Open access • Journal Article • DOI:10.1080/19401493.2020.1770861

Identification of multi-zone grey-box building models for use in model predictive control — Source link

Javier Arroyo, Fred Spiessens, Lieve Helsen

Institutions: Flemish Institute for Technological Research, Katholieke Universiteit Leuven

Published on: 27 May 2020 - Journal of Building Performance Simulation (Taylor & Francis)

Topics: Model predictive control and Building model

Identification of Multi-Zone Grey-Box Building Models for Use in Model Predictive Control

Javier Arroyo[a,b,c,*], Fred Spiessens[b,c], Lieve Helsen[a,b]

[a] *KU Leuven, Department of Mechanical Engineering, Leuven, Belgium*
[b] *EnergyVille, Thor Park, Waterschei, Belgium*
[c] *VITO NV, Boeretang 200, Mol, Belgium*

Abstract

Predictive controllers can greatly improve the performance of energy systems in buildings. An important challenge of these controllers is the need of a building model accurate and simple enough for optimization. Grey-box modeling stands as a popular approach, but the identification of reliable grey-box models is hampered by the complexity of the parameter estimation process, specifically for multi-zone models. Hence, single-zone models are commonly used, limiting the performance and applicability of the predictive controller.
This paper investigates the feasibility of identification of multi-zone grey-box building models and the benefits of using these models in predictive control. For this purpose, the parameter estimation process is split by individual zones to obtain an educated initial guess. A virtual test case from the BOPTEST framework is contemplated to assess the simulation and control performance. The results show the relevance of modelling thermal interactions between zones in the multi-zone building.

Keywords: Model predictive control, building modeling, grey-box modeling, multi-zone, BOPTEST

1. Introduction

Buildings use over 30% of the total final energy for all sectors with space heating accounting for 36% of this amount in the countries belonging to the Major Economies Forum [1]. Global building stock in 2040 could be 60% larger than today in floor area for no increase in overall energy demand [2]. This translates into a necessary enhancement of energy efficiency in buildings where digitalization is a key factor as it supports advanced control strategies in building energy management systems.
Advanced controllers like model predictive control (MPC) have proven to outperform the traditional rule based controllers by enabling energy savings and higher thermal comfort [3, 4]. Moreover, MPC can be used to provide demand response services to the electric grid when power-to-heat resources, like heat pumps, are available at building level [5, 6]. The main challenge when setting up an MPC controller is to identify a model of the building system [7, 8] that is accurate and fast enough at the same time, to be used for optimization.
Grey-box modeling has been extensively used for MPC, some examples are [7, 9–16]. Their popularity arises from their convenient trade-off between accuracy and computational tractability [17, 18]. This modeling approach starts by building the model structure from physical knowledge and then estimating its parameters by using building operational data. The physics of these models are simplified by using lumped parameters, i.e., parameters that represent the physical properties of not only one specific component but of several components together. Therefore, grey-box models combine the strengths of both worlds: physics and data. Their physical properties give them physical consistency and intelligibility, and their data-driven nature gives them adaptability.

However, the identification of individual parameters in grey-box building models is hampered by the nature of the associated non-convex parameter estimation process that is very likely to drive the solution towards local optima. Hence, an accurate initial guess of the parameters is crucial, but its calculation requires expert involvement and building data that is often not available. Moreover, parameter estimation of differential equations has been claimed to be computationally expensive and therefore unsuitable for large-scale problems [19]. The problem can be reformulated as convex [20], but only at the cost of losing physical interpretability of the parameters that need to be merged together. This challenge is even larger for multi-zone grey-box models where the number of parameters to be estimated increases. For this reason, researchers often decide to adopt single-zone grey-box models or black-box modeling approaches instead [21]. However, multi-zone grey-box models are able to output one temperature estimate per zone with the important advantage that it can be used by a predictive controller to optimize the heating inputs to each zone individually,

*Corresponding author
Email addresses: javier.arroyo@kuleuven.be (Javier Arroyo), fred.spiessens@vito.be (Fred Spiessens), lieve.helsen@kuleuven.be (Lieve Helsen)

This is an Accepted Manuscript of an article published by Taylor & Francis in *Journal of Building Performance Simulation* on *13 May 2020*, available online: http://www.tandfonline.com/10.1080/19401493.2020.1770861.

instead of committing the heat distribution to any predefined and non-optimal sub-controller.

Two main approaches can be followed to identify multi-zone models: first, a decentralized approach identifies a single-zone grey-box model for each zone independently of the others. This approach does not consider thermal interaction between zones, which can hamper the accuracy of the models if there is a strong inter-zonal thermal coupling. Second, a centralized approach identifies a multi-zone model directly, which allows to take into account thermal interactions of adjacent zones. A decentralized model allows to decouple the parameter and predictive controller problems among the different zones. Therefore, if a centralized model shows no added value, the decentralized approach may be more convenient to implement. Henceforth, the term *centralized model* is used to refer to a multi-zone grey-box building model according to the centralized approach, and the term *decentralized model* to refer to a multi-zone grey-box building model according to the decentralized approach.

Building models with a wide range of complexities have been implemented and compared for MPC in buildings. Verhelst et al., [22] showed that low-order models provide similar accuracy to higher order models for both building and borehole heat exchanger modeling. On the contrary, Picard et al. [23] showed that low order models may not be able to capture complex multi-zone dynamics. A direct comparison of the grey- and white-box paradigms can be found in [24], where more detailed white-box models are advised over low order grey-box models. Arendt et al., [25] compared white-, grey-, and black-box models in a real test case and across different cross-validation data sets. In this study the black-box models outperformed the grey- and white-box models in 7 out of 8 considered simulation tests. It was also found that low-order RC models are not reliable for long-term prediction due to the error accumulation over time. Additionally, restricting the model to a single or just a few zones limits the controllability without harnessing the full potential of the optimal controllers. Still, single-zone grey-box models are commonly implemented in practice. Some examples are [14, 26, 27].

Although extensive research has been carried out on grey-box building modeling in general, only a few studies focus on the feasibility and benefits of identifying higher-order grey-box models for multi-zone buildings. Agbi et al. [28] identified a 13-zone grey-box building model, but the model parameters were calculated from construction data and thermal properties of the materials. Additionally, their work was focused on the identifiability of the parameters and optimal experimental conditions for training, paying little attention to the model structure. An insightful methodology was developed in [29] where the coupling strength between zones was quantified to identify potential grouping of zones in order to break down the problem into more sub-problems with lower estimate dimensions. Furthermore, they reduced the parameter space of each zone by analyzing the information matrix. This sophisticated method was effectively validated in a nine-zone office building, but it required a considerable amount of expert involvement and was difficult to automate. In [27], a toolbox for grey-box system identification was developed, but it was not tested for models with more than two thermal zones. A survey of control-oriented thermal modeling of multi-zone buildings can be found in [30] which highlights the advantages and disadvantages of each modeling technique. In all cases, the performance was investigated for simulation only, even though control performance is a crucial test for the assessment of a controller model. This paper explores the feasibility and benefits of the identification of multi-zone grey-box models in an automated way, without too much engineering effort. Additionally, it compares the identified models in simulation and control performance. The main research questions to be answered are:

- Is it possible to identify a higher-order grey-box model for a multi-zone building without extensive expert involvement, to be used in predictive control?
- What is the added value of a centralized multi-zone grey-box model when compared against a decentralized and a single-zone model?
- How sensitive is a higher-order grey-box model to the training data length, to the boundary conditions, and to variations in the estimated parameter values?

