UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title

{Control of Residential Load Management Networks Using Real Time Pricing

Permalink

https://escholarship.org/uc/item/97j8z6ng

Author

Burke, William Jerome

Publication Date

2010

Peer reviewed|Thesis/dissertation

Control of Residential Load Management Networks Using Real Time Pricing

by

William Jerome Burke

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering – Mechanical Engineering

in the

Graduate Division
of the
University of California, BERKELEY

Committee in charge:
Professor David Auslander, Chair
Professor Edward Arens
Professor Paul Wright

Spring 2010

Control of Residential Load Management Networks Using Real Time Pricing

Abstract

Control of Residential Load Management Networks Using Real Time Pricing

by

William Jerome Burke

Doctor of Philosophy in Engineering – Mechanical Engineering

University of California, Berkeley

Professor David Auslander, Chair

The goal of residential load management is to manipulate the electricity demand of a group of consumers to match some desired shape. Effective load management needs two types of controls: systemic and local. Systemic control manipulates the aggregate power by signaling the consumers. Local control determines how the power consumer reacts to the signals. I present a price responsive intelligent thermostat for use in local load management. Further, I explore two types of systemic control using electricity price as the control variable. The first approach is to identify the system using a low order model and apply traditional control techniques. The second approach makes use of game theory to auction the electricity to the consumers. Both methods enable reliable and robust high resolution load management.

For
The Love Of My Life,
Krista

Contents

List of Figures

List of Tables

Acknowledgments

Many thanks go to my research advisor, Professor David Auslander. He has been a wonderful mentor. I most appreciate that he largely allowed me to set my own research direction. This enhanced my creativity and helped deepen my self motivation.

Thanks to Ron Hofmann for supporting this work from its inception. He saw the value of this work when no one else was discussing systemic modeling and control of load management. Further, without his assistance procuring funding, my research would have been much more difficult to complete.

Thanks also go to Joe Eto for excellent advice and encouragement. Joe was the first person to suggest that I consider what I would do if I was an electricity aggregator, and that simple idea spawned much of the price response research. Joe also helped with funding, and I owe him a debt of gratitude for that.

Most of all, I want to recognize the unselfish love of my wife. Even though it would delay our financial and familial progress, Krista encouraged me to strive for my dream of returning to school for a PhD. She is the most wonderful person I know, and I am very fortunate to be married to her.

Chapter 1

Introduction To Load Control

1.1   Motivation of Load Control

The interconnected smart grid promises many advantages to reliable and inexpensive energy distribution, and load management plays a key role in the smart grid's ability to deliver on these promises. Load management modifies power consumption in order to better suit supply constraints. It consists of many different techniques for both commercial and residential energy customers, and a good overview is given in [1].

Generally, load management programs are differentiated by the scale of the electricity demand under control. Commercial and industrial load management aims to manipulate the demand of large consumers, like office building lighting loads and factory processes. Because each individual participating in the program has a high demand, each individual *can* have a sizable effect on the overall load shape. Residential load management projects reduce the demand of residential customers. By stark contrast with commercial load management, each individual consumer affects the demand very little, and therefore the aggregate demand is more important than any individual's demand.

Residential load management has been deployed throughout the electricity generation and distribution system (the grid) for many years. Thermostatically controlled devices, such as heating ventilation and air conditioning (HVAC), refrigerators, and water heaters, are particularly conducive to load management because they store (thermal) energy and contribute heavily toward peak demand. Of these loads, HVAC systems have received the most attention, and consequently load management technologies for these devices have proliferated.

Direct load control (DLC) is the oldest and most widely used residential load management technique. DLC directly manipulates the on/off cycles of a residential HVAC unit via a radio operated switch attached between the compressor and mains power. The participants of DLC programs are split in to discrete groups that cycle on and off at the same time. Different groups are controlled in different ways in order to obtain a somewhat controllable aggregate demand. One common approach to this is as a scheduling problem, of which there has been considerable research. The authors of [2] present a means of scheduling DLC to obtain desired load shapes.

Using linear programming to schedule DLC for profit maximization is examined in [3].

One of the main problems of DLC is that the original switches controlled the HVAC units in an open-loop manner, completely disregarding the natural cycling of the units. The resulting predictability of the DLC cycling simplified control but created other problems, like loss of system diversity and free-riders. Periodically forcing large numbers of HVAC systems off at the same time synchronizes their cycling, and that synchronization persists after the DLC period ends, resulting in a loss of system diversity that could cause massive spikes in the demand. The term "free-riders" refers to units that are in a DLC program but operate at a duty cycle lower than the DLC duty cycle. The free-rider problem as been addressed through adaptive load switches as presented in [4].

The programmable communicating thermostat (PCT) is of growing interest because of the possibility that it could obviate these concerns about DLC. A PCT has temperature set-point programs like a traditional programmable thermostat, but it also includes a radio for load management communications. Our work developing the proof-of-concept prototype of the PCT is reported in [5]. Exactly how to manage load using PCTs is a topic of current interest, but the most popularly mentioned technique is by making static adjustments to the set-point temperature, as in [6, 7]. In [8], the author develops more advanced feedback techniques using set-point control of PCTs.

A few other researchers have taken a different approach to load management. In [9], they discuss autonomous agents bidding for energy in a distributed fashion. Appliances that automatically respond to fluctuations in grid frequency were presented in [10]. The article [11] presented a residential devices that responds to real-time pricing.

Most generally, load management control falls into two distinct yet coupled categories – systemic control and local control. Systemic control aims to modulate the aggregate power consumption to achieve some goal. It assumes that a super-agent, such as an Independent System Operator (ISO), power company, or commercial aggregator, communicates with each consumer in order to direct the total consumption. Local control considers each consumer connected to the load management network as an autonomous agent that makes decisions about how best to consume electricity. The coupling between the two types of controls occurs through communications between agents and super-agent.

Every load management technique fits into the systemic/local control hierarchy. Direct load control gives all of the control authority to the super-agent, and the local agent follows precise orders. A network of PCTs, share the control authority – the super-agent directly controls the initial thermostat set-back but the homeowner can override it. A network of intelligent agents bidding for energy on an open market entrusts all of the control authority to the local controllers.

The goal of our research is to control the aggregate response of a large network of PCTs. Toward this aim, we expanded the field of residential load management by introducing a new local control strategy and two new systemic control strategies. The following sections summarize the remainder of the thesis.

1.2   Systemic Simulation

To ensure the smart grid truly is "smart" before deployment, the load management control algorithms must be fully vetted. Unfortunately, deployment of smart load management technologies is very expensive. Further, real world experimentation does not allow testing of extreme circumstances until they are experienced in the real world, when unexpected behavior could result in catastrophe.

In order to enable safe and inexpensive experimentation, Chapter 2 presents a modular and extensible dynamic simulation of an advanced load management system capable of examining the response to different systemic and local demand response control strategies. The load group simulation considers local and systemic control directly by modeling large groups of agents separately from the super-agent. The agents are fully independent of one another (and the super-agent), and each one consists of a dynamic simulation, addressable communications, and discrete controls. The super-agent uses discrete control logic and communications with the agents to implement systemic control. The advantages of this architecture are as such:

- High resolution dynamic modeling yields accurate dynamic responses of the agents at a small sample time.

- Independence between individual agents and the super-agent provides appropriate load diversity.

- Communications modeling allows experimentation with different levels of agent information awareness and super-agent control.

- Discrete control logic and modular software allows quick and simple changes to the control algorithms inside the agents and super-agent.

1.3   Local Control Using Low-Frequency PWM

The most common type of residential heating ventilation and air conditioning compressor is single speed, meaning it is either off or on at full power. Because of maintenance, reliability, and efficiency concerns, the compressors must cycle at relatively low frequencies. Furthermore, the low cycle rate introduces considerable residual in the system output (inside temperature). Traditionally, these systems use a non-linear hysteresis controller for temperature set-point following, where the cycle rate is not directly defined. Instead, the width of the hysteresis band determines the duty cycle. Hysteresis control is very simple to implement, model free, and robust. Unfortunately, it has a number of disadvantages when viewed from a modern perspective.

Chapter 3 proposes a technique in which the single speed compressor can be treated as a variable power unit using low frequency pulse width modulation (PWM). Providing that the continuous system responds slowly, the discrete time PWM system can still be considered linear.

The difficulty arises in error measurement because the states of the system can change considerably from the start of the PWM time period to the end. Consequently, the main design effort comes from appropriate filter design.

Low frequency PWM control has a number of advantages over traditional control of HVAC compressors. Firstly, any linear or non-linear control design technique producing a proportional input signal can be used to control the unit. Another advantage is that the power consumption of the unit can be explicitly controlled using tunable saturation limits, which is particularly important for local control of load management and dynamic electricity pricing. Finally, operation of multi-stage and variable HVAC compressors becomes much easier with a proportional control signal.

1.4   Auction Response Using Synchronized PWM

The goal of the work presented in Chapter 4 is to develop a framework for distributing shared scarce resources amongst intelligent autonomous agents. In particular we are interested in modulating the total power consumption of a group of independent agents responsible for residential HVAC operation. Our system is hierarchical, consisting of independent home agents responsible for local control and a super-agent responsible for the power regulation. The coupling between the home agents and the super-agent occurs through shared communications.

We propose a market based approach to load management using a scarce resource auction. Each home agent knows a demand function that defines its price versus demand desires, but in order to effectively operate in the auction, it must be able to predict its future power consumption. By synchronizing the low frequency PWM of the home agents with the auction time windows, the complexity for learning and predicting local power consumption is drastically reduced. Additionally, adjusting the power consumption in response to the time varying price is simplified.

1.5   Fast Auction Clearing Algorithm

In Chapter 4, we introduce the notion of auctioning electricity, but the actual auction mechanism leaves much to be desired. We used the Tâtonnement Process, but this mechanism is a poor fit for this application because convergence requires a large number of messages. The proposed automatic auction process needs a fast auction mechanism that is guaranteed to converge with minimum number of messages.

In Chapter 5 we develop an improved auction mechanism – the Soft Budget Constrained Mechanism. This mechanism takes bids and a desired aggregate power demand, and computes the uniform price that guarantees the aggregate demand will match the desired value. This mechanism has several excellent properties. Firstly, it is fast, meaning it is computable in polynomial time. Secondly, it is communication efficient, in that only one message per bidder is needed to clear the auction. Most importantly, the mechanism is policy consistent. Policy consistency is

a powerful property that indicates that the bidder will always tell the truth because they cannot get a better outcome from lying.

1.6   Systemic Control Using Sliding Controller

Chapter 6 is concerned with controlling HVAC energy consumption, but we come at it with a different approach. Instead of auctioning the electricity, we have approached this problem from a traditional controls perspective – model it and apply feedback controls.

One key challenge of controller design is the high system complexity. This is a very large order non-linear system. Just considering the reduced order software-in-the-loop simulation, each house has four states, three uncontrolled inputs (conduction with outside air, infiltration of outside air, and solar radiation), and is controlled by a non-linear controller for both inside temperature and price. This system is nearly impossible to exactly model for systemic control purposes.

Another challenge that a feedback controller must meet is robust stability. This system would be used to improve the stability of the grid by reducing the demand in times of high stress on the grid. If the system were to become unstable, it could cause massive disruptions.

The final challenge is that we must use a low bandwidth input that changes on a 15 minute period. HVAC systems have slow time constants because the thermal dynamics of the house are slow. There are also reliability issues with cycling HVAC equipment too quickly. Furthermore, for infrastructure cost reasons, this system is designed to work with a slow communications system, like the digital sub-band on broadcast FM radio.

In this chapter, we develop a discrete time observer based sliding controller that adjusts the price of electricity in order to regulate the power demand. Using the software-in-the-loop simulation, we demonstrate the performance and robustness of the controller.

1.7   Implications

In Chapter 7, we conclude the thesis by summarizing the results and outlining the implications of the research. We discuss the unmet challenges that the research still faces. Most of the technical challenges are within grasp, but the policy implications of automatic price response are still unknown. Further, we outline a business case for this technology. Our research could almost immediately be used by an electricity aggregation company to manage demand for a group of houses. The controllable demand would enable the aggregator to buy energy on the wholesale markets and sell various reserve products. Finally, this research could be used to improve the efficiency and reliability of the grid.

Chapter 2

Modular and Extensible Systemic Simulation of Demand Response Networks

2.1 Introduction

Thermostatically controlled devices, such as heating ventilation and air conditioning (HVAC), refrigerators, and water heaters, are particularly conducive to demand response technologies because they store energy and contribute heavily toward peak loads. This segment has been studied for a relatively long time, and consequently demand response technologies have proliferated. A relatively recent technology is the programmable communicating thermostat (PCT) that has all of the normal functions of programmable thermostats, but also has a means to receive load management control signals.

In order to enable safe and inexpensive experimentation we constructed and verified a modular and extensible dynamic simulation of an advanced load management system capable of examining the response to different systemic and individual demand response control strategies. The load group simulation considers individual and systemic control directly by modeling large groups of agents separately from the super-agent. The agents are fully independent of one another (and the super-agent), and each one consists of a dynamic simulation, addressable communications, and discrete controls. The super-agent uses discrete control logic and communications with the agents to implement systemic control. The advantages of this architecture are as such:

- High resolution dynamic modeling yields accurate dynamic responses of the agents at a small sample time.

- Independence between individual agents and the super-agent provides appropriate load diversity.

- Communications modeling allows experimentation with different levels of agent information awareness and super-agent control.

- Discrete control logic and modular software allows quick and simple changes to the control algorithms inside the agents and super-agent.

We organized this chapter into five main sections. Section 2.2 gives a detailed overview of the simulation from details of the house model to the super-agent control. In Section 2.3 we explain the process by which we chose model parameters. Section 2.4 outlines some sample results to demonstrate the capabilities of the simulation. We examine a few different individual and systemic control scenarios applied to a network of residential intelligent thermostats. Our first demonstration shows the basic case of a PCT with static thermostat setback control. From there, we add dynamics to the agents and super-agent in order to implement payback smoothing. Then, we demonstrate a preliminary version of an intelligent agent responding to energy price sent from the super-agent (this will be further developed in later chapters). We conclude in Section 2.5 by outlining the power of our simulation and its future potential advanced control techniques.

2.2   Load Group Simulation Overview

This simulation focuses on simulating residential HVAC systems controlled by smart thermostats, but it could easily be modified to accommodate other types of thermostatically controlled devices. In the 1980's many researchers focused on modeling thermostatically controlled devices (HVAC, water heater, refrigerator, etc), and an excellent treatment of their work is given in [12]. Some of the physically based models are treated in [13, 14, 15], and a Markov based approach is given in [14]. More recently, researchers presented a State Queuing model for thermostatically controlled devices in [16].

The load group simulation was written using the TranRunC architecture, which was invented by David Auslander and detailed in [17]. TranRunC provides an object oriented approach to programming real time systems within the C programming language. The style utilizes strict task/state hierarchy as developed in [18].

The simulation consists of four main tasks, as illustrated in Figure 2.1. The Neighborhood Task is the heart of the simulation, as it contains a dynamic model of a large population of independent and random PCT controlled houses. The Measurement task performs the mundane task of aggregating the load. The Controller task sends messages to the smart thermostats, allowing examination of demand response events. Finally, the Master Task simply makes sure the other tasks behave during start-up and shutdown.

2.2.1   Neighborhood Task

The Neighborhood Task consists of a collection of house models with each house having unique thermal parameters and individualized thermostat settings. Each house is subject to the same outside temperature and solar gains. The outdoor environment forms the only coupling between the houses, apart from the demand response messages.
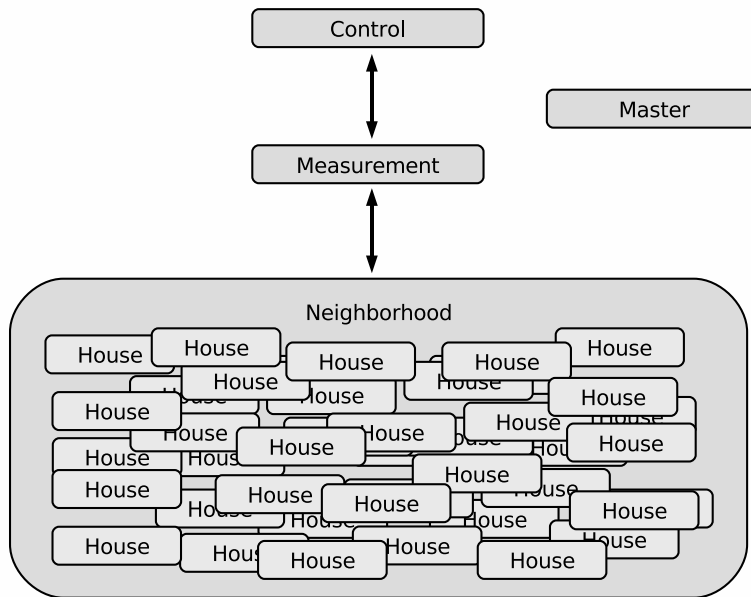
Figure 2.1: Load Group Simulation Task Diagram

The thermal model includes five states and a multitude of inputs. The states are the temperatures of the indoor air, internal walls, external walls, heater mass, and cooler mass. Table 2.1 lists the key model variables and parameters with descriptions.

As with a real HVAC system, a single speed blower fan is controlled independently from the heater or cooler compressor so that it can extract the energy from the still warm or cool compressor coils after the compressor has been shut off. For the case when the fan is on, an exponential model has been used to describe the heat transfer to the air moving from the inlet of the heater and cooler to the outlet (Equations 2.1 and 2.2). Further, Equation 2.3 accounts for thermal losses associated with the HVAC unit. The differential equation for the heater and cooler temperatures are given by Equation 2.4. The adjusted thermal input to the units ($Q_{hout/cout}$) takes the variation of thermal efficiency with temperature into account and is calculated using the method outlined in [19]. Finally, a perfect, instantaneous mixing process provides the mode for convective heat transfer between the supply air and the indoor air (Equation 2.5).

$$Q_{h2in/c2in} = k_1 (T_{hm/cm} - T_{air})(1 + k_3 [e^{\frac{-1}{k_3}} - 1]) \tag{2.1}$$

$$T_{hsup/csup} = T_{air} + (T_{hm/cm} - T_{air})(1 + e^{\frac{-1}{k_3}}) \tag{2.2}$$

$$Q_{hloss/closs} = k_2 (T_{hm/cm} - Tamb) \tag{2.3}$$

Table 2.1: Thermal Model Variables

Variable	Description
T_{air}	Temp. of Indoor Air
T_{iw}	Temp. of Internal Wall
T_{xw}	Temp. of External Wall
T_{hm}	Temp. of Heater Mass
T_{cm}	Temp. of Cooler Mass
Q_{h2in}	Heater to Inlet Air Conduction
Q_{hloss}	Heater to Ambient Conduction
Q_{hin}	Heat Input to Heater (Tonnage Rating)
q_{h2air}	Heater Supply Air and Indoor Air Convection
Q_{c2in}	Cooler to Inlet Air Conduction
Q_{closs}	Cooler to Ambient Conduction
Q_{cin}	Heat Input to Cooler (Tonnage Rating)
Q_{cout}	Adjusted Heat Input to Cooler
q_{c2air}	Cooler Supply Air and Indoor Air Convection
Q_{int}	Internal Gains to Indoor Air Conduction
q_{inf}	Infiltration Air Convection
Q_{iw2air}	Internal Walls to Indoor Air Conduction
Q_{xw2air}	External Walls to Indoor Air Conduction
Q_{xw2out}	External Walls to Outside Conduction
Q_{wincon}	Through Windows to Indoor Air Conduction
Q_{winrad}	Through Windows to Indoor Air Radiation
T_{out}	Temp. of Outside Air
T_{amb}	Temp. of Space Where Blower Resides
T_{hsup}	Temp. of Heater Supply Air
T_{csup}	Temp. of Cooler Supply Air
k_{xx}	Thermal Conductivity Constants
m_{xx}	Mass of Item XX
c_{pxx}	Specific Heat of Material XX

$$\dot{T}_{hm/cm} = \frac{Q_{hout/cout} - Q_{hloss/closs} - Q_{h2in/c2in}}{c_{ph/pc}m_{h/c}} \tag{2.4}$$

$$q_{h/c} = \frac{V_{h/c}}{m_{air}}(T_{hsup/csup} - T_{air}) \tag{2.5}$$

External walls exchange heat through conduction with the indoor air and outdoor air (Equations 2.6, 2.7, 2.8).

$$Q_{xw2air} = k_4(T_{air} - Txw) \tag{2.6}$$

$$Q_{xw2out} = k_4(T_{out} - Txw) \tag{2.7}$$

$$\dot{T}_{xw} = \frac{Q_{wx2out} + Q_{wx2out}}{c_{pxw}m_{xw}} \tag{2.8}$$

In the model, the internal wall elements are simply used to represent the thermal storage of anything solid inside the house – furniture, floors, walls, etc. Equations 2.9 and 2.10 illustrate the conduction between the indoor air and the internal walls.

$$Q_{iw2air} = k_4(T_{air} - Tiw) \tag{2.9}$$

$$\dot{T}_{iw} = \frac{Q_{iw2air}}{c_{piw}m_{iw}} \tag{2.10}$$

Windows allow a great deal of heat transfer in the forms of conduction and radiation that is very important to model. Solar radiation becomes very powerful later in the afternoon when the sun strikes the windows more directly, causing more heat input than the outdoor temperature would predict. Windows also have higher thermal conductivity than walls, and therefore allow much more conduction. Equations 2.11 and 2.12 were derived directly from [20], which elucidates methods of accounting for heat transfer through windows. The variable C_{di} changes with the time and date to account for different solar conditions throughout the year, and it is also computed in accordance with [20].

$$Q_{winrad} = \frac{1}{3600}A_{win}C_{di}C_{IAC} \tag{2.11}$$

$$Q_{wincon} = \frac{1}{3600}A_{win}C_w(T_{out} - T_{air}) \tag{2.12}$$

Infiltration is the process of unconditioned outside air leaking into the house. The leaks often occur around windows and doors, and as the insulation level of houses increases the leaks correspondingly decrease. The convective model shown in 2.13 accounts for the infiltration.

$$q_{inf} = \frac{V_{inf}}{m_{air}}(T_{out} - T_{air}) \tag{2.13}$$

Internal heat sources constitute the final input. Household objects that produce heat, usually as a by-product, constitute the internal heat sources modeled using Q_{int}. A few examples are lights, refrigerators, and computers. The computation of indoor air temperature couples all of the states together with the inputs through the windows and internal gains (Equation 2.14).

$$\dot{T}_{air} = q_{h2air} + q_{c2air} + q_{inf} + \frac{Q_{int} - Q_{iw2air} - Q_{xw2air}}{c_{pair}m_{air}} \tag{2.14}$$

Each house in the Neighborhood task contains thermostat software that exactly mimics a smart thermostat. By using the object oriented TranRunC programming style the thermostat software becomes modular and easily modified. Figure 2.2 shows the task diagram. Below is a brief description of each task:

- Master Task – Performs bookkeeping by starting all of the tasks upon initialization of the simulation.

- HVAC Com Task – Turns the simulated HVAC system on and off and relays the current indoor temperature from the simulated air.

- Heater and Cooler Control Tasks – Perform the temperature regulation calculations to determine the running state of their respective components (heater or AC).

- Coordinator Task – Ensures that the heater and cooler are on in accordance with the operating state (off, heat, cool) of the thermostat.

- Supervisor Task – Determines the current set-point temperature by implementing adjustable set-point tables that can be different for every day of the week.

