

UC Berkeley

Building Efficiency and Sustainability in the Tropics (SinBerBEST)

Title

Total and Peak Energy Consumption Minimization of Building HVAC Systems Using Model Predictive Control

Permalink

<https://escholarship.org/uc/item/81m4p0g8>

Journal

IEEE Design & Test of Computers, 29(4)

ISSN

0740-7475

Authors

Maasoumy, Mehdi
Sangiovanni-Vincentelli, Alberto

Publication Date

2012-08-01

DOI

10.1109/MDT.2012.2200871

Peer reviewed

Total and Peak Energy Consumption Minimization of Building HVAC Systems Using Model Predictive Control

**Mehdi Maasoumy and
Alberto Sangiovanni-Vincentelli**
University of California

Editors' notes:

This article addresses the challenge of realizing the building automation and control system using a distributed network of embedded computers. A specification methodology and design space exploration framework are proposed to raise the level of abstraction at which building control systems are designed, to reduce design effort, and to lower implementation cost.

—Yuvraj Agarwal, University of California, and
Anand Raghunathan, Purdue University

■ **ADVANCED CONTROL ALGORITHMS** are considered critical enablers to achieve low energy consumption in commercial buildings. Entire sections of the ASHRAE 90.1 standard [1] are dedicated to the specification of control requirements. Although the optimal control of an HVAC system is a complex multi-variable problem, it is standard practice to rely on simple control strategies that include PID and bang-bang controllers with hysteresis. In most cases, standard sequences of operations for *typical* installations are used by control contractors. Each sequence controls the HVAC

his operations.

This short summary of design and validation practices in the building industry shows the importance of a model-based design flow for building controls. To attain energy efficiency, control algorithms need to be tailored to the physical properties of the building at hand rather than being an adaptation of a standard sequence designed for a “typical” building. To design an optimal controller that balances comfort and energy usage, a thermal model of the building is needed. To achieve building-level energy-optimality, the model should be able to capture the interaction between physically connected spaces in the building, occupancy schedules, and state and input constraints.

A variety of approaches can be found in the literature. The authors of [2] proposed a nonlinear

equipment during an operation phase such as optimal start, safety shutdown and normal operation. After installation and tuning, the building is inspected by a commissioning agent that verifies that the building satisfies the owner's expectations. The commissioning agent does not only verify the expected performance right after installation, but also after the building has started

Digital Object Identifier 10.1109/MDT.2012.2200871

Date of publication: 08 June 2012; date of current version:

05 October 2012.

model of the overall cooling system including chillers, cooling towers and thermal storage tanks, and developed an MPC scheme for minimizing energy consumption. In [3] and [4], the authors use a model of the building which is bilinear in inputs, states and weather parameters. The approach is a form of Sequential Quadratic Programming (SQP) for solving nonlinear problems in which they iteratively linearize the non-convex constraints around the current solution, solve the optimization problem and repeat until a convergence condition is met. In [5], the authors proposed a building thermal model based on an RC-network, with a large number of coupled linear differential equations and reduce the order of the model via aggregation of states. In [6], the authors proposed a simple model where the model inputs are divided into manipulated variables and disturbance inputs.

The idea of modeling building thermal behavior using RC networks was introduced in [7]. The model was validated against historical data and a tracking LQR was proposed as controller in cases where tight temperature tracking is required (e.g., data centers). In this paper, we introduce controllers also for commercial buildings where there is a wider acceptable range of temperature especially during unoccupied hours that allow for more relaxed temperature boundaries.

This paper, building upon the results described above, achieves two major contributions.

- Estimating the “unmodeled dynamics of the system.” In most climates the unmodeled dynamics account for a significant contribution in the thermal dynamics of buildings. The authors of [8] propose a method for estimating the unmodeled dynamics of a *linear* building thermal model. However, the building thermal models proposed in the literature are mostly nonlinear models with a bilinear term, for example, the multiplication of the state and input. We propose a method for estimating the unmodeled dynamics of the nonlinear system which comprises the effect of the external heat loads from solar radiation and outside air temperature as well as the internal heat gains from occupants using a parameterization of the unmodeled dynamics with two quantities being the *outside temperature* and *CO2 sensor* data.

