staging, and part-load ratio. However, in the context of the transactive market considered here, such a detailed model is not necessary,

and adds unnecessary complexity to the demand curve estimation described previously. Our rationale is as follows:

- The air handler considered here is just one consumer of the chilled water in the district system, thus the impact it has on chiller power consumption changes as a function of the total chiller load. While this information could be obtained at the time of electric demand estimation, the actual impact will vary with the operation of other chilled water consumers, of which the air handler has no knowledge. This dependence on the behavior of other consumers makes the problem inseparable as discussed previously. The simplification adopted here, i.e., that the performance of the chiller is independent of part-load ratio, allows power to be simply estimated regardless of the actions of other consumers in the system.
- At the building level, the power used by the chiller to satisfy the building's cooling load is significant relative to power used by the AHU fan. Therefore, the higher-order effects introduced by environmental conditions are arguably less important than the gross demand of the chiller. At the building level, fan electric and chilled water demand are the only quantities over which the market has any control. Chiller electric power, while important, is only necessary in the market to estimate the impact from a change in chilled water demand; the simple model simply scales and converts the units of chilled water demand into electric demand so that it may be bid with the fan.

3.3.2 Parameter Identification

Chiller plant COP was simply determined from historical data by dividing total cooling delivered by total electric power and averaging over a month of data recorded during the summer. An average value of 5.69 was found by fitting a linear model to the data; the R^2 value of 0.924 shows a linear model to be a good approximation. The standard deviation of the COP measures 0.65, or 11% of the average. Figure 3.6 shows a scatter plot of chiller plant cooling load and chiller plant electric demand during the month of July and the linear model fit used to estimate COP.

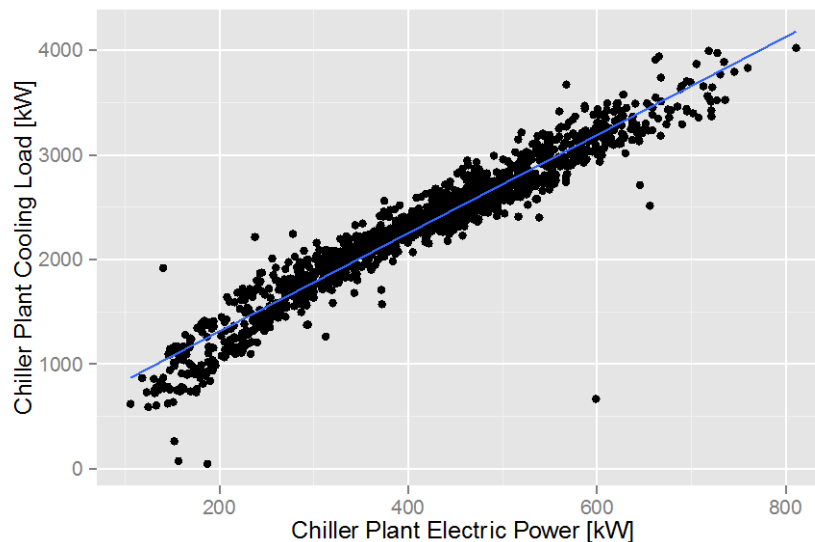


Figure 3.6. Scatter plot of total chiller plant cooling load versus plant electric demand. The blue linear regression line represents the average COP of the system.

3.3.3 Model Limitations

The chiller model is a gross simplification of the factors affecting chiller performance. The model described here is, at best, a first-order approximation of the relationship between load and electricity demand. Obviously, such a simple model is limited in its ability to accurately predict power consumption under changing conditions. As the current project scales to additional buildings, an improved model may be required to fully exploit the flexibility of chiller power demand, and critically, the accuracy of power predictions. The model may be improved significantly with the introduction of environmental conditions and detailed modeling of subcomponents of the plant, e.g., condenser system and pumps. However, the introduction of part-load characteristics remains problematic as the transactive market comes to include a greater portion of the chiller demand. Future investigations may consider solutions to this problem.

3.4 Rooftop Unit Models

A set of rooftop unit models were developed for the future deployment of transactive controls to commercial buildings that have one or more zones conditioned by packaged rooftop heat pumps. Although the models were not deployed or tested during the project, this section describes systems used to develop the models, data gathered, and modeling methods considered and adopted. In future work, these models will be deployed across the PNNL campus.

3.4.1 Introduction

Rooftop units (RTUs) are among the building systems that can be used to participate in transactive mechanisms to reduce demand or absorb extra generation. Heat pumps (HPs) have been more widely used in RTUs in the past decade because of their high efficiency and convenience in providing both heating and cooling in a single packaged RTU. Although HPs are energy efficient at the site level (i.e., building scale), they increase electricity consumption and contribute to electric power peak demand. This is because electricity is used to provide heating in addition to cooling. Hence, it is important to control RTU-HPs demand to not just conserve energy but also reduce power especially during high demand hours. To integrate RTU-HPs in a transactive mechanism and control their demand, it is important to estimate their consumption at a given time. Therefore, we require simple but representative control-oriented RTU-HP models to predict their capacity and power use.

Two buildings on the PNNL campus are likely candidates for RTU-HP deployment; their rooftop systems have been used to inform the development of the models discussed here. Both buildings, referred to in this section as BUILDING6 and BUILDING10, are similar in size at approximately 20,000 gross ft² (1858.0608 m²), and consist of office spaces, conference rooms, and a kitchen. Each is served by 11 RTU-HPs.

3.4.2 BUILDING6 Rooftop Unit Characteristics

Heat pumps installed in BUILDING6 consist of four different *Trane*[™] models. Here, these heat pumps are categorized as models A, B, C, and D. There are two model As, two model Bs, one model C, and five model Ds whose cooling and heating specifications are listed in Table 3.2.

Table 3.2. BUILDING6 heating and cooling systems.

Model	Net Cooling Capacity (kW)	Cooling Power (kW)	Cooling EER	Nominal Supply Airflow (m ³ /h)	Net Heating Capacity (kW)	Heating Power (kW)	Heating COP
A	22.0	6.58	11.4	4,080	21.0	5.95	3.50
B	17.6	6.74	8.90	3,400	17.3	5.76	3.50
C	23.0	8.40	9.41	3,567	21.4	6.60	3.27
D	18.6	5.13	12.09	3,400	17.1	4.94	3.50

EER = energy efficiency ratio

The BUILDING6 heat pump fan types and specifications are listed in Table 3.3 for both the indoor and outdoor fans. BUILDING10's system characteristics are summarized in Appendix B (Section B.1, BUILDING10 System Characteristics).

Table 3.3. BUILDING6 indoor and outdoor fan specifications.

Model	Outdoor Fan			Indoor Fan	
	Type	Airflow (m ³ /h)	Motor Power (kW)	Type	Motor Power (kW)
A	Propeller	5,800	0.52	FC Centrifugal	0.75
B	Propeller	3,200	0.30	FC Centrifugal	1.10
C	Propeller	4,550	0.25	FC Centrifugal	0.75
D	Propeller	5,138	0.30	FC Centrifugal	0.75

3.4.3 Model Formulation

Existing RTU-HP modeling methods can be classified in two broad categories of empirical models (e.g., regression-based polynomial curve-fit models) and physics-based models (e.g., gray-box and white-box models) (Gayeski et al. 2010). Each modeling approach assumes a different set of inputs and/or training data to characterize the performance of the RTU.

We consider two RTU modeling methods, selected based on the data sources commonly available for rooftop heat pumps: a simple nominal performance model based on manufacturer nameplate values; and an inverse model developed from either manufacturer's catalogs/specifications or measured performance data. These methods and their corresponding data sources are summarized in Table 3.4. We categorize the data source by level of effort and/or expense to obtain.

Table 3.4. Data and modeling methods classification.

Level of Data	Source of Data	Type of Data	Modeling Method Considered
1	Manufacturer	Nominal	Nominal
2	Manufacturer	Performance data (as-designed)	Inverse
3	Measured	Time-series data (as-installed)	Inverse

Level 1 data are easily gathered from manufacturer's catalogs and specifications based on a systems' model/nameplate number. A sample of such data is shown in Appendix B (Section B.2, Nameplate Equipment Characteristics). Level 2 data may or may not be available on the model and vintage. Level 2 data consist of capacity and/or power data measured at different ambient temperature, entering wet bulb temperature, entering dry-bulb temperature, and airflows. A sample of such data is shown in Appendix B (Section B.3, Catalog Performance Data).

Field, or *as-installed* performance data, compose Level 3 data and are less likely to be available, or only available with considerable effort or expense. Data points that must be collected to characterize an RTU are summarized in Table 3.5.

Table 3.5. Data points required to characterize a rooftop heat pump using the inverse method.

BUILDING10	Typical Value/Unit
Auxiliary Heat Command	0=inactive, 1=active
Electric Heat Command	0=inactive, 1=active
Electric Power Demand	kW
Compressor Command	0=inactive, 1=active
Cooling Temperature Set Point	°F, °C
Discharge Air Temperature	°F, °C
Heating Cooling Mode	1=COOL, 2=HEAT, 3=AUTO
Heating Temperature Set Point	°F, °C
Occupancy Mode	0=inactive, 1=active
Outdoor Air Temperature	°F, °C
Supply Fan Command	0=inactive, 1=active
Unoccupied Cooling Temperature Set Point	°F, °C
Unoccupied Heating Temperature Set Point	°F, °C
Zone Temperature	°F, °C

Note that fan power consumption (both indoor and outdoor) is typically included in the total system performance data (i.e., total electric power) provided by each manufacturer. This applies to both Level 1 and Level 2 data sources.

3.4.3.1 Fan Heat Correction

Not all manufacturer catalogs report net cooling capacity, requiring that it be calculated by subtracting indoor fan motor heat from gross cooling capacity. Fan heat is calculated at nominal airflow and different external static pressures. For example, the manufacturer suggests using the following equations to calculate fan motor heat for the units in BUILDING6:

Table 3.6. Fan motor heat calculation for BUILDING6 RTUs

Model	Fan Motor Heat (kW)
A	$(2.829 \times \text{Fan bhp} + 0.4024) \times 0.293$
B	$1.144 \times \text{Fan kW} + 0.132$
C	$1.375 \times \text{Fan kW}$
D	$(2.7672 \times \text{Fan bhp} + 0.4705) \times 0.293$

Additionally, manufacturers may provide fan power (kW or bhp) at different external static pressures and airflow rates. To estimate fan motor power, nominal airflow may be used at different external static pressures. Table 3.7 lists fan powers calculated at a range of external pressures.

Table 3.7. Fan motor power calculated at different external static pressures.

	External Static Pressure (Pa)				
	100	150	200	250	300
Model A	0.62	0.74	0.85	0.96	1.05
Model B	0.76	0.88	0.99	1.12	1.28
Model C	0.66	0.74	0.81	0.87	0.91
Model D	0.51	0.61	0.71	0.82	0.93

This correction may be done once, as in the case of Level 1 data, or for each gross cooling capacity value in the manufacturer supplied performance data, as in the case of Level 2 data. In Level 3, net capacity is calculated directly from measured supply temperature, return temperature, outdoor air temperature, outdoor air fraction and supply flow rate.

3.4.3.2 Nominal Performance Model

Model number, nominal electric power, and capacity data were first extracted from manufacturer's catalogs (see tables included in Appendix B). From nominal data, the power consumption, $P_{cool,elec}$, and system cooling capacity, \dot{Q}_{cool} , can be estimated for cooling operation as shown in Equations (3.8) and (3.9)

$$\dot{Q}_{cool} = \begin{cases} \dot{Q}_{cool,nom}, & \text{if } C = 1 \\ \dot{Q}_{vent} & \text{if } C = 0 \end{cases} \quad (3.8)$$

$$P_{cool,elec} = \begin{cases} P_{cool,nom}, & \text{if } C = 1 \\ P_{fan} & \text{if } C = 0 \end{cases} \quad (3.9)$$

where $\dot{Q}_{cool,nom}$ = the nominal net cooling capacity of the unit,
 $P_{cool,nom}$ = the nominal electric power of the unit during cooling operation,
 P_{fan} = the nominal fan power, and
 C = a flag indicating if the compressor is on.

The equivalent nominal performance models for heating operation are given by:

$$\dot{Q}_{heat} = \begin{cases} \dot{Q}_{heat,nom}, & \text{if } C = 1 \text{ and } A = 0 \\ \dot{Q}_{heat,nom} + \dot{Q}_{aux}, & \text{if } C = 1 \text{ and } A = 1 \\ \dot{Q}_{vent} & \text{if } C = 0 \text{ and } A = 0 \end{cases} \quad (3.10)$$

$$P_{heat,elec} = \begin{cases} P_{heat,nom}, & \text{if } C = 1 \text{ and } A = 0 \\ P_{heat,nom} + P_{aux}, & \text{if } C = 1 \text{ and } A = 1 \\ P_{fan} & \text{if } C = 0 \text{ and } A = 0 \end{cases} \quad (3.11)$$

where $\dot{Q}_{heat,nom}$ = the nominal net heating capacity of the unit,
 \dot{Q}_{aux} = the nominal auxiliary heating coil capacity,

- $P_{heat,nom}$ = the nominal electric power of the unit (including the fans) during heating operation,
 P_{aux} = the nominal auxiliary heating coil electric power, and
 A = a flag indicating if the auxiliary heat is on.

In both models, \dot{Q}_{vent} is an estimate of the heat gain or loss due to the outdoor air brought into the zone when the compressor is off:

$$\dot{Q}_{vent} = \dot{m}_{air} C_p (T_{mixed} - T_{zone}) \quad (3.12)$$

- where
- \dot{m}_{air} = the nominal air flow provided to the zone,
 - C_p = the heat capacity of air,
 - T_{supply} = the nominal supply air temperature provided by the unit and,
 - T_{zone} = the zone dry-bulb temperature.

Mixed air temperature, T_{mixed} , if not measured directly, requires either a measurement of outdoor air flow rate and temperature, or outdoor air fraction and temperature, in order to calculate T_{mixed} . In some cases, only supply, return, and outdoor air temperatures are available, in which case a simple mixing equation can be used to estimate the outdoor air fraction.