- DR Com Task – Processes communications over the simulated communications network, providing the link between other agents and the super-agent.

- Goal Seeker Task – The most important task because it determines the response to communications received from the DR Com Task.

The Control Tasks are responsible for temperature regulation in the house. In the case of the most basic thermostat, it implements discrete hysteresis control. The pseudo-code below explains the algorithm for the cooler unit.

The variable 'T_s' is the current set-point temperature, and 'T_air' is the current indoor temperature. The variable 'C_a' is the anticipator constant (0.1^oF), 'C_c' and 'C_h' are the cooling and heating max bounds (0.7^oF each). In the simulation, the hysteresis band comes out to about 1.2^oF centered about the set-point temperature. Of course, the modularity of the code allows other temperature regulation schemes.
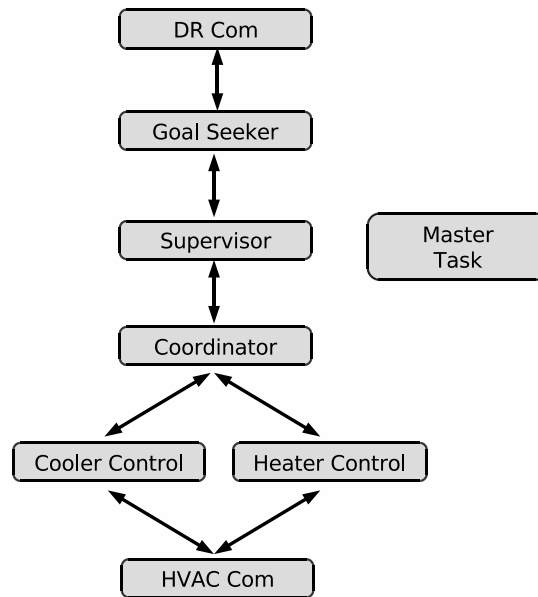
Figure 2.2: PCT Task Diagram

The Goal Seeker Task allows great flexibility in our response to demand response events. Set-point modifications are the simplest form of thermostat based demand response, and they consists of increasing or decreasing the Supervisor's current set-point value, defined in the table, by the amount specified in the message from the super-user. Further, the Goal Seeker can be easily modified to accommodate different responses to demand response events.

2.2.2   Measurement Task

The Measurement Task simulates the distribution substation in the power system architecture. It acts as a hub between the super-user and the consumers by combining the loads from the consumer tasks and relaying information about the loads to the Control Task. In this simulation, the Measurement Task simply reads the aggregate load from the Neighborhood Task on a set time interval and sends it to the Control Task, but if there were other consumers in the network it would account for their power as well.

2.2.3   Control Task

The Control Task assumes the role of the super-agent by performing systemic control for the demand response network. It receives the system power from the Measurement Task and distributes demand response messages to the agents in the network (the Neighborhood Task in this case). The Control Task controls the content of the message as well as when it is sent. In

the cases presented in Section 2.4 the messages are broadcast to every agent, but an addressing structure allows for individual or small group actions as well. Obviously, the agents and super-agent must be in agreement about what the message structure means, but the variable usually contains the event start time, end time, event type identifier, and event data fields. In the simplest case, thermostat setback is distributed in the message, but many other control variables could be used instead, e.g. duty cycle or price.

2.3   Simulation Parameter Verification

Verifying results is the major problem with simulating *only* thermostatically controlled devices. Power meters do not directly measure the HVAC, they measure the total power consumed by the house, which includes many random power sources. It is possible to instrument individual units in order to get their consumption, but this quickly becomes expensive for large groups of houses. The problem was bypassed by utilizing a widely used house/HVAC model to compare against our base case house. Following this strategy, the base parameters were tuned to closely match a single zone house simulation performed using the Multi-Zone Energy Simulation Tool (MZEST) [21], which is an extension of the California Non-Residential Engine (CNE). MZEST uses the same simulation engine as the widely used Energy-10 simulator.

Figure 2.3 shows the indoor temperature of our base simulations and MZEST simulations (indicated by CNE in the legend) under different conditions. It compares all permutations of well insulated (post 1991, California Title 24 compliant) house, poorly insulated (pre 1991, non Title 24) house, AC off all day, and AC only on from noon to 5:00pm without shutting off. The uncooled simulations (AC off all day) show that the thermal dynamics of the house simulation responds very much like the MZEST model. The cooled simulations (AC on from noon to 5:00pm) indicate a bit more deviation from the MZEST model while the AC is running. In the end, the results are very close considering the major difference in complexity between the two models.

Table 2.2: House Extents Testing

Parameter	Range	Scale
House Size (ft^2)	1661 - 3222	1x - 2x
AC Size (ton)	2 - 10	0.5x - 1.25x
Slab Construction	Y/N	

With the base case verified, the extents of the parameter range needed to be relatively simple to implement and seem reasonable compared with typical housing stock. In order to simplify implementation, the number of degrees of freedom were reduced to four, house size, insulation level, AC size, and slab construction, and a multiplicative modification strategy was used to obtain the variations in construction. The house size modifies the mass of indoor air, mass of interior and exterior walls, window area, AC size, and quantity of infiltration. Insulation level modifies the conductivity of the walls, the window area (as a proxy for R-Value), and
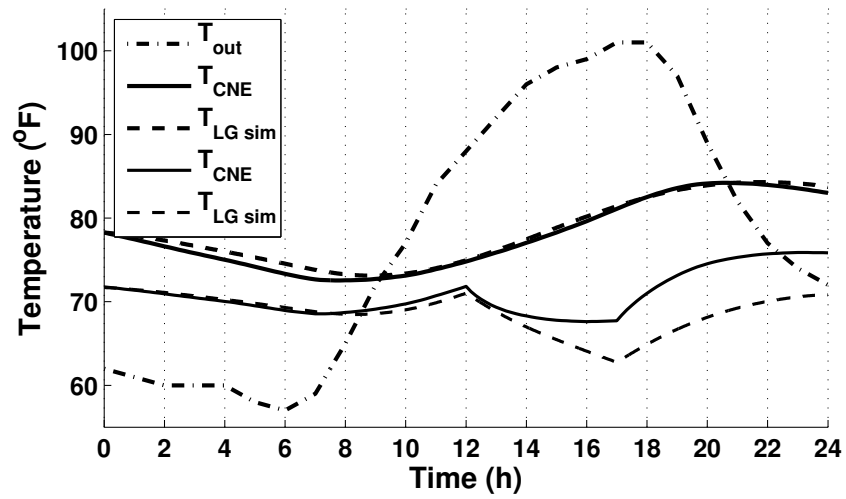
(a) Well Insulated
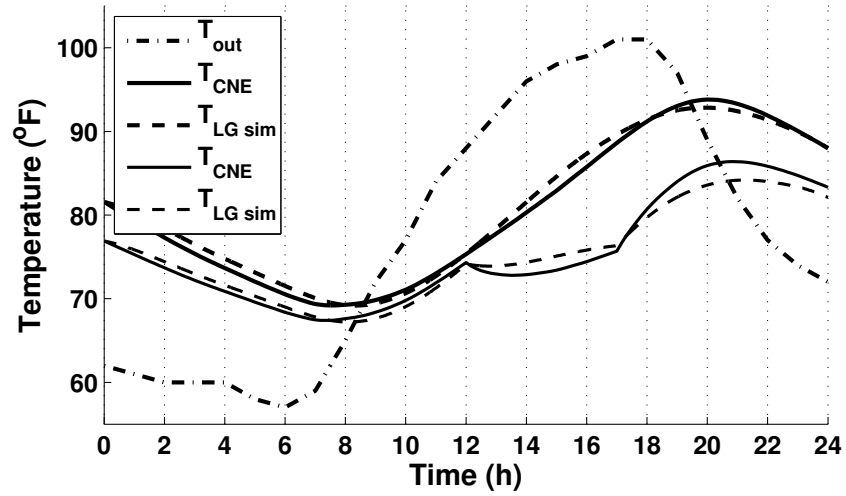
(b) Poorly Insulated

Figure 2.3: Base Simulation
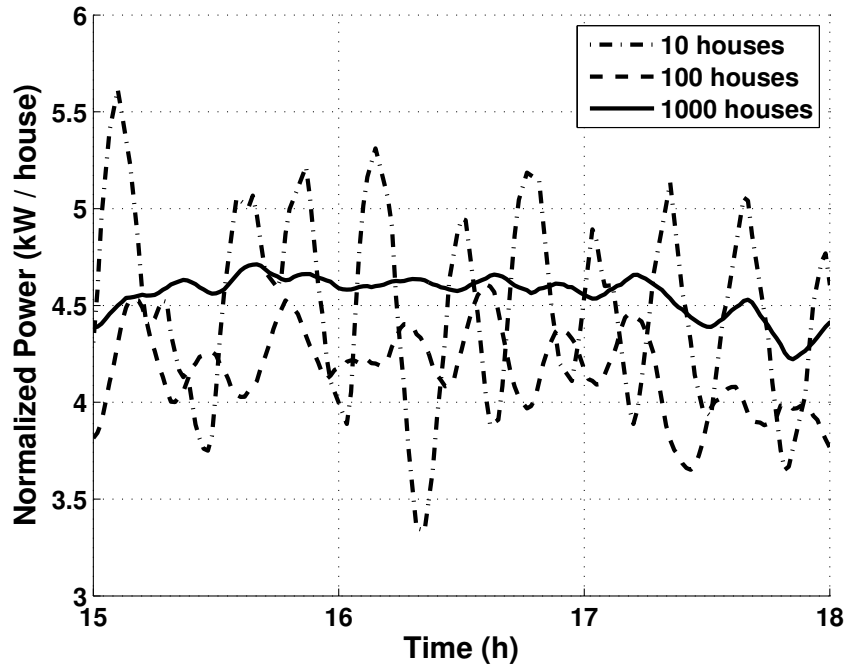
Figure 2.4: Population Testing

infiltration. The AC size is then applied on top of the house size modification because many houses of the same size do not have the same size AC. Finally, slab construction increases the mass of the internal walls to simulate the additional thermal storage associated with the slab. Table 2.2 shows the chosen range. Note that the range shown is repeated for both well and poorly insulated houses. The range was chosen independently, but it seems to fit nicely (though not exactly) with previous work done in [22].

In general, we do not intend to use the Load Group Simulation to perform exact simulations of particular neighborhoods (although it certainly could be used that way). Instead, we intend to obtain a representative sample that approximately matches average housing stock. A great deal of testing was completed to determine the best population size, and a sampling of the results can be found in Figure 2.4. The general trend is that more houses produce smoother more diversified aggregate power. Unfortunately, larger simulations take longer to complete. In the end, 1000 house simulations were found to be a good compromise between accurate data and reasonable compute times.

2.4 Example Controls Simulations

The modular nature of the simulation allows the testing of many different types of demand response controls. To illustrate the power of the simulation, we highlight a few example experi-
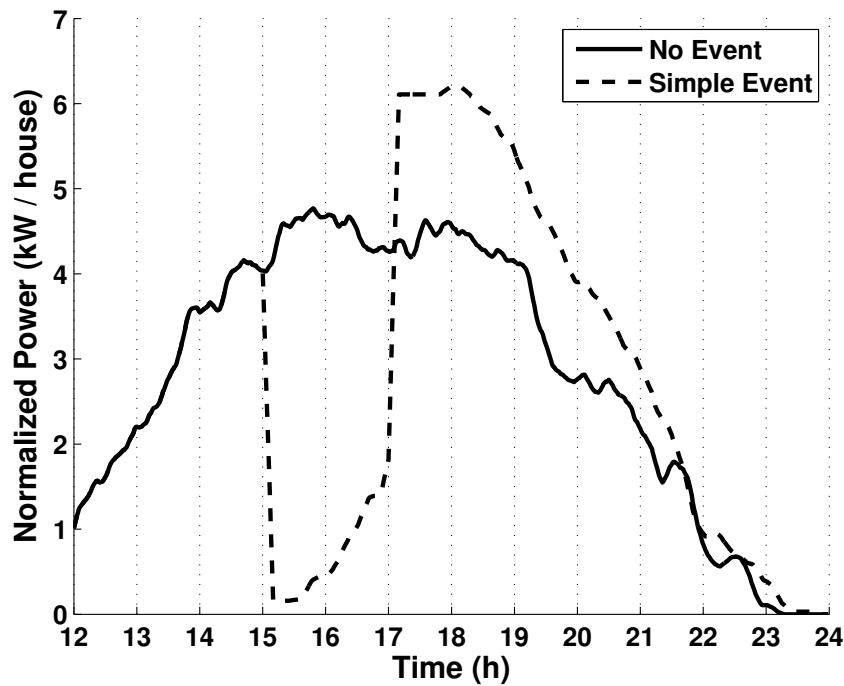
Figure 2.5: Static DR Event

ments.

2.4.1   Static Setback Demand Response

The most simple experiment is the response to static setback events. Many different setback quantities and durations could easily be studied, and quite a few were examined during the course of validating the model. A popularly talked about choice for DR events is a static setback of $4^o F$. The set-point profile of this type of event steps from the programmed value to the total setback at the beginning of the event. Scheduled set-point changes still occur during the event, but the scheduled change is modified by the setback. At the end of the event, the set-point steps back to the normal value. In this case the setback was applied from 3:00pm to 5:00pm. Figure, 2.5 shows the simulated response.

2.4.2   Payback Mitigation Demand Response

It is well known that the end of a static setback will result in a large rebound peak, or payback, as all of the ACs in the controlled area turn on simultaneously. System designers try hard to reduce the payback by shaping the event ending conditions. The Load Group Simulation allows easy experimentation with rebound mitigation techniques because the timing of messages and thermostat response can be tailored to the designers needs. We examined three distinct types
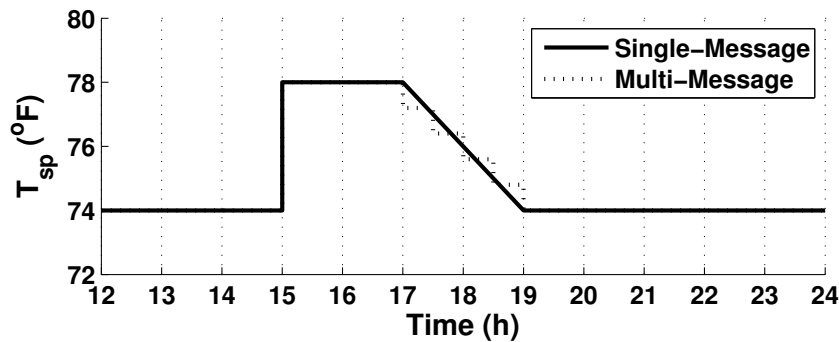
Figure 2.6: Ramped Setpoint

of rebound mitigation – random end times, multiple message based setback ramps, and single message based setback ramps.

The random end time strategy is a popularly talked about and simple method for smoothing the rebound peak. With it, a static setback DR event ends at different times for each house. Since each house begins cooling at different times, the power raises more smoothly at the end of the event.

There are a number of ways to implement the random end time. In this case we minimized the communications overhead by using a message that contains a field indicating that the end time should be randomized and each thermostat needs to compute the end time within the prescribed window. This experiment demonstrates slightly more advanced local control because the Goal Seeker Task inside each thermostat must make some decisions about how to use power.

A DR event using the ramped exit strategy starts like a simple event, with a two hour fixed setback. At the point when the simple event would have ended with a jump back to normal, the ramped strategies begin linearly changing the set-point from the maximum value back to the normal value over a time window. The researchers tested two different types of ramped exit strategies – single message ramp and multi-message ramp. Figure 2.6 shows the implementation difference between single and multi-message ramped exits.

Single-message ramped events demonstrate more advanced local control because the changing set-point (Figure 2.6) is implemented inside the thermostat software. A special DR message that specifies the start and duration of the ramp is decoded by the DR Com Task of each thermostat. From the decoded message, the Goal Seeker Task implements the ramp using a linear interpolation algorithm.

The multi-message ramp (Figure 2.6) is implemented by a series of separate DR events consisting of static set-point modifications that occur in a sequence. To achieve an exit ramp the setback in each event should be smaller than in the previous event. The transition to a new event causes a step in the set-point, and the time between the beginnings of each event cause the flat unchanging set-point. The multi-message ramp illustrates an example of systemic control because the timing and duration of the setback steps are calculated by the Control Task (from Figure 2.1).
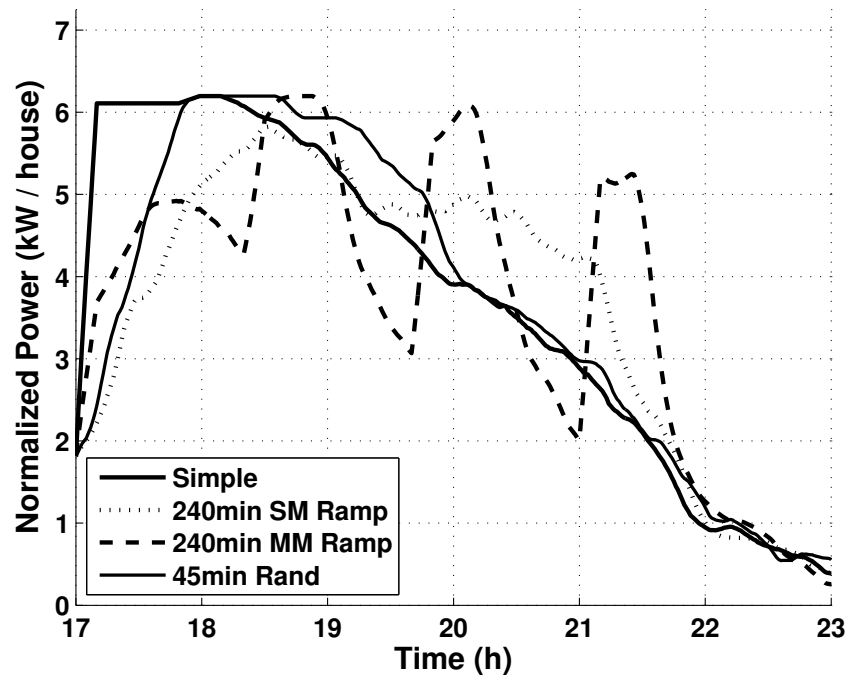
Figure 2.7: Payback Mitigation Comparision

Figure 2.7 shows characteristic results of the three types of payback mitigation techniques. Of primary interest is that the same program was able to simulate each of these cases.

2.4.3   Cost Ratio Demand Response

Cost Ratio Demand Response shifts the paradigm from directly controlling the thermostat setback to allowing the thermostat (and ultimately, customer) to decide how much energy to save by providing a framework for the autonomous use of energy price to control consumption. The first key idea is representing the energy price as a normalized quantity that allows straight-forward temporal comparison of energy costs. The second is introducing the concept of a cost tolerance that numerically illustrates a customers cost/comfort preferences. Using historical energy costs, normalized price, and a prediction for future energy consumption, the thermostat decides how best to cool (in the case of AC) the home while still meeting the cost tolerance relative to past consumption.

Cost Ratio Demand Response further flexes the muscles of the load group simulation by demonstrating advanced local control. In this case, the Heater and Cooler Control Tasks perform the price based HVAC modulation. The system still tries to maintain the set-point with hysteresis control, but it also wants to keep the incremental energy costs below the homeowner prescribed cost tolerance. Under normal conditions with the normal energy price, the set-point should be maintained as normal, but when the price increases, the system will relax the set-point in order

(a) Day 1: Temperature

(b) Day 2: Temperature

(c) Day 1: Duty Cycle

(d) Day 2: Duty Cycle

(e) Day 1: Price

(f) Day 2: Price
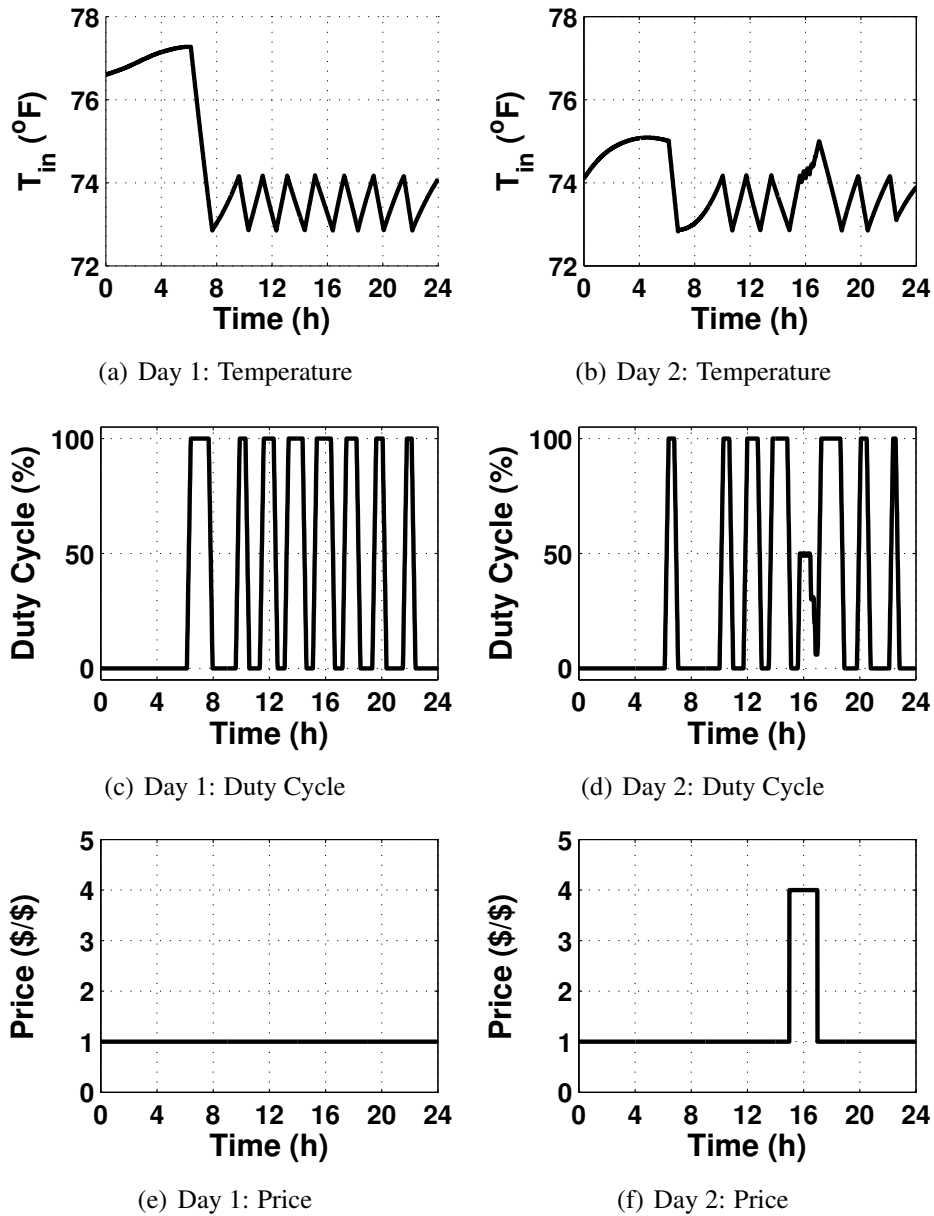
Figure 2.8: Cost Ratio Demand Response with cost tolerance of 2
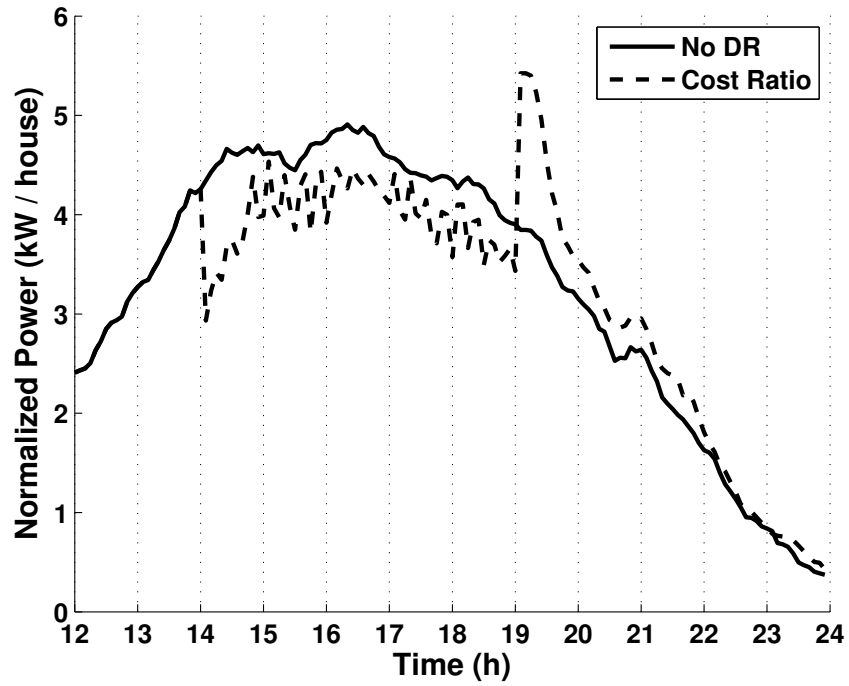
(a) Aggregate Power
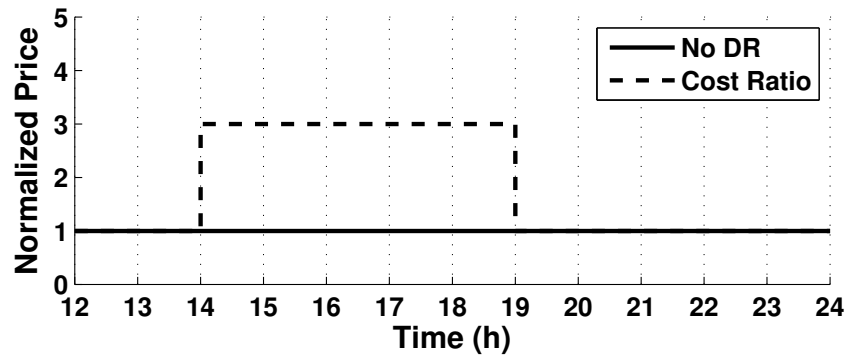
(b) CPP Price Signal

Figure 2.9: Cost Ratio Systemic Control Simulation

to remain below the cost tolerance. This requires processing historical HVAC actuation in order to maintain the total cost of energy consumed below the threshold.

