Identifying Models of HVAC Systems Using Semiparametric Regression

Anil Aswani, Neal Master, Jay Taneja, Virginia Smith, Andrew Krioukov, David Culler, and Claire Tomlin

***Abstract*— Heating, ventilation, and air-conditioning (HVAC) systems use a large amount of energy, and so they are an interesting area for efficiency improvements. The focus here is on the use of semiparametric regression to identify models, which are amenable to analysis and control system design, of HVAC systems. This paper briefly describes two testbeds that we have built on the Berkeley campus for modeling and efficient control of HVAC systems, and we use these testbeds as case studies for system identification. The main contribution of this work is that the use of semiparametric regression allows for the estimation of the heating load from occupancy, equipment, and solar heating using only temperature measurements. These estimates are important for building accurate models as well as designing efficient control schemes, and in our other work we have been able to achieve a reduction in energy consumption on a single room testbed using heating load estimation in conjunction with the learning-based model predictive control (LBMPC) technique. Furthermore, this framework is not restrictive to modeling nonlinear HVAC behavior, because we have been able to use this methodology to create hybrid system models that incorporate such nonlinearities.**

I. Introduction

Heating, ventilation, and air-conditioning (HVAC) systems are responsible for a major percentage of energy consumption in buildings, and as a result are a large constituent of overall energy usage [1], [2]. HVAC systems are important components of buildings because they regulate conditions related to occupant health such as carbon dioxide and humidity levels as well as occupant comfort such as temperature and airflow. With the recent interest in improving energy efficiency, HVAC is one area in which significant improvements are likely possible.

Research and development in improving HVAC efficiency can be roughly classified into two areas. One direction of work concerns the design of more efficient HVAC equipment as well as better architectural designs that require less regulation of environmental conditions, such as in [3]. The other direction of work concerns retrofitting HVAC through better configuration and new control schemes [4], [5], [6], [7], [8], [9]. Since buildings and HVAC equipment are slowly replaced [10], this approach is valuable because it saves energy with minimal infrastructure investments.

Having a mathematical model that describes the HVAC system can be instrumental in the retrofitting procedure. Most obviously, the design of typical control schemes greatly benefits from having a dynamical (e.g., ordinary differential equation (ODE) or difference equation (DE)) model. Modeling HVAC systems of buildings, individual rooms, and vehicles has a long history through simulation software such as EnergyPlus [11] and TRNSYS [12]. These models are very important for the design of new buildings and HVAC systems, as well as analysis and configuration of existing HVAC systems. However, the models generated by building simulation software are not amenable to control design because the models are high-dimensional from their inclusion of complex physical effects.

Identifying models that are amenable to controller design is challenging because of the complex physics [13], [14]. The modeling is further complicated by the fact that heating load due to occupants, equipment, and solar heating is difficult to capture. Some work in this area has been conducted [15], [16]. This paper describes one approach for modeling building and room HVAC systems that generates equations that can be used for control, while encapsulating the highly time-varying heating load from exogenous sources.

Our approach is to start with a macro-scale model of the physics (e.g., Newton's law of cooling) and then consider the heating load as a time-varying quantity. More specifically, we do not model heating from solar effects, equipment, or occupants; this is not restrictive, because this framework can be extended to the case in which these are modeled. The benefit of not modeling these effects is a simpler dynamical system from which we can design the control. From the perspective of energy efficiency, experiments that we have conducted on a simple testbed indicate that reductions in energy usage can be generated with this approach [17].

The statistical technique that enables this modeling approach is semiparametric regression. The advantage of this is that we do not need to explicitly model the heating load effects; rather those effects have a non-parametric representation. In contrast, the physical effects comprise the parametric portion of the model. This is an established technique, but its application to HVAC modeling provides an important benefit: We can estimate the heating load from only temperature measurements, which are already available from the thermostat.

This paper begins by describing two testbeds that we have built on the Berkeley campus for studying techniques

This material is based upon work supported by the National Science Foundation under Grants CNS-0931843 (CPS-ActionWebs) and CNS-0932209 (CPS-LoCal) and 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.