3.4.3.3 Inverse Model

A comprehensive literature review was performed to find a model that best matches the data available in this project. It was found that there are several equation forms used in the literature to model heat pumps. However, most of them are limited to compressor performance and/or rely on the measurements of suction and discharge pressure, or evaporation and condensation temperatures, which are associated with the refrigerant properties. Stoecker and Jones (1982) presented an empirical bi-cubic curve-fit model of compressor power consumption as a function of refrigerant condensing temperature and evaporating temperature. Duprez et al. (2007, 2010) used evaporation temperature, condensation temperature, saturated vapor pressure and compressor inlet temperature to predict the mass flow rate, compressor power use, and heat flow rate. In Air-Conditioning, Heating, Refrigeration Institute (AHRI) Air-Conditioning (2004), the compressor power is expressed using a polynomial with variables of suction and discharge dew point temperature. Kim and Braun (2012) and Shao et al. (2004) presented similar models by integrating the model given in ANSI/AHRI Standard 540 with an additional coefficient, which was associated with the variable operation frequency and rated frequency. Aprea and Renno (2009) applied the same approach to predict cooling capacity for variable-speed reciprocating compressors. These models are not complicated but evaporation temperature and condensation pressure are required as inputs. Cheung and Braun (2010) proposed a power model that was associated with the indoor fan speed and system load. In their model, the system load was linearly related to the outdoor air temperature. However, the power measured in this study was the sum of compressor and outdoor fan power and they say further investigation is needed to separate the compressor power from this sum. Gayeski et al. (2010) identified empirical, regression-based curve-fit models of heat pump power consumption, cooling capacity, and the COP that represent heat pump performance over the full range of test conditions. The curve-fit models developed by them are four-variable cubic polynomial functions as a function of outdoor air temperature, zone air temperature, compressor speed, and condenser fan speed. Li. et al. (2015) developed a semi-theoretical compressor power model for single-stage RTUs equipped with a variable-speed compressor and variable-speed indoor fan, based on the theoretical analysis and experimental studies. According to them, under normal conditions, the compressor power is correlated to the outside air temperature and compressor speed with a relative error of $\pm 8\%$. Li. et al. (2015) also found that compressor speed and

condenser inlet air temperature had a strong impact on the compressor power, while evaporator inlet air temperature and humidity had a weak impact on compressor power. Perers et al. (2012) developed a very simple model to describe the heat pump. In their model, the electric power input and thermal power output were correlated with inlet temperature to the heat pump evaporator and inlet temperature to the condenser.

The literature mentioned represents a sample of the models reviewed. Considering data commonly available from the manufacturer and/or through building automation and monitoring systems, we propose a set of second-order polynomials of the form shown in Equations (3.13) through (3.16). These models are based on the model developed by Gayeski et al. (2010) with two independent variables of indoor and outdoor air temperatures. This equation form is used to estimate heating and cooling power consumption and capacities.

$$\dot{Q}_{cool} = \begin{cases} a_1 T_{oa}^2 + a_2 T_{oa} + a_3 T_{za}^2 + a_4 T_{za} + a_5 T_{oa} T_{za} + a_6, & \text{if } C=1 \\ \dot{Q}_{vent} & \text{if } C=0 \end{cases} \quad (3.13)$$

$$P_{cool,elec} = \begin{cases} b_1 T_{oa}^2 + b_2 T_{oa} + b_3 T_{za}^2 + b_4 T_{za} + b_5 T_{oa} T_{za} + b_6, & \text{if } C=1 \\ P_{fan} & \text{if } C=0 \end{cases} \quad (3.14)$$

The equivalent inverse models for heating operation are given by:

$$\dot{Q}_{heat} = \begin{cases} c_1 T_{oa}^2 + c_2 T_{oa} + c_3 T_{za}^2 + c_4 T_{za} + c_5 T_{oa} T_{za} + c_6, & \text{if } C=1 \text{ and } A = 0 \\ c_1 T_{oa}^2 + c_2 T_{oa} + c_3 T_{za}^2 + c_4 T_{za} + c_5 T_{oa} T_{za} + c_6 + \dot{Q}_{aux}, & \text{if } C=1 \text{ and } A = 1 \\ \dot{Q}_{vent} & \text{if } C=0 \text{ and } A = 0 \end{cases} \quad (3.15)$$

$$P_{heat,elec} = \begin{cases} d_1 T_{oa}^2 + d_2 T_{oa} + d_3 T_{za}^2 + d_4 T_{za} + d_5 T_{oa} T_{za} + d_6, & \text{if } C=1 \text{ and } A = 0 \\ d_1 T_{oa}^2 + d_2 T_{oa} + d_3 T_{za}^2 + d_4 T_{za} + d_5 T_{oa} T_{za} + d_6 + P_{aux}, & \text{if } C=1 \text{ and } A = 1 \\ P_{fan} & \text{if } C=0 \text{ and } A = 0 \end{cases} \quad (3.16)$$

where T_{oa} is the outdoor air temperature, T_{za} is the zone temperature, and $a_1 \dots a_6$, $b_1 \dots b_6$, $c_1 \dots c_6$, and $d_1 \dots d_6$ are coefficients found through linear least-squares regression. \dot{Q}_{vent} , \dot{Q}_{aux} , P_{aux} , and P_{fan} are defined as described in the nominal performance model.

3.4.4 Parameter Identification Using Catalog Performance

Regression coefficients $a_1^i \dots a_6^i$, $b_1^i \dots b_6^i$, $c_1^i \dots c_6^i$ and $d_1^i \dots d_6^i$ in the preceding inverse models were estimated by fitting the model to the performance data extracted from manufacturer's catalogs. The fitted line is estimated by least-squares regression. The following sections include a summary of coefficients found and R^2 (i.e., the coefficient of determination) results for cooling and heating capacities and power consumption using manufacturer's performance data for BUILDING6. Parameter identification results for BUILDING10 are summarized in Appendix B (Section B.4, BUILDING10 Model Coefficients).

3.4.4.1 BUILDING6 Cooling Capacity Coefficients

The cooling capacity coefficients and R^2 values found for different systems in BUILDING6 are tabulated in Table 3.8. A sample of fitted line plots is shown in Figure 3.7. R^2 values suggest a strong correlation between the estimated and actual data indicating the model captures the behavior of heat pumps well.

Table 3.8. BUILDING6 cooling capacity coefficients and R^2 s.

System	a_1	a_2	a_3	a_4	a_5	a_6	R^2
Model A	-0.0043	-0.0109	0.0089	-0.3891	0.0036	29.1383	1.00
Model B	-0.0193	1.0339	0.0257	-1.5766	0.0084	20.5929	0.89
Model C	-0.0044	-0.3941	0.0353	-2.3167	0.0138	66.2300	0.99
Model D	-0.0017	-0.1956	0.0200	-1.0647	0.0032	38.7117	0.99

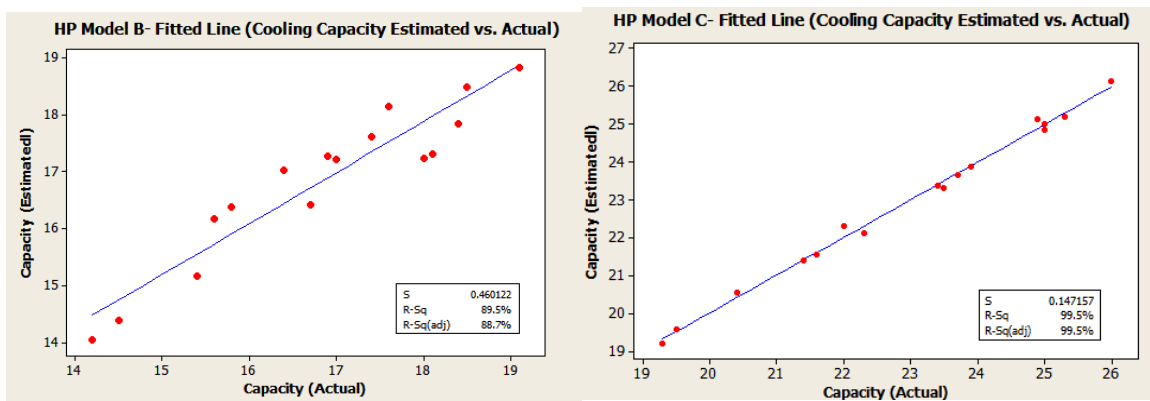


Figure 3.7. Fitted line plots showing estimated and actual cooling capacity for Models B and C.

3.4.4.2 BUILDING6 Cooling Power Consumption Coefficients

The manufacturer does not provide cooling power data for these models. In this case, we simply model the cooling performance of the unit using nominal characteristics.

3.4.4.3 BUILDING6 Heating Capacity Coefficients

Heating capacity coefficients are shown in Table 3.9 and a sample of fitted line plots are illustrated in Figure 3.8 showing estimated vs. actual heating capacity. Here, too, R^2 values suggest a very good model fit.

Table 3.9. BUILDING6 heating capacity coefficients and associated R^2 s.

System	c_1	c_2	c_3	c_4	c_5	c_6	R^2
Model A	0.0061	0.5543	-0.0002	-0.1014	-0.0027	17.0765	0.99
Model B	0.0036	0.4488	0.0000	-0.0627	-0.0024	12.9197	0.99
Model C	0.0034	0.6578	-0.0005	-0.2001	-0.0019	20.0594	0.99
Model D	0.0024	0.4524	0.0002	-0.1081	-0.0022	15.6753	0.99

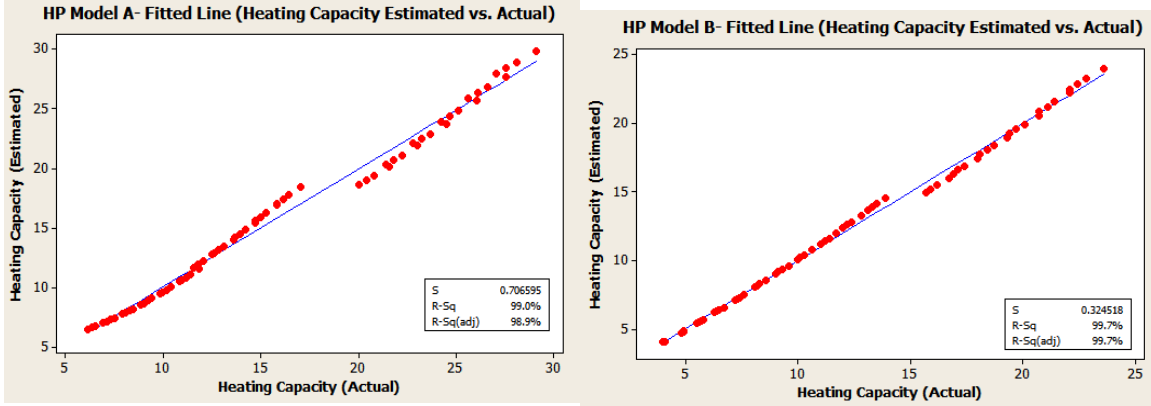


Figure 3.8. Fitted line plots showing estimated and actual heating capacity for Models A and B.

3.4.4.4 BUILDING6 Heating Power Consumption Coefficients

Heating electric power consumption coefficients are tabulated in Table 3.10 and a sample of fitted line plots are shown in Figure 3.9.

Table 3.10. BUILDING6 heating electric power consumption coefficients and associated R^2 s.

System	d_1	d_2	d_3	d_4	d_5	d_6	R^2
Model A	0.0005	0.0197	0.0010	0.0365	0.0005	4.2728	0.98
Model B	0.0004	0.0326	0.0012	0.0282	0.0002	3.2502	0.99
Model C	0.0005	0.0571	-0.0004	0.0352	0.0021	5.1339	0.99
Model D	0.0003	0.0235	0.0011	0.0457	0.0003	3.2271	0.99

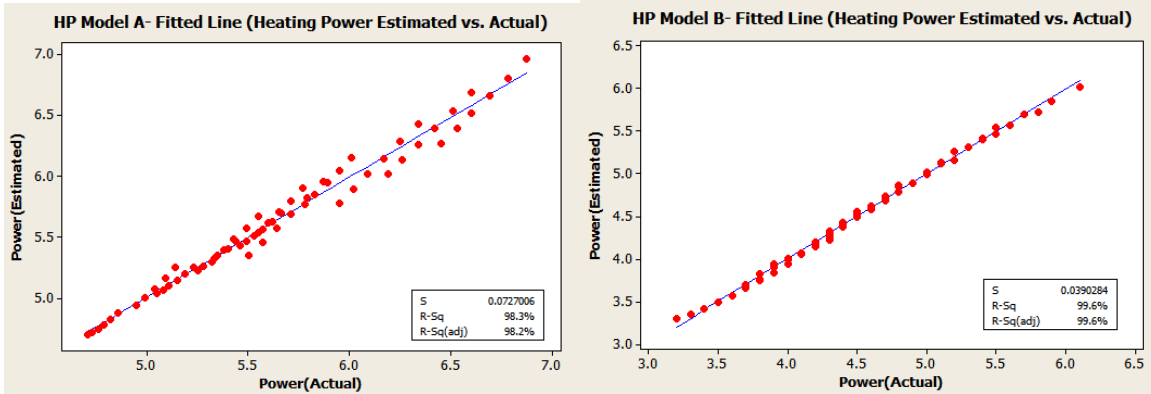


Figure 3.9. Fitted line plots showing estimated and actual heating electric power for Models A and B.

4.0 Market Simulation and Cross Validation

This section describes the simulation of transactive market controls using a detailed building model. Simulations are identical to those previously reported by Hao et al. (2016). Here, simulations are performed using a new co-simulation capability developed for this project. In this section, we summarize the building model used for the study, the co-simulation capabilities developed for this work, and results of simulations. We conclude by cross-validating these results with the previous work.

4.1 Building Model

The building model used for validation, and simulated in EnergyPlus (EnergyPlus 2016), is a representation of a newly constructed, 24,000 ft², small office building located on the PNNL campus. It is the same as that previously investigated by Hao et al. (2016). Only a portion of the building, representing approximately 50% of the conditioned floor area, and 90% of its occupants, is considered in this study. Heating and cooling are delivered by a VAV system served by a single built-up AHU located on the roof. The AHU contains supply and return fans, hot water and chilled water coils, and an outdoor air economizer. The AHU serves 17 VAV boxes, each with a hot water reheat coil. A boiler supplies hot water to the VAV boxes and AHU coil. Chilled water is supplied by chilled water from a central chiller plant.

The EnergyPlus model was developed from as-built engineering drawings and equipment specifications; the latter are summarized in Table 4.1. Operation schedules and set points were obtained from the building manager, and performance data collected by the BAS were used to calibrate the model. Fan parameters identified for the simple control-oriented models described previously were used to specify the AHU fan curve. Lighting and internal loads were estimated from a building walk-through conducted immediately after the building became occupied in the spring of 2015. A diagram showing the portion of the building simulated in this study is shown in Figure 4.1.

Table 4.1. Summary of HVAC equipment characteristics and operating parameters.

