Title

Piecewise Linear Thermal Model and Recursive Parameter Estimation of a Residential Heating System

Authors

Burger, Eric M.
Perez, Hector E.
Moura, Scott J.

Publication Date

2016-01-28

Piecewise Linear Thermal Model and Recursive Parameter Estimation of a Residential Heating System

Eric M. Burger, Hector E. Perez, Scott J. Moura

Energy, Control, and Applications Lab, University of California, Berkeley

Abstract

Model predictive control (MPC) strategies show great potential for improving the performance and energy efficiency of building heating, ventilation, and air-conditioning (HVAC) systems. A challenge in the deployment of such predictive thermostatic control systems is the need to learn accurate models for the thermal characteristics of individual buildings. This necessitates the development of online and data-driven methods for system identification. In this paper, we propose a piecewise linear thermal model of a building. To learn the model, we present a Kalman filter based approach for estimating the parameters. Finally, we fit the piecewise linear model to data collected from a residential building with a forced-air heating and ventilation system and validate the accuracy of the trained model.

Keywords: building thermal model, heating and air-conditioning, model predictive control, piecewise linear model

1. Introduction

Heating, ventilation, and air-conditioning (HVAC) account for 43% of commercial and 54% of residential energy consumption [1]. Space heating alone accounts for 45% of all residential energy use. HVAC systems are an integral part of buildings responsible for regulating temperature, humidity, carbon dioxide, and airflow, conditions which directly impact occupant health and comfort. Estimates suggest that component upgrades and advanced HVAC control systems could reduce building energy usage by up to 30% [2]. Such intelligent systems can improve the efficiency of building operations, better regulate indoor conditions to improve air quality and occupant comfort, and enable buildings to participate in demand response services to improve power grid stability and reduce energy related carbon emissions [3].

*Corresponding author: Eric M. Burger, Email: ericburger@berkeley.edu

**Affiliation Address: Energy, Control, and Applications Lab (eCAL), 611 Davis Hall, Department of Civil and Environmental Engineering, University of California, Berkeley, Berkeley, CA 94720, USA.

To effectively control the operation of an HVAC system, it is essential that a model predictive controller incorporate an accurate mathematical representation of a building's thermal dynamics. The processes that determine the evolution of temperatures within a building are complex and uncertain. A reliable model improves the ability of a controller to forecast conditions and meet efficiency and comfort objectives. Simulation software, such as EnergyPlus and TRNSYS, is capable of high fidelity modelling of building HVAC systems. These mathematical models play a crucial role in the architectural and mechanical design of new buildings, however, due to high dimensionality and computational complexity, are not suitable for incorporation into HVAC control systems [4].

The American Society of Heating, Refrigeration, and Air-Conditioning Engineers (ASHRAE) handbook [5] describes how to determine the thermal resistance values of a building surface given it materials and construction type. However, for existing buildings, details about the materials in and construction of walls and windows may be difficult to obtain or non-existent [6]. Additionally, modifications to the building or changes brought about by time and use (e.g. cracks in windows or walls) further diminish the potential for characterizing a building based on design or construction information.

Therefore, an ideal control-oriented model would capture the predominant dynamics and disturbance patterns within a building, enable accurate forecasting, adapt to future changes in building use, provide a model structure suitable for optimization, and be amenable to real-time data-driven model identification methods. For these reasons, low order RC models are widely employed for control-oriented thermal building models [6][7][8]. Such models trade complexity and accuracy for simplicity and efficiency.

In this paper, we present a piecewise linear RC model for the thermostatic control of buildings and a recursive Kalman filter method for parameter estimation. The piecewise model structure enables the approximate identification of unmodeled dynamics, in particular higher-order dynamics and time delays related to changes in the mechanical state of the system. By employing a recursive parameter estimation technique, we are able to perform online data-driven learning of the model.

We do not model heating from solar gain, building occupants, or equipment. This does not restrict the applicability of this work because the model structure can be extended for such cases. By estimating these effects with a single time-varying gain, we produce a simpler model better suited for predictive control.

This paper is organized as follows. Section 2 presents our piecewise thermal model and Section 3 overviews the parameter estimation problem. Section 4 provides background for the Kalman filter (KF) and Section 5 formulates a filter-based method for recursive parameter estimation of the piecewise thermal model. Section 6 provides numerical examples of our proposed model and algorithm for the parameter estimation of an apartment with a forced-air heating and ventilation system. Finally, Section 7 summarizes key results.

