

Reducing Transient and Steady State Electricity Consumption in HVAC using Learning-Based Model Predictive Control

Anil Aswani, *Member, IEEE*, Neal Master, *Student Member, IEEE*, Jay Taneja, *Student Member, IEEE*, David Culler, *Fellow, IEEE*, and Claire Tomlin, *Fellow, IEEE*

Abstract—Heating, ventilation, and air conditioning (HVAC) systems are an important target for efficiency improvements through new equipment and retrofitting because of their large energy footprint. One type of equipment that is common in homes and some offices is an electrical, single-stage heat pump air conditioner (AC). To study this setup, we have built the Berkeley Retrofitted and Inexpensive HVAC Testbed for Energy Efficiency (BRITE) platform. This platform allows us to actuate an AC unit that controls the room temperature of a computer laboratory on the Berkeley campus that is actively used by students, while sensors record room temperature and AC energy consumption. We build a mathematical model of the temperature dynamics of the room, and combining this model with statistical methods allows us to compute the heating load due to occupants and equipment using only a single temperature sensor. Next, we implement a control strategy that uses learning-based model predictive control (MPC) to learn and compensate for the amount of heating due to occupancy as it varies throughout the day and year. Experiments on BRITE show that our techniques result in a 30-70% reduction in energy consumption as compared to two-position control, while still maintaining a comfortable room temperature. The energy savings are due to our control scheme compensating for varying occupancy, while considering the transient and steady state electrical consumption of the AC. Our techniques can likely be generalized to other HVAC systems while still maintaining these energy saving features.

Index Terms—Air conditioning, building automation, energy efficiency, learning, model predictive control.

I. INTRODUCTION

BUILDINGS account for 73% of the electricity and 40% of greenhouse gas emissions in the United States [1], [2]. Heating, ventilation, and air conditioning (HVAC) compose 33% of building energy usage, making this an attractive target for reductions [1]. Several parallel directions are being taken towards the aim of reducing HVAC energy, one of which is the design of new, more efficient equipment. However, buildings and equipment are often slowly replaced [3]. This has led to interest in retrofitting HVAC to improve efficiency.

A. Aswani, N. Master, J. Taneja, D. Culler, and C. Tomlin are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94709 USA e-mail: {aaswani,tomlin}@eecs.berkeley.edu; {taneja,culler}@cs.berkeley.edu; neal.m.master@berkeley.edu.

This material is based upon work supported by the National Science Foundation under Grants #CNS-0931843 (CPS-ActionWebs) and #CNS-0932209 (CPS-LoCal) as well as the California Energy Commission Cory Hall Building-to-Grid Testbed project. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the National Science Foundation.

The simplest way to retrofit is to only change the software that controls the HVAC, but this is a challenging problem because of the large variety of physical effects that are used by HVAC equipment. Many homes use a single-stage heat pump that cools air at a constant rate for the entire building. In contrast, some large buildings use variable air volume (VAV) systems to centrally cool air that is partitioned into different amounts for each room. Some HVAC systems incorporate thermal storage tanks that freeze liquid at night and then provide cooling by allowing it to melt during the day.

HVAC equipment requires a separate design process tailored to its particular physical modalities. Within the cyber-physical system (CPS) community, the focus of research has been on modeling and control for VAV systems [4]–[7] or thermal storage tanks [8]. Occupants and equipment generate heat that raises the temperature of rooms, and existing HVAC control struggles with these effects because of their significant variation over time. Current work in this area concerns combining occupancy sensors with models of human behavior to estimate the number of occupants in different rooms [9].

We focus on reducing the energy usage for an electric, single-stage heat pump air conditioner (AC) that cools a single area, and our work is distinguished from past work by three aspects. First, this HVAC equipment is common in homes and has not been extensively studied by the CPS community. Second, modeling and statistics are used to estimate the heating load (i.e., amount of thermal energy transfer to the building) of occupants and equipment using only a temperature sensor. Third, we design a control scheme that improves efficiency by explicitly adapting to the occupancy heating load. These techniques are expected to generalize to other HVAC systems, though implementation and modeling details will vary depending upon equipment physics.

The second point is important from a general CPS viewpoint. One approach to solving CPS problems is to study the integration and communication of large numbers of sensors. For this particular application, we take an alternative approach. We evaluate how more intelligent computation may enable a reduction in the number of sensors needed to achieve a given task. This is done by constructing mathematical models that incorporate the physical aspects of the system, and then designing statistical schemes that combine these models with measurements to reduce the amount of needed infrastructure.

In order to conduct experiments, we have built the Berkeley Retrofitted and Inexpensive HVAC Testbed for Energy

Efficiency (BRITE). The BRITE testbed controls the room temperature of a computer laboratory on the Berkeley campus that is actively used by students. A computer actuates the AC unit by relaying computed control actions through a local area network (LAN) to the thermostat. Sensors are able to measure the room temperature and power consumption of the AC.

BRITE is a living laboratory, and so its large variations in weather and occupant behavior make it difficult to directly compare different control strategies. To overcome this challenge, we cyber-physically compare control methods using a mixture of experiments and simulations. This allows for a more fair comparison by using identical weather and occupancy conditions, but this does introduce some error into the comparisons because of modeling mismatch. To alleviate these issues, we alternate between the control schemes we use for experimentation and simulations.

We implement a new control technique on BRITE known as learning-based model predictive control (MPC) [10], which has provable properties considering its safety and robustness. It combines models with statistics to estimate occupancy heating load from only temperature measurements and then compensate for it within the control action, thereby reducing the amount of room overcooling and thus saving energy. Our experiments show that learning-based MPC reduces energy consumption by 30-70% compared to two-position control, which is the control scheme used by a typical thermostat [11].

A. CPS Aspects of Single-Stage Heat Pump Control

Heat pumps [12] are commonly used in homes to provide AC, and they use electrical energy to run a motor that compresses gas. It is the subsequent expansion of this compressed gas that is able to provide cooling for air that is then delivered to the entire building. Most heat pumps have motors with one fixed speed and are called single-stage. Multi-stage heat pumps can run the compressor motor at different speeds, but they are less common. We focus on single-stage equipment in our work because (a) this is the existing equipment in the room that we use in BRITE, and (b) the control scheme we develop can be extended to multi-stage equipment through the use of appropriate pulse-width modulation (PWM).

The physics of heat pumps leads to particular energy characteristics, and similar behavior occurs in other HVAC systems. Understanding these features is important for the design of efficient HVAC systems, and is reflective of insights gained from a CPS viewpoint. A tighter integration of the software to the physics of the HVAC allows for improved performance by reducing the conservativeness of the control schemes. In the case of HVAC, comfort is equivalent to keeping the room temperature within a range of temperatures [13], and conservativeness is how close the temperature is kept towards the boundaries of this range.

Keeping temperature near the boundaries of comfort uses less energy, because less heat transfer is required. In the case of AC, two-position control means that the AC is turns on when the room temperature exceeds T_{on} and turns off when the temperature is below T_{off} . Because of its physics, the AC continues to cool for a few minutes even after it is turned

off [14]. This actually represents overcooling and is a major cause of inefficiency. The obvious fix is to set T_{off} to be nearly identical to T_{on} ; however, this is not practical because of the physics of the equipment.

A heat pump has high transient power consumption when it is turned on, and then it uses lower amounts of power at steady state. This transient power is due to inrush current (a brief period of high current flow when turning on the electric motor that drives the compressor in the heat pump) and an increased load on the electric motor at startup (the pressure of the gasses in the heat pump are initially out of equilibrium [15]). The transient power consumption of a heat pump puts a limit on its switching frequency, because otherwise the equipment behaves inefficiently and can also be damaged.

