LOCAL PEER-TO-PEER COMMUNICATION TO IMPROVE DEMAND RESPONSE IN RESIDENTIAL NEIGHBORHOODS

by

Ryan Schwartz

A thesis

submitted in partial fulfillment

of the requirements for the degree of

Master of Science in Mechanical Engineering

Boise State University

December 2018

BOISE STATE UNIVERSITY GRADUATE COLLEGE

DEFENSE COMMITTEE AND FINAL READING APPROVALS

of the thesis submitted by

Ryan Schwartz

Thesis Title: Local Peer-to-Peer Communication to Improve Demand Response in Residential Neighborhoods

Date of Final Oral Examination: 30th July 2018

The following individuals read and discussed the thesis submitted by student Ryan Schwartz, and they evaluated the presentation and response to questions during the final oral examination. They found that the student passed the final oral examination.

John F. Gardner, Ph.D.	Chair, Supervisory Committee
Thad B. Welch, Ph.D.	Member, Supervisory Committee
Hoda Mehrpouyan, Ph.D.	Member, Supervisory Committee
Aykut Satici, Ph.D.	Member, Supervisory Committee

The final reading approval of the thesis was granted by John F. Gardner, Ph.D., Chair of the Supervisory Committee. The thesis was approved by the Graduate College.

To my fiancée, Taylor, for her love and unwavering support throughout this process.

ACKNOWLEDGMENTS

ABSTRACT

Increased electricity production from intermittent renewables presents a challenge to utilities. Since the grid has little ability to store energy, fluctuations in solar and wind generation require either an increase in generation from expensive sources or a reduction of demand. Demand Response (DR) programs focus on the latter and are designed to increase grid flexibility by allowing grid operators to modify when or how customers use electricity. For residential customers, this typically means shedding load during periods of high demand through a central controller temporarily shutting off air conditioning (AC) compressors. This type of DR can cause spikes in demand after the units come back online.

As the communication and computational capabilities of smart meters and smart thermostats grow, so does the potential to create more decentralized approaches to DR programs. This thesis presents novel thermostat on/off criteria that rely on limited peer to peer communication between a network of residential thermostats. Agent based modeling (ABM) software was used to simulate the emergent behavior that results from thermostat interactions. To demonstrate the benefit of communicating thermostats, the criteria were tested as a means to improve the response following an AC shut off DR event and as an alternative to such events.

The introduced criteria, by sharing only the state of neighboring compressors, improved the overall demand profile following a DR event by reducing peak demand up to 21%. However, it was also found to increase the number of cycles an AC unit experiences by 36%, which can reduce its lifetime. Additionally, the stability implications of this approach are explored.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

ABM – Agent Based Modeling

AC – Air Conditioning

CAES – Compressed Air Energy Storage

COP – Coefficient of Performance

DR – Demand Response

DSM – Demand Side Management

EV – Electric Vehicle

HVAC – Heating, Ventilation, and Air Conditioning

MAS – Multi-agent System

PHES – Pumped Hydroelectric Energy Storage

RTP – Real Time Pricing

TCL – Thermostatically Controlled Load

TOU – Time-of-Use

CHAPTER 1

INTRODUCTION

In response to continued and increasing climate change, the United States and the world continue on a path toward a carbon free energy system, at least as generation capacity growth suggests [2]. According to the Bloomberg New Energy Finance report, renewables will account for 50% of global electricity generation by 2050 due to continued reduction in the cost of solar and wind technology [3]. In the US, 55% of electricity generation will come from renewable sources.

While increasing the use of solar and wind power are an integral part of the solution to climate change, increased penetration of these particular sources introduces other challenges. As a result, it is well understood that a truly carbon-free electric grid will require both more energy storage and greater flexibility from energy consumers [4]. Increased storage, in the form of grid-scale batteries, continues to be the focus of significant investment from both the public and private sectors. On the other hand, efforts to increase grid flexibility, often in the form of programs to encourage and enable electric users to alter their usage in a way to improve overall grid performance, have seen limited investment or deployment.

Demand response (DR), is an approach to increasing grid flexibility by allowing the grid operator to control when or how electricity is consumed by certain customers. DR is well-proven yet its deployment has been limited to programs that shed load during periods of very high demand, when additional generating capacity is scarce

and therefore expensive. The rollout of the Advanced Metering Initiative (so-called smart meters), the rapid growth of computational and communication capabilities in everyday devices (i.e. the Internet of Things), and the increased public support for action to combat climate change, creates an environment conducive to the development of a more robust set of programs that can engage a broader sector of society, greatly increase grid flexibility and by doing so, enable a much higher penetration of wind and solar generation into the grid.

1.1   Challenges to Utilities

Electric utilities have their origins in the late 19th century and have evolved in the technological and social environment of the early 20th century. This has led to the current situation in which consumers have come to expect that electricity will always be available, regardless of when and how a consumer requires it, and utilities accept the responsibility of meeting those requirements. Utilities pull from a variety of generation sources to produce the required power. A subset of these generation sources (i.e nuclear, coal) continually operate, utilizing large thermal masses that are cost efficient at supplying a constant amount of power to the grid, but are unable to respond quickly to changes in demand.

The rapid deployment of intermittent renewable energy generation on the electric grid has presented significant challenges to the entities that are responsible for maintaining the reliability of our electric distribution system. Since the grid has very little inherent ability to store energy, the variable nature of wind and solar energy forces the utility to constantly compensate for those variations by modifying the output of the resources they do control. Typically, this compensation is done by changing the output of hydropower or natural gas generators, which can reach full generation in a

matter of minutes, allowing for tracking of peak loads or filling in the valleys caused by a renewable generation source suddenly losing power production (e.g. cloud covers solar farm, wind quickly dies down, etc.).

To better explain the nature of the problem, Figure 1.1, which is commonly referred to as the duck curve [1] explains one scenario that causes concern for the California Independent System Operator (CAISO, an independent system operator responsible for the California electric grid). While results come from the specific California study, it is well-understood that the challenges revealed here are widely applicable to any region where solar generation is being added to the system. For this reason, the duck curve has become emblematic of the problems associated with high penetration of wind and solar generation in general [5].
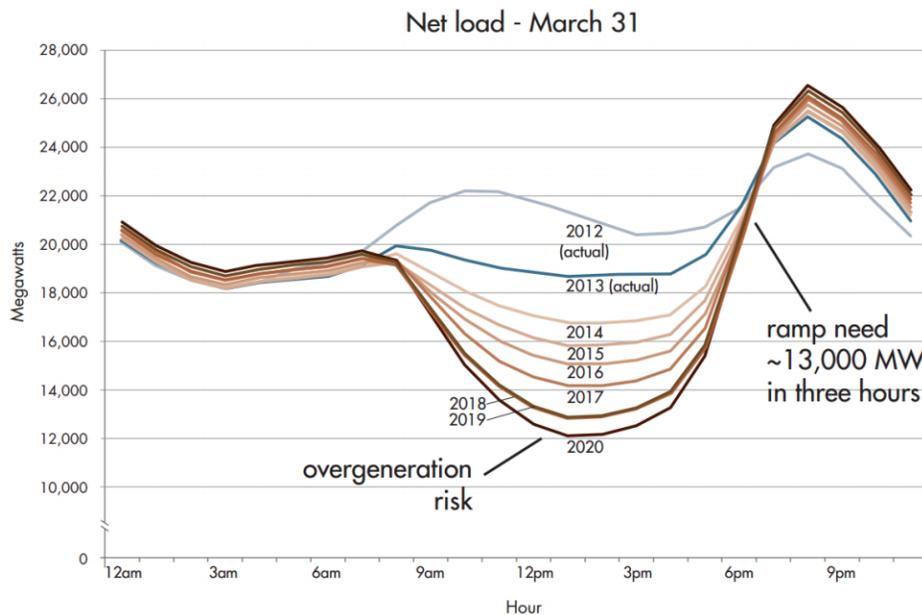
Figure 1.1: The Duck Curve [1]

Figure 1.1 shows a series of lines, each representing the net energy demand in the California grid during a spring day (31 March) on each of 8 years. Net energy is the energy provided by the electric utilities, which is equal to the energy consumed

minus the solar energy produced by distributed (non-utility owned) solar resources. The base line of the graph (labeled 2012 (actual)) is prior to California's incentives to promote installation of customer-owned solar generation. The general shape of that line, with a minimum around 3 AM, a peak mid morning, a lull in mid-afternoon and a high peak around dinner time, is typical of an aggregated electrical demand curve throughout the developed world. Additional lines show the impact that the increasing penetration of solar (2013) and that of the growth of solar penetration predicted at the time of the study. As solar generation increases (because more facilities are built) the net demand that utilities must meet changes drastically during the daylight hours. However, the load required at night is not impacted.

There are two major areas of concern that the duck curve highlights. One is that the greatly reduced net demand at mid-day might lead to over generation due to other constraints on grid operation. Over generation is an issue because the so-called base load generating stations cannot be economically ramped up and down to meet this relatively short period of high solar generation. The second issue is a rapid change in demand that utilities must compensate for as the sun sets across the region. Even those generators that can change their output have limitations as to how fast those changes can occur. Therefore, the concern is that there can come a time when the steepness of the ramp exceeds the ability of utilities to meet rapidly changing needs. The duck curve was first introduced in 2013, and experience since then has shown that it has actually under-predicted the nature of the problems [6].

In addition to these problems exposed by California and other system operators relative to solar energy, other utilities and independent operators find related challenges to the growing impact of wind generation capacity. These concerns are the main motivations for developing grid-level energy storage solutions in the form of electro-chemical batteries. Grid-scale battery deployment is seen as the main solution

to these problems, although they come at considerable expense and the long-term environmental impacts and costs associated with these technologies remain uncertain [7].

1.2 New Technologies and Electrical Distribution Models

Batteries are only one way to tame the duck curve. A growing number of energy observers are pointing out that by enlisting the cooperation of electricity users, and incentivizing changes in their behavior, we can impact the problem by changing consumption at critical times of the day [8]. Increasing development of the electrical smart grid offers unprecedented opportunity for more complex electrical supply and demand interactions in a relationship that has been historically unilateral. The smart grid allows for the application of modern communication technology, such as the internet of things, to improve or modify widespread electrical transmission and distribution.

In addition to developments related to the smart grid, utilities have begun to embrace Demand Side Management (DSM) as a way of helping it meet customer needs. In the broadest sense, DSM refers to those programs implemented by a utility (or independent grid operator) in which they endeavor to modify (or manage) customer use of electricity in some way. DSM programs mainly fall into two categories: (1) Energy Efficiency programs such as LED lighting or insulation upgrades, and (2) Demand response programs utilities use to incentivize deferral of energy consumption during times of high demand.

This shift in the utility-customer relationship has led observers to speculate that we are at the dawn of a new era in which the relationship between energy providers and customers is re-defined in a way that empowers individual consumers to provide

services back to the grid in addition to simply consuming. The terms used to describe this new relationship include the producing consumer, or *prosumer* [9], and the *Energy Internet* [10]. The speculation around the future of the prosumer is, until now, limited to the relationships between individual prosumers and the utility (or with third party aggregators acting on behalf of the utility).

Much of this activity is focused on thermostatically-controlled loads (TCLs) such as those systems used for space heating and cooling, hot water or refrigeration and food storage. Systems that use electricity in this manner are normally designed to maintain temperature, not at a single constant set point, but within a range of temperatures, known as the thermostat deadband.

In a typical DR application, residential air conditioning (AC) compressors (but not the circulating fan) are temporarily turned off under control of the utility, allowing the inside temperature to rise above the thermostat setting. The grid is relieved of the load that the AC compressor would have drawn and the homeowner (if present) experiences a small and possibly noticeable increase in indoor temperature. These programs are nearly universally used to shed load at times of very high demand, but they can also be used to increase consumption at times of energy surplus, resulting in a somewhat cooler home than the set point would imply. A common interpretation of this effect is that energy is being stored in the thermal mass of the home. In this respect, the home acts like a thermal battery, albeit a leaky one with limited storage capacity.

The ability of any individual TCL to impact the energy balance of the grid is limited. Therefore, most applications entail the aggregation of many hundreds and thousands of loads, coordinated by a central controller. The dynamics and control of such aggregated loads is the subject of significant research [8], [11], [12] and grid operators across the country have significant experience using aggregated DR to

manage peak loads.

Agent-based Modeling (ABM) has been a fixture of the social sciences for many years and has proven to be an effective means of understanding emergent behavior from a large number of individual actors. More recently, ABM has been identified as a means of understanding complex physical systems, such as the electric grid, where large numbers of individual homes or even appliances can be modeled in a coherent framework that allows for easy control of the statistical distribution of agent parameters and behavior [13].

The thesis of this work is to use ABM to develop distributed decentralized criteria that rely on limited communication between neighboring residential thermostats to implement a robust DR program for residential AC systems. While this work could be applied to commercial or industrial settings, that is outside the scope of this research.

CHAPTER 2

BACKGROUND

2.1 Demand Response Methods

2.1.1 Economic

One method of demand response is to influence residential customers to shift their electricity usage from peak times to off-peak hours by offering consumers real-time pricing (RTP) or Time-of-Use (TOU) pricing. Instead of charging the customer a flat rate for electricity, the price changes throughout the day, encouraging consumers to use energy during times of less load. The effectiveness of this method has been widely studied in a range of different settings. The results have been mixed. Typically, RTP does reduce peak demand, but the amount differs.

Sweden reached full smart meter coverage in 2009, which resulted in regulations giving consumers the option to have their electricity metered hourly [14]. This hourly data led utilities to offer variable-pricing contracts, which can be economically beneficial to the electricity consumer compared to flat rate contracts. They also enable greater adoption of other smart grid storage technologies, such as electric vehicle (EV) charging control schemes, since owners can charge overnight while prices are lower. In the study by Nilsson et al. [15], half of the participating Swedish households were given RTP, and the other half had a conventional fixed rate. The results showed that the RTP households shifted an average of 5% of their total daily energy consumption

to off peak hours. However, during non-peak night hours, Sweden utilizes more coal and gas to meet energy needs. This resulted in an increase of almost 3% in annual CO2eq emissions for the RTP households.

In Japan, Zhang et al. [16] found dynamic electricity pricing can reduce peak demand by 6% to 14%, depending on factors such as the outside temperature and floor area of the households.

There has also been some research done on this topic in the United States, where a majority of customers pay a fixed rate for their electricity. Allcott [17] found that RTP results in households reducing overall energy consumption by about 5% through energy conservation during peak hours. However, Lutzenhiser et al. [18] examined the results of a pilot program in California and found that implementing a TOU price scheme resulted in minimal load reduction.

2.1.2 Energy Storage

Smoothing the demand curve can be achieved by utilizing energy storage methods. While demand is low, energy storage systems can draw power to "charge" their reservoir. This stored energy can respond quickly when extra power is needed during peak demand.

The most common utility-scale method to store energy is in the form of water, using pumped hydroelectric energy storage (PHES). During periods of low demand, water is pumped uphill into a reservoir. At times of high demand, that water is dropped through turbines to generate power. To be economical, these reservoirs must be large in size, which limits PHES to mountainous or hilly regions. Despite this geographical constraint, PHES accounts for about 98% of bulk electricity storage globally, according to the Electric Power Research Institute (EPRI) [19].

Another large scale energy storage technique is compressed air energy storage (CAES). A CAES system utilizes gas turbine technology. When the cost of electricity is low, the system compresses air and pumps it into a large storage reservoir. When additional power is desired during peak demand, the compressed air is heated and expanded through the turbine. Since being introduced in the 1940s, a number of CAES systems have been installed around the world, with the largest plants able to generate around 300 MW [20].

Electric battery storage became a viable option for energy storage in recent years as the cost of batteries continues to decline. California is leading the way, with three large storage facilities installed within the last year totaling 77.5 MW [21]. While charging and discharging these batteries is an effective means of DR, the battery production process can negatively impact the environment [7].

2.1.3 Thermostatically Controlled Loads

A well-established tool to provide DR is through TCLs, which can be manipulated such that they can serve as thermal batteries. This type of control usually refers to space conditioning appliances, such as air conditioning units and furnaces [22], but it can be extended to smaller demand appliance such as water heaters and refrigerators [23]. By specifying when a population of TCLs are consuming power, demand can be controlled. As discussed later in section 2.3, this can be achieved through a centralized controller directing all TCLs within the population or through decentralized control, where the individual TCLs coordinate directly with each other.

2.2 Agent Based Modeling

Agent Based Models (ABMs) of Multi-Agent Systems (MAS) have been used as a simulation tool to model a variety of problems, covering a wide array of disciplines. An ABM consists of a population of unique and autonomous entities, known as agents, that interact with each other and the environment through a set of rules. These agents may be animals, humans, institutions, vehicles, or anything whose behavior can be described by a mathematical model. Agents act as individuals and can communicate and exchange information with each other, but typically the interaction is local, meaning the agents only share information with their neighbors rather than the whole population.

One of the key features of an ABM is the ability to capture and reproduce the phenomenon of weak emergence, which is the development of complex and unexpected behaviors that result from the interaction of a number of simple (but often nonlinear) agents and the enveloping system [24]. Unlike the concept of strong emergence, the aggregate behavior of weak emergence can be derived with adequate computational capabilities [25].

In recent years, ABMs have been used to model how diseases like cholera [26] and Zika [27] are spread in population centers. To develop drought management plans, researchers have used ABMs to simulate an urban water supply [28]. Sociologically, ABMs have been used to simulate problems from tax compliance and evasion [29] to the spread of opinions among members of a society [30]. Another common area of agent based modeling is to model market behavior, whether that be how the behavior of stock traders influences the overall market [31] or determining the cause of the housing bubble that led to the US financial crisis [32].

