Lawrence Berkeley National Laboratory

Practical Factors of Envelope Model Setup and Their Effects on the Performance of Model Predictive Control for Building Heating, Ventilating, and Air Conditioning Systems

D. H. Blum[a]
K. Arendt[b]
L. Rivalin[a,c]
M.A. Piette[a]
M. Wetter[a]
C.T. Veje[b]

[a] Lawrence Berkeley National Laboratory
[b] Center for Energy Informatics
[c] Engie Axima

Energy Technologies Area
February, 2019

For citation, please use 10.1016/j.apenergy.2018.11.093

Disclaimer:

Practical Factors of Envelope Model Setup and Their Effects on the Performance of Model Predictive Control for Building Heating, Ventilating, and Air Conditioning Systems

D. H. Blum[a,*], K. Arendt[b], L. Rivalin[a,c], M.A. Piette[a], M. Wetter[a], C.T. Veje[b]

[a] *Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA, USA*
[b] *Center for Energy Informatics, University of Southern Denmark, Campusvej 55, Odense, Denmark*
[c] *Engie Axima, 15 rue Nina Simone, 44000 Nantes, France*

Abstract

Model predictive control (MPC) for buildings is attracting significant attention in research and industry due to its potential to address a number of challenges facing the building industry, including energy cost reduction, grid integration, and occupant connectivity. However, the strategy has not yet been implemented at any scale, largely due to the significant effort required to configure and calibrate the model used in the MPC controller. While many studies have focused on methods to expedite model configuration and improve model accuracy, few have studied the impact a wide range of factors have on the accuracy of the resulting model. In addition, few have continued on to analyze these factors' impact on MPC controller performance in terms of final operating costs. Therefore, this study first identifies the practical factors affecting model setup, specifically focusing on the thermal envelope. The seven that are identified are building design, model structure, model order, data set, data quality, identification algorithm and initial guesses, and software tool-chain. Then, through a large number of trials, it analyzes each factor's influence on model accuracy, focusing on grey-box models for a single zone building envelope. Finally, this study implements a subset of the models identified with these factor variations in HVAC MPC controllers, and tests them in simulation of a representative case that aims to optimally cool a single-zone building with time-varying electricity prices. It is found that a difference of up to 20% in cooling cost for the cases studied can occur between the best performing model and the worst performing model. The primary factors attributing to this were model structure and initial parameter guesses during parameter estimation of the model.

Keywords: model predictive control, building simulation, HVAC, system

*Corresponding author
Email address: dhblum@lbl.gov (D. H. Blum)

identification

1. Introduction

1.1. Background

Model Predictive Control (MPC) of building heating ventilating and air conditioning (HVAC) systems is a control strategy that can help buildings meet many forthcoming challenges, including reducing energy consumption and carbon emissions [1], integrating with electric grid and other district-scale thermal network operation [2], and integrating occupant behavior [3]. It is a strategy that utilizes a model to predict building performance and optimize control, given information about operating conditions, constraints, and objectives. While there are many advantages to MPC, many challenges exist as well. A critical challenge to be addressed is the implementation cost, particularly labor time and required expertise of the implementer [4, 5]. Central to implementing an MPC controller is the setup of the model used for solving the optimal control problem, including generation and calibration. Studies indicate this task could take 70-75% of the implementation effort [6]. Therefore, reducing the time and expertise required for model setup can improve the scalability of MPC.

There are many factors that contribute to model setup, including the available implementation tools, model type, model complexity, interaction with the optimization solver, data availability and quality, calibration method, and the expertise of the implementer. However, despite a large number of case studies indicating that MPC is a viable approach to meeting new control objectives, a literature review presented in the next section reveals there is a lack of studies on how practical issues affecting model setup influence model accuracy and, especially, MPC performance. This information is difficult to infer from the many case studies in which applications and methods are varied simultaneously. Identifying the relationship between variation of a single factor and the outcome of interest, such as model accuracy and final MPC controller performance, enables identifying the extent to which each individual factor affects the outcome. Such study on MPC modeling approaches can help research and industry highlight the most important considerations, validate robust approaches, and identify needs for further study.

1.2. Objectives

There are four main objectives of this paper:

1. Review previous work that identifies and studies the various factors affecting model setup and their influence on model accuracy and MPC performance. Contribute to the literature by providing a comprehensive summary.
2. Study the effect of the identified factors on model accuracy through a number of simulation trials that apply ranges of each factor independently. Contribute to the literature by presenting the relative impact of each factor.

3. Study the effect of model accuracy on final MPC controller performance through a number of simulation trials that use models setup in the previously described accuracy study. Contribute to the literature by presenting the relative performance of each model in terms of cost savings relative to a conventional controller in the presence of time-varying energy prices, a common use-case for MPC.
4. Perform and document these studies using publicly available models, data, and approaches so that additional factors and variations may be studied and subsequently compared in the future.

To complete the objectives, the paper is structured as follows. Section 2 presents a literature review and summary of practical factors that affect model setup. Section 3 presents the methods used to study the effect of these factors on model accuracy and resulting MPC controller performance. Section 4 presents the results of model accuracy impacts, while Section 5 presents the results of MPC controller performance impacts. Section 6 provides a general discussion of the results and how they lead to future work. Finally, Section 7 concludes the paper with a summary and broader takeaways.

2. Practical factors of model setup

Central to the development of models for MPC is the choice of model structure. Many have categorized models into three broad categories, which indicate the level of physics represented in the underlying model equations. From most physics to least, these are white-box, grey-box, and black-box. Readers are referred to [7, 8] for more detail, however, generally speaking, white-box models are building simulation models, black-box models are identified from building data, and grey-box, or hybrid, are simplified physical models identified using building data. The black-box and grey-box models will be referred to here as MPC-oriented models. The choice of model has profound implications on the available control optimization algorithm and process for calibration. Controllers using white-box models typically require global, numerical optimization algorithms, while those using MPC-oriented models may also allow for local, gradient-based algorithms, which can be more efficient [9]. White-box model parameters and equations need to be tuned manually or using auto-calibration techniques for large systems and are often only suitable for a single building, while MPC-oriented models can use system identification techniques that could be adapted for use on multiple buildings. In [4], it is argued that white-box models are already developed as part of the building design process and can be reused for control, commissioning, and fault detection and diagnosis. In addition, advancements in simulation-based optimization techniques could make the model calibration and control optimization problems more computationally tractable. A number of studies have used white-box building energy simulation programs for MPC, including EnergyPlus [10] and TRNSYS [11], and their use in the future may grow. However, after a review of building modeling for use during operation, [12] still argue that the practical use of a white-box model during operation is limited by the effort required to build the model and calibrate the large number of parameters. The majority of MPC studies to date choose to identify MPC-oriented models, citing the efficiency of available control optimization algorithms and potential adaptability to multiple buildings and applications [7, 8], leading to a decrease in development and implementation costs. Despite their initial attractiveness, identifying MPC-oriented models comes with many challenges.

The first of these challenges is associated with the complexity of the model, including the order, or number of states, and lumping. Citing that lower order, yet accurate, models decrease computational requirements for optimization algorithms and large-scale simulation, many researchers [13, 14] have proposed and studied model-order-reduction techniques for buildings. Most recently, [15] studied the influence of linear time invariant (LTI) state space (SS) model order reduction on performance of the MPC controller, including resulting heating cost, thermal comfort, and computation time. The study found that higher order models greatly improved thermal comfort with slight decreases in heating cost and minimal increase in computation time, with a dense solver approach, indicating their value for an MPC controller. Additionally, higher order models were needed for buildings with heavier construction. While the study demonstrates the controller value and computational practicality of higher order models, it does not factor in the influence of model calibration. Specifically, it, and [13, 14], assumed that the underlying white-box model used for the initial linearization was an accurate representation of the physical building, with parameters and equations calibrated a priori. This white-box model development and calibration may be difficult and costly to ensure.

The second challenge of identifying MPC-oriented models is calibrating them to real building operation. The calibration accuracy highly depends on the training data set chosen, which must excite the model dynamics individually and across the range of operation [8] as well as include the presence of unmeasured disturbances such as internal heat gains [16]. These requirements are often difficult to meet during typical building operation, which follows highly repetitive and specific schedules to provide services to occupants. Therefore, [7] suggests the use of white-box models to create sufficient identification data sets. However, as indicated before, this again requires that the white-box model used for identification is an accurate representation of the physical building, which may be difficult and costly to ensure. Additionally, the model order must be sufficient to capture the dynamics of the building. For the commonly used grey-box resistance-capacitance (RC) models, this includes the number of resistors and capacitors [17, 18], while in a discrete linear transfer function model it can be the number of historic states considered [19].

A third challenge of identifying MPC-oriented models is the quality and accessibility of the data used for calibration. Issues occur at all stages of data collection; sensing, transmitting, and storing. Sensors may fail to report values, be improperly positioned or calibrated, have improper resolution, or be subject to anomalies, such as shading, animal activities, and electrical fields [20, 21]. Communication network or power failure may prevent the transmission of data. Database sizing errors, management errors, and software updates may cause stored data to be lost [22]. From these issues, the final impact on model calibration can be separated into three causes; missing data, noisy data, and biased data. In [23], missing data is classified into three categories depending on whether or not there is a dependency of the missing data on specific data values. Many missing values leads to an inability to use a model for system identification and control optimization, as model simulation requires boundary condition inputs and system identification requires sufficient measurements to compare with model outputs. Data collected through building automation systems are subject to multiple points of failure, which may occur at different points in time. For these reasons, it is common practice to impute missing values, especially when the data does not have chaotic variability between samples. For example, [24] compared multiple imputation strategies for filling missing outside air temperature data. Noise is an undesired modification that the signal may suffer during data acquisition, transmission, or even storage and can be harder to distinguish than missing data. Short-period noise may be caused by internal elements of the sensor, such as by electrical malfunctions or heat generation by resistors and transistors, by external electromagnetic fields, or by other interferences [25]. Long-period noise is mostly generated by the position of the sensor. One example is if an illuminance sensor is influenced by a shadow at particular times every day and another example is if an indoor temperature sensor is located next to a window or HVAC diffuser. A final issue with data is its availability in terms of having the correct sensors in the right locations to take key measurements for the MPC model and final control. An example of this would be flow measurements for air being delivered to each thermal zone to account for the amount of heating or cooling delivered by the HVAC system to the specific zone.

A fourth challenge of identifying MPC-oriented models is the choice and implementation of identification algorithm. Black-box models can be trained with system-identification or other data-driven techniques, see for example [7, 8]. They can be, though, over-fitted to training data and could potentially provide unphysical predictions. Grey-box models are arguably less prone to over-fitting due to the presence of fundamental physics, though the grey-box parameter estimation problem is often non-convex, prompting some caution when accepting a given solution without a global search space. Grey-box model parameters have been trained with non-linear programming (NLP) [17, 26, 18], global optimization [27, 28, 29], agent-based optimization [30], linear regression [31, 32, 19], maximum likelihood estimation (MLE) [33], online estimation [34], and MPC relevant identification (MRI) [21].

A fifth challenge of identifying MPC-oriented models is the software used to implement the model, identification algorithm, and final MPC controller, which also includes a control optimization algorithm, state estimator, and general data handling scripts. While most studies described previously fail to make their software implementation available for others to use, a number of toolboxes have been made to aid others in the model identification process. [31, 35] rely on MATLAB [36], [29] uses the open-source Functional Mock-up Interface (FMI)

Table 1: Summary of factors influencing model setup for MPC in buildings

Factor	Examples
Building Design	Envelope construction, HVAC system design
Model Structure	White-, grey-, black-box
Model Order	Number of states in RC network
Data Set	Training period length and data variability
Data Quality	Noisy or missing measurement data
Identification Algorithm	Batch optimization, online estimation
Software Tool-Chain	MATLAB, Python, R, Modelica

Standard [37] with Dakota optimization toolkit [38], and [39] uses the Python NumPy and SciPy packages [40, 41]. In [17, 42], the open-source Modelica modeling specification [43] and JModelica Modelica compiler and optimization toolkit [44] are utilized.