Figure 2.8 shows the effect of the Cost Ratio algorithm. Energy price is normal all day on the first day in the simulation, but on the second day the price increases to four times the normal price from 3:00pm to 5:00pm. In this case the cost tolerance is 2, meaning that the homeowner is willing to pay up to two times the normal incremental energy cost. The Cost Ratio algorithm reduces the energy consumption and maintains the cost tolerance.

In order to demonstrate the systemic effect of Cost Ratio Demand Response, we used the California Critical Peak Pricing pilot study as a prototype. As outlined in [7], the critical rate is about three times greater than the normal rate (Time of Use peak rate), and it occurs from 2-7pm. Figure 2.9 illustrates Critical Peak Pricing applied to a network of thermostats using the Cost Ratio algorithm.

2.5   Conclusion

We constructed and verified a modular and extensible dynamic simulation of an advanced load management system. The model simulates the thermodynamics of a random group of thermostatically controlled devices in order to determine the characteristic aggregate power consumption subject to demand response control.

The main advantages of our simulation are fourfold. High resolution dynamic modeling yields accurate dynamic response at a small sample time. Independence between individual agents and the super-agent provides load diversity. Communications modeling allows experimentation with different levels of agent information awareness and super-agent control. Finally, modular and discrete control software allow quick changes to the local and systemic control algorithms.

Chapter 3

Low-Frequency Pulse Width Modulation Design for HVAC Compressors [1]

3.1 Introduction

The most common type of residential heating ventilation and air conditioning (HVAC) compressor is single speed, meaning it is either off or on at full power. Because of maintenance, reliability, and efficiency concerns, the compressors must cycle at relatively low frequencies. Furthermore, heat transfer dictates that the system, i.e. house, reacts slowly to the HVAC input and environmental inputs, meaning there is considerable residual in the system output (inside temperature). Traditionally, these systems use a non-linear hysteresis controller for temperature set-point following. The cycle rate is not directly defined, instead the width of the hysteresis band indirectly determines it. Hysteresis control is very simple to implement, model free, and robust. Unfortunately, it has a number of disadvantages when viewed from a modern perspective.

In this chapter, we are proposing a technique in which the single speed compressor can be treated as a variable power unit using low frequency pulse width modulation (PWM). Providing that the continuous system responds slowly, the discrete time PWM system can still be considered linear. The difficulty arises in error measurement because the states of the system could change considerably from the start of the PWM time period to the end. Consequently, the main design effort comes in appropriate filter design.

Low frequency PWM control has a number of advantages over traditional control of HVAC compressors. Firstly, any linear or non-linear control design technique producing a proportional input signal can be used to control the unit. Another advantage is that the power consumption of the unit can be explicitly controlled using tunable saturation limits, which is particularly

---

important for load management and real time energy pricing. Finally, operation of multi-stage and variable HVAC compressors becomes much easier with a proportional control signal.

Very few people have written about low frequency PWM for residential HVAC compressors. A similar idea was presented in [23] for modulation of peak load in a multi-unit facility.

3.2   Motivation

The key motivation of this advancement came from the desire to simplify load management using thermostatically controlled devices. Firstly, load management is systematic modification of the load on the electricity generation and distribution system. Load management comes in many different flavors, but one key type is *peak shaving* which is needed when the power demand peaks above the generation capacity. Furthermore, load management could (and will) be extended to provide marketable products like load following and spinning reserve.

Programmable Communicating Thermostats (PCTs) were recently proposed as a method to provide load management. PCTs were envisioned to be low cost residential thermostats with the ability to communicate with some central authority for the purpose of reducing power when needed. Using the model of the PCT, we want to extend their capability by providing more intelligence while keeping their price low.

Traditional non-linear control of thermostatically controlled devices complicates energy consumption analysis and manipulation. The inherent non-linearities make system identification and prediction difficult and unreliable. Furthermore, controlling the electricity demand is complicated with hysteresis control. Traditionally, load control has been provided by communicating switches placed directly on the compressor to deny it power (called Direct Load Control or DLC), bypassing the temperature controller altogether as in [2].

3.3   Theoretical Basis

Pulse width modulation is a common technique for obtaining quasi-continuous output from an on/off type actuator. In its most common form, a PWM signal is a series of pulses produced on a fixed period (T). The on time (T_{on}) of the pulses varies between zero and full period. The varying pulse produces a variable output from the on/off actuator. If possible, a separate signal is used to control the direction of the actuator. PWM is generally specified in percent of period or duty ratio given by Equation 3.1.

$$\phi(kT) = \left\{ \begin{array}{ll} \frac{t_{on}}{T} & \text{for positive actuation} \\ -\frac{t_{on}}{T} & \text{for negative actuation} \end{array} \right. \tag{3.1}$$

The input to the system can be represented by the following:

$$u(t) = \begin{cases} U_{max} sgn(\phi) & \text{for } kT \leq t < kT + |\phi(kT)|T \\ 0 & \text{for } t \geq kT + |\phi(kT)|T \end{cases} \tag{3.2}$$

Now consider the linear time invariant continuous system represented by the standard state space formulation with n states, p inputs, and m outputs.

$$\begin{aligned} \tfrac{d}{dt}x &= Ax(t) + Bu(t) & A \in \Re^{nxn} & \quad B \in \Re^{nxp} \\ y &= Cx(t) + Du(t) & C \in \Re^{mxn} & \quad D \in \Re^{mxp} \end{aligned} \tag{3.3}$$

Let's say that one of the inputs to the system is given in terms of PWM. The system response to the discontinuous input $\phi(k)$ can be represented as in Equation 3.4.

$$x(t) = \begin{cases} e^{A(t-kT)}x(kT) + \int_{kT}^{t} e^{A(t-\tau)}BU_{max}sgn(\phi(k))d\tau \\ \qquad \text{for } kT < t \leq kT + |\phi(k)|T \\ e^{A(t-kT-|\phi(k)|T)}x(kT - |\phi(k)|T) \\ \qquad \text{for } t > kT + |\phi(k)|T \end{cases} \tag{3.4}$$

By discretizing the system at the PWM sample time T, a single non-linear equation describes the response to the PWM input at the instances $t = kT$.

$$\begin{aligned} x((k+1)T) &= A_d x(kT) + h(kT, u) \\ A_d &= e^{AT} \\ h(kT, \quad u) &= e^{AT}(I - e^{-AT|\phi(k)|})A^{-1}BU_{max}sgn(\phi(k)) \end{aligned} \tag{3.5}$$

Linearizing the non-linear function $h(kT, u)$ yields Equation 3.6.

$$\begin{aligned} x(k+1) &= A_d x(k) + \hat{B}_d \phi(k) \\ A_d &= e^{AT} \\ \hat{B}_d &= (e^{AT} - I)A^{-1}BU_{max} \end{aligned} \tag{3.6}$$

Traditionally, Equation 3.6 is only considered a valid approximation when the PWM sample time (T) is small, but that is not entirely true. Actually, the matrix quantity AT must be small, giving rise to the possibility that the system matrix A is small and the sample time T is large. Note that similar analysis of PWM systems can be found in [24, 25] among others.

The resulting discrete time linear system is only so useful though. The value of the states and output at instant Tk, are just that, and the value of the states and output between $T(k-1)$ and Tk are not considered. However, given the long PWM period, the values between discrete instances are still important. This gives rise to a filter design problem. What is the best output or state filter considering controller performance and objectives?

3.4   System Design

For the design of the low frequency PWM controller, we will use two different house models. For the final design, we apply the controller, and do final tuning, on a relatively complicated house model used for load management experimentation. We previously outlined the model in Chapter 2. For rough design, we consider the following first order continuous time multi-input single output system in state space form (Equation 3.7).

$$
\begin{aligned}
\dot{x}(t) &= Ax(t) + Bu(t) \\
y(t) &= Cx(t) + Du(t) \\
A &= [-4E-4] \quad B = [4E-4, -2.5E-6] \\
C &= [1] \qquad\quad D = [0,0] \\
u(t) &= [T_{out}, P_{ac}]^T \\
y(t) &= T_{in} \\
M &= 4000
\end{aligned}
\tag{3.7}
$$

The inputs to the system are the outside temperature and the instantaneous power from the single-speed HVAC compressor, $[T_{out}, P_{ac}]$. In this model, we approximate the compressor as producing full power ($M = 4000$) instantly. Further, the compressor produces the identical power regardless of outside temperature. The output of the system is the indoor temperature, T_{in}. The goal is for the output to track a set-point temperature, $y_{ref} = T_{sp}$.

For reliability, maintenance, and efficiency reasons, HVAC compressors should not be cycled too often, 4 to 6 times an hour. Considering this against the desires to have accurate inside temperature reference tracking and load management, we decided on a fifteen minute PWM period. Using this sample rate ($T = 900s$), we can discretize the continuous time system using Equations 3.5 and 3.6. Equation 3.8 describes the linear approximation of the system with the the variable u_{dr} representing the duty ratio of the HVAC compressor on the interval $[0,1]$. Without loss of generality, we have assumed that the first input, outside temperature (T_{out}), fluctuates slowly with respect to the sampling interval. (If the outside temperature were to fluctuate quickly, $\hat{B}_d(1,1)$ would be different and the signal would need to be appropriately filtered. However, no changes would be made to the PWM input or its effect on the system.) Figure 3.1 shows the simulation response to an open loop PWM signal for the continuous time system, sampled non-linear system, and sampled linear approximation.

$$
\begin{aligned}
x(k+1) &= A_d x(k) + \hat{B}_d \phi(k) \\
A_d &= [0.69768] \\
\hat{B}_d &= [0.30232, -7.5581] \\
\phi(k) &= [T_{out}, u_{dr}]^T
\end{aligned}
\tag{3.8}
$$

The previous analysis was simply to show that given a very small A matrix, the low fre-
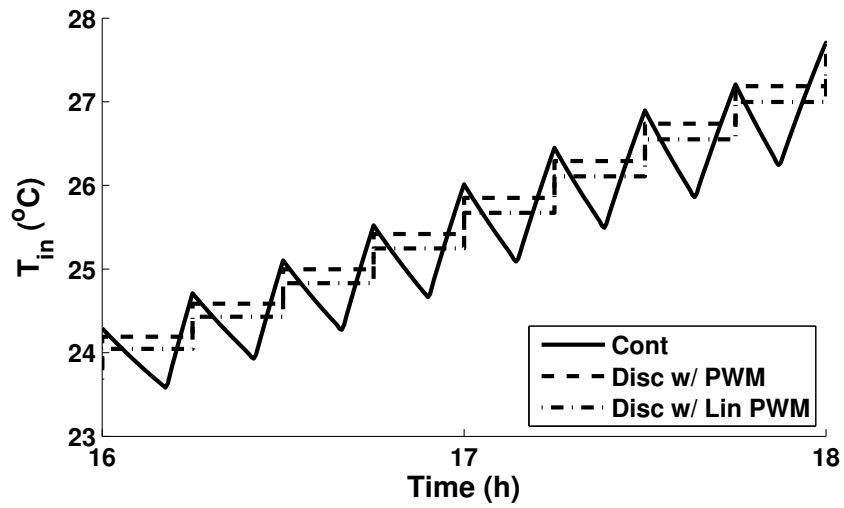
(a) Output
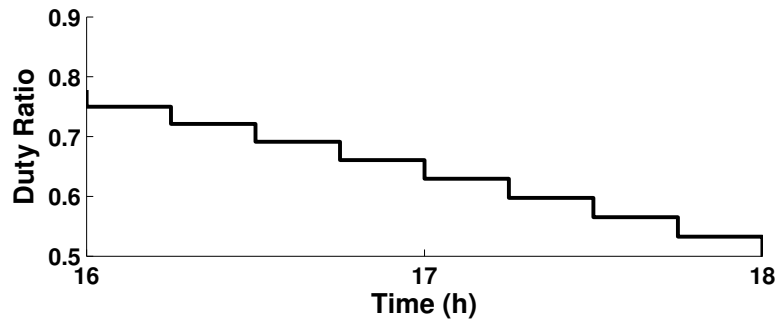
(b) Input

Figure 3.1: Discretization of Continuous System

quency PWM system becomes approximately linear.

Low frequency PWM requires a couple of changes from its high frequency counterpart. In the normal case of high frequency PWM, the system acts as a low pass filter attenuating the high frequency changing input. In the case of low frequency PWM, there is a large residual that necessitates a synchronous filter on the feedback path that operates on the (approximately) continuous signal but is downsampled at the controller sample period. Note that filtering with a continuous time linear filter does not change the linearity of the PWM system as long as the filtered system matrix remains small. Furthermore, the PWM sample rate is so slow that in order to obtain good controller performance, the controller must run at the same rate as the PWM. With high frequency PWM, the PWM sample rate and controller sample rate can be chosen somewhat independently.

3.4.1 Output Filter Design

It is certainly possible to design an analog filter to meet our requirements, but considering the long time scales, the electronic components would need to be very large and expensive. Therefore we will leave exact linearizability behind and restrict ourselves to approximation with digital filters. The filter sample rate (T_s) must be small compared to the PWM sample rate (T). Further, the filter should be synchronized with the PWM sample rate to ensure that the control receives the most up to date information. Hence, the filter sample rate must be chosen so that the PWM sample rate is an integer multiple of it.

Our goal for filtering is to attenuate the residual caused by the low frequency pulsed input. Explicitly, the synchronous filter should minimize the error over the previous time step (Equation 3.9). The synchronous filter signal is given by $y_f(kT)$, and the continuous time signal is given by $y(t)$.

$$\| e_f \| = \left(\sum_{k=0}^{\infty} \sum_{i=0}^{T/T_s-1} \left(y(Tk - T_s i) - y_f(Tk) \right)^2 \right)^{1/2} \tag{3.9}$$

Note that the output signal will be fluctuating in response to the PWM input at the the known PWM frequency. The PWM forced fluctuation is exactly the signal we want to attenuate. Therefore, we need a low pass filter with a cutoff frequency below the PWM frequency.

There are two main classes of filters – infinite impulse response (IIR) and finite impulse response (FIR). The Butterworth filter is an example of a simple IIR filter. A discrete time Butterworth filter is designed using two parameters, cutoff frequency (ω_n) and system order (n). The Boxcar filter is an example of a simple FIR filter. We performed a parametric study of the Butterworh filter with varying parameters and the Boxcar filter with varying orders. Each of the filters is sampled at 5 second sample rate and downsampled to the PWM sample rate. Table 3.1 shows the norm error for each filter tested.

As expected, the choice of error norms played a large part in selection of the filter. The Butterworth filter is very sensitive to the choice of cutoff frequency, and it seemed to perform best

Table 3.1: Parametric Filter Study

Type	n	ω_n	$\parallel e_f \parallel$
butter	3	0.0007937	99.721
butter	3	0.0009259	92.55
butter	3	0.0011111	84.306
butter	3	0.0013889	75.853
butter	3	0.0018519	69.395
butter	3	0.0027778	67.069
butter	3	0.0055556	79.135
butter	1	0.0013889	65.648
butter	2	0.0013889	67.298
butter	4	0.0013889	87.528
butter	5	0.0013889	99.994
butter	6	0.0013889	111.29
boxcar	30	-	102.76
boxcar	60	-	84.633
boxcar	90	-	73.561
boxcar	120	-	68.748
boxcar	150	-	65.659
boxcar	180	-	**64.042**

at around $\omega_n = 0.0028 rad/s$. With a fixed cutoff frequency, the Butterworth filter performed worse as the order increased, mainly because of the increasing delay. The high order ($n = 180$) boxcar filter yielded the minimum norm error, and this is not surprising given the specification of the norm error. Another great advantage of this filter is ease of computation. It only requires storage of one variable.

3.4.2   Controller Design

The filtered system is approximately linear, and a controller can be designed using any linear (or non-linear) design technique. We chose to design a proportional plus integral (PI) controller. This type of controller is quite simple to design and program (Equation 3.10).

$$
\begin{aligned}
e(k) &= y_{ref}(k) - y_f(k) \\
e_{int}(k) &= e_{int}(k-1) + e(k)T \\
P(k) &= k_p e(k) + k_i e_{int}(k)
\end{aligned}
\tag{3.10}
$$

The only trouble is how to deal with saturation. An HVAC compressor represents a single sided input, i.e. cooling or heating. Additionally, residential HVAC systems rarely control both the heating and the cooling systems at the same time. For instance, in cooling mode the temperature is allowed to drop well below the set-point. In these situations, the integrator in the

(a) Output

(b) Input
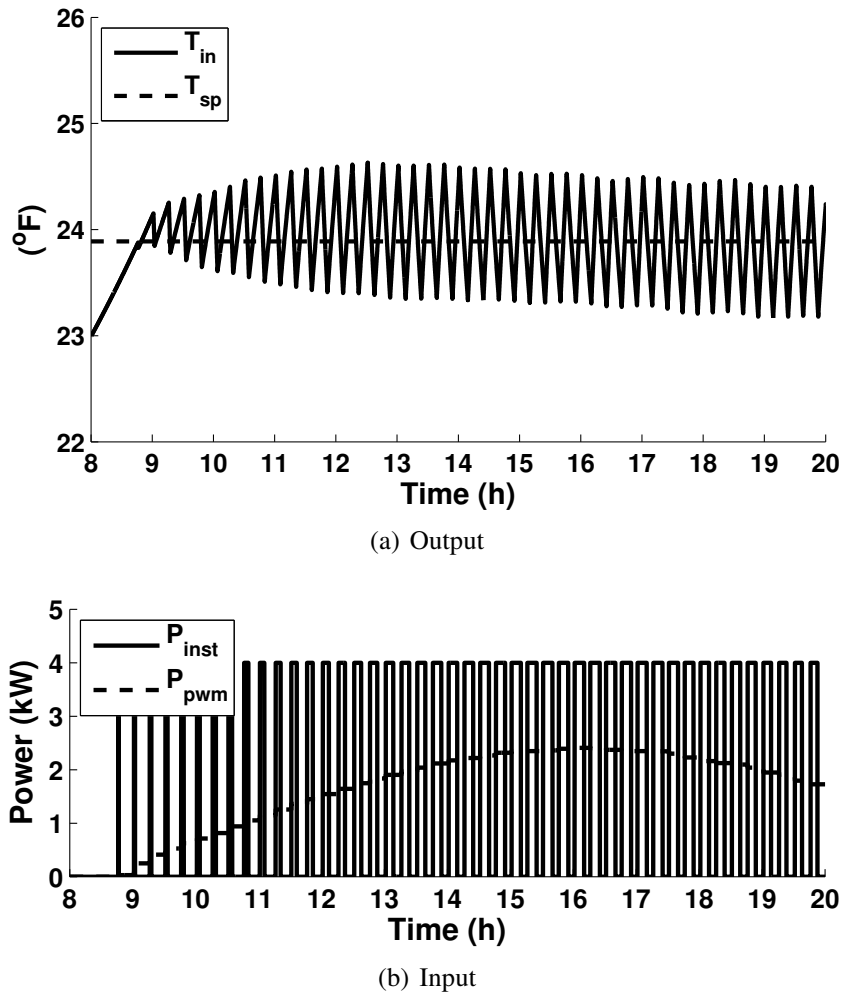
Figure 3.2: PWM Control of First Order System

PI controller winds up, creating poor performance. This necessitates an anti-windup mechanism, as illustrated with Equation 3.11.

$$e_{int}(k) = \begin{cases} \frac{P_{max} - k_p e(k)}{k_i} & P(k) > P_{max} \\ \frac{P_{min} - k_p e(k)}{k_i} & P(k) < P_{min} \end{cases} \qquad (3.11)$$

In order to tune the controller, we used an iterative process on the first order system to get the order of magnitude of the gains. The final tuning was completed, iteratively, on the more complicated system. Figure 3.2 shows the results of a PI control on the first order system. In this and all of the first order simulation results, the outside temperature is time varying with a sine wave that peaks at 32.2^oC at 4:00pm and has a minimum value of 21.1^oC at 4:00am.