A. Aswani, N. Master, J. Taneja, A. Krioukov, 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,krioukov,culler}@cs.berkeley.edu; neal.m.master@berkeley.edu

V. Smith is with Department of Computer Science, University of Virginia, Charlottesville, VA 22904 USA e-mail: vs2z@virginia.edu

of improving building automation. Next, we provide an introduction to semiparametric regression. This technique is applied to construct a model of the air-conditioning (AC) system in a single room testbed; we then analyze the ability of this technique to estimate heating load. Results from the application of the single room model to do control [17] using learning-based model predictive control (LBMPC) [18] are described. Then, we use semiparametric regression to identify a model for the building-wide testbed. We provide a comparison of the results from these two case studies, and then conclude by describing future plans.

II. Testbeds on Berkeley Campus

We have built and operated two testbeds for studying the efficiency of HVAC on the Berkeley campus. The first is named BRITE (for the Berkeley Retrofitted and Inexpensive HVAC Testbed for Energy Efficiency), which consists of a single room with an electric air conditioning (AC) unit. The other testbed, called BRITE-S, is a building-wide system with variable air volume (VAV) boxes; the S in the name stands for Sutardja Dai Hall – the building in which the testbed resides. Both platforms utilize technology developed under the LoCal project, which 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 [19].

A. BRITE Testbed

The BRITE platform allows actuation of a single-stage electric AC, which cools a student computing laboratory on the ground floor of a large engineering building; this type of AC is commonly found in homes. A CAD drawing of the room is shown in Fig. 1a; the room is 640 square feet and contains 16 desktop computer workstations along with two laser printers. Occupancy of the room peaks at over 20 occupants, and significantly varies depending on the time of day as well as the day of the semester. The room is equipped with a networked thermostat, whose temperature measurements are stored in a database; the thermostat can also receive commands for external actuation of the AC.

B. BRITE-S Testbed

Sutardja Dai Hall is a modern, 141,000-square-foot building that is divided between a four-floor nanofabrication laboratory (NanoLab) and seven floors of general space (including office space, classrooms, and a coffee shop with dining area). An example of a floor of general space, primarily offices and cubicles, is shown in Fig. 1b. The building automation equipment can be measured and actuated through a BACnet protocol interface.

The HVAC system in this building consists of a 650-ton chiller that cools water, over which air is driven by large supply fans in air-handler units (AHUs) to distribute to VAV boxes throughout the building, governing the airflow of each building of 130 building zones. Though the NanoLab and the general space have their own sets of AHUs, they share a common chiller. Because the NanoLab must operate within tight environmental tolerances, our control design can only modify the operation of the general space AHUs and VAV boxes, with no modification of chiller or chilled water settings.

III. Introduction to semiparametric Regression

In this section, we provide a brief introduction to semiparametric regression of partially linear models, a strategy used in both BRITE and BRITE-S. We specialize the presentation to the case of identification of dynamical models. The interested reader can refer to [20], [21] for more information about the general theory. The model assumes a linear dependence between the input and the change in temperature over time, and this is a reasonable assumption for the single room AC. It is also reasonable for a VAV box under the assumption that the supply air temperature is fixed. This is not restrictive, though, because we could identify a hybrid system model that consists of a different linear model for each supply air temperature.

The setup is as follows: Suppose there are $N+1$ measurements of temperature $T[n] \in \mathbb{R}^p$, where $n = 0, \dots, N$ indexes time and p is the number of different points in the building where temperature is measured. Furthermore, suppose that there are also $N+1$ measurements of weather $w[n]$ and inputs to the HVAC system $u_i[n] \in \mathbb{R}^m$, where m is the number of inputs for the i-th temperature.

Without loss of generality, the dynamic evolution of the i-th temperature is given by the partially linear model

$$T_i[n+1] = a_i' T[n] + b_i' u_i[n] + c_i' w[n] + q_i[n] + \epsilon_i[n], \quad (1)$$

where $q_i[n]$ is an unknown function that includes the heating load from occupants, equipment, and solar heating; and $\epsilon_i[n]$ is assumed to be independent and identically distributed zero mean noise with constant variance that is also conditionally independent of T_i, u_i, w, and n. The particular form of (1) is inspired by Newton's law of cooling.