2. Building Thermal Model

2.1. Linear Thermal Model

In this paper, we focus on the modeling of an apartment with a forced-air heating system. We employ a first order RC model to represent the predominant thermal characteristics of the conditioned space, specifically the heat transfer between the interior and the environment and power delivered by the mechanical system. Solar gain and radiative heat transfer from the ambient are not considered. The capacitive elements within the apartment (e.g. air, walls, furniture) are lumped into a single capacitor. Likewise, the resistance values of exterior surfaces are aggregated into a single resistance. Therefore, the change in temperature within the apartment can be represented by the continuous time state equation [9][10][11]

$$\dot{T}^t = \frac{T^t_\infty - T^t}{RC} + \frac{Pm^t}{C} \tag{1}$$

where $T^t \in \mathbf{R}$, $T^t_\infty \in \mathbf{R}$, and $m^t \in \{0,1\}$ are the indoor air temperature (state, °C), outdoor air temperature (disturbance input, °C), and heater state (control input, On/Off), respectively. The parameters R (°C$/kW$), C ($kJ/$°C), and P (kW) represent the thermal resistance, thermal capacitance, and forced-air heater power, respectively.

The model can be expressed in the state-space form

$$\dot{T}^t = A_c T^t + B_c u^t \tag{2}$$

where

$$\begin{aligned} A_c &= \begin{bmatrix} \frac{-1}{RC} \end{bmatrix} \\ B_c &= \begin{bmatrix} \frac{1}{RC} & \frac{P}{C} \end{bmatrix} \\ u^t &= \begin{bmatrix} T^t_\infty \\ m^t \end{bmatrix} \end{aligned} \tag{3}$$

Assuming a zero-order hold on the input u, the model can be discretized using the transforms

$$\begin{aligned} A_d &= e^{A_c \Delta t} \\ B_d &= A_c^{-1}(A_d - I)B_c \end{aligned} \tag{4}$$

where Δt defines the length in hours between each time step. In this paper, we define this as $\Delta t = 1/60$ (hours). Therefore, the state-space model becomes

$$T^{k+1} = A_d T^k + B_d u^k \tag{5}$$

where

$$\begin{aligned} A_d &= \begin{bmatrix} e^{-\frac{\Delta t}{RC}} \end{bmatrix} \\ B_d &= \begin{bmatrix} (1 - e^{-\frac{\Delta t}{RC}}) & (1 - e^{-\frac{\Delta t}{RC}})RP \end{bmatrix} \\ u^k &= \begin{bmatrix} T^k_\infty \\ m^k \end{bmatrix} \end{aligned} \tag{6}$$

and $k = 1, 2, \ldots, n$ denotes the integer-valued time step.

Finally, we can express the discrete time model in a form that is linear in both the states and the parameters.

$$T^{k+1} = \theta_a T^k + (1-\theta_a) T^k_\infty + \theta_b m^k + \theta_c \qquad (7)$$

where $T^k \in \mathbf{R}$, $T^k_\infty \in \mathbf{R}$, and $m^k \in \{0,1\}$ are the indoor air temperature (state, °C), outdoor air temperature (disturbance input, °C), and heater state (control input, On/Off), respectively.

The parameter θ_a corresponds to the thermal characteristics of the conditioned space as defined by $\theta_a = \exp(-\Delta t/RC)$, θ_b to the energy transfer due to the systems mechanical state as defined by $\theta_b = (1 - \exp(-\Delta t/RC))RP$, and θ_c to an additive process accounting for energy gain or loss not directly modeled.

As noted in [12][10], the discrete time model implicitly assumes that all changes in mechanical state occur on the time steps of the simulation. In this paper, we assume that this behavior reflects the programming of the systems being modeled. In other words, we assume that the thermostat has a sampling frequency of $1/(3600\Delta t)$ Hz or once per minute.

2.2. Piecewise Linear Thermal Model

The linear discrete time model (7) is capable of representing the predominant thermal dynamics within a conditioned space. Unfortunately, because it does not capture any higher-order dynamics or time delays related to changes in the mechanical state of the system, the model is fairly inaccurate in practice. Research into higher-order models, in particular multi-zone network models and the modeling of walls as 2R-1C or 3R-2C elements, have shown potential for producing higher fidelity building models [6][7][8]. However, this comes at the cost of increasing the model complexity and the need for temperature sensing (in particular, within interior and exterior walls).