(a) Output

(b) Input
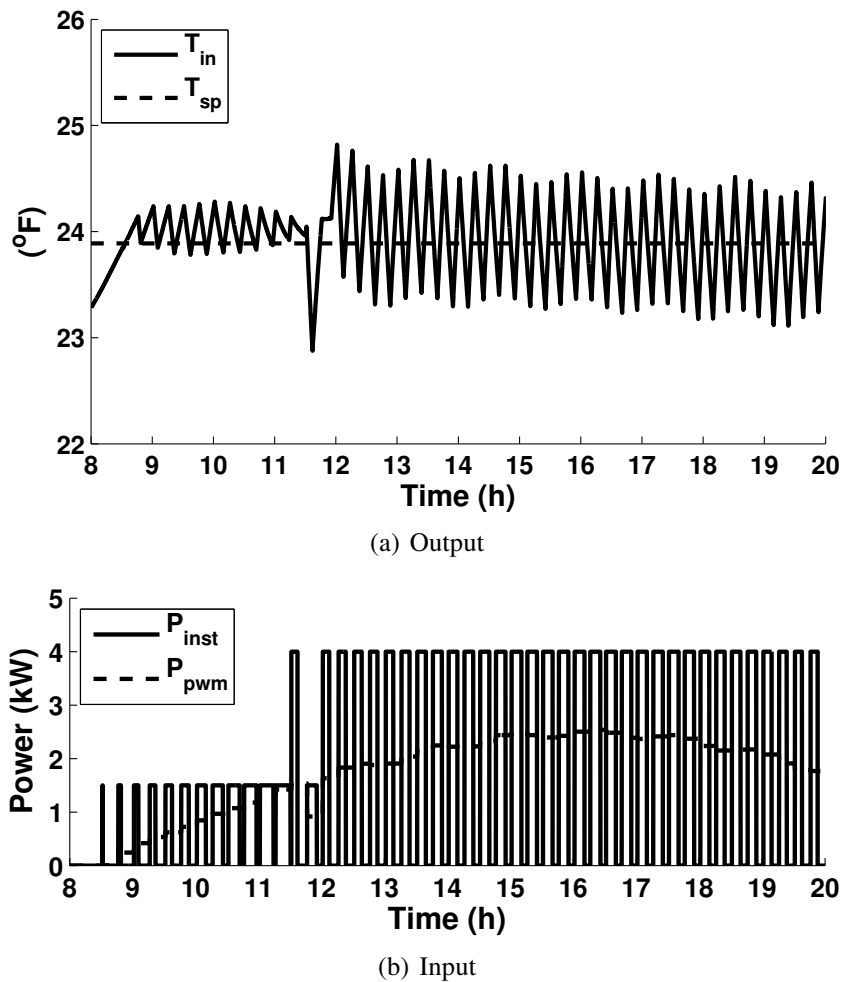
Figure 3.3: PWM Control with 2 Stage Compressor

3.4.3   On/Off Time Limits

For maintenance and reliability reasons, typical HVAC compressors need to be in the on state or off state for a certain amount of time before switching. This requirement is typically accomplished using on/off timers directly on the compressor unit. These traditional cycling timers will still work with low frequency PWM actuated units, but the control will be slightly biased as a result. Further, low frequency PWM can account for these timers directly by making use of slightly more complicated saturation guidelines that round the low and high PWM to ensure the on/off times are met.

3.4.4 Multi Stage Units

Multi-stage compressors have been commonplace for years. Without going into the mechanical design of these units, multi-stage compressors essentially have more than one output power that can be switched between. In general, they offer pretty significant efficiency advantages over traditional single-stage units.

Unfortunately, operation of multi-stage units using tradition hysteresis control is cumbersome at best. Using hysteresis control, the controller has no way of judging how much power is needed to control the system. Some additional rate detection, etc. needs to be implemented in order to make hysteresis control feasible.

Control of a multi-stage unit using low frequency PWM is very simply accomplished. When creating the PWM duty ratio from the control calculation, $P(k+1)$, the largest stage with power less than $P(k+1)$ should be used. Equation 3.12 illustrates the calculation for a two stage unit. Figure 3.3 shows simulation of the same first order system as in Figure 3.2, with the identical controller (gains included), but with a two stage compressor.

$$
\begin{aligned}
P_{cur}(k+1) &= \left\{ \begin{array}{ll} P_1 & 0 < P(k+1) \leq P_1 \\ P_2 & P_1 < P(k+1) \leq P_2 \end{array} \right. \\
u_{dr}(k+1) &= P(k+1)/P_{cur}(k+1)
\end{aligned} \tag{3.12}
$$

3.4.5 Tunable Saturation

Control of HVAC power consumption usually takes the form of radio operated direct load control (DLC) switches attached to the compressor, as in [2]. A DLC switch bypasses the temperature controller and shuts off the compressor for a specified interval of time. In general, the switch does not consider the cycling characteristics of the unit and simply shuts off when commanded. This results in dramatically different responses for different compressors, ranging from almost no change at all if the natural cycling period is below the switch setting, to dramatic changes when the natural cycling period is much greater. Some DLC manufactures have tried to remedy this problem by introducing adaptive switches that reduce the power proportionally, as in [4].

One of the key advantages of PWM actuation is direct control over power consumption using tunable saturation. Figure 3.4 illustrates this concept using the same first order system as used in the previous examples. Here, the saturation level was statically set at 50%, but it would be straightforward to extend the system with a higher level controller manipulating the saturation in real time.
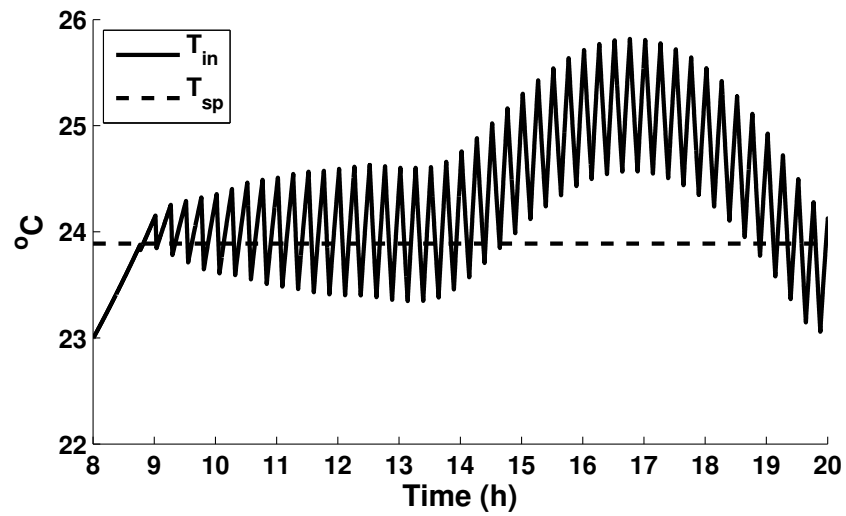
(a) Output
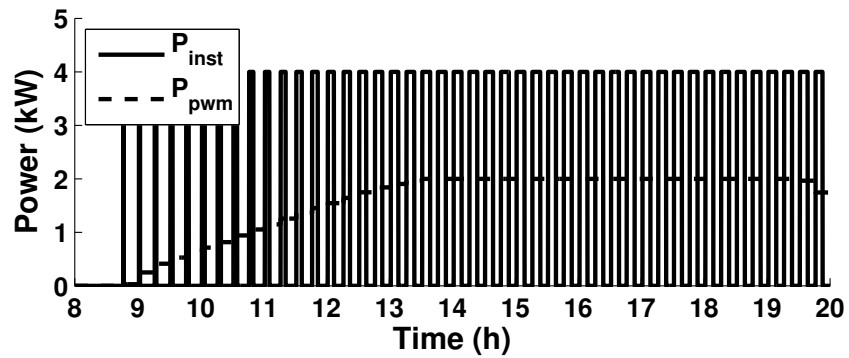
(b) Input

Figure 3.4: PWM Control with Tunable Saturation
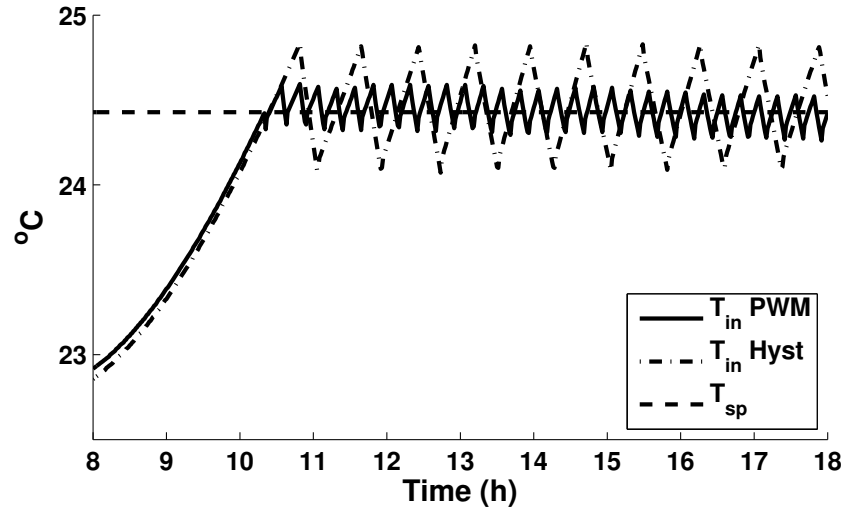
(a) Output
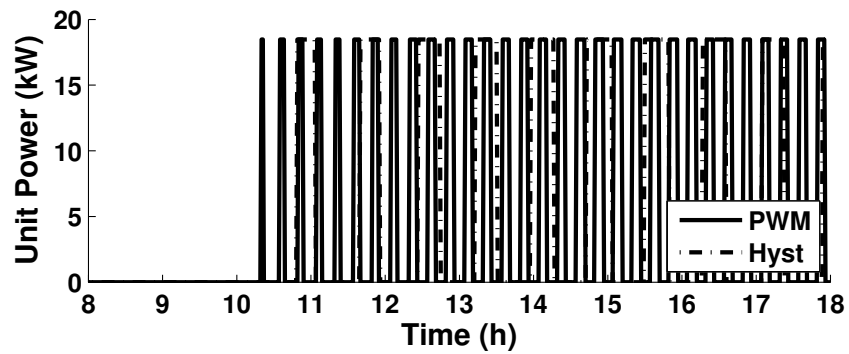
(b) Input
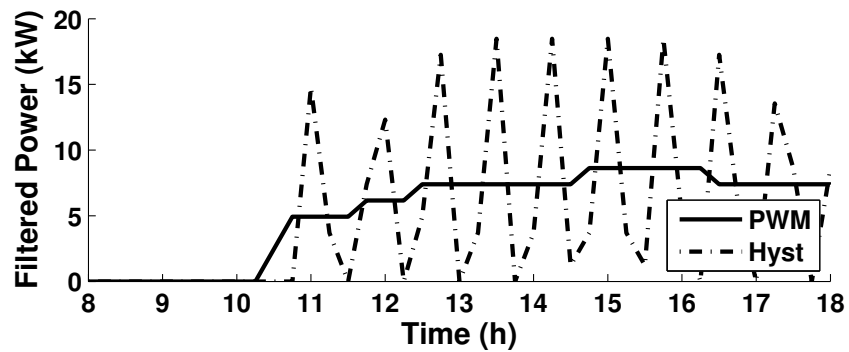
(c) Filtered Input

Figure 3.5: PWM Control of Full Simulation

3.5 Results

In order to verify the performance of the PWM actuated system, we performed experiments using a more complicated model that better simulates the dynamics of a house. We previously described this model in Chapter 2. Two tests were performed on an identical house model under identical environmental conditions, i.e. outside temperature and solar radiation from a hot summer day in Fresno California. The first test controlled the HVAC using a hysteresis control, and the second test used the PI controller with PWM actuation as designed previously. Figure 3.5 illustrates the results of the two controllers tracking an identical set-point temperature.

The first obvious difference between the two controller types is the difference in error between the on and off peaks – the error band. For the PWM actuation, the error band is "set" by the PI-controller *and* the choice of PWM frequency. If it were hotter outside, causing the temperature in the house to rise more quickly, the unit would cycle at the same rate because the PWM frequency is fixed, but the error band would be larger. At lower outside temperatures, the frequency would still be the same and the band would be smaller. Alternatively, the error band for the hysteresis controller is set *directly* by the controller *only*. Regardless of the outside temperature, the error band will always be the same (if the compressor has the capacity to cool the house), but the cycling frequency fluctuates. If it is hotter outside, the unit cycles more quickly, and cooler temperatures result in less frequent cycling. This fluctuating cycling rate makes prediction and analysis difficult because of the lack of time consistency.

The linearizing quality is the main advantage of PWM actuation. Figure 4.1(c) plots the filtered power using a boxcar filter over a fifteen minute interval similar to the one used in the controller. This treatment of the system input clearly shows the discontinuous hysteresis control and the smooth PI control.

3.6 Conclusion

We demonstrated, through simulation, how low frequency PWM simplifies control of multi-stage compressors. The compressor stage is stepped based on a simple set of PWM rules. The simplicity advantage extends for variable-speed compressors a well. It is in fact simpler as the compressor speed is determined directly via the controller.

Low frequency PWM control dramatically simplifies the analysis and control of HVAC compressors when viewed through the lens of load management. Hysteresis control results in a difficult to analyse highly non-linear system. System identification is difficult, meaning that state prediction is unreliable. PWM control linearizes the system, simplifying not only controller analysis but system identification and prediction as well. In later chapters, we take advantage of these key properties of the PWM actuated system.

Assume that energy consumption is roughly proportional to the compressor on-time. This is mostly true except that the efficiency (and therefore power) varies somewhat with outdoor temperature. With low frequency PWM, the control signal is calculated at the start of the PWM

period, and therefore the energy consumption for the period is known in advance. This results in the ability to artificially limit the power consumption using a simple tunable saturation variable. This is a major advantage for PWM actuation that, when coupled with the linearizing qualities, will enable more intelligent load management systems than could be designed using hysteresis control. The next few chapters are aimed exactly at this.

Chapter 4

PWM Synchronization for Intelligent Agent Scarce Resource Auction [1]

4.1 Introduction

The goal of this work is to develop a framework for distributing shared scarce resources amongst intelligent autonomous agents. In particular we are interested in modulating the total power consumption of a group of independent agents responsible for residential HVAC operation. Our system is hierarchical, consisting of independent home agents responsible for comfort and a super-agent responsible for the power regulation. The coupling between the home agents and the super-agent occurs through shared communications.

In this chapter, I propose a market based approach to load management using a scarce resource auction. Each home agent knows a demand function that defines its price versus demand desires, but in order to effectively operate in the auction, it must be able to predict its future power consumption. By synchronizing the low frequency PWM of the home agents with the auction time windows, the complexity for learning and predicting local power consumption is drastically reduced. Additionally, adjusting the power consumption in response to the time varying price is simplified.

Previous researchers have also proposed market based approaches to load management, [26, 27, 28]. We build on their work by presenting an inexpensive autonomous method for information poor residential HVAC systems.

---

[1]Reprinted, with permission, from "PWM SYNCHRONIZATION FOR INTELLIGENT AGENT SCARCE RESOURCE AUCTION," by William J. Burke, David M. Auslander, Proceedings of the ASME 2009 Dynamic Systems and Control Conference, Paper Number DSCC2009-2689

4.2   Motivation

Thermostatically controlled devices are well suited for load management because they are ubiquitous and heavy electricity consumers. Specifically, air conditioning contributes heavily to the summer peak loads throughout the United States. Programmable Communicating Thermostats (PCTs) were recently proposed as a method to provide load management. PCTs were envisioned to be low cost residential thermostats with the ability to communicate with some central authority for the purpose of reducing power when needed. Using the model of the PCT, we want to extend their capability by providing more intelligence while keeping their price low.

Traditional non-linear control of thermostatically controlled devices complicates energy consumption analysis and manipulation. The inherent non-linearities make system identification and prediction difficult and unreliable. Furthermore, direct load control is crude and complicated with hysteresis control. Traditionally, it has been provided by communicating switches placed directly on the compressor to deny it power, bypassing the temperature controller altogether as in [2].

With the coming advances of the "Smart Grid" communications and control are taking a prominent role in load management. By distributing inexpensive communicating thermostats and developing appropriate controls, effective and inexpensive load management could be implemented with little effect on personal comfort.

4.3   Scarce Resource Auction

The market operates using the Tâtonnement Process [29] – the auctioneer (operations agent) suggests a price, and the bidders (consumption agents) respond with their expected average power during that period conditioned on the suggested price. Starting with the lowest allowed price, the auctioneer raises the price through successive suggestions until the expected power meets or exceeds the objective. By starting bidding at the lowest allowed price, this mechanism is an ascending price auction with the price suggestions never decreasing. The final price is set once the objective is achieved.

The market operates only sporadically with the notion that there are "normal" periods and "control" periods. When closed, the cost of the resource (energy) is fixed at the normal price. When the market is open, price control periods last 15 minutes and are consecutive. The price during every period is determined during the 15 minutes prior to the start of the period.

Throughout this paper we will abstract exact price away and consider the price ratio instead. The price ratio is the ratio of the current price to the normal price (which does not have to have the same dollar value for every period, it is simply "normal" for that period). Therefore, the normal price has a price ratio of 1, and higher prices have price ratios greater than one.

During the "normal" period, the agent can consume as much resource as they like. Alternatively, during the "control" periods, the agent cannot consume more energy than bid under fear of a heavy penalty. The notion of this penalty will be kept vague at this point, and we will just

assume that nobody invokes it. The exact penalty is more a question of market design and is left for related works.

4.4 Auction Synchronized Home Agents

The home agents interacting within this market could take many forms, but we have a few design constraints. Primarily, we want to keep the cost spent on computing resources for each agent low, which means very little additional sensing and an inexpensive (i.e. slow) processor. Two-way communications is the only luxury we have. Therefore, we need a robustly simple design that includes the following elements:

- Temperature control to maintain, or seek, comfort.

- On-line system identification and prediction to enable power demand bidding.

- Computable demand function to direct the cost to comfort decision making.

Our solution to this design criteria is synchronized PWM, and in the following sub-sections we elucidate why synchronized PWM is such an elegant solution.

4.4.1 PWM Control and Synchronization

Traditionally, temperature control using HVAC systems is accomplished with non-linear hysteresis controllers. Recently we developed an alternative to this that treats the unit as a proportional actuator using low frequency PWM as in Chapter 3. With low frequency PWM the on/off HVAC unit is operated proportionally as a fraction of a long period. As an example, using our PWM period of 15 minutes, a 30% duty cycle would result in the unit being on for only 5 minutes of the period. Low frequency PWM enables the use of linear control laws for temperature regulation. We use a simple PI controller, but any linear (or non-linear) controller with proportional output would work. Figure 4.1 shows the difference between low frequency PWM and traditional HVAC controls.

The major advantage of using low frequency PWM to operate the HVAC system is the ability to synchronize the actuation period (PWM period) with the auction period. Synchronization offers up two key advantages. First, the power consumed during the next auction period is known at the start of the period making power limiting very simple using tunable saturation. Second, prediction of power consumption for the auction reduces to a single step look ahead.

The only drawback to synchronized PWM is the possible load diversity issues. In general, the power grid relies on all of the different loads to operate asynchronously, and in particular, it relies on HVAC systems to operate at random times. If all of the HVAC systems turned on at the same time, the peak power would be tremendous, but with them operating, more-or-less, randomly, a much lower peak is maintained. This is referred to as load diversity.
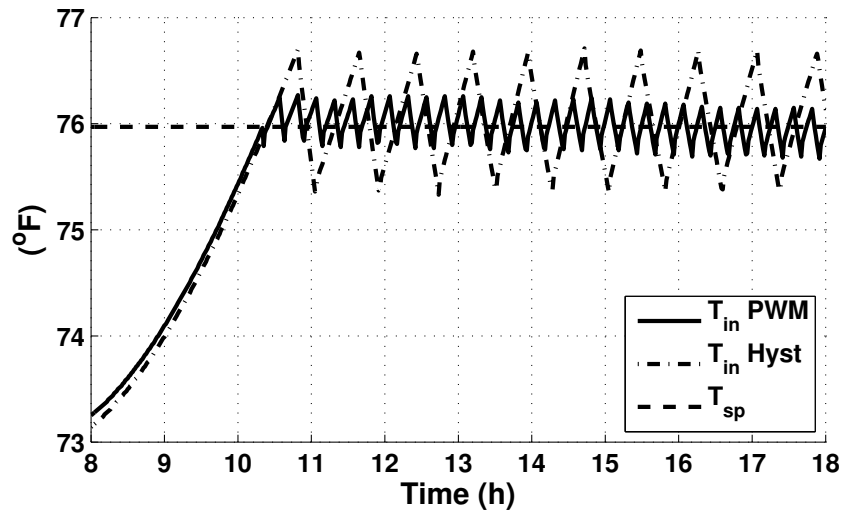
(a) Output
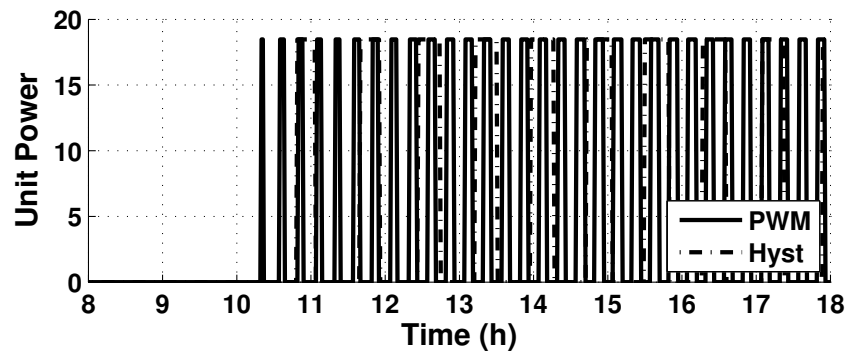
(b) Input
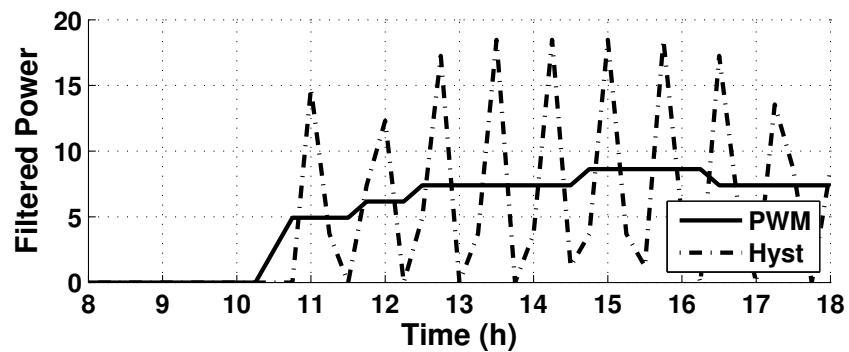
(c) Filtered Input

Figure 4.1: Simulation of the same house/HVAC under hysteresis control and low frequency PWM with PI control

With synchronized PWM, the load diversity must be forced upon the system. There are at least two ways to handle this. The first way randomizes the start time of the on-pulse for every new period. While simple to implement, this method suffers from random controller bias and reduced performance. The second way randomizes the middle time of the on-pulse at the initialization of the controller. From initialization forward, the middle time of every on-pulse is at the same time in the period. If the pulse is so big that the specified middle time would result in the pulse extending outside of the PWM period, the middle time is shifted to obey the period boundaries for that pulse. Luckily, PWM theory does not really care when during the period the on-pulse occurs as long as it occurs at the same time each period. Therefore, this method results in very little controller bias and good performance.

4.4.2 System Identification and Prediction

The primary desire for system identification and prediction is to estimate power consumption correctly. With the simplest view, we need to estimate power given our available information – inside temperature, set-point temperature, previous power consumption, and outside temperature. With our design goal of reduced cost and minimized additional sensing, outside temperature and previous power consumption become slightly more difficult to come by. We get the outside temperature through communications with the super-agent. Previous power consumption is estimated from the duty cycle and rudimentary knowledge about the HVAC compressor – SEER rating and compressor size.

Synchronized PWM makes this task easier in two ways. First, the linear control turns the previously on/off operation of the HVAC system into a proportional, nearly linear, signal sampled at the PWM period. Second, since the sample period is synchronized with the auction period, one-step look ahead is all that is required to participate in the auction.