In summary, we chose to categorize the challenges of generating a model for MPC into the factors defined in Table 1.

With this number of factors facing an implementer of MPC, and with each factor largely addressed in individual studies for specific cases, it is difficult to discern the relative impact of each on final model performance, and ultimately MPC controller performance. There are, however, a few studies that have looked at the relative effect of these factors on MPC-oriented model performance. In [26], the MPC performance was analyzed for concrete core activation systems using 2nd and 4th-order grey-box models and five different training data sets, ranging from inputs created for effective parameter estimation to realistic inputs. MPC performance was evaluated based on thermal discomfort and energy consumption over a one-year simulation using a TRNSYS [45] emulation model. Parameter estimation was performed using a MATLAB interface to the ACADO Toolkit [46], which performs parameter estimation through non-linear programming. The study found that a second-order model excluding solar and internal heat gains provides similar MPC performance to a fourth-order model including solar and internal heat gains if a proper model error correction scheme is applied, a similar idea to modeling unmeasured disturbances in [16]. In [33], the effects of data set length, season, RC model order, noise, and measured inputs on two single-zone buildings (insulated and uninsulated), was studied with measurements of the test buildings emulated using the IDEAS Modelica library [47] and parameter estimation performed using MLE in R [48]. The study found that a 4th-order RC model showed sufficient accuracy to be used as control for both buildings. The addition of heat flux to building elements as an observation variable, to internal air temperature, significantly reduced the uncertainty in the estimated parameters. Data sets of one, two, and four weeks showed small improvements on model accuracy. Noise in measurements generally increased the uncertainty of estimated parameters, though only biased noise, as opposed to unbiased noise, impacted the model accuracy. Final MPC control performance was not evaluated for any of the cases. [18] tested the identification of first, second, and third order RC models using interior-point non-linear optimization in the MATLAB optimization toolbox on three real case study buildings of different size, type, and heating system. The study found that a 4R2C produced acceptable model accuracy using zone-average indoor temperature, though final MPC performance was not evaluated.

3. Methods

3.1. Software tool-chain

The analysis in this paper uses Modelica [43] and FMI [37] in order to ease the development of models, properly initialize states during MPC performance simulation, and compare gradient-based and derivative-free identification algorithms on the same model implementation. In particular, the JModelica [44] software is used for compiling models into Functional Mock-up Units (FMUs), simulating the FMUs, and solving the dynamic optimization problems for parameter estimation and optimal control, while ModestPy [49] is used to solve the parameter estimation problem with derivative-free optimization methods and FMU simulation. More details on the parameter estimation algorithms are given in Section 3.4. The analysis is conducted using the open-source MPC framework, MPCPy, described in [42], commit `e31ea84`, which implements the required modeling, simulation, and optimization tools described above with common interfaces, enabling the setup and execution of the many test cases considered.

3.2. Building design

The building envelopes used in this study are those described in the ASHRAE / ANSI BESTEST [50] methodology as implemented in the Modelica Buildings Library v. 5.0.1. [51], referred to hereafter as the emulation models. We use the BESTEST building descriptions because they serve as a benchmark model for building energy simulation for which the dimensions and construction are publicly available and implementations are available in several building simulation environments. These envelopes consist of a single zone with a window on the south facade and a constant infiltration mass flow rate. These single-zone benchmark models also prevent inclusion of additional confounding factors associated with multi-zone modeling. Our tests analyze two variants of the BESTEST envelope, the lightweight Case 600 (LW) and heavyweight Case 900 (HW). The difference between the LW and HW cases is the construction. The exterior walls and roof of the LW case are constructed of plaster board and fiberglass insulation, while those of the HW case are constructed of concrete block and foam insulation. The floor of the LW case is timber construction, while the floor of the HW case is concrete slab. The EPW weather file DRYCOLD.epw, also publicly available as part of the BESTEST methodology, is used for all tests. The data used for the trials presented later in this study are shown in Figure 1. For more information about the BESTEST building description, see [50].

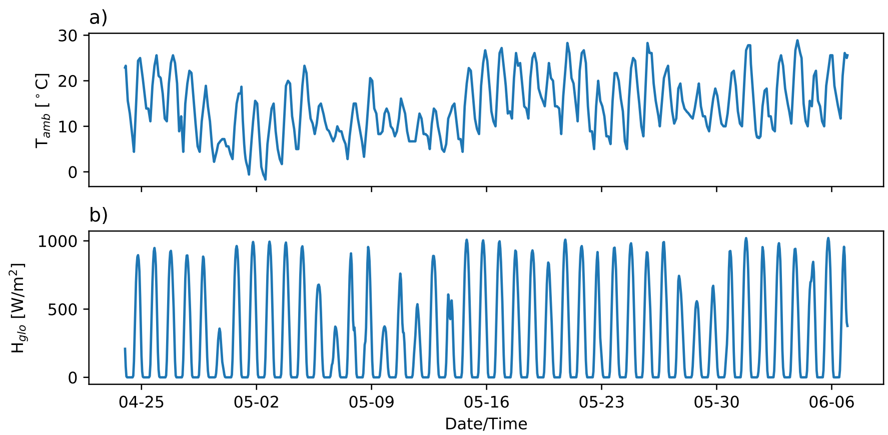

Figure 1: (a) Outdoor temperature T_{amb} and (b) horizontal global irradiance H_{glo} used in the study time period from the "DRYCOLD.epw".

In addition to the envelope, our emulation models utilize a simple HVAC system, consisting of a proportional dual-setpoint feedback controller that controls room temperature according heating and cooling setpoints, heater with constant efficiency of 0.99, and cooler with a constant coefficient of performance (COP) of 3.0. The heater and cooler add or extract energy from the room air mass, which is assumed to be well-mixed. As this study focuses on the identification of the room model, the HVAC model is kept simple in order to not compound the complexity of the whole building model. In this fashion, the effects of the room model and HVAC system model on MPC performance is separated.

The operation of the building represents a typical office, with the following load assumptions coming from [52]. The operating hours of the building are assumed to be from 8 AM to 6 PM. The heat gains for office equipment is assumed to be 5.4 W/m^2 with 30% radiative and six workstations per 92.9 m^2, which corresponds to a light office density. The heat gain from lighting is assumed to be 11.8 W/m^2 with 58% radiative, corresponding to an open office plan. The heat gain from people include both sensible and latent, affecting room air temperature and humidity respectively. They are assumed to be 73.3 W sensible with 60% radiative and 58.6 W latent, corresponding to moderate office work. With six workstations per 92.9 m^2 and a zone floor area of 48 m^2, there are three people assumed to be in the zone. In total, the radiative heat gain is 11.2 W/m^2, the convective heat gain is 10.6 W/m^2, and the latent heat gain is 3.67 W/m^2. These heat gains are active during the occupied period and zero during unoccupied periods. The heating and cooling occupied and unoccupied setpoint temperatures are (21 °C, 16 °C) and (24 °C, 29 °C) respectively.

3.3. Model structure and order

Three models are considered to analyze the effect of model complexity, each a variant of an RC network model commonly found in literature described previously. The models are presented in Figure 2 and gradually increase the number of capacitive and resistive components in order to account for thermal mass and resistances within the room, walls, floor, and roof with increasing accuracy. The models are R1C1, R3C3, and R5C4. For the R5C4 models, the extra resistance is in parallel with the wall in order to account for infiltration and window conduction gains separately. The inputs to the models are outside air temperature T_{amb} [K], global horizontal irradiance H_{glo} [W/m^2], radiative and convective internal heat gains $q_{occ,r}$, $q_{occ,c}$ [W/m^2], and HVAC heating (q_h) and cooling (q_c) power $q_{hvac} = q_h - q_c$ [W]. The output of the models are the temperature of the capacitance representing the zone air T_i [K]. For the R1C1 model, this is the only capacitance. The model parameters to be estimated are each of the resistance and capacitance values of the model, as well as gain parameters that scale the total global horizontal irradiance incident on the floor, α, and exterior walls (for R3C3 and R5C4), α_e. These solar irradiance parameters have the units m^2, although represent a number of factors, including sun-exposed areas, absorptivity and transmissivity factors, and orientations. The wall and floor areas are assumed known, meaning the resistance and capacitance parameters to be estimated have the units of m^{2}K/W and J/(m^{2}K) respectively, except for the capacitance representing the air volume, which has units of J/K.

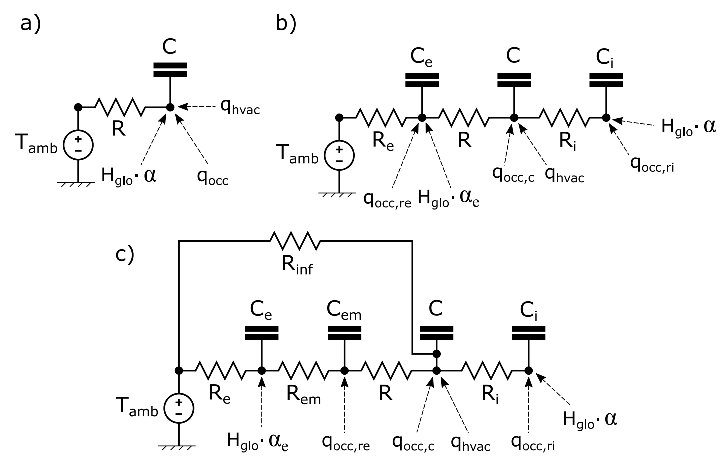

Figure 2: RC thermal network models considered in the study: a) R1C1, b) R3C3, c) R5C4.

3.4. Identification algorithm

A common formulation of the model identification problem as an optimization problem is defined by Equations 1-4:

$$\min_{\boldsymbol{\theta}\in\mathbb{R}^n} J = \int_{t_s}^{t_f} (y-\hat{y})^2 dt \tag{1}$$

s.t.

$$f(\dot{\boldsymbol{x}}, \mathbf{x}, \mathbf{u}, \boldsymbol{\omega}, \boldsymbol{\theta}) = 0 \tag{2}$$

$$\mathbf{x}(t_s) = \mathbf{x}_0 \tag{3}$$

$$\theta_i^l \leq \theta_i \leq \theta_i^u \quad \forall i \in \{1, ..., n\} \tag{4}$$

Here, the model dynamics are represented by f, which is twice-differentiable, with parameter vector $\boldsymbol{\theta}$ of n real-valued numbers. Each individual parameter is indexed by $i \in \{1, ..., n\}$. The state vector is $\mathbf{x}$ and its derivative with respect to time is $\dot{\boldsymbol{x}}$. The control vector is $\mathbf{u}$, which contains signals for the heater and cooler. The disturbances of weather and internal loads are contained in $\boldsymbol{\omega}$. The minimum and maximum values for any given parameter are θ_i^l and θ_i^u respectively. The period of training is from a start time, t_s, to a final time, t_f. All variables are real-valued functions of time except $\boldsymbol{\theta}$. The model output, y, in this study is the zone air temperature, with its measured value being $\hat{y}$.

Three identification algorithms, which represent a common subset of the approaches found in the literature, are analyzed to solve this problem along with two sets of parameters for each model, represented by an initial guess, minimum, and maximum. The first identification algorithm is implemented using JModelica, which uses the model and optimization equations defined in Optimica [53] with a direct collocation [54] discretization method to setup an NLP, using CasADi for algorithmic differentiation [55], which is then solved by IPOPT [56] with the MA27 linear solver [57]. The objective of the NLP is to minimize the sum of the squared errors between measured and modeled air temperature at each discretization point, subject to the minimum and maximum constraints on parameter values. In this study, the measured values are the emulated values. Overall, this method represents an NLP approach to find the local minimum of the parameter estimation problem, which has been used previously in [17, 26, 18], and is referred to from here on as the NLP method. Due to the number of simulation cases to be processed in this study, and after experienced was gained in determining how long it took a solution to converge, the maximum CPU time per optimization was constrained to 150s.