The remaining part of the paper proceeds as follows: Section 2 details the methodology that is followed for the identification of multi-zone building models; Section 3 describes the simulation building test case from the BOPTEST framework that is used to implement the methodology; Section 4 elaborates on the results and benchmarks the obtained centralized model with a decentralized and a single-zone model. Finally, the main conclusions are drawn in Section 5.

2. Methodology

2.1. Description of the methodology

The building models consider the thermal powers to each zone as inputs and the indoor temperatures of each zone as outputs. Hence, the modeling of Heating Ventilation and Air Conditioning (HVAC) systems are out of the scope of this work. A building zone is defined as a building space with a temperature measurement and a controllable heating input, i.e., a thermal power that can be directly or indirectly controlled. The coupling between the building envelope model and the model of the HVAC system, if required, can be done in a later step, using the zone thermal powers as interface between both.

The approach used to have a good initial guess of the model parameters is to systematically split the complexity of the problem by zones while keeping the adjacent zone temperatures as exogenous variables. In this way, the problem is broken down into tractable subproblems, as in a decentralized approach, while keeping the thermal influence of adjacent zones, as in a centralized approach. Then, the zones are coupled together and the results of each subproblem are used as an educated initial guess that allows to estimate the parameters of the full model merged. The approach incorporates scalability and is based on physical checks and heuristics.

Only general building knowledge that can be provided by a non-specialist is required to launch the process, namely the zone volumes and the physically connected zones. Operational system data is required as well to identify and train the model parameters. Particularly, it is assumed that historical data of the following variables are available.

- Air temperature of each zone, T_z.
- Heat released to each zone, $\dot{Q}_z$.
- Internal gains of each zone, $\dot{G}_z$.
- Ambient temperature, T_a.
- Solar global horizontal irradiation, $\dot{Q}_{Sun}$.

A forward selection process compares the results among models of increasing level of complexity for each individual zone. The RC architectures used in the forward selection inputs are shown in Figure 1 without any adjacent zone temperature input. These RC architectures represent a gradual refinement of the zone model. The assumption that a limited set of RC models can capture the zone dynamics is supported by Reynders et al. [11] who argued that only few model types are required to represent the majority of buildings. The model complexities range from only one thermal state to four thermal states, being the zone air, wall, internal and embedded temperatures: T_z, T_w, T_i and T_e, respectively, where internal refers to internal walls, and embedded refers to embedded pipes in the emission system. The parameters to be estimated are the air, internal, wall and embedded thermal capacitances, namely C_z, C_w, C_i and C_e; as well as the thermal resistors for the wall, internal, infiltration and embedded: R_w, R_i, R_{inf} and R_e; and the solar transmittance gA. The initial thermal states are also allowed to vary in the set of optimization variables.

Having a good initial guess for the optimization variables is critical, even for the identification of each single-zone model. A Latin Hypercube Sampling (LHS) method is used as a multi-start method to explore the parameter space in the identification of each single zone model. Each parameter of the sample gets a random initial guess around a predefined physical value assigned from the little physical information provided. In particular, the initial guess of the air capacitance is estimated from the volume of the zone. The rest of the parameters are given an initial value based on their nature and a much wider search space. The thermal capacitances representing the air of the zones are allowed to vary only one order of magnitude around the computed physical value used as initial guess, whereas the other parameters are allowed a much wider search space, namely of six orders of magnitude around their initial guesses. Therefore, only the air zone thermal capacitances are expected to keep their strict physical meaning whereas the other parameters are more abstract and often lumped representations of the building components.

The parameter estimation problem minimizes the objective formulated in Equation 1 and subject to the algebraic differential equations derived from the model dynamics.

$$J = \sum_{z=1}^{\mathcal{Z}} \int_{t_0}^{t_f} e_z(t)^2 dt \tag{1}$$

In Equation 1, t_0 and t_f are the start and final time of the training period and e is the model deviation from the real system.

The Grey-Box Toolbox [27] is chosen to solve the parameter estimation problem. The main reason for this choice is that this toolbox relies on the JModelica platform [31] that allows optimization from Modelica models. Mature libraries exist for the development of complex building and HVAC components in Modelica, like the Buildings [32] or the IDEAS [33] libraries. Predefined models of different building components can be integrated from these libraries enabling a direct coupling between the grey-box building envelope model and any HVAC model available. Additionally, the Grey-Box Toolbox offers a convenient workflow for comparison of different model structures and integrates the necessary functionality to implement the aforementioned forward selection and LHS. The underlying algorithm of JModelica uses direct collocation with CasADi [34]. For more information on how the parameter estimation problem is solved, see [27].

In the centralized approach, contrary to the decentralized approach, the physical connection of zones is introduced as an input to the model identification process using a sparse connectivity matrix. This information allows each zone to be optimized with the temperatures of the neighboring zones as inputs in every step of the forward selection process. The temperature of each adjacent zone is included as an exogenous input and considered perfectly known. A pure resistance branch R1C0 is used to represent the coupling effect between these two zones. A wide search space is allowed for the boundary thermal resistors enabling large values for zones with low thermal coupling.

Each model attempt undergoes two validation tests that filter out models that make no sense from a physical point of view. The first test checks whether the heat flow is being used dynamically and the second test checks whether very large or very small capacities are estimated. It is worth noting that the physical tests are former features of the Grey-Box Toolbox as well.

To estimate the model accuracy of each valid grey-box model attempt of zone z, the Root Mean Square Error (RMSE) of the simulation is calculated in auto- and cross-validation as defined in Equation 2.

$$RMSE_z = \sqrt{\frac{\sum_{k=1}^{\mathcal{M}} (e_{z,k})^2}{\mathcal{M}}} \tag{2}$$

Where:

$$e_{z,k} = y_{z,k} - m_{z,k} \qquad \forall k \in 1, ..., \mathcal{M} \tag{3}$$

In Equation 2, $e_{z,k}$ are the residuals and $\mathcal{M}$ is the number of measurements. In Equation 3, $y_{z,k}$ indicates the model output, $m_{z,k}$ the measurement. All terms refer to zone z at time index k. From all single-zone model attempts, the model showing the lowest RMSE in auto-validation is chosen as the candidate to represent the thermal zone. A summary of the process to identify a single-zone model is represented in Algorithm 1.