Unfortunately, the system is still not *easy* to identify. The PWM does not totally linearize the system, because saturation is still present. The power consumption of the HVAC system will never be less than zero or greater than some maximum value. Further, the actual system is of huge order and has many unmodeled inputs, like solar radiation.

Traditional, on-line linear least squares resulted in erratic and unstable performance. A non-linear least-squares like identification system resulted in much better performance. The non-linear system is shown in Equation 4.1 with parameter vectors ψ and θ. The index, j, of parameter ψ is chosen based on the time of day, with each 15 minute interval receiving its own entry in the vector. Further, saturation is applied to the estimated power to keep the signal greater than zero.

$$\hat{P}(k+1) = \psi_j + (T_{out}(k+1) - \theta_2)\theta_1 + (T_{in}(k) - T_{sp}(k))\theta_3 \qquad (4.1)$$

The parameter update law is shown in Equation 4.2 with parameter convergence variable vector
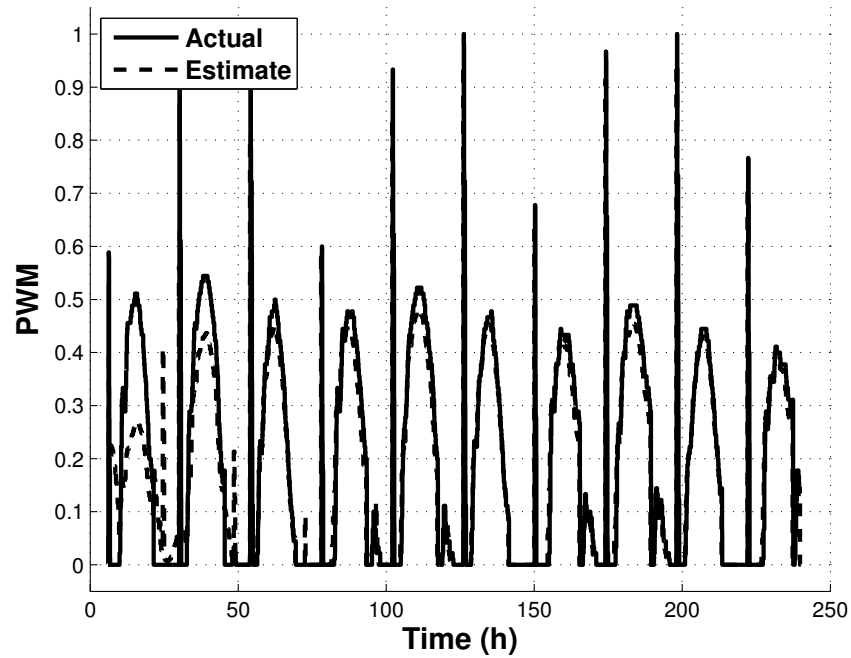
(a) PWM
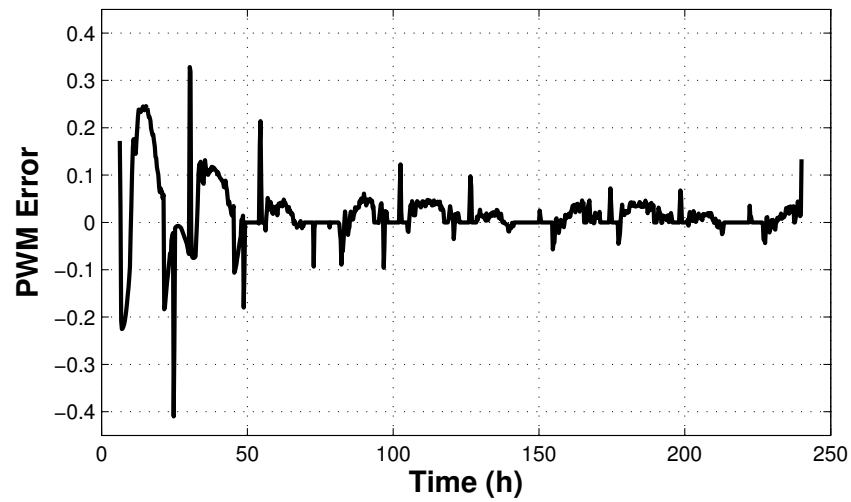
(b) PWM Error

Figure 4.2: Simulation results showing the convergence of the power estimation to the actual value.

γ.

$$
\begin{aligned}
\text{for} \quad P(k) \quad &> \quad 0 \\
\theta_1(k+1) \quad &= \theta_1(k) + \gamma_1(P(k) - \hat{P}(k))(T_{out}(k) - \theta_2(k)) \\
\theta_2(k+1) \quad &= \theta_2(k) + \gamma_2(P(k) - \hat{P}(k))\theta_1(k) \\
\theta_3(k+1) \quad &= \theta_3(k) + \gamma_3(P(k) - \hat{P}(k))(T_{in}(k) - T_{sp}(k)) \\
\psi_j(k+1) \quad &= \psi_j(k) + \gamma_4 * (P(k) - \hat{P}(k)) \\
\text{for} \quad P(k) \quad &\leq \quad 0 \\
\theta_1(k+1) \quad &= \theta_1(k) \\
\theta_2(k+1) \quad &= \theta_2(k) \\
\theta_3(k+1) \quad &= \theta_3(k) \\
\psi_j(k+1) \quad &= \psi_j(k)
\end{aligned}
\tag{4.2}
$$

Figure 4.2 shows the convergence of the power estimate with time. The identification system was initialized at $t = 0.0$, and the error converged pretty quickly. This obviously does not prove stability, but the testing shows no significant stability issues.

4.4.3   Demand Function

The consumption agents are autonomous and intelligent, and in general, each agent is interested in achieving its own goal that does not align perfectly with the super-agents (and could potentially be orthogonal). The primary objective of our home agents is to maintain comfort of the house. However, in the presence of time varying energy price they have an additional objective to manage cost verses comfort. Toward the later goal, our home agents use a cost limiting demand function to regulate their energy costs. The demand function (Equation 4.3) uses an estimate of the power needed to regulate the temperature (P_{est}), a user input neutral factor (f_n), and the energy price ratio (p_r) to calculate the power demand (P_d) during the bidding period.

$$
P_d = \min \left\{ \frac{P_{est} f_n}{p_r}, P_{est} \right\}
\tag{4.3}
$$

The cost limiting demand function guarantees that the total cost of energy used during an auction period is bounded by the Neutral Factor times the the normal energy consumption. Figure 4.3 shows a normalized demand curve for a neutral factor of 5 ($f_n = 5$).

With the use of synchronized PWM control, the power demand is easily regulated using a tunable saturation limit on the temperature controller. With traditional temperature regulation schemes, like hysteresis control, power limiting is considerably more difficult.
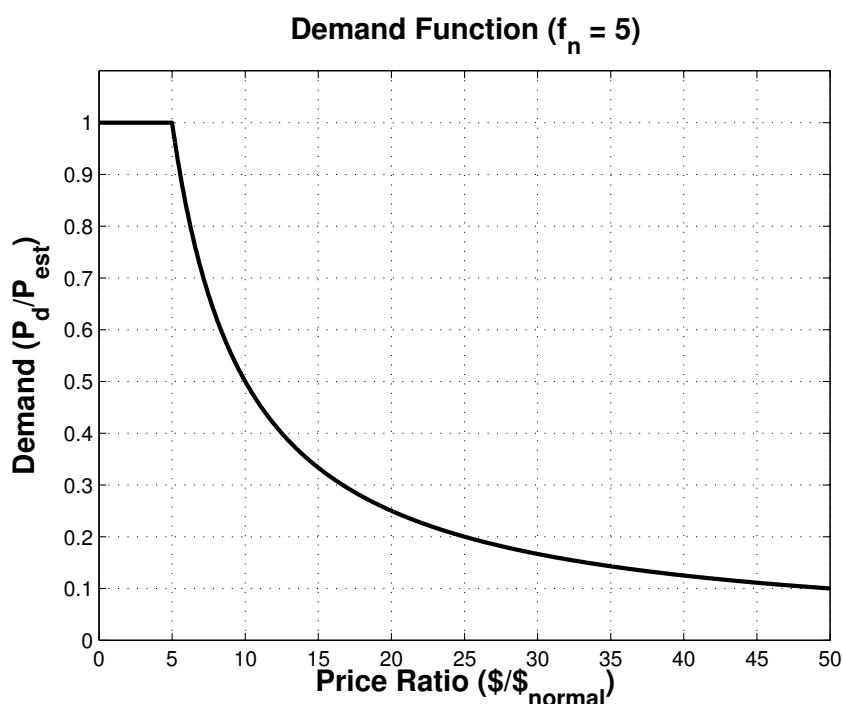
Figure 4.3: Characteristic cost limiting demand function with neutral factor of 5.

4.5 Results

In order to inexpensively and safely test our intelligent agent scarce resource auction system, we built upon the systemic control simulation we previously outlined in Chapter 2. The simulation makes independent houses with randomly chosen properties, including the neutral factor. Simulation is advantageous for this testing because of the ability to experiment with identical networks with and without control to see the exact difference that the control makes.

This testing was conducted on three houses to make the data visualization easier. We initialized the group of houses and let them operate for 7 days in order for the learning to settle. Then, at midnight the auction began with the intention of keeping the normalized time average power below 4000kW. Figure 4.4 shows the filtered power and price ratio with and without control.

Notice that the power does not exactly track 4000kW when the price exceeds 1. During this period, the price was chosen so that the bid power is lower than the target power, and in response to the auction each house is limiting its power (if necessary) to its bid demand using tunable saturation. Therefore, if the total system power is above the total estimate, then the power estimate does not match with the saturation limit. This mismatch is primarily due to the rudimentary PWM to Power conversion that each house uses for estimation. While SEER rating and unit size are primary factors in the total power consumption of a compressor, they are not the sole factors. The outside temperature, age, and service history also play a role in the power
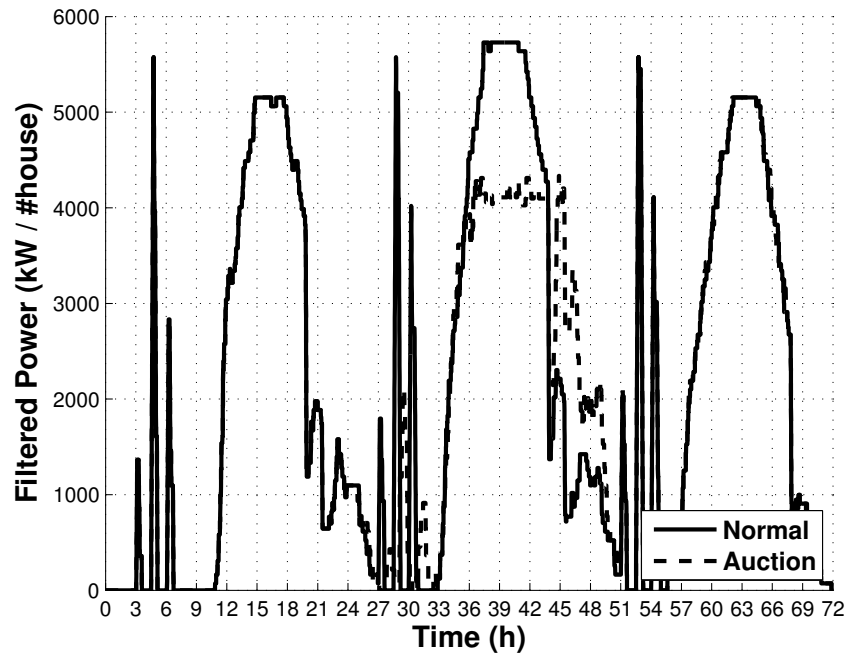
(a) Power
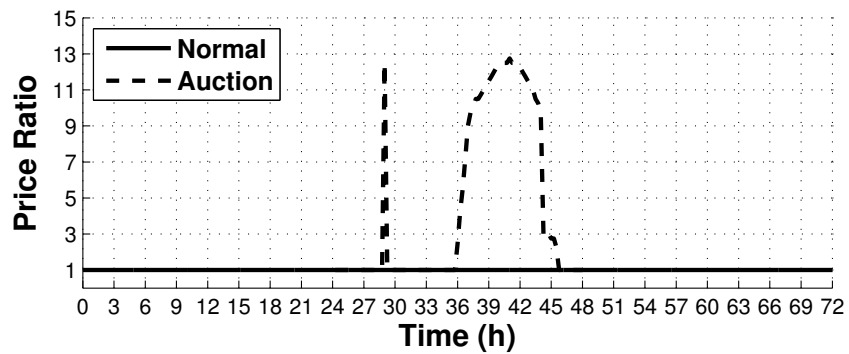
(b) Price

Figure 4.4: Simulation results showing the filtered power for 3 houses under auction control. The auction began at hour 24.

consumption of the unit.

Figure 4.5 shows the inside temperature with and without control for each of the houses in this simulation. Notice that the inside temperature deviates before the price goes above 1. This is due to inaccuracies in the power prediction. When the price does move above 1, each house responds differently because each house had a different neutral factor.
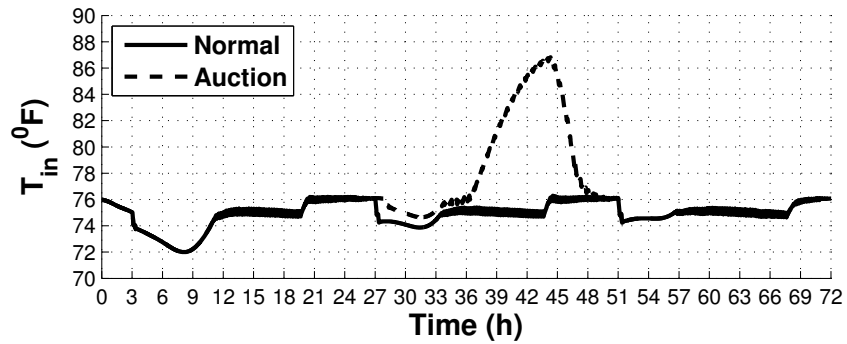
(a) House 1
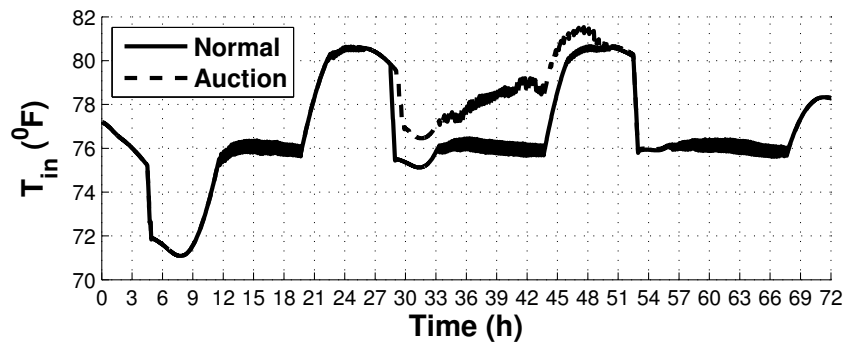
(b) House 2
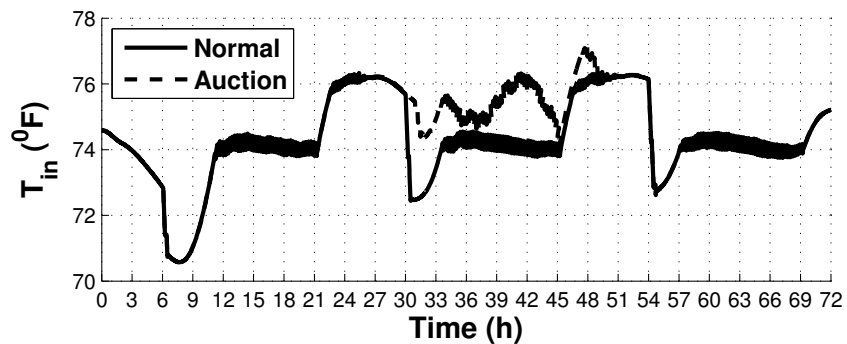
(c) House 3

Figure 4.5: Simulation results showing the inside temperature of the houses with and without control. Each house is different with different neutral factors.

4.6 Discussion

Synchronization of PWM with the auction periods yields a huge simplicity advantage with very little cost. PWM control allows the use of linear control which is well studied and easy to implement. It helps to linearize the system, making identification and prediction simpler. Further, synchronization enables power modulation with a simple tunable saturation on the controller.

The primary open controls issue relates to refinement of the system identification. By improving identification and prediction accuracy, the effect on home comfort will be reduced, especially during low price auction control times. Further, a systematic stability analysis is needed for any identification system to be distributed throughout such a diverse systems as residences.

There are also a number of open issues related to the market design that are outside the scope of this analysis. We assume that each home is honest with its power bid, and we skirted the over-consumption penalty issue. These two issues are interwoven and in need of further study. Also, this mechanism was chosen from the standpoint of simplicity, but it does not scale well with increasing number of market participants. Are there better ways to compute the market clearing price that are less communication intensive? Game theory and economics helped us answer this important question, and we show a truthful clearing algorithm using only one message per user agent in Chapter 5.

Another major open issue is the effect these control strategies have on personal comfort. With our simulation, we model occupants, and their personal preferences, directly. We hope to model our way out of this issue.

4.7 Conclusion

The major conclusion is that implementation of intelligent agent auctions is simplified by synchronizing the agent's control with the auction times. This paradigm eases identification and prediction and makes resource modulation as simple as adjusting saturation limits.

For this treatment, we assumed that the power reference is known a-priori, but from an power systems perspective, the reference is as important as how to maintain it. In future work we are directly targeting the controls/operations research issue of choosing the power reference.

Chapter 5

Residential Electricity Auction with Uniform Pricing and Cost Constraints [1]

5.1 Introduction

A primary tenet of our economic system is the allocation of goods and services based on cost. There are many cases where an individual may want a certain quantity of an item, but if he cannot get all of what he wants because of budget constraints, he will accept as much as he can afford. Consumable natural resources, such as electricity, gas, and water, are a prime example of such goods. In general, the value they possess comes through achievement of some external goal, and that goal is often proportional, meaning that all is best, but some is good. For instance, consider an increase in the price of gasoline. Generally, the price increases because the supply cannot keep pace with the demand. People do not completely stop driving their cars, most just drive less, therefore consuming less gasoline. Anecdotally at least, they buy as much fuel as they can afford (e.g. Bill buys $20 of gasoline regardless of the volume $20 buys). In this vein, we study auctions of an infinitely divisible good to bidders with soft constraints on total cost.

The primary motivation for this work is real time residential electricity pricing. In [30], Severin Borenstein advocates for consumer real time electricity pricing that couples the consumer price to the generation and distribution markets. In California there are a number of time varying pricing programs in place, but the pricing usually follows some predetermined schedule with the possibility of a few extreme pricing instances each year. Further, the likelihood of such a program becoming commonplace in the state of California is very high. We want to take it a step further, and our goal, through this work, is to enable automatic residential real time electricity pricing.

---

[1]Copyright 2009 IEEE. Reprinted, with permission, from 2009 North American Power Symposium, "Residential Electricity Auction with Uniform Pricing and Cost Constraints", by William J. Burke and David M. Auslander

5.2   Problem Description

We have removed the complexity of how, and at what price, the product has been created. We simply assume that the product is divisible and scarce, in that we only have a finite quantity to be allocated in whatever way we choose. Further, we assume that there are many bidders desiring the item. Each bidder desires a certain quantity of the item, and they are willing to pay some maximum cost. If they cannot get their desired amount because the price is too high, then they want to get as much as they can afford.

The goal of this research is to uncover an auction mechanism for allocating the scarce resource to the bidders adhering to their cost constraints. The mechanism should compute the clearing price off-line using only a single bid from each bidder. For reasons of simplicity and perceived fairness we use uniform pricing, so that each bidder pays the same price. Further, we want a policy-consistent algorithm that elicits the bidders true desires.

5.3   Previous Work

Our treatment of this issue is from a game theoretic perspective, and for an excellent modern discussion of algorithmic game theory see [31]. Game theory formalizes the tools for analyzing how rational agents interact to achieve their goals from an economic point of view. The game theoretic issue of algorithmic mechanism design is concerned with developing the framework for agent interaction such that certain criteria are met.

A search of the literature uncovered no previous research on this exact problem. There is a wealth of literature on auctioning divisible goods, mainly dealing with treasury bonds. Further, a few researchers have worked on budget constraints, but they usually take the form of hard constraints where the item is not desired above a certain price. Electricity auctions are also a hot topic, but mainly focusing on generation scheduling.

The auctioning of multi-unit and divisible goods has been discussed for quite some time. In [32], the authors considered both uniform pricing and discriminatory pricing of divisible goods. Under uniform pricing, all winners pay the same unit price for the item, and under discriminatory pricing, each winner potentially pays a different price, as with a Vickrey-Clark-Groves (VCG) mechanism. They concluded that discriminatory pricing is not better than uniform pricing, and it could in fact be worse. But their's is not the final word on this issue.

The article [33] described the possibility of "collusion like equilibria" in divisible goods auctions under uniform pricing. With these detrimental equilibria, the auction revenue could be arbitrarily low, but later the researchers figured out a way around this issue by allowing the auctioneer the option of strategic supply withdrawal, [34]. Further, other researchers wrote about treasury auctions in [35]. With treasury auctions, each bidder receives the same value of the item, but the true value is not known a prioi. They also found undesirable revenue with uniform pricing. The unwanted collusion like equilibria were found to only exist under a continuous bidding strategy, [36]. If the bidding is restricted to finite steps, the underpricing can be well

controlled.

There is limited literature on budget constrained auctions, and as mentioned previously, most research has focused on hard constraints where the bidder does not want the item for greater than a certain cost. Budget constrained single good auctions were investigated in [37]. An interesting finding was that these auctions loose efficiency under budget constraints. Researchers examine budget constrained auctions of multiple goods in [38]. They consider a model where there are many bidders and only two non-divisible items. They noticed that cross bidding can effect the price paid for each item. Other researchers considered the problem of multiple identical items for auction to a multitude of budget constrained bidders, [39]. The bidders either want a subset of items below their budget, or they want no items. The classic VCG mechanism is not incentive compatible in this case because the utilities are not quasi-linear. They found that there is no deterministic truthful mechanism that can allocate all of the units, but they demonstrated a method to allocate some of the units truthfully with maximized revenue. Recently (2008) a group published an impossibility proof for multi-unit constrained budget auctions, [40]. They consider a budget in which each bidder has a value for the item and a maximum payment they are willing to make. They will take some fraction of the items at any price below the budget, but they do not want any if the price exceeds the budget. They find that it is impossible to get pareto-optimality and truthfulness without public budgets.

Electricity market deregulations of the late 1990's and the California electricity crisis of 2000 and 2001 sparked considerable research regarding electricity auctions. The paper, [41], discusses the design of electricity transmission markets. A description of a VCG mechanism for electricity supply and demand markets was given in [42]. Modeling of electricity markets is discussed in detail in [43]. The researchers discuss uniform pricing and discriminatory pricing in the forms of VCG and pay-as-bid mechanism, among other things. They also describe how the collusion like equilibria can be thwarted with discrete bidding functions in the context of electricity markets. They seem to prefer VCG auctions to the alternatives. The debate over uniform pricing versus discriminatory pricing continues with [44], and they found no clear cut advantages either way. Finally, an excelent treatise on the problems with electricity markets is given with [30], where the author advocates for, among other things, consumer real-time pricing.

5.4   Policy-Consistent Clearing Algorithm

We have formulated a policy-consistent auction mechanism for allocating all of the goods under soft budget constraints. The clearing price is computed in polynomial time, and it only requires a single bid from each bidder. The mechanism was inspired by the Shapley Value cost sharing mechanism described in [45] and the real-time dynamic wireless spectrum auctions in [46], but it bears little actual resemblance to either.