As it relates to energy, ABMs can be used to optimize building climate control

strategies by predicting occupant behavior and room usage [33]. Concerning demand response, there are a number of papers that focus on centrally controlled aggregated models of HVAC (Heating, Ventilation, and Air Conditioning) MAS [13, 22, 34]. These models focuses on keeping track of the overall system behavior, rather than looking at how individual loads interact with each other. An ABM was used to develop decentralized DR through price signaling that led to "emergent coordination" among the agents, even though the agents don't interact with each other directly [35]. Alternatively, communication between agents has been used to model the charging scheduling for electric vehicles as a means to reduce peak demand and follow renewable generation sources [36].

2.3 Control of Distributed Systems

While centralized control of TCLs is more common, due to simpler implementation, experts have noted a number of potential problems with grid stability and higher peaks before or after a DR event [37]. Also, as the number of loads being controlled increases, so does the computational power and required infrastructure to effectively communicate in a reliable and timely manner. In centralized control, cyber security is an issue because of its vulnerability to single point failure such as a targeted cyber attack against the central controller that can disrupt the control system and expose customer data [38]. Decentralized control has the potential to overcome these shortcomings.

2.3.1 Centralized Control

Grid operators typically exert central control over large populations of TCLs through radio links or power line carrier protocols. For example, Idaho Power has

an AC Cool Credit program where individual consumers, in exchange for a small reduction in their monthly electric bill, give the utility permission to install equipment on their AC unit [39]. This equipment can receive a signal from Idaho Power that shuts off the unit's compressor for a short period of time. After that time has expired, the unit is allowed to turn back on while another population is disconnected, thus reducing the overall demand for as long as required and preventing uncomfortable conditions from occurring within any individual homes.

Another common method of controlling TCLs is through set-point control. Callaway [8] makes the case that changing thermostat set-points of a population of TCLs can be used to follow the variability of wind generation. Building on this work, Bashash and Fathy [22, 40] developed a model that uses a centralized controller to broadcast a uniform signal to vary the thermostat set-point temperature of the population of TCLs. This enables the tracking of a real wind power trajectory.

A centralized controller can coordinate many agents without knowing individual agent states. By comparing a reference to a received aggregate output value, the controller determines what signal to broadcast to all agents. Each agent then makes a decision based on the signal and defined probabilities [41].

Similarly, Zhang et al. [13] developed a control scheme where the centralized control signal is broadcast to all agents. The agents then decide how to implement the signal based on their local temperature and power state.

A priority-stack-based control strategy can be an effective way to control TCLs [34]. Sorting the population of TCLs by temperature into two stacks, one where the TCLs are off and the other where the TCLs are on, allows for the most appropriate selection of the next TCL to turn on or off. When the grid has excess power, perhaps due to an increase in wind generation, the centralized controller can send a signal to the TCLs. The TCL with the highest priority in the off stack will turn on first and

then continue down the stack until the excess power is being utilized. This also works the other way. When the grid is trying to reduce load, the TCLs in the on column turn off sequentially, in order of priority, until the desired reduction has been met. The downside of this model is the need to have the information of all agents in the system accessible to the central controller to sort correctly.

2.3.2 Decentralized Control

In decentralized control, agents take in information and make decisions as individuals, rather than being told what to do by one controller.

One method of decentralized demand response is to use price signalling and adaptive mechanisms, coupled with smart meters, to prevent loads from syncing up and creating high peak demand [35].

With the rising number of electric vehicles, an opportunity for demand response has been created. EV charging management allows for the vehicles to charge at times of low demand or to match renewable energy production. Xydas et al. [36] developed a model to demonstrate the effectiveness of "responsive" EVs. These vehicles determine their charging schedule according to a signal that takes power demand and generation forecasts into account. Their model demonstrated that responsive EVs could reduce the peak charging demand of all EVs, including unresponsive EVs, by shifting demand to a time when the unresponsive EVs were finished charging. They also demonstrated the ability of responsive EVs to charge in response to a real time photovaltic (PV) generation profile.

CHAPTER 3

AGENT BASED MODEL DEVELOPMENT

The agents in the model are residential houses equipped with a smart thermostat that controls the air conditioning unit. These agents interact with the environment (outside temperature) and can communicate with neighboring houses.

3.1 Thermal Dynamics of HVAC Systems

In order to understand how the environment and the house interact, the heat transfer mechanics between them must be examined. Building off of the work of Bashash and Fathy [22], the thermal dynamics of a house can be approximated as a first-order ordinary differential equation:

$$\dot{T}(t) = \frac{1}{RC}\left(T_\infty - T(t) + R\left(Q_I - m(t)Q\right)\right)$$

(3.1)

Here, T and T_∞ correspond to the internal and the ambient temperatures (°C), respectively. The thermal capacitance, C (kWh/°C), and thermal resistance, R (°C/kWh), are properties related to factors such as building insulation and materials. Q_I (kW) is the heat generated by internal loads and Q (kW) is the load's cooling power when ON. The ON/OFF signal $m(t)$ is controlled by the thermostat and the corresponding temperature limits

$$m(t) = \begin{cases} 0, & \text{if } m(t) = 1 \quad \& \quad T(t) \leq T_{min} \\ 1, & \text{if } m(t) = 0 \quad \& \quad T(t) \geq T_{max} \\ \text{unchanged}, & \text{otherwise} \end{cases} \qquad (3.2)$$

where T_{min} and T_{max} are the lower and upper limits of the thermostat deadband, δ. The setpoint temperature, T_{sp}, is related to these limits as follows:

$$T_{min} = T_{sp} - \frac{\delta}{2}, \qquad T_{max} = T_{sp} + \frac{\delta}{2} \qquad (3.3)$$

Considering a population containing N TCLs, the total load can be expressed as

$$P_{\text{TCL}}(t) = \sum_{i=1}^{N} \frac{1}{\eta_i} Q_i m_i(t) \qquad (3.4)$$

where η_i is the coefficient of performance (COP) of the i^{th} load.

3.2   Agent Development

There exist numerous ABM software currently available to the general public, such as AnyLogic, NetLogo, MASON, Repast, and AgentSheets [42]. While the mathematics that drive an agent based model can be developed using line-by-line code (eg. MATLAB), these platforms are useful as they allow for rapid development of a large number of simple agents, control of the statistical spread of parameters, and ease of result reporting. For this research, AnyLogic is used because it allows for combination of agent-based modeling, discrete events, and system dynamics in a visually clear manner [43]. First, the individual agent is modeled and then a population of agents can be generated to simulate the aggregate dynamics.
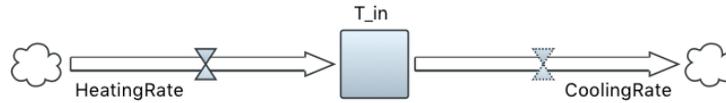
Figure 3.1: AnyLogic Stock and Flow Diagram

3.2.1 System Dynamics

To represent differential equations, AnyLogic makes use of "Stock and Flow Diagrams" which are typical in economics, accounting, and related disciplines. The Stock represents a reservoir, or continuous state, that is changing with the incoming and outgoing Flows. Figure 3.1 shows the house dynamics described in Equation 3.1, where the continuous state is the indoor temperature, which increases due to the incoming flow of heat transfer and decreases due to the AC unit's ability to remove heat. Therefore, the heating rate consists of the heat transfer due to the difference between outdoor and indoor temperature, $\frac{1}{RC}(T_\infty - T)$, and the internal heat generation, Q_I/C. For this model, the internal loads are treated as a disturbance similarly to Hao et al. [34]. The cooling rate resulting from the AC unit is P/C when ON and 0 when OFF.

3.2.2 State Charts

Individual AC unit state behavior is modeled using a state chart. The state chart keeps track of what state the AC unit is in and transitions between states if certain thermostat criteria are met. As seen in Figure 3.2, the AC has three states (ON, OFF, Locked). The Locked state represents the compressor time relay delay, which prevents damage due to rapid ON/OFF cycling. The compressor "locks" for 3-10 minutes, during which a switching signal from the thermostat is ignored [13]. A
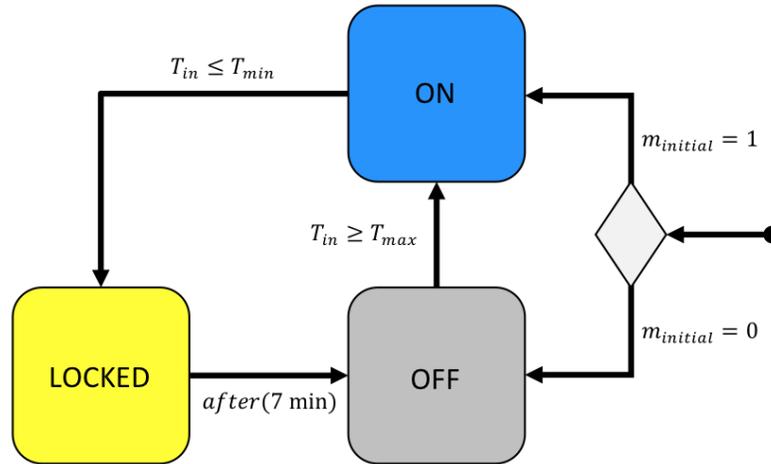
Figure 3.2: AC unit state chart

lockout time of 7 minutes, the middle of this range, is used for this research. This locked state does not impact the building comfort level, but it can affect the aggregate response if the control frequently switches states of participating units. In normal thermostat operation, the transition from OFF to ON and ON to Locked occurs when the indoor temperature reaches the upper or lower limit of the thermostat deadband. The additional transitions from the diamond state to ON and OFF determine the initial state of the AC unit.

3.3 Agent Parameters

Both homogeneous and heterogeneous population of agents are modeled in this research. Bashash and Fathy [22] assumed homogeneous house parameters (Table 3.1). This a good representation assuming the modeled population of houses is tract housing, sometimes referred to as cookie cutter neighborhoods, where all the houses are very similar in design (Figure 3.3a).

Many neighborhoods contain a mix of houses that vary in age, size, and building

<div align="center">(a)   (b)</div>

Figure 3.3: Example of (a) Homogeneous Neighborhood and (b) Heterogeneous Neighborhood in Boise, Idaho via Google Earth.

Table 3.1: Population Parameter Values

Parameter	Value	Standard Deviation (Heterogeneous)
R, Thermal resistance	2 °C/kW	0.1 °C/kW
C, Thermal capacitance	10 kWh/°C	0.5 kWh/°C
P, Energy transfer rate	14 kW	-
η, Load efficiency	2.5	0.125
T_{sp0}, Initial setpoint temperature	20 °C	1 °C
T_∞, Ambient temperature	32 °C	-
δ, Thermostat deadband	0.5 °C	0.025 °C

materials (Figure 3.3b). To develop a model for these neighborhoods, a heterogeneous set of parameters must be used. This can be done by creating a statistical distribution around the homogeneous values. A parameter spread of $\pm 15\%$ is desired, which results in the normal standard deviations found in Table 3.1.

As a precaution, a lognormal distribution is calculated from the normal distribution values. For this application, these lognormal distributions are essentially the same as the normal distributions upon which they are based, but they ensure that all parameters are positive, since it doesn't make sense for certain parameters to be

negative (e.g thermal resistance and deadband). Details of this distribution can be found in Appendix B.

The energy transfer rate of a house's AC unit is sized depending upon the thermal dynamics of the house. The homogeneous population of houses' 14 kW is equivalent to a 4 ton unit (1 ton = 3.5 kW_{th}), which, for these parameters, means that the cooling rate is 0.8 °C/hr, or the temperature moves from the upper limit of the deadband to the lower limit in about 37.5 minutes. The necessary tonnage to achieve this cooling time for the heterogeneous population was calculated and then rounded up to the nearest half-ton to reflect sizes commercially available. The resulting range in unit sizes is 3.5-5 tons (12.25-17.5 kW_{th}). Rounding up of the unit size results in slight over sizing, which means some houses will cooler faster than 37.5 minutes and therefore cycle more often than their homogeneous counterpart. The minimum cooling time for a heterogeneous house is 30.4 minutes.

Bashash and Fathy's model [22] maintained a constant ambient temperature, which is not realistic, since temperature can drastically increase from morning to evening during the summer months (e.g. Figure 3.4). To improve upon their model,
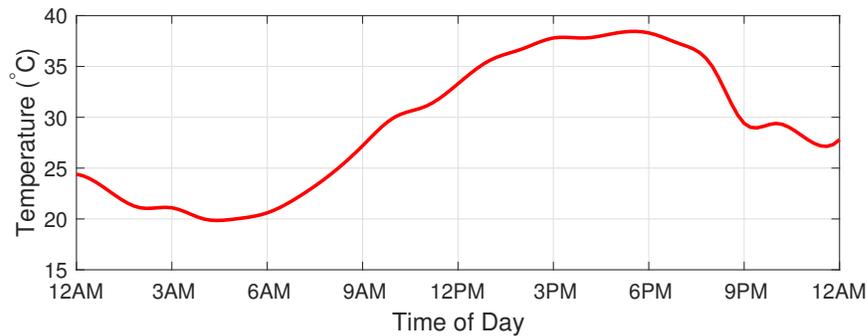
Figure 3.4: Example of hourly TMY data for Boise (July 21st) fitted with a spline.

hourly typical meteorological year (TMY) data for Boise, Idaho [44] for the week

of July 21 to July 28 is used to represent a realistic summer temperature profile. AnyLogic applies a spline fit to the data and passes it to the model.

3.4   Agent Population

An agent population of 100 houses is used in this research. This is a large enough number to produce meaningful results, but small enough for rapid simulation testing in AnyLogic. Each house is randomly assigned parameters and the start state of ON or OFF is evenly split among them.

3.4.1   Baseline Behavior

Baseline behavior occurs when each house maintains its desired temperature, independently from the rest of the population. As seen in Figure 3.5, the indoor temperature oscillates back and forth between the limits of the deadband in response to the changing ambient temperature.
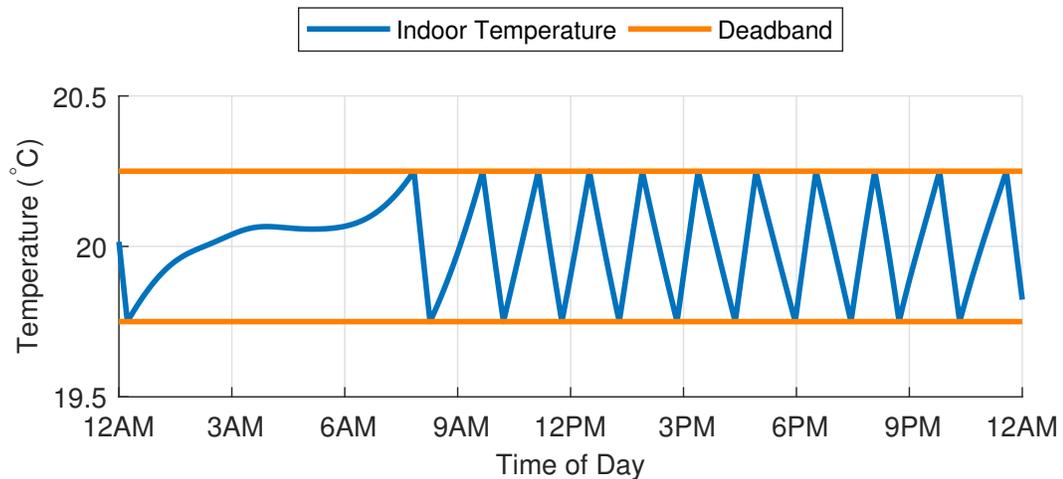
Figure 3.5: Example of individual homogeneous house indoor temperature profile for the day of July 21st.
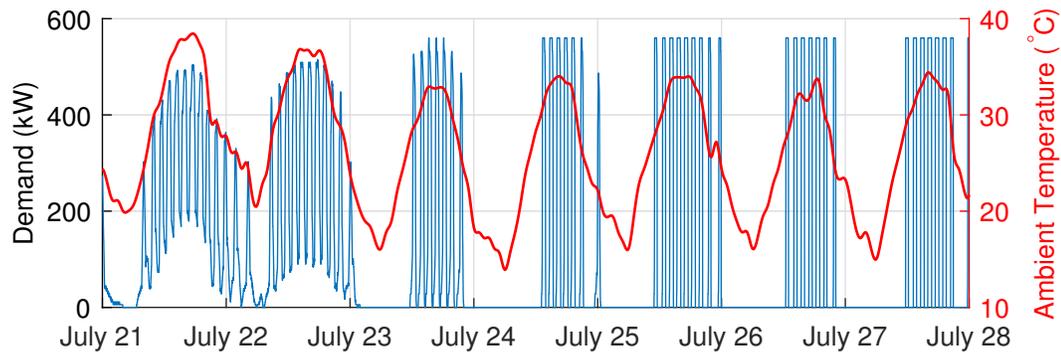
Figure 3.6: Demand profile for population of homogeneous houses over the course of one week.
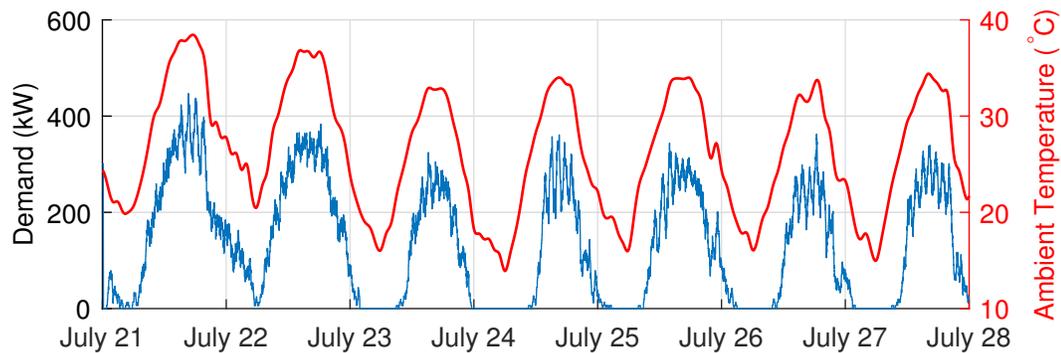
Figure 3.7: Demand profile for population of heterogeneous houses over the course of one week.