Algorithm 1 starts by assigning the set of inputs $\mathcal{I}_z$ to the zone, which include controllable inputs, disturbances, and the set of adjacent boundary zone temperatures $\mathcal{T}_z$. Line 3 starts the forward selection by looping over the set $\mathcal{S}$ of all RC model structures shown in Figure 1. In line 4, a Latin Hypercube Sample $\mathcal{LHS}$ of initial guesses is generated that inherits the parameters from the set of parameters of the best model obtained so far

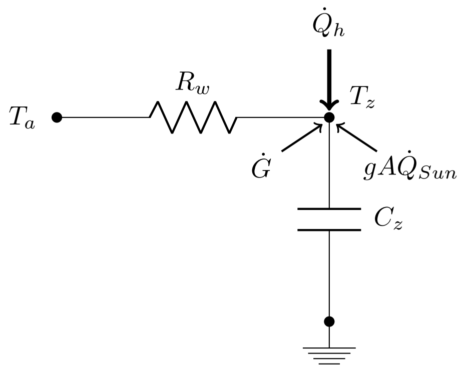

(a) RC circuit of first order: Zone temperature.

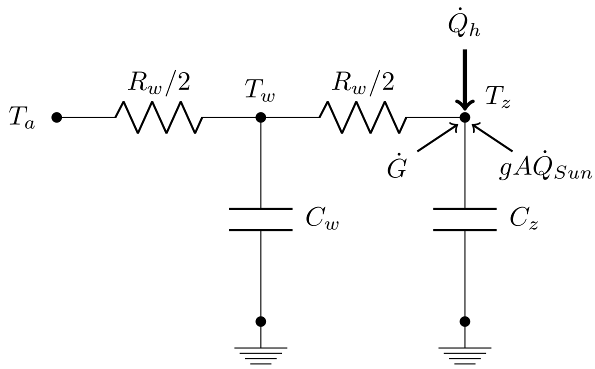

(b) RC circuit of second order: Zone and Wall temperatures.

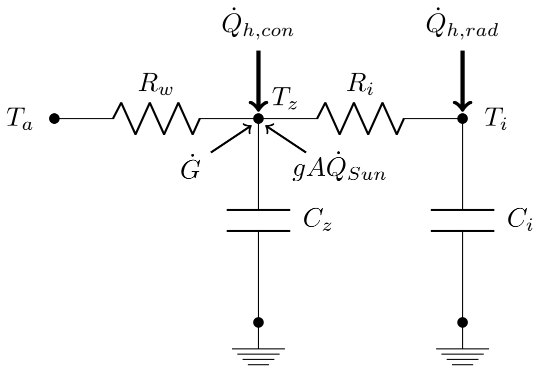

(c) RC circuit of second order: Zone and Internal temperatures.

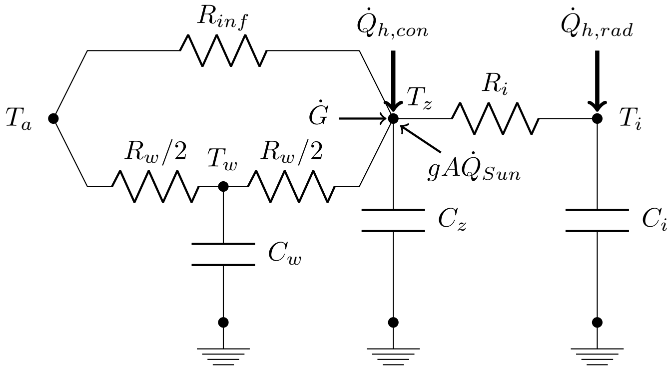

(d) RC circuit of third order: Zone, Wall and Internal temperatures.

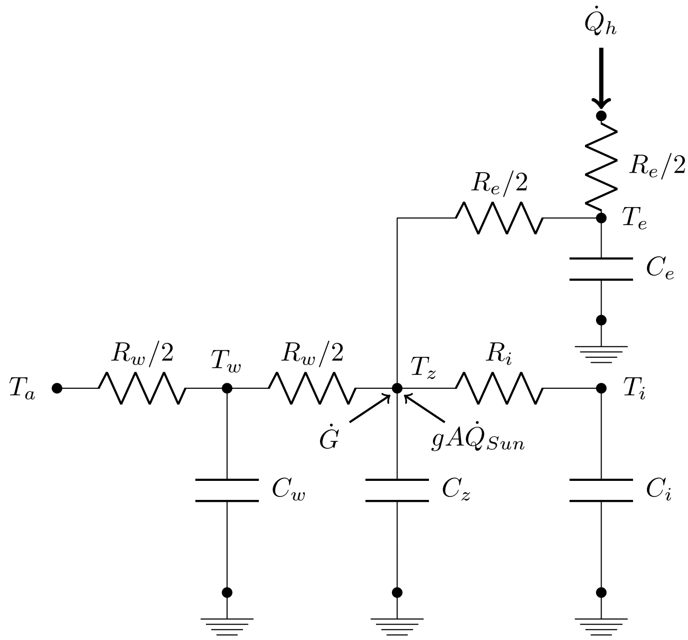

(e) RC circuit of fourth order: Zone, Wall, Internal and embedded temperatures.

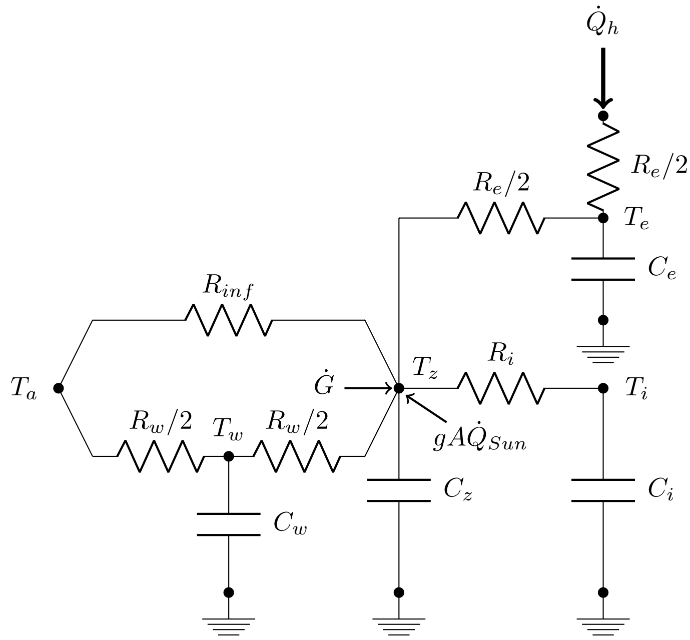

(f) RC circuit of fourth order: Zone, Wall, Internal and embedded temperatures. Zone infiltration is modelled as well in this case.

Figure 1: Grey-box model structures used for the forward selection.