The high transient power usage acts as a penalty for turning the AC on. Efficient control schemes need to balance the efficiency from turning the heat pump on and off frequently (this reduces overcooling) with the added energy consumption and physical fatigue of frequent switching. This tradeoff can be handled by the learning-based MPC technique, which picks control actions for the AC that minimize a cost (consisting of steady state energy consumption, transient energy consumption from switching, and deviation from desired temperature) subject to the thermal dynamics of the room and constraints on the allowed temperatures of the room.

II. BERKELEY RETROFITTED AND INEXPENSIVE HVAC TESTBED FOR ENERGY EFFICIENCY

BRITE is a system for testing different control strategies on an AC unit that cools a computer laboratory on the Berkeley campus, and it is shown in Fig. 1. Though it is built using commodity parts, the computers can be replaced with microcontrollers. The strength of this structure is that it scales to building-wide systems. Moreover, our MPC schemes are computationally scalable because of their convexity.

In this testbed, the LoCal server gathers sensor data and stores this in a Simple Measurement and Actuation Profile (sMAP) database [16]. A control computer accesses the Internet and LoCal server to get weather forecasts and sensor data, and it runs a learning-based MPC scheme that computes a control input that is sent through the LoCal server to the thermostat. The thermostat transmits a corresponding signal to the AC.

A. LoCal

The Berkeley LoCal project aims to produce a network architecture for localized electrical energy reduction, generation, and sharing by examining how pervasive information can fundamentally change the nature of these processes [17]. A key component of this is the use of sMAP [16] to exchange physical data about the systems involved. This allows producers of physical information to directly publish their data in a format for consumption by a diverse set of clients. We use temperature measurements from a networked thermostat in BRITE, though we also have the capability to measure plug-load [18] and wireless temperature readings. The ability to easily integrate streams of sensor data is critical to the scalability of BRITE to entire buildings.




Fig. 1. The Berkeley Retrofitted and Inexpensive HVAC Testbed for Energy Efficiency (BRITE) is a system built on the Berkeley campus that allows testing of different control strategies for controlling an AC in order to explore tradeoffs between energy consumption and tracking a temperature setpoint.

B. Room, Air Conditioner, and Thermostat

The BRITE testbed shown in Fig. 1 is deployed in a student computing laboratory on the ground floor of a large engineering building. The room is 640 square feet, has external windows on its south and west walls, and contains 16 desktop computer workstations and two laser printers. Occupancy of the room peaks at over 20 individuals and is constantly varying depending upon the due dates of projects, assignments, and exams. An important reason that this room was chosen for the testbed is that it has its own HVAC equipment, which allows us to do experiments independent of other rooms.

A Proliphix brand NT160e model thermostat controls the AC. It is a modern thermostat with networking functionality that allows computers on a shared network to communicate with it and control it. We can transmit commands to the thermostat and also receive diagnostic information on the thermostat settings, current HVAC state, and room temperature.

C. Control Computer

We use a dedicated control computer to avoid disrupting processes running on the LoCal computer, though their combined functionality can be implemented on a single microcontroller. The control computer runs a 64-bit version of the Ubuntu operating system, and the control loop is implemented in MATLAB: The learning-based MPC [10] uses the SNOPT solver [19] from the TOMLAB library, and polytopes are handled using the MPT toolbox [20]. A Python script downloads weather forecasts from NOAA's National Weather Service.

D. Metrics for Human Comfort

The objective of BRITE is to minimize energy consumption while keeping occupants comfortable. However, there are multiple ways to quantify comfort. The ANSI/ASHRAE

standards [13] are defined in terms of the predicted mean vote (PMV), which is a complex function of indoor air temperature, human activity, relative air velocity, the occupants' clothing, and other variables that are difficult to measure [21]. The Occupational Safety and Health Administration (OSHA) [22] does not have regulations but provides guidelines of $68\text{--}76^\circ\text{F}$ (about $20\text{--}24.4^\circ\text{C}$). Alternative metrics are defined in terms of on temperature deviation: Category A, B, and C thermal requirements [13] respectively dictate a temperature range of 2°C , 4°C , and 6°C .

These alternatives specify temperature bands and simplify the design of HVAC systems. In experiments on the BRITE testbed, we keep the temperature near the middle of comfort (22°C) and try to satisfy Category A requirements, because these are the strictest and consume the most energy. More specifically, Category A is used as a range preference for the learning-based MPC and Category B are hard constraints on the temperature range. Future directions can consider smart methods for switching between different Category requirements based on, for instance, network-level load and demand signaling or occupancy estimates.

III. ELECTRICAL CHARACTERISTICS AND ENERGY CONSUMPTION OF A SINGLE-STAGE HEAT PUMP

Fig. 2 shows experimentally measured data of a typical power consumption profile for the HVAC in BRITE. A striking feature is that there are both transient and steady state behaviors. There is a transient spike in power consumption immediately after the heat pump is turned on that lasts for about one minute, before the power consumption reaches a steady state. Intuitively, the large transient is a penalty for turning the heat pump on. Physically, the transient power consumption is due to inrush current drawn by the electric motor in the compressor of the heat pump, as well as due to non-equilibrium pressure conditions in the heat pump [15].

As mentioned earlier, this profile highlights some important issues regarding energy usage. The transient spike in power consumption suggests an objective that minimizes switching. This is not typically considered in an explicit manner, and in fact makes the implementation of a controller in digital hardware difficult unless some approximation is used. Furthermore, the steady state energy usage is linear in the control, which matches the cost used in [5] and differs from the more commonly used quadratic cost [6].

A. Pulse Width Modulation Control

The single-state heat pump is strictly speaking a hybrid system [23], [24] because it has two modes corresponding to the pump being on or off. Fortunately, we can considerably simplify the design of a controller by considering sampled control. As such, we use MPC to compute a new control action $u[k]$ at intervals of once every 15 minutes. We chose this rate because switching more frequently than once every 10 to 15 minutes can physically damage the heat pump. Pulse-width modulation (PWM) is used to convert the discrete time control $u[k]$ into a continuous signal that turns the AC on or off [25], and so $u[k]$ can also be interpreted as a duty cycle.

There is an important note to make regarding why the constraints on the input for the MPC are $u[k] \in [0, 0.5]$. The reason for the choice of 0.5 as an upper value is because the thermostat does not stop cooling the room when it is turned off. This is discussed in more detail in Sect. IV, but the choice of 0.5 ensures that the control action at one discrete time sample does not affect the control action at the next one.

B. Measuring the Electrical Energy Consumption in BRITE

It is important to be able to compute the energy consumed by the AC in the BRITE platform, given the input that the AC receives. An estimate of the energy consumption is used in the cost function of MPC, and it is useful for being able to compare different control schemes. To be able to provide this equation, we need to make a few definitions. Define the vector $\mathbf{u}_m = (u[m] \dots u[m + N - 1])$. The term $\|\mathbf{u}_m\|_0$ counts the number of nonzero entries in the vector \mathbf{u}_m . Also, the values r, λ are constants which are used to compute the energy consumption. The value N is the number of discrete time steps (recall that each time step corresponds to 15 minutes) over which the energy consumption is to be computed.

The steady state energy consumption of the AC over $N/4$ hours in units of (kWh) is given by $\sum_{k=0}^{N-1} r/4 \cdot u[m+k]$, where $r = 3.7kW$ is the average rate of steady state energy consumption in the BRITE platform (compare to Fig. 2) and the value 4 is used to compensate for the fact that $u[m+k]$ is the control for 1/4 of an hour. Furthermore, the AC consumes $\lambda = 0.015kWh$ of energy every time we turn the AC on; this corresponds to the area of the triangle in Fig. 2 formed by the transient energy. The total energy used over N time steps is given by

$$E_{actual} = \sum_{k=0}^{N-1} r/4 \cdot u[m+k] + \lambda \|\mathbf{u}_m\|_0. \quad (1)$$




Fig. 2. Experimental data of a typical power consumption profile during actuation in the BRITE testbed is shown. The first vertical dashed line indicates the time when the heat pump turns on, the second indicates the time when the power reaches steady state, and the third indicates the time when the heat pump turns off. For this particular power profile, the amount of total energy consumed by the transient and steady state is labeled in the legend.