Running this simulation over the course of one week results in a measurement of what would be considered normal operation for a neighborhood where there are no attempts to regulate demand. For a homogeneous set of houses (Figure 3.6), the identical system dynamics of the thermal load result in a syncing up of AC units. The heterogeneous population's spread of parameters creates system dynamics with various time constants that prevents all units from cycling ON at the same time (Figure 3.7). In both cases, the demand profile follows the changing ambient temperature, peaking in the early evening and reaching a minimum of zero at night when the temperature drops below the thermostat set point.

Assessing the performance of the simulations is based upon four indices, the peak demand, the total energy, the number of times an AC unit cycles ON, and the time spent above the deadband. The indices for the baseline week can be found in Table 3.2. The number of cycles is important because an increase in cycling can result in a reduced lifetime of the unit. The average number of cycles of all the AC units and the value of the AC unit that cycles the most are included. The time above the deadband gives a rough idea of the cost to comfort level. The longer a house stays above the desired set point temperature, the more likely the occupants' dissatisfaction. All of these indices are calculated for a given time window: a week for typical operation and 24 hours (6am-6am) for DR events. For the baseline week and DR event with no participation, the time spent outside the deadband is zero due to typical thermostat operation.

Table 3.2: Indices of Performance - Baseline Week

Population	Peak Demand (kW)	Total Energy (kWh)	Cycles/week Avg	Max
Homogeneous	504	21738	62.78	63
Heterogeneous	446.9	21974	64.89	92

3.4.2   Demand Response Event

Typically DR events are scheduled during times of peak load. Focusing on the first day (July 21) of baseline results, a peak of 504 kW occurs at 5:31 PM for the homogeneous population. For the heterogeneous population, a peak of 447 kW occurs at 4:42 PM. To prevent these peak demand values, a DR event is initiated six minutes before each peak time and lasts fifteen minutes. During this time a specified percentage of the houses are forced OFF. A DR program is simulated using three levels of participation (10%, 20%, 30%). The extreme limit of DR event participation

is 100%, which also represents complete neighborhood power failure. The demand profile post DR event for the population of homogeneous houses is shown in Figure 3.8.

(a) 10% OFF

(b) 20% OFF

(c) 30% OFF

(d) 100% OFF
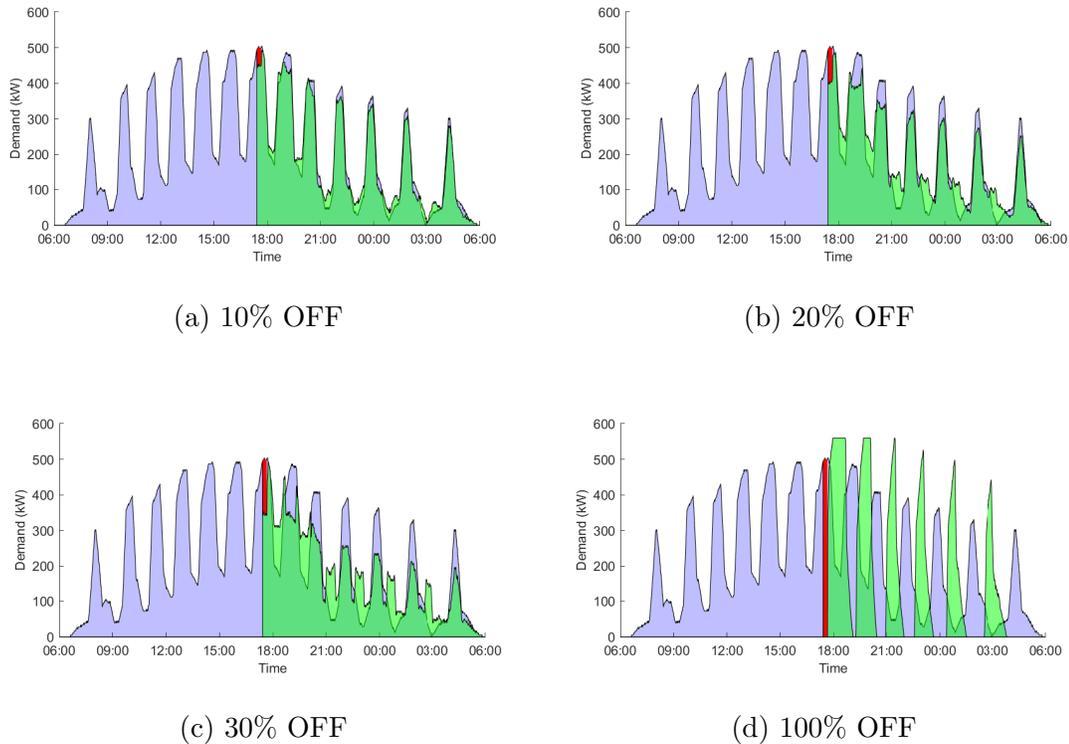
Figure 3.8: Homogeneous population DR event response for various percentages of participating homes. Blue is the baseline demand profile, the red area represents the energy saved during the DR event, and green is the demand profile post DR event.

The data compiled in Table 3.3 shows the homogeneous population DR program responses reduce the peak demand and total energy during the 6am-6am time window when compared to no participation. The full participation event causes the houses to sync up even further, resulting in a larger peak demand, but because all houses were off for 15 minutes the energy consumed was 104 kWh less than the baseline. All events had similar average number of times the units turned ON.

The heterogeneous population DR event responses (Figure 3.9) have a damping

Table 3.3: Indices of Performance - Homogeneous DR Events

Turn OFF	Peak Demand (kW)	Total Energy (kWh)	Cycles/day Avg	Max	Time above δ (min) Avg	Max
None	504.0	4918.5	12.94	13	-	-
10%	492.8	4907.1	12.93	13	1.8	40.1
20%	492.8	4897.0	12.93	13	1.4	37.0
30%	492.8	4882.9	12.93	13	2.2	41.6
100%	560.0	4803.7	12.89	13	8.1	41.6

effect due to the parameter spread. This is most noticeable in the 100% participation case, where the response returns to baseline levels within a few hours of spiking.

(a) 10% OFF

(b) 20% OFF

(c) 30% OFF

(d) 100% OFF
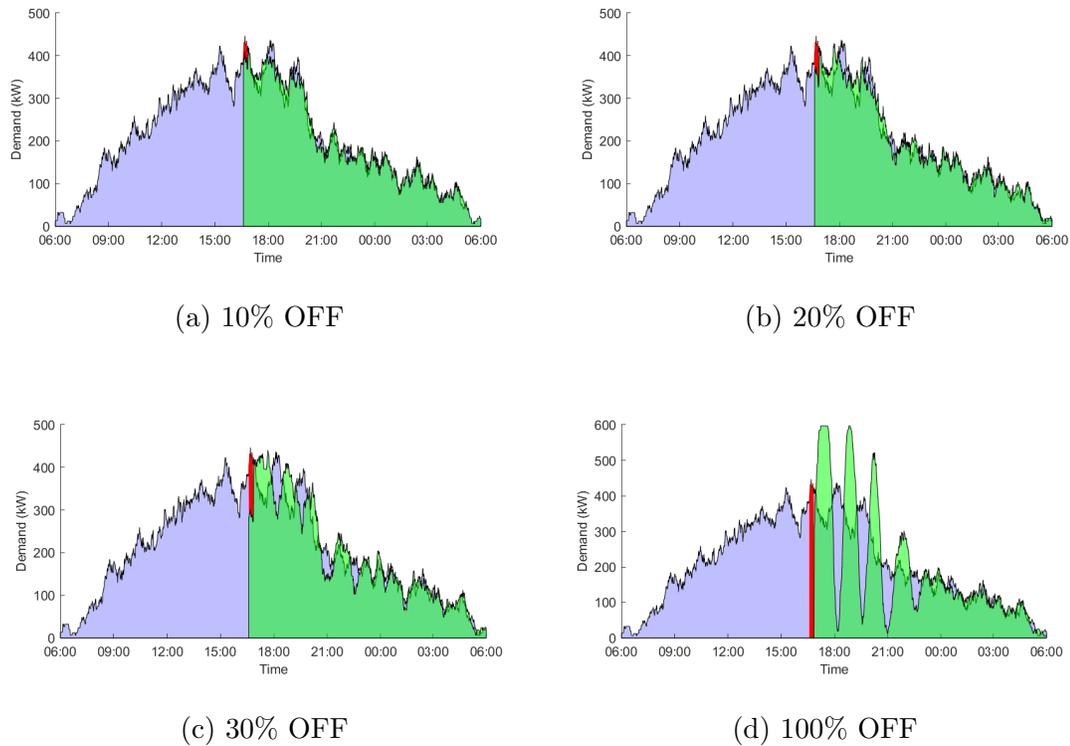
Figure 3.9: Heterogeneous population DR event response for various percentages of participating homes. Blue is the baseline demand profile, the red area represents the energy saved during the DR event, and green is the demand profile post DR event.

The indices of performance for the heterogeneous population (Table 3.4) show the DR program simulations resulted in a smaller peak demand and the total energy

stayed approximately the same. As expected, due to the AC sizing differences the number of cycles per day is increased relative to the homogeneous population.

Table 3.4: Indices of Performance - Heterogeneous DR Events

Turn OFF	Peak Demand (kW)	Total Energy (kWh)	Cycles/day		Time above δ (min)	
			Avg	Max	Avg	Max
None	446.9	4901.6	13.27	18	-	-
10%	423.5	4905.2	13.31	18	1.4	32.2
20%	428.0	4904.5	13.28	18	2.5	36.5
30%	440.6	4906.0	13.36	18	3.1	33.4
100%	597.1	4902.6	13.50	18	9.8	42.0

CHAPTER 4

DYNAMICS AND STABILITY OF

AGGREGATED AGENTS

4.1 Agent Connection

There are various network types for connecting agents within AnyLogic, such as random, distance based, ring lattice, and scale free. The ring lattice network (Figure 4.1) is used for this model because an equal number of connections per agent is desired. A ring lattice is also an approximation of a nearest neighbors network, where each house is connected to the specified number of closest agents.
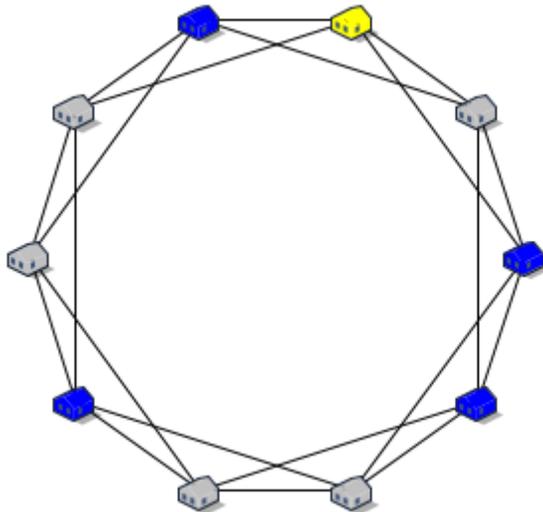
Figure 4.1: Ring lattice example containing ten houses, each connected to its four closest neighbors. Colors correspond to the state of that agent (See Figure 3.2).

4.1.1 Agent Communication

AnyLogic has built-in messaging ability, which allows information to be shared among connected agents. For this research, messages containing relevant parameters are sent at a rate of once per minute to ensure behavior reacts in a timely manner to changes in the system. The agent can then use that information to drive its behavior.

4.1.2 Graph Theory Basics

Graph theory is the mathematical lens through which this network of connected houses is viewed. As Barabási [45] explains, in graph theory the network (or graph) is described as a set of nodes (the agents) and edges (links between agents). The degree, d, of a node describes the number of connections that node has to other nodes. These connections between agents can be *directed* or *undirected*. In a directed link, connection is established in one direction from one agent to another, similar to citations in a paper or a webpage linking to another webpage. Other networks utilize undirected links, like the power grid where transmission line current can flow both directions. The connections between residential thermostats in this model are assumed to be undirected because connected houses know the ON/OFF state of each other's AC units.

Networks of connections are often represented as an adjacency matrix, $\mathscr{A}$. For a network containing N nodes, the adjacency matrix has N rows and N columns containing elements that follow the rules:

$$\mathscr{A}_{ij} = \begin{cases} 0, & \text{if nodes } i \text{ and } j \text{ are not connected to each other} \\ 1, & \text{if nodes } i \text{ and } j \text{ are connected to each other} \end{cases} \tag{4.1}$$

For an undirected network the adjacency matrix is symmetric, $\mathscr{A}_{ij} = \mathscr{A}_{ji}$, and since a

house is not connected to itself, the diagonal consists of zeros. The adjacency matrix can be used to find the degree of house i by summing either the column or the row corresponding to that house:

$$d_i = \sum_{j=1}^{N} \mathscr{A}_{ij} = \sum_{j=1}^{N} \mathscr{A}_{ji} \tag{4.2}$$

4.1.3 Network Connectedness

A graph is considered connected in graph theory, when a path exists between every pair of nodes. A lack of connectedness means there is more than one graph with no means of information sharing between them. Whether or not a graph is connected can be determined by examining the eigenvalues of the graph's Laplacian matrix:

$$L = D - \mathscr{A} \tag{4.3}$$

where D is the matrix whose diagonal contains the degree of each node, $d_1, ..., d_n$. By definition, the Laplacian matrix always has a zero eigenvalue. The second smallest eigenvalue, known as the algebraic connectivity of the graph, is greater than zero if and only if the graph is connected [46]. For the purposes of this research, a ring lattice network with $d = 4$ (e.g. Figure 4.1) and 100 houses is used. The second smallest eigenvalue for this network configuration is 0.0197, so the graph is connected.

4.2 Additional Parameters

Consider a new non-dimensional temperature parameter, θ_i, where the bottom of the deadband is $\theta_i = 0$ and the top of the deadband is $\theta_i = 1$.

$$\theta_i = \frac{T_i - T_{min,i}}{T_{max,i} - T_{min,i}} = \frac{T_i - (T_{sp,i} - \frac{\delta_i}{2})}{\delta_i} \tag{4.4}$$

Now typical thermostat behavior can be described in terms of this normalized parameter instead of individual house temperatures and deadbands:

$$m(t) = \begin{cases} 0, & \text{if} \quad \theta \leq 0 \\ 1, & \text{if} \quad \theta \geq 1 \\ \text{unchanged}, & \text{otherwise} \end{cases} \tag{4.5}$$

Consider a situation in which the m of each house can be broadcast from its thermostat to connected thermostats, allowing for the calculation of an average ON/OFF state of connected houses. This value is calculated by each load, excluding its own m value. The adjacency matrix representing connected agents can be used to easily calculate these values:

$$\tilde{m} = \frac{1}{d}\mathscr{A}m \tag{4.6}$$

Let us also consider the variable, λ, which represents a willingness to participate. The λ parameter allows the homeowner to change their participation level in the same manner as changing their thermostat set point. This parameter has a value from 0 to 1, where a setting of 0 means the homeowner is not willing to participate and setting of 1 corresponds to maximum participation level. In order to get customers to increase their willingness to participate, the utility may have to use incentives such as different price schedules depending on participation levels (similar to TOU pricing). Each house knows the state and willingness to participate of its connected neighbors. Developing state transition criteria based on this information is used to drive desired aggregate responses.

4.3 Criteria Development

4.3.1 Emergent Behavior

One of the goals of this work is to use ABM to look for beneficial emergent behavior from this connected network. By developing criteria that rely on information sharing between the network of thermostats, there is potential for complex aggregate behavior to emerge that reduces peak demand while smoothing the demand profile. The key to unlocking emergent behavior is the non-linear nature of the thermostats. Linear systems produce predictable results, since superposition applies. Non-linear systems are less predictable, which can lead to unexpected results.

4.3.2 The $\tilde{m}$ Model

Here, a new addition to the thermostat model is proposed which uses the average state of the surrounding units, $\tilde{m}$, to inhibit operation based on the number of connected units that are operating.

$$
m(t) = \begin{cases} 0, & \text{if} \quad \theta - k\tilde{m} \leq 0 \\ 1, & \text{if} \quad \theta - k\tilde{m} \geq 1 \\ \text{unchanged}, & \text{otherwise} \end{cases} \tag{4.7}
$$

The addition of the average ON/OFF state of connected neighbors allows agents to reduce overall demand by causing an earlier entry to the OFF state if a larger number of neighbors turn ON. For example, consider a house where two of its neighbors are ON, resulting in a $\tilde{m}$ of 0.5. Assuming $k = 1$, this house will turn OFF as soon as $\theta = 0.5$, instead of the standard $\theta = 0$.