- The implementation of two controllers, different from the original controller of the real building, on the building model. After extensive comparison, we demonstrated significant achievements in reducing the total and peak energy consumptions for the HVAC system.

The paper is organized as follows. Section I presents the proposed high-level thermal model for buildings. The parameterizations of the external and internal loads are laid out in this section as well. We present the identification results in the second sections. We explain the linearization and the division of inputs into *manipulated variables* and *disturbance inputs* in the third section. The fourth section describes the two controllers designed for the system. Finally, the fifth section presents results obtained from simulation of controllers and a comparison of their performance.

Modeling and validation

The detail of building thermal modeling is presented in [7]. Here we focus on the *model manipulation* and the estimation of *parameters* and *unmodeled dynamics* using historical data.

Mathematical model of the building

We model the thermal dynamics of the building climate by first analyzing the corresponding *thermal network*. Each room or wall is represented by a *node* in the thermal network. The nodes are connected via resistors and capacitors to the neighboring nodes and to the ground, respectively. The capacitors model thermal capacitance of walls and rooms, and resistances model thermal resistances between the nodes. We have approximated the temperature of each wall with the temperature of its centerline T_w and the temperature of each room is assumed to be the average temperature of the room, T_i . Each T_w is connected with two series resistors with resistances $R_{wall}/2$ and R_{int} to the neighboring node which is a room temperature or for peripheral walls with two series resistances $R_{wall}/2$ and R_{out} to the outside air temperature. Windows are modeled as a parallel resistance with the rest of the wall, with lower thermal resistance.

We derive a mathematical model of the thermal behavior of buildings that can be effectively used in control design. The model is for rooms on the same

floor and for the sake of simplicity are considered isolated from rooms on adjacent floors. We use a lumped model where the air in a room has one temperature across its volume, and the temperature of a wall across its volume is assumed to be equal to its centerline temperature. We assume that all rooms are at the same pressure which is equal to the pressure used in the heating and cooling ducts. Air exchange between a room and a vent is then *isobaric*, so the air mass in the room will not change in the process. We denote the air mass in the room by m , the rate of air mass entering the room and also leaving the room by m_r , and the temperature of the conditioned air entering room i by T_{sr} . The temperature of the air leaving the room equals the current temperature of the room. We ignore the capacitance of windows since their mass is negligible compared to the mass of walls (i.e., windows are modeled as pure resistances in the thermal circuit). The radiative heating for each building face is regarded as a disturbance to the plant model. Our network consists of two types of nodes: walls and rooms. There are in total n nodes, m of which represent rooms and the remaining $n - m$ nodes represent walls.

The temperature of the i th wall is governed by the following equation:

$$C_{w_i} \frac{dT_{w_i}}{dt} = \sum_{j \in \mathcal{N}_{w_i}} \frac{T_j - T_{w_i}}{R'_{ij}} + r_i \alpha_i A_i q''_{rad_i}$$

where T_{w_i} , C_{w_i} , α_i , and A_i are the temperature, heat capacity, absorption coefficient, and area of wall i , respectively. R'_{ij} is the total resistance between wall i and node j . q''_{rad_i} is the radiative heat flux density on wall i . \mathcal{N}_{w_i} is the set of all of neighboring nodes to node w_i . r_i is equal to 0 for internal walls, and to 1 for peripheral walls.

The temperature of the i th room is governed by the following equation:

$$C_{r_i} \frac{dT_{r_i}}{dt} = \sum_{j \in \mathcal{N}_{r_i}} \frac{T_j - T_{r_i}}{R'_{ij}} + \dot{m}_{r_i} c_a (T_{s_i} - T_{r_i}) + w_i \tau_{win_i} A_{win_i} q''_{rad_i} + \dot{q}_{int}$$

where T_{r_i} , C_{r_i} , and \dot{m}_{r_i} are the temperature, heat capacity, and air mass flow into the room i , respectively. c_a is the specific heat capacity of air, A_{win_i} is

the total area of window on walls surrounding room i , τ_{win_i} is the transmissivity of glass of window i , q''_{rad_i} is the radiative heat flux density radiated to the area of node i and \dot{q}_{int} is the internal heat generation in thermal zone i . \mathcal{N}_{r_i} is the set of all of the neighboring nodes to room i and, w_i is equal to 0 if none of the walls surrounding room i has window, and is equal to 1 if at least one of them has.