Unfortunately, the $\|\mathbf{u}_m\|_0$ term is not convex in \mathbf{u}_m . Convexity is important for the computations of the MPC that has to solve an optimization problem at each step. To simplify the computations, we make a standard convex relaxation [26] and replace the term $\|\mathbf{u}_m\|_0$ with $\|\mathbf{u}_m\|_1$. This relaxation is powerful: When it is used in the cost function of an optimization problem, it actually leads to having many of the $u[m+k]$ be equal to exactly zero [26]. In this way, it can reduce switching of the AC.

This approximation yields a convex equation for the energy consumed $\sum_{k=0}^{N-1} r/4 \cdot u[m+k] + \lambda \|\mathbf{u}_m\|_1$. However, we have $u[m+k] \geq 0$, and so we can further simplify the convex cost for energy consumption to

$$E_{convex} = \sum_{k=0}^{N-1} (r/4 + \lambda) \cdot u[m+k]. \quad (2)$$

What is surprising about this is that a cost for energy that is linear in the length of control action automatically considers a cost for switching, as long as the inputs $u[m+k]$ are constrained to be non-negative. Stated in another way, this means that a cost that is linear in the duty cycle of the control inherently considers the tradeoff between switching too frequently and the length of the duty cycle.

In practice, (1) is used if the actual energy needs to be computed. On the other hand, (2) is used if a control action needs to be computed by the MPC. Having these two formulations gives considerable flexibility.

IV. SYSTEM IDENTIFICATION OF COOLING DYNAMICS

An important step towards realizing efficient control schemes for the BRITE testbed is building a mathematical model that describes the impact of weather, occupancy, and AC operation on the temperature of the room. It is important because all MPC schemes inherently require a nominal model in order to be able to optimize system performance. More importantly, identifying a model allows us to estimate the heating load due to occupancy from only temperature measurements.

This enables evaluation of the importance of occupancy [6], [27] and techniques that compensate for it.

Though building simulator software [28], [29] models complicated thermodynamic and fluid effects, experimental data collected from buildings show that linear models with exogenous inputs [4]–[6], [27] can often be used to describe many rooms. The main physical effect is convective heat transfer and is described by Newton’s law of cooling. This is a linear ordinary differential equation (ODE), and so it may be abstracted as a resistor-capacitor network [4]–[6].

A. Discrete Time Model

There is a “delay” from when the AC is turned off and when it stops cooling the room, due to the dynamics of the heat pump. Specifically, the evaporator which cools the air does not instantly warm up and continues to cool air for some time after the heat pump is turned off [14]. We begin with a discrete time model where each time step is separated by $T_s = 15$ minutes. The advantage of this approach is that the AC behavior gets lumped into a single term that encompasses the modes where the AC is on and then turned off but still cooling. This makes it easier to do the modeling.

With this approach and inspired by the physics of convective heat transfer, we start with the model

$$T[n + 1] = k_r T[n] - k_c u[n] + k_w w[n] + q[n], \quad (3)$$

where $T[n] \in [15, 30]$ is the temperature of the room in ($^{\circ}C$), $k_r > 0$ is the time constant of the room, $k_c > 0$ is the change in temperature over 15 minutes in ($^{\circ}C$) caused by cooling for a duty cycle of $u[n] \in [0, 0.5]$, $k_w > 0$ is the time constant for heat transfer from the room to the outside, $w[n]$ is the outside temperature in ($^{\circ}C$), and $q[n]$ is change in temperature over 15 minutes due to the heating from occupancy of humans and equipment within the room, as well as other external inputs, in ($^{\circ}C$). The time constants here are dimensionless.

B. Parameter Identification

We collected data from 12:00PM to 5:30AM on a weekday using the BRITE testbed. This portion of the day was used because it exhibits a variety of occupancy levels. The room is actively used by students during the afternoon and evening, with fewer students using the room late at night and early in the morning.

Generally speaking, parameter estimation is usually more accurate when the control inputs are independent of the system states or external inputs (i.e., weather and occupancy). To ensure that this was the case, we actually applied a random input with uniform distribution over $[0, 0.5]$ at each discrete time step; the corresponding PWM control is shown in Fig. 3a. Because this only needs to be done once and over a span of about a day, it may be reasonable to allow the temperature in actual implementations to be unregulated for this day. Future work includes designing methods that keep the temperature in a comfortable range while still sufficiently exciting the system.

Because we have measurements of T , $w[n]$, and $u[n]$, the model is linear with respect to the parameters k_r , k_c , and k_w . On the other hand, $q[n]$ is not known and is expected to be

highly nonlinear with respect to time, because it incorporates heating due to human occupancy and equipment in the room. Consequently, standard linear system identification techniques cannot be used. Identification of models with the form given in (3) more generally falls into the class of problems known as semi-parametric regression of partially linear models [30], [31]. An alternative approach is to parametrize $q[n]$, with say a polynomial or spline basis, and then identify all parameters using nonlinear regression. The difficulty with this is the uncertainty associated with $q[n]$.

Using the technique given in [31], we identified the parameters of the model

$$T[n + 1] = 0.64 \cdot T[n] - 2.64 \cdot u[n] + 0.10 \cdot w[n] + q[n], \quad (4)$$

where $q[n]$ is shown in Fig. 3b. The experimental room temperature is the solid line in Fig.3c. Similarly, the temperature simulated by the model (4) is the dashed line shown in Fig.3c, and the initial condition for the simulation was taken from the experimental measurements. Furthermore, the simulation was conducted with the same inputs as were applied to the real BRITE system. The root-mean-squared (RMS) error of the simulation is $0.10^{\circ}C$. The plots show that the model fits reasonably well to the measured temperature data.

C. Impact of Occupancy

The identified model (4) shows that the role of occupancy is significant in the temperature dynamics of the room, confirming the intuition of [6] and results of [27]. The function $q[n]$ has an average value of $6.98^{\circ}C$, and it is highly nonlinear with respect to time: It varies by up to $0.61^{\circ}C$ depending on what time of day it is. Furthermore, there are fluctuations over both long and short time horizons.

The heat generated by occupancy and equipment $q[n]$ displays interesting features. The room is a computer laboratory used by students at their own convenience and shows characteristics consistent with this role. The heat input $q[n]$ increases from lunchtime and peaks at 1PM, while the outside temperature peaks at 2PM. The occupancy has quick changes in its direction at 3PM and 5PM. Finally, it is relatively constant from 8PM to 5AM, which is typically when there are few or no students in the room.

The large fluctuations have a major impact on the design process of a control scheme. This is because the nominal model for which a controller is designed can be inaccurate by $0.61^{\circ}C$ (in our case) because of varying levels of occupancy. This causes issues with respect to efficiency, because standard MPC requires accurate models to provide high performance. It is for this reason that we make use of learning-based MPC [10] to design the controller. It will estimate occupancy by measuring the temperature of the room and comparing it to what is expected by the model (4).