The second identification algorithm is implemented using ModestPy [49], which uses an FMU representation of the model to simulate the model within a hybrid Genetic Algorithm-Pattern Search approach. The algorithm begins by using a Genetic Algorithm (GA) to explore the solution space (global search), and uses the Generalized Pattern Search algorithm to converge to a minimum (local search). The implemented hybrid approach is similar to the Particle Swarm optimization combined with Hooke-Jeeves algorithm available in GenOpt [58]. Note that the FMU is generated from the same Modelica equations used to generate the Optimica code for the gradient-based JModelica approach in order to ensure the same dynamic model is used. Overall, the method represents a global optimization approach to parameter estimation, which has been used previously in [27, 28, 29]. This method is referred to from here on as the GA+PS method.

The third algorithm considered results from the acknowledgement that the NLP approach is a local search method and the parameter estimation problem may be non-convex. This approach applies a global start algorithm to the NLP approach described previously. Here, Latin Hypercube Sampling is used to identify a number of initial guesses within the allowed parameter space. Then, the NLP approach is run for each of the initial guesses and the results with the minimum objective value is chosen as the final result. The chosen number of initial guesses for this study is 20. This method is referred to from here on as the NLP+LHS method.

Finally, for each identification algorithm, two parameter sets are considered. One parameter set has wide minimum and maximum values, representing a case where the implementer does not have a good guess as to what the parameters should be. A second set has minimum and maximum values set according to a range that limits the parameters to physically plausible values, representing a case where the implementer has a good idea of what the parameters should be from experience or access to building information. Both parameter sets are shown in Table 2

All computations that were carried out for which computational time is reported were performed on a Linux virtual machine (Ubuntu 16.04), hosted on a workstation running on two Intel E5-2630 v4 processors (20 cores in total), 192 GB RAM, with SSD disks for system data (host and guests) and HDD disks for user data (guests). The workstation, however, hosts multiple virtual machines, and the resources are shared and scaled dynamically. Each simulation case was analyzed on a single core, but several cases were run simultaneously. Since the resources of the computer are shared dynamically depending on other simultaneous users, the reported CPU time for each algorithm should be taken as a general means of comparison, rather than exact expectations of computation time with the stated resources. With thousands of cases run, the authors expect that any random fluctuation in computation performance occurred in all cases and do not bias the results towards any.

3.5. Training data length

The analysis of training data length is conducted for a number of consecutive days from May 15 to May 30. For each day, the training data are collected for a different number of preceding days, n = 1, 2, 3, 5, 7, 10, 14, 21. Also, for each day and for each number of preceding days, two validation periods are assessed: 1) one following day and 2) seven following days.

3.6. Noisy data

In this study, noise is added to sensor data for the outside air temperature, solar irradiance, and zone air temperature during training periods, then the noisy data is used to estimate the models of the parameter, and then the model is used to produce new data for validation during the validation period. No noise

Table 2: Parameter sets (wide bounds, narrow bounds)

Model	Parameter	Unit	Wide bounds		Narrow bounds	
			Lower	Upper	Lower	Upper
R1C1	R	m^2K/W	0.001	100	1	10
	C	J/K	1	5e7	1e4	15e6
	α	m^2	0	100	0	10
R3C3	R	m^2K/W	0.001	100	0.1	3
	C	J/K	1	5e7	1e4	5e6
	α	m^2	0	100	0	10
	α_e	m^2	0	100	0	10
	Ri	m^2K/W	0.001	100	0.01	3
	Ci	$J/(m^2K)$	1	1e7	200	4e5
	Re	m^2K/W	0.001	100	0.1	10
	Ce	$J/(m^2K)$	1	1e7	200	4e5
R5C4	R	m^2K/W	0.001	100	0.1	3
	C	J/K	1	5e7	1e5	5e6
	α	m^2	0	100	0	10
	α_e	m^2	0	100	0	10
	Ri	m^2K/W	0.001	100	0.01	3
	Ci	$J/(m^2K)$	1	1e7	200	4e5
	Re	m^2K/W	0.001	100	0.1	5
	Ce	$J/(m^2K)$	1	1e7	200	2e5
	Rem	m^2K/W	0.001	100	0.1	5
	Cem	$J/(m^2K)$	1	1e7	200	2e5
	Rinf	m^2K/W	0.001	100	0.1	10

is added to validation data. To add the noise, a random value is drawn from a Gaussian distribution with a mean of zero and particular standard deviation [59] and added to particular data points of the chosen sensor according to the period of the noise. The value is drawn independently for each noisy data point. The study is carried out for various standard deviations, representing noise magnitude, and a number of noise periods for each of the three variables. For noise applied to temperature data, standard deviations ranging from 0 (clean) to 10 °C are applied [14], while for solar irradiance data, standard deviations ranging from 0 to 500 W/m^2 are applied. A limit is applied to keep solar irradiance values greater than or equal to zero after noise application. As the sensor period is assumed to be 1 hour, noise periods varying from 1 hour to 1 day are applied, with the noise lasting for the full hour. Note that noise periods higher than the sensor period, which may represent electromagnetic interference, are not detected at a higher rate than the sensor period. Noise periods on the order of one day can represent daily occurring phenomenon such as sudden shadows or daylight on a sensor. Finally, since the noise values are generated randomly, each case is tested 10 times to obtain a distribution of results.

3.7. Missing data

Two aspects of missing data are studied for each of the three sensors for outside air temperature, solar irradiation, and zone air temperature. A first aspect has to do with the amount of missing data, and is characterized by two variables; the percent of total missing data and the length of each missing data gap. A second aspect has to do with strategies for filling in the missing data for use in the parameter estimation problem. Four strategies were tested. First, the strategy of replacing the missing data with zeros. While this is likely not a strategy to be chosen intentionally by an implementer, it represents a possible default configuration on data collection and management systems. The second strategy is filling missing data with the last known value, also known as forward padding. This is a simple strategy that fills data with a more reasonable value than zero. The third strategy is linear interpolation, where the last and next valid values are linearly interpolated to fill the missing data in between. Finally, the fourth strategy is a more advanced technique that replaces missing data with the average of values from the previous two days and next two days at the same time steps. If some of this data are also not available, the average of whatever data are available is taken. If none of this data are available for averaging, the next valid value is used. Since the actual data points selected to be missing are chosen randomly according to the percent of total missing data and length of missing data gaps, each case is tested 10 times to obtain a distribution of results.

3.8. Control optimization

A simulation of the MPC controller is used to evaluate the influence of model accuracy on performance for a range of cases as presented in the previous sections. In this way, the overall effect of model accuracy on MPC performance can be evaluated as well as the relative effects of the previously described influential factors. MPC provides benefits when there are incentives to shift load in time, such as to improve HVAC equipment efficiency according to the changing outside air temperature or reduce electricity bill costs according to dynamic pricing schemes. In this study, since a detailed model of the HVAC system is not included, we implement a simple dynamic electricity price, where the price is five times higher from 2-6pm than all other hours. Such a price increase is reasonable based on realized differentials in both wholesale market prices (see for example the PJM Interconnection hourly real-time locational marginal price from September 9-23, 2018 [60] as well as current commercial customer Time-Of-Use (TOU) retail tariffs (see for example the E19 Schedule for PG&E in California [61]). This will incentivize a shift in HVAC load away from these hours and utilization of thermal mass present in the structure of the building. Therefore, the control optimization problem is defined by Equations 5-11:

$$\min_{u_c,u_h} J = \int_{t_s}^{t_f} \pi_e P dt \tag{5}$$

s.t.

$$f(\dot{\boldsymbol{x}}, \mathbf{x}, \mathbf{u}, \boldsymbol{\omega}, \boldsymbol{\theta}) = 0 \tag{6}$$

$$\mathbf{x}(t_s) = \mathbf{x}_0 \tag{7}$$

$$P = u_c \frac{Q_c}{COP} + u_h \frac{Q_h}{\eta} \tag{8}$$

$$T_h \leq T_i \leq T_c \tag{9}$$

$$0 \leq u_c \leq 1 \tag{10}$$

$$0 \leq u_h \leq 1 \tag{11}$$

Here, u_c is the cooling control signal and u_h is the heating control signal, contained in $\mathbf{u}$. The model dynamics for a given model are represented by f, which is twice-differentiable, with parameter vector $\boldsymbol{\theta}$. The model dynamics in Equation 6 are the same as Equation 2. The state vector is $\mathbf{x}$, within which is T_i, the zone internal air temperature. The disturbances of weather and internal loads are contained in $\boldsymbol{\omega}$. The upper and lower temperature limits are set to the cooling and heating setpoints respectively T_c and T_h. P is the HVAC power consumption, Q_c and Q_h are the maximum cooling and heating capacities and COP and η are the cooling coefficient of performance and heating efficiency respectively. Finally, π_e is the electricity price. All variables are real-valued functions of time except θ_i, COP, and η, which are constant for the problem time horizon. The time period t_s to t_f is the problem time horizon.

For each case, the model and set of identified parameters are used in the optimization, represented by f with $\boldsymbol{\theta}$ in Equation 6, to produce an optimal control solution for a given time horizon. The solution is then implemented in the emulator model by setting the heating and cooling temperature setpoints equal to the temperature trajectory in the optimal solution for the control step of the simulation. At the end of the control step, a state estimator is used to estimate the states of control model before the process is repeated for the next control step and until the end of the simulation time. For each control step optimization, the result of the previous control step optimization is used as an initial guess. In practice, an envelope around the optimal temperature trajectory will likely be required for supervisory control implementation. However, in this simulated case study, good results are obtained with zero deadband and the tuning of this deadband would add additional variability in the final performance of the MPC controller. While various methods exist for state estimation, a simple method was used for this study, where the measured states are set to the corresponding emulation model measurements and the unmeasured states are set to the values of the corresponding control step from the previous control step optimization. The only measured state for this study is the zone air temperature, and the unmeasured states correspond to the temperatures of extra capacitances representing internal and wall thermal mass. Finally, for this study, the optimization horizon is 24 hours, the control step is one hour, and each simulation occurs for one week. The optimization problem is implemented and solved using the open-source platform for MPC in buildings MPCPy [42].

It is worth a note on the choice of optimal control implementation to be as temperature setpoints for the HVAC controller, rather than open-loop injection of optimal heating and cooling signals. The former we call the supervisory method and the latter the open-loop method. From a research perspective, the performance of the open-loop method could be evaluated by the resulting objective cost and the magnitude of constraint violations in the emulator model. In this case, this would be in the form of thermal comfort violations, likely requiring both maximum violation over the time of simulation in [K] and total discomfort in [K-h]. In total, this leads to performance evaluation requiring the use of three metrics, for which the relative importance can be disputed. In contrast, the supervisory method can have an added feature that the conventional temperature setpoints are used as absolute boundaries of setpoint implementation, ensuring no more thermal discomfort than the conventional controller. In this case, under the assumption of sufficient capacity, any model error that would cause thermal comfort violation in the open-loop method instead leads to increased objective cost (in the form of unanticipated additional heating or cooling) to maintain thermal comfort in the supervisory case. Therefore, the supervisory method collapses the three performance metrics required in the open-loop method into a single performance metric, the objective cost. This makes it easier to evaluate the performance of a given case relative to others. Additionally, from a practical perspective, it is often easier to implement the optimal control as supervisory setpoints to a BMS system that already exposes available setpoints, rather than directly control a number of local actuators. Finally, by applying the absolute boundary conditions to the supervisory setpoints, building occupants and facility operators can be assured that the conventional service of the building will not be compromised.

A final important consideration is how to initialize the optimization for each control step, including the handling of initial time constraint violation. In some cases, the resulting initial state of the zone air in the model may violate constraints, leading to an infeasible optimization problem. To handle this, the constraints of the optimization are expanded by a small envelope around this initial state.