In this paper, we present a piecewise linear model capable of approximating dynamics related to changes in the mechanical state of the system. Our piecewise modelling approach is related to linear parameter-varying (LPV) systems which employ a linear model whose parameters change according to a time-varying state. This parameter dependency enables LPV systems to approximate nonlinear dynamics.

In our piecewise thermal model, the number of time steps since the system turned on or off serves as the time-varying state with which the parameters are determined. Specifically, we define N_a models for when the mechanical system is off ($m^k = 0$) and N_b model for when the mechanical system is on ($m^k = 1$). Each of the $i = 1, \ldots, N_a$ and $j = 1, \ldots, N_b$ submodels describe a particular range of time steps after the mechanical system has switched from an on to an off state or vice versa. When the system is off, we define the length of each range as δ_a, the number of ranges as N_a, and the number of time steps since the system was last on before switching off as λ_a (i.e. if $m^{k-1} = 1$ and $m^k = 0$ then $\lambda_a = 1$). Likewise, when the system is on, we define the length of each range as δ_b, the number of ranges as N_b, and the number of time steps since the system was last off as λ_b (i.e. if $m^{k-1} = 0$ and $m^k = 1$ then $\lambda_b = 1$). Thus, the piecewise thermal model is given by

$$T^{k+1} = \begin{cases} \theta_{a,1} T^k + (1-\theta_{a,1}) T^k_\infty + \theta_{c,1} \\ \quad \text{if } m^k = 0 \\ \quad \text{and } \lambda_a \leq \delta_a \\ \theta_{a,2} T^k + (1-\theta_{a,2}) T^k_\infty + \theta_{c,2} \\ \quad \text{if } m^k = 0 \\ \quad \text{and } \delta_a < \lambda_a \leq 2\delta_a \\ \qquad \vdots \\ \theta_{a,N_a} T^k + (1-\theta_{a,N_a}) T^k_\infty + \theta_{c,N_a} \\ \quad \text{if } m^k = 0 \\ \quad \text{and } \lambda_a > (N_a - 1)\delta_a \\ \theta_{a,N_a} T^k + (1-\theta_{a,N_a}) T^k_\infty \\ + \theta_{c,N_a} + \theta_{b,1} \\ \quad \text{if } m^k = 1 \\ \quad \text{and } \lambda_b \leq \delta_b \\ \theta_{a,N_a} T^k + (1-\theta_{a,N_a}) T^k_\infty \\ + \theta_{c,N_a} + \theta_{b,2} \\ \quad \text{if } m^k = 1 \\ \quad \text{and } \delta_b < \lambda_b \leq 2\delta_b \\ \qquad \vdots \\ \theta_{a,N_a} T^k + (1-\theta_{a,N_a}) T^k_\infty \\ + \theta_{c,N_a} + \theta_{b,N_b} \\ \quad \text{if } m^k = 1 \\ \quad \text{and } \lambda_b > (N_b - 1)\delta_b \end{cases} \qquad (8)$$

where $\theta_{a,i}$ and $\theta_{c,i}$ are the parameters for the i-th model $i = 1, \ldots, N_a$ and $\theta_{b,j}$ is the parameter for the j-th model $j = 1, \ldots, N_b$. When the system is on, we employ the θ_{a,N_a} and θ_{c,N_a} parameters regardless of λ_a. In the following sections, we describe a recursive method for estimating the parameters in (8) using a Kalman filter.

3. Parameter Estimation Background

A fundamental machine learning problem involves the identification of a linear mapping

$$y^k = \theta^T x^k \tag{9}$$

where variable $x^k \in \mathbf{R}^X$ is the input, $y^k \in \mathbf{R}^Y$ is the output, and the linear map is parameterized by $\theta \in \mathbf{R}^{X \times Y}$. Additionally, X and Y are the number of inputs and outputs, respectively.

3.1. Batch Parameter Estimation

Learning can be performed in a batch manner by producing estimates of the parameters $\hat{\theta}$ given a training set of observed inputs and desired outputs, $\{x, y\}$. The goal of a parameter estimation algorithm is to minimize some function of the error between the desired and estimated outputs as given by $e^k = y^k - \hat{\theta}^T x^k$.