To define the problem more succinctly, consider that there is a divisible item E^* with many bidders ($n > 1$). Each bidder ($i = 1, 2, ..., n$) has a private evaluation of the maximum desired quantity of the item, α_i, and a maximum unit price for that quantity, ρ_i. The evaluations are

such that the total desires of the bidders cannot be met with the supply, that is, $\sum_{i=1}^{n} \alpha_i > E^*$. Furthermore, the bidder exhibits a *soft budget constraint* (see Definition 5.4.1) in that if the price (P) is too high ($P > \rho_i$), then the bidder wants the largest amount she can get given her budget constraint $\alpha_i \rho_i$. The goal is to allocate the full quantitiy to the bidders, $A = a_1, a_2, ..., a_n$, at a uniform price P.

Definition 5.4.1 *Soft Budget Constraint*

A budget in which the bidder desires a given quantity α_i of an item, for a maximum price ρ_i, but will accept a proportionally lower quantity at a higher price exhibits a soft budget constraint. The actual quantitiy desired is represented by a_i, and the actual price is given by P. A soft budget constraint function, $\beta_i : P \rightarrow a_i$, is succinctly represented by Equation 5.1.

$$\beta(P) = \left\{ \begin{array}{ll} \frac{\rho_i \alpha_i}{P} & ; P \geq \rho_i \\ \alpha_i & ; P < \rho_i \end{array} \right. \tag{5.1}$$

5.4.1 Soft Budget Constrained Mechanism

The soft budget constrained mechanism accepts a single bid, consisting of the maximum quantity and the maximum price, from each bidder. It returns the clearing price and allocations to the bidders. Proposition 5.4.2 describes the mechanism in full.

Proposition 5.4.2 *Soft Budget Constrained Mechanism*

The soft budget constrained mechanism, f, accepts bids and returns a clearing price and allocation for the item: $f : \{E^*, b_1, b_2, ..., b_n\} \rightarrow \{P, a_1, a_2, ..., a_n\}$. The bids consist of the maximum quantity and maximum price for that quantity, $b_i = \{\alpha_i, \rho_i\}$. There are three steps to the soft budget constrained mechanism:

1. *Order the bids based on the maximum price in ascending order placing a fictitious bid, $b_{n+1} = \{0, \infty\}$, at the end, Equation 5.2.*

$$\rho_1 \leq \rho_2 \leq ... \leq \rho_n \leq \rho_{n+1} \tag{5.2}$$

2. *Iterate on Equations 5.3, starting from $k = 1$ until $\rho_{k-1} < P \leq \rho_k$. (Note that there could be numerical concerns because the denominator starts out negative and through successive iterations becomes positive. A divide by zero could occur. This is easily enough remedied by checking feasibility before dividing.) The fictitious bid ensures that the algorithm stops.*

$$k \leftarrow \quad k + 1 \tag{5.3}$$
$$P \leftarrow \quad \frac{\sum_{i=1}^{k-1} \rho_i \alpha_i}{E^* - \sum_{i=k}^{n} \alpha_i}$$

3. Compute the allocations using Equations 5.4 and 5.5.

$$a_i = \frac{\rho_i \alpha_i}{P} \quad for\ i = 1, 2, ..., k-1 \tag{5.4}$$

$$a_i = \alpha_i \quad for\ i = k, k+1, ..., n \tag{5.5}$$

The first step of the mechanism orders the bids by the maximum bid price. Next, it constrains the allocation to the lower priced bidders ($i < k$, the numerator) so that they get the maximum that their soft budget constraint allows. It leaves the higher priced bidders' allocation at the maximum ($i \geq k$, the denominator), and it recomputes the price. This continues until the price only constrains the $k-1$ bidders. Finally, the allocation is computed for the constrained bidders, and the unconstrained bidders receive their full allocation.

To see why the algorithm works, we need to take a look at the computation for the total allocation under uniform pricing with constrained allocation to some of the bidders, Equation 5.6. The bidders $i = 1, 2, ..., k-1$ have their allocation reduced by their budget constraint, and the remaining $i = k, k+1, ..., n$ bidders obtain their full desired allocation. A good allocation occurs when the price P only constrains the $i = 1, 2, ..., k-1$ bidders.

$$E^* = \sum_{i=k}^{n} \alpha_i + \frac{1}{P} \sum_{i=1}^{k-1} \alpha_i \rho_i \tag{5.6}$$

It is obvious that this mechanism can be computed in polynomial time. Further, it only requires a single bid from each bidder, making it communication efficient.

5.4.2 Policy Consistency

An algorithm is policy-consistent if it coaxes the bidders to directly reveal their true valuations of the item. This coaxing comes in the form of their not being able to get a more individually favorable outcome by revealing false information. The following formal definition is taken from [47].

Definition 5.4.3 *Incentive Compatibility*
A mechanism $(f, b_1, b_2, ..., b_n)$ is called incentive compatible if for every player i, every $v_1 \in V_1, v_2 \in V_2, ..., v_n \in V_n$ and every $\tilde{v}_i \in V_i$, if we denote $A = f(v_i, v_i)$ and $\tilde{A} = f(\tilde{v}_i, v_i)$, then $v_i(A) - p_i(v_i, v_i) \geq \tilde{v}_i(\tilde{A}) - p_i(\tilde{v}_i, v_i)$.

The soft budget constrained mechanism is policy-consistent, but before we get to the formal proof we need a few more concepts. Each bidder gains some utility from the allocation, and the utility takes the normal form of valuation, v_i, minus cost, p_i, as in Equations 5.7, 5.8, and 5.9.

$$u_i = v_i - p_i \tag{5.7}$$

$$v_i(a_i, P) = \begin{cases} \min\{a_i\rho_i, \alpha_i\rho_i\} & ; P \leq \rho_i \\ \min\{a_iP + a_i - \frac{\alpha_i\rho_i}{P}, \alpha_i\rho_i\} & ; P > \rho_i \end{cases}$$ (5.8)

$$p_i = Pa_i$$ (5.9)

Because of the inherent nonlinearities, finding an appropriate valuation for the soft budget constraint was not straightforward. In general, this valuation ensures that the utility is zero along the budget constraint. To see how it works, let's look at a few cases of the total utility under varying circumstances:

- When the the price is below the maximum per-unit price with full or partial allocation, the utility evaluates to 0.

- When the price is below the maximum per-unit price with greater than full allocation, the utility is given by $u_i = \alpha_i\rho_i - Pa_i$. This becomes negative if the total cost exceeds the budget constraint and is positive if the bidder get more for less. A bidder does not need any more than her maximum, but she will accept more if the cost is low enough.

- If the price is higher than the maximum per-unit price, and the allocation is along the budget constraint given in Equation 5.4, then the utility is zero.

- If the per-unit price is higher, and the allocation is below the budget constrained allocation, then the utility is negative because a bidder wants as much as she can get.

- If the price is higher, and the allocation is greater than the budget constrained allocation, the utility is negative because the total cost exceeds the budget.

With the goal of maximizing the utility of each agent, we are now ready for the theorem.

Theorem 5.4.4 *Policy Consistency of Soft Budget Constrained Mechanism*
The Soft Budget Constraint Mechanism as outlined in Proposition 5.4.2 is policy-consistent when the bidders have soft budget constraints as defined in Definition 5.4.1.

Proof The proof consists of a series of inequalities evaluated for the cases of constrained allocation and unconstrained allocation. Throughout the proof, we will consider the case when the bid consists of truthful revelation with a hat and any other bid with a tilde.

- Truthful revelation variables: $\hat{\rho}_i, \hat{\alpha}_i, \hat{P}, \hat{a}_i, \hat{u}_i$

- Untruthful revelation variables: $\tilde{\rho}_i, \tilde{\alpha}_i, \tilde{P}, \tilde{a}_i, \tilde{u}_i$

Case 1 ($i \geq k$): We begin with the unconstrained allocation where each unit receives the maximum allocation at a price below their maximum price. The truthfully revealed allocation is $\hat{a}_i = \hat{\alpha}_i$, and the utility is given by Equation 5.10.

$$
\begin{aligned}
\hat{u}_i &= \min\{\hat{\alpha}_i \hat{P}, \hat{\alpha}_i \hat{\rho}_i\} - \hat{P}\hat{\alpha}_i \\
\hat{P} &\leq \hat{\rho}_i \\
\implies \hat{u}_i &= 0
\end{aligned}
\tag{5.10}
$$

For the false bid, the utility is given as follows:

$$
\tilde{u}_i = \min\{\tilde{\alpha}_i \tilde{P}, \hat{\alpha}_i \hat{\rho}_i\} - \tilde{P}\tilde{\alpha}_i
$$

For $\tilde{\alpha}_i > \hat{\alpha}_i \implies \tilde{P} > \hat{P}$

$$
\begin{aligned}
\tilde{\alpha}_i \tilde{P} &\geq \hat{\alpha}_i \hat{P} \\
\implies \tilde{u}_i &= (0 \text{ or } \hat{\alpha}_i \hat{\rho}_i - \tilde{P}\tilde{\rho}_i) \leq 0
\end{aligned}
\tag{5.11}
$$

For $\tilde{\alpha}_i < \hat{\alpha}_i \implies \tilde{P} < \hat{P} < \hat{\rho}_i$

$$
\begin{aligned}
\implies \tilde{\alpha}_i \tilde{P} &\leq \hat{\alpha}_i \hat{\rho}_i \\
\implies \tilde{u}_i &= \tilde{\alpha}_i \tilde{P} - \tilde{\alpha}_i \tilde{P} = 0
\end{aligned}
\tag{5.12}
$$

Equations 5.10, 5.11, and 5.12 $\implies \tilde{u}_i \leq \hat{u}_i \implies$ no better off by lying!

Case 2 ($i < k$): We continue the proof for the case when the allocation is constrained. In this case, the price is larger than the maximum bid price, and the allocation is constrained.

The truthfully revealed allocation is constrained with the utility given by Equation 5.13.

$$
\begin{aligned}
\hat{u}_i &= \min\{\hat{a}_i \hat{P} + \hat{a}_i - \frac{\hat{\alpha}_i \hat{\rho}_i}{\hat{P}}, \hat{\alpha}_i \hat{\rho}_i\} - \hat{P}\hat{a}_i \\
\hat{a}_i &= \frac{\hat{\alpha}_i \hat{\rho}_i}{\hat{P}} \\
\implies \hat{u}_i &= 0
\end{aligned}
\tag{5.13}
$$

For the false bid, the allocation and utility are given as follows:

$$
\begin{aligned}
\tilde{a}_i &= \frac{\tilde{\alpha}_i \tilde{\rho}_i}{\tilde{P}} \\
\tilde{u}_i &= \min\{\tilde{\alpha}_i \tilde{\rho}_i + \frac{\tilde{\alpha}_i \tilde{\rho}_i}{\tilde{P}} - \frac{\hat{\alpha}_i \hat{\rho}_i}{\tilde{P}}, \hat{\alpha}_i \hat{\rho}_i\} - \tilde{P}\tilde{a}_i
\end{aligned}
\tag{5.14}
$$

For $\tilde{\alpha}_i \tilde{\rho}_i \geq \hat{\alpha}_i \hat{\rho}_i \implies \tilde{P} \geq \hat{P}$.

$$
\begin{aligned}
\tilde{\alpha}_i \tilde{\rho}_i + \frac{\tilde{\alpha}_i \tilde{\rho}_i}{\tilde{P}} - \frac{\hat{\alpha}_i \hat{\rho}_i}{\hat{P}} &\geq \hat{\alpha}_i \hat{\rho}_i \\
\implies \quad \tilde{u}_i = \hat{\alpha}_i \hat{\rho}_i - \tilde{\alpha}_i \tilde{\rho}_i &\leq 0
\end{aligned}
\tag{5.15}
$$

For $\tilde{\alpha}_i \tilde{\rho}_i < \hat{\alpha}_i \hat{\rho}_i \implies \tilde{P} < \hat{P}$.

$$
\begin{aligned}
\tilde{\alpha}_i \tilde{\rho}_i + \frac{\tilde{\alpha}_i \tilde{\rho}_i}{\tilde{P}} - \frac{\hat{\alpha}_i \hat{\rho}_i}{\hat{P}} &< \hat{\alpha}_i \hat{\rho}_i \\
\frac{\tilde{\alpha}_i \tilde{\rho}_i}{\tilde{P}} - \frac{\hat{\alpha}_i \hat{\rho}_i}{\hat{P}} &< 0 \\
\implies \quad \tilde{u}_i < \hat{\alpha}_i \hat{\rho}_i - \tilde{\alpha}_i \tilde{\rho}_i &= 0
\end{aligned}
\tag{5.16}
$$

Finally Equations 5.13, 5.15, and 5.16 $\implies \tilde{u}_i \leq \hat{u}_i \implies$ again the bidder is no better off by lying!

5.5 Application

The soft budget constrained mechanism was envisioned with residential real time electricity markets in mind. The key idea is to treat electricity consumption over a fixed time period as a scarce resource and auction the resource to individual energy consumers. The auction would operate as follows:

1. Bid Call: At a predefined time before the start of the auction period, the units (e.g., homes, refrigerators, thermostats, etc.) in the residential market must submit bids. The bids consist of the maximum expected energy consumption and the maximum price the unit is willing to pay.

2. Clearing: The soft budget constrained mechanism computes the clearing price and returns the price and allocation to the units before the start of the auction period.

3. Auction Period: At the start of the auction period, the price goes into effect and lasts the predefined length of time. The units are charged the clearing price for their electricity consumption.

For a concrete example of how residential real time electricity auctions could work using this scheme, consider a market for intelligent programmable communicating thermostats. The market could operate with real time prices determined using the soft budget constrained mechanism on 15 minute intervals. The bids would be due 5 minutes before the start of each auction period. The quantity of constrained resource, that is, MWh electricity, for each 15 minute period would be determined based on supply and demand in the generation and distribution markets. The thermostats need two things to compete in this auction setting: price responsiveness and energy consumption prediction.

In Chapter 3 we outlined a method of controlling heating ventilation and air conditioning (HVAC) compressors using low frequency pulse width modulation (PWM). Low frequency PWM turns the on/off control signal to an HVAC compressor into a proportional signal at a low sample rate, typically 15min. Temperature control can be accomplished using any type of controller that outputs a proportional signal.
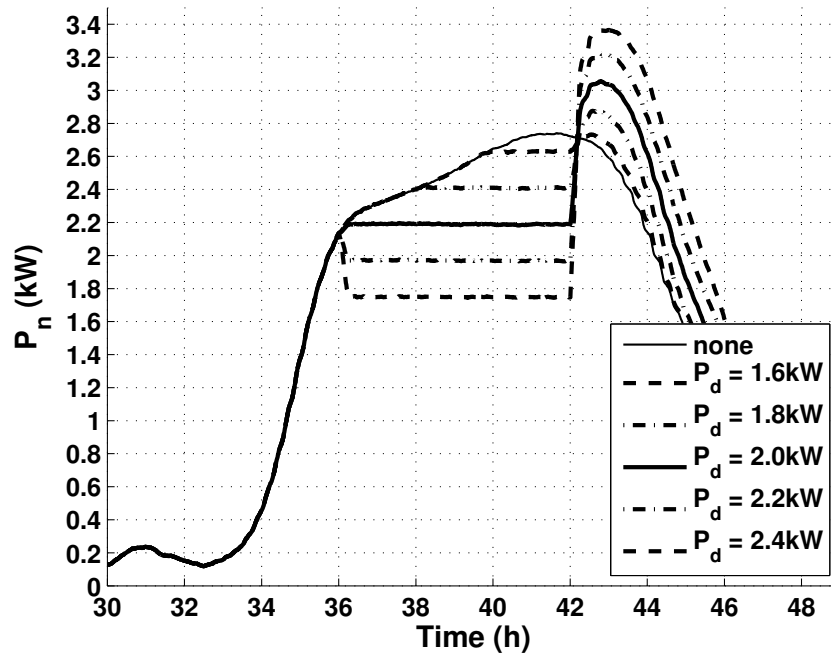
(a) Power
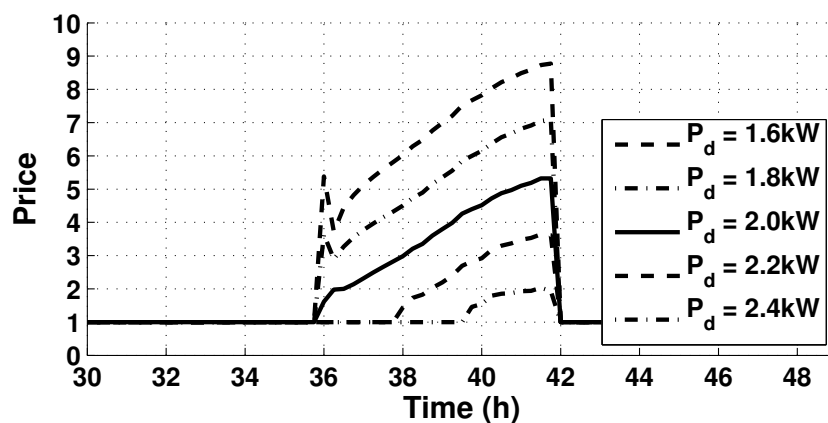
(b) Price

Figure 5.1: Nominal Performance

Low frequency PWM enables thermostat price responsiveness. The power consumption of HVAC compressors is approximately proportional to the compressor duty cycle. If the PWM of the HVAC system is synchronized with the auction period, the energy consumption can be controlled using tunable saturation of the control signal, giving the unit price responsiveness, as in Chapter 4.

Energy prediction is also enhanced by low frequency PWM. Low frequency PWM has a linearizing effect on HVAC systems that makes on-line system identification simpler than with traditional control strategies. With an accurately identified system, the expected energy consumption could be bid based on one step look ahead using the model. In an even more simple scenario, the unit could operate with a time delay equal to the time between the bid call and the start of the auction period (5min). That is, the control action would be computed before the bid due time, submitted for maximum power in the bid, and not performed until the start of the auction period. Either method would result in accurate power prediction.

The soft budget constrained mechanism is perfectly suited to residential real time electricity markets for a number of reasons. The bidding language is simple enough to require very little communication expense yet expressive enough to define the demand over a continuous range of prices. Further, the bidding language relates directly to human notions of cost and, therefore, is easy for users to understand. The mechanism is fast enough and communication efficient enough that communication and computation take a trivial amount of time. Finally, the speed allows the bids to be accepted very close to the start of the auction period, which improves the bidding unit's energy prediction.

Figure 5.1 shows simulation results of 1000 houses participating in an electricity auction with different desired power levels (P_d). The controlled system responds with very flat demand, but there is a large (approximately 10%) steady state error present in the response. This error is solely due to the simple manner in which the individual units predict their power consumption from their desired duty cycle. The duty cycle to power calculation does not account for the typical changes in thermodynamic efficiency that result from different outside temperatures. Most of this error could be resolved by linking the thermostats with the electricity meter so that the actual power consumed could be measured, thereby improving the accuracy.

Another exciting result of using this control is the ability to produce rebound-free response. Typical load management system exhibit large post-control power peaks that are difficult to control. Figure 5.2 shows perfect peak clipping with no rebound peak.

5.6 Discussion

Previous research may seem to suggest a policy-consistent divisible item auction with uniform pricing and bidder cost constraints is impossible, but we have shown that it, on the contrary, is possible. As noted in [33] and others, collusion like equilibria are a strong possibility with uniform pricing in divisible goods markets. Up until this point, the two primary defenses to these disastrous equilibria have been discrete bidding functions and the option of strategic reduction
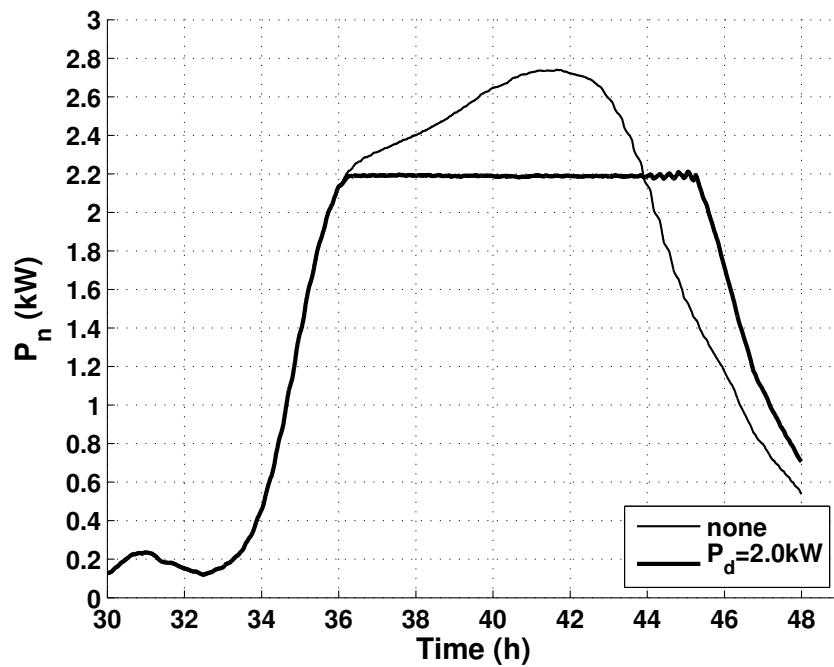
(a) Power
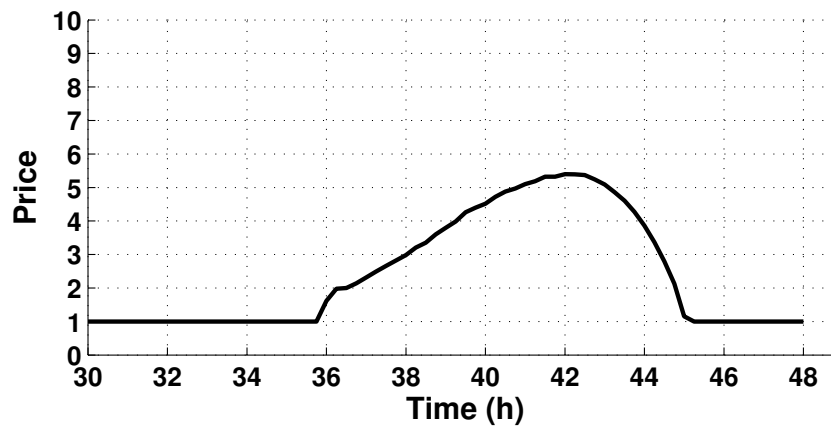
(b) Price

Figure 5.2: Rebound-Free Control

of the item's quantity.

Our positive result is based on the assumption that the bidders exhibit soft cost constraint. The collusion like equilibria were seen when the bidding strategy was arbitrary and continuous, which allowed the bidders to reveal demand curves that were very steep around equilibrium. The soft cost constraint is explicitly not arbitrary because it fixes the strategy over a continuum based on two variables. Further, it is not steep around the equilibrium point. Therefore, in this

special case, collusion like equilibria do not apply.

This treatment of the issue did not directly consider the main criticism of uniform pricing, low revenue. In our formulation there is no explicit way to examine revenue because the true value of the item is not considered, only that it is scarce. Therefore, the value is strictly determined by the bidders' willingness to pay. On the other hand, it is unlikely that low revenue would be a problem unless the bidders strictly undervalued the item, and if that were the case, the item would not be scarce.