We use the historical data of zone 8 in the UC Berkeley Bancroft library, a conference room, to estimate these loads and identify the parameters of the model. This thermal zone is on the corner of the building surrounded from the West by thermal zone 10 and from the South by thermal zone 7 and from North and East to the outside environment. The walls on the Northern and Eastern sides have windows allowing external radiation to penetrate into the room.

By writing the heat transfer equation for every wall and room in the building and representing the equations in a state space form we get the following form of equation:

$$\dot{x}_t = Ax_t + f(x_t, u_t, d_t) \quad (1)$$

$$y_t = Cx_t \quad (2)$$

where $x_t \in^{n \times 1}$ is the state vector representing the temperature of the nodes in the thermal network, $u_t \in^{(r \times m) \times 1}$ is the input vector which in this case is the mass flow rate and discharge air temperature of conditioned air into each thermal zone, and $y_t \in^{m \times 1}$ is the output vector of the system which represents the temperature of the thermal zones. A is a square $n \times n$ matrix and $C \in^{m \times n}$ determines which states are used as the output of the system. Vector $f(x_t, u_t, d_t) \in^{n \times 1}$ is composed of the states, inputs, and disturbance to the system.

Estimating external loads. Heat flux is radiated from the sun to the exposed walls and to the peripheral rooms through windows. This heat flux is a hard-to-estimate function of several variables including the *altitude* and *azimuth* angle of the location of the building on the Earth, orientation of the considered wall or window, day of the year, time of the day, outside weather, and sky condition. However, to be able to estimate the heat flux at each time, we approximate it by assuming that this

quantity is an affine function of the outside air temperature¹ given by

$$q''_{rad_i}(t) = \lambda T_{out}(t) + \gamma \quad (3)$$

where λ and γ will be obtained by the parameter estimation algorithm detailed later. Hence, we parameterize $q''_{rad_i}(t)$ and then identify all of the parameters using nonlinear regression. Note that this method does not cover all the uncertainties associated with external loads to the building, however it leads to a decent estimation as shown later in the second section.

Estimating internal loads. Internal loads in the building are usually related to occupants and electrical devices. The heat emitted from electrical devices is easy to identify based on the electrical characteristics of the device with high precision; the main uncertainty in identifying the internal loads is due to the load associated to the building occupants. We propose a parameterization of the internal loads by occupants using the CO_2 sensor data (current CO_2 concentration in the room) which yields

$$\dot{q}_{int}(t) = \mu \Psi(t) + \nu \quad (4)$$

where $\Psi(t)$ is the CO_2 concentration in the room in (ppm). μ and ν are constants to be obtained from the identification process.

Parameters and unmodeled dynamics identification

We use historical data to identify the parameters of the system along with the unmodeled dynamics described in the first section. The identification process is done through an optimization problem over the parameters given in (5)

$$\begin{aligned} & \min_{C_{[1]}, R_{[1]}, \lambda, \gamma, \mu, \nu} \|Y^m - Y^s\|_2^2 \\ \text{s.t.} \quad & \begin{cases} x_{t+1}^s = Ax_t^s + f(x_t^s, u_t^m, d_t^m) & t = 0, \dots, N-1 \\ y_t^s = Cx_t^s & t = 0, \dots, N \end{cases} \end{aligned} \quad (5)$$

where the subscript t refers to time and the superscript m and s refer to measured and simulated

¹Note that other quantities such as global horizontal irradiance (GHI) data either from the *California Irrigation Management Information System* (CIMIS) [9] or METAR data [10] from nearby airports archived in the *National Climatic Data Center*, can be used for the purpose of this parameterization as well.

data, respectively. The vector $Y \in^{N \times 1}$ stores the values of y_t for $t = 1, \dots, N$.