Equipment	Characteristics/Parameters
HVAC System	Multi-zone VAV
Design Supply Airflow	7.400 m ³ /s
Design Return Airflow	6.938 m ³ /s
Design Outdoor Airflow	0.4625 m ³ /s
Supply Fan Motor Nominal Power	14.9 kW
Zone Terminal Boxes	17
Airside Economizer	Fixed dry-bulb
Nominal Cooling Capacity	153 kW
Nominal Heating Capacity	84 kW
Occupied Hours	06:00-18:00 Mon-Fri
Occupied Cooling Temperature Set Point	22.8°C
Occupied Heating Temperature Set Point	20.0°C

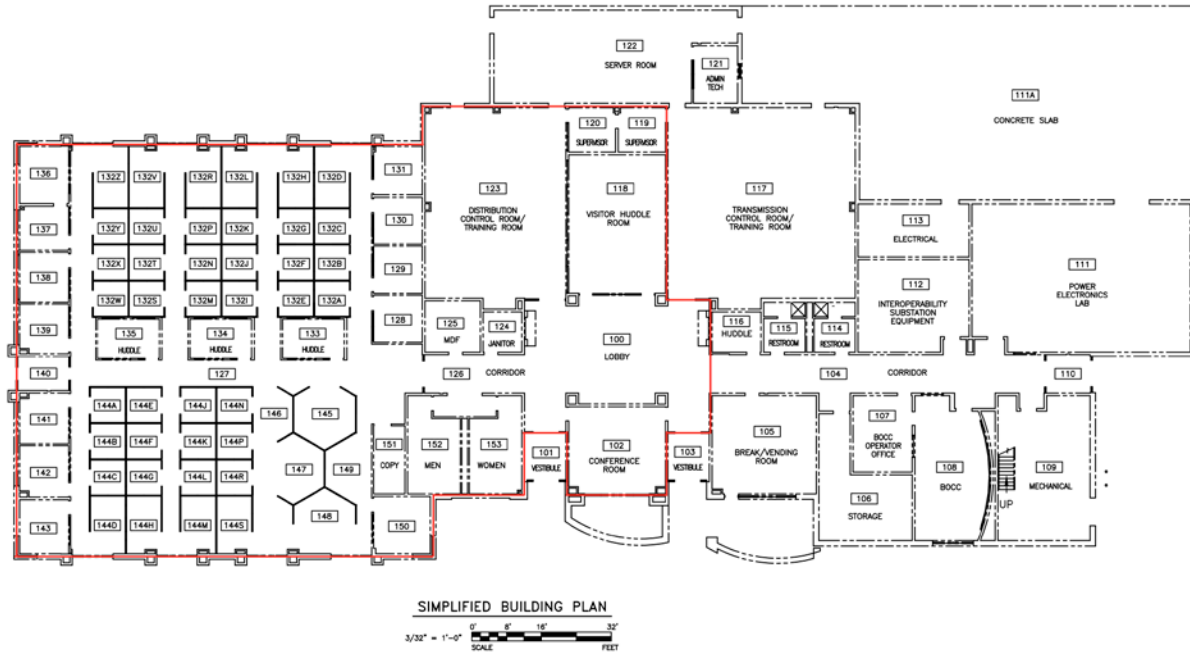


Figure 4.1. Floor plan of the modeled building. Red outline indicates the area to which transactive market control was applied.

4.2 VOLTRON Co-Simulation

Previous simulation studies were performed using the Building Control Virtual Test Bed co-simulation software developed by the Lawrence Berkeley National Laboratory (LBNL 2016). This environment allowed for fast development and testing of control algorithms against our building model in simulation. Results produced by Hao et al. (2016) are a reference for validating the algorithms after their having been migrated to VOLTRON agents. This new co-simulation capability, built upon the VOLTRON platform, allows for testing and validation of the developed algorithms against our building simulation model within the target deployment environment.

4.2.1 EnergyPlus Agent

To allow VOLTRON agents to be co-simulated with EnergyPlus, a VOLTRON agent was developed that manages communication between the VOLTRON message bus and the EnergyPlus simulation engine. This agent manages the initialization and execution of the building model simulation and masks the details from the other agents running on the platform. In essence, the EnergyPlus agent appears, from the perspective of other agents running on the platform, as a physical building. This allows VOLTRON applications to be developed and tested without requiring changes to agents for deployment. Technical details of the agent are documented in the companion user guide (Corbin 2016).

4.2.2 Market and HVAC Agents

In addition to the EnergyPlus agent that ties the transactive market simulation to the EnergyPlus building simulation, agents were developed to represent the markets, thermal zones, and HVAC equipment in the transactive market system. The transactive system, depicted in Figure 2.1, consists of 17 VAV agents, a single combined AHU/Chiller agent, a meter agent, two market agents, and a director agent.

4.2.2.1 VAV Agent

This agent represents the VAV box and the thermal zone it serves. It is responsible for estimating the cooling demand curve for the zone using the zone model described in Section 3.1. The agent submits its cooling demand curve to the air market as a buy bid at each market interval, and responds to the cleared air price by resetting its thermostat.

4.2.2.2 AHU/Chiller Agent

This agent represents the VAV system air handler and chiller plant. It is responsible for supplying cool air to the VAV boxes. It estimates the electricity required to power the fan and chiller—using models described in Sections 3.2 and 3.3—in order to serve the cooling load bid by the VAV agents. It submits the resulting electricity demand curve as a buy bid to the electricity market.

4.2.2.3 Meter Agent

The meter agent simply sells electricity to the electricity market. It is configurable to submit sell bids for electricity based on a fixed price, dynamic price, or demand limit.

4.2.2.4 Market Agent

The market agent collects bids from buyers and sellers, determines the clearing price and quantity of its commodity, and broadcasts the cleared values to its buyers and sellers. In this transactive system, there are two market agents, each responsible for a separate commodity, either cool air or electricity.

4.2.2.5 Director Agent

The director synchronizes the bidding and clearing of the markets. It simply broadcasts a message to the markets at each market interval to initiate the bidding and clearing process.

4.3 Case Studies

Three case studies explored in the previous work were used to validate the new market application built on the VOLTTRON platform (Hao et al. 2016). These cases, described below, represent energy conservation, demand-limiting, and demand flexibility applications of transactive control.

All simulation settings and assumptions from Hao et al. (2016) have been adopted in this study to allow a direct comparison of results from the VOLTTRON transactive market implementation to those reported in the previous work. Parameters for the simulation are summarized in Table 4.2 and Table 4.3. The clearing interval is set to 5 minutes, consistent with the simulation time-step.

Table 4.2. Simulation parameters used in transactive market simulations.^{(a)1}

Parameter	Value
Upper Cooling Temperature Set Point	23°C
Lower Cooling Temperature Set Point	19°C
Baseline Case Cooling Temperature Set Point	21°C
Upper Price Limit	\$100/MW
Lower Price Limit	\$10/MW
Occupied Hours	06:00-18:00, Mon-Fri
Simulation Timeframe	August 20-24
Weather File	Pasco, WA (TMY3)
Simulation Time-Step	5 minutes

(a) Upper and lower price limits vary in the dynamic price case according to the method described in Section 4.4.3

Table 4.3. Minimum and maximum zone airflow configuration for VAV agents in simulation experiments.

VAV Terminal Box	Minimum Airflow (m ³ /s)	Maximum Airflow (m ³ /s)
VAV100	0.1363	0.4544
VAV102	0.1471	0.4904
VAV118	0.0557	0.1857
VAV119	0.01031	0.034388
VAV120	0.01056	0.03520
VAV121	0.009593	0.03198
VAV123A	0.0680	0.2268
VAV123B	0.06998	0.2333
VAV127A	0.1936	0.6452
VAV127B	0.3360	1.120
VAV129	0.03582	0.1194
VAV131	0.01461	0.04870
VAV133	0.03208	0.1069
VAV136	0.04000	0.1333
VAV142	0.2222	0.7405
VAV143	0.04204	0.1401
VAV150	0.04552	0.1517

4.3.1 Fixed Price

This case tests the behavior of the market when the electricity price is fixed, and it provides insight into the transactive market operation independent of changes in price from one clearing interval to the next. Recall that the ElectricityMeter is responsible for bidding a supply curve into the electricity market. In this simulation case, this curve is defined as having a fixed price for all possible demand values. When this curve is intersected by the aggregate electricity demand curve, the demand is fixed; notionally, this is the demand that must be divided among the zones.

Results from our investigation in the previous work suggested the flat price case was an effective mechanism for encouraging electric energy reduction through conservation and improved balancing of the cooling load. Specifically, the higher the electricity price, the more the market control tends to reduce and flatten electric demand over the course of each day.

¹ Upper and lower price limits vary in the dynamic price case according to the method described in Section 4.3.3

Simulations for fixed prices of \$55, \$60, \$65, and \$70/MW were repeated using the VOLTRON-based market simulation. Figure 4.2 shows the result of one such simulation in which a fixed price of \$65/MW was used.




Figure 4.2. Comparison of \$65/MW fixed price simulation to baseline.

As noted in the previous work, cooling set points (not shown) tend to be depressed in the morning hours when cooling loads are low, then gradually rise throughout the day, shifting cooling load to morning hours and reducing peak demand during the afternoon. We also observe that zone set point trajectories differ from one another, illustrating how the transactive controls balance the needs of the zones throughout the day by allowing the zones to negotiate the amount of cooling they need through the market.

4.3.2 Demand Limit

In this simulation case we impose a demand limit on the total (fan plus chiller) electric demand at each market-clearing interval. Here, the curve bid by the ElectricityMeter is defined as having a fixed demand value for all price values. When this fixed demand curve is intersected by the aggregate electricity demand curve, a clearing price is established. If the aggregate electricity demand curve does not intersect the supply curve, the electricity and air markets fail to clear. A failure of the markets to clear signals the VAV boxes to set their cooling set points to their nominal baseline value. Effectively, this prevents the zones from over-cooling and wasting electricity, and thereby ensures that the markets only clear if the demand limit has been exceeded.

Figure 4.3 shows the results from the demand limit simulation. We observe that peak demand is mostly maintained below the imposed demand limit. Small excursions above the demand limit can be seen in several instances. Further investigation reveals that these excursions are due to model mismatch. That is, the demand cleared in the market, being a function of the demand calculated by our simplified models, does not exactly match that simulated by EnergyPlus. The degree to which the aggregate HVAC electric demand calculated by the simplified models agrees with the EnergyPlus simulation is encouraging; it suggests that (in the context of simulation, at least) the simplified models are adequate for the purpose of limiting HVAC demand.

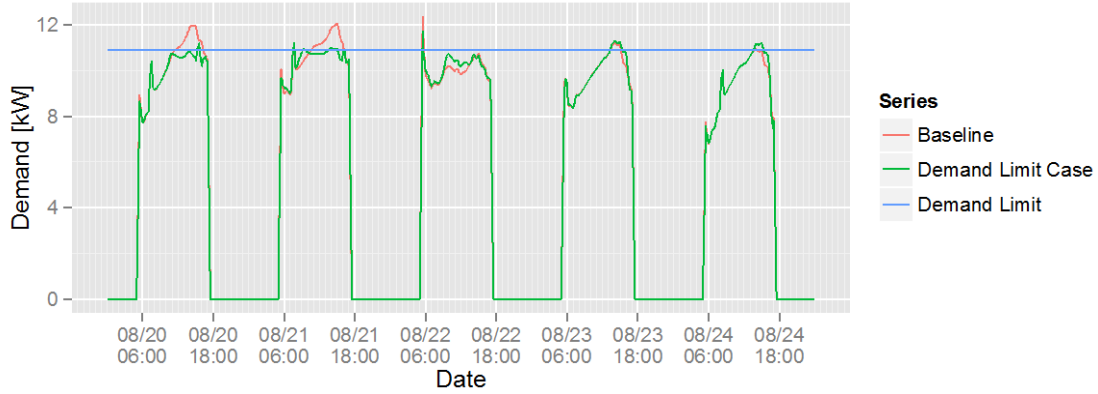


Figure 4.3. Comparison of demand limit simulation to baseline.

4.3.3 Dynamic Price

This simulation tests the ability of the market and building to react to price signals that change over time. In this case, the ElectricityMeter bids a supply curve similar to that in the fixed price case, except that the price varies at each market-clearing interval. The source of the price is a time series of real-time prices obtained from the California Independent System Operator (CAISO 2016).

The general market bidding and clearing process is the same as in previous cases. However, in contrast to the previous cases, upper and lower bid prices submitted by the VAVs are calculated as follows:

$$p_{high} = \bar{p}_{24} + \sigma_{24} \quad (4.1)$$

$$p_{low} = \bar{p}_{24} - \sigma_{24} \quad (4.2)$$

where p_{high} = the price associated with the maximum cooling demand calculated by the zone,
 p_{low} = the price associated with the minimum cooling demand calculated by the zone,
 \bar{p}_{24} = the average price over the preceding 24 hours, and
 σ_{24} = the standard deviation of price over the preceding 24 hours.

These points define the curve bid by zone agents into the air market. This bidding strategy is consistent with that implemented by Hao et al. (2016). The strategy allows the zone to adjust bids over time as prices fluctuate, and it prevents the situation in which zones are continually being controlled to upper set point limits, as would happen during a week of high prices.

Simulated electricity demand from the dynamic price case is compared to the baseline in Figure 4.4. High prices tend to drive thermostat set points up, and therefore cooling demand down, as the thermostats react. Low prices have the opposite effect. It is important to observe that dynamic prices can create new demand peaks. With the exception of the startup peaks in the morning, this can be seen in each daily demand profile. Also worth noting are the large increases in demand on days four and five. These excessive increases in demand, associated with aggressive cooling, are the result of the 24-hour averaging in the bidding logic. Essentially, electricity appears to be very inexpensive on these days compared to the previous days. This is especially true on the fifth day as the price drops to near zero.

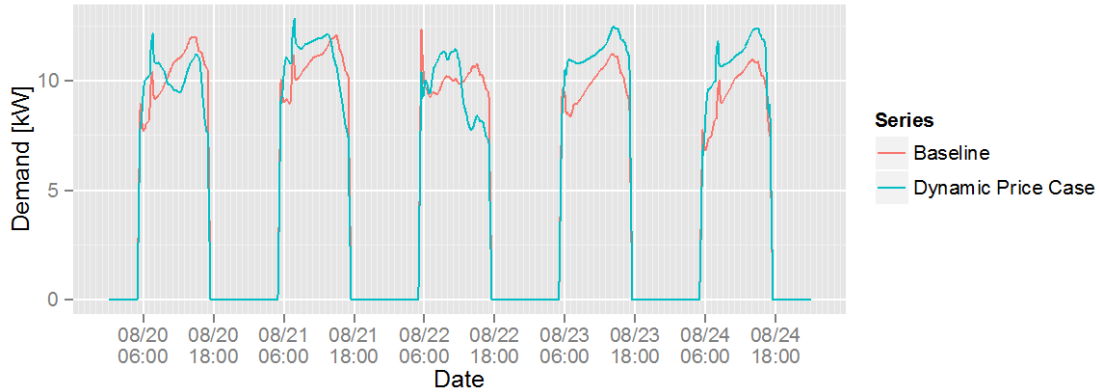


Figure 4.4. Comparison of dynamic price simulation to baseline.