D. Modeling the Two-Position Control of Thermostat

For the purpose of comparing the energy consumption of different control strategies, it is useful to identify a model of the two-position control of the thermostat. The thermostat does its control in continuous time, and so this model is an ordinary

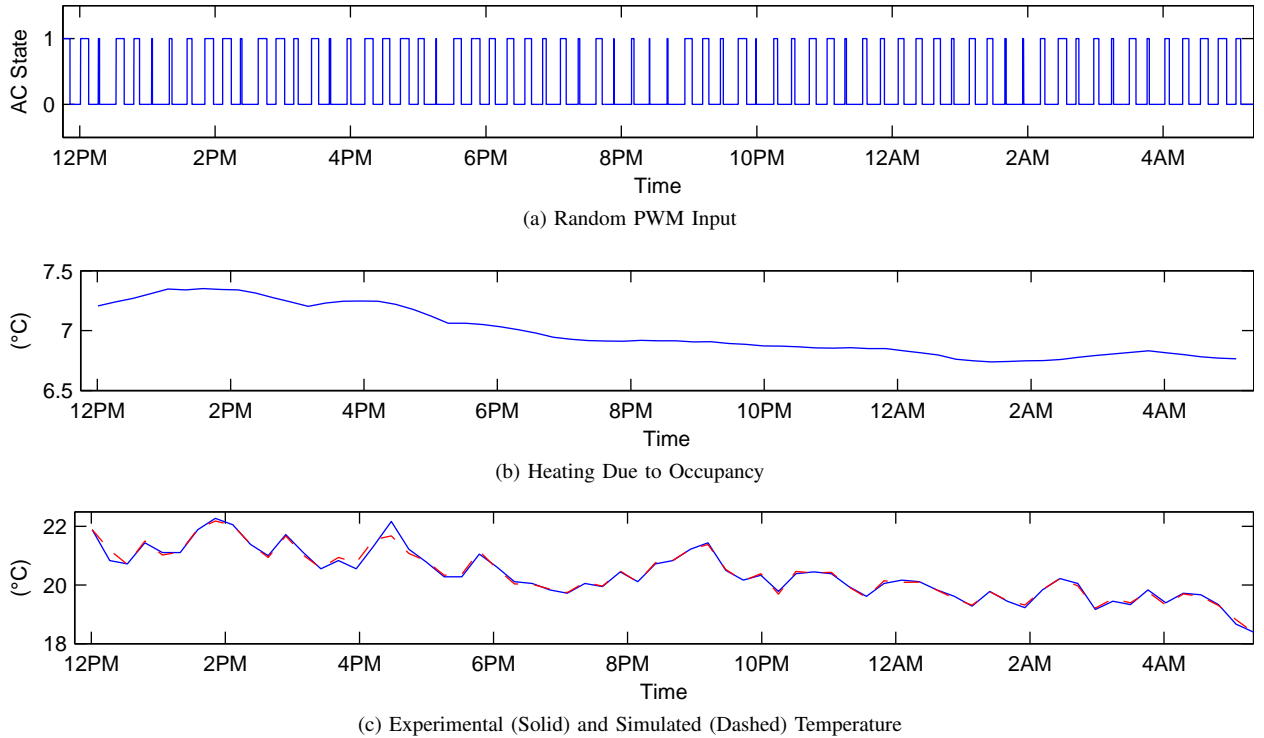


Fig. 3. At each discrete time step, we applied a randomly generated input, which is the duty cycle of the PWM over 15 minute periods, taken from a uniform distribution ranging over $[0, 0.5]$. This was done over a period of the day (12:00AM to 5:30AM) during which the room is both in and not in use. Using semi-parametric regression [31], we identified both a discrete time model (4) and the term $q[n]$ which is given in units of ($^{\circ}C$) and includes heating due to occupancy, equipment, and other external inputs. The measured room temperature is given in units of ($^{\circ}C$) by the solid line, and a simulation of our model in units of ($^{\circ}C$) is shown by the dashed line. The simulation uses the same inputs as provided to the BRITE platform over this range, and the initial condition of the simulation is taken to be the experimentally measured temperature. The simulation has a root-mean-squared (RMS) error of $0.10^{\circ}C$.

differential equation. Part of the model is derived from a statistical analysis of temperature data from BRITE gathered over a 20 hour period. On average, the thermostat turns the AC on when the temperature reaches $22.8^{\circ}C$ (standard deviation of less than $0.1^{\circ}C$), and it turns the AC off when the temperature reaches $22.4^{\circ}C$ (standard deviation of $0.1^{\circ}C$). The thermostat has a feature called a heat anticipator that adjusts the top and bottom temperature thresholds, in an effort to conserve energy and reduce overcooling. We do not model this behavior. Furthermore, it takes the AC an average of 354 seconds (standard deviation of 75 seconds) to stop cooling the room after it is turned off. Though this is due to the internal dynamics of the heat pump, we approximate this by assuming that the AC stops cooling after a fixed time.

We again used semi-parametric regression on data from BRITE to estimate a continuous time model for the two-position control. The time constants for the room and heat transfer to the outside were taken from the discrete time model (4) and converted into continuous time constants by doing the reverse of an exact discretization. The model identified is

$$\dot{T} = -5.0 \times 10^{-4} \cdot T + 1.4 \times 10^{-4} \cdot w(t) - 1.2 \times 10^{-3} + q(t), \quad (5)$$

if the AC is turned on or for the first 354 seconds after it is turned off. Otherwise, the dynamics are given by

$$\dot{T} = -5.0 \times 10^{-4} \cdot T + 1.4 \times 10^{-4} \cdot w(t) + q(t). \quad (6)$$

In our model, the AC turns on when the temperature exceeds $T_{on} = 22.8^{\circ}C$, and it turns off when the temperatures goes below $T_{off} = 22.4^{\circ}$.

Visually examining the measured (Fig. 4a) and simulated (Fig. 4b) temperature under two-position control indicates that there are several modeling errors; many of these are previously mentioned, but we collect them into one location for clarity. The temperature in the simulation rises slower than on BRITE, and this indicates that the identified time constant is slower than it should be. Furthermore, the model does not incorporate the internal dynamics of the heat pump or the thermostat's heat anticipator logic. Also, there is variation in the steady-state and transient energy consumption that is not captured in (1), which is used to make energy estimates.

Despite the modeling errors and simplifications, the simulation and (1) are reasonable proxies. The true (Fig. 4a) and simulated (Fig. 4b) temperature of the BRITE platform under two-position control over a period starting at midnight share similar qualitative features. The occupancy heating $q(t)$ is not shown because it displays characteristics similar to Fig. 3b and that in [27]. Moreover, the true energy consumed by the AC was measured to be $8.6kWh$, computed by (1) to be $8.6kWh$, and simulated to be $9.0kWh$. This represents an error of less than 1% and 5%, respectively.

The overshoot of going below T_{off} seen in both the measured and simulated temperature is in some sense wasted energy because it represents overcooling of the BRITE space. And even though the thermostat in BRITE has a heat an-