Parameter identification

For identifying the parameters of the model we have used the data of zone temperature of a specific zone at Bancroft library of UC Berkeley campus along with airflow, discharge air temperature (DAT) and outside air temperature (OAT) data to simulate the thermal behavior of that specific zone and then compare the simulation results with the measured temperature of the zone. We have used the WebCTRL of Automated Logic Corporation (ALC) to download the temperature data. The results of model validation is shown in Figure 1. We store the unmodeled dynamics in a time-varying vector called disturbance d_t . From now on we use the terms “disturbance” and “unmodeled dynamics” interchangeably.

Note that there are two peaks in the unmodeled dynamics values, one around 9 a.m. and another around 3 p.m. which are due to occupants and outside radiation. The interesting observation is that the first peak of disturbance load does not cause as much temperature increase in the room as opposed to the second disturbance peak. The reason is that in the morning the walls which represent the slow-dynamic masses in the system are cold due to low temperature at night. Therefore, part of the heating load, earlier in the day, goes toward warming up the slow-dynamic thermal masses in the building (e.g., walls and furniture). However, in the afternoon the slow-dynamic thermal masses absorb heat at a slower rate and therefore, cause faster increase to the temperature of the fast-dynamic thermal mass of the system which is the air in the room. Also a decrease in the values of the unmodeled dynamics is observed which can be due to the people leaving the room around noon for lunch and/or cloudy sky. The optimal parameters of the model are reported in [7].

Parameter validation

In order to validate the parameters of the model, we have simulated the temperature of the same thermal zone using the data of next weekend. The results of the simulation are presented in [7].

Linearization

We linearize the system dynamics around the nearest equilibrium point to the specified operating




Figure 1. Simulated temperature and measured temperature of zone 8 of Bancroft library (Oct. 24, 2010) and unmodeled dynamics.

point of the system (details in [11]). The algorithm to find the equilibrium point of the system starts from an initial point and searches, using a *sequential quadratic programming algorithm*, until it finds the nearest equilibrium point. First we linearize the model considering all the inputs to the model. Once the linearization is done, we divide the inputs into *manipulated variables* and *disturbance variables*. Discretizing the state space realization leads to $x_{k+1} = Ax_k + Bu_k + Ed_k$ where d_k stores the disturbance at time k and the original B obtained from linearization process is split into two parts. The new B keeps the columns corresponding to the manipulated variables and E stores the columns of the original B corresponding to the disturbance variables. In this study we have kept the air flow as a manipulated variable and we regard the rest of the inputs as the disturbance input on which we don't have control. Note that since the range of the variations of inputs as shown in Figure 2 during day (on-mode) and also the thermal zone temperature that the system experiences in the course of a day is not so wide (usually 20 °C–22 °C), linearizing about the equilibrium point does not introduce significant error as shown in Figure 3. On the other hand dealing with a linear system dramatically decreases the computa-

tional efforts. The results of the linearization is shown in Figure 3.

Controller design

We presented the results of an optimal controller for the tracking case and compared the results with a PID controller in [7]. The introduced tracking LQR in [7] fits best to cases where a tight temperature tracking is required, for example, data centers. However, the wider acceptable range of temperature for commercial buildings, especially during unoccupied hours, allows for more relaxed temperature boundaries. Accordingly, we have introduced new controllers to make the best use of this flexibility and achieve more savings. In this paper we study the performance of two different controllers on the temperature regulation problem. Here we have implemented two different controllers on the system and compared their performance with the original controller on the existing building. The way the original controller works is that the controller opens the valves of conditioned air flow to the thermal zones at 5:00 am and keeps it fully opened till 5:00 pm. The discharge air temperature is also kept constant at 47 °C during that time period.




Figure 2. Measured data of air flow and discharge air temperature.

We implement two controllers: an *on-off controller* and a *Model Predictive Controller*. For controller implementation, we assume that we can only manipulate the air flow valve while the discharge air

temperature remains the same as before. Note that due to the weather condition at Berkeley, where the considered building is located, the HVAC system only needs to provide heating (according to




Figure 3. Temperature of room from measured data, nonlinear model and the linearized model.