3.2. Recursive Parameter Estimation

The parameter estimation problem can be expressed in a recursive form using a discrete-time state-space model representation

$$\theta^k = \theta^{k-1} + n^k \tag{10a}$$

$$y^k = (\theta^k)^T x^k + e^k \tag{10b}$$

where θ^k represents the parameter estimates at time step k and $n^k \in \mathbf{R}^X$ corresponds to the parameter update noise (i.e. change in parameter values). The goal of a recursive parameter estimation algorithm is to produce $\hat{\theta}^k$ so as to minimize some function of the error e^k.

4. Kalman Filter Background

The Kalman filter (KF) is a recursive estimator for linear models such as the discrete-time state-space model

$$x^k = Ax^{k-1} + Bu^k + v^k \tag{11a}$$

$$y^k = Cx^k + Du^k + w^k \tag{11b}$$

where variable $x^k \in \mathbf{R}^X$ is the state of the system, $u^k \in \mathbf{R}^U$ is the known exogenous input, and $y^k \in \mathbf{R}^Y$ is the observed measurement signal. The state transition model is given by $A \in \mathbf{R}^{X \times X}$ and the control-input model by $B \in \mathbf{R}^{X \times U}$. The process noise $v^k \in \mathbf{R}^X$ has covariance $Q_v \in \mathbf{R}^{X \times X}$, $v^k \sim N(0, Q_v)$. The observation model is given by $C \in \mathbf{R}^{Y \times X}$ and the feedthrough model by $D \in \mathbf{R}^{Y \times U}$. The measurement noise $w^k \in \mathbf{R}^Y$ has covariance $Q_w \in \mathbf{R}^{Y \times Y}$, $w^k \sim N(0, Q_w)$. The variances of v^k and w^k (i.e. diagonal elements of Q_v and Q_w, respectively) must be known in order to implement a Kalman filter.

The Kalman filter (KF) algorithm consists of a prediction step and an update/correction step. The KF will model x^k as a Gaussian random variable (GRV) with estimated mean $\hat{x}^k \in \mathbf{R}^X$ and covariance $Q_x^k \in \mathbf{R}^{X \times X}$. To provide clarity, it is helpful to expand the k notation to distinguish between the state estimates produced before and after the KF correction step. Therefore, at each time step k, the predicted (a priori) state estimate, denoted as $\hat{x}^{k|k-1}$, is the mean estimate of x^k given measurements $y^0, \ldots, y^{k-1}$. The corrected (a posterior) state estimate, $\hat{x}^{k|k}$, is the mean estimate of x^k given measurements $y^0, \ldots, y^k$. To reiterate, throughout this paper, the uncorrected predictions (a priori) are denoted by $k|k-1$ or $k+1|k$ whereas the corrected predictions (a posterior) are denoted by $k|k$, $k-1|k-1$, or $k+1|k+1$.

The KF prediction step is given by

$$\hat{x}^{k|k-1} = A\hat{x}^{k-1|k-1} + Bu^k \tag{12a}$$

$$Q_x^{k|k-1} = AQ_x^{k-1|k-1}A^T + Q_v \tag{12b}$$

and the update/correction step by

$$\hat{y}^k = C\hat{x}^{k|k-1} + Du^k \tag{13a}$$

$$Q_y = CQ_x^{k|k-1}C^T + Q_w \tag{13b}$$

$$\mathcal{K} = Q_x^{k|k-1}C^T Q_y^{-1} \tag{14a}$$

$$r^k = y^k - \hat{y}^k \tag{14b}$$

$$\hat{x}^{k|k} = \hat{x}^{k|k-1} + \mathcal{K}r^k \tag{14c}$$

$$Q_x^{k|k} = Q_x^{k|k-1} - \mathcal{K}Q_y\mathcal{K}^T \tag{14d}$$

Figure 1 illustrates the KF algorithm. The block TD represents a time delay (commonly denoted in controls literature by z^{-1} or $1/z$, the Z-transform of the delay operator).