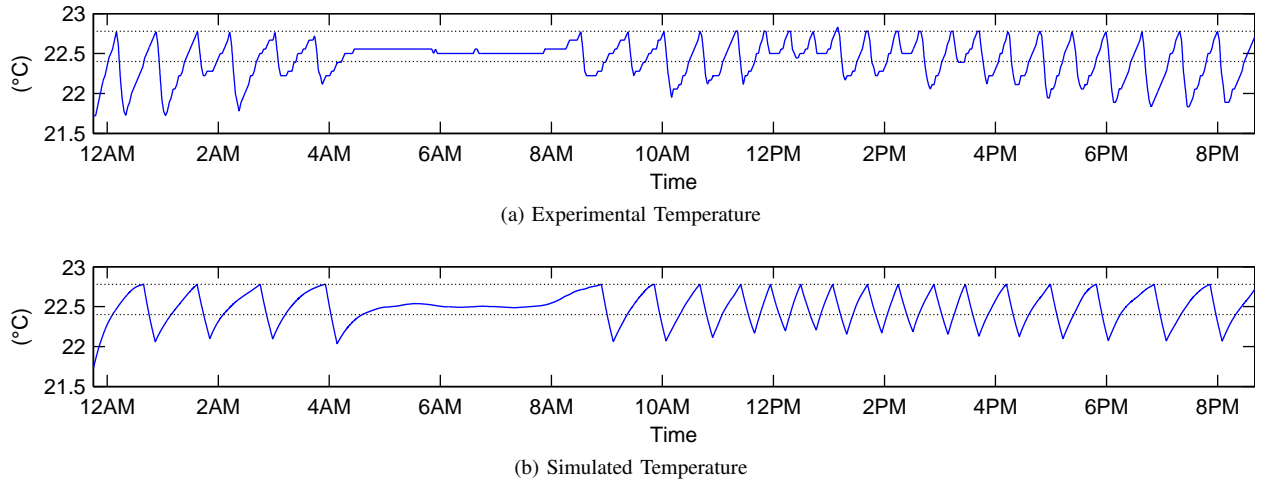


Fig. 4. The thermostat used two-position control to maintain the temperature. Its average on and off temperatures were $T_{on} = 22.8^\circ C$ and $T_{off} = 22.4^\circ C$, and these are shown by the solid, horizontal lines. The experimentally measured temperature is shown in ($^\circ C$). Semi-parametric regression was used to identify a continuous time model, and a simulation of this model using experimentally measured temperature as the initial condition is shown in ($^\circ C$). The control of the simulation is different than the experimental control, and it was determined using the T_{on} and T_{off} values. The energy estimated by the simulation is $9.0kWh$, and (1) applied to the measured inputs computes $8.6kWh$; this is in contrast to the measured consumption of $8.6kWh$. Despite modeling errors, the energy estimates differ from the true value by only 5% and less than 1%.

ticipator that adjusts the T_{on} and T_{off} of the two-position control, it cannot adequately compensate for variations in weather and occupancy. The learning-based MPC scheme we have developed can compensate for these factors, and so it can prevent overcooling and thus save energy.

V. LEARNING-BASED MPC OF BRITE

Safety and robustness can be guaranteed with approximate models, but maximum efficiency requires accurate models. This tradeoff has driven research in adaptive control [32], [33] and learning-based control [34]–[36]. Statistical methods by themselves cannot ensure robustness [37], [38], and so the approach of learning-based MPC [10] is to begin with an approximate model of the system and refine it with statistical methods. It is a rigorous control method that (a) handles state and input constraints, (b) optimizes system performance with respect to a cost function, (c) uses statistical identification tools to learn model uncertainties, and (d) provably converges.

The control situation is as follows. We have a model (4) for the cooling dynamics of the BRITE room, and we have constraints on the maximum ($24^\circ C$) and minimum temperature ($20^\circ C$) to ensure comfort for people in the room. Preliminary experiments [27] made use of tube-MPC [39]–[41] (a form of robust MPC [42]) to ensure that these constraints were never violated despite varying occupancy and uncertainties in the weather forecast. However, testing over an extended period of time showed that the robust MPC described in [27] was too conservative because, when tracking a desired temperature of ($22^\circ C$), the temperature rarely approached the constraints.

Consequently, we began to test standard linear MPC for its ability to stay within the desired temperature range. This control scheme had the same property in our tests—that it kept the room temperature within the constraints. It is important to remember this fact that a standard linear MPC ensures constraint satisfaction. However, the energy efficiency

of this base scheme was lacking. It was unable to stay close to the desired ($22^\circ C$), and it could use more energy than two-position control of the thermostat. Because of this, we implemented a learning-based MPC technique to control the room temperature.

A. Special Case of Learning-Based MPC

The main idea of this technique [10] is that we decouple performance from robustness. By robustness, we mean whether an MPC scheme can ensure constraint satisfaction despite modeling errors and other uncertainties. Linear MPC itself has certain robustness properties [43]. As a practical issue, our tests on the BRITE testbed show that standard linear MPC gives sufficient robustness. In more general cases, we would need to use tube-MPC to ensure enough robustness for learning-based MPC [10].

We use a tilde to denote the temperature predicted by the learning-based model, an overline denotes the temperature predicted by the linear model with constant occupancy term $6.98^\circ C$, and no overline indicates the measured temperature. The control action at time m , with temperature $T[m]$, control horizon $N = 20$ (5 hours), weight $p = 0.075$, and desired temperature $T_d = 22^\circ C$ is given by the minimizer to the following optimization problem

$$\min_{u[\cdot]} \sum_{k=0}^N p \cdot (\tilde{T}[m+k] - T_d)^2 + \sum_{k=0}^{N-1} (r + \lambda) \cdot u[m+k] \quad (7)$$

$$\text{s.t. } \tilde{T}[m+i] = 0.64 \cdot \tilde{T}[m+i-1] - 2.64 \cdot u[m+i-1] \quad (8)$$

$$\begin{aligned} &+ 0.10 \cdot w[m+i-1] + \hat{q}[m+i-1] \\ \bar{T}[m+i] &= 0.64 \cdot \bar{T}[m+i-1] - 2.64 \cdot u[m+i-1] \\ &+ 0.10 \cdot w[m+i-1] + 6.98 \end{aligned} \quad (9)$$

$$\bar{T}[m+i] \in [20, 24] \quad (10)$$

$$u[m+i-1] \in [0, 0.5] \quad (11)$$

for $i = 1, \dots, N$. The problem (7) generates an input $u[m]$ that minimizes the expected future performance of BRITE with respect to a cost function that encodes energy consumption and temperature deviation. Here, the term $\hat{q}[n]$ represents the predicted amounts of occupancy and is computed using learning. In our unoptimized MATLAB code, this computation (7) takes between one to two seconds. It is easily scalable to larger problems because (7) is simply a quadratic program.

The optimization problem (7) decouples performance and robustness in the following manner. Robustness is due to the constraints (9) which are nothing more than the identified model (4) with constant occupancy. Performance is due to the use of (8) in the cost function. The intuition is that the cost depends on the learned occupancy \hat{q} through (8), and the control is chosen such that the MPC without learning that is sufficiently robust (4) would satisfy the temperature and control constraints.

There are several important things to note about this formulation (7). The cost function contains (a) $p \cdot (T[m+k] - T_d)^2$ that represents deviation from the desired temperature and (b) the convex energy (2). This explicitly controls the tradeoff between keeping the temperature close to a comfortable value and the amount of energy used, and the value $p = 0.075$ was chosen because it gives a good tradeoff. Also, the convex energy (2) encourages a tradeoff between minimizing switching and duration of keeping the AC on, as discussed in Sect. III-B.

B. Learning Occupancy

Estimating occupancy is a detailed process that requires combining models of human behavior with sensors [9]. The BRITE platform faces an additional challenge because the occupancy varies immensely over the span of days and weeks, depending upon when assignments and projects are due. Some of the occupancy, such as for assignments, will likely be periodic in nature; other occupancy, like for projects, is more irregular and harder to predict. Furthermore, we need to know the heat generated by occupants and their use of computer equipment in the room for the purposes of energy efficient control. The correlation between the number of individuals in the room and the heat load will likely vary depending upon how many computers are in use.

Instead of relating the number of individuals in the BRITE room to the heat load $q[n]$, we focus our efforts on estimating this $q[n]$ directly from the temperature measurements and our model (4). We use the estimate

$$\hat{q}[m+i] = T[m] - \left(0.64 \cdot T[m-1] - 2.64 \cdot u[m-1] + 0.10 \cdot w[m-1] \right), \quad (12)$$

for $i = 0, \dots, N-1$. The intuition is that the occupancy heating $q[n]$ is the discrepancy in what the linear model without the occupancy term 6.98 predicts the temperature at the next time step is and what the actual temperature is.