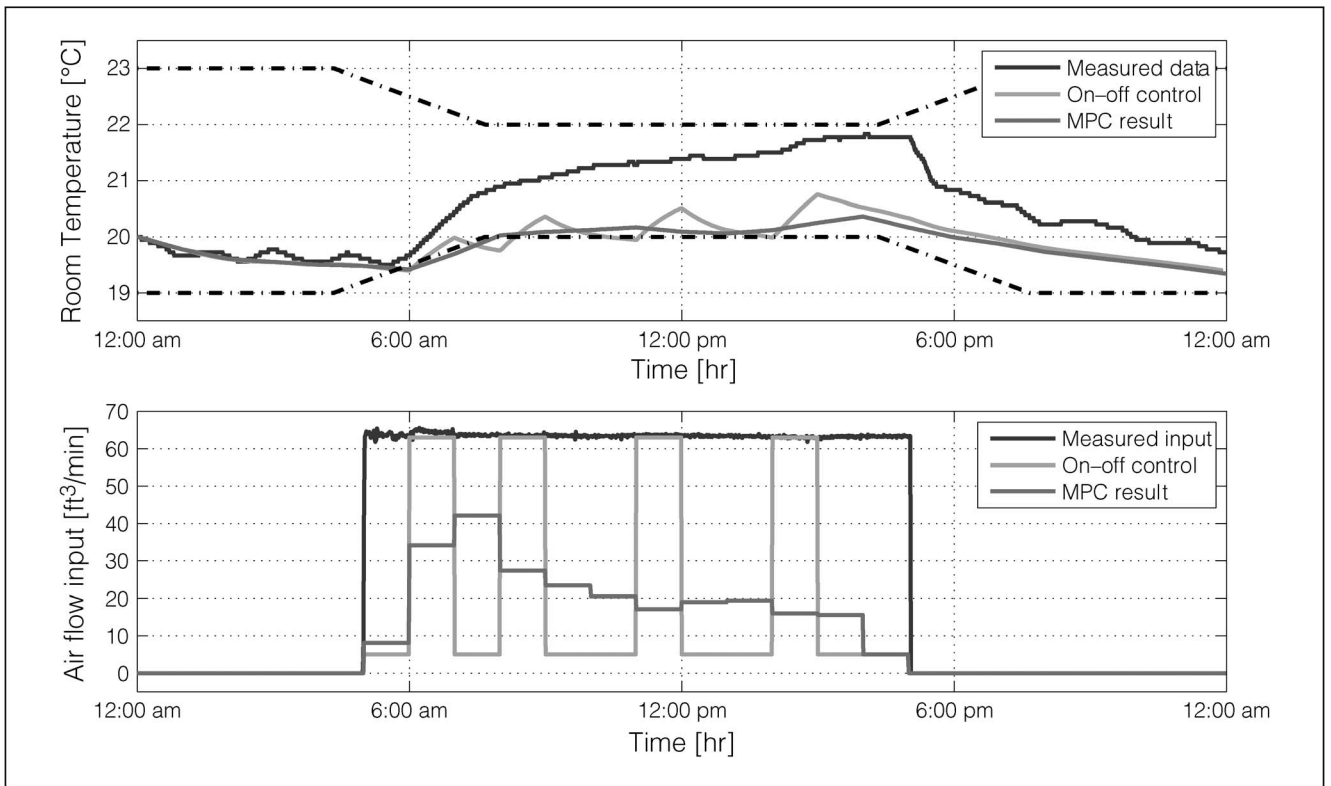


Figure 4. Measured data and the result of on-off controller and MPC.

Figure 4), and cooling is done naturally by turning off the heating mode of the HVAC system and by running the ventilation system. Note that even if no cooling is needed the rooms have to be ventilated for air quality reasons. For the two following controllers we consider a time-varying lower bound and upper bound for temperature which define the comfort zone to be between 20 °C and 22 °C during day and between 19 °C and 23 °C at night as shown in Figure 4.