4. Impact of factors on model accuracy

This section presents the results of analyzing the impact of the factors presented in Table 1 of Section 2 on model accuracy, according to the methods described in Section 3. The analysis is split into two main components. The first component studies the impacts of building design, model order, identification algorithm, and training data length, while the second component uses a subset of the first component cases to study the impact of noisy and missing data. The accuracy indicator used in this study is inside air temperature Root Mean Square Error (RMSE), which is the most common indicator found in the

Table 3: Summary of cases for the Component 1 analysis

Category	Cases	Description
Building Design	2	Lightweight (LW), Heavyweight (HW)
Model Order	3	R1C1, R3C3, R5C4
Identification Algorithm	3	NLP, NLP+LHS, GA+PS
Parameter Bounds	2	Wide Bounds, Narrow Bounds
Training Data Length	8	1, 2, 3, 5, 7, 10, 14, 21 Days
Beginning of Validation	15	May 15-30

literature and is a single variable that can be used to represent the accuracy of each model when comparing MPC performance in Section 5.

4.1. Building design, model order, identification algorithm, and training data length

Component 1 of the analysis presents the impacts of building design, model order, identification algorithm, and training data length on resulting model accuracy according to the methods described in Sections 3.2-3.5. The variation of cases is summarized in Table 3 and leads to a total of 4320 model setup permutations considered. Due to a significant number of outliers in the results, which result from convergence to local minima during model identification, we focus mainly on the analysis of medians instead of means. However, the spread of data are discussed and presented as well in some cases. The RMSE is calculated for the training period as well as two validation periods of 1-day and 7-day lengths.

First, Figure 3 presents the median RMSE of all cases by building type and model order. We can notice that the control models were more accurate for the heavyweight building than the lightweight building. This may be attributed to the dampening of model errors associated with fast dynamics by the increased thermal mass in the heavyweight building compared to the lightweight building. For both types of buildings, the R3C3 model was most accurate for all three of training, 1-day, and 7-day validation periods. In addition, for the lightweight building, the R5C4 model had lower training error than the R1C1 model, but higher 1-day and 7-day validation error, indicating the tendency of the R5C4 model to be overfitted to training data.

Next, Figures 4-5 present the performance of each identification algorithm and parameter set. Figure 4 shows the distributions of training, 1-day validation, and 7-day validation RMSEs for all cases by identification algorithm and parameter set, while Figure 5 presents the mean computational time for all cases by identification algorithm and parameter set. According to the data in Figure 4, NLP+LHS outperforms or is similar to NLP in all of the analyzed cases. NLP+LHS particularly outperforms NLP for the parameter set with wide bounds, and models with higher orders, with a difference in median RMSE ranging from 0.3 to 1.0 °C depending on the case. This highlights the non-convexity

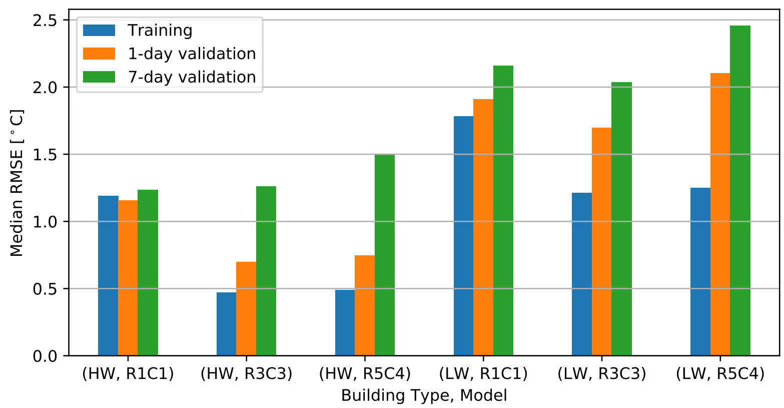

Figure 3: Median RMSE [°C] by building type (HW and LW) and RC model type.

of the parameter estimation problem, particularly for higher-order models. The NLP+LHS also outperforms or is similar to the GA+PS method, again in particular with cases of wide bound parameter sets and models with higher orders, with a difference in median RMSE of up to 1.2 °C. The worse performance of GA+PS on wide bound parameter sets and higher-order models was likely due to the chosen evolution settings. As shown in Figure 5, the computational time of GA+PS is lower for wide bound parameter sets than for narrow bound (conversely to NLP+LHS), meaning that it switches too quickly to the local search. We believe better results could be achieved by extending computational time of the GA+PS, for the GA global search portion, however, due to the significant number of cases to run, and successful results with NLP+LHS, we decided not to further tune GA+PS. In summary, with respect to identification algorithms and parameter sets, when dealing with twice-differentiable models where the directional derivatives with respect to the parameters being estimated are available (as in this study), the NLP+LHS approach provides the best balance between the computational demand and accuracy. Meanwhile, an advantage of the GA+PS algorithm not studied any further here is the ability to be applied to discontinuous functions.

The parameter bounds had a significant effect on the estimation results. The median training and validation errors were lower in all cases with narrow bounds, although the R1C1 model performed similarly with the NLP algorithm for both parameter sets, according to median RMSEs. However, Figure 6 shows that even the R1C1 estimation problem is not trivial to solve, as even with the NLP solver different parameters were obtained depending on the exact case, indicating possible non-convexity of the problem. In the cases of R3C3 and R5C4 model orders, the cases with narrow parameter sets outperformed those with wider parameter sets by differences in median RMSEs of 0.86 and 1.72 °C respectively. In general, the R3C3 model presented the most attractive balance

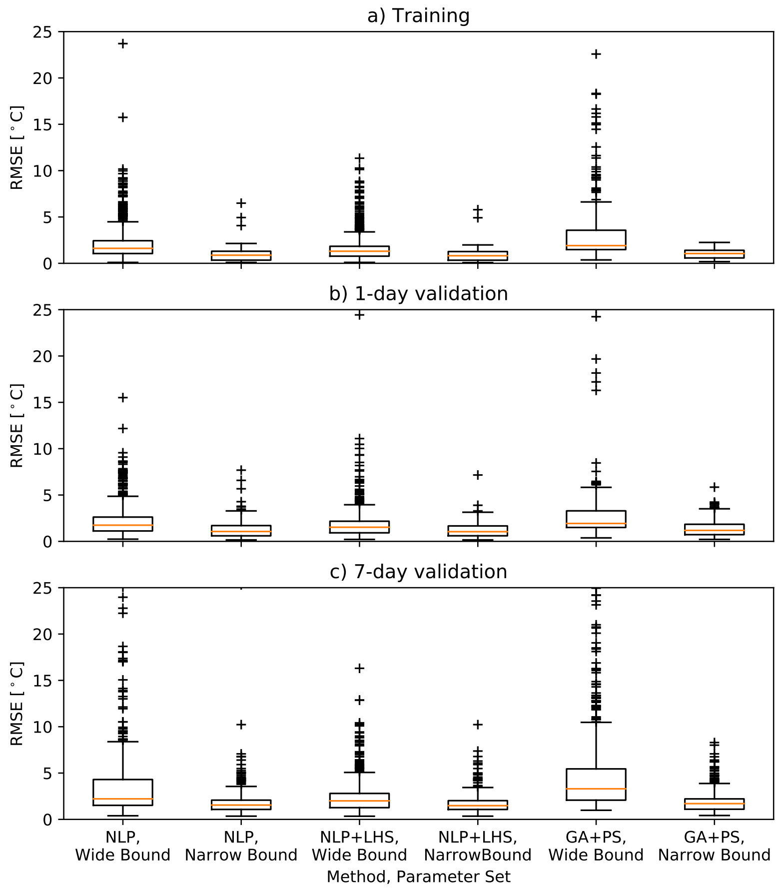

Figure 4: RMSE box (quartiles) and whisker (1st/3rd quartile +/- 1.5 IQR) distribution with outliers by estimation method (NLP, NLP+LHS, GA+PS) and parameter set (wide bounds, narrow bounds).

between estimation difficulty and model accuracy, performing best on both wide and narrow parameter bounds.

In addition, the results indicate that parameter sets impact not only the median model accuracy, but also the spread of the results. In the case of NLP with wide parameter sets, for all validation cases, most of the errors are below 8 °C, but outliers have errors up to 400 °C (not shown in Figure 4, as it is zoomed to 0-25 °C). For GA+PS, most of the errors are below 5 °C, but outliers have errors up to 12,000 °C. The large errors (over 100 °C) yielded by the outliers appear only in the cases with wide parameter bounds, and are due to extremely overestimated solar gains, underestimated thermal mass of the building, and overestimated thermal resistance of external walls. Since the considered RC models do not include separate resistors representing the window, the excessive

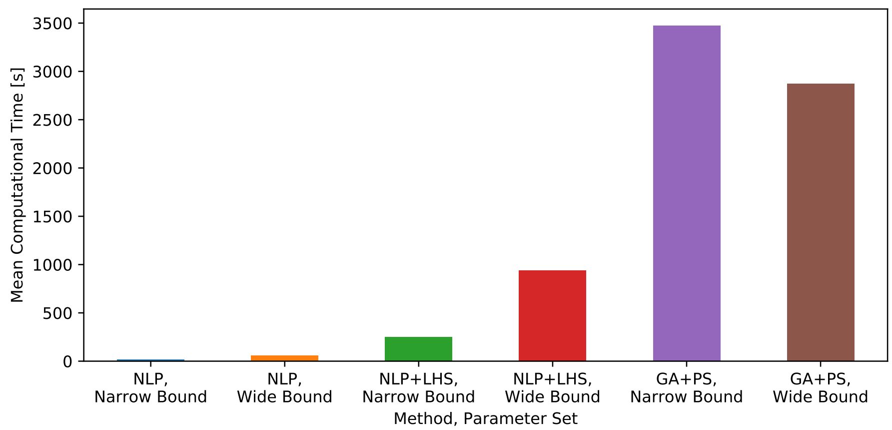

Figure 5: Mean computation time by estimation method (NLP,NLP+LHS, GA+PS) and parameter set (narrow bound, wide bound).

heat gains have no means of dissipating to the ambient environment. In practice, and as mentioned previously, additional tuning of the GA+PS can help rectify the outliers. Meanwhile, NLP+LHS yielded all RMSEs below 13 °C. It is worth noting for the narrow parameter sets that, while there are minimal outliers for training data in the cases using NLP+LHS and GA+PS, there still exists several validation outliers, especially for 7-day validation.

Finally, Figure 7 presents the median RMSE for all cases by training period length. The results indicate that the optimum length of training depends on the considered length of the validation period. In the case of 1-day validation, the best results were achieved with training on just the previous day and continues to decrease as more training days are added. In the case of 7-day validation, the optimum length of the training period was between 3-7 days. Training period RMSE increases with training length because with each additional training day, the parameters are estimated for an average day instead of any one particular day. Similar is true for 1-day validation when conditions from one day to the next tend to be similar. 7-day RMSE decreases at first with training length because with each additional training day, the parameters are estimated over a wider range of conditions, which may be more similar to the validation period. However, the RMSE begins increasing after a certain point for a similar reason as why training period and 1-day validation RMSE increases. The training period data becomes too general for the specific validation period. This data suggests that, when using grey-box models, an adaptive model identification framework, where the model parameters are estimated often using a few days of previous data, would produce the most accurate MPC-oriented models for MPC control.