Algorithm 1 Forward selection for zone z

1: $\mathcal{I}_z \leftarrow \mathcal{T}_z, \dot{Q}_z, \dot{G}_z, T_a, \dot{Q}_{Sun}$ ▷ Set inputs to the zone
2: Initialize Θ_z^*, $RMSE^*$
3: **for** s in $\mathcal{S}$ **do** ▷ Every RC structure
4: $\quad \mathcal{LHS} \leftarrow \Theta_z^*$ ▷ Inherit parameters
5: $\quad$ **for** Θ_0 in $\mathcal{LHS}$ **do**
6: $\quad\quad \Theta_z \leftarrow$ Estimate from Θ_0
7: $\quad\quad$ Calculate RMSE
8: $\quad\quad$ **if** $RMSE < RMSE^*$ **then**
9: $\quad\quad\quad \Theta_z^* \leftarrow \Theta_z$
10: $\quad\quad\quad RMSE^* \leftarrow RMSE$
11: $\quad\quad$ **end if**
12: $\quad$ **end for**
13: **end for**

Θ_z^*. Notice that the initial guess values of the air thermal capacitances are always generated from the provided zone volumes. The zone parameters Θ_z are estimated from each initial guess of the sample Θ_0 and lines 7-10 validate and assess the result.

In the centralized approach, the outcoming parameters Θ_z^* are a result of the optimization of each zone independently and considering the temperature input of the adjacent zones as perfectly known. When simulating the model with all zones coupled together, the temperature inputs from each adjacent zone are not an input anymore, but are subject to model mismatch. Therefore, a final parameter estimation of the whole model merged is necessary in the centralized case to balance the error propagation that may be induced by the interaction between the zones. Even though this optimization is expensive because of the large number of parameters to be estimated, it can be performed from the educated initial guess of the parameters obtained from Algorithm 1. For the thermal resistors coupling two zones, the average of the values obtained from each zone is taken.

An example of a possible expected outcome of a three-zone model is shown in Figure 2. In the example, a single state RC structure represents the dynamics of zone 1. The model of zone 2 includes the air temperature and the wall temperature. The model of zone 3 includes the air temperature and the internal temperature. The decision on which RC structure to keep for each zone is based on the process described by Algorithm 1. The parameters R_{12}, R_{13} and R_{23} are the thermal resistors coupling the zones. Their presence enables heat exchange between zones and constitute the difference with a decentralized approach.

2.2. Run-time platform architecture

The aforementioned methodology accounts for a large number of operations that linearly grows with the number of zones. A LHS performs several optimizations per zone structure and multiple zone structures are attempted per zone. Therefore, the total amount of optimizations needed to identify a building model is of the order of nx$\mathcal{S}$x$\mathcal{Z}$, being n the number of initial guesses attempted in the LHS of each RC model architecture, $\mathcal{S}$ the number of RC model architectures tested per zone, and $\mathcal{Z}$ the number of zones in the building. In addition to the optimizations, it is needed to perform the physical tests and the accuracy comparison across all the different attempted zone grey-box cases.

One of the main advantages of the approach followed by Algorithm 1 is that the process of identifying each building zone can be completely decoupled. This allows to run these processes in parallel, removing the computational time dependency on the number of zones. For this purpose, the functionality of identifying a single zone is encapsulated within a Docker container image which is a standard for environment specification to integrate all software dependencies in a lightweight fashion, without the overhead of a virtual machine. Kubernetes is a system that enables orchestration among deployed Docker images. In the envisaged application, one Kubernetes pod is a deployed Docker image assigned to each building zone. An http request sent to each pod with a zone identifier triggers the identification of each building zone. The IPOPT software library [35] for large scale nonlinear optimization is used to solve the parameter estimation problem at each pod. IPOPT implements an interior point line search filter method leading to a sparse linear system. The linear solver used to obtain the solution to the latter system is key for the overall computational performance. For this, the MA57 [36] solver from the HSL library [37] is used since it is the most advanced in-core serial solver provided by HSL [38]. For an in depth comparison between linear solvers for IPOPT see [39].

The fact that the process runs in Kubernetes proves not only that the process computational time is independent of the number of zones, but also that the algorithm could be easily accommodated as a web service. A graphical user interface could serve as a gateway to the performed architecture. In this scenario any user, not necessarily an expert, could provide the few parameters necessary to identify the building model, namely the approximate air volume of each zone, and the physical connections between zones.

Figure 3 shows the run-time process. In the figure, the discontinuous arrows indicate the http request that trigger the identification of a certain zone. The continuous arrows represent the transfer of objects with information defining a building model: thin arrows for single-zone objects and thick arrows for multi-zone.