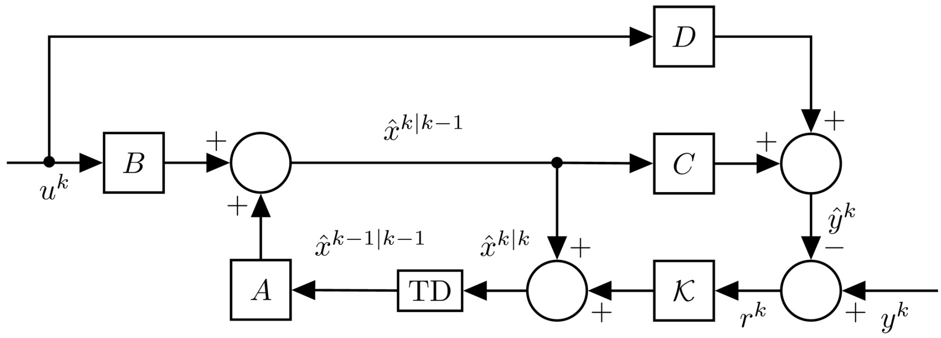

Figure 1: Kalman Filter Diagram

5. Recursive Parameter Estimation of Piecewise Thermal Model

In this section, we present a parameter estimation method for the piecewise linear thermal model (8) using the Kalman filter (KF) algorithm. To begin, we consider the recursive parameter estimation problem for the linear model (7), which can be expressed with the state-space model

$$\theta^k = \theta^{k-1} + n^k \tag{15a}$$

$$y^k = \begin{bmatrix} T^k - T^k_\infty & m^k & 1 \end{bmatrix} \theta^k + T^k_\infty + q^k \tag{15b}$$

where (15b) combines (7) with the observation model $y^k = T^{k+1} + w^k$. The noise term $q^k \in \mathbf{R}$ has covariance $Q_q \in \mathbf{R}$, $q^k \sim N(0, Q_q)$, and represents the sum of the process noise v^k and the measurement noise w^k. The parameter update noise $n^k \in \mathbf{R}^3$ has covariance $Q_n \in \mathbf{R}^{3\times3}$, $n^k \sim N(0, Q_n)$.

If we model (15) using the Kalman filter with

$$\begin{aligned} A &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} & B &= \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \\ C &= \begin{bmatrix} T^k - T^k_\infty & m^k & 1 \end{bmatrix} & D &= \begin{bmatrix} 1 \end{bmatrix} \\ u^k &= \begin{bmatrix} T^k_\infty \end{bmatrix} & y^k &= \begin{bmatrix} T^{k+1} \end{bmatrix} \end{aligned} \tag{16}$$

then the estimated state $\hat{\theta}^{k|k} \in \mathbf{R}^3$ will be a time-varying estimate of the thermal model parameters θ_a, θ_b, and θ_c. Note that the values of the C matrix will also be time-varying.

To learn the piecewise thermal model, we define a Kalman filter for each of the $N_a + N_b$ models in (8). When the mechanical system is off ($m^k = 0$), the state-space model is

$$\theta_i^k = \theta_i^{k-1} + n_i^k \tag{17a}$$

$$y^k = \begin{bmatrix} T^k - T^k_\infty & 1 \end{bmatrix} \theta_i^k + T^k_\infty + q_i^k \tag{17b}$$

and thus the filter models take the form

$$\begin{aligned} A &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} & B &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ C &= \begin{bmatrix} T^k - T^k_\infty & 1 \end{bmatrix} & D &= \begin{bmatrix} 1 \end{bmatrix} \\ u^k &= \begin{bmatrix} T^k_\infty \end{bmatrix} & y^k &= \begin{bmatrix} T^{k+1} \end{bmatrix} \end{aligned} \tag{18}$$

for each of the $i = 1, \ldots, N_a$ models. Therefore, the estimated state $\hat{\theta}_i^{k|k} \in \mathbf{R}^2$ is a time-varying estimate of the thermal model parameters $\theta_{a,i}$ and $\theta_{c,i}$. Additionally, $n_i^k \in \mathbf{R}^2$ and $q_i^k \in \mathbf{R}$ are noise terms for each model.