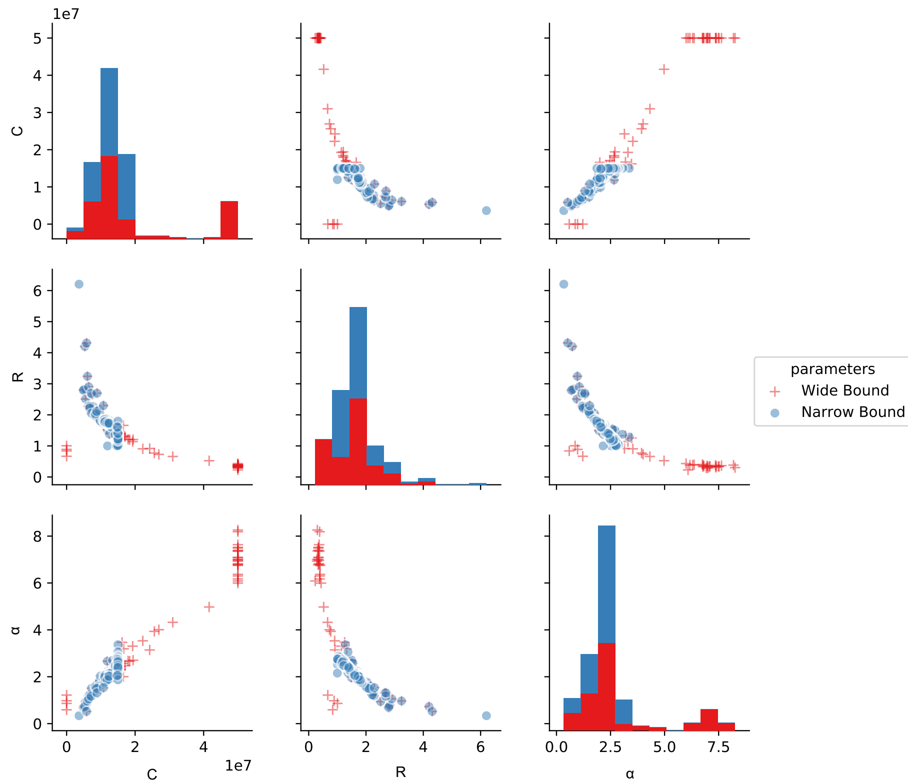

Figure 6: R1C1 parameters estimated for the HW building with NLP on wide parameter bounds and narrow bounds. The figures on the diagonal are histograms for how often the x-axis value of the parameter was estimated. Note that the histogram bars are stacked. The figures on the off-diagonal plot the y-axis parameter value against the x-axis parameter value for each estimation case.

4.2. Noise and missing data

Component 2 of the analysis presents the impacts of noisy and missing data on resulting model accuracy according to the methods described in Sections 3.6-3.7. The variation of cases is summarized in Table 4 and leads to a total of 5760 model setup permutations considered. Each variation was performed on a single case of the Component 1 analysis: the use of the NLP identification algorithm, narrow parameter bounds, seven day training length starting on May 15, and 7-Day validation. With reference to the maginitude of noise tested, we note that in the clean case, the standard deviation for outside dry bulb temperature is 5.51 °C and global horizontal radiation is 362 W/m^2. Zone temperature standard deviation is 3.60 °C for the LW building and 1.44 °C for the HW building.

Figure 8 presents the results for noise applied at various periods and magnitudes for each of the three data variables and building cases. Clean cases

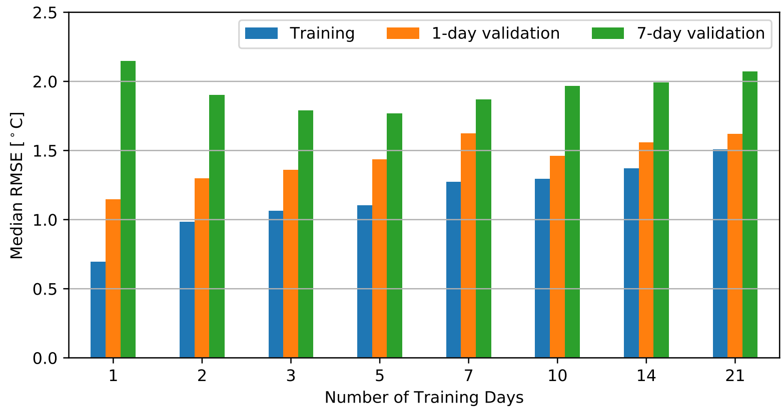

Figure 7: Median RMSE for all cases by training period length.

Table 4: Summary of cases for the Component 2 analysis.

Category	Cases	Description
Building Design	2	Lightweight (LW), Heavyweight (HW)
Model Order	1	R3C3
Affected Data Variables	3	T_{amb}, H_{glo}, T_i
Noise Period	4	1 Hour, 4 Hours, 12 Hours, 1 Day
Noise Magnitude	4	0,2,5,10 °C and 0,100,200,500 W/m^2
Missing Data Percent	4	0%, 25%, 50%, 75%
Missing Data Gap Length	5	1 hour, 3 hours, 10 hours, 15 hours, 1 day
Missing Data Correction	4	Zeros, Interpolate, Last value, Average

without any noise are represented by a grey line in the Figure. Outside dry bulb temperature noise has the least effect on RMSE, especially on the HW building, due to the exterior wall thermal mass acting as a filter. Meanwhile, noise in the solar radiation, due to the window transmitting heat gain directly to the interior mass, and zone air temperature measurement have greater effects on RMSE. We observe that the estimated total thermal mass in the models, by sum of all thermal capacitances, increases as noise standard deviation increases, as the estimation algorithm attenuates the impact of the noisy signal on inside air temperature. This would lead to an overestimation of thermal mass. Finally, we observe that in some cases, adding noise can improve validation RMSE compared to the clean case. Indeed, adding some noise to training data has been used as a technique to avoid over-fitting in data-intensive analytics, such as machine learning [62].

Figure 9 presents the results for missing data studied at varying degrees with different cleaning strategies, as described in Section 3.7. Clean cases without

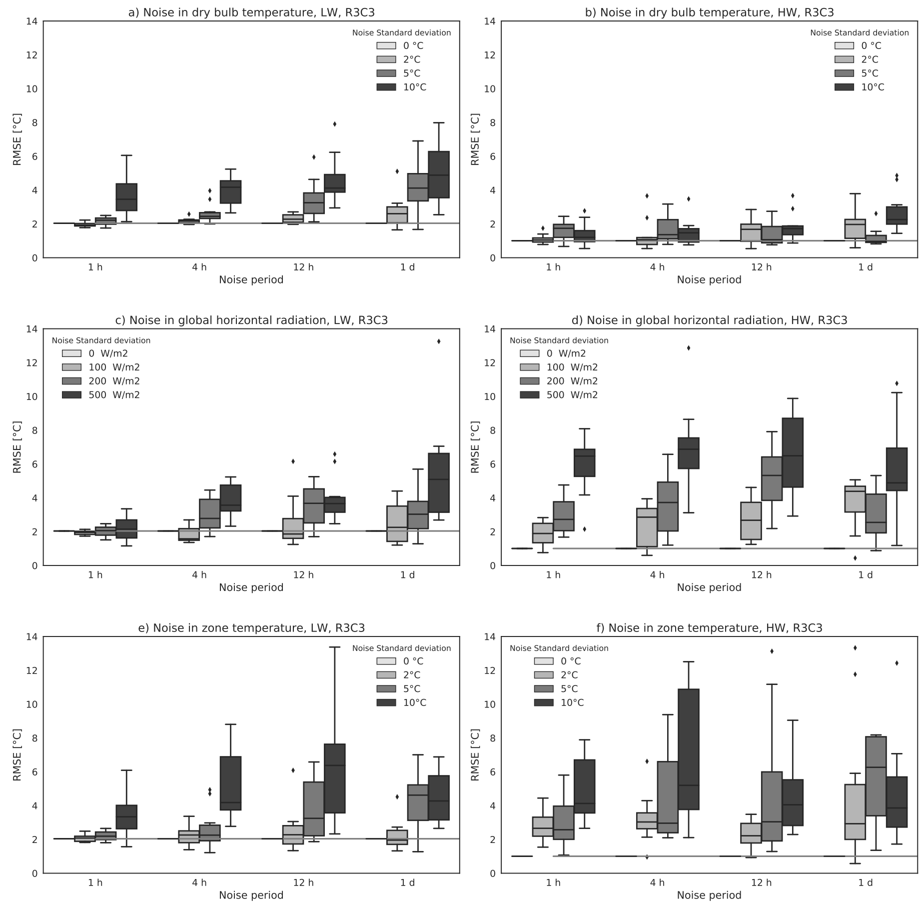

Figure 8: 7-Day validation RMSE boxes (quartiles) and whiskers (1st/3rd quartile +/- 1.5 IQR) distributions with outliers of the impact of gaussian white noise applied to dry bulb temperature (row 1), horizontal radiation (row 2) and zone temperature (row 3) applied to LW (column 1) and HW (column 2) buildings with NLP estimation method, narrow bound parameter set, and an R3C3 model.

any missing data are represented by a grey line in the Figure. It was found that the percent of missing data and cleaning strategy had significantly more impact on the accuracy of the models than length of gaps of missing data. Therefore, results are shown including all gap length cases. From the Figure, it is clear that replacing missing data with a zero is the worst strategy to use, particularly at high rates of missing data and for the zone temperature measurement. While this is not a strategy likely employed on purpose by an implementer, it may be a default setting on data collection and processing tools and is therefore worth ensuring another strategy is used. Overall, the average strategy works the best, while the interpolation and last valid value strategies seem to work reasonably well for all cases, with the interpolation strategy performing slightly better in many of the cases. The average strategy is a particularly good strategy for zone temperature in the LW building, even for more than 50% missing data, due to the consistent day-to-day operation at the cooling or heating setpoints. Meanwhile, for the HW building, which spends more time in free-float, the average strategy does not perform as well due to the variability in day-to-day temperature profiles. The average or interpolation strategies can work well for low percentages of missing data, underlining the importance of keeping missing data to less than 25%.

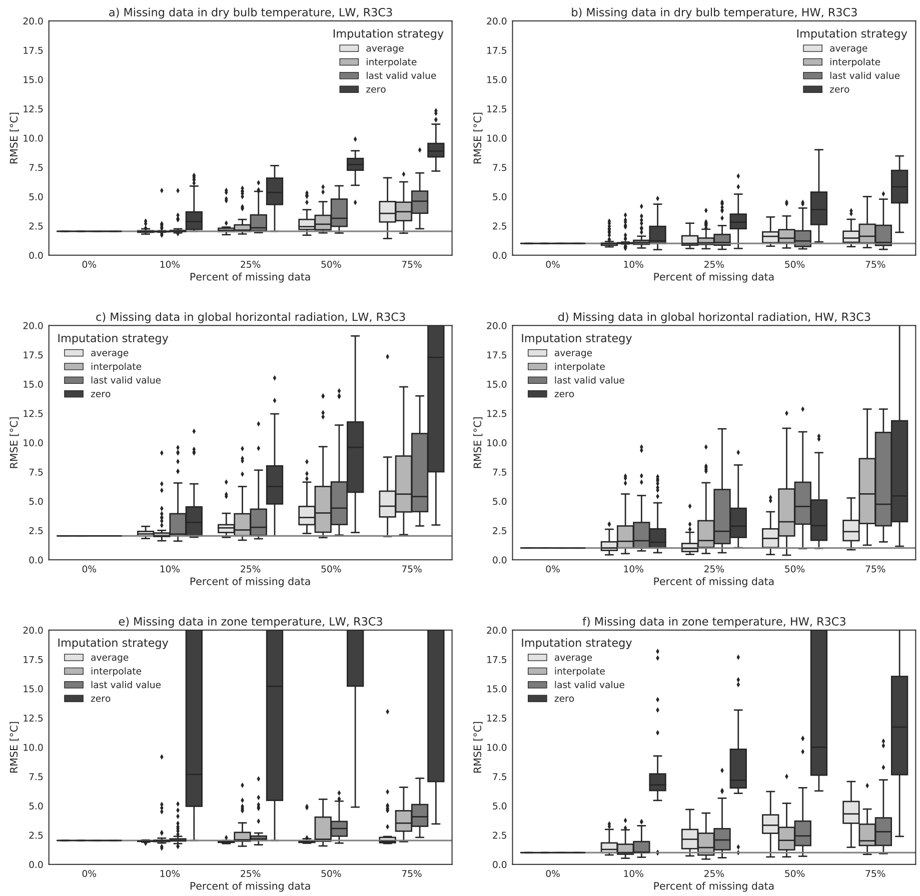

Figure 9: 7-Day validation RMSE boxes (quartiles) and whiskers (1st/3rd quartile +/- 1.5 IQR) distributions with outliers of the impact of missing data applied to dry bulb temperature (row 1), horizontal radiation (row 2) and zone temperature (row 3) applied to LW (column 1) and HW (column 2) buildings with NLP estimation method, narrow bound parameter set, and an R3C3 model.