It is important to note that the sparsity pattern of a_i can often be pre-determined before the regression procedure is carried out. This reduces the number of parameters that need to be identified and protects against overfilling. The sparsity pattern can be selected, for instance, by the physical proximity of different temperature measurements. An example of this is given when the model for BRITE-S is identified.

IV. Identification of Coefficients

We begin by explaining the intuition behind how the coefficients of the linear portion of the model are computed. First, define the conditional expectations: $\hat{T}[n] = \mathbb{E}\Big[T[n]\Big|n\Big]$, $\hat{u}_i[n] = \mathbb{E}\Big[u_i[n]\Big|n\Big]$, $\hat{w}[n] = \mathbb{E}\Big[w[n]\Big|n\Big]$. Using these, we take the conditional expectation of both sides of (1) to get

$$\hat{T}_i[n+1] = a_i'\hat{T}[n] + b_i'\hat{u}_i[n] + c_i'\hat{w}[n] + \mathbb{E}\Big[q_i[n]\Big|n\Big] + \mathbb{E}\Big[\epsilon_i[n]\Big|n\Big]. \quad (2)$$

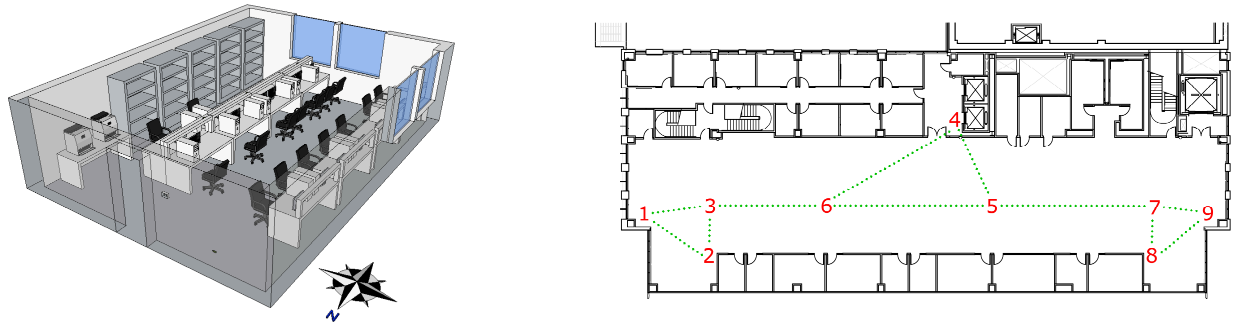

(a) Room in BRITE

(b) Single Floor of BRITE-S

Fig. 1. A CAD model of the room used in the BRITE platform and a schematic floor plan of a single floor (out of seven total floors) in the BRITE-S platform are shown. The numbers mark the location of VAV boxes that cool the public areas of the floor, and the dashed lines indicate a set of nearest neighbors for these VAV boxes.

However, we have that $\mathbb{E}\Big[q_i[n]\Big|n\Big] = q_i[n]$ and $\mathbb{E}\Big[\epsilon_i[n]\Big|n\Big] = 0$. And so subtracting (2) from (1) gives

$$T_i[n+1] - \hat{T}_i[n+1] = a_i'(T[n] - \hat{T}[n]) + b_i'(u_i[n] - \hat{u}_i[n]) + c_i'(w[n] - \hat{w}[n]) + \epsilon_i[n]. \quad (3)$$

The significance of (3) is that the relationship no longer contains the $q_i[n]$ term, and so the coefficients a_i, b_i, c_i can be computed using ordinary least squares (OLS): $(\hat{a}_i, \hat{b}_i, \hat{c}_i) = \arg\min L(a_i, b_i, c_i)$, where

$$L(a_i, b_i, c_i) = \|T_i[n+1] - \hat{T}_i[n+1] - a_i'(T[n] - \hat{T}[n]) - b_i'(u_i[n] - \hat{u}_i[n]) - c_i'(w[n] - \hat{w}[n])\|^2. \quad (4)$$

Using the estimates $\hat{a}_i, \hat{b}_i, \hat{c}_i$; the $q[n]$ term can be estimated through $\hat{q}[n] = \hat{T}_i[n+1] - \hat{a}_i'\hat{T}[n] - \hat{b}_i'\hat{u}_i[n] - \hat{c}_i'\hat{w}[n]$.