When the mechanical system is on ($m^k = 1$), the state-space model is

$$\theta_j^k = \theta_j^{k-1} + n_j^k \tag{19a}$$

$$y^k = \theta_j^k + (T^k - T^k_\infty)\theta_{a,N_a} + T^k_\infty + \theta_{c,N_a} + q_j^k \tag{19b}$$

and the filter models are given by

$$\begin{aligned} A &= \begin{bmatrix} 1 \end{bmatrix} \quad B = \begin{bmatrix} 0 \end{bmatrix} \\ C &= \begin{bmatrix} 1 \end{bmatrix} \quad D = \begin{bmatrix} 1 \end{bmatrix} \\ u^k &= \begin{bmatrix} \theta_{a,N_a}(T^k - T^k_\infty) + T^k_\infty + \theta_{c,N_a} \end{bmatrix} \\ y^k &= \begin{bmatrix} T^{k+1} \end{bmatrix} \end{aligned} \tag{20}$$

for $j = 1, \ldots, N_b$ where θ_{a,N_a} and θ_{c,N_a} are the current parameter estimates of model $i = N_a$. Therefore, the estimated state $\hat{\theta}_j^{k|k} \in \mathbf{R}$ is an estimate of θ_b. Additionally, $n_j^k \in \mathbf{R}$ and $q_j^k \in \mathbf{R}$ are noise terms for each model.

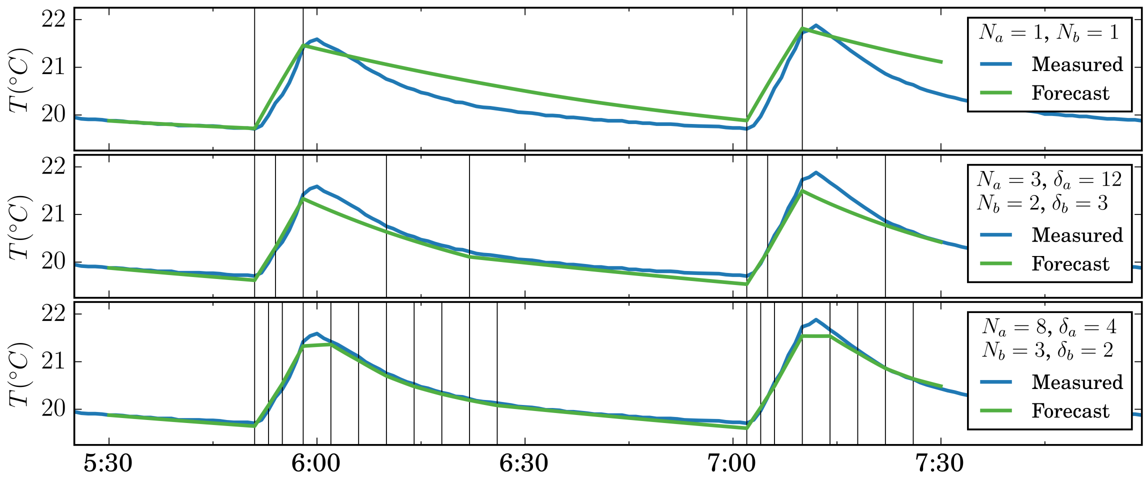

Figure 2: Examples of piecewise thermal models with $N_a = N_b = 1$ (top), $N_a = 3$ and $N_b = 2$ (middle), and $N_a = 8$ and $N_b = 3$ (bottom) used to produce a 2 hour forecast.

When training the models, the temperature measurement T^{k+1} at each time step is used as an observation to train only one of the $N_a + N_b$ submodels as given by (8). In this way, each submodel is learning to represent a particular characteristic of the thermal dynamics of the system. For the remainder of the submodels, the parameter estimates are unchanged (i.e. $\hat{\theta}_i^k = \hat{\theta}_i^{k-1}$ for each filter i where i is not the observed model).

With respect to the covariances $Q_\theta^{k|k}$, there are two ways of updating the matrices. The first is to set each covariance matrix to the previous value. This expresses that, even though we did not observe the model in the current time step, we have not lost confidence in the parameter estimates. Alternatively, we can add the process noise covariance (as done in the Kalman filter prediction step (12b)), expressing a increasing loss of confidence in the parameter values. In this paper, we assume the former and only alter the covariance matrix when the model is observed.