These results highlight an important finding regarding dynamic price response: dynamic pricing may result in increased demand during peak demand times if pricing signals are not well aligned with the peak cooling loads.²

4.3.4 Summary and Discussion of Results

Table 4.4 summarizes the impact from the three incentive signals on three metrics of interest: building peak demand, total energy consumption, and zone temperature deviation. Reduction in peak load is calculated as a percentage of baseline peak load measured in the baseline for the same day. Total cooling energy consumption is defined as the percentage change from the baseline over the entire 5-day simulation. Finally, temperature deviation is measured by calculating the average temperature difference between the case and the baseline.

The fixed price case suggests that both energy savings and—perhaps more interestingly—peak load reduction may result. The fixed price tends to level demand compared to the baseline, shifting demand to the morning hours. This tends to precool the zones, while allowing a reduction in demand during the baseline afternoon peak. This is observed in the temperature profiles in Figure 4.2. This case provides some justification for market-based control without an external signal, because savings in both total electricity use and demand may be realized.

In the demand-limiting case, we see that the market is able to attenuate the peak demand. While demand cleared in the market remains at or below the demand limit, a few excursions exist due to model mismatch. We observe that the transactive market is able to reduce peak demand and save energy by raising zone temperatures by only a small amount on average. A number of factors determine the peak demand, which may be naturally much lower or higher on any given day. Therefore, the selection of an appropriate limit is important.³ Prediction of daily demand profiles is an area of ongoing work.

² Although not the focus of this report, the implications for large-scale deployments are important. When groups of buildings, e.g., in a campus deployment, are subjected to the same signal, peak demand may be increased, stressing campus infrastructure. In this example, the price is given, resulting in an open feedback loop; the feedback loop must be closed so that prices adjust in response to high demand bids.

³ When considering that demand charges are typically based on the peak 15 or 30 minutes in a given calendar month, avoidance of a few days or hours of peak demand are all that is necessary to avoid these charges. A simple heuristic for the demand limit would be: multiply the peak demand measured in the trailing month by a scaling factor less than one, e.g., 0.9.

Dynamic prices drive significant changes in both demand and temperature set points. Zone temperatures are generally reduced on average, and both energy consumption and peak demand increase during most days. Changes in demand are strongly tied to price. As noted previously, the behavior on days four and five is a side effect of the adaptive bidding strategy in which prices appear to be lower relative to the trailing average, thus causing the VAVs to lower set points and consume more cooling energy. This behavior has been seen in previous simulation studies involving price responsive residential thermostats (Widergren et al. 2014), and underscores the complex interactions between bidding strategy and system behavior.

Table 4.4. Results of transactive market simulation cases.

Case		Peak Load (%)			Energy (%)	Temperature (°C)
		Mean	Min	Max		
Fixed Price	\$55	5.6	-0.6	11.4	4.7	-0.07
	\$60	1.1	-4.2	6.9	1.0	0.30
	\$65	-4.2	-9.3	1.4	-3.1	0.61
	\$70	-8.5	-12.3	-5.1	-7.1	0.85
Demand Limit		-6.2	-6.3	-6.1	-1.5	0.17
Dynamic Price		5.8	-4.2	14.2	3.1	-0.03

Results from these simulations have been compared with those reported by Hao et al. (2016) and found to be identical for each of the 5-minute iterations of the building simulation and market. Values analyzed for differences include zone temperature, building HVAC electricity demand, market-clearing quantity, and price. This cross validation confirms that the VOLTTRON-based market performs identically to the previous version.

5.0 Physical Experiments

In this section we describe testing of transactive controls on a small commercial building located on the PNNL campus. The building used for these experiments is the same as that described in Section 4.1. We begin with an overview of the experiments performed, then summarize the deployment and integration process, follow with a description of the analysis methodology we have adopted, and then present the results of the experiments. Challenges encountered and next steps conclude this section.

5.1 Experiment Description

Physical experiments are modeled after the simulations described in Section 4.0. Experimentation began first with flat price and progressed to demand limiting experiments. A fully dynamic price experiment was not performed for reasons discussed in Section 5.2.3. A summary of the experiments is provided in Table 5.1.

Table 5.1. Summary of experiments performed on the physical testbed.

Date	Start Time	Stop Time	Experiment	Notes
1-Jun-2016	08:00	12:00	Flat Price	Dry run. Did not control building. Observed significant fan power oscillations. \$65/MW flat price.
8-Jun-2016	14:00	17:00	Flat Price	First test. Observed data latency. Flat price of \$65/MW flat price.
16-Jun-2016	07:00	17:00	Flat Price	Observed BAS truncating set points. \$75/MW flat price.
29-Jun-2016	07:00	17:00	Flat Price	Confirmed truncation was removed, but still observed data latency. \$75/MW flat price.
14-Jul-2016	07:00	17:00	Demand Limit	Lost connection to building at 11:30. 10 MW demand limit.
21-Jul-2016	07:00	17:00	Demand Limit	No data from AHU until 10:30. 10 MW demand limit.
28-Jul-2016	07:00	17:00	Demand Limit	Observed strange bidding behavior, traced back to zone model coefficients. 9.5 MW demand limit.
11-Aug-2016	07:00	17:00	Flat Price	Experiment terminated due to configuration error.
16-Aug-2016	05:30	NA	Flat Price	Experiment interrupted by operator overrides around noon. Experiment terminated. \$65/MW flat price.
17-Aug-2016	06:00	NA	Flat Price	Operator did not release overrides from previous day. Experiment terminated. \$65/MW flat price.
18-Aug-2016	06:00	16:00	Flat Price	High temperature complaints resulted in early termination. \$65/MW flat price.
25-Aug-2016	14:30	16:00	Demand Limit	Platform install problems. 10 MW demand limit.
31-Aug-2016	06:00	18:00	Demand Limit	10 MW demand limit.
1-Sep-2016	06:00	18:00	Flat Price	\$65/MW flat price.
2-Sep-2016	06:00	18:00	Demand Limit	6.5 MW demand limit.
8-Sep-2016	06:00	18:00	Flat Price	\$77.5/MW flat price.
13-Sep-2016	12:30	18:00	Flat Price	\$65/MW flat price.
14-Sep-2016	06:00	18:00	Flat price	\$65/MW flat price.

5.2 Deployment and Integration

The BAS of the physical testbed consists of Alerton (2016) field controllers networked through an Alerton Ascent Control Module (ACM). The system is managed via a Johnson Controls Metasys Network Automation Engine (NAE) using the Metasys graphical user interface (Johnson Controls 2016).

Software agents composing the market-based control system were deployed to a VOLTTRON instance running on an Intel NUC mini PC (Intel 2016). This device was physically located on the testbed premises and connected to a subnet of the physical building’s Internet Protocol (IP) network reserved for the BAS. The device was configured to monitor points (see Table 5.5 and Table 5.6 in Section 5.2.2) through the Metasys frontend over the BACnet (Building Automation and Control Network) IP. The device was later reconfigured to communicate directly with the field controllers through the ACM after initial testing revealed delays in updates from the field controllers. This issue is discussed in more detail in Section 5.2.3

Heating and cooling set points were originally controlled through the Metasys front end through a combination of seven points for each VAV: a single base set point, an occupied heating offset, an occupied cooling offset, a standby heating offset, a standby cooling offset, an unoccupied heating set point, and an unoccupied cooling set point. Our initial integration with Metasys allowed only the base set point and occupied offsets to be modified by VOLTTRON. To allow control during standby and unoccupied times, we subsequently configured two virtual points in the ACM representing the effective heating and cooling set points and modified the VOLTTRON configuration to write to these points instead. Similarly, the Metasys frontend was reconfigured to write to the effective set points.

Points in the ACM use a BACnet priority array. During experiments, VOLTTRON would write to the point at a priority of 9. This allowed the building operator to use a standard override (priority 8) to selectively override set points during experiments. Although planned, a global kill-switch to revert all set point values to their original state at the loss of heart beat from the VOLTTRON instance or operator command was not implemented.

5.2.1 Model Tuning

Fan, chiller, and zone thermal models were tuned using the methods described in Sections 3.1.2, 3.2.2, and 3.3.2. Training data were collected from the BAS over the month of July. Zone models were periodically re-tuned every month¹. The identified model parameters are listed in Table 5.2, Table 5.3 and Table 5.4.

Table 5.2. Chiller model parameter for physical experiments.

Parameter	Value
Coefficient of Performance	5.69

¹ Small variations in zone model parameters could be observed from one tuning to the next. While the impact of changing parameters is not the focus of this particular work, this observation suggests that ongoing automated tuning may be necessary to keep models up to date. At the very least, zone models may need to be re-tuned seasonally.

Table 5.3. Fan model parameters for physical experiments.

Coefficient	Value
a_0	5.071e-3
a_1	2.605e-4
a_2	8.062e-8
a_3	-2.841e-12
a_4	7.858e-1
a_5	4.382e-1

Table 5.4. Example of zone model parameters for a single zone used in physical experiments.

Model	Coefficient	Value
Temperature	c_0	1.77e-01
	c_1	9.88e-01
	c_2	5.50e-03
	c_3	1.87e-04
	c_4	-1.07e-04
Cooling Demand	x_0	-1.66e+01
	x_1	-5.16e+01
	x_2	-2.58e-01
	x_3	5.25e+01
	x_4	9.90e-01

5.2.2 Agent Configuration

Each agent is configured to receive measured real-time values from the BAS. These values set the current states of the models so that model calculations—and thus market bids—reflect current conditions. Points supplied to the AHUChiller and VAV agents are summarized in Table 5.5 and Table 5.6, respectively. Full details regarding agent configuration for all agents in the market are provided in the implementation guide (Corbin 2016).

Table 5.5. Input points configured for the AHUChiller agent.

Point	Model Parameter	Units
AHU Discharge Air Temperature	T_{supply}	°F
AHU Return Air Temperature	T_{return}	°F
AHU Mixed Air Temperature	T_{mixed}	°F
AHU Discharge Air Flow	\dot{m}_{air}	ft ³ /min
Duct Static Pressure	ρ	inH ₂ O

Table 5.6. Input points configured for each VAV agent.

Point	Model Parameter	Units
AHU Discharge Air Temperature	T_{supply}	°F
VAV Discharge Air Temperature	T_{sa}	°F
Zone Temperature	T_{za}	°F
Outdoor Temperature	T_{oa}	°F
VAV Discharge Air Flow	\dot{m}_{sa}	ft ³ /min
Zone Occupancy		[0]1
Standby Status		[0]1
AHU Occupancy		[0]1

Zone occupancy, standby status, and AHU occupancy points are binary flags that alter the behavior of the VAV agent. When zone occupancy is 1, the desired cooling set points used to calculate minimum and maximum desired cooling demand are constrained to a narrow range in order to minimize impacts on occupant comfort. An occupancy sensor in each zone determines whether the flag is 1 or 0. When standby status is 1 this range is widened. Standby mode is entered immediately after zone occupancy is no longer detected and the flag has a value of 1 for the following 15 minutes, at which time it changes to 0. When both zone occupancy and standby status flags are 0, a wider unoccupied set point range is used. The AHU occupancy flag indicates that the system is off, in which case the VAV agent does not bid. Temperature limits for the three occupancy states are summarized in Table 5.7.¹

Table 5.7. Temperature limits for VAV agents.

Parameter	Value
Occupied Cooling Temperature Set Point Upper Limit	23.89°C
Occupied Cooling Temperature Set Point Lower Limit	20.00°C
Standby Cooling Temperature Set Point Upper Limit	25.56°C
Standby Cooling Temperature Set Point Lower Limit	20.00°C
Unoccupied Cooling Temperature Set Point Upper Limit	26.67°C
Unoccupied Cooling Temperature Set Point Lower Limit	19.44°C

VAV agents control the heating and cooling set points in occupied, standby, and unoccupied modes strictly within these limits. Effective heating and cooling set points controlled by the VAV agents are listed in Table 5.8.

Table 5.8. Control points configured for each VAV agent.

Point	Units
Effective Cooling Set Point	°F
Effective Heating Set Point	°F

Recall that VAV agents bid their minimum and maximum desired cooling demand, corresponding to a maximum and minimum temperature deviation. Minimum and maximum cooling demand is bounded by the minimum and maximum airflow configured for a VAV; values for each VAV are listed in Table 5.9.

¹ These values are typical for the majority of the zones. In fact, some zones have slightly different limits based on occupant preference.

Table 5.9. Minimum and maximum zone airflow configuration for VAV agents in physical experiments.

VAV Terminal Box	Minimum Airflow (m ³ /s)	Maximum Airflow (m ³ /s)
VAV100	0.6359	1.386
VAV102	0.3469	0.6938
VAV118	0.1156	0.4625
VAV119	0.02891	0.05781
VAV120	0.02891	0.05781
VAV121	0.05781	0.1734
VAV123A	0.3469	1.388
VAV123B	0.6938	1.388
VAV127A	0.5781	1.156
VAV127B	0.5781	1.156
VAV129	0.06938	0.1388
VAV131	0.02313	0.04625
VAV133	0.1301	0.2602
VAV136	0.04336	0.1445
VAV142	0.1156	0.5781
VAV143	0.04336	0.1445
VAV150	0.07516	0.2023

5.2.3 Challenges and Resolution

Initial data analysis revealed delays in data collected by the VOLTTRON platform from the BAS. Although points in the BAS were polled every minute, values would not change for up to 15 minutes. Interestingly, different points on the same device, e.g., discharge air flow and discharge air temperature at the VAV, would be delayed by different amounts of time. We confirmed that change-of-value and update interval logic was not the cause; points would stay stagnant for several minutes, then jump by several degrees (or hundreds of cubic feet per minute), remain stagnant for a different interval, then change by fractions of the previous change. A test was performed in which VOLTTRON was configured to poll the field controllers directly. In this test, the previously observed delays were not present, indicating that the interaction between the ACM and NAE were likely the cause. Further root cause analysis was not performed because of time and budget constraints.