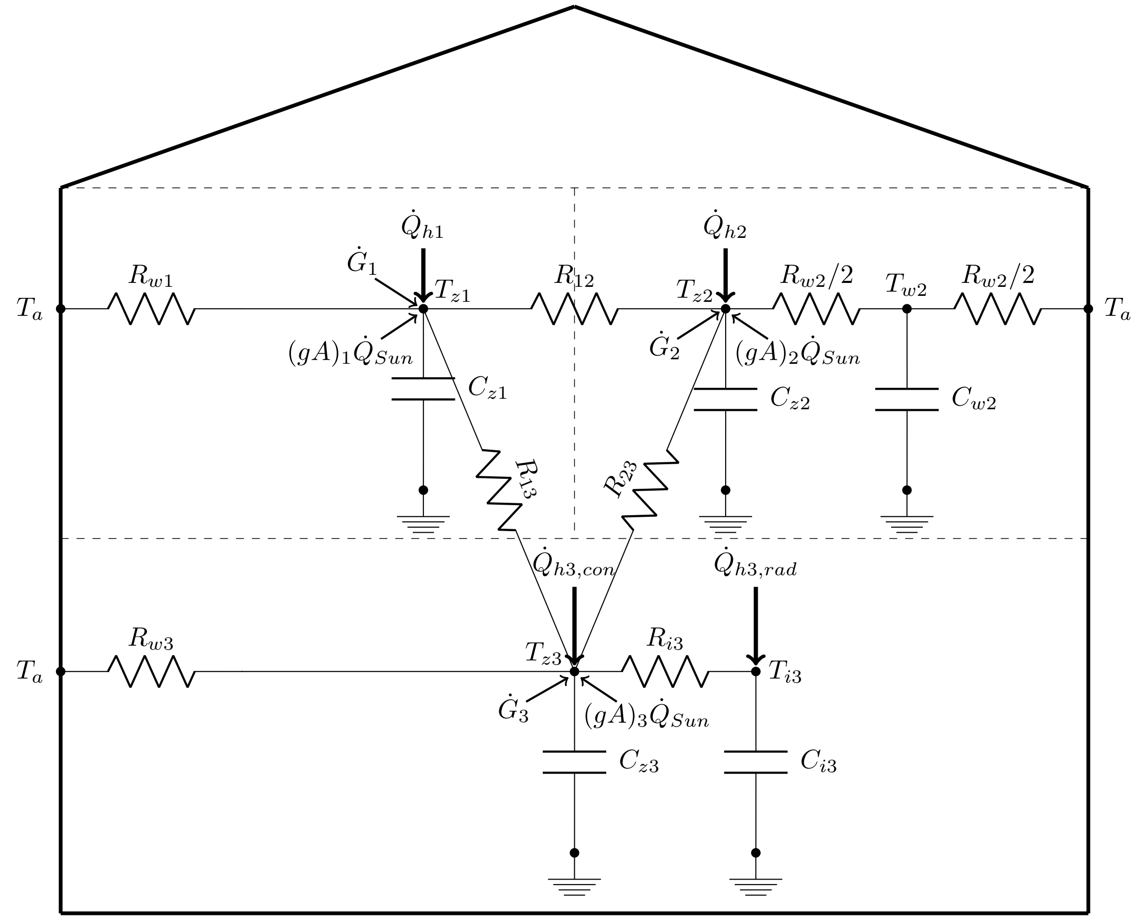

Figure 2: Example of a centralized three-zone grey-box building model.

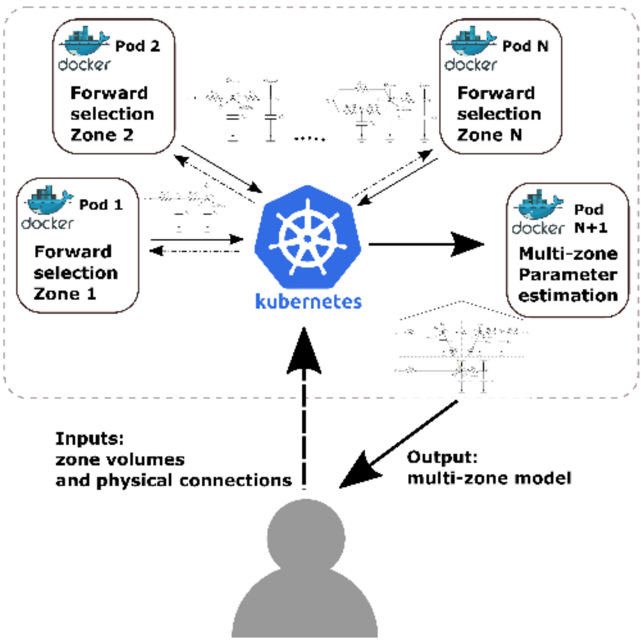

Figure 3: Run-time platform architecture.

In a final step, the obtained parameters of each zone are gathered and coupled together. For the centralized case, this latter model merging all zones is used as an educated guess for the final parameter estimation process in Pod N+1 of Figure 3. In a decentralized approach there is no thermal coupling between zones, and therefore there is no need to optimize the parameters of a building model that integrates all zones, the first step that optimizes the zones independently is thus sufficient in this approach. As a results the output is a multi-zone model where each building zone works independently from the others and its indoor temperature estimate does not affect the estimate of the other zones. Hence, in the decentralized approach the zones simply co-exist in the same model without interaction.

3. Test case description

3.1. The BOPTEST framework

The Building Optimization Performance Test (BOPTEST) is an open source initiative that facilitates the comparison of advanced control strategies in buildings. It provides a software platform architecture that enables to read and overwrite signals in detailed simulation building models by an external controller. It also features data forecast retrieval from the building test cases and a module to compute key performance indicators. The set of test cases is a common playground where researchers and any advanced controller developer can test and unambiguously asses their control strategies. The final goal is to throw light on the best practices for advanced building control. More

details about the BOPTEST initiative can be found in [40].
The BOPTEST framework will provide ten emulator building models, five with an air-based emission system more representative of the American building stock, and five with an hydronic emission system more representative of the European building stock. Each subset widely ranges in complexity from very simple residential single-zone buildings to very complex commercial multi-zone buildings. At the time of writing, the BOPTEST platform is under development and only two very simple test cases are accessible through the BOPTEST repository. However, the authors have already access to the multi-zone residential hydronic building model that will be integrated within the set of BOPTEST test cases. It is important to note that this model is not yet officially integrated within the repository. This means that slight changes could still be made towards the final BOPTEST version.
In any case, the emulator envisaged in this paper has been prepared following the BOPTEST standard and principles. This means that all inputs to the buildings are meant to be realistic inputs that could be accessible from a real building and that the physics behind the model are meant to be representative of a real building system.

3.2. Multi-zone residential hydronic building model

The emulator building model represents a residential French dwelling compliant with the French Thermal regulation of 2012, i.e., the French national building energy regulation. Therefore, the typology is defined to be representative of French new dwellings. Its area is approximately 120 m^2 and consists of six thermal zones that are actively controlled and two unheated zones. The actively controlled thermal zones are: a living room, three bedrooms, a bathroom and a hallway. The unheated zones are a garage and an attic. Figure 4 shows the building layout and sketch of the hydraulic system.
Each building zone in the model includes conventional occupancy, heating, cooling, ventilation, lighting, and internal loads. The building is considered occupied continuously by four adults from 19:00 h until 10:00 h during four weekdays, from 15:00 h until 10:00 h during Wednesdays, and all day long during the weekends. The temperature set-points are defined accordingly to provide thermal comfort during the occupied period and include a night set-back. The internal loads considered are mainly due to lighting and appliances and they are defined according to the buildings thermal regulation of the CSTB (Centre Scientifique et Technique du Bâtiment). There is 30% reduction of the internal loads during nighttime.
The multi-zone residential building is heated using a gas boiler and a hydronic distribution system. There is one radiator per zone with a motorized valve controlled by a PI controller that follows the air zone temperature set point. Only the hallway zone has no valve and its hydronic circuit remains always open to ensure that there is water flow. This is a typical layout that avoids vacuum failures when all valves are closed while the distribution pump is working. The distribution pump frequency and the gas boiler load ratio are controlled through two independent PI controllers to follow the temperature set point of the living room. The boiler also includes an on-off controller with a bandwidth of 0.5°C around the reference signal, an anti-legionella security system, and an anti-short-cycling controller that prevents the boiler from quickly switching states. No domestic hot water is considered in this model.