On-off controller

For the *on-off controller*, we assume that the valves can have three states: fully opened, minimally opened (not fully closed due to air quality reasons) during occupied hours, or fully closed at night, and the duration of each state of the valve cannot be less than 1 hour, for example, when the valve is set to open it has to remain open for at least 1 hour before it can close. The controller turns on the heating mode when the room temperature falls below the lower limit and turns it to either minimally open or fully closed (depending on whether it is occupied or unoccupied hours of the day) when the temperature

is within the comfort zone. The performance of this controller is depicted in Figure 4.

Model predictive controller

A model predictive control problem is formulated with the objective of minimizing a linear combination of the total energy consumption and the peak airflow. An MPC with similar cost function is also used in [12] with a simple linear model. However, here we have modified the controller in order to reflect the constraints of the considered system, the system dynamics and etc. Also we have implemented the control inputs obtained from the MPC which utilizes the linearized system dynamics of the model on the original nonlinear model.

Note that the fan energy consumption is proportional to the *cubic* of the airflow. Hence minimizing the peak airflow would dramatically reduce fan energy consumption. We have considered a cost function for the MPC which comprises linear combination of the total heating power consumption ($l-1$ norm of input) and the peak of airflow ($l-\infty$ norm of input). Also in order to guarantee feasibility (constraint satisfaction) at all times we have

implemented soft constraints. The predictive controller solves at each time step the following problem:

$$\min_{U_t, \underline{\varepsilon}, \bar{\varepsilon}} \{ |U_t|_1 + \kappa |U_t|_\infty + \rho (|\bar{\varepsilon}_t|_1 + |\underline{\varepsilon}_t|_1) \} = \quad (6a)$$

$$\min_{U_t, \underline{\varepsilon}, \bar{\varepsilon}} \left\{ \sum_{k=0}^{N-1} |u_{t+k|t}| + \kappa \max(|u_{t|t}|, \dots, |u_{t+N-1|t}|) + \rho \sum_{k=1}^N (|\bar{\varepsilon}_{t+k|t}| + |\underline{\varepsilon}_{t+k|t}|) \right\} \quad (6b)$$

$$\text{s.t. } x_{t+k+1|t} = Ax_{t+k|t} + Bu_{t+k|t} + Ed_{t+k|t}, \quad k = 0, \dots, N-1 \quad (6c)$$

$$y_{t+k|t} = Cx_{t+k|t}, \quad k = 1, \dots, N \quad (6d)$$

$$\underline{u}_{t+k|t} \leq u_{t+k|t} \leq \bar{u}, \quad k = 0, \dots, N-1 \quad (6e)$$

$$\underline{T}_{t+k|t} - \underline{\varepsilon}_{t+k|t} \leq y_{t+k|t} \leq \bar{T}_{t+k|t} + \bar{\varepsilon}_{t+k|t}, \quad k = 1, \dots, N \quad (6f)$$

$$\underline{\varepsilon}_{t+k|t}, \bar{\varepsilon}_{t+k|t} \geq 0, \quad k = 1, \dots, N \quad (6g)$$

where $U_t = [u_{t|t}, u_{t+1|t}, \dots, u_{t+N-1|t}]$ is the vector of control inputs, $\underline{\varepsilon} = [\underline{\varepsilon}_{t+1|t}, \dots, \underline{\varepsilon}_{t+N|t}]$ is the temperature violations from the lower bound, $\bar{\varepsilon}$ is the temperature violation from the upper bound, $y_{t+k|t}$ is the thermal zone temperature vector, $d_{t+k|t}$ is the disturbance load prediction, and $\underline{T}_{t+k|t}$ and $\bar{T}_{t+k|t}$ for $k = 1, \dots, N$ are the lower and upper bounds on the zone temperature, respectively. $\underline{u}_{t+k|t}$ and \bar{u} are the lower and upper limit on the airflow input by the VAV damper, respectively. Note that based on ASHRAE requirements for *Air Change per Hour* (ACH) of rooms, there has to be a minimum non-zero airflow during occupied hours for ventilation purposes. ρ is the penalty on the comfort constraint violations, and κ is the penalty on peak power consumption.