This unbounded $\tilde{m}$ criteria (Equation 4.7) is the more general approach and lends itself to linearization. However, this criteria also changes the upper limit at which a unit will turn ON, from 1 to $1 + k\tilde{m}$, which results in the houses spending more time at warmer temperatures above the deadband. Replacing $\theta - k\tilde{m} \geq 1$ with the $\theta \geq 1$ criterion from the standard thermostat criteria (Equation 4.5) bounds the response and reduces the time spent outside the deadband. The lower deadband limit, $\theta \leq 0$, is redundant and therefore not included in the bounded $\tilde{m}$ criteria (Equation 4.8).

$$m(t) = \begin{cases} 0, & \text{if} \quad \theta - k\tilde{m} \leq 0 \\ 1, & \text{if} \quad \theta \geq 1 \\ \text{unchanged}, & \text{otherwise} \end{cases} \qquad (4.8)$$

The bounded $\tilde{m}$ criteria produces a more desirable demand response and therefore is the criteria used for simulation testing.

4.3.3 Determination of k

In order to implement the bounded $\tilde{m}$ criteria, a value for the gain, k, must be decided upon. This was done by running multiple simulations of the 100% participation DR event for the homogeneous population, where the deadband bounded $\tilde{m}$ criteria was implemented upon completion of the DR event. Table 4.1 shows the results of the simulation for a spread of k values, indicating that the region from 0.8 to 0.9 is worthy of further exploration.

As k increases, the number of cycles and time spent above the deadband increases, but the peak demand levels out to 492.8 kW somewhere between $k = 0.8$ and $k = 0.9$. The demand value of 492.8 kW is the value at initiation of the DR event and the peak proceeding the DR event. A peak demand value of 492.8 kW demonstrates successful

Table 4.1: Indices of Performance - k values from 0.5 to 1.5

k	Peak Demand (kW)	Total Energy (kWh)	Cycles/day Avg	Max	Time above δ (min) Avg	Max
-	560.0	4803.7	12.89	13	8.1	41.6
0.5	560.0	4875.0	14.47	15	8.1	41.7
0.8	520.8	4862.2	16.70	19	8.5	41.7
0.9	492.8	4865.2	17.99	21	16.5	71.0
1.0	492.8	4858.9	19.42	22	50.1	124.7
1.1	492.8	4846.3	21.03	24	73.1	114.4
1.5	492.8	4822.6	32.41	44	258.9	376.7
2.0	509.6	4718.1	68.43	85	605.5	754.7

criteria implementation, because the response to the DR event is prevented from spiking above the pre-DR event levels. Since fewer cycles and less time spent outside the deadband is desired, further gain values were tested. As shown in Table 4.2, the desired peak demand of 492.8 kW is reached at $k = 0.86$, but the Total Energy, average and max Cycles ON over the course of the day, and time spent above the deadband are lower for $k = 0.87$. Therefore, a gain constant of $k = 0.87$ is used for this research.

Table 4.2: Indices of Performance - k values from 0.8 to 0.9

k	Peak Demand (kW)	Total Energy (kWh)	Cycles/day Avg	Max	Time above δ (min) Avg	Max
0.8	520.8	4862.2	16.70	19	8.5	41.7
0.85	498.4	4868.4	17.39	19	11.4	41.7
0.86	492.8	4857.4	17.47	19	12.5	52.5
0.87	492.8	4856.2	17.45	19	12.4	45.4
0.88	492.8	4862.5	17.70	19	12.8	51.2
0.89	492.8	4852.9	17.71	20	14.0	59.0
0.9	492.8	4865.2	17.99	21	16.5	124.7

4.3.4 Stability

The addition of the $\tilde{m}$ criteria introduces a potentially destabilizing influence and rigorous stability criteria need to be developed. Attempts to mathematically find a stability limit for the gain, k, were unsuccessful. Simulations don't show a clear stability limit, only that there is a degradation in performance as k increases, which can be seen in the demand profiles following the 100% DR event for the homogeneous population (Figure 4.2). When $k = 1.0$, the demand profile stays below the pre-DR event peak value with little oscillation. At $k = 2.0$ the demand spikes above the pre-DR event level and the rapid oscillations cause the profile to appear as a thick band.
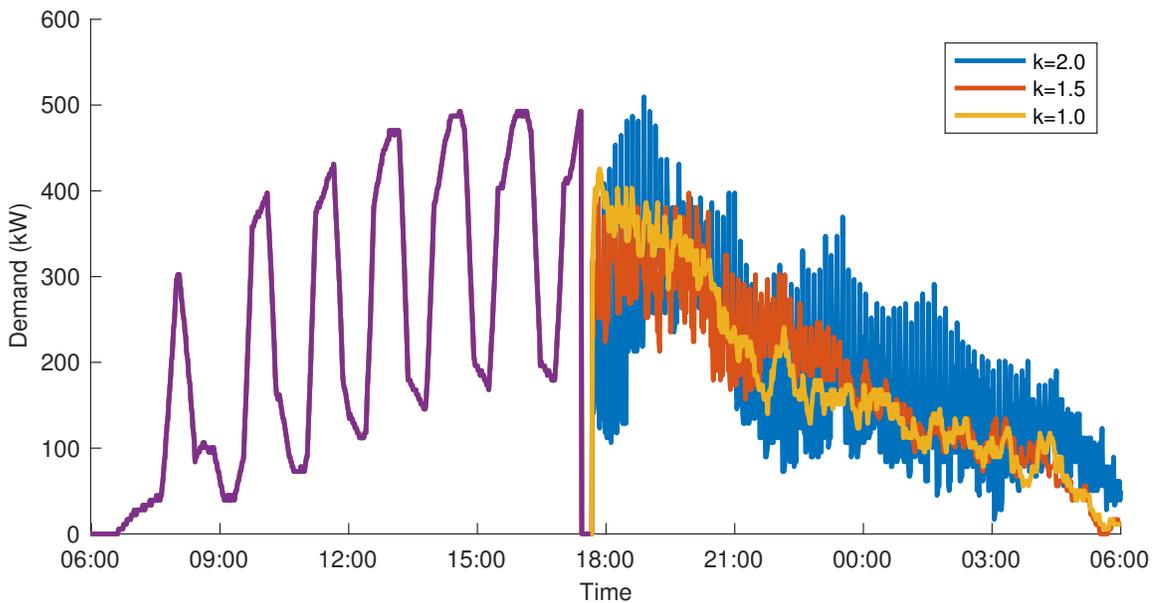
Figure 4.2: 100% DR event for homogeneous population at various gain values.

This degradation is also visible when examining the average time spent above the deadband and the average number of cycles (Figure 4.3). Around $k > 1$, the number of cycles and time above the deadband increase at a faster rate.
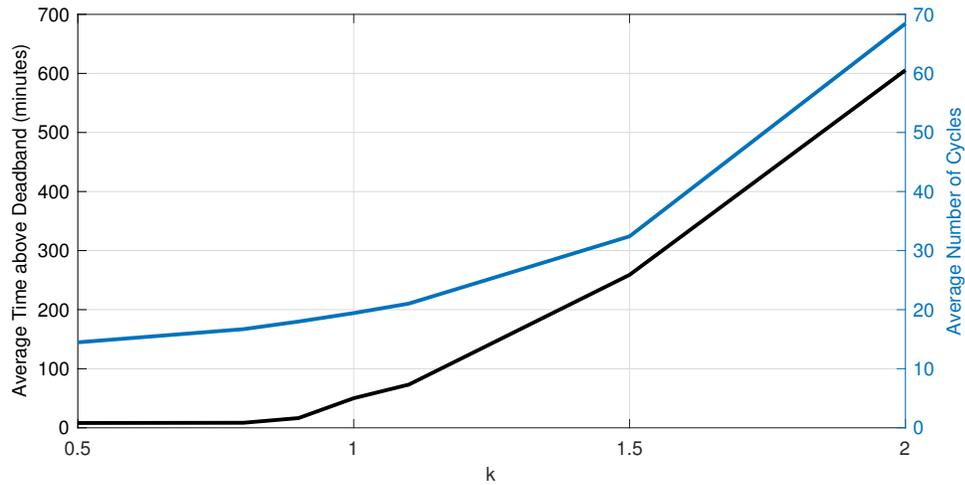
Figure 4.3: Average cycles and average time above deadband over course of 24 hours for various gain values.

4.3.5 The Addition of λ

Allowing homeowners to change their willingness to participate, λ, reduces the effectiveness of the $\tilde{m}$ criteria. λ is added to the $\tilde{m}$ criteria to dilute the effect the gain has on the response.

$$m(t) = \begin{cases} 0, & \text{if} \quad \theta - k\tilde{m}\lambda \leq 0 \\ 1, & \text{if} \quad \theta \geq 1 \\ \text{unchanged}, & \text{otherwise} \end{cases} \tag{4.9}$$

Completely opting out, $\lambda = 0$, cancels out the $\tilde{m}$ criteria and reverts the thermostat criteria to Equation 4.5. Any level of willingness less than $\lambda = 1$ lowers the θ at which early entry to the OFF state occurs, meaning the house spends more time concurrently cooling with its neighbors, resulting in higher peak demands than full participation.

CHAPTER 5

SIMULATION RESULTS

The results of testing the $\tilde{m}$ criteria are broken into three categories: homogeneous DR events, heterogeneous DR events, and continual operation of the criteria. These $\tilde{m}$ criteria results are compared to the baseline results from Chapter 3.

First, simulations which support the claim that the bounded $\tilde{m}$ criteria produces better results than the unbounded $\tilde{m}$ criteria are presented. Here the criteria is implemented directly after conclusion of the DR event shut off period and compared to the baseline response. This section is followed by examination of the effectiveness of implementing the $\tilde{m}$ criteria in place of homogeneous population DR events. An abrupt switch from typical thermostat criteria to the $\tilde{m}$ criteria produces a reduction in demand, which allows for the possibility of using this criteria as an alternative to a typical DR event. Then various levels of willingness to participate are simulated to get a better understanding of the parameter's importance to the criteria response.

Next, the heterogeneous population undergoes the the same simulations: implementation of criteria after DR event, implementation as an alternative to such events, and examination of willingness to participate.

Finally, the results of implementing the $\tilde{m}$ criteria continually over the course of a week are simulated for both the homogeneous and heterogeneous populations. This provides insight into the ability of the $\tilde{m}$ to improve typical demand behavior for both populations in the absence of DR events.

5.1 Homogeneous Population DR Results

5.1.1 Implementation of $\tilde{m}$

Focusing on the 100% particpation DR event for the homogeneous population, the emergent behavior that results from the unbounded $\tilde{m}$ criteria greatly smooths and reduces the demand following the DR event (Figure 5.1).
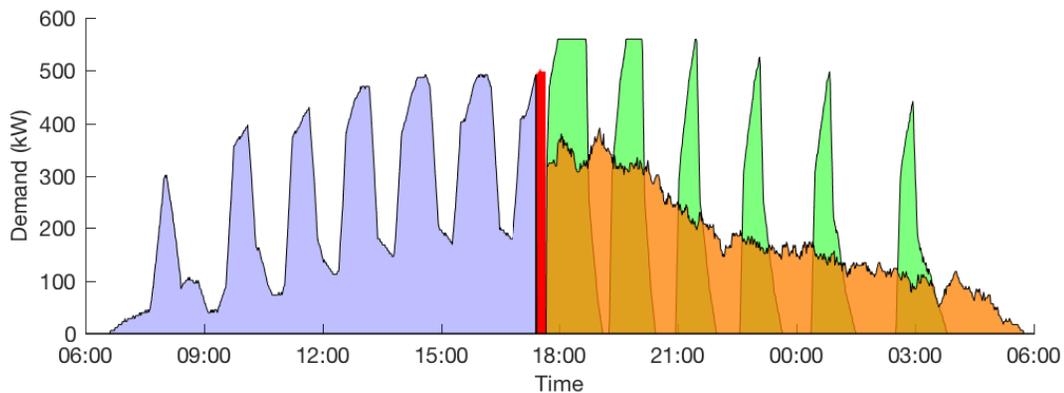
Figure 5.1: 100% DR event profile comparison between normal thermostat operation (Green) and implementation of unbounded $\tilde{m}$ criteria from Equation 4.7 (Orange).

This is possible due to individual houses spending a considerable amount of time above their deadband in the hours immediately following the DR event, as demonstrated by the indices in Table 5.1. Applying the bounded $\tilde{m}$ criteria (Equation 4.8) limits the time spent above the deadband. Looking at the position of the population's θ values within the deadband over the course of the 24 hour time period offers the best contrast of this difference. The unbounded $\tilde{m}$ criteria (Figure 5.2) maximum θ values spike above the upper deadband limit after the DR event begins, before slowly returning to the deadband. Compare this to the bounded criteria (Figure 5.3) θ values that briefly spike during the forced off-period, but quickly return to the deadband.
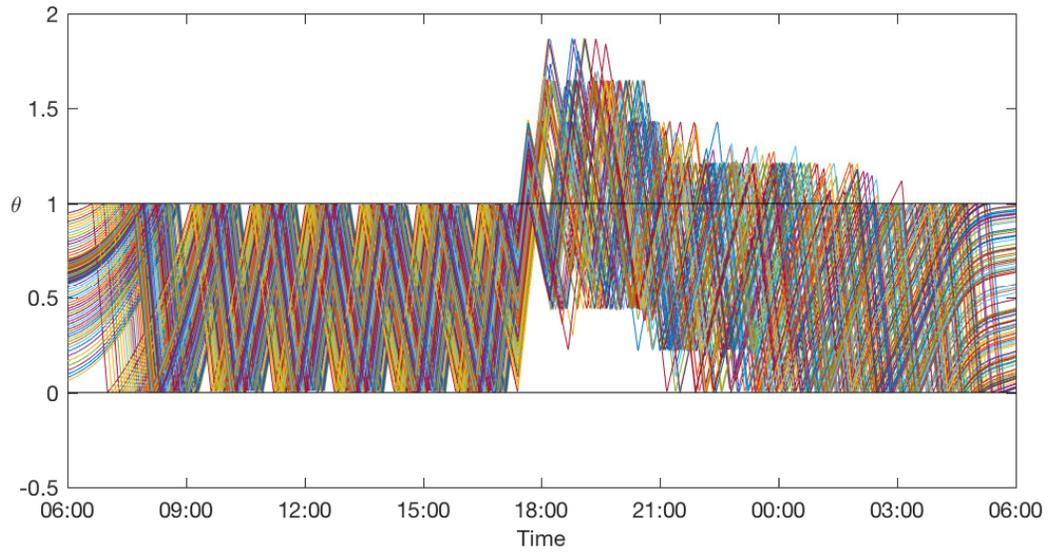
Figure 5.2: θ response of all houses where unbounded $\tilde{m}$ criteria is applied after 100% DR event occurs at 17:31.
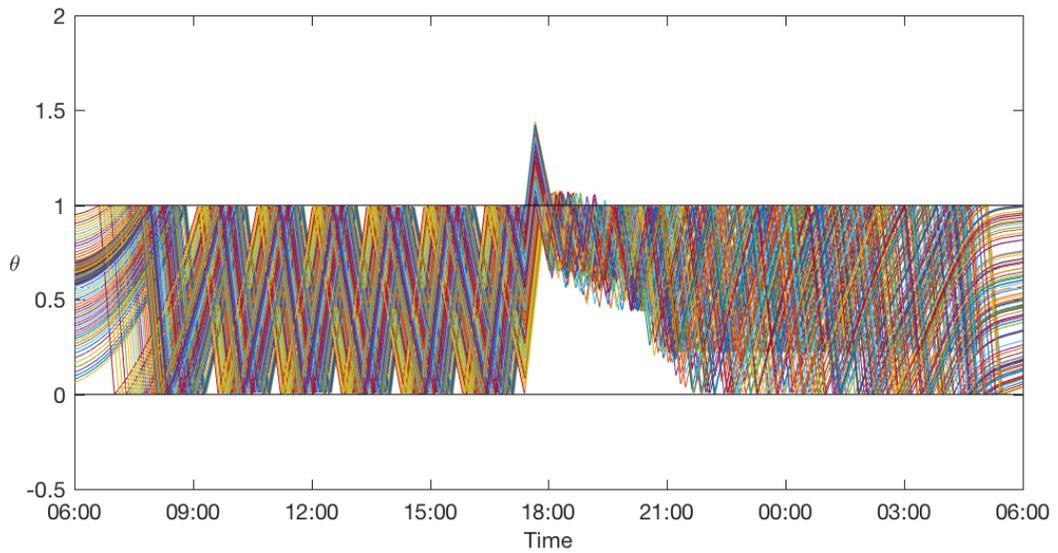
Figure 5.3: θ response of all houses where bounded $\tilde{m}$ criteria is applied after 100% DR event occurs at 17:31.

While both sets of $\tilde{m}$ criteria have the same peak demand over the course of the the day, it is worth noting that after the DR event, the unbounded $\tilde{m}$ criteria reaches a maximum demand of 392 kW, whereas the bounded criteria reaches 492.8

kW (Figure 5.4). However, the unbounded criteria are less desirable, since it causes houses to spend, on average, approximately 2.5 hours above the deadband, while the bounded average is just over 12 minutes spent outside the deadband. Based on this data, the bounded $\tilde{m}$ criteria will be used for the remainder of this research.
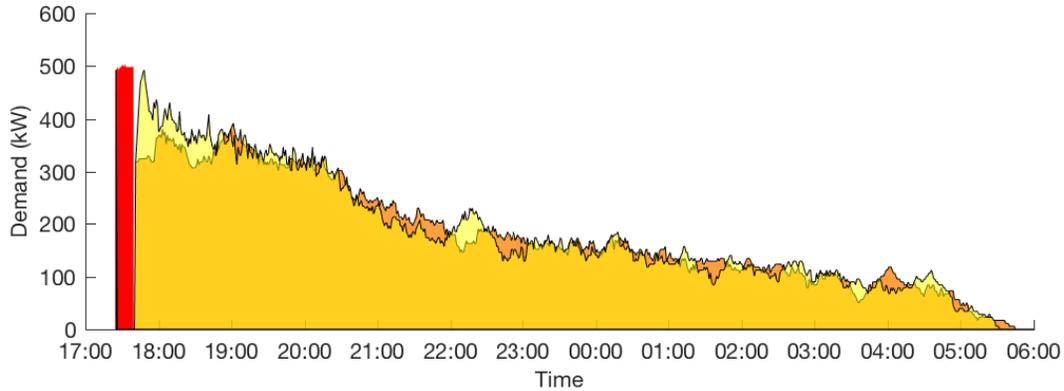
Figure 5.4: 100% DR event profile comparison between unbounded $\tilde{m}$ criteria from Equation 4.7 (Orange) and bounded $\tilde{m}$ criteria from Equation 4.8 (Yellow).