Table 5: Summary of cases for the controller performance analysis.

Category	Cases	Description
Building Design	2	Lightweight (LW), Heavyweight (HW)
Model Order	3	R1C1, R3C3, R5C4
Identification Algorithm	1	NLP+LHS
Parameter Bounds	2	Wide Bounds, Narrow Bounds
Training Data Length	3	1, 3, 7 Days
Beginning of Validation Date	3	May 15, 20, 25
Electricity Prices	1	Five times higher from 2-6pm

5. Impacts on controller performance

This section analyzes the results of simulating the MPC controller using many of the models that have been identified through the previous studies as described in Section 3.8. In this way, we can compare the performance of the resulting MPC controllers in terms of operating cost, the minimization of which is the objective of the control optimization, as a function of the accuracy of the model that is used as reported during the identification process. Since the simulation occurs for one week for each controller, the 7-day validation RMSE is used as a proxy of model accuracy, as was done in the previous analyses in this paper and often considered in the literature. The cases considered in this analysis are pulled from the Component 1 analysis in this study and are detailed in Table 5. In summary, they are a subset of models that represent a cross section of model setup methods, and do not include the noisy or missing data models. There are 108 total cases represented.

Figure 10 summarizes the results of the analysis. The y-axis is the cost of HVAC operation with the MPC controller as a percent of the conventional feedback controller for the same week of simulation, while the x-axis is the seven-day validation RMSE for the model setup with the same end of training/beginning of validation date. Figure 10(a) presents the results for the LW building, while 10(b) presents the results for the HW building. Each color represents a different starting day of the simulation week (May 15: red, May 20: green, May 25: blue), each shape represents a different model structure (R1C1: circle, R3C3: triangle, R5C4: square), and each filling represents a different parameter bound (wide bounds: white, narrow bounds: filled). Readers are referred to Figure 1 for outside conditions. To promote clarity, training data length is not differentiated. Therefore, there are three of any given symbol on the Figures. Finally, some cases had virtually the same performance, meaning their points overlap and some may be hidden.

The first thing to note in Figure 10 is that greater cost savings are achieved with the HW building than with the LW building. Less thermal mass in the LW case leaves little ability to shift load, and little benefit provided by the MPC controller. Meanwhile, the thermal mass of the HW building enables shifting of load, leading to benefit of the MPC controller to save cost over the conventional

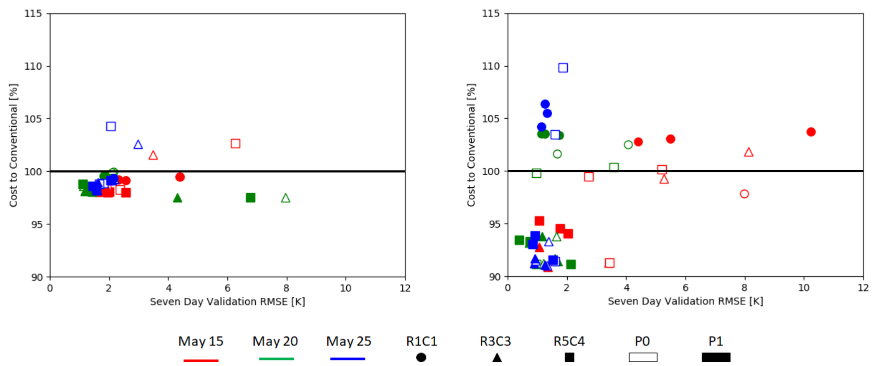

Figure 10: MPC performance vs. accuracy for the LW building (a) and HW building (b) including three starting days (May 15: red, May 20: green, May 25: blue), three model variants (R1C1: circle, R3C3: triangle, R5C4: square), and two parameter sets (wide bounds: white, narrow bounds: filled), and three learning period lengths (1 day, 3 days, 7 days: not differentiated).

controller. In the best case for each week tested, these savings are approximately 9%. It is important to consider that the relative performance of an MPC controller, and thereby deemed savings, compared to a more conventional controller is highly dependent on an applications incentive and ability to shift load, as well as the specific setup of the MPC controller. In this study, the performance of the MPC controllers relative to the conventional controller could be adjusted by adding or removing thermal mass or energy storage elements to the building case as well as providing more or less incentive to shift load through the diurnal price disparity or additional objectives like peak demand costs. It could also be adjusted with more fine tuning on the MPC controller for each individual case, for instance the length of the control step and the method of state estimation. MPC controllers that, in the particular situation shown here, do not appear to save on energy costs may actually still save on energy costs in situations with stronger incentives or more finely tuned MPC controllers. However, as indicated in Section 1 of this paper, it is not in the scope to consider these additional controller variations. Instead, we focus here on the model setup. These other variations should be the subject of future work to further identify minimum building design and application requirements for MPC to be effective. Therefore, the remainder of this analysis will focus on the performance of the MPC controllers relative to each other, which indeed is the primary purpose of this study. In particular, we focus on the HW case, where sufficient thermal mass is present to take advantage of the use of MPC.

In general, the data indicates that the strongest predictor of MPC performance is model structure and parameter set. As can be seen in Figure 10(b), all of the points that represent controllers that performed as well as or worse than the conventional are those cases that used the R1C1 model and/or a wide parameter bound set for estimation. Otherwise, for the models that used a model structure with a complexity of at least three states and a narrow parameter bound, savings of 5-9% were achieved. This can be seen in the Figure 10(b) by the group of points clustered in this savings range and with an RMSE ranging from 0.5 to 2 °C. Notice, however, that RMSE is a better predictor for performance for the week of May 15, which experiences the highest variation in RMSE, and the worst for the week of May 25. This indicates that a relatively high model accuracy, in terms of RMSE, could be necessary for good MPC controller performance, but not sufficient.

Figure 11 shows the time series data for three cases for the week of May 25. The Figure shows zone temperature and HVAC cooling power consumption for the emulator under conventional control (blue) and MPC (red). Note that no heating is used in any of the cases and is, therefore, omitted from presentation. In addition, the optimal predicted temperature and HVAC performance produced by the MPC controller for one specific hour, 24, is also shown (black). The first case, Figure 11(a), shows the best performing controller model, using an R3C3 model with narrow parameter sets and one day of training data, which saved approximately 9% energy cost compared to the conventional. This controller shifts cooling to the morning hours in order to reduce cooling during the periods of high prices. Overall, more cooling energy is used than the conventional control to generally reduce the need to cool during high price periods, indicated by consistently lower zone temperatures and greater area under the HVAC power trajectory. This is a natural consequence of minimizing energy cost and not energy alone, as is discussed elsewhere [63, 32].

The second case Figure 11(b) shows a poorly performing controller, using an R5C4 model with wide parameter sets and seven days of training data, which had an energy cost approximately 10% more than the conventional. This optimization solver in this controller often finished with the detection of an infeasible problem, perhaps due to parameters that are over-trained to the training operating conditions and unphysical temperature response predictions to cooling and heating control signals during optimization iterations. The returned temperature trajectory was, therefore, not the optimal for the building.

Finally, the third case Figure 11(c) shows a poorly performing controller, using an R1C1 model with narrow parameters sets and one day of training data, which had an energy cost approximately 5% more than the conventional. Here, the controller suffered from an inability to correctly predict the effect of thermal mass. Lumping all of the thermal mass into a single state did not allow for the split of fast and slow dynamics present in the air volume and structure of the building. As the HVAC system was assumed to provide heating and cooling directly to the air, the resulting lumping of slow and fast dynamics led to an over-prediction of the effect of thermal mass on the air temperature. Therefore, this controller overestimates the ability to reduce cooling during high price hours. While this controller is stable, the resulting control solution is not optimal for the emulated building, resulting in performance that is worse than conventional control.

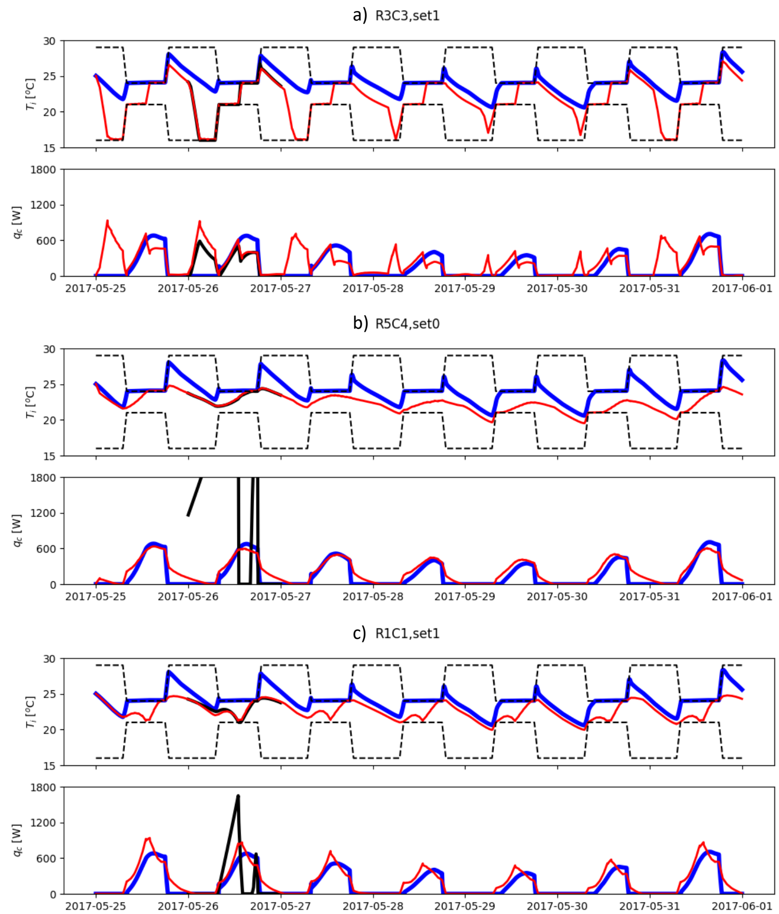

Figure 11: Conventional (blue) and MPC (red) controller performance for the HW building over the week of May 25 for three model cases. The optimal predicted temperature and HVAC performance produced by the MPC controller for one specific hour, 24, is also shown (black). Zone temperature (top) and cooling power (bottom) is shown for each subplot. Zone temperature limits are shown as dashed black lines. No heating power is used in any case and is not shown. The first model case represents the best performer (a), the second model case represents a performer with poor parameter estimation (b) leading to difficulty for the optimization solver to find a feasible solution, and the third model case represents a performer with poor model structure (c) leading to over-estimation of the effect of thermal mass.

6. Discussion

This analysis has identified and tested many factors that play a role in determining the accuracy of models for MPC, and quantified the impacts of them according to a single accuracy metric, RMSE, and MPC performance metric, energy cost. It distinguished the effects of factors that would otherwise compound each other in studies presenting a single implementation. Despite the scope of the analysis being limited to a single zone building, simple HVAC system, and envelope model, it can be expected that the conclusions are applicable to MPC in buildings in general because of the systematic approach taken to analyzing a large number of variations on a representative case.