There is one last point concerning how $\hat{T}, \hat{u}, \hat{w}$ are computed. Because these quantities are indexed by time, it turns out that this is equivalent to smoothing over time. There are a variety of methods for doing so. We make use of the Nadaraya-Watson estimator.

V. Bayesian Identification of Coefficients

The procedure outlined in the previous section assumes identifiability of the coefficients. Such an assumption may not be valid if the system is not sufficiently excited. This situation unfortunately occurs quite often in HVAC systems. For the BRITE testbed, there is only one input, and so we could randomly vary this to provide the desired excitation. On the other hand, the BRITE-S testbed involves a building-wide system with a large number of inputs, requiring a significant amount of time to generate enough temperature data from randomly varied inputs.

In light of this difficulty, we can use a Bayesian procedure for the semiparametric regression of the partially linear model. Assuming a prior distribution on the coefficients: $a_i \sim \mathcal{N}(\overline{a}_i, A_i)$, $b_i \sim \mathcal{N}(\overline{b}_i, B_i)$, $c_i \sim \mathcal{N}(\overline{c}_i, C_i)$, where the notation $\mathcal{N}(\mu, \Sigma)$ indicates a set of jointly Gaussian random variables with a vector of means μ and covariance Σ. The intuition is that we use prior knowledge about the HVAC system to compensate for the fact that the measured system is not sufficiently excited.

For these given prior distributions, the coefficients of the linear portion of the model are given by

$$(\hat{a}_i, \hat{b}_i, \hat{c}_i) = \arg\min L(a_i, b_i, c_i) + \|A_i^{-1/2}(a_i - \overline{a}_i)\|^2 + \|B_i^{-1/2}(b_i - \overline{b}_i)\|^2 + \|C_i^{-1/2}(c_i - \overline{c}_i)\|^2, \quad (5)$$

where the additional terms (cf., (4)) are a Tikhonov regularization that corresponds to the prior distributions defined earlier. The remainder of the semiparametric regression procedure is unmodified.

VI. Identification and Control for BRITE

For the BRITE platform, we identified a model and implemented a control scheme using LBMPC that leads to a 30-70% reduction in energy usage while maintaining occupant comfort. We use a discrete time model with a sampling period of $T_s = 15$ minutes, and an input $u[n] \in [0, 0.5]$ indicates the fraction T_s over which the AC is kept on.

A. Modeling the AC and Room in BRITE

Data were collected on a weekday over a timespan in which students are and are not using the room. To ensure sufficient excitation, a random input with uniform distribution over $[0, 0.5]$ was applied at each discrete time step. Since this process would be done once, it may be acceptable for most implementations to allow the room temperature to be unregulated for this one day. We chose to do this process over only one day because we did not want to let the temperature be uncontrolled for a very long period of time.

We used semiparametric regression to identify the model

$$T[n+1] = 0.64 \cdot T[n] - 2.64 \cdot u[n] + 0.10 \cdot w[n] + \hat{q}[n], \quad (6)$$

where $\hat{q}[n]$ is shown in Fig. 2a. The experimental room temperature is the solid line in Fig.2b. Similarly, the temperature simulated by the model (6) is the dashed line shown in Fig.2b, 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 well to the measured temperature data.