Finally, the impossibility proof of multi-unit auctions with constrained budgets described in [40] has a foreboding presence because it says that optimality and truthfulness are impossible. Dobzinski et al., however, considered a very different type of cost constraint than considered here. Their constraint was hard, meaning that above a maximum unit price, the bidder no longer wanted the item. The soft budget constraint considered in our treatment proportionally reduces the desired amount above the price threshold. The difference in the constraints most likely explains our conflicting results.

5.7 Conclusion

The overreaching goal of this research was to enable automatic real-time consumer electricity pricing. With that goal in mind we studied scarce resource auctions, that is auctions of a divisible good when the demand outstrips supply. Further, we introduced a new budget constraint concept – the soft budget constraint – where, if the price is too high, the bidder wants as much of the item as he can afford.

We have described a polynomial time computable mechanism to allocate a scarce divisible good under soft budget constraints. Further, the mechanism only requires a single bid from each of the bidders making it communication efficient. The key positive result is that our mechanism elicits truthful responses from the bidders.

The soft budget constrained mechanism offers much promise for enabling automatic real time consumer electricity pricing. Revenue concerns are the key unanswered question with this scheme, and they are a quite interesting and nuanced topic. Residential real time pricing and electricity auctions should be introduced to reduce generation costs by lowering the peak load, among other things. Further, under real time pricing the revenue must be sufficient to maintain (or improve) profitability. We do not address these issues here, but this line of reasoning points our direction for future research.

Chapter 6

Robust Systemic Control of Load Management Network Using Observer Based Sliding Control

6.1 Introduction

This chapter is also concerned with systemic control of HVAC energy consumption, but we come at it with a different approach. Instead of auctioning the electricity, we have approached this problem from a traditional controls perspective – model it and apply feedback controls. There are a number of challenges to implementing feedback control in such a complex and important system, though.

The first key challenge of controller design is the high system complexity. This is a very large order non-linear system. Just considering the reduced order software-in-the-loop simulation, each house has four states, three uncontrolled inputs (conduction with outside air, infiltration of outside air, and solar radiation), and is controlled by a non-linear controller for both inside temperature and price. This system is nearly impossible to exactly model for controller purposes.

Another challenge that a feedback controller must meet in this environment is robust stability. Load management is used to improve the stability of the grid by reducing the demand in times of high stress. If the system were to become unstable, it could have disastrous consequences, like blackouts and failure of expensive equipment.

The final challenge is that we must use a low bandwidth input that changes on a 15 minute period. HVAC systems have slow time constants because the thermal dynamics of the house are slow. Consequently, there are reliability issues with cycling HVAC equipment too quickly. Furthermore, for infrastructure cost reasons, this system is designed to work with a slow communications requirements like the digital sub-band on broadcast FM radio.

Because of these difficult challenges, we chose to implement a discrete time observer based sliding controller. In Section 6.2, we discuss some background of sliding control. Section 6.3

describes our procedure for designing this controller. Then Section 6.4 presents simulation results that include robustness experiments. Section 6.5 discusses the implications of our design choices, and we conclude the paper in Section 6.6.

6.2 Background

Sliding mode control, or variable structure control, has been discussed extensively in the literature, with a good overview given in [48]. Sliding control is a non-linear control system that commonly uses state-feedback, although output feedback is possible when the plant possess a particular structure. The key idea is to construct an invariant manifold ($s(t) = Gx(t)$), on which the plant is asymptotically stable. Once the manifold reaches equilibrium ($s(t) = 0$), the plant states are guaranteed to converge. The key benefit of sliding control is that it can guarantee robust stability when the plant model has known error bounds.

Discrete sliding control developed slightly slower than the continuous time version because digital implementation required different mathematics in order to guarantee stability. By the late 1980's researchers were beginning to figure out robust discrete sliding control ([49, 50, 51, 52]). Instead of smooth sliding on a manifold, discrete sliding control uses a switching hyper-plane with bounded (or shrinking) oscillations about the equilibrium point when converged. Unlike continuous time sliding control, there are a number of different convergence criteria, and our controller was developed using the criteria from [52].

Using state feedback, discrete sliding control can guarantee robust stability when the parameter and disturbance errors can be bounded by known constants (or bounded functions). Often times (and in our case), the controller states cannot be measured directly, requiring the use of an observer in the loop. In [53], they develop a proof of robust stability using observer state feedback, but the proof is somewhat weak because it relies on the idea that the observer converges fairly quickly. Observer based sliding control was also examined in [54, 55].

More recently, research has focused on improving the design of sliding controllers. In [56], an LQR based method for designing discrete sliding surfaces is presented. Sliding control for uncertain time delay systems was developed in [57].

6.3 Control Design

My goal is to control the aggregate electricity demand by manipulating the electricity price on a 15 minute cycle period similar to that done in the preceding chapters. We have normalized the input variable such that the normal price of electricity during any period of time is 1. The outside temperature is an uncontrollable, but measurable, disturbance input. The output of the system is electricity demand of the aggregate system, in units of $kW/house$.

The design of the controller required four overlapping steps. First the system needed to be identified as a low order system. Sliding control requires state feedback, so an observer was designed. Since we are interested in controlling the output of the system, we needed to generate

the state reference based on the desired output trajectory. Finally, the sliding controller was designed using robustness criteria.

6.3.1   System Identification

As already discussed, the actual system is a very large order non-linear system, and exact modeling is clearly not feasible. Therefore we attempt to capture the aggregate dynamics of the system using a low order linear system (with saturation below 0kW). Since we can only update our control variable every 15 minutes, we are interested in the 15 minute running average of electricity demand. Explicitly, we want to identify a discrete time state-space system with one input (price, $u(k)$), one disturbance input (outside temperature, $d(k)$), and one output (energy, $y(k)$), as in Equation 6.1.

$$x(k + 1) = Ax(k) + Bu(k) + \beta d(k)$$
$$y(k) = Cx(k) + \delta d(k) \tag{6.1}$$

The system identification was completed in two steps. First, a partial-system was identified using data of the response to a constant input and time varying disturbance. The partial-system identification fixed the following parameters: A, β, C, δ. Then, using the partial-system as the base, the remaining parameter, B, was identified using the response to identical disturbance and randomly varying input. The physical reasoning for this is that the identified system should track the real system when the price is at the normal value, and this procedure ensures this outcome.

Table 6.1 shows the error for different system orders (n). The second column ($||e||_\infty$) shows the maximum error. The third column ($cov(0)$) shows the value of the unbiased auto covariance of the error at zero lag. The final column ($||e||_2$) illustrates the mean sum of squares of the error.

Table 6.1: System ID Error

n	$||e||_\infty$	cov(0)	$||e||_2$
2	3.1032	0.47336	8.8381
3	2.5101	0.39662	8.09
4	2.2225	0.34755	7.5787
5	1.2861	0.22367	6.2027
6	1.1898	0.24101	6.4105
7	1.3499	0.22531	6.2181
8	3.504	0.60646	10.007

A system of order 5 was chosen because it exhibited the best error properties. Unfortunately, the error is still relatively high because of inherent non-linearities and un-modeled inputs in the actual system. Equation 6.2 gives the values of the identified model.

$$A = \begin{bmatrix} 0.919, & -0.197, & -0.123, & -0.0970, & -0.110 \\ 0.170, & 0.794 & 0.0244, & -0.376, & -0.190 \\ 0.0378, & -0.144, & -0.462, & -0.0822, & 0.264 \\ -0.0168, & 0.614, & -0.292, & 0.0684, & 0.267 \\ 0.0588, & -0.0564, & 0.131, & -0.400, & 0.750 \end{bmatrix}$$

$$B^T = [0.0561, 0.0551, -0.0287, -0.0537, 0.00735]$$

$$\beta^T = [0.0998, -0.0544, 1.69, 0.586, -0.157,]$$

$$C = [-3.09, 1.47, 0.325, -0.311, 0.105] \qquad \delta = 0.0191 \qquad (6.2)$$

6.3.2 Observer Design

We have modeled this as a linear system with no noise in either the output or the states. Furthermore, the system is observable and controllable. Therefore the Luenberger Observer is an appropriate state estimator. The a-posteriori state observer takes the form specified in Equation 6.3, with the estimated values given by hats ($\hat{o}$).

$$\hat{x}^o(k) = A\hat{x}(k-1) + Bu(k-1) + \beta d(k-1)$$
$$\hat{x}(k) = \hat{x}^o(k) + L\left[y(k) - (C\hat{x}^o(k) + \delta d(k))\right] \qquad (6.3)$$

The observer error dynamics are defined by the eigenvalues of $(A - LCA)$, and when the system is observable and the A matrix is non-singular the eigenvalues can be placed arbitrarily.

6.3.3 Reference Generator

We are interested in controlling the output to maintain a desired reference trajectory, but sliding control is a state feedback process that requires a state reference for tracking control. Therefore, we need to generate the state reference from the output. This procedure is generally called "x_d -generation" as discussed in [58].

The process of x_d-generation uses a simulation of the nominal system to generate a state reference. For our system, the basic structure of the x_d-generator is shown in Figure 6.1.

Interestingly, the x_d-generator is a feedforward simulation with is no uncertainly in the plant parameters. Therefore, the control that produces perfect output reference ($y_d(k)$) tracking can be determined by inverting the dynamics, Equation 6.4.

$$u_d(k) = (CB)^{-1}\left[C(Ax(k) + Bd(k)) + Dd(k+1) - y_d(k+1)\right] \qquad (6.4)$$

This control law is anti-causal because it requires that the output reference and measurable disturbance be known one step ahead of time. In general, it is not a problem to know the next
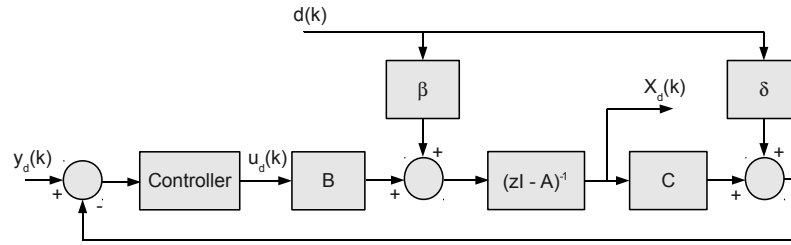
Figure 6.1: x_d-genetator

step of the output reference, but the disturbance may be more difficult. If the disturbance is slowly changing, like outdoor temperature, an estimate of the next temperature could result in adequate performance. (Note that if the disturbance is identically zero or does not contribute directly to the output via the δ matrix, this control law is computable and results in perfect tracking!)

6.3.4 Sliding Control Design

The switching function is the most difficult design element of the sliding controller because it has many parameters and directly effects performances. We followed the form specified in [52], with necessary modification for state reference following. Our switching surface is linear (Equation 6.5) with the quasi-sliding mode satisfying Equation 6.6.

$$s(k) = G\left(x_d(k) - x(k)\right) \qquad (6.5)$$
$$s(k) = s(k+1) = 0 \quad \forall k \qquad (6.6)$$

The article [52] gives a constructive method for choosing the $n-1$ eigenvalue of the sliding surface (the n-th eigenvalue is fixed at the origin). The steps are as follows:

1. Using a similarity transform (Equation 6.7), transform the system into a modified Controllable Canonical form (Equation 6.8).

$$x = Q\bar{x} \qquad (6.7)$$

$$\bar{A} = \begin{bmatrix} \bar{A}_{11} & \bar{A}_{12} \\ \bar{A}_{21} & \bar{A}_{22} \end{bmatrix}, \qquad \bar{B} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

$$\bar{x}^T = [\bar{x}_1^T, \bar{x}_2], \qquad \bar{x}_1 \in \Re^{n-1 x 1}, \qquad \bar{x}_2 \in \Re^1$$

$$\bar{G} = \begin{bmatrix} \bar{G}_1, 1 \end{bmatrix}, \qquad \bar{G}_1 \in \Re^{1 x n-1} \qquad (6.8)$$

2. The dynamics on the sliding surface are given by Equation 6.9. Using any pole placement algorithm (like Matlab place()) the eigenvalues of $\left(\bar{A}_{11} - \bar{A}_{12}\bar{G}\right)$ can be placed arbitrarily

inside the unit circle.

$$\bar{x}_1(k+1) = \left(\bar{A}_{11} - \bar{A}_{12}\bar{G}_1 \right) x_1(k) \tag{6.9}$$

3. Transform G back into the original coordinate system using Equation 6.10.

$$G = \bar{G}Q^{-1} \tag{6.10}$$

Robust stability is directly considered when designing a sliding controller. Our design used multiplicative state and disturbance errors and additive input error (Equation 6.11), with the *tilde* ($\tilde{\circ}$) variables indicating actual error.

$$
\begin{aligned}
x(k+1) =&(A + B\tilde{A})x(k) + B(\gamma + u(k)) + \beta(1+\alpha)d(k) \\
&\tilde{A} \in \Re \quad \gamma \in \Re \quad \alpha \in \Re
\end{aligned} \tag{6.11}
$$

This combination of multiplicative and additive errors leads to a natural choice for symmetrical error bounds ($\bar{A}, \bar{\gamma}, \bar{\alpha}$). Using similar notation to that in [52], the controller with error bounds is shown in Equation 6.12. The variables q and ϵ are used to adjust the convergence of the reaching phase and the switching band. The variable T is the sample time of the discrete system.

$$
\begin{aligned}
u(k) =&(GB)^{-1}[Gx_d(k+1) - GAx(k) - (1 - qT)s(k) + \epsilon T \operatorname{sgn}(s(k)) \\
&- \bar{S}_1(k) - \bar{F}_1(k) - (\bar{S}_2(k) + \bar{F}_2(k)) \operatorname{sgn}(s(k))] \\
\bar{S}_1(k) =& \frac{GB\bar{A}x(k) - GB\bar{A}x(k)}{2} \\
\bar{S}_2(k) =& \frac{GB\bar{A}x(k) + GB\bar{A}x(k)}{2} \\
\bar{F}_1(k) =& \frac{GB\bar{\gamma} + G\beta(1+\bar{\alpha})d(k) - GB\bar{\gamma} + G\beta(1-\bar{\alpha})d(k)}{2} \\
\bar{F}_2(k) =& \frac{GB\bar{\gamma} + G\beta(1+\bar{\alpha})d(k) + GB\bar{\gamma} - G\beta(1-\bar{\alpha})d(k)}{2}
\end{aligned} \tag{6.12}
$$

Table 6.2: Control Parameters

Parameter	Magnitude
$\bar{A}$	$0.001[1, 1, 1, 1, 1]$
$\bar{\gamma}$	1.0
$\bar{\alpha}$	0.1
ϵ	0.0001
q	0.0000005

With the robustness bounds, the controller guarantees the stability of the system if the actual plant parameters fall within the criteria. The robustness parameters effect the performance of the system as well. When the bounds are set too large, the controller exhibits very poor performance because the gains become large. Therefore, we designed the controller with more conservative error bounds that cannot guarantee stability. Table 6.2 lists all of the controller parameters.

6.4   Results

In order to assess the performance of the sliding controller, we used the software-in-the-loop simulation shown in Chapter 2. The main advantage of using simulation results is that all aspects of the system, from house parameters to outside conditions, can be controlled. This level of control makes repeating experiments using identical systems very simple.

First, let us consider the tracking performance of the control system on the nominal plant for various output references. Figure 6.2 and Table 6.3 summarize the performance. The response should exhibit good transient response, total error, and steady state error. The quantity $\%e_{max}$ represents the percentage of the maximum tracking error. The norm $||e||_2$ is the euclidean norm of the tracking error. Finally, $\%e_{ss}$ illustrates the steady state tracking error as a percentage of the output reference.

Table 6.3: Nominal System Tracking Performance

$\mathbf{P_d\ (kW_n)}$	$\%\mathbf{e_{max}}$	$||\mathbf{e}||_\mathbf{2}$	$\%\mathbf{e_{ss}}$
1.6	15.2	3.25	1.10
1.8	9.33	2.49	1.76
2.0	9.76	2.73	2.57
2.2	4.93	2.41	3.30
2.4	7.78	2.86	3.82

I am also interested in the robust stability of the closed-loop system. Put simply, the system should remain stable when the nominal plant parameters are perturbed. By adjusting the distributions from which the house parameters in the software-in-the-loop simulation are selected, we experimentally examined the robustness of the controller. Table 6.4 describes the modifications to the base parameter range, with the number representing percent change. The last column in the table shows the relative magnitude of the uncontrolled peak demand (also as a percent difference).

Figure 6.3 illustrates the tracking performance of the perturbed plants along with the nominal plant with no changes to the control system. Note that each of the perturbed systems remains stable. Table 6.5 lists the performance of the control system on the perturbed plant.

The final performance criterion of interest is rebound free control. The goal of rebound free control it to extend the control so that there is no post control peak above the target power. Since the local control in each house is not sophisticated enough to consume extra energy when
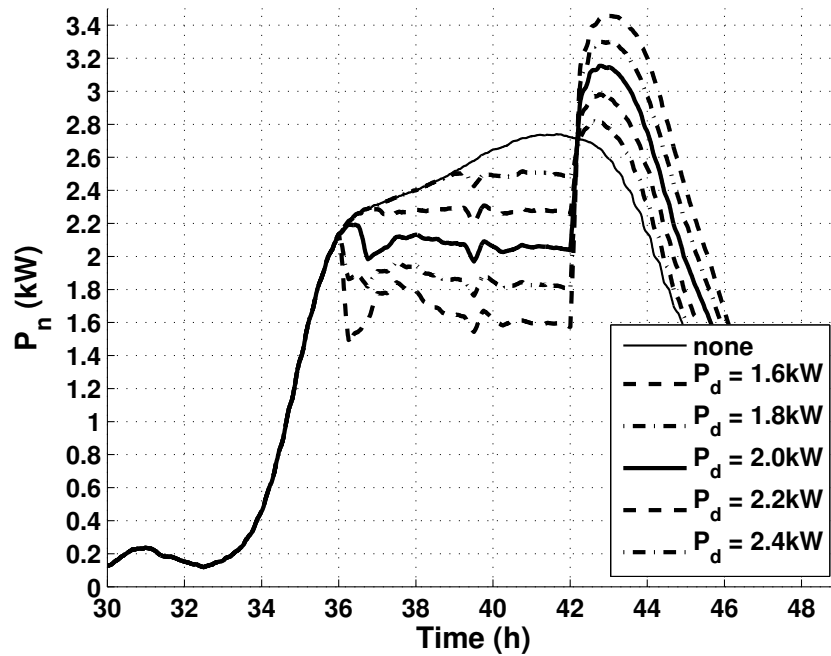
(a) Power
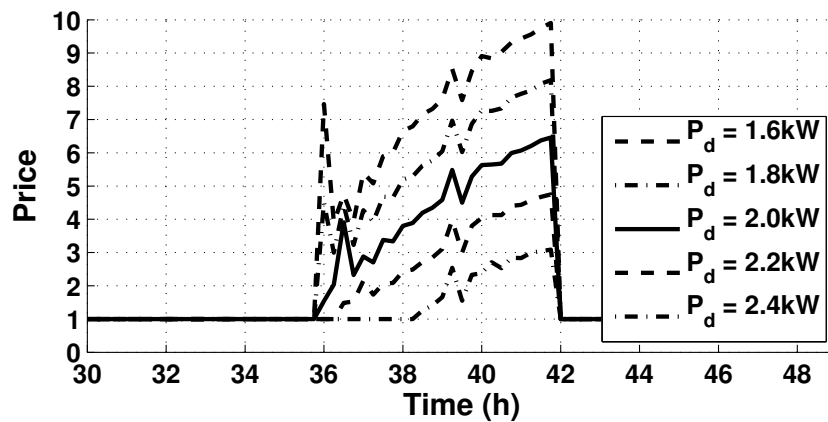
(b) Price

Figure 6.2: Nominal Performance

the price is below 1.0, the controller never issues a price below 1.0. Therefore, to test rebound free control, we just extend the control period until the controller issues a number of consecutive prices at the minimum. Figure 6.4 shows the performance of the controller (and nominal plant) under these conditions.

Table 6.4: House Parameter Modification

ID	House Size	Insulation Quality	Window Size	Peak Power
hhl	25%	25%	-25%	-11.8%
hlh	25%	-25%	25%	32.3%
hll	25%	-25%	-25%	6.53%
lhl	-25%	25%	-25%	-20.4%
llh	-25%	-25%	25%	3.38%

Table 6.5: Perturbed Performance

ID	$\%e_{max}$	$\|e\|_2$	$\%e_{ss}$
nominal	9.76%	2.73	2.57%
hhl	24.0%	5.088	2.35%
hlh	39.0%	9.18	2.14%
hll	9.88%	3.26	4.14%
lhl	21.6%	5.77	1.30%
llh	16.2%	3.72	1.18%
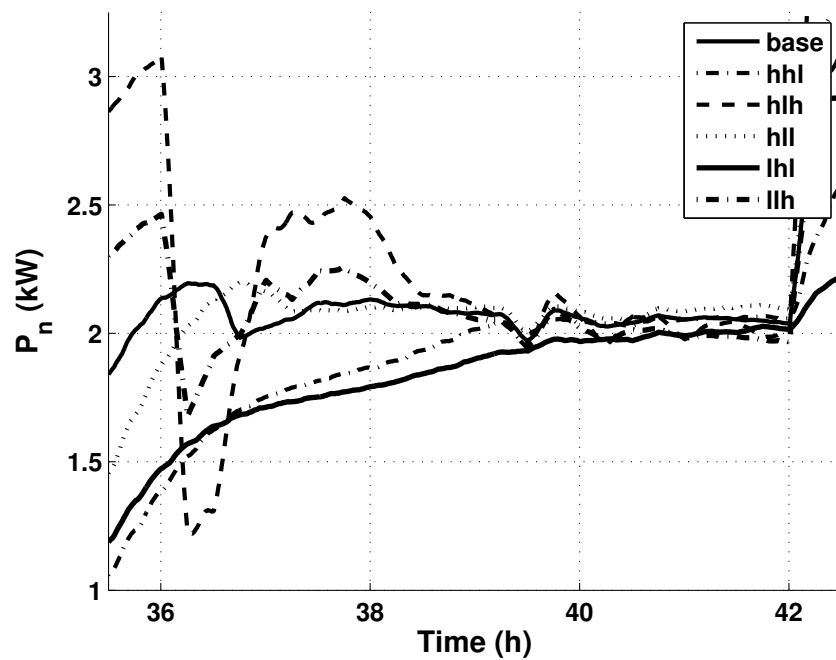
Figure 6.3: Tracking Performance of Perturbed Plants

6.5   Discussion

The nominal performance of the controller is quite good. The steady state error is low (below 4%) in all cases. Predictably, the overshoot is larger when the desired decrease in consumption is larger (i.e. when the P_d is low). The system takes a little while (about 4 hours) to settle, but this is primarily due to the low bandwidth of the control signal (4 hours is only 16 control updates).