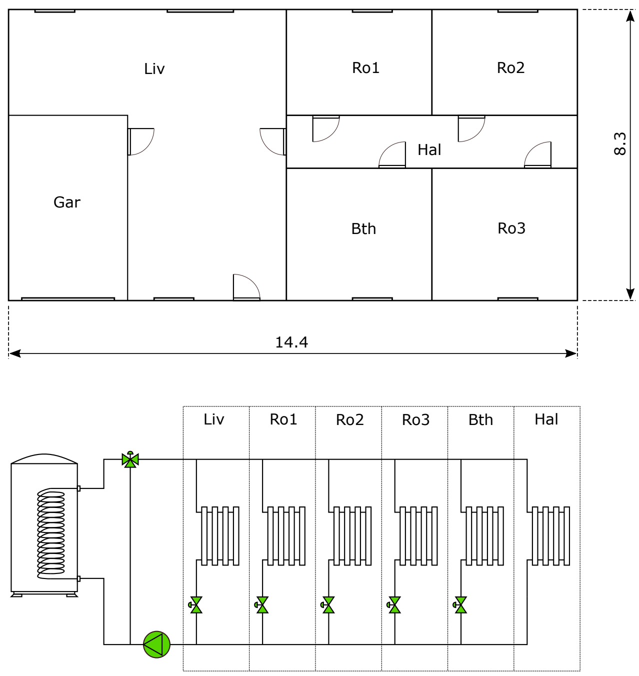

Figure 4: Layout of the multi-zone residential building and sketch of the hydraulic system for the heat distribution. Gar, Liv, Ro1, Ro2, Ro3, Bth, and Hal stand for garage, living room, bedroom 1, bedroom 2, bedroom 3, bathroom, and hallway respectively. The coloured elements in the scheme represent the controllable components through the BOPTEST interface. The dimensions are provided in meters.

The physics of this emulator are modelled using the Buildings Modelica library [32] and integrate: dynamic wall heat transfer, complex fenestration modeling, air infiltration to each zone from ambient air based on constant mass-flow, thermal bridges, inter-zone airflow exchange, inter-zone thermal conductivity, and non-linear convection and radiation models. Particularly, the main model components used from the library are: the `ThermalZones.Detailed.MixedAir` for each building zone, the `Fluid.Boilers.BoilerPolynomial` for the boiler, and the `Fluid.HeatExchangers.Radiators.RadiatorEN442` for the radiators. The insulation levels are based on material properties as defined in Table 1.

Wall type	Description	Transmittance $[W/(m^2K)]$
External	Brick (200mm) and polystyrene (80mm)	0.272
Floor	Hollow block (150mm) and polyurethane (60mm)	0.327
Ceiling	Glass wool (200mm)	0.193
Fenestration	Double glazing with argon and PVC	1.40

Table 1: Material properties of the building emulator model.

In conclusion, the emulator building model integrates a deep level of detail and is considered a high fidelity model able to represent the dynamics of a real building and, therefore, it is used as the real system to retrieve data and assess the control performance in the envisaged application.

4. Results

This section identifies and compares different types of grey-box building models for the test case described in Section 3. The focus is the identification of a centralized model, but a decentralized and a single-zone grey-box model are identified as well for the sake of benchmarking and comparison. First, the methodology explained in Section 2 is implemented to identify a centralized model that integrates the thermal interactions among the zones in the building. Second, the same methodology is implemented but this time without inter-zonal thermal interactions to identify a decentralized model. Finally, a single-zone grey-box building model is identified using only temperature measurements of the living room as representative for the overall building.

The performance of the models is assessed in both simulation and optimal control. An open-loop simulation is performed, i.e. without any state update during one month using cross-validation data. The outputs of the model are compared against the measurements and the residuals are evaluated. Additionally, the influence of the training data length is investigated. For the control performance assessment, the models are used as controller models in the same model predictive controller for one month of co-simulation. Key performance indicators obtained from BOPTEST are used to assess the performance of each optimal controller when using the respective models. Finally, the impact of making the controller models adaptive is studied by retraining their parameters daily.

4.1. Identification of the test case grey-box building models

Data of the heat inputs and temperature evolutions of each zone is required for the identification of grey-box building models. In this study one week of data during the heating season is used with the baseline controller working in the BOPTEST emulator. The measurement data are thus data generated by a detailed emulator model. No extra excitation is applied in order to study how the identification process behaves in the business-as-usual scenario where no additional costs are incurred due to building excitation.

It is decided to use n=20 initial guesses for the latin hypercube sample of each RC model structure attempted, and the $\mathcal{S}$=6 RC model structures of Figure 1 for the forward selection of each zone. In this building with $\mathcal{Z}$=6 thermal zones, the total number of optimization problems to be solved for the hyperparameter optimization of Algorithm 1 is of 20x6x6=720, together with all subsidiary operations related to physical checks, computation of key performance indicators and comparisons among model candidates. The run-time platform architecture described in Section 2.2 facilitates this task.

In the centralized multi-zone approach, a 36R20C model is identified from the hyperparameter optimization process. Adding up the solar admittance to the zones and the initial states, the total number of estimated parameters is of 73. Each of the parameters has a unique optimized value after parameter estimation from an educated initial guess with all zones coupled into a unique model.

In the case of the decentralized model, the identification process results into a 18R18C model for the whole building with a total of 60 parameters to train. In this case each parameter has a unique optimized value as well, but these values come directly from the optimization of the parameters of each individual room since no optimization of the whole model is required as there are no interactions among zones. Notice that the resulting decentralized model does not necessarily contain the same number of states than the centralized one since the forward selection procedure of Algorithm 1 may choose a different zone structure depending on whether adjacent zone temperatures are included as inputs or not.

Finally, the single zone model identified is a 4R3C model with 10 parameters in total, which corresponds to model "d" in Figure 1. Table 2 compares the identified models. It is worth noting that, in all cases, from the obtained parameter values, only the air thermal capacitances keep their strict physical meaning. Other parameters like the thermal resistors of the building envelope or the internal heat capacitances can lose their physical meaning due to their lumped nature and the wide freedom allowed in the search space.

Model	R	C	Total
Centralized	36	20	73
Decentralized	18	18	60
Single-zone	3	3	10

Table 2: Description of the grey-box models features, where R, C and Total stand for the number of thermal resistances, thermal capacitances and the total number of model parameters trained.

4.2. Simulation performance assessment

The simulation performance is assessed in cross-validation with a month of data that immediately follows the training period. Figure 5 shows the temperature evolution in each of the BOPTEST building zones along with the outputs in open loop simulation of the centralized model described in Section 4.1. These outputs come from an open loop simulation of the centralized model when reacting to the same set of inputs and disturbances that were also applied to the BOPTEST building emulator model. The vertical dashed line indicates the end of the training period and beginning of the cross-validation period. Figure 5 demonstrates that the obtained model is able to follow the main dynamics of the system and allows to output a temperature estimate for each of the six building zones. The fitness is steadily worsening while stepping away from the training period where different boundary conditions apply. This is caused by: 1. accumulation of the error over time due to the lack of a state update in the open loop simulation, and 2. the model is slightly over-fitted to the training data. The former can be solved by implementing a proper state estimator, the latter strongly motivates the implementation of an online updating of model parameters to adapt the model to the changing boundary conditions.