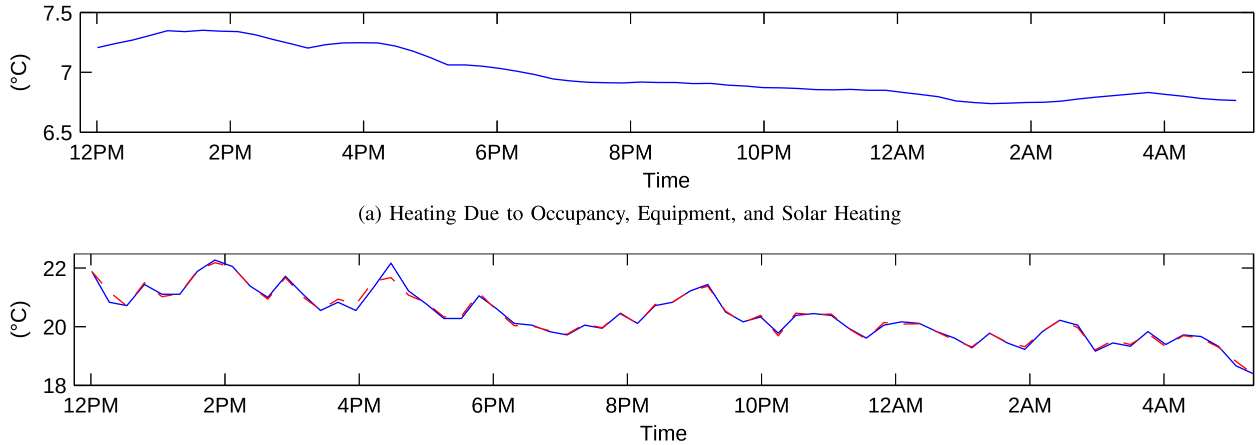

(a) Heating Due to Occupancy, Equipment, and Solar Heating

(b) Experimental (Solid) and Simulated (Dashed) Temperature

Fig. 2. semiparametric regression was used to identify a model for BRITE [17]. The first plot shows the increase in temperature in ($^\circ C$) over a span of 15 minutes due to occupancy, equipment, and solar heating. In the second plot, the solid line shows the measured room temperature ($^\circ C$), and the dashed line shows a simulation of our temperature model ($^\circ C$).

B. Impact of Heating Load

The model (6) shows that occupancy, equipment, and solar heating play a significant role in the temperature dynamics of the room, which is in line with [8]. The average of $\hat{q}[n]$ is $6.98^\circ C$, and it varies on both long and short time horizons by up to $0.61^\circ C$ depending on time of day.

The room is a computer laboratory used by students at their own convenience. The heat input $\hat{q}[n]$ increases from lunchtime and peaks at 1PM, while the outside temperature peaks at 2PM. It is fairly constant from 8PM to 5AM, which is typically when there are few or no students in the room.

The time-varying heating load makes the design of efficient control schemes difficult because the nominal model can be inaccurate by $0.61^\circ C$ (in our case). Standard MPC requires accurate models to provide high performance, and so instead we use LBMPC to do control on BRITE.

C. LBMPC on BRITE

Estimating current heating load requires combining models of human behavior with sensors [22], [15], and the BRITE testbed faces additional challenges because of the highly varying occupancy of the laboratory it manages. Some of the variations will likely be periodic in nature, while others are 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 heating load $q[n]$, we focus our efforts on estimating this $q[n]$ directly from the temperature measurements and our model (6). We use the estimate

$$\hat{q}[m+i] = T[m] - \Big(0.64 \cdot T[m-1] - 2.64 \cdot u[m-1] + 0.10 \cdot w[m-1]\Big), \quad (7)$$

for $i = 0, \dots, N-1$. The intuition is that heating load $q[n]$ is the discrepancy between (a) the temperature that the linear model without the average-heating-load term 6.98 predicts at the next time step and (b) the actual temperature. More accurate estimates of $q[n]$ taking into account specific models will only improve the energy efficiency of the BRITE testbed.

Application of LBMPC to BRITE resulted in an estimated 30-70% reduction of energy use as compared to the two-position control of the thermostat; linear MPC had inconsistent performance. Simulation studies from other work on the application of MPC to models typically show a 10-40% reduction, and the discrepancy between our experimental results and these simulations is that the HVAC equipment we experimented on has additional modes of energy wastage that is not found in the types of systems considered in other work. Details on the experiments and comparison method can be found in [17], though we summarize them here. Since it is critical to have identical conditions for each experiment, one technique was simulated while the other was experimentally run. In one experiment, the two-position control of the thermostat consumed 32.6 kWh of electricity over a day, while the LBMPC scheme was simulated and consumed 23.6 kWh. In another experiment, the LBMPC scheme consumed 11.8 kWh over a day, while the two-position control was simulated and consumed 34.5 kWh.