The approach we take in this paper is to use the simplest possible estimate — more accurate estimates of $q[n]$ taking into account specific models will only improve the energy efficiency of the BRITE testbed. An obvious extension is to

fit our estimates to curves of best fit (e.g., a line or parabola) to compensate for the time-varying nature of the occupancy. Other extensions are to incorporate models of human behavior and other sensors.

We used this estimate of the occupancy for several reasons. As mentioned, this is the easiest estimate in terms of modeling: We do not need to worry about how to model long and short term human behavior. Secondly, it is well-behaved. Extrapolations using curves of best fit can significantly over- and under-estimate on long time-horizons. Finally, this estimate is easy to compute and shows that the learning can be done in a scalable manner.

VI. EVALUATING THE ENERGY EFFICIENCY OF LEARNING-BASED MPC WITH BRITE

The original aim of building the BRITE platform [27] was to enable evaluation of existing methods and design new control schemes that minimize the energy consumption needed to maintain a comfortable temperature in the room. In our experiments, linear MPC had inconsistent performance due to its inability to compensate for the impacts of occupancy; it had difficulty with either saving energy as compared to two-position control [27] or maintaining temperature close to the desired value. This is related to a fundamental tradeoff in control systems between robustness to model uncertainty and performance due to model accuracy [10].

We implemented a learning-based MPC scheme [10] with the intuition that this could provide improved performance and reduce energy usage. Our experiments on the BRITE platform suggest that this is indeed the case. The energy improvements come from two features of the learning-based MPC. First, it can compensate for fluctuations in weather and occupancy through learning. For instance, the two-position control of the thermostat overcools the room when occupancy is low, and the heat anticipator in the thermostat does not adequately compensate. Second, it considers the penalty due to electrical consumption by the heat pump transient and tries to optimize the tradeoff between minimizing switching and AC on time.

A. Experimental Methodology on BRITE

BRITE is a living laboratory in which we cannot control weather and occupant behavior, and so we cannot make direct experimental comparisons between control methods. One potential solution is to run many experiments, but this is difficult due to the huge variability in weather and occupancy. Our approach is to run one control scheme on BRITE and simulate the others. This allows a comparison under identical weather and occupancy conditions, though with some error between the simulated and real energy consumptions because of modeling mismatch. To mitigate this, we alternate which method is simulated.

The results of two experiments are summarized in Table I. The first controlled BRITE with two-position control, and the second used learning-based MPC on BRITE; we do not include any experiments with linear MPC. The energy usage measured by BRITE and estimated by (1) are both provided for these experiments. These are compared to energy consumption

	Method	Switches	On Duration	Energy		Tracking Error	Temperature Variation	Average External Load
				Measured	Estimated			
Two-Position Control Experiment	Learning-Based MPC	94	6.0 hours		23.6kWh	0.75°C	0.13°C	11.0°C
	Linear MPC	96	7.9 hours		30.5kWh	0.62°C	0.30°C	11.0°C
	Two-Position	71	9.2 hours	32.6kWh	35.1kWh	0.61°C	0.20°C	11.0°C
Learning-Based MPC Experiment	Learning-Based MPC	81	3.3 hours	11.8kWh	13.3kWh	0.86°C	0.17°C	8.7°C
	Linear MPC	70	2.0 hours		8.6kWh	0.93°C	0.21°C	8.7°C
	Two-Position	38	9.2 hours		34.5kWh	0.55°C	0.19°C	8.7°C

TABLE I
SUMMARY OF EXPERIMENTAL AND SIMULATED ENERGY COMPARISONS ON BRITE

estimates, using (1), of simulations of other control schemes under identical weather and occupancy levels. The table lists the number of times the AC was turned on, the duty cycle of the AC, the tracking error as measured by the RMS error between the room temperature and the desired temperature $T_d = 22^\circ C$, and the variation in the room temperature as measured by its standard deviation. The external load corresponds to the average temperature increase over 15 minutes caused by the weather and occupancy.

B. Two-Position Control Experiment on BRITE

Over a 24 hour span beginning and ending on a weekday, we started running the two-position control of the thermostat on BRITE at 11PM. The experimentally measured temperature is shown in Fig. 5a. Using our models, we simulated the corresponding behavior of the learning-based MPC, which is shown in Fig. 5b. For our simulation, we used the stored weather forecasts, true weather temperature, and occupancy estimated using our model of two-position thermostat control. The learning-based MPC used an estimated 28% less energy than the two-position control. The PWM control actions corresponding to two-position control and learning-based MPC are shown in Fig. 5c and Fig. 5d, respectively. Moreover, Fig. 5e shows the change in temperature over 15 minutes corresponding to experimentally measured weather and occupancy (i.e., $k_w w[n] + q[n]$).

C. Learning-Based MPC Experiment on BRITE

We ran the learning-based MPC control on the BRITE platform over a time range that covered two weekdays, and started at roughly 1PM. The experimentally measured temperature is shown in Fig. 6a. Using our models, we simulated the corresponding behavior of the two-position control, which is shown in Fig. 6b. For our simulation, we used the true weather temperature and occupancy estimated using our model of the learning-based MPC. Our learning-based MPC approach on BRITE is estimated to reduce the energy consumption by 66%, when compared to the existing two-position control scheme. The control actions corresponding to the two-position and learning-based MPC are shown in Fig. 6c and Fig. 6d, respectively. The measured weather and occupancy for this experiment $k_w w[n] + q[n]$ is given in Fig. 6e.

D. Discussion of Results

Both comparisons show that significant energy is saved by the learning-based MPC scheme. It is useful to discuss what

features of our implementation and scheme contribute to this, because many of these principles may generalize to other HVAC systems. Broadly speaking, the improvements come about through the use of modeling and statistical techniques.

Identifying a discrete time version of a mathematical model taken from physics (3) helps to improve efficiency. There are complex dynamics in the heat pump, and the evaporator continues to briefly cool air after the heat pump is turned off [14]. The discrete time form of the model (3) accounts for this behavior by considering the AC behavior over a 15 minute span of time, rather than its instantaneous behavior.

Furthermore, identifying the parameters of the model allows us to be able to estimate occupancy through only temperature measurements, as in (12). These occupancy estimates are important because this feature of the system adds considerable variation in the temperature dynamics of the room. Whereas two-position control overcools the room when there is lower occupancy, learning-based MPC detects lower levels of occupancy and reduces the amount of cooling.

Lastly, the electrical energy characteristics of the heat pump are important to conserving energy. The transients of the heat pump effectively add a penalty, in terms of energy used, for switching too frequently. The learning-based MPC can make a tradeoff between how long the heat pump is turned on for and how often it switches, and it can dynamically adjust this tradeoff based on the estimated occupancy.

This tradeoff is actually very interesting, because it leads to counter-intuitive behaviors with the learning-based MPC. Examining the temperature of the learning-based MPC (i.e., Figs. 6a and 5b) shows that the bands within which the temperature is maintained actually vary over time. Generally speaking, when the outside temperature or occupancy are high, the learning-based MPC actually tightens the temperature bands. When the outside is cold or occupancy is low, the learning-based MPC widens the temperature bands.

These behaviors can be explained by thinking of the electrical behavior of the heat pump. When the temperature or occupancy is high, the AC needs to be turned on for a greater fraction of time. The steady state energy consumption is much higher than the transient energy consumption, and so the learning-based MPC does not penalize as much for frequently switching. In fact, it increases switching to prevent overcooling the room. In the opposite situation, the AC needs to do less total cooling. Here, the steady state energy consumption is smaller and so transient energy due to switching becomes important. The learning-based MPC reduces switching in these cases and allows for larger temperature variations.