The pre-existing BAS set point logic, which used seven points for each VAV, severely limited initial testing. This logic used a single set point to establish a baseline, then used offsets depending on occupancy status. For example, a base set point of 72°F was common for many of the VAVs. To this, a 2°F offset was added to establish the cooling set point of 74°F, and a separate 2°F offset was subtracted to establish the heating set point of 70°F. During standby, i.e., after 15 minutes of no occupancy sensor readings during normal operating hours, an additional 4°F were added to the cooling set point to obtain 78°F, and 4°F was subtracted from the heating set point to obtain 66°F. During times when the building was scheduled to be unoccupied, i.e., between 18:00 and 6:00, a separate set of heating and cooling set points were used, e.g., 65°F and 80°F. All of this logic existed in the BAS. To greatly simplify the logic of the market-based controllers, the BAS was modified to write the result of this logic, i.e., the effective heating and cooling set points, to a virtual point, which was then used by the VAV for control. This allowed us to then override these effective set points regardless of operating mode, eliminating the need to replicate the offset logic in the market-based controllers.

Unfortunately, the first implementation of the new BAS effective set points logic truncated set point values. This limited our control to whole-value increments. For example, if a VAV reported a zone temperature of 72.5°F, and our VAV agent calculated a new cooling set point of 72.9°F, the BAS would truncate the effective cooling set point to 72°F, resulting in additional cooling, instead of the intended

decrease in cooling. This issue was resolved, but not until after analyzing data from a previous experiment.

Once measurement delays were resolved, we observed data measurement problems previously masked by the delays. Several VAV sensors exhibited sporadic and erroneous readings. For example, one VAV (Zone 142) reported airflows that would oscillate every minute between zero and a value in the normal range of 100 to 300 CFM for the VAV. Another VAV (Zone 133) reported zone temperature oscillations of 2–4°F every minute. These errors had not been corrected by the time of testing, which required that we remove these zones from future experiments.

An analysis of zone minimum airflows indicated that many zones were overcooled prior to our control experiments. Conversations with building operators supported the finding. Operators mentioned that the boiler and hot water reheat system were active during peak cooling season in the previous summer. Although some adjustments had been made by the building operator to minimize reheat (by disabling the boiler), minimum airflow set points still remained at 50% of the cooling design maximum for many zones, leading to overcooled zones at minimum airflow. Although it is difficult for us to quantify the impact of these settings on our experiments, we believe that the high minimum airflow set points severely restricted subsequent tests.

Significant oscillations in AHU fan power were observed during initial tests. The largest, on the order of 2 kW or about 50% of peak, would consistently appear at regular 45-minute intervals. Further investigation revealed that several zones were operating at 100% of their configured cooling airflow values, triggering static pressure reset logic. This logic would gradually increase static pressure up to a pre-defined maximum, or until no more than one VAV reported discharge airflow of 100% of the configured maximum. At this point, the logic would gradually decrease static pressure down to a pre-defined minimum until no more than one VAV reported airflow at the configured minimum. Only later did we discover that the building operator had configured one zone's minimum airflow to be the same as its maximum, forcing the building to constantly cycle between minimum and maximum static pressure set points. The VAV in question was reconfigured to have a lower minimum air flow set point, but oscillations were not eliminated entirely, because this condition does occur naturally. Further retuning is required to eliminate the oscillations entirely.

Finally, we note that well-documented software release and change control processes are critical to ensuring successful testing. Several experiments were negatively affected or canceled because of updates to the underlying platform. Problems introduced by these changes included data missing from the central data store, which prevented analysis of experiment results; loss of measurements from devices during experiments, which resulted in poor or no control; changes to core platform functionality that resulted run-time errors and prevented agent execution during experiments.

5.3 Analysis Method

Measuring the impact of our market-based controls during physical experiments is fundamentally more difficult than doing so in simulation. This is because a number of factors—including weather conditions and occupant behavior that change from moment to moment and are beyond our ability to control—determine the energy use of a building. Researchers in this field have used a number of techniques to address this problem. In essence, these techniques attempt to predict the energy consumption of a building under a given set of conditions using a model trained with historical data. This is analogous to predicting the energy consumption for days identical to those on which an experiment was performed, as if it had not been performed. In this way, differences between days can be controlled for, allowing differences to be correctly attributed to the experiment.

In this section we describe two simple approaches for evaluating the performance of the market-based controls during the physical tests. These are:

- **Method 1** – a simple, naïve approach that compares experimental results from performance measured during a similar “baseline” day.
- **Method 2** – a statistical approach that attempts to model building performance based on a set of relevant predictors selected by their association with the underlying physical processes that determine building energy use.

Regardless of which method we apply, the first step in our analysis requires us to collect relevant data from the HVAC system and pre-process the data to remove erroneous reads. Data from the BAS are collected continuously at 1-minute intervals and are available from a historical database. Records date to early 2016 when VOLTTRON data trending was installed. Data exist for both experiment and non-experiment days.

The collected data include both system-level data (such as AHU supply air flow rate, supply and return air temperature, outdoor air temperature, etc.) and zone-level data (such as supply and return air temperature of each zone). For the proposed method, only system-level data are used, because the objective is to evaluate the whole HVAC system, instead of the interactions between zones. Specifically, the data points shown in Table 5.10 are used in this analysis.

Table 5.10. Data points required for analysis of physical tests.

Point	Analysis Method
Supply fan power	Method 1 & 2
Chilled water flow rate	Method 1 & 2
Cooling coil entering water temperature	Method 1 & 2
Cooling coil exiting water temperature	Method 1 & 2
Outdoor air temperature	Method 1 & 2
Mixed air temperature	Method 1 & 2
Return air temperature	Method 2
Supply air flow rate	Method 2
Duct static pressure	Method 2
Outdoor air flow rate	Method 2

5.3.1 Method 1

A simple method for measuring performance is to directly compare results from an experiment day to a non-experiment, i.e., baseline, day. Because building operation is a function of weather (among other factors), care must be taken when selecting the baseline. Specifically, the baseline should match the experiment day in temperature profile, solar insolation, humidity, and wind. Additionally, occupancy, operating schedules, and internal loads should match as closely as possible to control for their effects on building performance. Finding a baseline day that matches an experiment day can be challenging, and sometimes impossible, especially if all relevant factors are not readily available or measured.

As an example, we have selected the June 29 experiment to illustrate the method. This day is well matched in terms of outdoor air temperature to the previous day, June 28. This can be seen in the lower panel of Figure 5.1. A comparison of solar insolation and wind is not possible due to a lack of site-measured data. Operation schedules are similar for these days. A survey of occupancy was not performed;

we assume that the occupants and their behavior do not vary appreciably between these 2 days. A similar assumption is made regarding internal gains.

Of interest to our analysis are the fan and chiller electric power. Fan power can be measured directly from the BAS. To obtain the chiller power estimate, we first calculate cooling coil load from the measured cooling coil water flow rate, \dot{m}_{water} , the heat capacity of water, C_p , and the cooling coil entering and exiting water temperatures, $T_{water,in}$ and $T_{water,out}$, respectively:

$$\dot{Q}_{cool} = \dot{m}_{water} C_p (T_{water,in} - T_{water,out}) \quad (5.1)$$

From the cooling coil load, we can simply estimate chiller power by dividing by the nominal chiller COP, or 5.69, as found during parameter identification in Section 3.3.2.

In general, the experiment on June 29 results in a flatter chiller and fan power profile. When compared to the June 28 series, the June 29 series shows increased demand in the morning hours, and decreased demand in the afternoon hours. This trend is consistent with the simulation results, even if the absolute shape is very different.

Of particular note are oscillations observed in both curves starting at approximately 09:30. These oscillations were found to be the result of static pressure reset logic in the BAS. This logic is triggered when any VAV terminal box damper is 100% open (resulting in a static pressure set point increase), or when none of the dampers are at their minimum (resulting in a static pressure set point decrease). While the market-based control sees and responds to these changes in airflow and static pressure, it has no direct control over the latter. Specific logic to avoid this condition is an interesting future area of work.

Additionally, a dip in chiller power on June 29 can be seen at approximately 14:00, which is followed by a significant spike. Further analysis reveals that this is the result of a rapid change in chilled water supply temperature to the AHU cooling coil. Unfortunately, the cause for this dip and spike is unknown and we have no way to determine whether this was a side effect of our experiment.

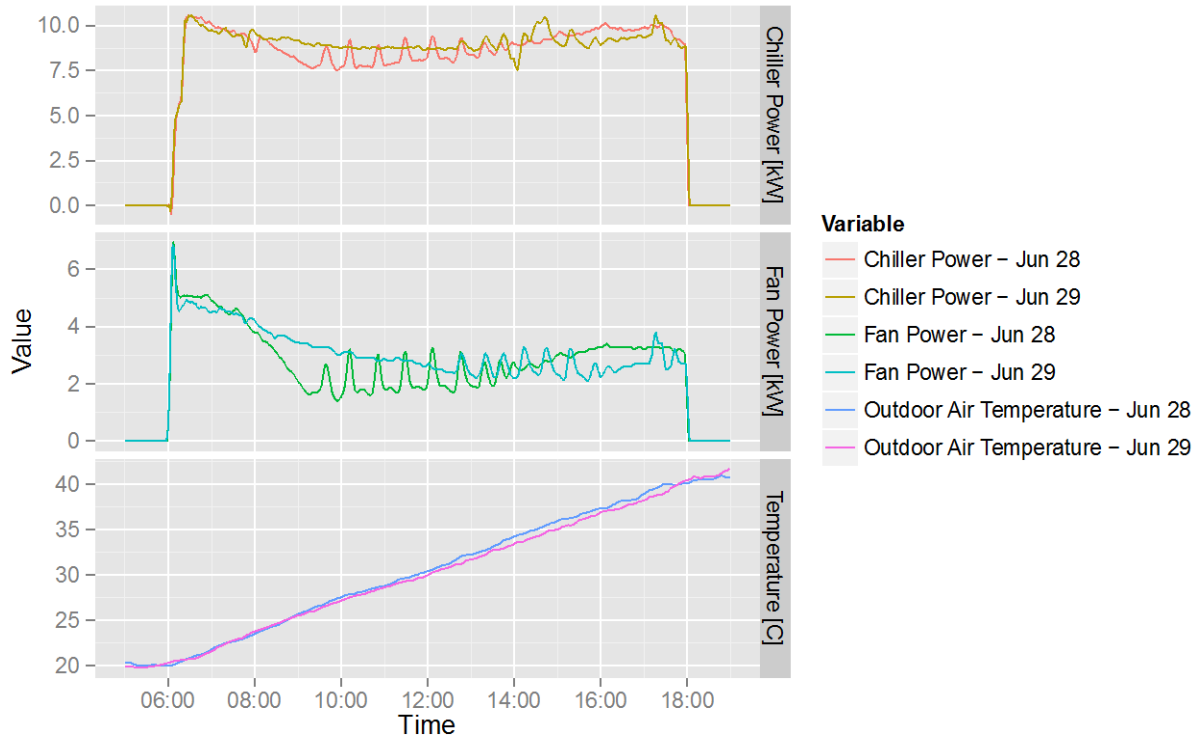


Figure 5.1. Direct comparison of chiller power, fan power, and outdoor air temperature for similar days.

To quantify the performance of the market-based control, we calculate two metrics from the measured fan power, calculated chiller power, and their sum (total). These metrics are shown in Table 5.11. In this table, “energy” is defined as the difference in total electric energy consumed between the experiment and baseline, and “demand” is the difference in peak demand measured between the experiment and baseline, where peak demand is defined as the maximum demand in a 30-minute rolling average of demand. The definition of peak demand is consistent with that defined by PNNL electric tariffs. Percentage values are calculated relative to the baseline day.

Table 5.11. Difference in energy and demand between June 28 and June 29.

Experiment Date	Series	Energy (kWh)	Energy (%)	Demand (kW)	Demand (%)
June 29	Fan	0.29	0.18%	-0.09	-1.90%
	Chiller	0.25	0.05%	0.07	0.70%
	Total	0.54	0.08%	-0.25	-1.76%

Metrics indicate a net increase in electricity consumption and a higher chiller peak demand during the experiment. Increased total consumption is insignificant at 0.09% of the baseline. Regarding demand, this increase can be attributed to the chiller spike that occurs at 14:30, and because there is no way to tell if this is a result of the experiment, it is hard to determine if this effect is real or simply coincidental. We note that at this time, a significant change in chilled water supply temperature (to the building) can be observed, suggesting that other causes may be implicated. Recall that the building is served by a central plant, which serves several other buildings, each significantly larger than the one studied here. This highlights a shortcoming of this analysis method: chiller and fan demand can be influenced by factors outside of our control, resulting in the observed increase in chiller power. In Method 2, these effects are, in theory, better controlled.

5.3.2 Method 2

In this method, we apply a simple machine-learning technique to predict building energy use in place of a comparable baseline. Because we cannot simply “replay” an experiment day as we can in simulation, this method allows us to predict building energy use as if the experiment had not been performed. In other words, this method allows us to predict the *counterfactual*.

The advantage of this method is that we do not need to carefully select a comparison day; the statistical model automatically accounts for changes in outdoor temperature, for example. It also controls for changes in performance outside of our control, i.e., exogenous effects that may or may not be measurable. This allows us to isolate the impact of our market-based system on building performance regardless of other effects.

Of interest to our analysis is the impact of our experiments on fan and chiller power. Each requires a separate model to estimate. To estimate fan power, P_{fan} , we construct a linear function in the following form:

$$P_{fan} = a_1 T_{supply} + a_2 T_{mixed} + a_3 E T_{mixed} + a_4 T_{oa} + a_5 \dot{v}_{oa} + a_6 \rho + C \quad (5.2)$$

where

- T_{mixed} = the mixed air temperature,
- T_{supply} = the supply air temperatures,
- T_{oa} = the outdoor air temperature,
- \dot{v}_{oa} = the volume flow rate of the outdoor air, and
- ρ = the static pressure in the air distribution system.

Variables $a_1 \dots a_7$ are the coefficients found through model training. The variable E is a binary value that is 1 when the experiment is running, and 0 when it is not. We interact this “dummy” variable with T_{mixed} , because T_{mixed} is affected by our experiment. That is, when the experiment is running, zone air temperatures are modified, resulting in a change in mixed air temperature. Thus, the coefficient, a_3 , associated with E allows us to estimate the average impact of the experiment on mixed air temperature. The last term, C , captures the differences between experiment and non-experiment days that are not accounted for by the experiment, i.e., the unobserved disturbances affecting fan power.

A similar model can be formulated for the AHU cooling coil load, \dot{Q}_{cool} :

$$\begin{aligned} \dot{Q}_{cool} = & b_1 T_{supply} + b_2 T_{mixed} + b_3 E T_{mixed} + b_4 T_{oa} + b_5 \dot{v}_{oa} + b_6 \rho + b_7 \dot{v}_{supply} \\ & + b_8 E \dot{v}_{supply} + C \end{aligned} \quad (5.3)$$

where the supply air volume flow rate, \dot{v}_{supply} , is interacted with the experiment dummy variable and variables $b_1 \dots b_9$ are the coefficients found through model training. The dependent variable, cooling coil load, and chiller power are calculated in the manner described previously.