Remark 1. Note that for predicting the *disturbance load*, we use the optimal solution $(\lambda^*, \gamma^*, \mu^*, \nu^*)$ to the optimization problem (5) along with the predictions of $T_{out}(t)$ and $\Psi(t)$ for future times. The predictive values for these two quantities can be obtained from weather forecasts and from the occupancy schedules of each thermal zone of the building, respectively. Although here we are assuming a *perfect* forecast for these quantities in this formulation, imperfect weather and occupancy predictions

can potentially deteriorate the performance of MPC. The results of our recent study on the performance and effectiveness of MPC in the presence of forecast errors can be found in [13].

At each time step only the first entry of U_t is implemented on the plant. At the next time step the prediction horizon N is shifted leading to a new optimization problem. The prediction horizon is $N = 24$, and at each time step only the first entry of the input vector U_t is implemented on the model. This process is repeated over and over until the total time span of interest is covered. We use YALMIP [14] to set up the MPC problem in MATLAB.

We implement the hierarchical control algorithm proposed in [7]. At the lower level, each thermal zone is controlled by a PID controller while a model-based optimal control is used at the higher level for a group of thermal zones.

Simulation results

Original controller. The air flow valve is turned on from 5:00 am until 5:00 pm and remains off for the rest of the day. This approach results in a total airflow input of 45360 [ft³] per day and a maximum air flow rate of 63 [ft³/min].

On-off controller. The airflow valve is turned on only for 4 hours in that specific day as shown in Figure 4. This reduces the total airflow input to 17,520 [ft³] per day, which is 61.4% less than the original case and the maximum airflow rate remains at 63 [ft³/min].

MPC. The air flow valve is not just an on-off switch; rather any intermediate values can be set for the air flow into the room. The performance of the controller is shown in Figure 4. Implementation of MPC results in a total airflow input of 14870 [ft³] per day, which is 67.2% less than the original case and a

Table 1 Comparison of performance for three different controllers. The inputs are heating air flow to the room.

Controller	Total input [ft ³]	Peak input [ft ³ /min]	Total energy [kWh]	Running time [s]
Original control	45360	63	12.46	-
On-off control	17520	63	4.62	1.8
MPC	14870	42	3.33	102.4

maximum airflow rate of 42 [ft³/min] that is 33.3% less than the original case.

Remark 2. Note that the peak airflow required for a building is an important design parameter for sizing the HVAC system components at the design stage. A smaller Air Handling Unit (AHU) package (and hence a smaller fan) would, on top of the energy savings due to advanced control algorithms introduced here, lead to more electric power consumption reductions.

To compare the overall energy performance of the controllers we calculate the *total energy consumption* for each controller. The total energy consumption is given by

$$E_{tot} = \int_{t=0}^{24} [P_c(t) + P_h(t) + P_f(t)] dt \quad (7)$$

in which, the cooling power P_c , the heating power P_h and the fan power P_f are defined by

$$P_c(t) = \dot{m}_c(t)c_p[T_{out}(t) - T_c(t)] \quad (8)$$

$$P_h(t) = \dot{m}_h(t)c_p[T_h(t) - T_{out}(t)] \quad (9)$$

$$P_f(t) = \alpha \dot{m}^3(t) \quad (10)$$

where $c_p = 1.012$ [kJ/kg.°C] is the specific heat capacity of air and $\alpha = 0.5$ [kW.s³/kg³] is the fan power constant. Using the above equations and constants results in fan power values in [kW].

The comparisons of results are summarized in Table 1.

We presented a model based hierarchical control strategy that balances comfort and energy consumption. The building is modeled by a thermal network that captures the relevant dynamics of the temperature for each room taking into account the interactions between rooms, external loads, and building occupants. The parameters of the model and the unmodeled dynamics are identified using historical data of the thermal zone. We implemented an on-off controller and a Model Predictive Controller with time-varying thermal comfort zone during day and compared their performances. The on-off controller achieves 61.4% reduction in total air flow into the thermal zone but does not reduce the maximum airflow rate, however the MPC is shown to reduce the total airflow into the thermal zone by 67.2% and the maximum air flow by 33.3%. The MPC also reduces the total energy consumption of the

HVAC system by 73.2% with respect to the current control algorithm of the building. The simulation times for on-off controller and for MPC are 1.8 seconds and 102.4 seconds, respectively.