Fig. 5. The AC was controlled by the two-position control of the thermostat, and the corresponding measured room temperature is shown in units of ($^{\circ}C$). A simulation of the learning-based MPC is given in ($^{\circ}C$). The two-position control uses $32.6kWh$ (estimated $35.1kWh$) of electrical energy, and the learning-based MPC is estimated to use $23.6kWh$. The PWM control generated by the two-position and learning-based MPC control are also shown. An AC state of 0 corresponds to the AC off, and AC state of 1 corresponds to the AC on. The external heating load over 15 minutes due to weather and occupancy $k_w w[n] + q[n]$ is given in ($^{\circ}C$).

VII. CONCLUSION

We have presented our BRITE platform, studied the transient and steady-state electrical characteristics of the heat pump in BRITE, identified a dynamical model of the system, explained the impact of occupants on the dynamics, and implemented a learning-based MPC scheme that estimates occupancy using only temperature measurements. Experiments show that learning-based MPC saves an estimated 30-70% of energy compared to two-position control. More sophisticated estimates of occupancy will likely yield further reductions.

One future direction is evaluating how the energy savings depend upon the outside temperature and occupancy levels;

the lower 28% savings occurred on a warmer day than the savings of 66%. It is not known how much of this is due to differences in weather versus simulation modeling errors. We are gathering more data to further evaluate these issues.

Another planned direction is the implementation of learning-based MPC on a larger testbed. We have studied a single-stage heat pump for a single room or small building; however, large HVAC systems for many rooms adds more challenges to the problem of saving energy [4]–[8]. Estimating and adjusting for occupancy, as well as the transient and steady-state electrical consumption of the HVAC equipment, will likely lead to real savings in energy.

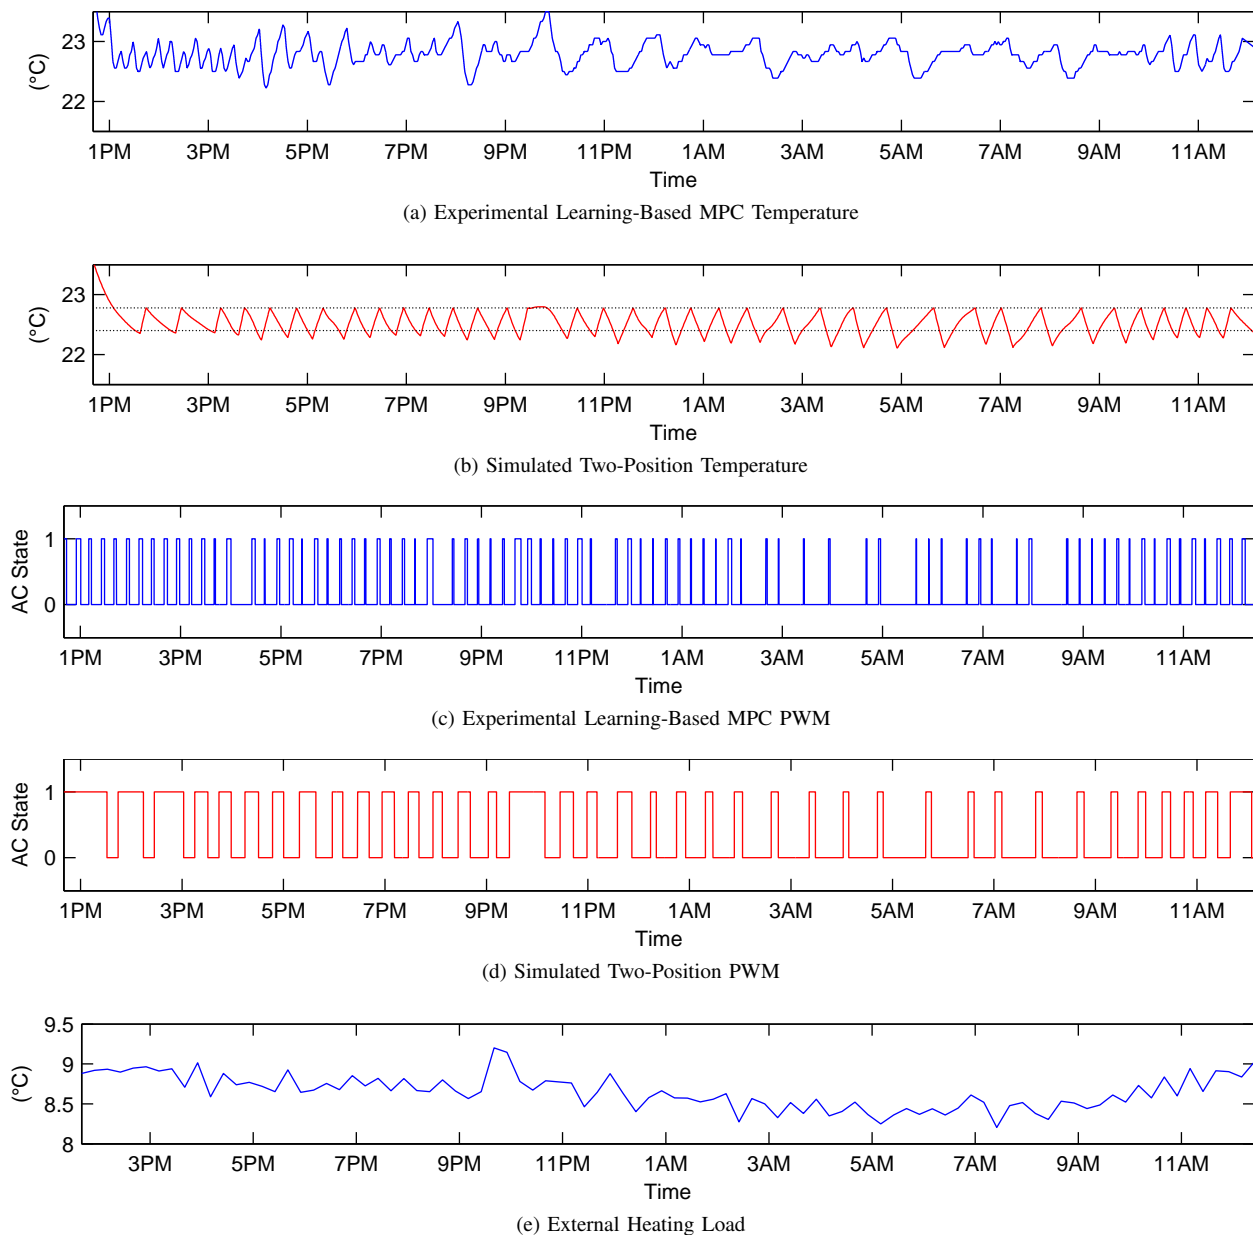


Fig. 6. The AC was controlled by the learning-based MPC, and the corresponding measured room temperature is shown in units of ($^{\circ}C$). A simulation of the two-position control is given in ($^{\circ}C$). The learning-based MPC uses $11.8kWh$ (estimated $13.3kWh$) of electrical energy, and the two-position control is estimated to use $34.5kWh$. The PWM control generated by the learning-based MPC and the two-position control are also shown. An AC state of 0 corresponds to the AC off, and AC state of 1 corresponds to the AC on. The change in temperature over 15 minutes corresponding to experimentally measured weather and occupancy $k_w w[n] + q[n]$ is provided in ($^{\circ}C$).

VIII. ACKNOWLEDGMENTS

We thank Albert Goto, Stephen Dawson-Haggerty, and Scott McNally for their networking and infrastructure support, as well as for useful discussions.