Once the models are trained, we use them to predict the dependent variable (either P_{fan} or \dot{Q}_{cool}) to produce a time-series estimate for the experiment day ($E = 1$). Next, we use the same model to predict the dependent variable for the same day, assuming the experiment was not performed ($E = 0$). A comparison between two predictions gives an estimate of the impact of our system on building energy consumption and demand.

Two separate machine-learning techniques were investigated for training a model: general linear regression and Gaussian process regression. The general linear regression model is a widely used method for representing the relationship between dependent and independent variables. Gaussian process regression uses Gaussian distribution and covariance functions to perform the prediction of data. In the end, general linear regression was selected because it best captured the temporal behavior of the dependent variables and resulted in improved residuals, R^2 , and p value compared to Gaussian regression. We observed that Gaussian regression models failed to capture trends when large changes in values from one sample to the next exist. A possible reason is that Gaussian process model may only represent a limited number of covariance functions, thus the model is not flexible enough to capture the data when there is a sudden change (Murphy 2012).

In the results that follow, we present analysis using Method 2. Because each experiment performed has unique settings, we select a separate training set from the data collected. For each experiment, we select the experiment day and approximately 1 week of non-experiment days both before and after the experiment for training.¹ Other experiment days lying within this range are eliminated.

5.3.3 Limitations and Assumptions

A few assumptions were made before applying the machine-learning techniques to the measured data. First, it is assumed that data measured at one time point are independent of data measured at another time point, and the measured data do not affect the next data point measured after the time interval of measurement. Second, a few factors that are not considered in the developed models may also affect HVAC operation such as wind pressure on building envelope, solar insolation, environmental humidity, occupant behavior, and internal thermal gains. The analysis used the developed models to predict both situations (with and without experiment flag), and the predicted results were compared based on the predicted results. Therefore, the model prediction error may be considered as a system error and should not affect the conclusions addressed from this comparison.

We recognize that the analysis methods described do not represent the state-of-the-art, and that other methods may yield different results. Although the results of model fitting show that the generalized linear regression method works fairly well for the current research problem, other machine-learning techniques are also widely used to solve nonlinear problems. We are continuing to explore additional methods for predicting building energy use at a range of time scales from minutes to hours. For example, neural networks are a nonlinear generalization of general linear models (Magoulès et al. 2013). A few studies described an artificial neural networks (ANNs) model for predicting the building thermal load (Forrester and Wepfer 1984; Kawashima et al. 1996; Kreider and Wang 1992; Zhuang et al. 2015). Zhuang et al. (2015) conducted a case study to predict the building cooling load based on neural networks. Anstett and Kreider (1993) examined the accuracy of the ANN model for energy predictions. While predicting building energy consumption, the outputs given by ANN may not be exactly as expected, Kajl et al. (1997) developed a fuzzy logic to correct the ANN outputs. Use of support vector machine (SVM) models is another popular approach for solving nonlinear problems by creating a sparse kernel machine (Murphy 2012). Previous studies show that SVMs work well in predicting lone terms (such as monthly and annual) building energy consumption (Dong et al. 2005; Lai et al. 2008). Li et al. (2009) used SVMs to predict the hourly cooling load of an office building and indicated that SVMs perform better than neural networks; a similar conclusion is also addressed in another study (Li et al. 2010). As stated above, we assumed independence between measured data points in the current research. For analysis of time-series data without this assumption, the autoregressive moving-average (ARMA) model, which is a

¹ During model training, certain predictors may be found to be statistically insignificant. These predictors are eliminated from the model when it is used to subsequently predict the dependent variable.

classical method, could be used to predict time-series data (Murphy 2012). MacArthur et al. (1989) and Spethmann (1989) developed a forecasting method based on the autoregressive integrated moving-average (ARIMA) model for an optimal cold storage controller. The Hidden Markov Model (HMM) is a stochastic model, which has an advantage over Markov models by representing long-range dependencies between observations (Murphy 2012). However, it was not much investigated in the field of building energy prediction (Zia et al. 2011). Mocanu et al. (2014) compared multiple machine-learning methods for estimating building energy consumptions. The authors investigated a newly developed stochastic model for time-series prediction (Continuous Restricted Boltzmann Machine [CRBM]), and concluded that CRBMs outperform ANNs and HMMs. As a summary, other machine-learning techniques listed above are robust models for solving nonlinear problems when generalized linear models may not provide good results. However, the disadvantages of using these models are that they may require sufficient and large amount of training data to achieve desired accuracy of predictions, and they have extremely high complexity compared to statistical models (Magoulès et al. 2013).

5.4 Results and Analysis

In this section we present results from physical testing of the market-based controls. We stress that these results have been significantly affected by the challenges already enumerated and should be interpreted as preliminary and not necessarily representative. Here, we analyze the results critically and provide further commentary on deficiencies identified in our tests.

5.4.1 Flat Price Experiments

Simulation experiments shown in our previously published work demonstrated an inverse relationship between electricity price and both daily HVAC electricity use and peak electric demand (Hao et al. 2016). In the fixed price experiments, the same observation generally holds with a few exceptions. A summary of these metrics is presented in Table 5.12. Note that negative values indicate a reduction over the baseline.

Table 5.12. Difference in energy and demand for flat price experiments.

Experiment Date	Series	Energy (kWh)	Energy (%)	Demand (kW)	Demand (%)
June 29	Fan	-0.27	-0.14%	-0.27	-6.48%
	Chiller	0.00	0.00%	-0.02	-0.20%
	Total	-0.27	-0.04%	-0.16	-1.15%
August 18	Fan	-0.14	-0.09%	-0.14	-2.99%
	Chiller	-0.21	-0.04%	-0.21	-2.13%
	Total	-0.35	-0.05%	-0.35	-2.43%
September 1	Fan	-0.28	-0.20%	-0.24	-6.59%
	Chiller	-0.14	-0.04%	-0.12	-1.43%
	Total	-0.42	-0.09%	-0.42	-3.71%
September 8	Fan	-0.03	-0.03%	-0.03	-1.17%
	Chiller	-0.12	-0.05%	-0.15	-2.20%
	Total	-0.15	-0.04%	-0.19	-1.99%
September 13	Fan	-0.19	-0.12%	-0.19	-6.57%
	Chiller	-0.03	-0.01%	0.08	0.98%
	Total	-0.22	-0.04%	-0.11	-1.02%
September 14	Fan	-0.21	-0.17%	-0.21	-7.55%
	Chiller	-0.04	-0.02%	-0.02	-0.26%
	Total	-0.25	-0.07%	-0.24	-2.35%

In contrast to the direct comparison method, the regression method shows reduced demand and energy during the June 29 experiment across both the fan and chiller. Predicted baseline and experiment results are shown in Figure 5.2. Consistent with the measured data, the model predicts a spike in chiller peak demand in the afternoon. Recall from the previous discussion that there is no way to determine whether this spike resulted from the experiment. By comparison, August 18 (Figure 5.3) is perhaps the “cleanest” result of all flat price experiments, with no unexplained chilled water spikes, and little fan power oscillation from static pressure reset logic. Coincidentally, the greatest reduction in chiller demand and second greatest total HVAC electric demand is measured during this experiment. The nearly constant fan power between mid-morning and early afternoon coincides with all zones requesting minimum airflow.




Figure 5.2. Comparison of measured and model-predicted chiller and fan power, June 29.




Figure 5.3. Comparison of measured and model-predicted chiller and fan power, August 18.

The last week of August marks the end of hot weather and its associated cooling loads. Beyond this date, economizer operation is observed with increasing regularity. In addition, estimated chiller power becomes much less consistent. This can be observed in the September 1 and September 8 experiments (Figure 5.4 and Figure 5.5). Recall that chiller power is estimated from chilled water load, which is a function of flow rate and coil supply/return temperature, and not a true measurement of electric power. The frequent changes in the estimated chiller power can be attributed to water flow adjustments resulting from low cooling loads and economizer operation, which are compounded by frequent changes in the static pressure set point. In general, all zones tend to request minimum airflow; the static pressure set point appears to be driven by a single zone, Zone 133. Recall that the temperature sensor in this zone appears to provide faulty measurements, resulting in frequent changes in airflow and subsequent changes in static pressure set point.




Figure 5.4. Comparison of measured and model-predicted chiller and fan power, September 1.




Figure 5.5. Comparison of measured and model-predicted chiller and fan power, September 8.

Results from experiments on September 13 and 14 (Figure 5.6 and Figure 5.7) show savings similar to previous days, with the exception of chiller demand on September 13. It is unclear why chiller demand increased; there is no change in supply water temperature as we observe in the June 29 case. We note that the outdoor air temperatures measured during these days were substantially lower than in previous experiments. This resulted in significant economizer operation, and perhaps most importantly, consistently low zone airflows that deviate little from the configured minimum values, except in cases where static pressure reset logic results in increased airflow.




Figure 5.6. Comparison of measured and model-predicted chiller and fan power, September 13.

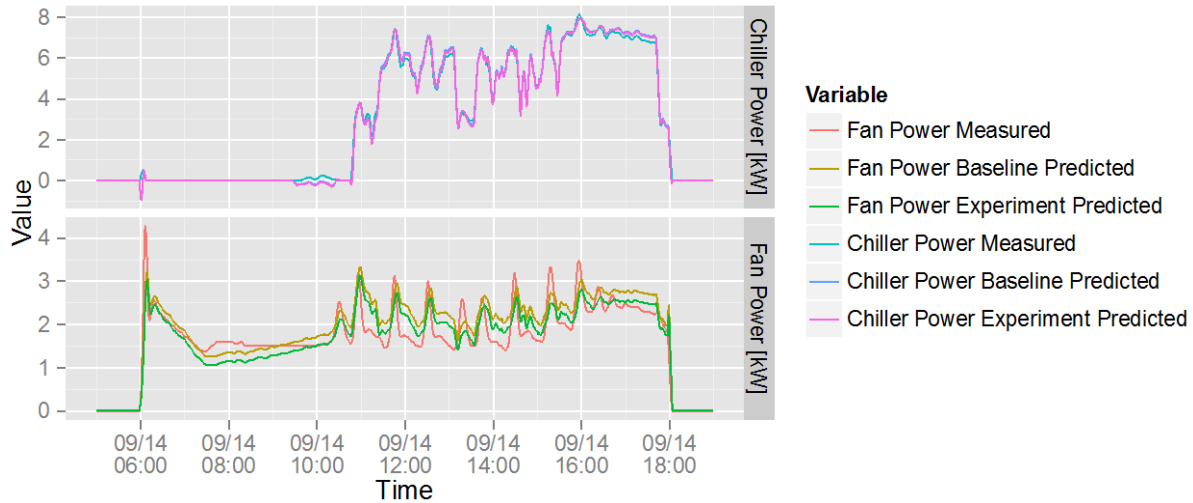


Figure 5.7. Comparison of measured and model-predicted chiller and fan power, September 14.

5.4.2 Demand-Limiting Experiments

Demand limit experiments show mixed results (Table 5.13) compared to expectations from simulations. Similar to flat price experiments conducted between late August and September, plots for these three experiments (Figure 5.8, Figure 5.9, and Figure 5.10) show low chiller load in the morning hours due to economizer operation, and frequent fan power oscillations due to static pressure reset. Results from August 25 indicate that market-based control results in an increase in demand. In this case, a demand limit of 10 MW is imposed, which is approximately 90% of the demand measured during the days prior to the test. Logs collected during the experiment indicate that the limit is exceeded consistently during the 1.5 hours in which the demand limit is imposed, and that set points are generally increased in an attempt to decrease load. One zone is observed to lower its set point during the experiment, which is not unusual as zones trade-off against each other in the market, but the total increase in demand this one zone creates does not account for the net increase in building HVAC demand. After analyzing the occupancy status for this day compared to that used in regression model training, we speculate that occupancy may play a role in these results. August 31 has a similar, but smaller, increase in chiller demand, which is offset by a decrease in fan demand that results in a net decrease in building HVAC demand. Interestingly, on this day, the demand limit (also 10 MW) was exceeded only a handful of times between 14:00 and 17:00, and for only 5–10 minutes at a time. Savings in both energy and demand are observed in the September 2 case. Though modest, these results are consistent with the simulation. On this day, a much lower demand limit of 6.5MW is imposed, which is exceeded six separate times throughout the day for durations of 15 minutes to 1 hour.

Table 5.13. Difference in energy and demand for demand limit experiments.

Experiment Date	Series	Energy (kWh)	Energy (%)	Demand (kW)	Demand (%)
August 25	Fan	0.00	0.00%	0.00	0.00%
	Chiller	0.23	0.04%	0.25	2.59%
	Total	0.23	0.03%	0.25	1.94%
August 31	Fan	-0.11	-0.07%	-0.11	-3.43%
	Chiller	0.01	0.00%	0.08	0.95%
	Total	-0.10	-0.02%	-0.03	-0.26%
September 2	Fan	-0.05	-0.04%	-0.05	-1.80%
	Chiller	-0.12	-0.05%	-0.13	-1.66%
	Total	-0.17	-0.05%	-0.18	-1.71%




Figure 5.8. Comparison of measured and model-predicted chiller and fan power, August 25.




Figure 5.9. Comparison of measured and model-predicted chiller and fan power, August 31.




Figure 5.10. Comparison of measured and model-predicted chiller and fan power, September 2.




Figure 5.11. Comparison of measured and model-predicted chiller and fan power, September 15.

6.0 Conclusions and Next Steps

Experimental results, though limited, showed promise for the application of market-based controls to commercial building HVAC systems. Savings were modest compared to results from previous simulations, but we believe they can be improved. Some differences between experiments and simulations were expected due to mismatches between the EnergyPlus model and the physical building, while many differences may likely be attributed to imperfect modeling of building dynamics. In simulation, decisions made by VAV agents were acted upon immediately and control was perfect; in a physical experiment, control decisions came with delays, and control was imperfect.

Issues encountered during early testing and integration (see Section 5.2.3) severely limited the number of tests we were able to perform. As a result, many of the flat price experiments occurred after peak cooling season. Additionally, we were unable to fully test demand-limiting and dynamic price cases.

Because of the high minimum airflows, tests performed later in the summer were unable to affect electricity demand in any meaningful way; zones stayed at minimum airflow during experiments due to low cooling demand. Although it was difficult to determine to what extent the minimum airflow configuration affected the experiments, we believe that with additional flexibility in airflow, additional demand and energy savings can be achieved, especially when loads are low.