Table 5.1: Indices of Performance - $\tilde{m}$ Criteria for 100% DR event (6am-6am)

Criteria	Peak Demand Post DR (kW)	Total Energy (kWh)	Cycles/day		Time above δ (min)	
			Avg	Max	Avg	Max
None	560.0	4803.7	12.89	13	8.1	41.6
Unbounded	392.0	4852.8	13.17	14	155.6	206.3
Bounded	492.8	4856.2	17.45	19	12.4	45.4

5.1.2 Criteria Created DR Event

The $\tilde{m}$ criteria has shown an ability to respond to a DR event where houses are forced off for a set period of time. Now, instead of reacting to a DR event, let the DR event be created by engaging the $\tilde{m}$ criteria. The initial start time of the DR event remains the same, but rather than turning off all AC units, the thermostats switch from typical thermostat criteria to $\tilde{m}$ criteria.

This switch in criteria results in many thermostats entering the off state earlier, which sheds approximately 80% of the load shed during a typical DR event (Figure 5.5). However, the load of the criteria created event starts increasing immediately and by the end of the 15 minute event reaches 50% of load typically shed. It takes the criteria created event 19 minutes to return to 75% of the pre-DR event load, while implementation of the criteria post-DR requires less than three minutes to reach the same level.

Figure 5.5: Comparison between $\tilde{m}$ criteria implemented after DR event (Yellow) and $\tilde{m}$ criteria implemented in place of DR event (Purple).

As seen in Table 5.2, compared to the baseline DR event with typical thermostat operation, the criteria created event reduces peak demand post DR event by 27% and only increases total energy by 1.07%. There is an uptick in the number of cycles, 32% increase from baseline, as thermostats are briefly restricted to the upper region of the deadband (Figure 5.6).

Previously, turning off 100% of homes for 15 minutes caused many to heat up past the upper limit of the deadband. Utilizing the $\tilde{m}$ criteria means the 7 minute compressor lockout is the only time a house can coast above the deadband. Since this time is shorter, the deviation from the deadband is less severe. With an average of
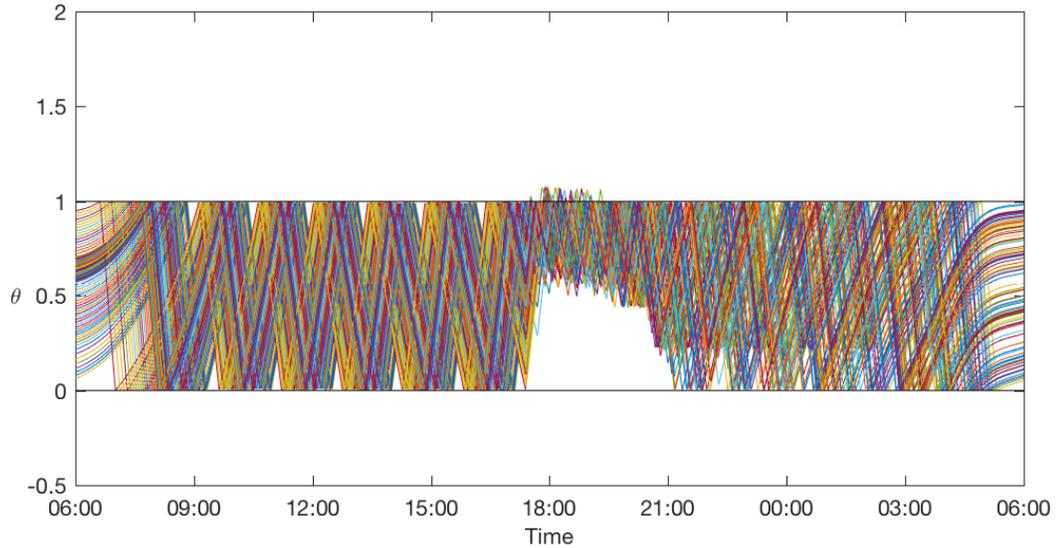
Figure 5.6: θ response of all houses where $\tilde{m}$ criteria is implemented in place of DR event.

1.9 minutes spent outside of the deadband over a period of one day, a 76% reduction from baseline, it can be argued that this slight deviation from normal thermostat behavior would go unnoticed by occupants.

Table 5.2: Indices of Performance - $\tilde{m}$ created DR event compared to $\tilde{m}$ 100% OFF DR event (6am-6am)

Simulation	Peak Demand Post DR (kW)	% from Baseline	Total Energy (kWh)	% from Baseline
Baseline DR	560.0	-	4803.7	-
$\tilde{m}$ - Post DR	492.8	-12.0%	4856.2	1.09%
$\tilde{m}$ - Replace DR	408.8	-27.0%	4855.1	1.07%
	Cycles/day		Time above δ (min)	
Simulation	Avg (Max)	% from Baseline	Avg (Max)	% from Baseline
Baseline DR	12.89 (13)	-	8.09 (41.6)	-
$\tilde{m}$ - Post DR	17.45 (19)	35.38% (46.15%)	12.37 (45.4)	52.92% (9.07%)
$\tilde{m}$ - Replace DR	17.07 (19)	32.43% (46.15%)	1.9 (18.5)	-76.48% (-55.67%)

5.1.3   DR Program Alternative

The utilization of criteria to instigate a DR event can be extended to the DR program percentages (10%, 20%, and 30%), where only the percentage of houses participating in the program switch to the $\tilde{m}$ criteria at the DR start time. All the homes still share state information, but the rest of the homes use typical thermostat criteria for the duration of the simulation. Table 5.3 and Figure 5.7 show the criteria initiated DR event has similar results to the criteria implemented post DR event.

Overall, the energy consumed by each method is approximately equal and the peak demand is less than or equal to the baseline DR program values. This shows that there is potential to use implementation of the criteria as an alternative to typical DR programs. Looking at the criteria created DR event alternative, the peak demand reduction increases with the percentage of homes utilizing the criteria. As expected, the main difference between the DR events is the time spent above the deadband. While the averages are all fairly low, with a maximum average of about 4 minutes spent outside the deadband, the maximum time a house spends outside the deadband is considerably lower for the criteria created DR event. The inherent cost of utilizing the $\tilde{m}$ criteria is visible in the 2.40% to 8.97% increase in the average number of times the AC units must turn on.

It should be noted that this is not a direct comparison to DR programs, because while only a certain percentage of the homes switch to the $\tilde{m}$ criteria, that criteria is based on state sharing of non-participating houses. In a typical DR program, non-participating houses do not allow the utility access to the AC unit.

However, it does demonstrate that incremental amounts of demand response can be produced by engaging only a portion of the population. It also acts as an introduction into willingness to participate. For the percentage of homes that
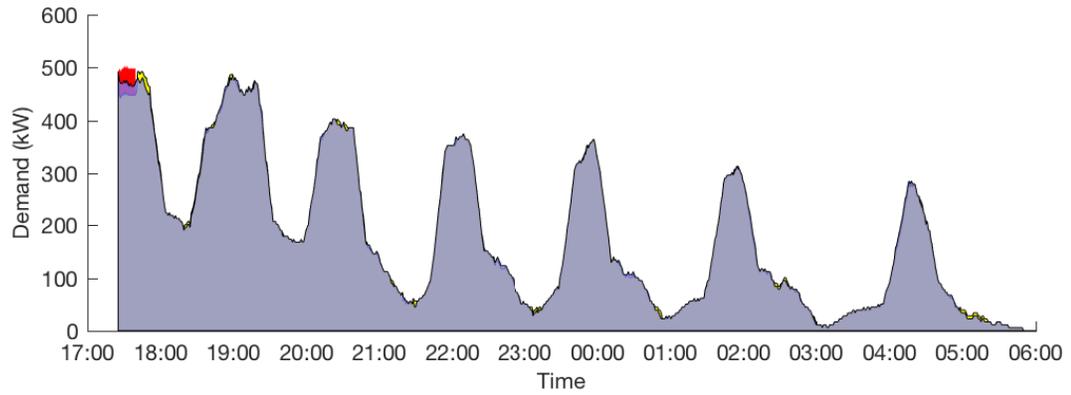
(a) 10% Participation
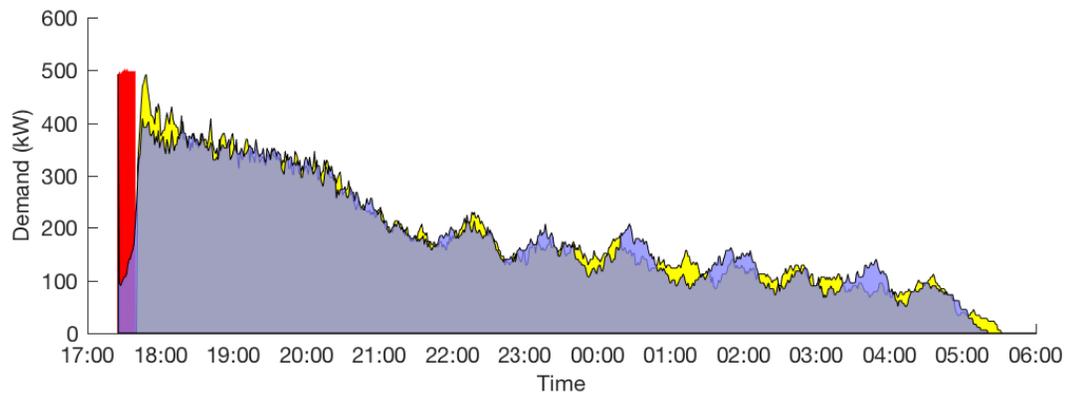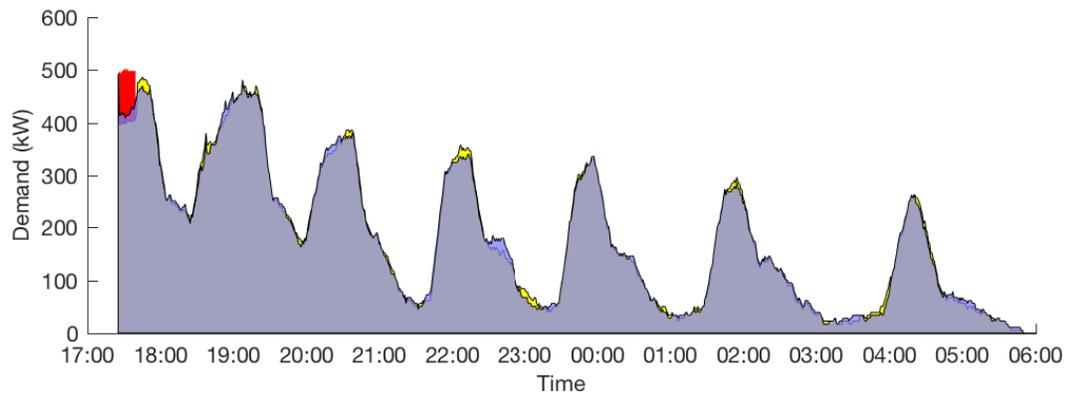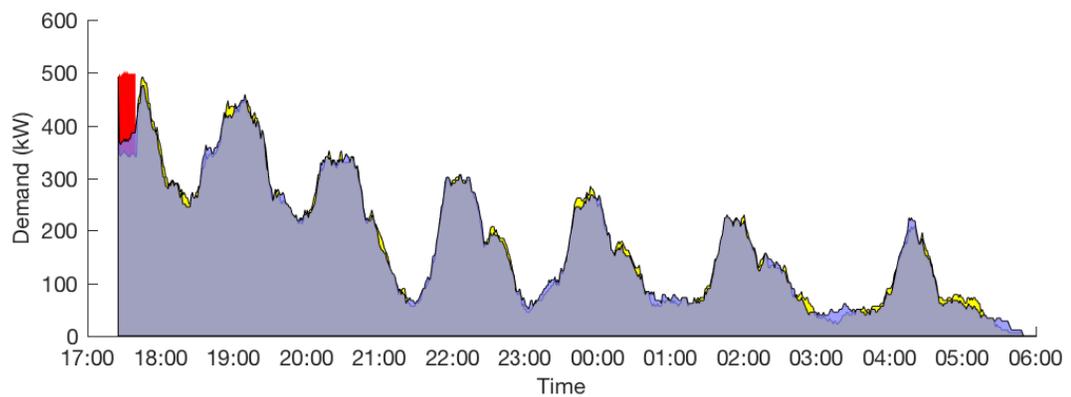

(b) 20% Participation

(c) 30% Participation

Figure 5.7: Comparison between $\tilde{m}$ criteria implemented after DR event (Yellow) and $\tilde{m}$ criteria implemented in place of DR event (Purple) for various percentages of participating homogeneous population homes.

switched to the $\tilde{m}$ criteria, $\lambda = 1$. All other homes ignored the criteria and continued operating with normal thermostat behavior ($\lambda = 0$).

Table 5.3: Indices of Performance - Homogeneous $\tilde{m}$ DR events (6am-6am)

% CC	Simulation	Peak Demand Post DR (kW)	% from Baseline	Total Energy (kWh)	% from Baseline
10	Baseline DR	492.8	-	4907.1	-
	$\tilde{m}$ - Post DR	492.8	0%	4909.1	0.04%
	$\tilde{m}$ - Replace DR	487.2	-1.14%	4907.9	0.02%
20	Baseline DR	487.2	-	4897.0	-
	$\tilde{m}$ - Post DR	487.2	0%	4910.8	0.28%
	$\tilde{m}$ - Replace DR	481.6	-1.15%	4910.7	0.28%
30	Baseline DR	492.8	-	4882.9	-
	$\tilde{m}$ - Post DR	492.8	0%	4902.6	0.40%
	$\tilde{m}$ - Replace DR	476.0	-3.41%	4900.9	0.37%

% CC	Simulation	Cycles/day		Time above δ (min)	
		Avg (Max)	% from Baseline	Avg (Max)	% from Baseline
10	Baseline DR	12.93 (13)	-	1.85 (40.1)	-
	$\tilde{m}$ - Post DR	13.24 (18)	2.40% (38.46%)	2.65 (51.7)	43.59% (28.96%)
	$\tilde{m}$ - Replace DR	13.27 (17)	2.63% (30.77%)	1.06 (26.8)	-42.61% (-33.29%)
20	Baseline DR	12.93 (13)	-	1.39 (37.0)	-
	$\tilde{m}$ - Post DR	13.65 (18)	5.57% (38.46%)	2.45 (37.1)	76.5% (0.04%)
	$\tilde{m}$ - Replace DR	13.73 (19)	6.19% (46.15%)	1.39 (19.3)	0% (-48.0%)
30	Baseline DR	12.93 (13)	-	2.24 (41.6)	-
	$\tilde{m}$ - Post DR	14.07 (18)	8.82% (38.46%)	3.86 (51.0)	72.51% (22.48%)
	$\tilde{m}$ - Replace DR	14.09 (18)	8.97% (38.46%)	1.52 (19.6)	-32.17% (-53.04%)

5.1.4   The Effect of Willingness to Participate

Now consider a population of houses that have a non-uniform value for willingness to participate. To examine how a decrease in the population's average willingness to participate influences results, three normal distributions of λ values were randomly assigned to houses. The averages of the distributions are 0.25, 0.50, and 0.75. Each distribution has the same standard deviation of 0.1. Simulations run using these distributions and the λ thermostat criteria described in Equation 4.9, show that as the level of willingness to participate falls, so does the performance of the criteria to prevent spiking immediately after the 100% shut off DR event (Figure 5.8). However, over time the responses seem to even out and approach similar behavior regardless of the average willingness to participate. The lower the average λ value, the longer the criteria takes to exert its damping effect on the population demand profile. This demonstrates the inherent benefit of having communication among houses, because even an average willingness to participate of 0.25 was able to prevent the large oscillations seen in the case where no $\tilde{m}$ criteria is active ($\lambda = 0$).
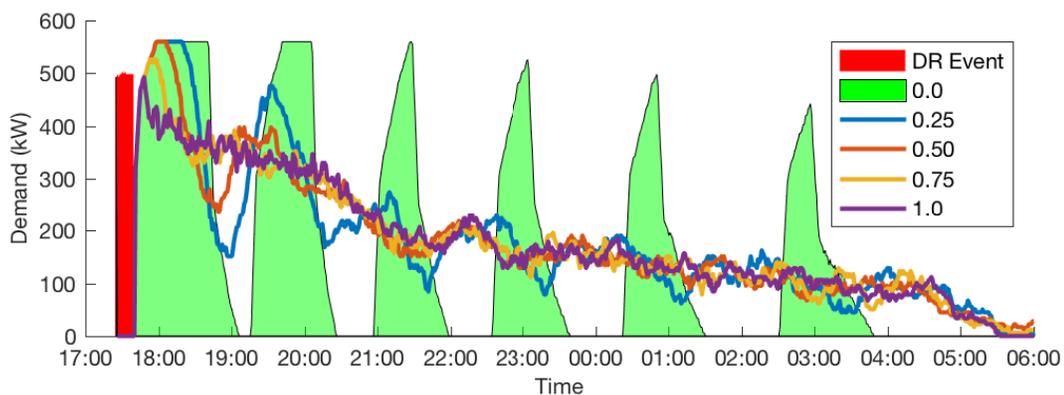
Figure 5.8: Comparison between various average λ values (0.25, 0.5, 0.75), each with a standard deviation of 0.1, the case of no participation ($\lambda = 0$), and full participation ($\lambda = 1$).