More importantly, the experiments show that the estimates of the heating load behave reasonably. Fig. 3 shows the heating load as estimated using (7). The estimates obviously have some noise with standard deviation of roughly $0.2\ ^\circ C$ because of modeling error. However, the more salient point to note is that the changes in the estimated heating load occur on the timescale of hours. This is what would be expected based on the sources of the load.

VII. Identification for BRITE-S

In the BRITE-S testbed, each VAV box regulates the temperature of a zone by (a) allowing varying amounts of

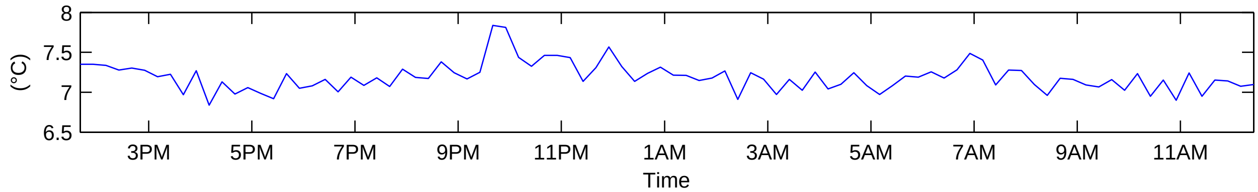

Fig. 3. As the LBMPC algorithm is run on the actual BRITE testbed, it computes an estimate of the heating load ($^\circ C$) using (7).

cold air into the room, and (b) heating the incoming cold air by operating a reheat coil that resides at the VAV box. The reason that reheating is often required is that each zone has a bound on the minimum amount of airflow in order to meet air quality regulations [23]. This is problematic because the temperature set point of the cold air is picked to be able to cool the hottest room in the building, and this can overcool other rooms because of their minimum air flow. So the VAV box can reheat the air to prevent overcooling.

There are two possible choices on what we take as the input $u_i[n]$ for the model. The first choice is that the control inputs for each VAV box are air flow rate $u_{i,1}[n]$ in units (m^3/s) and amount of reheating of air $u_{i,2}[n]$ in units of (percent valve opening). The second choice is that the control input is desired zone temperature $T_{d,i}[n]$ in ($^\circ C$), and air flow rate and reheat quantity are determined by the VAV boxes' onboard controller. Note that we do not take the set point of the supply air to be an input, as in [8]. This is not restrictive, because we can identify a hybrid system [24], [25] composed of a set of linear models for each supply air set point. Note that this model does not assume that the temperature of the air leaving a VAV box is fixed; in fact, the VAV box controller can continuously vary the temperature it provides to each zone. This behavior is modeled by our framework.

We have built a model using the second choice by first identifying a model with u_i as inputs and then modeling the onboard control of the VAV boxes. The onboard control at a sampling interval of 15 minutes approximately looks like a piecewise linear gain controller. For instance, if $e_i[n] = T_i[n] - T_{d,i}[n]$ and $k_p > 0$ is the proportional gain, then

$$u_{i,1}[n] = \begin{cases} 0, & \text{if } e[n] < 0 \\ k_p \cdot e[n], & \text{otherwise} \end{cases}. \tag{8}$$

Because we know the physical proximity of different zones, we can *a priori* define a sparsity structure for the a_i coefficients. On the floor of BRITE-S shown in Fig. 1b, the green lines indicate non-zero elements. For example, we model zone 1 so that it depends only upon its own temperature as well as those of zones 2 and 3. Mathematically, we would represent this as a_1 having the sparsity pattern

$$a_1 = \begin{bmatrix} \cdot & \cdot & \cdot & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \tag{9}$$

The main challenge in identifying a model for BRITE-S is that the system is not sufficiently excited. It is difficult to provide random control inputs, as we did for BRITE, because strict regulatory requirements must be satisfied during building operation. To overcome this, we used Bayesian semiparametric regression. The means of the prior distribution were chosen using physical intuition. For instance, the prior mean of the coefficient for $u_{i,1}[n]$ was selected to be $1/30^\circ C \cdot s/m^3$, because the VAV box can cool the zone by roughly $1^\circ C$ every 15 minutes when $30m^3/s$ of cold air is fed into the zone. The covariance of the prior distribution was chosen subjectively to ensure that the identified coefficients of the model had physically reasonable values.