We were unable to determine the impact of the static pressure reset logic on these experiments. In some experiments, the reset may have been triggered by adjustments made by the VAV agents; in others, zones outside of our control may have triggered this logic, e.g., Zones 133 and 142. An interesting question we would like to answer is whether or not the market system could be used to prevent a static pressure reset, and thus prevent a spike in fan power.

The role of occupancy in our experiments is yet another area in which we should focus future efforts. Recall that occupied and standby cooling set point limits were different. A difference in occupancy between days may have a significant impact on energy use and demand, irrespective of an experiment. At the least, our analysis methodology needs to control for occupancy so that differences can be correctly attributed. For example, differences in occupancy between days used to train the regression model and the experiment day could incorrectly attribute the difference in energy consumption to the experiment. Although we cannot definitively point to this as a cause for the increase we observed on August 25, occupancy is nevertheless a key factor that needs to be considered.

One key insight we have gained as a result of early testing regards the precision of the physical VAV controller. Specifically, early testing showed that VAV agents often make very small temperature set point adjustments, on the order of 0.1°F. The physical VAV controller is incapable of controlling to this level of precision, and thus cooling demand remains unaffected. To correct this, we made changes to the VAV logic such that the agent bids a wider range, and therefore makes larger temperature set point adjustments for a given clearing price. With additional testing, further refinement of this bidding strategy, and correction of a few other key issues, e.g., zone regression coefficients, we believe that we may show improved performance approaching that observed in previous simulation studies.

Future efforts will be focused on exploring dynamic signals and the extension of market-based control to heating operation. We will continue to adapt this method to additional HVAC system types. As discussed in Section 3.4, an RTU-HP model has been developed, which, along with the VAV system developed here, will be used in a campus-wide deployment of transactive controls during the second phase of this project.

7.0 References

- AHRI (Air-Conditioning, Heating, Refrigeration Institute). 2004. AHRI Standard 540, *Performance Rating of Positive Displacement Compressors and Compressor Units*. Arlington, Virginia.
- Alerton. 2016. <https://alerton.com/en-us/pages/default.aspx>.
- Anstett M and JF Kreider. 1993. Application of neural networking model to predict energy use. *ASHRAE Transactions 1*, 505–17.
- Apra C and C Renno. 2009. Experimental model of a variable capacity compressor. *International Journal of Energy Research 33*(1), 29–37. doi: 10.1002/er.1468
- CAISO (California Independent System Operator). 2016. <http://www.caiso.com/Pages/default.aspx>.
- Cheung H and JE Braun. 2010. Performance Characteristics and Mapping for a Variable-Speed Ductless Heat Pump. Paper presented at the 13th International Refrigeration and Air Conditioning Conference, West Lafayette, Indiana.
- Corbin CD and N Radhakrishnan. 2016. *Transactive Control of Commercial Building HVAC Systems: VOLTRON User Guide*. PNNL-XXXX, Pacific Northwest National Laboratory, Richland, Washington.
- DOE (U.S. Department of Energy). 2011. Buildings Energy Data Book. Retrieved May 15, 2014, from <https://catalog.data.gov/dataset/buildings-energy-data-book>
- Dong B, C Cao, and SE Lee. 2005. Applying support vector machines to predict building energy consumption in tropical region. *Energy and Buildings 37*(5), 545–553. doi: <http://dx.doi.org/10.1016/j.enbuild.2004.09.009>
- Duprez M-E, E Dumont, and M Frère. 2007. Modelling of reciprocating and scroll compressors. *International Journal of Refrigeration 30*(5), 873–886. doi: <http://dx.doi.org/10.1016/j.ijrefrig.2006.11.014>
- Duprez, M-E, E Dumont, and M Frère. 2010. Modeling of scroll compressors–Improvements. *International Journal of Refrigeration*, 33(4), 721-728. doi: <http://dx.doi.org/10.1016/j.ijrefrig.2009.12.025>
- EIA (U.S. Energy Information Administration). 2015. The State Energy Data System. Retrieved March 9, 2016, from <http://www.eia.gov/state/seds/>
- EnergyPlus. 2016. EnergyPlus Energy Simulation Software. <http://apps1.eere.energy.gov/buildings/energyplus/>.
- Forrester JR and WJ Wepfer. 1984. Formulation of a load prediction algorithm for a large commercial building. *Ashrae Transactions 90*, 536–551.
- Fraisse G, C Viardot, O Lafabrie, and G Achard.(2002. Development of a simplified and accurate building model based on electrical analogy. *Energy and Buildings 34*(10), 1017–1031. doi: [http://dx.doi.org/10.1016/S0378-7788\(02\)00019-1](http://dx.doi.org/10.1016/S0378-7788(02)00019-1)

- Gayeski N, T Zakula, PR Armstrong, and LK Norford. 2010. Empirical modeling of a rolling-piston compressor heat pump for predictive control in low-lift cooling. *Ashrae Transactions*, 116.
- Gouda MM, S Danaher, and CP Underwood. 2002. Building thermal model reduction using nonlinear constrained optimization. *Building and Environment*, 37(12), 1255–1265. doi: [http://dx.doi.org/10.1016/S0360-1323\(01\)00121-4](http://dx.doi.org/10.1016/S0360-1323(01)00121-4)
- Hao H, CD Corbin, K Kalsi, and RG Pratt. 2016. Transactive Control of Commercial Buildings for Demand Response. *IEEE Transactions on Power Systems PP(99)*, 1-1. doi: 10.1109/TPWRS.2016.2559485
- Huberman B and SH Clearwater. 1995. A Multi-agent System for Controlling Building Environments. Paper presented at the The 1st International Conference on Multiagent Systems, San Francisco, California.
- Intel. 2016. Mini PC: Intel® NUC: <http://www.intel.com/content/www/us/en/nuc/overview.html>.
- Johnson-Controls. 2016. <http://www.johnsoncontrols.com/buildings/building-management/building-automation-systems-bas>.
- Kajl S, M-A Roberge, L Lamarche, and P Malinowski. 1997. Evaluation of building energy consumption based on fuzzy logic and neural networks applications. Paper presented at the Proceedings of Clima 2000 Conference, Brussels.
- Karlsson M, F Ygge, and A Andersson. 2007. Market-Based Approaches to Optimization. *Computational Intelligence* 23(1), 92-109. doi: 10.1111/j.1467-8640.2007.00296.x
- Katipamula S, DP Chassin, DD Hatley, RG Pratt, and DJ Hammerstrom. 2006. *Transactive Controls: Market-Based GridWise™ Controls for Building Systems*. PNNL-15921, Pacific Northwest National Laboratory, Richland, Washington.
- Kawashima M, CE Dorgan, and JW Mitchell. 1996. Optimizing system control with load prediction by neural networks for an ice-storage system. *Ashrae Transactions* 102(1).
- Kim Wand JE Braun. 2012. Virtual refrigerant mass flow and power sensors for variable-speed compressors. Paper presented at the International Refrigeration and Air Conditioning Conference Purdue.
- Kreider JF and XA Wang. 1992. Improved artificial networks for commercial building energy use prediction. Paper presented at the the ASME Annual Solar Engineering Meeting, Maui, Hawaii.
- Lai F, F Magoulès, and F Lherminier. 2008. Vapnik's learning theory applied to energy consumption forecasts in residential buildings. *International Journal of Computer Mathematics* 85(10), 1563–1588. doi: 10.1080/00207160802033582
- LBNL (Lawrence Berkeley National Laboratory). 2016. Building Controls Virtual Test Bed. <https://simulationresearch.lbl.gov/bcvtb>.
- Li Q, Q Meng, J Cai, H Yoshino, and A Mochida. 2009. Applying support vector machine to predict hourly cooling load in the building. *Applied Energy* 86(10), 2249–2256. doi: <http://dx.doi.org/10.1016/j.apenergy.2008.11.035>

- Li Q, P Ren, and Q Meng. 2010. *Prediction model of annual energy consumption of residential buildings*. Paper presented at the Advances in Energy Engineering (ICAEE), 2010 International Conference, 19-20 June 2010.
- Li S, W Zhang, J Lian, and K Kalsi. 2016. Market-Based Coordination of Thermostatically Controlled Loads—Part II: Unknown Parameters and Case Studies. *IEEE Transactions on Power Systems* 31(2), 1179–1187.
- Li Y, M Liu, and J Lau. 2015. Development of a variable speed compressor power model for single-stage packaged DX rooftop units. *Applied Thermal Engineering* 78, 110-117. doi: <http://dx.doi.org/10.1016/j.applthermaleng.2014.12.038>
- MacArthur JW, A Mathur, and J Zhao. 1989. On-line recursive estimation for load profile prediction. *Ashrae Transactions* 95, 621–628.
- Magoulès F, H-X Zhao, and D Elizondo. 2013. Development of an RDP neural network for building energy consumption fault detection and diagnosis. *Energy and Buildings* 62, 133–138. doi: <http://dx.doi.org/10.1016/j.enbuild.2013.02.050>
- Mocanu E, PH Nguyen, M Gibescu, and WL Kling. 2014. Comparison of machine learning methods for estimating energy consumption in buildings. Paper presented at the Probabilistic Methods Applied to Power Systems (PMAAPS), 2014 International Conference, 7-10 July 2014.
- Murphy K. 2012. *Machine Learning — A Probabilistic Perspective*. The MIT Press.
- Perers B, E Anderssen, R Nordman, and P Kovacs. 2012. A Simplified Heat Pump Model for use in Solar Plus Heat Pump System Simulation Studies. *Energy Procedia* 30, 664–667. doi: <http://dx.doi.org/10.1016/j.egypro.2012.11.075>
- Shao S, W Shi, X Li, and H Chen. 2004. Performance representation of variable-speed compressor for inverter air conditioners based on experimental data. *International Journal of Refrigeration* 27(8), 805–815. doi: <http://dx.doi.org/10.1016/j.ijrefrig.2004.02.008>
- Spethmann DH. 1989. Optimal control for cool storage. *Ashrae Transactions* 95, 710–721.
- Stoecker WS and JW Jones. 1982. *Refrigeration and Air Conditioning*. USA: McGraw-Hill.
- VOLTTRON. 2015. VOLTTRON 3.0: User Guide: <http://bgintegration.pnnl.gov/volttron.asp>.
- Widergren S, K Subbarao, JC Fuller, DP Chassin, A Somani, C Marinovici, JL Hammerstrom. 2014. *AEP Ohio Gridsmart Demonstration Project Real-Time Pricing Demonstration Analysis*. PNNL-23192, Pacific Northwest National Laboratory, Richland, Washington.
- Zhuang J. Y Chen, X Shi, and D Wei. 2015. Building Cooling Load Prediction Based on Time Series Method and Neural Networks. *International Journal of Grid Distribution Computing* 8, 105–114.
- Zia T, D Bruckner, and A Zaidi. 2011. A hidden Markov model based procedure for identifying household electric loads. Paper presented at the IECON 2011 – 37th Annual Conference of the IEEE Industrial Electronics Society.

Appendix A

Supplement to Alternative Market Structures

Appendix A

Supplement to Alternative Market Structures

A.1 Markets and Optimization

In a market structure presented in this document, the ultimate goal is to maximize the total benefit, or social welfare, received by the agents. The total benefit, B_{tot} , can be described as follows:

$$B_{tot} = f_{glo,benefit}(x_1, \dots, x_n) \quad (A.1)$$

where x_i is the allocations of the commodities received by the i th agent and n is the number of the agents. The total benefit can be also expressed as

$$B_{tot} = \sum_{i=1}^n f_{loc,benefit,i}(x_i) \quad (A.2)$$

where $f_{loc,benefit,i}$ describes the benefit for the i th agent. When the benefit of an agent depends only on the allocation of the commodity it receives, Equations (A.1) and (A.2) are equivalent and the allocation problem is said to be separable (Karlsson et al. 2007). As we have shown in our previously published work (Hao et al. 2016), the maximum global benefit, i.e., the optimal solution to the social welfare maximization (SWM) problem, is equivalent to that obtained by maximizing the welfare of each agent individually when the problem satisfies the Karush-Kuhn-Tucker (KKT) optimality conditions. Furthermore, we have shown that the market-based method is equivalent to solving the SWM problem when KKT criteria are satisfied. Thus, the market system yields a globally optimal result if the optimization problem is separable and satisfies KKT.

When the problem is inseparable, optimality is not assured by maximizing the individual benefit of each agent. Therefore, the market system cannot guarantee that the settling price and quantity are globally optimal. As a result, an inseparable system is much more difficult to handle than a separable system that features market-based control. Because agents are designed to maximize their own benefits (local benefits) in the market system, they may make decisions that undermine the global objective. In order for agents to solve the global objective, they require information from others, thereby fundamentally breaking the “rules” of the market system. Furthermore, imbuing agents with this information inhibits scalability. This is not to say that inseparable problems cannot be solved using market methods, rather that some relaxation of the rules is necessary as is well-designed information sharing to support scaling of the system. Finally, it is not at all obvious whether the types of inseparable problems encountered in building systems cannot be solved by market methods in an optimal or near-optimal way. Much additional study is required to understand when and how markets can be used for these types of problems.

Appendix B

Supplement to Rooftop Unit Model

Appendix B

Supplement to Rooftop Unit Model

B.1 BUILDING10 System Characteristics

Eleven heat pumps serve this building. Heat pumps installed in BUILDING10 consist of four *Carrier* models, which are categorized here as models E, F, G, and H. There are eight model Es, one model F, one model G, and one model H whose cooling and heating specifications are listed in Table B.1.

Table B.1. BUILDING10 heating and cooling systems.

Model	Net Cooling Capacity (kW)	System Power (kW)	Cooling EER	Nominal Airflow (m ³ /h)	Net Heating Capacity (kW)	Heating COP
E	21.0	6.76	10.5	4080	20.2	3.4
F	16.7	6.34	9.0	3400	16.7	7.5 ^(a)
G	20.5	7.78	9.0	4077	21.1	3.2
H	10.4	4.04	9.0	2157	10.4	3.1

(a) COP (coefficient of performance) was not provided. Value shown is the Heating Seasonal Performance Factor (HSPF).

The specification for BUILDING10 heat pump fans are listed in Table B.2. The manufacturer-recommended fan motor heat is also included in the table.

Table B.2. BUILDING10 indoor and outdoor fan specifications.

Model	Outdoor Fan		Indoor Fan		
	Type	Motor power (kW)	Type	Motor power (kW)	Fan motor heat (kW)
E	Propeller	0.186	Centrifugal	1.79	0.781
F	Propeller	0.325	Centrifugal	0.9	0.723
G	Propeller	0.325	Centrifugal	1.79	0.723
H	Propeller	0.186	Centrifugal	0.373	NA ^(a)

(a) Performance data are given in net quantities and no information provided by the manufacturer about fan motor heat.