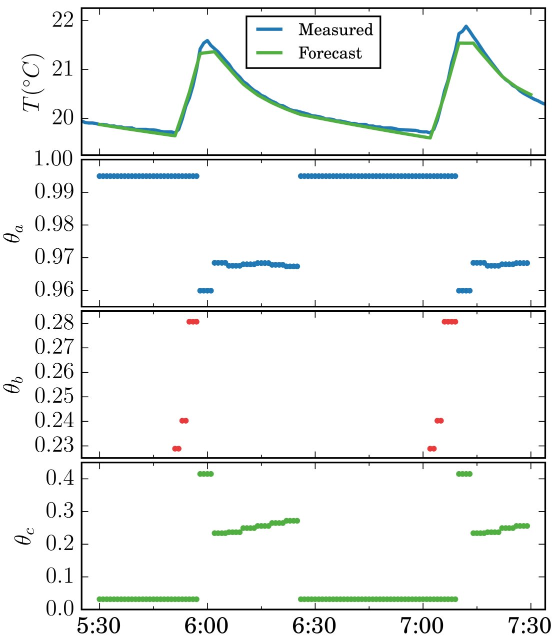

Figure 3: Piecewise thermal model parameters with $N_a = 8$ and $N_b = 3$ used to produce a 2 hour forecast.

6. Residential Heating System Parameter Estimation Experiments

In this section, we present parameter estimation results for an 850 sq ft apartment with a forced-air heating and ventilation system. The apartment is located in Berkeley, California and equipped with a custom thermostat designed and built for this research. Therefore, we are able to control the operation of the heating system and to measure the indoor air temperature. Local weather data, specifically ambient air temperature, is retrieved from the Internet service, Weather Underground [13].

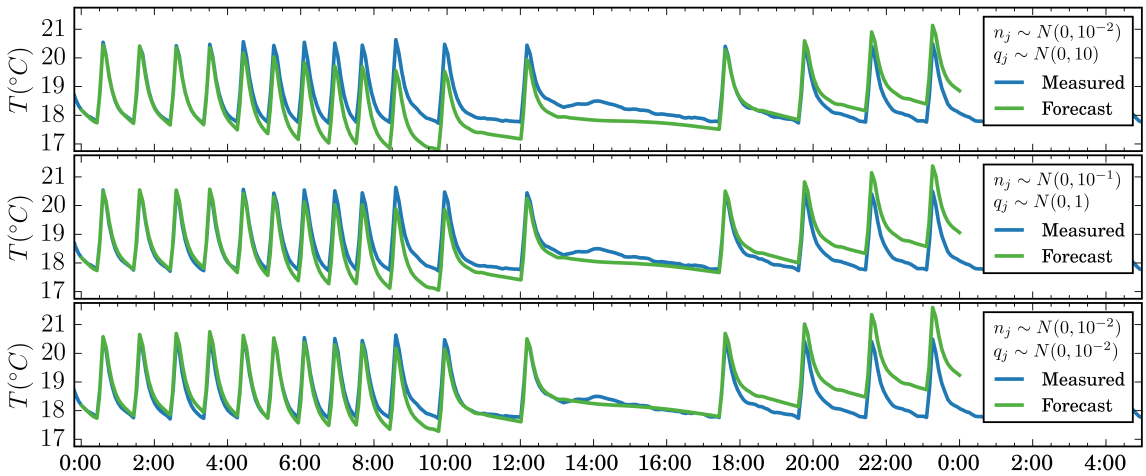

Figure 4: Examples of piecewise thermal models used to produce a 24 hour forecast with various covariance values for the n_j and q_j noise terms. In each case, $N_a = 15$, $\delta_a = 5$, $N_b = 2$, $\delta_b = 4$, $n_{i,1} \sim N(0, 10^{-4})$, $n_{i,2} \sim N(0, 10^{-4})$, $q_i \sim N(0, 10^{-1})$.

Data was collect at a time-scale of one minute for 6 weeks during December and January of 2015-2016. With this data, we are able to perform recursive parameter estimation of the piecewise thermal model (8). The results presented in this section focus of quantifying and qualifying the advantages of the piecewise model and the Kalman filter based learning method.