Using data spanning one week, we used Bayesian semiparametric regression to identify a model. The heating load for zone 1 of the floor is shown in Fig. 4a. A comparison of the measured (solid line) and simulated (dashed line) zone 1 temperature is in Fig.4b. The initial condition for the simulation was taken from the experimental measurements, and it was simulated taking the input as desired temperatures $T_{d,i}$ that were applied to the BRITE-S system. The root-mean-squared (RMS) error of the simulation is $0.27\ ^\circ C$, $0.17\ ^\circ C$, $0.23\ ^\circ C$, $0.21\ ^\circ C$, $0.22\ ^\circ C$, $0.18\ ^\circ C$, $0.22\ ^\circ C$, $0.24\ ^\circ C$, $0.18\ ^\circ C$; for zones 1 to 9, respectively.

The plots related to zone 1 show several interesting features. The estimated heating load tends to peak in the middle of the day, and the heating load decreases significantly over the weekend. This supports the notion that the contribution from occupants is also being estimated. The plot of temperature also shows that the model provides a reasonable fit to the measured temperature data.

VIII. Comparison of Case Studies

The two testbeds are significantly different. BRITE is a single room with a single use (i.e., computer lab) that is cooled by an electrical heat pump AC. On the other hand, BRITE-S consists of many rooms with different uses and a building-wide HVAC system. Because of these physical differences, these testbeds have distinct models. For BRITE, the AC is either on or off, and this is modeled by assuming the control is pulse-width modulated with cycle length of 15 minutes. In BRITE-S, the controls are continuous variables that are controlled in real time.

Yet despite these large differences, the two testbeds can be described by a similar mathematical equation (1). (Albeit, BRITE-S would be fully represented by a series of these models.) And more importantly, semiparametric regression is able to identify models that can fit the measured temperature data well. The heating load estimated using this technique matches the use of the physical space, and it provides a technique for handling changes in occupancy, equipment use, and solar heating without the need to add additional sensors.

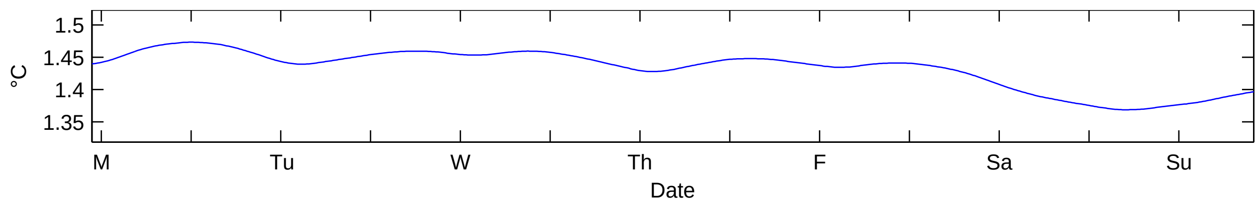

(a) Heating Due to Occupancy, Equipment, and Solar Heating

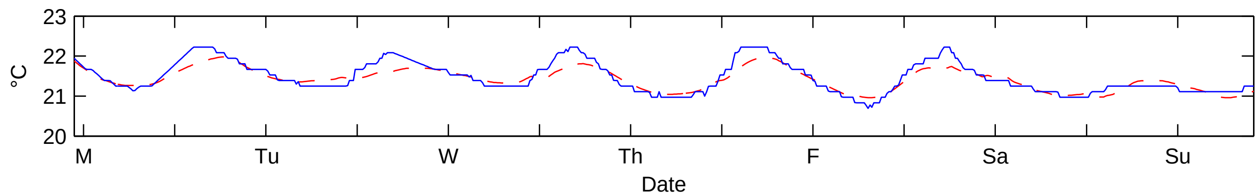

(b) Experimental (Solid) and Simulated (Dashed) Temperature

Fig. 4. semiparametric regression was used to identify a model for one floor on BRITE-S, and this figure shows plots for zone 1 of the floor shown in Fig. 1b. The first plot shows the increase in temperature in ($^\circ C$) over a span of 15 minutes due to occupancy, equipment, and solar heating. In the second plot, the solid line shows the measured temperature ($^\circ C$), and the dashed line shows a simulation of our temperature model ($^\circ C$) for this zone.