REFERENCES

- [1] *Buildings Energy Data Book*. U.S. DoE, 2009. [Online]. Available: <http://buildingsdatabook.eren.doe.gov/>
- [2] J. McQuade, "A system approach to high performance buildings," United Technologies Corporation, Tech. Rep., 2009.
- [3] S. Dawson-Haggerty, J. Ortiz, X. Jiang, J. Hsu, S. Shankar, and D. Culler, "Enabling green building applications," in *Sixth Workshop on Hot Topics in Embedded Networked Sensors*, 2010.
- [4] K. Deng, P. Barooah, P. Mehta, and S. Meyn, "Building thermal model reduction via aggregation of state," in *American Control Conference*, 2010, pp. 5118–5123.
- [5] F. Oldewurtel, A. Parisio, C. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth, "Energy efficient building climate control using stochastic model predictive control and weather predictions," in *American Control Conference*, 2010, pp. 5100–5105.
- [6] Y. Ma, G. Anderson, and F. Borrelli, "A distributed predictive control approach to building temperature regulation," in *American Control Conference*, 2011.
- [7] T. Nghiem and G. Pappas, "Receding-horizon supervisory control of green buildings," in *American Control Conference*, 2011.
- [8] G. Henze, C. Felsmann, and G. Knabe, "Evaluation of optimal control for active and passive building thermal storage," *International Journal of Thermal Sciences*, vol. 43, no. 2, pp. 173–183, 2004.
- [9] C. Liao and P. Barooah, "An integrated approach to occupancy modeling

- and estimation in commercial buildings,” in *American Control Conference*, 2010, pp. 3130–3135.
- [10] A. Aswani, H. Gonzales, S. Sastry, and C. Tomlin, “Provably safe and robust learning-based model predictive control,” submitted.
- [11] Honeywell, *Honeywell Engineering Manual of Automatic Control for Commercial Buildings*. Honeywell, 1998.
- [12] V. W. Goldschmidt, “Heat pumps: Basics, types, and performance characteristics,” *Annual Review of Energy*, vol. 9, no. 1, pp. 447–472, 1984.
- [13] ANSI/ASHRAE, “Standard 55, thermal environmental conditions for human occupancy,” 2008.
- [14] J. V. C. Vargas and J. A. R. Parise, “Simulation in transient regime of a heat pump with closed-loop and on-off control,” *International Journal of Refrigeration*, vol. 18, no. 4, pp. 235–243, 1995.
- [15] J. MacArthur, “Transient heat pump behaviour: a theoretical investigation,” *International Journal of Refrigeration*, vol. 7, no. 2, pp. 123–132, 1984.
- [16] S. Dawson-Haggerty, X. Jiang, G. Tolle, J. Ortiz, and D. Culler, “sMAP - a simple measurement and actuation profile for physical information,” in *Proceedings of the Eighth ACM Conference on Embedded Networked Sensor Systems (SenSys '10)*, November 2010.
- [17] R. Katz, D. Culler, S. Sanders, et al., “An information-centric energy infrastructure: The Berkeley view,” *Sustainable Computing: Informatics and Systems*, vol. 1, pp. 7–22, 2011.
- [18] X. Jaing, S. Dawson-Haggerty, P. Dutta, and D. Culler, “Design and implementation of a high-fidelity AC metering network,” in *The 8th ACM/IEEE International Conference on Information Processing in Sensor Networks Track on Sensor Platforms, Tools, and Design Methods (IPSN/SPOTS 09)*, April 2009.
- [19] P. Gill, W. Murray, and M. Saunders, “Snopt: An sqp algorithm for large-scale constrained optimization,” *SIAM Review*, vol. 47, no. 1, pp. 99–131, 2005.
- [20] M. Kvasnica, P. Grieder, and M. Baotić, “Multi-Parametric Toolbox (MPT),” 2004. [Online]. Available: <http://control.ee.ethz.ch/mpt/>
- [21] W. Guo and M. Zhou, “Technologies toward thermal comfort-based and energy-efficient hvac systems: A review,” in *Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics*, August 2009.
- [22] *Standards Interpretation - 29 CFR, Part 1910.1000*, OSHA Std.
- [23] C. Tomlin, J. Lygeros, and S. Sastry, “A game theoretic approach to controller design for hybrid systems,” *Proceedings of the IEEE*, vol. 88, no. 7, pp. 949–969, Jul. 2000.
- [24] J. Lygeros, C. Tomlin, and S. Sastry, “Controllers for reachability specifications for hybrid systems,” *Automatica*, vol. 35, no. 3, pp. 349–370, Mar. 1999.
- [25] W. Burke and D. Auslander, “Low-frequency pulse width modulation design for HVAC compressors,” in *The ASME 2009 International Design Engineering Technical Conferences*, August 2009.
- [26] E. Candès, M. Wakin, and S. Boyd, “Enhancing sparsity by reweighted ℓ_1 minimization,” *Journal of Fourier Analysis and Applications*, vol. 14, pp. 877–905, 2008.
- [27] A. Aswani, J. Taneja, N. Master, D. Culler, and C. Tomlin, “The berkeley retrofitting and inexpensive hvac testbed for energy efficiency,” submitted.
- [28] D. Crawley, C. Pedersen, L. Lawrie, and F. Winkelmann, “Energyplus: Energy simulation program,” *ASHRAE Journal*, vol. 42, pp. 49–56, 2000.
- [29] M. Duffy, M. Hiller, D. Bradley, W. Keilholz, and J. Thornton, “TRNSYS - features and functionality for building simulation 2009 conference,” in *IBPSA Conference*, 2009, pp. 1950–1954.
- [30] P. Bickel, C. Klaassen, Y. Ritov, and J. Wellner, *Efficient and adaptive estimation for semiparametric models*, ser. Johns Hopkins series in the mathematical sciences. Springer, 1998.
- [31] P. Robinson, “Root-n-consistent semiparametric regression,” *Econometrica*, vol. 56, no. 4, pp. 931–954, 1988.
- [32] S. S. Sastry and M. Bodson, *Adaptive Control: Stability, Convergence, and Robustness*. Prentice-Hall, 1989.
- [33] K. Åström and B. Wittenmark, *Adaptive control*. Addison-Wesley, 1995.
- [34] J. Xu and Y. Tan, *Linear and nonlinear iterative learning control*. Springer, 2003.
- [35] R. Tedrake, “LQR-trees: Feedback motion planning on sparse randomized trees,” in *Robotics: Science and Systems*, 2009, pp. 17–24.
- [36] P. Abbeel, A. Coates, and A. Ng, “Autonomous helicopter aerobatics through apprenticeship learning,” *International Journal of Robotics Research*, vol. 29, no. 13, pp. 1608–1639, 2010.
- [37] A. Aswani, P. Bickel, and C. Tomlin, “Statistics for sparse, high-dimensional, and nonparametric system identification,” in *International Conference on Robotics and Automation*, 2009.
- [38] ———, “Regression on manifolds: Estimation of the exterior derivative,” *Annals of Statistics*, vol. 39, no. 1, pp. 48–81, 2011.
- [39] L. Chisci, J. Rossiter, and G. Zappa, “Systems with persistent disturbances: predictive control with restricted constraints,” *Automatica*, vol. 37, pp. 1019–1028, 2001.
- [40] W. Langson, I. Chrysochoos, S. Raković, and D. Mayne, “Robust model predictive control using tubes,” *Automatica*, vol. 40, no. 1, pp. 125–133, 2004.
- [41] D. Limon, I. Alvarado, T. Alamo, and E. Camacho, “Robust tube-based MPC for tracking of constrained linear systems with additive disturbances,” *Journal of Process Control*, vol. 20, no. 3, pp. 248–260, 2010.
- [42] F. Borelli, A. Bemporad, and M. Morari, *Constrained Optimal Control and Predictive Control for linear and hybrid systems*, 2009, in preparation.
- [43] G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel, “Examples when nonlinear model predictive control is nonrobust,” *Automatica*, vol. 40, no. 10, pp. 1729–1738, 2004.