Fig. 2 presents a comparison of thermal models using varying numbers of N_a and N_b submodels. In each subplot, a 2 hour forecast is produced using the model parameters as estimated at the start of the time horizon. The vertical lines designate the start and end of each model's corresponding range. The top subplot shows the most basic case where $N_a = N_b = 1$, for a total of 2 submodels. As shown, the model is simply incapable of representing the evolution of the indoor air temperature. Most notably, the forecast poorly accounts for the thermal dynamics immediately after the heating system turns off. These dynamics are related to the interaction between the air and the other thermal masses (walls, furniture, etc.) within the conditioned space. These dynamics could, in theory, be captured by a higher order model, but this would increase the model complexity and the need for temperature sensor measurements.

By increasing the number of submodels, as shown in the second and third subplots of Fig. 2, the piecewise thermal model is able to better approximate the dynamics of the apartment and heating system without significantly increasing the model complexity. Fig. 3 presents a forecast produced by a piecewise model with $N_a = 8$, $\delta_a = 4$, $N_b = 3$, and $\delta_b = 2$. The top subplot shows the 2 hour forecast and the measured air temperature within the apartment. The remaining subplots show the θ_a, θ_b, and θ_c parameter values employed by the piecewise model at each time step of the forecast.

Fig. 4 illustrates the ability of the model to produce accurate multi-hour forecasts and the influence of the noise covariances on the parameter estimates. In each subplot, different covariance values are used to represent the n_j and q_j noise terms and the forecasts are produced using the model parameters as estimated at the start of the time horizon. In the top subplot, the model is very accurate for the first several hours before the forecasted temperature begins to drift downward. The root mean squared error (RMSE) over the first 3 hours is 0.039 °C and over the first 12 hours is 0.573 °C. In the bottom subplot, the error is less varied with an RMSE of 0.162 °C over the first 3 hours and 0.240 °C over the first 12 hours.

7. Conclusions

This paper addresses the need for control-oriented thermal models of buildings. We present a piecewise linear thermal model of a building that is suitable for model predictive control applications. To estimate the model parameters, we present a

Kalman filter based system identification method. Finally, we present experimental results using real temperature data collected from an apartment with a forced-air heating and ventilation system. These results demonstrate the potential of the model and parameter estimation method to produce accurate forecasts of the air temperature within the apartment.

References

[1] U.S. Department of Energy, 2010 Buildings Energy Data Book, accessed May. 2, 2014.
URL `http://buildingsdatabook.eren.doe.gov`

[2] R. Brown, Us building-sector energy efficiency potential, Lawrence Berkeley National Laboratory.

[3] E. M. Burger, S. J. Moura, Generation Following with Thermostatically Controlled Loads via Alternating Direction Method of Multipliers Sharing Algorithm (October 2015).
URL `http://escholarship.org/uc/item/2m5333xx`

[4] A. Aswani, N. Master, J. Taneja, V. Smith, A. Krioukov, D. Culler, C. Tomlin, Identifying models of hvac systems using semiparametric regression, in: American Control Conference (ACC), 2012, IEEE, 2012, pp. 3675–3680.

[5] A. Handbook-Fundamentals, American society of heating, refrigerating and air-conditioning engineers, Inc., NE Atlanta, GA 30329.

[6] Y. Lin, T. Middelkoop, P. Barooah, Issues in identification of control-oriented thermal models of zones in multi-zone buildings, in: Decision and Control (CDC), 2012 IEEE 51st Annual Conference on, IEEE, 2012, pp. 6932–6937.

[7] C. Agbi, Z. Song, B. Krogh, Parameter identifiability for multi-zone building models, in: Decision and Control (CDC), 2012 IEEE 51st Annual Conference on, IEEE, 2012, pp. 6951–6956.

[8] P. Radecki, B. Hencey, Online building thermal parameter estimation via unscented kalman filtering, in: American Control Conference (ACC), 2012, IEEE, 2012, pp. 3056–3062.

[9] S. Ihara, F. C. Schweppe, Physically based modeling of cold load pickup, IEEE Transactions on Power Apparatus and Systems 100 (9) (1981) 4142–4250.

[10] R. E. Mortensen, K. P. Haggerty, A stochastic computer model for heating and cooling loads., IEEE Transactions on Power Systems 3 (3) (1998) 1213–1219.

[11] E. M. Burger, S. J. Moura, Recursive Parameter Estimation of Thermostatically Controlled Loads via Unscented Kalman Filter (October 2015).
URL `http://escholarship.org/uc/item/7t453713`

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

[13] Weather Underground Web Service and API.
URL `wunderground.com/weather/api/`