5.2 Heterogeneous Population DR Results

5.2.1 Implementation of $\tilde{m}$

Similar to the homogeneous population results, the $\tilde{m}$ criteria prevents the oscillations in demand that results from thermostats following typical deadband criteria (Figure 5.9). The same gain value of $k = 0.87$ is also used for the heterogeneous population $\tilde{m}$ criteria. As seen in Table 5.4, the resulting spike in demand, once units turn back on after the 15 minute DR event, is slightly larger than the load pre-DR event, but still remains considerably less, about 21%, when compared the baseline DR peak value. The $\tilde{m}$ criteria response cycles approximately 36% more times and spends just over 3 more minutes above the deadband limit than the baseline results.
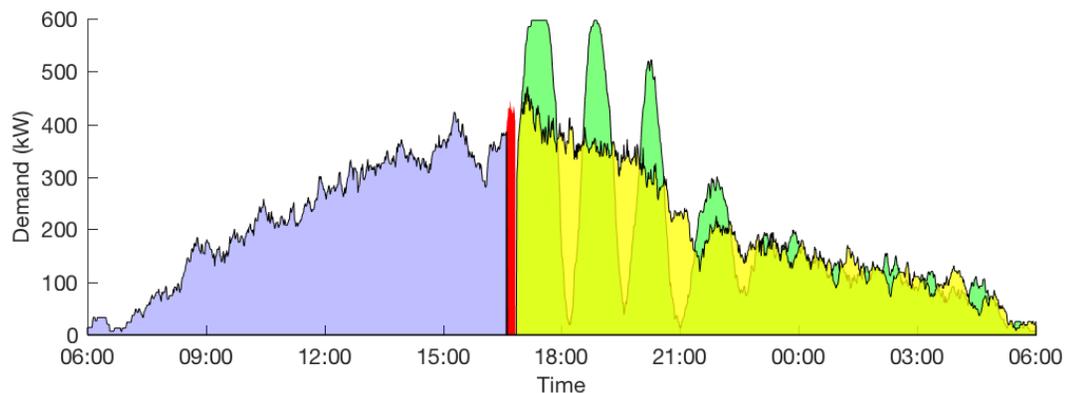
Figure 5.9: 100% DR event profile comparison between normal thermostat operation (Green) and implementation of $\tilde{m}$ criteria (Yellow).

5.2.2 Criteria Created DR Event

Again, switching the heterogeneous population to the $\tilde{m}$ criteria instead of forcing off the homes results in a drop in demand (Figure 5.10). This drop isn't as pronounced as the homogeneous population's response, only shedding approximately 45% of the

load shed during the typical DR event, which by the end of the 15 minute DR event has climbed to 35% of load typically shed. It takes the criteria created event 24 minutes to return to 75% of the pre-DR event load, while implementation of the criteria post-DR requires less than one minute to reach the same level.
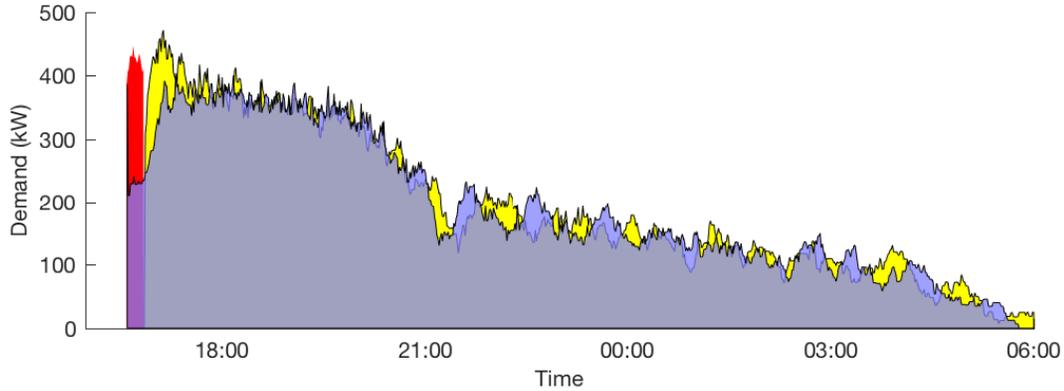
Figure 5.10: Comparison between $\tilde{m}$ criteria implemented after DR event (Yellow) and $\tilde{m}$ criteria implemented in place of DR event (Purple).

Table 5.4: Indices of Performance - Heterogeneous $\tilde{m}$ 100% DR event (6am-6am)

Simulation	Peak Demand Post DR (kW)	% from Baseline	Total Energy (kWh)	% from Baseline
Baseline DR	597.1	-	4902.9	-
$\tilde{m}$ - Post DR	472.0	-20.95%	4883.2	-0.40%
$\tilde{m}$ - Replace DR	404.2	-32.31%	4872.4	-0.62%

Simulation	Cycles/day		Time above δ (min)	
	Avg (Max)	% from Baseline	Avg (Max)	% from Baseline
Baseline DR	13.50 (18)	-	9.78 (42.0)	-
$\tilde{m}$ - Post DR	18.38 (23)	36.15% (27.78%)	13.1 (42.5)	34.28% (1.35%)
$\tilde{m}$ - Replace DR	18.05 (23)	33.70% (27.78%)	0.72 (12.0)	-92.59% (-71.50%)

Comparing relative to the baseline DR event with no $\tilde{m}$ criteria implementation, the $\tilde{m}$ criteria implemented instead of the DR event outperforms the post DR im-

plementation in all indices of performance, as seen in Table 5.4. This is especially noticeable in the average time spent above the deadband (-92.59% vs +34.28%). The criteria created event is also able to reduce the peak demand after the start of the DR event to 404.2 kW from the baseline's 597.1 kW, which is approximately a 32% reduction. The reduction in demand and energy does come at a cost to the number of cycles the AC unit experiences. To produce this desired response, the AC unit cycles almost 34% more times than normal daily operation.

5.2.3 DR Program Alternative

Now the $\tilde{m}$ criteria is implemented instead of the forced off DR program for the heterogeneous parameters (Figure 5.11). The same three levels of participation as the homogeneous results (10%, 20%, 30%) are simulated and the results compiled in Table 5.5. The same houses that participated in the DR shut off baseline results from Chapter 3 are the houses that switch to the $\tilde{m}$ criteria at the beginning of the DR event.

Overall, the two methods of implementing the $\tilde{m}$ criteria have similar results, with the exception being the time spent above the deadband, where implementing the $\tilde{m}$ criteria in place of the DR program shut off outperforms the post DR implementation. Otherwise, both $\tilde{m}$ criteria methods similarly reduce total energy and increase the number of cycles per day experienced by the AC units. At the lower levels of participation, implementing the criteria does not improve the peak demand following the DR event. It isn't until 30% of the houses turn off that a reduction in demand is observed. This could mean that a minimum level of participation is required to achieve a response with lower peak demand post DR event.
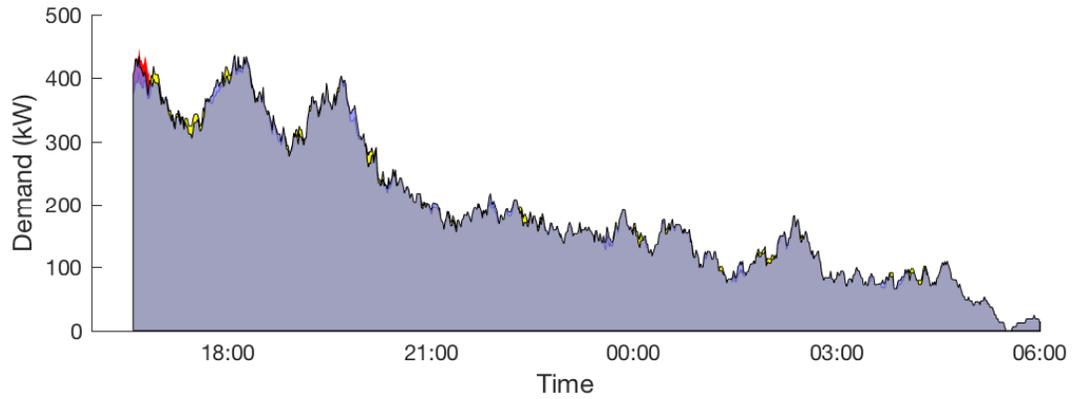
(a) 10% Participation
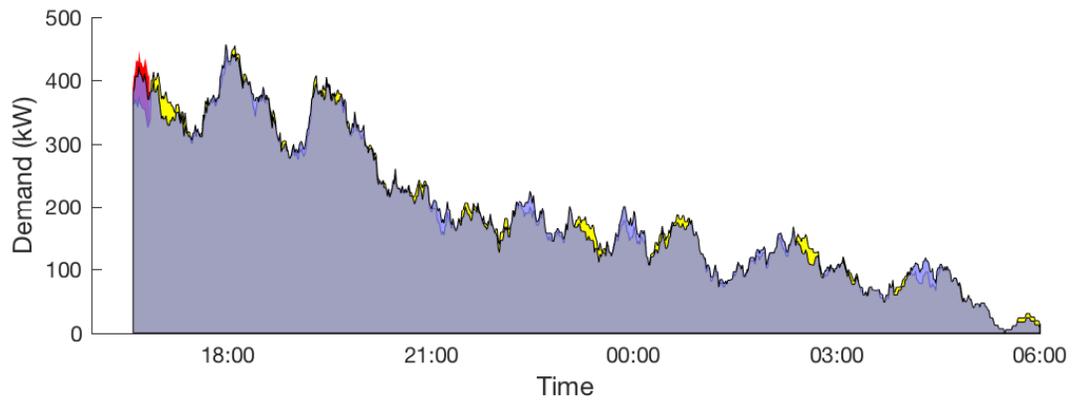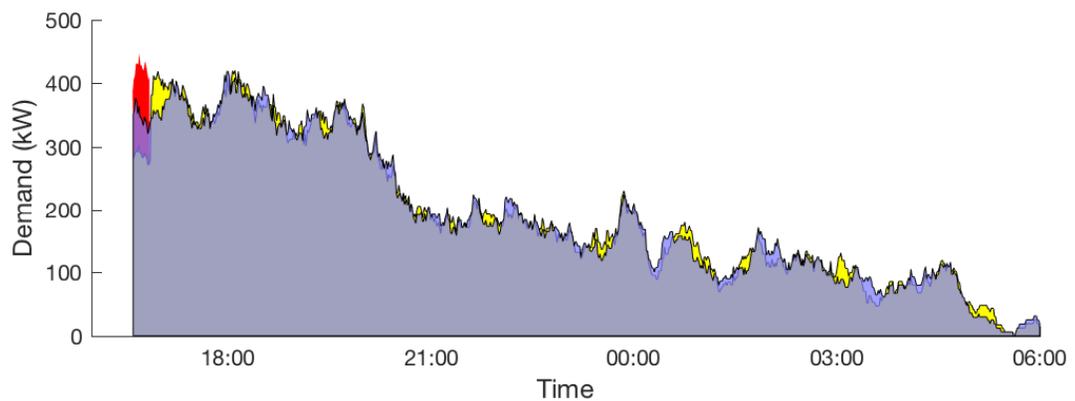
(b) 20% Participation

(c) 30% Participation

Figure 5.11: Comparison between $\tilde{m}$ criteria implemented after DR event (Yellow) and $\tilde{m}$ criteria implemented in place of DR event (Purple) for various percentages of participating heterogeneous population homes.

Table 5.5: Indices of Performance - Heterogeneous $\tilde{m}$ DR events (6am-6am)

% CC	Simulation	Peak Demand Post DR (kW)	% from Baseline	Total Energy (kWh)	% from Baseline
10	Baseline DR	419.0	-	4905.2	-
	$\tilde{m}$ - Post DR	428.1	2.19%	4899.9	-0.11%
	$\tilde{m}$ - Replace DR	436.6	4.20%	4899.8	-0.11%
20	Baseline DR	429.0	-	4904.5	-
	$\tilde{m}$ - Post DR	455.1	6.09%	4897.5	-0.14%
	$\tilde{m}$ - Replace DR	457.0	6.54%	4898.8	-0.12%
30	Baseline DR	440.6	-	4906.0	-
	$\tilde{m}$ - Post DR	420.5	-4.56%	4898.0	-0.16%
	$\tilde{m}$ - Replace DR	419.1	-4.88%	4897.0	-0.18%

% CC	Simulation	Cycles/day Avg (Max)	% from Baseline	Time above δ (min) Avg (Max)	% from Baseline
10	Baseline DR	13.31 (18)	-	1.41 (32.2)	-
	$\tilde{m}$ - Post DR	13.69 (20)	2.86% (11.11%)	2.12 (43.2)	50.63% (33.95%)
	$\tilde{m}$ - Replace DR	13.72 (20)	3.08% (11.11%)	0.95 (21.7)	-32.59% (-32.68%)
20	Baseline DR	13.28 (18)	-	2.49 (36.5)	-
	$\tilde{m}$ - Post DR	14.07 (20)	5.95% (11.11%)	2.98 (39.1)	19.43% (7.11%)
	$\tilde{m}$ - Replace DR	14.11 (20)	6.25% (11.11%)	0.48 (24.5)	-80.95% (-32.86%)
30	Baseline DR	13.36 (18)	-	3.14 (33.4)	-
	$\tilde{m}$ - Post DR	14.65 (22)	9.66% (22.22%)	4.53 (38.8)	44.31% (16.09%)
	$\tilde{m}$ - Replace DR	14.62 (21)	9.43% (16.67%)	1.17 (21.5)	-62.71% (-35.81%)

5.2.4 The Effect of the Willingness to Participate

The heterogeneous population was assigned the same distributions of λ values as the homogeneous population; three distributions with average values of 0.25, 0.50, and 0.75 and a standard deviation of 0.1. The results seen in Figure 5.12 are

nearly identical to the effect the willingness to participate had on the homogeneous population. The higher the average willingness to participate, the better the response. However, this response is only better for the initial post DR event response. By midnight, the heterogeneous population with no participation ($\lambda = 0$) is essentially the same as all other levels of particpation, which is a result of the heterogeneous parameters inherent ability to negate the effects of a DR event.
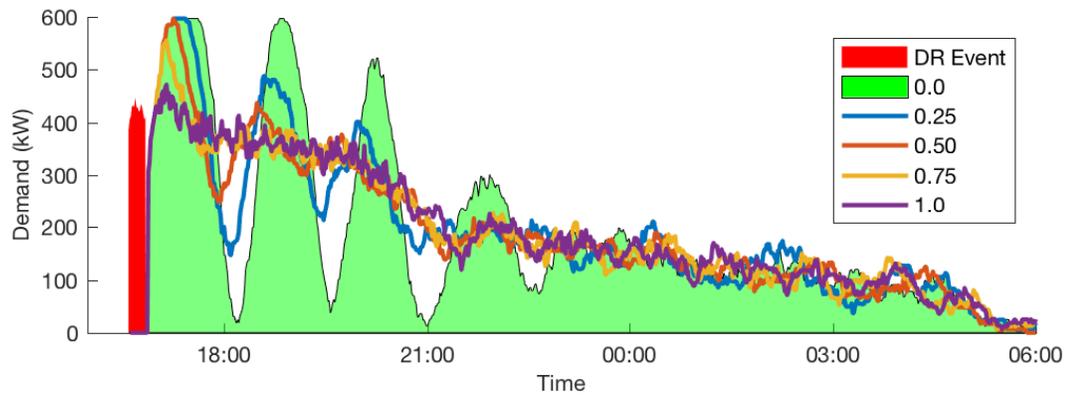
Figure 5.12: Comparison between various average λ values (0.25, 0.5, 0.75), each with a standard deviation of 0.1, the case of no participation ($\lambda = 0$), and full participation ($\lambda = 1$).

5.3 Continual Operation Results

Now let's consider a period of longer operation without DR events, where the $\tilde{m}$ criteria takes the place of typical thermostat criteria and the simulation is run over the course of the baseline week (July 21 - July 28).