In the case of BRITE, we were able to reduce electricity usage by using this approach.

IX. Conclusion

We have presented the BRITE and BRITE-S testbeds, and provided an introduction to semiparametric regression. The main contribution of using this technique is that heating load from occupancy, equipment, and solar heating can be estimated from only temperature measurements. A model for BRITE was used along with LBMPC to provide reductions in energy consumption. Our current work involves using our model of BRITE-S to similarly use LBMPC and reduce energy consumption. We are working on an implementation of LBMPC with a hybrid system model of the HVAC system in our building-wide testbed.

X. Acknowledgments

We thank Domenico Caramagno, Stephen Dawson-Haggerty, Albert Goto, and Scott McNally.

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. Schiavon, K. H. Lee, F. Bauman, and T. Webster, "Influence of raised floor on zone design cooling load in commercial buildings," *Energy and Buildings*, vol. 42, no. 8, pp. 1182–1191, 2010.

[4] K. Deng, P. Barooah, P. Mehta, and S. Meyn, "Building thermal model reduction via aggregation of state," in *ACC*, 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 *ACC*, 2010, pp. 5100–5105.

[6] M. S. Elliott, C. Estrada, and B. P. Rasmussen, "Cascaded superheat control with a multiple evaporator refrigeration system," in *ACC*, 2011, pp. 2065–2070.

[7] V. Chandan and A. G. Alleyne, "Optimal control architecture selection for thermal control of buildings," in *ACC*, 2011, pp. 2071–2076.

[8] Y. Ma, G. Anderson, and F. Borrelli, "A distributed predictive control approach to building temperature regulation," in *ACC*, 2011, pp. 2089–2094.

[9] T. Nghiem and G. Pappas, "Receeding-horizon supervisory control of green buildings," in *American Control Conference*, 2011.

[10] 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.

[11] D. Crawley, C. Pedersen, L. Lawrie, and F. Winkelmann, "Energyplus: Energy simulation program," *ASHRAE Journal*, vol. 42, pp. 49–56, 2000.

[12] 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.

[13] S. Goyal and P. Barooah, "A method for model-reduction of nonlinear building thermal dynamics," in *ACC*, 2011, pp. 2077–2082.

[14] S. Ahuja, A. Surana, and E. Cliff, "Reduced-order models for control of stratified flows in buildings," in *ACC*, 2011, pp. 2083–2088.

[15] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng, "Occupancy-driven energy management for smart building automation," in *Proc. of 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, ser. BuildSys, 2010, pp. 1–6.

[16] C. Liao and P. Barooah, "A novel stochastic agent-based model of building occupancy," in *ACC*, 2011, pp. 2095–2100.

[17] A. Aswani, N. Master, J. Taneja, D. Culler, and C. Tomlin, "Reducing transient and steady state electricity consumption in hvac using learning-based model-predictive control," *Proceedings of the IEEE*, no. 99, pp. 1–14, 2011.

[18] A. Aswani, H. Gonzales, S. Sastry, and C. Tomlin, "Provably safe and robust learning-based model predictive control," arXiv:1107.2487v1. [Online]. Available: http://arxiv.org/abs/1107.2487v1

[19] 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.

[20] P. Robinson, "Root-n-consistent semiparametric regression," *Econometrica*, vol. 56, no. 4, pp. 931–954, 1988.

[21] 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.

[22] C. Liao and P. Barooah, "An integrated approach to occupancy modeling and esimation in commerical buildings," in *ACC*, 2010, pp. 3130–3135.

[23] *2008 Building Energy Efficiency Standards for Residential and Nonresidential Buildings, Section 121*, California Energy Commission Std. [Online]. Available: http://www.energy.ca.gov/2008publications/CEC-400-2008-001/CEC-400-2008-001-CMF.PDF

[24] 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.

[25] J. Lygeros, C. Tomlin, and S. Sastry, "Controllers for reachability specifications for hybrid systems," *Automatica*, vol. 35, no. 3, pp. 349–370, Mar. 1999.