In this section it is examined whether this steady bias can be solved using a different training data length or if, on the contrary, a steady bias is inherent to the limited dynamics of grey-box models. According to [25], grey-box models do not usually require long data-sets, in fact, long training periods may lead to overestimation of the thermal mass if the RC model is unable to experience temperature variations. This section investigates whether multi-zone grey-box models can benefit from longer data-periods because of their higher model order. For this purpose, a centralized multi-zone grey-box model is trained starting from the educated initial guess already obtained in Section 4.1 and using a set of different training data lengths. Particularly 3, 5, 7, 10, and 14 days are envisaged as training periods.

For the sake of conciseness, the time-series outputs of each of the obtained multi-zone models are not shown. Instead, the residuals of the models are computed according to Equation 3 and displayed in the boxplots of Figure 6. It should be noted that the simulation period remains the same in all cases as for Figure 5, i.e., one month and one week. Keeping the same model structure, initial guess, and simulation period elucidates the impact of the training period length in the multi-zone grey-box model.

The residuals for all modelled building zones are shown in the boxplots of Figure 6. This means that each boxplot of a multi-zone model contains $\mathcal{Z} \times \mathcal{M}$ points, $\mathcal{Z}$ being the number of zones and $\mathcal{M}$ the number of measurements.

From Figure 6 it is clear that longer training data periods normally lead to smaller fitting errors. The models based on training sets of 10 and 14 days are less biased, but no significant difference is evident for training periods longer than a week. One unanticipated observation is that the obtained parameter set with a period of five days cannot properly follow the system dynamics. This suggests that special care is to be taken of the training set for multi-zone grey-box models, even when starting from an educated initial guess of the parameters. Because of the ill-posed optimization problem of the grey-box parameter estimation problem, the solution is not unique and it is very sensitive to data. Constraining the parameters with physics drives the problem towards a more confined solution space and therefore improves the optimization problem properties.

Finally, the results of the centralized model are benchmarked against those of the decentralized and single-zone models when trained with one week of data. Again, the outputs of each model are shown using a boxplot in Figure 7 during the training and cross-validation data periods. Note that the boxplots of the multi-zone models with $\mathcal{Z} = 6$ are showing six times more points than the one of the single-zone model where $\mathcal{Z} = 1$.

The comparison of the models reveals that the centralized model is slightly more accurate than the decentralized approach, suggesting that thermal interactions between zones should be taken into account in multi-zone grey-box models. The accuracy of the single-zone model does not significantly differ from the accuracy of the multi-zone centralized model in open loop simulation.

4.3. Control performance assessment

Only the simulation performance has been investigated so far. However, the acid test for a controller model is to test its performance when implemented into an actual building predictive controller. Hence, this section analyzes the control performance of a model predictive controller when using the centralized model identified in Section 4.1 as controller model. Then, the performance is compared to the same predictive controller when using the decentralized and single-zone models instead. The performance is also compared to the in-place baseline controller that comes along with the BOPTEST emulator building model, which is a rule based control.

A co-simulation between the MPC and the BOPTEST model is performed using the FastSim toolbox described in [16]. The co-simulation starts the second week of the year, i.e. right after the training period, and lasts for one month. A sampling time of $\Delta t = 900s$ and an optimization horizon of 6 hours are used. Shorter sampling times or longer optimization horizons have not shown any improvement. These and other controller parameters are tuned from experience and according to the time constant of the building. The temperature measurements retrieved from the BOPTEST framework are populated with artificial unbiased noise and a realistic standard deviation of 0.3°C, which is within the measurement error of a regular temperature sensor. A time-varying Kalman filter is implemented to update the controller model states from the new measurements every sampling time. Then, an optimization is performed, and the principle of receding horizon is applied. The set of Equations 4 define the optimal control problem that is solved every time step for the time period between the initial time t_i and the end of the prediction horizon t_h.

$$\min_{x,u} \sum_{z=1}^{\mathcal{Z}} \int_{t=t_i}^{t_h} (u_z(t)^2 + w\delta_z(t))dt \tag{4a}$$

$$\frac{dx(t)}{dt} = f(x(t), u(t), d(t)) \tag{4b}$$

$$\underline{x}_z(t) - \delta_z(t) \leq x_z(t) \leq \overline{x}_z(t) + \delta_z(t) \qquad \forall z \in 1, ..., \mathcal{Z} \tag{4c}$$

$$\underline{u}_z(t) \leq u_z(t) \leq \overline{u}_z(t) \qquad \forall z \in 1, ..., \mathcal{Z} \tag{4d}$$

$$\delta_z(t) \geq 0 \qquad \forall z \in 1, ..., \mathcal{Z} \tag{4e}$$

The optimal control problem is defined in Equations 4, where $u_z(t)$ is the thermal power to zone z that should remain within the lower $\underline{u}_z(t)$ and upper $\overline{u}_z(t)$ limits of the technical con-

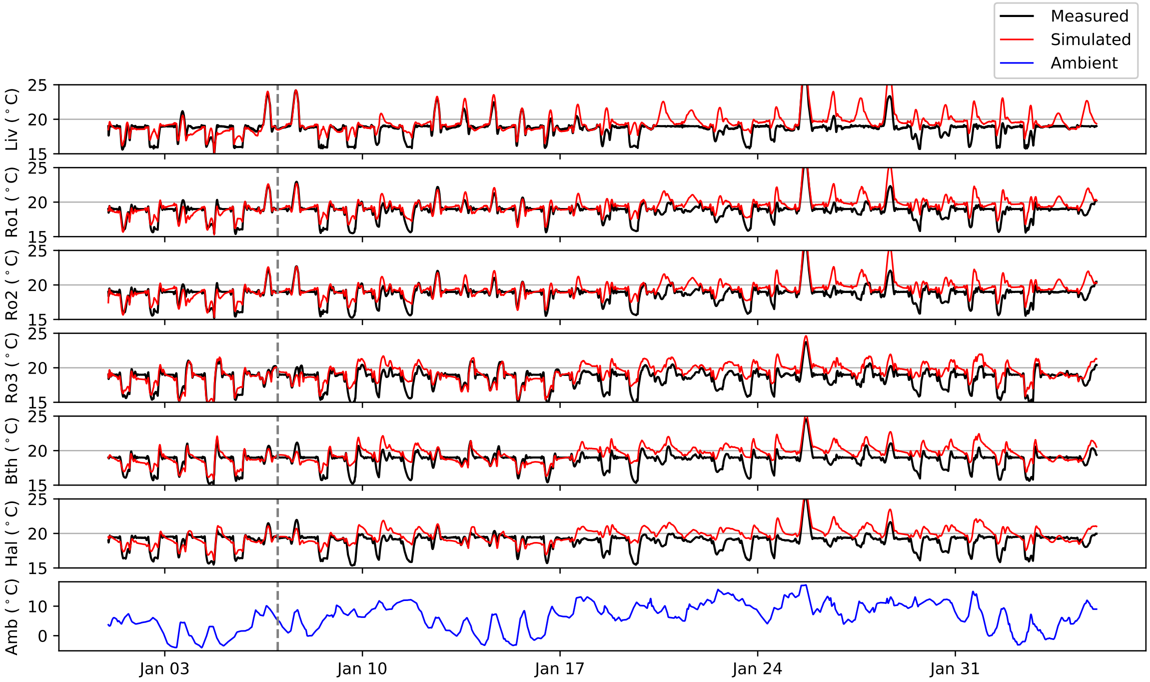

Figure 5: Open-loop simulation temperatures during training and cross-validation periods. The red curves show the output of the centralized model for each building zone. The Black curves show the outputs of the BOPTEST emulator model and are considered as measurements. The dashed vertical lines indicate the end of the training period and beginning of the cross-validation period. The last subplot shows the ambient temperature.