5.3.1 Homogeneous Population

Starting with the homogeneous set of houses, the continuous operation of the $\tilde{m}$ criteria drastically improves the demand profile by eliminating the large oscillations
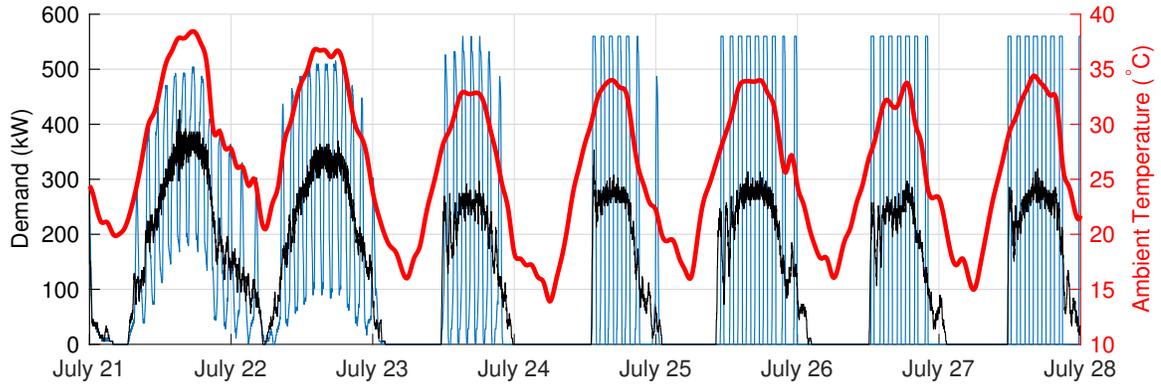
Figure 5.13: Demand profile for population of homogeneous houses over the course of one week.

Table 5.6: Indices of Performance - Homogeneous Baseline vs $\tilde{m}$ criteria (Week)

Simulation	Peak Demand (kW)	% from Baseline	Total Energy (kWh)	% from Baseline
Baseline	560.0	-	21738	-
$\tilde{m}$ Criteria	425.6	-24.0%	21554	-0.85%
	Cycles per Week			
Simulation	Avg	% from Baseline	Max	% from Baseline
Baseline	62.78	-	63	-
$\tilde{m}$ Criteria	96.66	53.97%	100	58.73%

of the baseline profile (Figure 5.13).

As seen in Table 5.6, comparison of the maximum demand over the course of the week shows the $\tilde{m}$ criteria is 24% less than the baseline. The $\tilde{m}$ peak demand value of 425.6 kW occurs on the first day of the week, but later in the week the daily maximum demand stays closer to 300 kW, so this reduction is greater on certain days. The total energy is approximately the same with a reduction of 0.85% for the $\tilde{m}$ criteria, but to obtain these power and energy results requires an average increase of about 54% in the number of cycles. The time outside the deadband was not included in Table 5.6 since the average total time the $\tilde{m}$ criteria spent above the deaband is right around 2.5 minute and can be considered negligible over the course of a week. The maximum

total time was closer to 20 minutes, which is less than one percent of the week.

5.3.2   Heterogeneous Population

The $\tilde{m}$ criteria has less of an effect on the heterogeneous population than it did on the homogeneous population. This is expected with the difference in heterogeneous time constants preventing the sizable oscillations caused by grouping up of the AC units. Looking at the demand profile (Figure 5.14) the criteria does prevent the larger oscillations that occur around late afternoon when the temperature is highest.
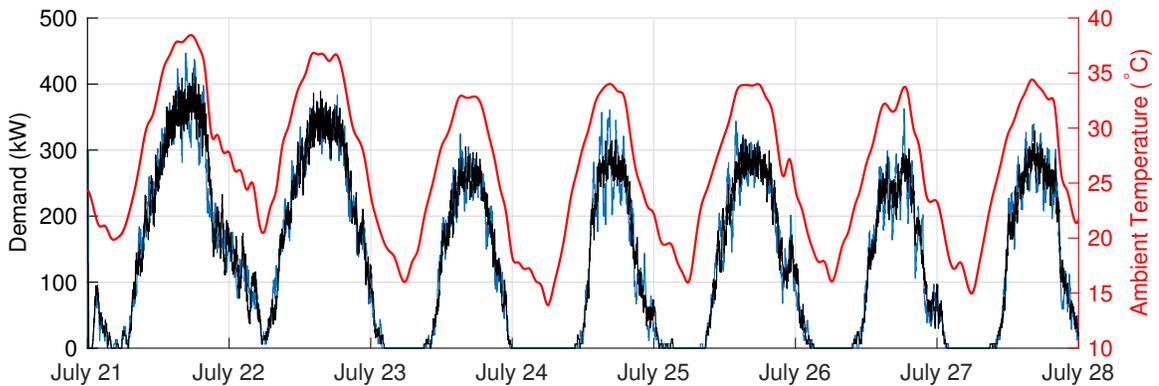
Figure 5.14: Demand profile for population of heterogeneous houses over the course of one week.

Table 5.7: Indices of Performance - Heterogeneous Baseline vs $\tilde{m}$ criteria (Week)

Simulation	Peak Demand (kW)	% from Baseline	Total Energy (kWh)	% from Baseline
Baseline	446.9	-	21974	-
$\tilde{m}$ Criteria	425.6	-6.76%	21827	-0.67%
	Cycles per Week			
Simulation	Avg	% from Baseline	Max	% from Baseline
Baseline	64.89	-	92	-
$\tilde{m}$ Criteria	96.23	48.30%	123	33.70%

When compared to the baseline data found in Table 5.7, the reduction in peak demand is 6.8%, while using 0.67% less energy. The increase in average number of cycles is 48.3%, which is less than the homogeneous increase of 54%. The average time spent outside the deadband is 0.99 minutes, which is even less than the homogeneous result and is considered negligible.

CHAPTER 6

DISCUSSION AND CONCLUSIONS

6.1 Summary of Findings

This research has demonstrated that limited local communication among neighboring thermostats produces emergent behavior, which has the potential to improve post DR event aggregate demand behavior by simply sharing the state of AC units with four connected neighbors. This novel criteria can be used in conjunction with current DR event methods or as an alternative to such programs. In the most extreme case, where all houses are shut off, the criteria is effective at preventing the large oscillations expected for typical thermostat operation.

This criteria is effective for both homogeneous and heterogeneous populations of houses. Looking at the 100% DR shut off event response, the criteria was more effective for the heterogeneous population, where the post DR event peak demand was 21% less than the baseline, while the homogeneous population only achieved a 12% reduction in peak demand. The heterogeneous population also used less energy than baseline (-0.4%) while the homogeneous population used more (+1.09%). Both sets of parameters resulted in an increase in the number of cycles of approximately 36%. One possible explanation for the better heterogeneous response is the spread of parameters inherently working to damp the oscillations that result from the shut off event. The heterogeneous population parameter spread results in different cycle times, which prevent syncing up. The homogeneous population, once synced up,

would stay synced up if not for the criteria. These findings point to both populations benefiting from implementation of the criteria.

It was also discovered that implementing the criteria, instead of shutting off all the houses, produced better results than implementing the criteria in response to the shut off event. For the heterogeneous parameters, the peak demand reduction from baseline was 32%, while it was 27% for the homogeneous. Since the houses aren't being forced off for 15 minutes, there is limited time spent above the deadband. Compared to baseline, the heterogeneous population average time above the deadband over the course of the day was 0.72 minutes or almost a 93% reduction. For the homogeneous population the average was 1.9 minutes or about 76% less. This is important because excessive amounts of time outside the deadband may be noticeable by the consumer and reduce their satisfaction and willingness to be a part of the program. This research shows that using the criteria to create a DR event, rather than respond to one, is a viable option.

As a result of the success of the criteria created DR event, the criteria was tested as a replacement to DR programs. Smaller percentages of the population initiated use of the criteria to create a DR event and it was assumed that even if a house was not switching to the $\tilde{m}$ criteria, their state was available to their neighbors. The results are mixed. The homogeneous population used the approximately same amount of energy and maintained a lower peak demand. On the other hand, the heterogeneous population had a higher peak demand for the lower participation levels (10% and 20%) and used approximately the same amount of energy. Perhaps this is a byproduct of the parameters of the random houses chosen to participate. It could also point to a certain level of participation (30%) is needed to produce a desired demand response.

Continual operation of the $\tilde{m}$ criteria, over the course of a week of summer temperatures, proved effective in reducing the peak demand and total energy relative

to typical thermostat operation. The homogeneous results show this most clearly with a reduction in peak demand of 24% and reduction in total energy of 0.85%. Looking at the demand profile, it can be seen that the criteria is able to prevent syncing up from occurring, instead providing smoother daily demand profiles. The heterogeneous population results are less pronounced with only a 6.8% reduction in demand and a 0.67% reduction in total energy. In both cases, the $\tilde{m}$ criteria has the potential to replace typical thermostat criteria.

The underlying cost that runs through all the simulations is the increase in the average number of cycles the AC units experience. Since the bounded $\tilde{m}$ criteria forces units off sooner than typical thermostat operation, this increase is to be expected. The percent increase varies depending on the simulation, but can get as high as the 54% increase for the continual operation of the homogeneous population. Since AC units are designed to operate under typical thermostat criteria, further examination is needed to determine how this increase affects the lifetime performance of the units. A reduction in lifetime of the units may offset any savings resulting from a reduction in demand.

Another variable, the willingness to participate, was introduced to represent a more realistic situation where consumers have some control over the extent to which the $\tilde{m}$ criteria affects their comfort. As would be expected, the effectiveness of the criteria deteriorates as the average willingness to participate drops. However, even the lower levels of willingness to participate achieved better results than simulations with no communication.

Efforts to mathematically prove stability for the new criteria were unsuccessful. Empirical data from simulations show that as the gain increases, the response deteriorates, but there is no clear stability limit. Further research will need to explore this area.

6.2   Future Work

Beyond what has already been discussed, there are a number of different ways to expand upon the work of this thesis. One of the larger changes could be to the model used to calculate the thermal dynamics of the home. This research assumed a first order model was adequate to describe the thermal dynamics of the houses. It would be worth checking to see if a second order model produces similar results, therefore supporting that claim.

Increasing the number of houses modeled will improve the results by reducing sensitivity to noise. Only 100 houses were used in this research, due to software limitations and the complexity of the model. Utilizing another software that can handle 10,000 houses may offer further insight into the emergent behavior.

The ring lattice with four connections per agent was the only graph used in this research, but different number of connections and other graph layouts, such as a square lattice, may show how the network shape affects performance. It may also be practical to explore a distance based graph, where agents are only connected to agents within a certain distance. This would be the network to use if the ability for thermostat to communicate depends on distance limited technology.

Specifics of the technology necessary to facilitate connections between thermostats requires further examination. This research assumed the technology was in place to allow for the sharing of information.

A network robustness analysis of the population network would provide insight into how lost connections impact the response. In this research it was assumed that all thermostats would be able to communicate with all its connected neighbors at all times, but this won't always be the case in real world implementation. Sometimes a thermostat may unexpectedly stop communicating with its neighbors due to cyber

security intrusions, equipment failure, weather, or other circumstances. When this happens, how will the system handle it? This analysis could point to a rate of degradation in performance as the graph deteriorates.

Along similar lines, the rate that thermostats send messages should be explored. This research had all thermostats send a message every minute and assumed 100% of messages were received to ensure behavior followed the criteria closely. But what happens if the message is sent at a slower or faster rate? Perhaps there is a rate that must be maintained to retain criteria effectiveness. Furthermore, perhaps the rate can be changed throughout the day depending on the need to for neighbor information, where rates would increase in the evening when temperature peaks. It would also be useful to explore the effects of only a certain percentage of messages reaching their intended target. Assigning probabilities that a message will be sent successfully will create a more realistic system.

Further research into the parameters of houses would help support this model. This thesis built on the parameters established by Basash and Fathy [22], which are rather constricted and are not representative of many houses. Empirical data of different houses would help solidify the model. Broadening the scope of the population could allow for the application of this model to other demand response options such as apartment or office buildings.

The proposed criteria could be integrated with machine learning, where thermostats learn over time how to optimize their behavior. There's also the potential to include renewable tracking into the criteria, which would allow the population to ramp up or down with changes in wind and solar production.

This research assumed constant set point temperature, which is not typically the case for smart programmable thermostats. Therefore, the impact of temperature setbacks on overall performance should be explored. Similarly, allowing individual

willingness to participate to be changed over the course of a day would be more realistic and should be explored.

As variable speed AC units become more popular this model will need to be updated. With variable speed units, instead of the common ON/OFF control of traditional units, a controller modulates the compressor output allowing for improved efficiency and comfort. This allows for the option to model m as a continuous variable on the range from 0 to 1.

REFERENCES

[1] P. Denholm, M. O'Connell, G. Brinkman, and J. Jorgenson, "Overgeneration from solar energy in california. a field guide to the duck chart," tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States), 2015.

[2] "Renewables 2017." https://www.iea.org/publications/renewables2017/. Accessed: 12-Mar-2018.

[3] M. Scott, "Half of all power set to come from renewable energy by 2050, while coal recedes to just 11%," *Forbes*, Jun 2018.

[4] O. Ma and K. Cheung, "Demand response and energy storage integration study," tech. rep., Office of Energy Efficiency and Renewable Energy, Office of Electricity Delivery and Energy Reliability, 2016.

[5] "Hawaiis Solar-Grid Landscape and the Nessie Curve — Greentech Media." https://www.greentechmedia.com/articles/read/hawaiis-solar-grid-landscape-and-the-nessie-curve#gs.X6TMXBk. Accessed: 13-Mar-2018.

[6] C. Vlahoplus, G. Litra, P. Quinlan, and C. Becker, "Revisiting the california duck curve: An exploration of its existence, impact, and migration potential," tech. rep., Tech. rep., Scottmadden, Inc.(Aug. 2016), 2016.

[7] L. Vandepaer, J. Cloutier, and B. Amor, "Environmental impacts of lithium metal polymer and lithium-ion stationary batteries," *Renewable and Sustainable Energy Reviews*, vol. 78, pp. 46 – 60, 2017.

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

[9] J. Rodríguez-Molina, M. Martínez-Núñez, J.-F. Martínez, and W. Pérez-Aguiar, "Business models in the smart grid: Challenges, opportunities and proposals for prosumer profitability," *Energies*, vol. 7, no. 9, pp. 6142–6171, 2014.

[10] K. Zhou, S. Yang, and Z. Shao, "Energy internet: The business perspective," *Applied Energy*, vol. 178, pp. 212 – 222, 2016.

[11] R. Malhame and C.-Y. Chong, "Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system," *IEEE Transactions on Automatic Control*, vol. 30, pp. 854–860, Sept. 1985.

[12] Z. Xu, J. Ostergaard, M. Togeby, and C. Marcus-Moller, "Design and Modelling of Thermostatically Controlled Loads as Frequency Controlled Reserve," in *2007 IEEE Power Engineering Society General Meeting*, pp. 1–6, June 2007.

[13] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi, "Aggregated Modeling and Control of Air Conditioning Loads for Demand Response," *IEEE transactions on power systems*, vol. 28, no. 4, pp. 4655–4664, 2013.

[14] J. Campillo, E. Dahlquist, F. Wallin, and I. Vassileva, "Is real-time electricity pricing suitable for residential users without demand-side management?," *Energy*, vol. 109, pp. 310–325, Aug. 2016.

[15] A. Nilsson, P. Stoll, and N. Brandt, "Assessing the impact of real-time price visualization on residential electricity consumption, costs, and carbon emissions," *Resources, Conservation and Recycling*, vol. 124, pp. 152–161, Sept. 2017.

[16] Y. Zhang, W. Gao, Y. Ushifusa, W. Chen, and S. Kuroki, "An Exploratory Analysis of Kitakyushu Residential Customer Response to Dynamic Electricity Pricing," *Procedia - Social and Behavioral Sciences*, vol. 216, pp. 409–416, Jan. 2016.

[17] H. Allcott, "Rethinking real-time electricity pricing," *Resource and Energy Economics*, vol. 33, pp. 820–842, Nov. 2011.

[18] S. Lutzenhiser, "Beyond the price effect in time-of-use programs: results from a municipal utility pilot, 2007-2008," 2009.

[19] B. Barker, "Storing in bulk - the need for large-scale energy storage grows, but the business case remains uncertain," *EPRI Journal*, Jan 2018.

[20] M. Budt, D. Wolf, R. Span, and J. Yan, "A review on compressed air energy storage: Basic principles, past milestones and recent developments," *Applied Energy*, vol. 170, pp. 250–268, May 2016.

[21] I. Penn, "Edison and tesla unveil giant energy storage system," *Los Angeles Times*, Jan 2017.

[22] S. Bashash and H. K. Fathy, "Modeling and Control of Aggregate Air Conditioning Loads for Robust Renewable Power Management," *IEEE Transactions on Control Systems Technology*, vol. 21, pp. 1318–1327, July 2013.

[23] M. A. Zehir and M. Bagriyanik, "Demand side management by controlling refrigerators and its effects on consumers," *Energy Conversion and Management*, vol. 64, pp. 238–244, 2012.

[24] S. Railsback and V. Grimm, *Agent-based and Individual-based Modeling: A Practical Introduction.* Agent-based and Individual-based Modeling: A Practical Introduction, Princeton University Press, 2012.

[25] M. A. Bedau, "Weak emergence," *Noûs*, vol. 31, pp. 375–399, 1997.

[26] A. T. Crooks and A. B. Hailegiorgis, "An agent-based modeling approach applied to the spread of cholera," *Environmental Modelling & Software*, vol. 62, pp. 164–177, Dec. 2014.

[27] C. Gunaratne, M. I. Akbas, I. Garibay, and O. Ozmen, "Evaluation of Zika Vector Control Strategies Using Agent-Based Modeling," *arXiv preprint arXiv:1604.06121*, 2016.

[28] A. Mashhadi Ali, M. E. Shafiee, and E. Z. Berglund, "Agent-based modeling to simulate the dynamics of urban water supply: Climate, population growth, and water shortages," *Sustainable Cities and Society*, vol. 28, pp. 420–434, Jan. 2017.

[29] A. L. Andrei, K. Comer, and M. Koehler, "An agent-based model of network effects on tax compliance and evasion," *Journal of Economic Psychology*, vol. 40, pp. 119–133, Feb. 2014.