Acknowledgment

The authors would like to thank A. Pinto of UTRC for his helpful feedbacks during the development of the project. This research has been sponsored by the MARCO FCRP MuSyC.

References

- [1] *Amer. Soc. Heating, Refrigerating, Air-Conditioning Engineers*, 90.1-1989, 1989.
- [2] Y. Ma, F. Borrelli, B. Hancey, B. Coffey, S. Bengea, and P. Haves, "Model predictive control for the operation of building cooling systems," in *Proc. Amer. Control Conf. (ACC), 2010*, 2010, pp. 5106–5111, IEEE.
- [3] 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 *Proc. Amer. Control Conf. (ACC), 2010*, 2010, pp. 5100–5105, IEEE.
- [4] D. Gyalistras and M. Gwerder, "Use of weather and occupancy forecasts for optimal building climate control (OptiControl): Two years progress report," in *Proc. Terrestrial Systems Ecology ETH Zurich, Switzerland and Building Technologies Division, Siemens Switzerland Ltd., Zug, Switzerland*, 2010.
- [5] K. Deng, P. Barooah, P. Mehta, and S. Meyn, "Building thermal model reduction via aggregation of states," in *Proc. Amer. Control Conf. (ACC), 2010*, 2010, pp. 5118–5123, IEEE.
- [6] M. Gwerder and J. Toedtli, "Predictive control for integrated room automation," in *Proc. 8th REHVAVorld Congress Clima*, 2005.
- [7] M. Maasoumy, A. Pinto, and A. Sangiovanni-Vincentelli, "Model-based hierarchical optimal control design for HVAC systems," in *Proc. Dynamic Syst. Control Conf. (DSCC), 2011*, 2011, ASME.
- [8] 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," *Proc. IEEE*, no. 99, pp. 1–14.
- [9] Cimis, *California Irrigation Management Information System*. [Online]. Available: <http://www.cimis.water.ca.gov>

- [10] Metar, *National Climatic Data Center*. [Online]. Available: <http://www.ncdc.noaa.gov>
- [11] M. Maasoumy, "Modeling and optimal control algorithm design for hvac systems in energy efficient buildings," Master's thesis, EECS Depart., Univ. California, Berkeley, Feb. 2011. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2011/EECS-2011-12.html>
- [12] A. D. Y. Ma, A. Kelman, and F. Borrelli, "Model predictive control of thermal energy storage in building cooling systems," *IEEE Control System Magazine*, pp. 1–65, 2011.
- [13] M. Maasoumy and A. Sangiovanni-Vincentelli, "Optimal control of HVAC systems in the presence of imperfect predictions," in *Proc. Dynamic System Control Conference (DSCC)*, 2012, ASME.
- [14] J. Lofberg, "Yalmip: A toolbox for modeling and optimization in MATLAB," in *Proc. CACSD Conf.*, TaipeiTaiwan2004. [Online]. Available: <http://users.isy.liu.se/johanl/yalmip>

Mehdi Maasoumy is a PhD Candidate at the Mechanical Engineering Department of UC Berkeley,

Berkeley, CA. His research interests include modeling and optimal control of linear, nonlinear and hybrid systems, and energy efficient building control systems. Maasoumy has a M.Sc. degree in mechanical engineering from UC Berkeley. He is a student member of IEEE and ASME.

Alberto Sangiovanni-Vincentelli is the Buttner Chair of the Electrical Engineering and Computer Science Department at UC Berkeley, Berkeley, CA. He is the Cadence and Synopsys Cofounder, and a member of the GM Science and Technology Board, the UTC Technology Advisory Council, the Italian Institute of Technology ExecCommittee, and NAE. He is also the recipient of the Kaufman Award for "pioneering contributions to EDA" and the IEEE/RSE Maxwell Medal "for groundbreaking contributions with exceptional impact on development of electronics."

■ Direct questions and comments about this article to Mehdi Maasoumy, Department of Mechanical Engineering, University of California, Berkeley, CA 94720; Mehdi@me.berkeley.edu.