The results of this study may be applied more generally in the following ways:

1. This study has identified many of the practical issues to be considered when implementing a model for MPC. This may then serve as a comprehensive list of factors that should be considered when designing and implementing MPC, as well as when analyzing the performance after implementation. Such a list can be particularly helpful as MPC gains attention in industry.

2. The results of the study indicate that model structure and data quality play significant roles in the success of the MPC controller. In particular, higher order models are able to more accurately account for dynamics occurring on different timescales, which improves controller performance, however, require good initial guesses for a larger number of parameters. While this notion is not necessarily new, we see from this study that the influence on a given MPC implementation could mean a performance difference of up to 20% for specific performance periods. In addition, this difference is not necessarily accounted for by the time series accuracy metric of RMSE. While this study indicates low RMSE is necessary for good performance, it is not sufficient, which points to the importance of using additional methods for determining accuracy, such as frequency response or other time series analysis methods.

   Data quality is represented in a few ways in this study. First, the initial guess of parameters being identified also plays a significant role in the accuracy of the model, as well as ultimate controller performance. Good initial guesses can only come from comprehensive and accurate documentation of building construction. They can also be aided by workflows to translate this information into digital forms or take digital forms of this information and communicate it to MPC controllers. Second, noisy and missing data can significantly influence the ability to identify models. Some noise can be tolerated, though too much can lead to over-estimation of thermal mass. In addition, some missing data can be handled by cleaning strategies. The cleaning strategy that works best is dependent on the nature of the data source. For instance, averaging data from similar times of different days works well if data are highly periodic from day to day, however, if data are not, simple interpolation can work best. In general, strategies work best if missing data is less than 25% so that the cleaning strategies maintain the step-by-step variability in the data.

3. This study opens up multiple avenues for future work. First, there are many parameters that additionally influence MPC controller performance other than model setup. The timestep of control, state estimator, method of implementation of optimal control strategy to the real system, control optimization algorithm, degree of incentives for MPC, and missing data for state estimation will all affect performance relative to a more conventional control. Second, multi-zone buildings and HVAC systems add additional models for which the practical factors discussed in this study apply. Third, the analysis of controller performance as a function of accuracy could lead to better understanding and eventual predictions of theoretical limits of performance under MPC with perfect models or of performance with less accurate models. While such sensitivity studies may be difficult to generalize for all applications, characterizing the performance of the controller based on the model accuracy required for particular applications, and further identifying the effort required to reach a particular accuracy, would help determine the cost-benefit ratio of implementing MPC controllers before embarking on implementation.

7. Conclusions

MPC for buildings is attracting significant attention in research and industry due to its potential to address a number of challenges facing the building industry, including energy cost reduction, grid integration, and occupant connectivity, as indicated by a number of research demonstrations. However, the strategy has not yet been implemented at any scale. One major reason for this is the significant effort required to configure the model used in the MPC controller. While many studies have focused on methods to improve model accuracy, few have studied their impact, along with other factors affecting model accuracy, on a wide range of possible cases. In addition, few have continued on to analyze their impact on ultimate MPC controller performance. Therefore, this study has set out to identify the practical factors affecting model accuracy, analyze their individual influence, and then analyze the impact on MPC controller performance.

The factors analyzed included building design, model structure, identification algorithm, initial parameter guesses, training data length, training data noise, training data gaps, and strategies for accounting for missing data. The ASHRAE / ANSI BESTEST lightweight Case 600 and heavyweight Case 900 building envelopes were used along with a simple HVAC system with direct heating and cooling to the building air and constant efficiency and COP. RMSE of zone temperature was used as an indicator of model accuracy, as is done in previous literature. Data was presented for the many trials of parameter estimation that were performed in a way that indicates the relative effect of each factor on model accuracy. Finally, a subset of models resulting from the many parameter estimation trials were implemented in an MPC controller, tested in simulation, and compared to a conventional feedback controller as well as each other.

The first main conclusion from the study is that low RMSE is necessary for good MPC control, but not sufficient. Over-training of parameters and improper model structure can lead to MPC controllers where the optimal control solution determined by the controller is actually a bad control function for the actual building. For the specific cases studied in this paper through building emulation, these issues lead to a difference in energy cost of up to 20% between the best and worst performing controllers, and an increase of 10% energy cost compared to conventional control for the worst performing controller. Such issues can be avoided with the appropriate model complexity and access to comprehensive building design documentation for good initial parameter estimates. A second conclusion is that the optimal length of training data for the single zone envelope models studied is a few days, indicating promise for adaptive parameter estimation frameworks that update parameters often. A third conclusion is that a gradient-based optimization algorithm with an associated global start algorithm, such as latin hybercube sampling, can estimate parameters as well as a global optimization algorithm with less computational cost. Finally, significant noise and missing data can reduce the ability to identify models, particularly if applied to variables used to measure the accuracy of the model, such as indoor air temperature, and those that directly affect such variables, such as solar gain through a window. However, moderate amounts of missing data can be handled by simple interpolation or more complex averaging strategies.

This study forms a foundation of future work on how to analyze MPC performance. While it is necessary to show successful demonstrations and case studies, it is important to begin characterizing controller performance with respect to the many factors that play a role. This paper begins with model setup factors and relatively simple case studies. However, future work can expand to other factors, including control time steps, state estimators, control optimization algorithms, objective functions, and building designs. In addition, other research pathways involve quantifying the effort and cost required to achieve the data and control integration with building automation systems necessary for MPC implementation, as well as other advanced building control strategies and analytics. Combined, such work would enable better evaluation of the effort required to implement MPC control that is good enough, the marginal gain in performance resulting from extra efforts, and ultimately the ability to predict the cost-benefit ratio before embarking on implementation. Such an upfront evaluation is critical for adoption at scale.

Acknowledgements

This research was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technologies of the U.S. Department of Energy, under Contract No. DE-AC02-05CH11231.

This work was supported by the U.S.-China Clean Energy Research Center (CERC) 2.0 on Building Energy Efficiency (BEE).

This work was supported by the Innovation Fund Denmark for the project COORDICY (4106-00003B).

This work was supported by Engie Axima.

This work emerged from the IBPSA Project 1, an international project conducted under the umbrella of the International Building Performance Simulation Association (IBPSA). Project 1 will develop and demonstrate a BIM/GIS and Modelica Framework for building and community energy system design and operation.

References

[1] A. Afram, F. Janabi-Sharifi, Theory and applications of HVAC control systems: A review of model predictive control (MPC), Building and Environment 72 (2014) 343–355. doi:10.1016/j.buildenv.2013.11.016.
URL https://www.sciencedirect.com/science/article/pii/S0360132313003363

[2] Y. Kim, L. K. Norford, Optimal use of thermal energy storage resources in commercial buildings through price-based demand response considering distribution network operation, Applied Energy 193 (2017) 308–324. doi:10.1016/J.APENERGY.2017.02.046.
URL https://www.sciencedirect.com/science/article/pii/S0306261917301897{#}e0145

[3] A. Mirakhorli, B. Dong, Occupancy behavior based model predictive control for building indoor climateA critical review, Energy and Buildings 129 (2016) 499–513. doi:10.1016/J.ENBUILD.2016.07.036.
URL https://www.sciencedirect.com/science/article/pii/S0378778816306338

[4] G. P. Henze, Model predictive control for buildings: a quantum leap?, Journal of Building Performance Simulation 6 (3) (2013) 157–158. doi:10.1080/19401493.2013.778519.
URL http://www.tandfonline.com/doi/abs/10.1080/19401493.2013.778519

[5] D. Sturzenegger, D. Gyalistras, M. Morari, R. S. Smith, Model Predictive Climate Control of a Swiss Office Building: Implementation, Results, and Cost-Benefit Analysis, IEEE Transactions on Control Systems Technology 24 (1). doi:10.1109/TCST.2015.2415411.

[6] P. Rockett, E. A. Hathway, Model-predictive control for non-domestic buildings: a critical review and prospects, Building Research and Information 45 (5). doi:10.1080/09613218.2016.1139885.

[7] S. Prívara, J. Cigler, Z. Vána, F. Oldewurtel, C. Sagerschnig, E. Žáčeková, Building modeling as a crucial part for building predictive control, Energy and Buildings 56. doi:10.1016/j.enbuild.2012.10.024.

[8] E. Atam, L. Helsen, Control-Oriented Thermal Modeling of Multizone Buildings: Methods and Issues: Intelligent Control of a Building System, IEEE Control Systems 36 (3). doi:10.1109/MCS.2016.2535913.

[9] M. Wetter, M. Bonvini, T. S. Nouidui, Equation-based languages a new paradigm for building energy modeling, simulation and optimization, Energy and Buildings 117 (2016) 290 – 300. doi:https://doi.org/10.1016/j.enbuild.2015.10.017.
URL http://www.sciencedirect.com/science/article/pii/S0378778815303315

[10] P. May-Ostendorp, G. P. Henze, C. D. Corbin, B. Rajagopalan, C. Felsmann, Model-predictive control of mixed-mode buildings with rule extraction, Building and Environment 46 (2) (2011) 428–437. doi:10.1016/J.BUILDENV.2010.08.004.
URL https://www.sciencedirect.com/science/article/pii/S0360132310002477

[11] B. Coffey, F. Haghighat, E. Morofsky, E. Kutrowski, A software framework for model predictive control with GenOpt, Energy and Buildings 42 (7) (2010) 1084–1092. doi:10.1016/J.ENBUILD.2010.01.022.
URL https://www.sciencedirect.com/science/article/pii/S0378778810000290{#}bib1

[12] X. Li, J. Wen, Review of building energy modeling for control and operation, Renewable and Sustainable Energy Reviews 37 (2014) 517 – 537. doi:https://doi.org/10.1016/j.rser.2014.05.056.
URL http://www.sciencedirect.com/science/article/pii/S1364032114003815

[13] Y. Lin, T. Middelkoop, P. Barooah, Issues in identification of control-oriented thermal models of zones in multi-zone buildings, in: 2012 IEEE 51st IEEE Conference on Decision and Control (CDC), IEEE, 2012, pp. 6932–6937. doi:10.1109/CDC.2012.6425958.
URL http://ieeexplore.ieee.org/document/6425958/

[14] D. Kim, J. E. Braun, A general approach for generating reduced-order models for large multi-zone buildings, Journal of Building Performance Simulation 8 (6) (2015) 435–448. doi:10.1080/19401493.2014.977952.
URL http://www.tandfonline.com/doi/full/10.1080/19401493.2014.977952

[15] D. Picard, J. Drgona, M. Kvasnica, L. Helsen, Impact of the controller model complexity on model predictive control performance for buildings, Energy and Buildings 152 (2017) 739–751. doi:10.1016/j.enbuild.2017.07.027.
URL http://www.sciencedirect.com/science/article/pii/S0378778817302190

[16] D. Kim, J. Cai, K. B. Ariyur, J. E. Braun, System identification for building thermal systems under the presence of unmeasured disturbances in closed loop operation: Lumped disturbance modeling approach, Building and Environment 107 (2016) 169–180. doi:10.1016/J.BUILDENV.2016.07.007. URL http://www.sciencedirect.com/science/article/pii/S0360132316302542

[17] R. De Coninck, F. Magnusson, J. Åkesson, L. Helsen, Toolbox for development and validation of grey-box building models for forecasting and control, Journal of Building Performance Simulation 9 (3). doi:10.1080/19401493.2015.1046933.

[18] H. Harb, N. Boyanov, L. Hernandez, R. Streblow, D. Müller, Development and validation of grey-box models for forecasting the thermal response of occupied buildings, Energy and Buildings 117. doi:10.1016/j.enbuild.2016.02.021.

[19] P. Armstrong, S. B. Leeb, L. Norford, Control with building mass part i: Thermal response model, ASHRAE Transactions 112.

[20] T. Muneer, S. Younes, S. Munawwar, Discourses on solar radiation modeling, Renewable and Sustainable Energy Reviews 11 (4) (2007) 551–602. doi:10.1016/j.rser.2005.05.006.