[30] H. Alatas, S. Nurhimawan, F. Asmat, and H. Hardhienata, "Dynamics of an agent-based opinion model with complete social connectivity network," *Chaos, Solitons & Fractals*, vol. 101, pp. 24–32, Aug. 2017.

[31] Y. K. Hessary and M. Hadzikadic, "Role of behavioral heterogeneity in aggregate financial market behavior: An agent-based approach," *Procedia Computer Science*, vol. 108, pp. 978 – 987, 2017. International Conference on Computational Science, ICCS 2017, 12-14 June 2017, Zurich, Switzerland.

[32] J. Ge, "Endogenous rise and collapse of housing price," *Computers, Environment and Urban Systems*, vol. 62, pp. 182–198, Mar. 2017.

[33] V. L. Erickson, Y. Lin, A. Kamthe, R. Brahme, A. Surana, A. E. Cerpa, M. D. Sohn, and S. Narayanan, "Energy efficient building environment control strategies using real-time occupancy measurements," in *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pp. 19–24, ACM, 2009.

[34] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Aggregate Flexibility of Thermostatically Controlled Loads," *IEEE Transactions on Power Systems*, vol. 30, pp. 189–198, Jan. 2015.

[35] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings, "Agent-based control for decentralised demand side management in the smart grid," in *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pp. 5–12, International Foundation for Autonomous Agents and Multiagent Systems, 2011.

[36] E. Xydas, C. Marmaras, and L. M. Cipcigan, "A multi-agent based scheduling algorithm for adaptive electric vehicles charging," *Applied Energy*, vol. 177, pp. 354–365, Sept. 2016.

[37] M. Ruth, A. Pratt, M. Lunacek, S. Mittal, H. Wu, and W. Jones, "Effects of home energy management systems on distribution utilities and feeders under various market structures," in *International Conference on Electricity Distribution*, 2015.

[38] I. Ghansah, "Smart grid cyber security potential threats, vulnerabilities and risks," *California Energy Commission, PIER Energy-Related Environmental Research Program, CEC-500-2012-047*, 2009.

[39] "A/C Cool Credit." `https://www.idahopower.com/ways-to-save/savings-for-your-home/rebates-and-offers/ac-cool-credit/`. Accessed: 05-Mar-2018.

[40] S. Bashash and H. K. Fathy, "Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads," in *American Control Conference (ACC), 2011*, pp. 4546–4553, IEEE, 2011.

[41] L. B. Wood and H. H. Asada, "Cellular Stochastic Control of the Collective Output of a Class of Distributed Hysteretic Systems," *Journal of Dynamic Systems, Measurement, and Control*, vol. 133, p. 061011, Nov. 2011.

[42] S. Abar, G. K. Theodoropoulos, P. Lemarinier, and G. M. OHare, "Agent based modelling and simulation tools: A review of the state-of-art software," *Computer Science Review*, vol. 24, pp. 13 – 33, 2017.

[43] "AnyLogic Homepage." `https://www.anylogic.com/`.

[44] "National Solar Radiation Data Base — 1991-2005 Update: Typical Meteorological Year 3." `http://rredc.nrel.gov/solar/old_data/nsrdb/1991-2005/tmy3/by_state_and_city_old.html#I`.

[45] A. Barabási and M. Pósfai, *Network Science*. Cambridge University Press, 2016.

[46] M. Fiedler, "Algebraic connectivity of graphs," *Czechoslovak mathematical journal*, vol. 23, no. 2, pp. 298–305, 1973.

APPENDIX A

Homogeneous Parameters

House	StartState	StartTemp (°C)	CoolCredit10	CoolCredit20	CoolCredit30	House	StartState	StartTemp (°C)	CoolCredit10	CoolCredit20	CoolCredit30
1	1	20.015	1	0	0	51	1	19.801	1	0	0
2	0	20.138	0	0	0	52	1	19.954	0	0	1
3	1	20.230	0	0	0	53	1	20.107	0	1	0
4	1	20.138	0	1	1	54	0	19.913	1	0	0
5	1	20.056	0	0	0	55	0	19.995	0	0	0
6	1	19.995	0	0	1	56	0	20.066	0	0	1
7	0	20.148	0	0	0	57	0	20.036	0	0	1
8	1	19.760	0	0	0	58	0	19.954	0	0	0
9	1	20.168	0	0	1	59	0	19.964	0	1	0
10	1	19.923	0	1	0	60	0	20.230	0	0	0
11	0	19.760	0	0	0	61	1	19.934	0	0	0
12	0	19.974	0	0	0	62	0	19.781	0	0	0
13	0	19.811	0	1	0	63	0	19.893	0	0	0
14	1	19.985	0	0	0	64	1	19.832	1	0	1
15	0	19.791	0	0	0	65	0	19.872	0	1	0
16	0	19.801	0	0	0	66	1	19.893	0	0	0
17	1	20.046	0	0	1	67	1	19.913	0	0	1
18	0	19.750	0	0	0	68	1	20.026	0	0	1
19	1	20.077	0	1	1	69	0	20.179	0	0	1
20	1	20.189	0	0	0	70	0	19.832	0	1	0
21	1	19.811	0	1	1	71	0	20.250	0	0	0
22	1	20.240	1	0	0	72	0	20.005	0	0	0
23	1	20.148	0	0	0	73	0	20.128	0	1	1
24	0	20.168	0	0	0	74	1	20.158	0	1	0
25	1	19.974	0	0	0	75	1	20.209	0	1	1
26	0	19.985	0	0	1	76	1	19.964	0	0	0
27	1	19.852	0	0	1	77	0	20.077	0	0	0
28	0	19.842	0	0	0	78	0	20.026	1	0	0
29	1	20.005	0	0	0	79	1	20.250	0	0	0
30	0	20.199	0	0	0	80	1	20.199	0	0	0
31	0	19.923	0	0	0	81	0	20.015	0	1	0
32	1	20.066	0	0	0	82	1	19.781	0	0	0
33	1	20.087	0	1	0	83	0	20.107	1	1	0
34	0	20.189	0	1	0	84	1	19.750	0	0	0
35	0	19.883	0	0	0	85	0	20.087	1	0	0
36	0	20.056	0	0	0	86	1	19.791	0	0	0
37	0	19.770	0	1	0	87	1	20.219	0	0	0
38	0	20.209	0	0	1	88	1	19.903	0	0	0
39	1	20.036	0	0	0	89	0	19.852	0	1	1
40	0	19.821	0	0	0	90	0	20.240	0	0	0
41	0	20.097	1	0	0	91	0	20.219	0	0	1
42	0	19.934	0	0	1	92	0	20.158	0	0	0
43	1	19.842	0	0	1	93	1	20.117	0	0	0
44	0	19.862	0	0	0	94	1	19.770	0	1	1
45	1	19.872	0	0	1	95	0	19.903	0	1	0
46	1	20.097	0	0	0	96	1	20.179	0	0	0
47	0	20.117	0	0	1	97	1	19.862	0	0	1
48	1	19.944	0	0	1	98	1	19.821	0	0	1
49	0	19.944	0	0	1	99	1	20.128	0	0	1
50	1	19.883	0	0	0	100	0	20.046	1	0	0

Figure A.1: Homogeneous Parameters

APPENDIX B

Heterogeneous Parameters

The parameters for the heterogeneous population of houses are randomly generated using a lognormal distribution and MATLAB's *lognrnd* function, which uses the lognormal standard deviation, σ_l, and lognormal mean, μ_l:

$$\sigma_l = \sqrt{\ln\left(1 + \frac{\sigma_n^2}{\mu_n^2}\right)} \tag{B.1}$$

$$\mu_l = \ln \mu_n - \frac{1}{2}\sigma_l^2 \tag{B.2}$$

where σ_n and μ_n are the normal standard deviation and mean, respectively.

House	StartState	C	R	Deadband	TempSet	SystemSizeTons	SystemSizekW	COP	CoolCredit10	CoolCredit20	CoolCredit30
1	1	602692.39	0.001913	0.5157	19.328	4.5	15750	2.332	0	0	0
2	1	607666.35	0.001911	0.5044	18.836	4.5	15750	2.606	0	0	1
3	1	622200.09	0.001879	0.4987	20.384	4.5	15750	2.589	0	0	0
4	1	591508.76	0.001979	0.4752	20.543	4	14000	2.587	0	0	0
5	1	598473.57	0.001833	0.5514	20.752	4.5	15750	2.579	0	0	0
6	1	602154.70	0.002013	0.4765	19.726	4	14000	2.340	0	0	0
7	1	603583.18	0.002081	0.4608	20.563	4	14000	2.396	0	0	1
8	1	657492.75	0.002056	0.4939	19.135	4.5	15750	2.427	0	1	0
9	1	633122.31	0.002107	0.4925	19.735	4.5	15750	2.468	0	0	0
10	1	588447.21	0.001919	0.4779	19.741	4	14000	2.606	0	0	0
11	0	613552.70	0.002079	0.5142	19.464	4.5	15750	2.340	0	0	1
12	0	574644.77	0.001966	0.4904	19.049	4.5	15750	2.488	0	0	1
13	1	583775.15	0.002119	0.5099	20.216	4	14000	2.237	0	1	0
14	1	559076.46	0.002006	0.5325	20.541	4	14000	2.607	0	0	1
15	0	599378.62	0.001860	0.4655	20.328	4	14000	2.579	0	1	0
16	0	558629.20	0.001858	0.5106	20.252	4	14000	2.352	0	0	0
17	0	683379.31	0.001825	0.5367	18.924	5	17500	2.493	0	0	0
18	1	624300.67	0.001857	0.5078	18.501	4.5	15750	2.410	0	0	0
19	1	557551.97	0.001984	0.5124	19.999	4	14000	2.846	0	0	1
20	1	584472.86	0.001900	0.4936	21.607	4	14000	2.871	0	1	0
21	1	583940.70	0.001742	0.5146	18.041	5	17500	2.577	0	0	0
22	0	547630.14	0.002174	0.5149	19.771	4	14000	2.735	0	0	0
23	0	629592.09	0.001969	0.4897	19.797	4.5	15750	2.478	1	0	1
24	0	588545.85	0.002113	0.5174	19.965	4	14000	2.638	0	0	1
25	0	634145.76	0.001967	0.4543	18.771	4.5	15750	2.454	0	0	0
26	0	611625.01	0.001981	0.5344	19.284	4.5	15750	2.354	1	0	1
27	1	624287.95	0.001989	0.4771	18.912	4.5	15750	2.707	0	0	1
28	0	608900.16	0.001994	0.4957	21.134	4	14000	2.549	0	0	0
29	1	676110.24	0.002218	0.4895	21.123	4	14000	2.334	0	0	0
30	0	595846.56	0.001957	0.4852	19.751	4	14000	2.676	1	0	0
31	1	602795.70	0.002316	0.5520	20.047	4.5	15750	2.586	0	0	1
32	1	629946.29	0.001836	0.4868	18.645	4.5	15750	2.537	0	0	0
33	0	588636.84	0.002027	0.5016	20.005	4	14000	2.519	0	0	0
34	0	597605.97	0.001856	0.4989	19.582	4.5	15750	2.407	0	0	0
35	0	606525.33	0.002043	0.4813	21.737	4	14000	2.525	1	0	0
36	0	522699.93	0.001835	0.4714	20.743	4	14000	2.651	0	0	0
37	0	605074.07	0.002004	0.4531	20.505	4	14000	2.324	0	0	1
38	0	594913.47	0.001886	0.5000	19.507	4.5	15750	2.309	0	0	0
39	0	644428.97	0.002207	0.4792	19.955	4	14000	2.644	0	0	0
40	1	587489.94	0.002092	0.4818	23.366	3.5	12250	2.554	0	0	0
41	1	653323.51	0.002166	0.4532	21.587	4	14000	2.515	0	1	0
42	1	606841.71	0.001949	0.5246	20.037	4.5	15750	2.399	0	1	0
43	0	581012.84	0.001882	0.4596	20.425	4	14000	2.357	0	0	1
44	0	629211.33	0.001861	0.5536	20.192	4.5	15750	2.613	0	0	0
45	0	603186.63	0.002045	0.4500	19.811	4	14000	2.457	1	0	0
46	1	588439.86	0.002135	0.5062	19.716	4	14000	2.332	0	1	1
47	1	596463.78	0.002241	0.4435	20.104	4	14000	2.567	0	0	0
48	1	614099.06	0.002084	0.5089	20.834	4	14000	2.410	0	0	0
49	1	642719.92	0.002050	0.4752	20.753	4	14000	2.644	0	0	1
50	0	591096.10	0.002344	0.5017	20.196	4	14000	2.675	0	0	0

Figure B.1: Heterogeneous Parameters Houses 1-50

House	StartState	C	R	Deadband	TempSet	SystemSizeTons	SystemSizekW	COP	CoolCredit10	CoolCredit20	CoolCredit30
51	1	619833.53	0.001965	0.5157	21.041	4.5	15750	2.495	1	0	1
52	0	604985.39	0.001910	0.4802	18.696	4.5	15750	2.531	0	0	0
53	1	619664.66	0.001901	0.5031	20.620	4.5	15750	2.501	0	0	0
54	0	605523.01	0.001819	0.4769	18.369	4.5	15750	2.546	1	1	0
55	1	609150.11	0.002153	0.5300	19.664	4.5	15750	2.584	0	0	0
56	1	575101.78	0.001841	0.5245	19.388	4.5	15750	2.390	0	0	0
57	0	609272.81	0.001914	0.4816	20.396	4	14000	2.634	0	0	0
58	0	609638.51	0.002108	0.5612	17.802	5	17500	2.518	1	0	1
59	1	610731.10	0.001954	0.4396	19.216	4	14000	2.330	0	0	0
60	1	632761.89	0.002359	0.5224	18.421	4.5	15750	2.376	0	0	0
61	0	629794.22	0.002008	0.5051	20.406	4.5	15750	2.545	0	1	0
62	1	605530.65	0.001922	0.4907	20.212	4.5	15750	2.340	0	1	1
63	0	587416.60	0.002066	0.5127	19.797	4	14000	2.377	0	0	1
64	0	614596.51	0.001929	0.5217	19.770	4.5	15750	2.374	1	1	1
65	0	595437.86	0.001990	0.5821	21.229	4.5	15750	2.703	0	1	0
66	0	591666.50	0.002000	0.4842	19.936	4	14000	2.558	0	1	0
67	1	596919.03	0.002208	0.5186	20.054	4	14000	2.620	0	0	0
68	0	589903.27	0.001888	0.5257	19.304	4.5	15750	2.430	0	0	1
69	0	590086.96	0.002064	0.5007	20.083	4	14000	2.530	0	0	1
70	0	611388.20	0.002144	0.5047	19.783	4	14000	2.539	0	0	0
71	1	595388.55	0.001917	0.5290	19.794	4.5	15750	2.544	0	0	0
72	1	640355.79	0.002003	0.4907	19.579	4.5	15750	2.369	0	0	0
73	1	619677.76	0.001914	0.5124	19.767	4.5	15750	2.635	0	1	0
74	0	578892.65	0.002182	0.5240	19.872	4	14000	2.678	0	1	0
75	0	597779.42	0.002016	0.5091	21.348	4	14000	2.507	0	0	1
76	0	589547.25	0.002041	0.4938	19.441	4	14000	2.362	0	1	0
77	0	608400.98	0.001975	0.5056	20.977	4	14000	2.616	0	0	0
78	1	664106.18	0.002037	0.4538	20.252	4	14000	2.327	0	0	1
79	0	652615.08	0.001946	0.5174	21.475	4.5	15750	2.436	0	1	0
80	1	614741.81	0.001768	0.5399	19.671	5	17500	2.598	0	0	0
81	0	593657.54	0.002013	0.4792	19.263	4	14000	2.537	0	0	1
82	0	581858.99	0.001916	0.4854	21.293	4	14000	2.735	0	0	0
83	1	634482.28	0.002005	0.4816	17.398	4.5	15750	2.513	0	0	1
84	0	603167.86	0.002133	0.4748	19.416	4	14000	2.662	0	0	0
85	1	617072.21	0.001984	0.5179	22.933	4	14000	2.617	0	0	0
86	1	602355.59	0.001862	0.5015	20.036	4.5	15750	2.386	0	0	0
87	1	607225.38	0.001981	0.4888	18.018	4.5	15750	2.431	0	0	0
88	0	600894.55	0.002220	0.4847	20.006	4	14000	2.531	0	0	0
89	1	605913.31	0.002149	0.4811	20.303	4	14000	2.473	0	0	0
90	1	601982.68	0.002091	0.5157	20.719	4	14000	2.255	0	0	0
91	0	594773.18	0.002154	0.4946	21.374	4	14000	2.566	0	0	0
92	1	582677.53	0.001953	0.4962	19.761	4	14000	2.376	0	1	0
93	0	547727.62	0.001887	0.4975	20.518	4	14000	2.330	1	0	0
94	1	585320.49	0.001928	0.5282	20.557	4.5	15750	2.632	0	0	1
95	0	572824.83	0.001965	0.4931	19.960	4	14000	2.286	0	0	0
96	0	615882.49	0.001851	0.4896	20.035	4.5	15750	2.385	0	0	0
97	0	577635.92	0.002014	0.5093	19.080	4.5	15750	2.196	0	1	0
98	1	645114.82	0.001949	0.4827	19.459	4.5	15750	2.681	0	0	1
99	1	614590.74	0.002057	0.4953	18.954	4.5	15750	2.509	0	0	1
100	0	620355.56	0.001909	0.4868	20.283	4.5	15750	2.415	0	1	1

Figure B.2: Heterogeneous Parameters Houses 51-100