B.2 Nameplate Equipment Characteristics

Table B.3. A sample of general data provided by the manufacturer (Trane 2004).

	WSC060AD,T	WSC072AD,T	WSC090AD,T	WSC120AD,T
Cooling Performance¹				
Gross Capacity - kW (MBh)	17.3 (59.0)	22.6 (77.0)	27.8 (95.0)	34.6 (118.0)
COP (EER) ²	2.72 (9.3)	3.31 (11.3)	3.20 (10.9)	3.28 (10.1)
Nominal Airflow - m3/h (cfm)	3400 (2000)	4080 (2400)	5100 (3000)	6800 (4000)
Rated Airflow - m3/h (cfm)	3400 (2000)	3570 (2100)	4460 (2625)	5950 (3500)
Net Capacity - kW (MBh)	16.7 (57.0)	21.4 (73.0)	26.4 (90.0)	33.1 (113.0)
System Power - kW	6.13	6.46	8.26	11.1
Heating Performance¹				
High Temperature Capacity - kW (MBh)	16.0 (54.5)	19.9 (68.0)	24.0 (82.0) 3	1.9 (109.0)
COP	3.4	3.45	3.54	3.36
System Power - kW	4.71	5.76	6.78	9.48
Compressor				
Number - Type	1-Climatuff Scroll	1-Climatuff Scroll	1-Trane 3-D Scroll	2-Climatuff Scroll
Outdoor Sound Rating - dB³				
	80	85	85	79
Outdoor Fan - Type				
No. Used / Diameter (in.)	Propeller 1 / 22	Propeller 1/26	Propeller 1/26	Propeller 1/26
Drive Type / No. Speeds	Direct/1	Direct/1	Direct/1	Direct/1
CFM	2900	5100	5200	5800
No. Motors / kW (HP)	1/0.30(0.40)	1/.56 (0.75)	1/.56 (0.75)	1/.56 (0.75)
Motor RPM	950	950	950	950
Belt Drive Indoor Fan - Type				
No. Used	FC Centrifugal 1	FC Centrifugal 1	FC Centrifugal 1	FC Centrifugal 1
Fan Diameter x Width - mm - in.	280 X 280 (11 x 11)	305 X 305 (12 x 12)	305 X 305 (12 x 12)	(381 X 381) 15 x 15
Drive Type / No. Speeds	Belt / Variable Speed	Belt / Variable Speed	Belt / Variable Speed	Belt / Variable Speed
No. Motors	1	1	1	1
Standard Motor Power - kW (HP)	1.1 (1.5)	1.1 (1.5)	1.5 (2.0)	2.2 (3.0)
Oversized Motor Power - kW (HP)	-	1.5 (2.0)	2.2 (3.0)	-
Motor RPM - Standard / Oversized	1450 / -	1450 /1450	1450/ 2850	2850 / -
Motor Frame Size	56	56	56	56

NOTES:

1. Cooling Performance is rated at 35.0°C (95°F) ambient, 26.7°C (80°F) entering dry bulb, 19.4°C (67°F) entering wet bulb. Heating Performance is rated at 20.0°C (68°F) ambient, 8.3°C

(47°F) entering dry bulb, 6.1°C (43 F) entering wet bulb. Gross capacity does not include the effect of fan motor heat. Net capacity includes the effect of fan motor heat. Units are suitable for operation to + 20 % of nominal airflow.

2. EER are rated at ARI conditions.

3. Outdoor Sound rating shown is tested in accordance with ARI Standard 270. For more information refer to Performance Data Table “Sound Power Level”.

B.3 Catalog Performance Data

Table B.4. Sample of manufacturer's performance data (Trane 2004).

Table PD-1 — Gross Cooling Capacities (kW) WSC060AD,T (SI)

m3/h Air- flow	Dry Bulb Temp (°C)	Enter. Ambient Temperature (°C)																											
		30				35				40				45															
		Entering Wet Bulb Temperature (°C)																											
		16	19	22	16	19	22	16	19	22	16	19	22	16	19	22	16	19	22	16	19	22	16	19	22	16	19	22	
		TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC	TGC	SHC
3060	24	14.5	13.2	16.6	10.7	18.3	7.6	15.6	13.8	17.7	11	19	8	13.3	12.6	15.3	10.1	17.4	7.2	12.1	12	13.9	9.5	16.3	6.6				
	27	15.3	15.3	16.7	13.4	18.4	10.4	16.3	16.3	17.8	13.9	19.1	10.8	14.3	14.3	15.4	12.8	17.5	10	13.3	13.3	14.1	12.2	16.3	9.5				
	30	16.7	16.7	17	16.3	18.5	13.2	17.6	17.6	18	16.8	19.3	13.4	15.7	15.7	15.8	15.8	17.5	12.8	14.7	14.7	14.7	14.7	16.4	12.3				
	33	17.9	17.9	17.9	17.9	18.7	15.9	18.6	18.6	18.6	18.6	19.5	16	17	17	17	17	17.7	15.6	16	16	16	16	16.6	15.2				
3400	24	14.9	14.1	16.9	11.3	18.5	7.8	16.1	14.7	18	11.5	19.2	8.2	13.7	13.5	15.6	10.4	17.6	7.4	12.5	12.5	14.2	10.1	16.5	6.9				
	27	16	16	17	14.2	18.6	10.9	17	17	18.1	14.7	19.3	11.5	15	15	15.8	13.6	17.7	10.5	13.9	13.9	14.5	13	16.6	10				
	30	17.4	17.4	17.4	17.4	18.8	13.8	18.2	18.2	18.4	17.8	19.5	13.9	16.4	16.4	16.4	16.4	17.8	13.6	15.4	15.4	15.4	15.4	16.7	13.2				
	33	18.5	18.5	18.5	18.5	19	16.7	19.1	19.1	19.1	19.1	19.8	16.7	17.6	17.6	17.6	17.6	18.1	16.6	16.7	16.7	16.7	16.7	17	16.3				
3740	24	15.3	15	17.2	11.5	18.7	8.1	16.5	15.6	18.2	12	19.3	8.4	14.1	14.1	15.9	11	17.8	7.6	12.9	12.9	14.5	10.8	16.7	7.1				
	27	16.6	16.6	17.4	15	18.8	11.2	17.5	17.5	18.3	15.4	19.6	11.3	15.6	15.6	16.2	14.5	17.9	11	14.4	14.4	14.8	13.9	16.8	10.5				
	30	17.9	17.9	17.9	17.9	19	14.4	18.6	18.6	18.7	18.7	19.7	14.4	17	17	17	18	14.2	15.9	15.9	15.9	15.9	16.9	13.9					
	33	18.9	18.9	18.9	18.9	19.3	17.5	19.5	19.5	19.5	19.5	20	17.4	18.1	18.1	18.1	18.1	18.4	17.5	17.2	17.2	17.2	17.2	17.3	17.2				
4080	24	15.6	15.6	17.4	12.1	18.8	8.3	16.8	16.4	18.4	12.4	19.5	8.6	14.5	14.5	16.1	11.5	17.9	7.9	13.4	13.4	14.7	10.8	16.9	7.3				
	27	17.1	17.1	17.6	15.8	18.9	11.6	18	18	18.5	16	19.7	11.8	16.1	16.1	16.5	15.3	18.1	11.4	14.9	14.9	15.1	14.7	17	11				
	30	18.3	18.3	18.3	18.3	19.2	14.9	19	19	19	19	19.9	14.8	17.4	17.4	17.4	17.4	18.2	14.9	16.4	16.4	16.4	16.4	17.1	14.6				
	33	19.2	19.2	19.2	19.2	19.5	18.2	19.8	19.8	19.8	19.8	20.2	17.9	18.5	18.5	18.5	18.5	18.6	18.2	17.6	17.6	17.6	17.6	17.6	17.6				

Notes

- 1. All capacities shown are gross and have not considered indoor fan heat. To obtain net cooling, subtract indoor fan heat.
- 2. TGC = Total Gross Capacity
- 3. SHC = Sensible Heat Capacity

B.4 BUILDING10 Model Coefficients

B.4.1 Cooling Capacity Coefficients

Cooling capacity coefficients of Models E through H are tabulated in Table B.5. An example of a fitted line plot is shown in Figure B.1.

Table B.5. BUILDING10 cooling capacity coefficients and associated R^2 s.

Systems	a_1	a_2	a_3	a_4	a_5	a_6	R^2
Model E	-0.0032	0.0403	0.0307	-1.3346	0.0013	36.1903	0.97
Model F	-0.0011	-0.0584	0.0002	0.4656	0.0001	7.7568	0.99
Model G	-0.0024	0.0409	-0.0010	0.6125	0.0000	6.7663	0.99
Model H	-0.0001	-0.0352	0.0009	0.3576	-0.0022	2.6543	0.99

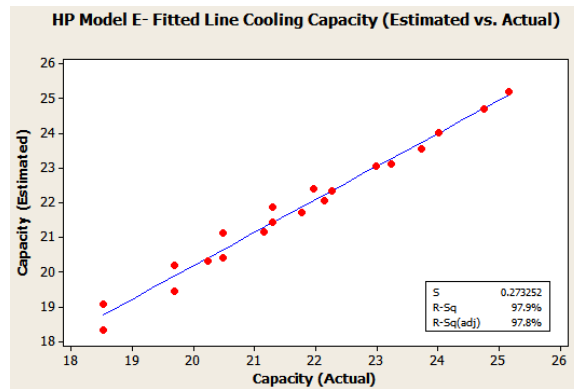


Figure B.1. Fitted line plot showing the estimated vs. actual cooling capacity for Model E.

B.4.2 BUILDING10 Cooling Power Consumption Coefficients

Cooling electric power consumption coefficients of Models E through H are tabulated in Table B.6, and an example of a fitted line plot is shown in Figure B.2.

Table B.6. BUILDING10 cooling electric power coefficients and associated R^2 s.

System	b_1	b_2	b_3	b_4	b_5	b_6	R^2
Model E	0.0009	0.0449	0.0038	-0.1917	0.0002	5.4305	1.00
Model F	0.0010	0.0137	-0.0006	0.0708	0.0003	1.5879	1.00
Model G	0.0004	0.0965	0.0005	0.0105	0.0001	1.7897	1.00
Model H	0.0000	0.0029	0.0009	-0.0565	0.0012	3.5832	1.00

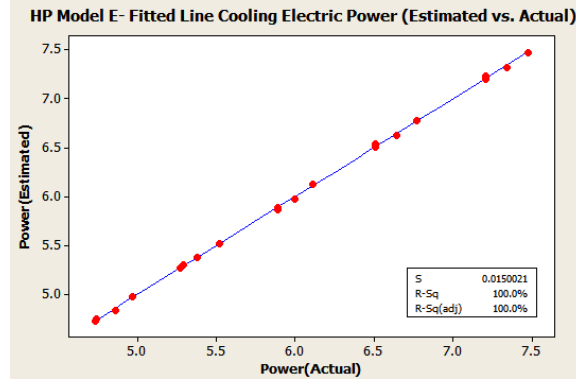


Figure B.2. Fitted line plot showing estimated vs. actual cooling electric power for Model E.

B.4.3 BUILDING10 Heating Capacity Coefficients

Heating capacity coefficients for Models E through H are tabulated in Table B.7, and examples of fitted line plots for Models E and F are shown in Figure B.3.

Table B.7. BUILDING10 heating capacity coefficients and associated R^2 s.

System	c_1	c_2	c_3	c_4	c_5	c_6	R^2
Model E	0.0048	0.4781	-0.0026	0.0506	0.0019	14.8444	0.97
Model F	0.0060	0.3949	-0.0378	1.4607	-0.0003	0.0001	0.95
Model G	0.0033	0.4926	-0.0001	-0.0310	-0.0007	17.4735	0.99
Model H	0.0035	0.3553	-0.0024	0.0281	-0.0010	7.6246	0.99

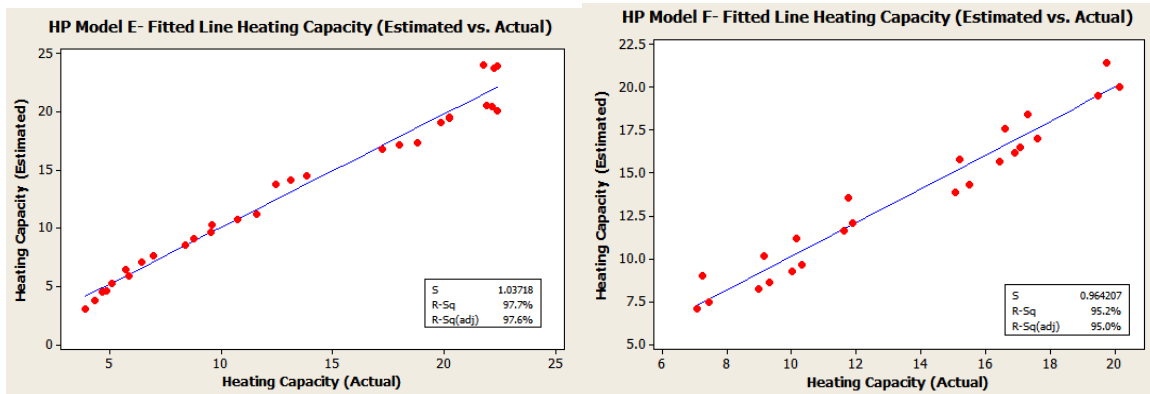


Figure B.3. Fitted line plots showing estimated vs. actual heating capacity for Models E and F.

B.4.4 BUILDING10 Heating Power Consumption Coefficients

Heating electric power consumption coefficients for Models E through H are tabulated in Table B.8 and examples of fitted line plots for Model E and F are illustrated in Figure B.4.

Table B.8. BUILDING10 heating electric power consumption coefficients and associated R^2 s.

System	d_1	d_2	d_3	d_4	d_5	d_6	R^2
Model E	0.0004	0.0293	0.0001	0.0783	0.0010	3.5541	0.98
Model F	0.0007	0.0362	-0.0103	0.4748	0.0007	0.0000	0.89
Model G	0.0008	0.0433	-0.0002	0.0984	0.0010	4.0722	0.99
Model H	0.0001	0.0151	0.0001	0.0117	0.0007	2.7920	0.99




Figure B.4. Fitted line plots showing estimated vs. actual heating electric power for Models E and F.



Pacific Northwest
NATIONAL LABORATORY

*Proudly Operated by **Battelle** Since 1965*

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)

U.S. DEPARTMENT OF
ENERGY

www.pnnl.gov