Unfortunately, this controller is not provably robustly stable for two reasons. While there exists a proof for the stability of observer based discrete sliding control, it relies on the tenuous notion that the observer estimation error converges fairly quickly. While considerable effort was put into identifying the system, experimental results indicate that the modeling errors are too
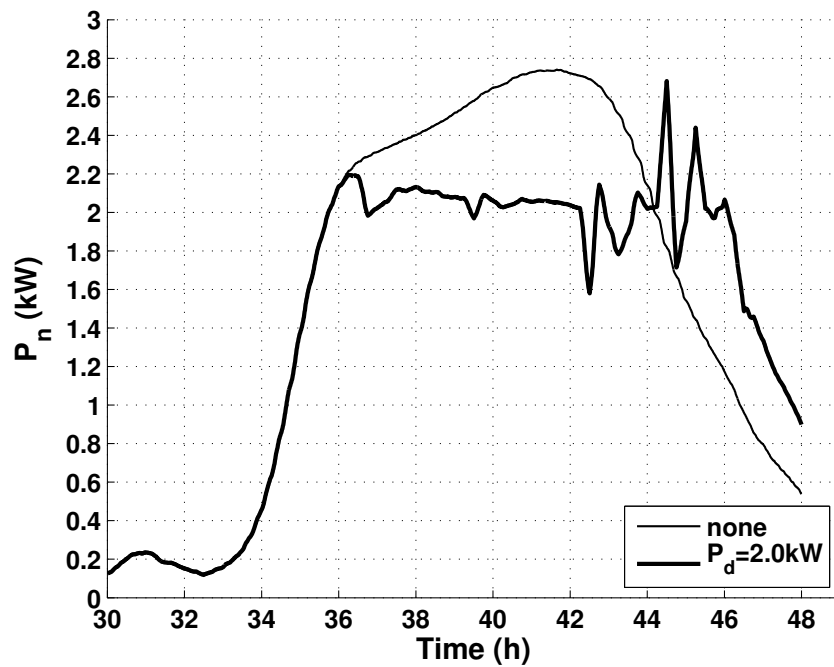
(a) Power
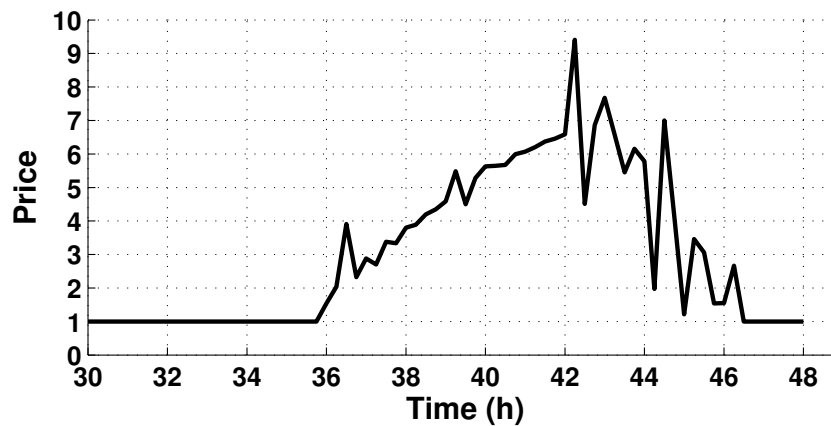
(b) Price

Figure 6.4: Rebound-Free Control

large to rely to this condition.

Further, sliding control can be designed to directly compensate for soft non-linearities in a system, but it cannot compensate for hard non-linearities like input saturation. The design of the thermostat network is such that price signals below 1.0 have no affect on the system because the thermostats do not know how to over-consume in times of low price. This lack of control authority at low prices restricted the controls gains, thereby reducing the error bounds. The

restricted bounds could not even appropriately bound the output error in the identified system.

Regardless of our lack of provable robustness, the controller experimentally exhibits excellent robustness properties. Given a plant that produces peak power over 30% higher than the nominal plant, an unmodified control system brought the measured output to within 2.14% of the desired output (closer than the nominal plant). A similar result is found with smaller than nominal plants, indicating excellent robustness properties.

The control system did not produce particularly good rebound-free performance. Figure 6.4 shows large output chatter past hour 42. At this time, the power consumption in the uncontrolled case begins to decrease, and this inflection must introduce large modeling errors and unmodeled dynamics. Therefore, the chatter is likely caused by modeling errors.

6.6   Conclusion

I have demonstrated systemic control of a residential load management system using observer based sliding control. The controller does an excellent job tracking different constant output references. Further, the controller is robust to system perturbations, producing very good steady state error for plants over 30% larger than the design system. Unfortunately, the errors in the identified model cause unacceptably large oscillations when attempting rebound-free control. Further, the slow settling time indicates that the controller is ill-suited to time varying output references. Regardless, the controller is excellently suited to short term robust following of constant power references.

Chapter 7

Implications of Load Control

7.1 Research Overview

In the preceding chapters, we have described the development of both local and systemic control of residential load management systems. The goal of local control is to modulate the power consumption of an individual consumer (generally an autonomous device) in such a way that achieves the prescribed goals of the human. Whereas, systemic control aims to manipulate the aggregate power consumption of the group of consumers. Both levels in this hierarchical control scheme are important to achieving the goals of load management.

Experimental verification is a key difficulty in developing advanced load management techniques. Deploying experimental devices is costly, and large scale control experiments could have detrimental impact on the grid. To solve these problems, we developed a software-in-the-loop simulation of a load management network (Chapter 2). The simulation contains smart thermostat devices controlling dynamic models of homes, and it enables inexpensive and safe experimentation with local and systemic control techniques.

In Chapter 3, we demonstrated local control via an alternative thermostat design using low-frequency pulse width modulation (PWM). Traditional thermostats use a hysteretic controller to regulate inside temperature, but we propose operating the HVAC system with low-frequency PWM and using a linear controller instead. The key advantages of low-frequency PWM are the following: Power consumption can be directly modulated using tunable saturation, and it is much simpler to predict future power consumption using either delay control (for one step look ahead) or on-line identification. Low-frequency PWM also reduces the complexity of controlling multi-stage HVAC compressors.

Chapter 4 deepened the local control and provided the foundation for systemic control by outlining how price responsiveness can be achieved in thermostats using synchronized low-frequency PWM. Two things are needed for price responsiveness – control and timing. First, the thermostat must be able control power consumption in response to price. The soft budget constraint demand curve determines the power consumption in response to price, with low-frequency PWM enabling power modulation. Second, the thermostat needs to time the power

consumption to conform to the periods of time varying price. For this, we suggest synchronizing the low-frequency PWM period with the price control period.

Using a game theoretic approach, we develop auction based systemic control in Chapter 5. The Soft Budget Constrained Mechanism computes a uniform price that elicits a desired aggregate power demand, and the main result is that the Soft Budget Constrained Mechanism is policy consistent when all bidders (i.e. price responsive thermostats) have a soft budget constrained demand curve. Policy consistent is a powerful property that means that the bidders will bid their true valuation because they cannot get a better outcome by lying. Further, the mechanism has excellent computation and communication qualities. It is fast, meaning that it is computable in polynomial time. Further, the mechanism is communications efficient, in that it only requires a single message from each bidder to compute the clearing price.

Taking a traditional feedback controls view of systemic control in Chapter 6, we presented a discrete time observer based sliding controller to regulate aggregate power demand using price. The main challenge of systemic control design is insuring robust stability and performance using the low bandwidth price input. The sliding controller achieves good performance with less than 4% steady-state tracking error. Further, it maintained robust stability with aggregate demands over 30% larger than the design nominal.

I have demonstrated systemic and local control of a load management network. But what is this technology good for? In the following sections, we will outline our vision for where this work could have impact and the challenges it faces.

7.2   Research Challenges

In its current form, this work is only deployable in a limited sense, and to gain wider impact, there are a number of challenges to overcome. From an engineering perspective, three things need to be worked out – "real" real-time pricing, "full" truthfulness of mechanism, and synchronization randomization.

"Real" real-time pricing alludes to the fact that our local price control assumes a normalized price of electricity that does not translate well to a pricing scheme with constant fluctuations. The main problem with "real" real-time pricing is that the price of electricity varies seasonally, making it difficult to set a normal price. We believe that this problem can be solved through forward contracts using look-ahead control.

The policy consistency of the Soft Budget Constrained Mechanism relies on the assumption that the bidders actually have soft budget constraint demand curves and does not consider the case when the demand curve is arbitrary, i.e. does not conform to the generalized Soft Budget Constraint. It is not clear what impact arbitrary demand curves have on truthfulness when constrained to bid with a soft budget constraint, but this is a topic of current interest. Regardless of the outcome of this question, it will not change the usefulness of this mechanism for engineered systems.

Synchronization randomization refers to the means by which we impart artificial load di-

versity when the PWM periods are actually synchronized. Without proper randomization, the units could cycle in a pattern that causes large variation of instantaneous demand within the synchronization period. As yet, we have not identified the optimal distribution for which to initialize the PWM start time to insure that the instantaneous demand remains approximately constant over the period.

The preceding are engineering problems and, as such, can most likely be solved through effort. The more daunting challenge is of the policy form. When will dynamic pricing gain wide spread adoption? Or maybe a more apposite question: Will regulators allow consumers to be subjected to prices set by control algorithms? We do not have the answers to these questions, but in many respects the stability and efficiency of the grid depend on the answers.

7.3 Research Vision

Traditionally, electricity generation has been bought and sold through a number of products that account for reliability of the end service. Transactions for energy occur in the wholesale market through long term contracts, in day ahead trading, and at the spot price in the real-time market. Ancillary services, of which there are many different types, are products that attempt to maintain the reliability of the grid. Reserve, a class of ancillary services, provides a kind of insurance against changes in the supply or load by being ready to provide additional power when needed. The grid regulating organizations have rules that dictate the minimum reserves that must be maintained by the electricity providers. The seller of reserve is paid a capacity payment based on the amount of power they *could* supply, and they are paid the spot price for the energy they do provide.

My vision for this work is to start an electricity aggregation service that manages the electricity for a large network of consumers. Each consumer would have a price responsive thermostat responsible for local control in the home, and each consumer would buy electricity directly from the aggregator. The aggregator determines the price of electricity sold to the consumers, thereby providing systemic control. The goal of the aggregator is to provide mutual benefits throughout the system. Therefore, the aggregator must be able to make a profit while not disappointing the consumers.

Given the particulars of the wholesale electricity markets, and the fact that the aggregator can control price and demand, the aggregator could potentially sell a number of electricity products. Essentially, demand control makes the aggregator look exactly like a generator, so any product a generator could supply, the aggregator could supply too. The key constraint is that the aggregator must keep the customers happy or they will leave the service. Therefore, the aggregator has "production" constraints.

The aggregator would buy *and sell* energy in the wholesale markets like generators do now. The power to control demand effectively removes the forecasting error typical in demand prediction. The energy for this deterministic demand could be bought directly on the day ahead market. Alternatively, the aggregator could make an opportunistic bid into the day ahead market

with the ability to buy and sell (depending on whether the bid was high or low) at the spot prices.

Additionally, the aggregator could provide different reserve products. The ability to provide reserve also stems from the fact that demand is controllable. The capacity payments could then help subsidize the price of electricity each consumer pays, thereby helping to lower their electricity bill.

My modeling and systemic control research provides excellent tools for predicting and controlling power demand. The software-in-the-loop simulation would be used to predict aggregate demand for bidding into the day ahead market. The aggregator would then set the dynamic price of electricity via auctions using the Budget Constrained Mechanism. This fast and communication efficient mechanism would allow the aggregator to set arbitrary desired power demands. Using this level of information and control, the aggregator would devise a beneficial wholesale market trading strategy.

The price responsive thermostats would be wholly based on the work presented in preceding chapters. Since all of the experimentation up until this point took place in a software-in-the-loop simulation, most of the software for the intelligent thermostat is already written. The thermostat code would be ported over to the Android operating system. Android is a free Linux based operating system for use in mobile devices. The main advantage of using Android is that the networking software already exists, and Android can be updated over-the-air. Finally, this device would contain a 802.11b/g (wi-fi) radio for communications via the home's high-speed Internet.

This smart grid strategy leverages communications and controls to provide benefits to the consumer and society. The individual consumer would experience smaller electricity bills because they could shift around when they consumer electricity to take advantage of lower off-peak prices, and they would receive capacity payments that lower their bills. The benefits to society would be multi-fold. The average cost of electricity would decrease with the reduction in expensive peak production and the availability of less expensive ancillary services. Further, the reserve capacity would actually increase when the aggregate demand increased, thereby alleviating concerns over reserve capacity and improving reliability of the electricity grid.

Bibliography

[1] Bellarmine, G. T., 2000. "Load management techniques". In IEEE SoutheastCon 2000 'Preparing for the New Millennium', Apr 7-Apr 9 2000, Conference Proceedings - IEEE SOUTHEASTCON, Institute of Electrical and Electronics Engineers Inc., Piscataway, NJ, USA, pp. 139–145.

[2] Navid-Azarbaijani, N., and Banakar, M., 1996. "Realizing load reduction functions by aperiodic switching of load groups". *IEEE Transaction on Power Systems, Vol 11, No. 2*.

[3] Ng, K., and Sheble, G., 1998. "Direct load control-A profit-based load management using linear programming". *Power Systems, IEEE Transactions on,* **13**(2), pp. 688–694.

[4] Moe, M. L., and Hedman, T. A., 1997. Adaptive load cycler for controlled reduction of energy use, Oct. U.S. Classification: 364/492; 364/483; 364/557; 307/32; 307/39 ; International Classification: H02J 318; G05F 166.

[5] Do, A., Auslander, D. M., Burke, W. J., White, R. M., and Wright, P. K., 2007. Technical review of residential programmable communicating thermostat implementation for title 24-2008. draft final report CEC-500-2007-XXX, California Energy Commission, PIER Buildings End-Use Energy Efficiency Program.

[6] Katipamula, S., and Lu, N., 2006. "Evaluation of residential hvac control strategies for demand response programs". *ASHRAE Transactions Volume 112, Part1*.

[7] Herter, K., McAuliffe, P., and Rosenfeld, A., 2007. "An exploratory analysis of california residential customer response to critical peak pricing of electricity". *Energy 32*.

[8] Callaway, D. S., 2009. "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy". *Energy Conversion and Management,* **50**(5), May, pp. 1389–1400.

[9] Kok, J., Warmer, C., and Kamphuis, I., 2005. "PowerMatcher: multiagent control in the electricity infrastructure". In 4th International Conference on Autonomous Agents and Multi agent Systems, AAMAS 05, Jul 25-29 2005, Proceedings of the International Conference on Autonomous Agents, Association for Computing Machinery, New York, NY 10036-5701, United States, pp. 115–122.

[10] Lu, N., and Nguyen, T., 2006. "Grid friendly [trademark] appliances - load-side solution for congestion management". In IEEE Power Engineering Society Transmission and Distribution Conference, PES TD 2005/2006, May 21-24 2006, Proceedings of the IEEE Power Engineering Society Transmission and Distribution Conference, Institute of Electrical and Electronics Engineers Inc., Piscataway, NJ 08855-1331, United States, pp. 1269–1273.

[11] Zhang, W., and Feliachi, A., 2003. "Residential load control through real-time pricing signals". In System Theory, 2003. Proceedings of the 35th Southeastern Symposium on, pp. 269–272.

[12] Mortensen, R., and Haggerty, K., 1990. "Dynamics of heating and cooling loads: Models, simulation, and actual utility data". *IEEE Transactions on Power Systems, Vol. 5 No. 1.*

[13] Calloway, T., and C.W. Brice, I., 1982. "Physically-based model of demand with applications to load management assessment and load forcasting". *IEEE Transactions on Power Apparatus and Systems, Vol. PAS-101, No.12.*

[14] Mortensen, R., and Haggerty, K., 1988. "A stochastic computer model for heating and cooling loads". *IEEE Transactions on Power Systems, Vol. 3, No. 3.*

[15] Ihara, S., and Schweppe, F. C., 1981. "Physically based modeling of cold load pickup". *IEEE Transactions on Power Apparatus and Systems, Vol. PAS-100, No. 9.*

[16] Lu, N., Chassin, D. P., and Widergren, S. E., 2005. "Modeling uncertainties in aggregated thermostatically controlled loads using a state queuing model". *IEEE Transactions on Power Systems, Vol 20, No. 2.*

[17] Auslander, D. M., 2007. Tranrunc: Realizing task/state mechanical system control software in c. UC Berkeley ME135 Class Notes.

[18] Auslander, D., Ridgely, J., and Ringgenberg, J., 2002. *Control Software for Mechanical Systems: Object-Oriented Design in a Real-Time World*. Prentice Hall PTR.

[19] Huizenga, C., and Barnaby, C. S., 1985. "Methodology for hourly simulation of residential hvac equipment". In The First National Conference on Microcomputer Applications for Conservation and Renewable Energy, pp. 337 – 343.

[20] ASHRAE, 2001. *ASHRAE Fundamentals Handbook*. American Society of Heating, Refrigerating, and Air Conditioning Engineers, pp. 30.13 – 30.14.

[21] LaRue, A., 2006. "Distributed Sensing and Controlling of Residential HVAC Systems for Thermal Comfort, Demand Response, and Reduced Annual Energy Consumption". Master's thesis, UNIVERSITY OF CALIFORNIA.

[22] Katipamula, S., and Lu, N., 2006. "Evaluation of residential hvac control strategies for demand response programs". *ASHRAE Transactions, Volume 112, Part 1.*

[23] Federspiel, C., Lanning, S. D., Li, H., and Auslander, D. M., 2001. "Coordinated control of HVAC systems". In ICEBO - International Conference for Enhanced Building Operations, Energy Systems Laboratory (http://esl.tamu.edu).

[24] Sira-Ramirez, H., 1989. "On the design of nonlinear pulse-width-modulation controllers". In Control and Applications, 1989. Proceedings. ICCON '89. IEEE International Conference on, pp. 397–402.

[25] Hou, L., and Michel, A. N., 2001. "Stability analysis of pulse-width-modulated feedback systems". *Automatica,* **37**(9), Sept., pp. 1335–1349.

[26] Akkermans, H., Ygge, F., and Gustavsson, R., 1996. "HOMEBOTS: intelligent decentralized services for energy management". *In: Proceedings of the Fourth International Symposium on the Management of Industrial and Corporate Knowledge, ISMICK'96*, pp. 128—142.

[27] Gustavsson, R., 1999. "Agents with power". *Commun. ACM,* **42**(3), pp. 41–47.

[28] Kok, K., Derzsi, Z., Hommelberg, M., Warmer, C., Kamphuis, R., and Akkermans, H., 2008. "Agent-Based electricity balancing with distributed energy resources, a multiperspective case study". In Hawaii International Conference on System Sciences, Proceedings of the 41st Annual, p. 173.

[29] Codenotti, B., and Varadarajan, K., 2007. "Computation of market equilibria by convex programming". In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, eds. pp. 135–158.

[30] Borenstein, S., 2002. "The trouble with electricity markets: Understanding california's restructuring disaster". *Journal of Economic Perspectives,* **16**(1), pp. 191–211.

[31] Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V., 2007. *Algorithmic Game Theory*. Cambridge University Press, Sept.

[32] Wilson, R., 1979. "Auctions of shares". *The Quarterly Journal of Economics,* **93**(4), Nov., pp. 675–689.

[33] Back, K., and Zender, J. F., 1993. "Auctions of divisible goods: On the rationale for the treasury experiment". *The Review of Financial Studies,* **6**(4), pp. 733–764.

[34] Back, K., and Zender, J. F., 2001. "Auctions of divisible goods with endogenous supply". *Economics Letters,* **73**(1), oct, pp. 29–34.

[35] Wang, J. J., and Zender, J. F., 2002. "Auctioning divisible goods". *Economic Theory,* **19**(4), June, pp. 673–705.

[36] Kremer, I., and Nyborg, K. G., 2004. "Underpricing and market power in uniform price auctions". *Rev. Financ. Stud.,* **17**(3), July, pp. 849–877.

[37] Maskin, E. S., 2000. "Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers". *European Economic Review,* **44**(4-6), May, pp. 667–681.

[38] Benot, J., and Krishna, V., 2001. "Multiple-Object auctions with budget constrained bidders". *The Review of Economic Studies,* **68**(1), pp. 155–179.

[39] Borgs, C., Chayes, J., Immorlica, N., Mahdian, M., and Saberi, A., 2005. "Multi-unit auctions with budget-constrained bidders". In Proceedings of the 6th ACM conference on Electronic commerce, ACM, pp. 44–51.

[40] Dobzinski, S., Lavi, R., and Nisan, N., 2008. "Multi-unit auctions with budget limits". In Foundations of Computer Science, 2008. FOCS '08. IEEE 49th Annual IEEE Symposium on, pp. 260–269.

[41] Chao, H., and Peck, S., 1996. "A market mechanism for electric power transmission". *Journal of Regulatory Economics,* **10**(1), July, pp. 25–59.

[42] Hobbs, B. F., Rothkopf, M. H., Hyde, L. C., and ONeill, R. P., 2000. "Evaluation of a truthful revelation auction in the context of energy markets with nonconcave benefits". *Journal of Regulatory Economics,* **18**(1), July, pp. 5–32.

[43] Fabra, N., von der Fehr, N., and Harbord, D., 2002. "Modeling electricity auctions". *The Electricity Journal,* **15**(7), pp. 72–81.

[44] Zhang, Y., Jiao, L., Chen, S., and Ni, Y., 2003. "Information, incentive and efficiency in power market design". In Advances in Power System Control, Operation and Management, 2003. ASDCOM 2003. Sixth International Conference on (Conf. Publ. No. 497), Vol. 2, pp. 587–592.

[45] Moulin, H., and Shenker, S., 1996. "Strategyproof sharing of submodular access costs: Budget balance versus efficiency". *SSRN eLibrary*, May.

[46] Gandhi, S., Buragohain, C., Cao, L., Zheng, H., and Suri, S., 2008. "Towards real-time dynamic spectrum auctions". *Computer Networks,* **52**(4), Mar., pp. 879–897.

[47] Nisan, N., 2007. "Introduction to mechanism design (for computer scientists)". In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, Éva Tardos, and V. V. Vazirani, eds. Cambridge University Press, pp. 209–241.

[48] Young, K., Utkin, V., and Ozguner, U., 1999. "A control engineer's guide to sliding mode control". *Control Systems Technology, IEEE Transactions on,* **7**(3), pp. 328–342.

[49] Drakunov, S. V., and Utkin, V. I., 1992. "Sliding mode control in dynamic systems". *International Journal of Control,* **55**(4), p. 1029.

[50] Furuta, K., 1990. "Sliding mode control of a discrete system". *Systems & Control Letters,* **14**(2), Feb., pp. 145–152.

[51] Sarpturk, S., Istefanopulos, Y., and Kaynak, O., 1987. "On the stability of discrete-time sliding mode control systems". *Automatic Control, IEEE Transactions on,* **32**(10), pp. 930–932.

[52] Gao, W., Wang, Y., and Homaifa, A., 1995. "Discrete-time variable structure control systems". *Industrial Electronics, IEEE Transactions on,* **42**(2), pp. 117–122.

[53] Misawa, E., 1995. "Observer-based discrete-time sliding mode control with computational time delay: the linear case". In American Control Conference, 1995. Proceedings of the, Vol. 2, pp. 1323–1327 vol.2.

[54] Richter, H., 1997. "Hyperplane design in Observer-Based Discrete-Time sliding mode control". Masters thesis, Oklahoma State University, May.

[55] Richter, H., and Misawa, E. A., 2006. "Boundary layer eigenvalues in SISO Discrete-Time sliding mode control with observer". *Journal of Dynamic Systems, Measurement, and Control,* **128**(3), pp. 729–730.

[56] Tang, C. Y., and Misawa, E. A., 2002. "Sliding surface design for discrete VSS using LQR technique with a preset real eigenvalue". *Systems & Control Letters,* **45**(1), pp. 1–7.

[57] Xia, Y., and Jia, Y., 2003. "Robust sliding-mode control for uncertain time-delay systems: an LMI approach". *Automatic Control, IEEE Transactions on,* **48**(6), pp. 1086–1091.

[58] Misawa, E. A., 1997. "Discrete-Time sliding mode control: The linear case". *Journal of Dynamic Systems, Measurement, and Control,* **119**(4), Dec., pp. 819–821.