[21] E. Žáčeková, Z. Vána, J. Cigler, Towards the real-life implementation of MPC for an office building: Identification issues, Applied Energy 135 (2014) 53–62. doi:10.1016/J.APENERGY.2014.08.004. URL https://www.sciencedirect.com/science/article/pii/S0306261914008071

[22] D. F. M. Cabrera, H. Zareipour, Data association mining for identifying lighting energy waste patterns in educational institutes, Energy & Buildings 62 (2013) 210–216. doi:10.1016/j.enbuild.2013.02.049. URL http://dx.doi.org/10.1016/j.enbuild.2013.02.049

[23] A. Rogier, T. Donders, G. J. M. G. van der Heijden, T. Stijnen, K. G. M. Moons, Review: A gentle introduction to imputation of missing values, Journal of Clinical Epidemiology 59 (10) (2006) 1087–1091. doi:10.1016/j.jclinepi.2006.01.014.

[24] J. Hu, O. T. Ogunsola, L. Song, R. A. McPherson, M. Zhu, Y. Hong, S. Chen, Restoration of 1-24 hour dry-bulb temperature gaps for use in building performance monitoring and analysis - Part i, HVAC and R Research 20 (6) (2014) 594–605. doi:10.1080/10789669.2014.925347.

[25] G. Rodriguez, C. Catalita, Qualité des données capteurs pour les systèmes de surveillance de phénomènes environnementaux, Ph.D. thesis (2010).

[26] M. Sourbron, C. Verhelst, L. Helsen, Building models for model predictive control of office buildings with concrete core activation, Journal of Building Performance Simulation 6 (3) (2013) 175–198. doi:10.1080/19401493.2012.680497.
URL http://www.tandfonline.com/doi/abs/10.1080/19401493.2012.680497

[27] K.-h. Lee, J. E. Braun, Model-based demand-limiting control of building thermal mass, Building and Environment 43 (10) (2008) 1633–1646. doi:10.1016/J.BUILDENV.2007.10.009.
URL http://www.sciencedirect.com/science/article/pii/S0360132307001850

[28] J. Cai, J. E. Braun, EFFICIENT AND ROBUST TRAINING METHODOLOGY FOR INVERSE BUILDING MODELING, Proceedings of SimBuild 5 (1) (2012) 456–461.
URL http://www.ibpsa-usa.org/index.php/ibpusa/article/view/463

[29] J. Bonilla, J. A. Carballo, L. Roca, M. Berenguel, Development of an open source multi-platform software tool for parameter estimation studies in FMI models, in: Proceedings of the 12th International Modelica Conference, 2017. doi:10.3384/ecp17132683.
URL https://www.modelica.org/events/modelica2017/proceedings/html/submissions/ecp17132683{_}BonillaCarballoRocaBerenguel.pdf

[30] J. Joe, P. Karava, Agent-based system identification for control-oriented building models, Journal of Building Performance Simulation 10 (2) (2017) 183–204. doi:10.1080/19401493.2016.1212272.
URL https://www.tandfonline.com/doi/full/10.1080/19401493.2016.1212272

[31] T. Zakula, P. Armstrong, L. Norford, Modeling environment for model predictive control of buildings, Energy and Buildings 85 (2014) 549–559. doi:10.1016/J.ENBUILD.2014.09.039.
URL https://www.sciencedirect.com/science/article/pii/S0378778814007737

[32] D. H. Blum, T. Zakula, L. K. Norford, Opportunity Cost Quantification for Ancillary Services Provided by Heating, Ventilating, and Air-Conditioning Systems, IEEE Transactions on Smart Grid 8 (3) (2017) 1264–1273. doi:10.1109/TSG.2016.2582207.
URL http://ieeexplore.ieee.org/document/7494593/

[33] G. Reynders, J. Diriken, D. Saelens, Quality of grey-box models and identified parameters as function of the accuracy of input and observation signals, Energy and Buildings 82 (2014) 263–274.

doi:10.1016/J.ENBUILD.2014.07.025.
URL https://www.sciencedirect.com/science/article/pii/S0378778814005623

[34] M. Maasoumy, M. Razmara, M. Shahbakhti, A. S. Vincentelli, Handling model uncertainty in model predictive control for energy efficient buildings, Energy and Buildings 77 (2014) 377–392. doi:10.1016/J.ENBUILD.2014.03.057.
URL https://www.sciencedirect.com/science/article/pii/S0378778814002771{#}bib0205

[35] D. Sturzenegger, D. Gyalistras, V. Semeraro, M. Morari, R. S. Smith, BRCM Matlab Toolbox: Model generation for model predictive building control, in: 2014 American Control Conference, IEEE, 2014, pp. 1063–1069. doi:10.1109/ACC.2014.6858967.
URL http://ieeexplore.ieee.org/document/6858967/

[36] MATLAB, version 7.10.0 (R2010a), The MathWorks Inc., Natick, Massachusetts, 2010.

[37] T. Blochwitz, M. Otter, J. Akesson, M. Arnold, C. Clau, H. Elmqvist, M. Friedrich, A. Junghanns, J. Mau, D. Neumerkel, H. Olsson, A. Viel, Functional mockup interface 2.0: The standard for tool independent exchange of simulation models, in: Proceedings of the 9th International MODELICA Conference; September 3-5; 2012; Munich; Germany, no. 76, Linkping University Electronic Press; Linkpings universitet, 2012, pp. 173–184.

[38] B. M. Adams, M. S. Ebeida, M. S. Eldred, J. D. Jakeman, L. P. Swiler, J. A. Stephens, et al., Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.1 User's Manual, Sandia National Laboratories, Albuquerque, New Mexico, USA, 2014.

[39] M. Bonvini, M. Wetter, M. D. Sohn, An FMI-based Framework for State and Parameter Estimation, in: Proceedings of the 10th International Modelica Conference, Lund, Sweden, 2014, pp. 647–656. doi:10.3384/ECP14096647.
URL http://www.ep.liu.se/ecp/096/068/ecp14096068.pdf

[40] T. E. Oliphant, Guide to NumPy, 2nd Edition, CreateSpace Independent Publishing Platform, USA, 2015.

[41] E. Jones, T. Oliphant, P. Peterson, et al., SciPy: Open source scientific tools for Python, [Online; accessed 8/6/2018] (2018).
URL http://www.scipy.org/

[42] D. H. Blum, M. Wetter, MPCPy: An Open-Source Software Platform for Model Predictive Control in Buildings, in: Proceedings of the 15th

IBPSA Conference, San Francisco, CA, 2017, pp. 1381–1390. doi:10.26868/25222708.2017.351.
URL http://www.ibpsa.org/proceedings/BS2017/BS2017{_}351.pdf

[43] S. E. Mattsson, H. Elmqvist, Modelica - an international effort to design the next generation modeling language, in: 7th IFAC Symposium on Computer Aided Control systems Design, Gent, Belgium, 1997.
URL https://www.modelica.org/publications/papers/CACSD97Modelica.pdf

[44] J. Åkesson, M. Gäfvert, H. Tummescheit, JModelica—an Open Source Platform for Optimization of Modelica Models, 2009, 6th Vienna International Conference on Mathematical Modelling ; Conference date: 11-02-2009.

[45] U. of Wisconsin-Madison. Solar Energy Laboratory, TRNSYS, a Transient Simulation Program, Report, The Laboratory, 1975.
URL https://books.google.dk/books?id=16BFnQEACAAJ

[46] B. Houska, H. J. Ferreau, M. Diehl, ACADO toolkit-An open-source framework for automatic control and dynamic optimization, Optimal Control Applications and Methods 32 (3) (2011) 298–312. doi:10.1002/oca.939.
URL http://doi.wiley.com/10.1002/oca.939

[47] F. Jorissen, G. Reynders, R. Baetens, D. Picard, D. Saelens, L. Helsen, Implementation and verification of the ideas building energy simulation library, Journal of Building Performance Simulation 0 (0) (2018) 1–20. arXiv:https://doi.org/10.1080/19401493.2018.1428361, doi:10.1080/19401493.2018.1428361.
URL https://doi.org/10.1080/19401493.2018.1428361

[48] R Core Team, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria (2018).
URL https://www.R-project.org

[49] K. Arendt, M. Jradi, M. Wetter, C. T. Veje, ModestPy: An Open-Source Python Tool for Parameter Estimation in Functional Mock-up Units, in: Proceedings of the American Modelica Conference 2018, Cambridge, MA, USA, 2018, to appear.

[50] ANSI/ASHRAE, Standard Method of Test for the Evaluation of Building Energy Analysis Computer Programs, Standard 140-2007, The American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2007.
URL https://www.techstreet.com/ashrae/standards/ashrae-140-2007?product_id=1588605

[51] M. Wetter, W. Zuo, T. S. Nouidui, X. Pang, Modelica Buildings library, Journal of Building Performance Simulation 7 (4) (2014) 253–270. doi:10.1080/19401493.2013.765506.

URL http://www.tandfonline.com/doi/abs/10.1080/19401493.2013.765506

[52] J. D. Spitler, Load Calculation Applications Manual, The American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2008.

[53] J. Åkesson, Optimica-An Extension of Modelica Supporting Dynamic Optimization, in: Proc. 6th international modelica conference, Bielefeld, Germany, 2008, pp. 57–66.
URL https://pdfs.semanticscholar.org/8efe/348a3ffa632aaddf07a9ecae9e4a6bde79af.pdf

[54] F. Magnusson, J. Åkesson, Collocation Methods for Optimization in a Modelica Environment, in: In Proceedings of the 9th International Modelica Conference, 2012, pp. 649–658.
URL https://lup.lub.lu.se/search/ws/files/7670906/2972281.pdf

[55] J. Andersson, J. Åkesson, M. Diehl, CasADi: A Symbolic Package for Automatic Differentiation and Optimal Control, Springer, Berlin, Heidelberg, 2012, pp. 297–307. doi:10.1007/978-3-642-30023-3_27.
URL http://link.springer.com/10.1007/978-3-642-30023-3{_}27

[56] A. Wächter, L. T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, Mathematical Programming 106 (1) (2006) 25–57. doi:10.1007/s10107-004-0559-y.
URL http://link.springer.com/10.1007/s10107-004-0559-y

[57] HSL, a collection of Fortran codes for large-scale scientific computation.
URL http://www.hsl.rl.ac.uk/

[58] M. Wetter, GenOpt(R), generic optimization program, User Manual, Version 2.0.0, Report LBNL-54199, Lawrence Berkeley National Laboratory, 2003.
URL https://escholarship.org/uc/item/5dp8q7m1

[59] P. Bacher, H. Madsen, Identifying suitable models for the heat dynamics of buildings, Energy & Buildings 43 (7) (2011) 1511–1522. doi:10.1016/j.enbuild.2011.02.005.
URL http://dx.doi.org/10.1016/j.enbuild.2011.02.005

[60] P. Interconnection, Data Miner 2: Real-time hourly LMPs, [Online; accessed 10/28/2018] (2018).
URL http://dataminer2.pjm.com/feed/rt_hrl_lmps

[61] P. Gas, E. Company, Electric Schedule E-19 : Medium general demand-metered TOU service, [Online; accessed 10/28/2018] (2010).
URL https://www.pge.com/tariffs/assets/pdf/tariffbook/ELEC_SCHEDS_E-19.pdf

[62] M. Rodriguez-Garcia, M. Batet, D. Sánchez, A semantic framework for noise addition with nominal data, Knowledge-Based Systems 122 (2017) 103–118. doi:10.1016/j.knosys.2017.01.032.
URL http://dx.doi.org/10.1016/j.knosys.2017.01.032

[63] R. De Coninck, L. Helsen, Quantification of flexibility in buildings by cost curves Methodology and application, Applied Energy 162 (2016) 653–665. doi:10.1016/J.APENERGY.2015.10.114.
URL http://www.sciencedirect.com/science/article/pii/S0306261915013501