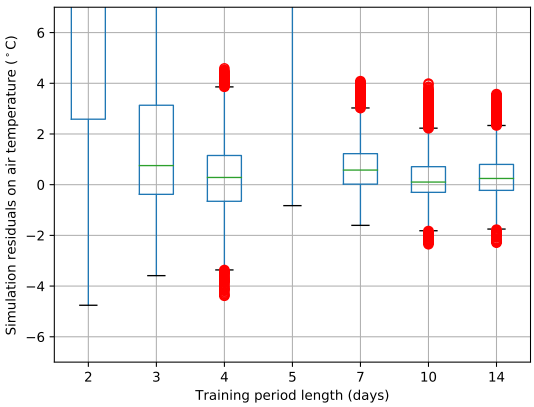

Figure 6: Simulation residuals in auto- and cross-validation for the centralized model identified with different training data period lengths. Notice that the residuals for all six zones are grouped in each of the boxplots.

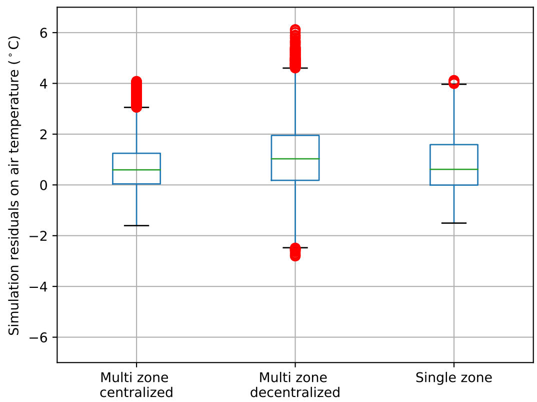

Figure 7: Simulation residuals in auto- and cross-validation when using 7 days of training data for both multi-zone models: centralized and decentralized, as well as for the single-zone model.

straints; $x_z(t)$ is the vector of thermal states of zone z in the controller model; and $\delta(t)$ is the vector of slack variables for the thermal discomfort, i.e., the allowed deviations from the lower $\underline{x}_z(t)$ and upper $\overline{x}_z(t)$ comfort bounds. Notice that, even though discomfort is enabled by Equation 4e, it is penalized with a weight w that accounts for the different orders of magnitude in the terms of the multi-objective function. This weight also prioritizes the comfort within the building rather than the minimization of the energy use. The heat inputs are squared in the objective function to avoid, to the extent possible, peaks and abrupt variations in the heat inputs. This facilitates the post-processing and enhances the overall efficiency of the controller. Finally, f is the model of the building envelope identified in Section 4.1 that estimates the evolution of the thermal states for

a given set of inputs u and disturbances d. For more information on how this optimization problem is solved see [16].
The BOPTEST interface can interact with the emulator model from an external controller by overwriting a set of predefined signals. In this emulator, the signals that are accessible to the external controller are the opening of the valves, and the desired mass-flow-rate of the pump, as explained in Section 3.2. A considerable effort is required to translate the optimized heat inputs into the signals sent to the plant actuators. For this purpose, the outputs of the optimization are post-processed in a subcontroller to provide the heat requested in each zone. Specifically, the pump is switched on/off to work at its nominal value when heat is requested and a PI controller in each zone valve tracks the heat input computed by the MPC.
Again in the interest of benchmarking, a comparison of performance is conducted when using a decentralized multi-zone grey-box model, a single-zone grey-box model, and the in-place baseline controller of the emulator model. This baseline controller is based on a set of predefined heuristic rules and is representative of a typical controller for this type of building. The controller using the decentralized model optimizes the heat input to each zone, but the internal model structure is different. In the case where the single-zone is used as controller model, only the heat input to the overall building is optimized. Then, this heat is distributed among the zones assuming that the valves in the building are static and that their opening position provides a heat distribution proportional to the volume of each individual zone. Therefore, perfect hydraulic balancing in the multi-zone building is assumed for this particular case. The BOPTEST setup allows a fair comparison under the same boundary conditions.
The temperature evolution in each of the six heated zones for each control strategy is shown in Figure 8. Particularly, Figures 8a, 8b, and 8c show the zone temperatures of the MPC when using the centralized, decentralized, and single-zone models, respectively; Figure 8d provides the evolution of the zone temperatures when the baseline controller is implemented. The most striking observation from the figures is that the least amount of comfort violations is achieved with the MPC using the centralized model. The centralized model can see the effect of the thermal interaction among the zones and therefore the MPC overheats some of the zones to keep all of them above the desired minimum temperature. On the contrary, the MPC using the decentralized model tries to maintain the temperature of each zone as close as possible to the allowed minimum individually, resulting in many more thermal violations that are particularly critical if the model overestimates the temperature of any of the zones leading to severe discomfort in that zone. This is the case for the living-room, that shows a permanent discomfort of between 1-2°C, even though the whole heating capacity of this zone is not entirely exploited.
Interestingly, the results of the MPC using the single-zone model are particularly well-performant: it presents few comfort violations compared to the MPC with the decentralized model, and uses less overheating than the MPC with the centralized model. The MPC with the single-zone model proves to accurately determine the overall heat required to heat up the building, most likely because of a confident estimation of the few parameters of this model. This finding suggests that single-zone grey-box models are suitable for MPC in buildings. Nonetheless, strong assumptions are taken. First, the heat distribution to each individual zone is perfectly balanced among the zones in the building which requires a hydraulic balance that may never be perfect in reality. Moreover, the heat distribution will always be static unless a sub-controller is implemented increasing the overall control complexity. Second, all zones follow the same thermal regime meaning that the same zone temperature set points are followed in all zones and variations would only be possible again with the aid of a subcontroller. The MPC implementation with a multi-zone model can automatically adjust the optimal heat distribution towards the zones, even for cases with different zone temperature set points.
A second co-simulation during the month of November is performed to test the models in a validation period with boundary conditions that differ more significantly to those of the training period. Figure 9 shows the ambient temperature and the global solar horizontal irradiation for the two envisaged cross-validation months.
A closer inspection of the results is also carried out for both validation months by quantifying the discomfort and total operational cost. The BOPTEST framework is used to calculate these key performance indicators according to Equations 5.

$$D(t_0, t_f) = \sum_{z \in \mathcal{Z}} \int_{t_0}^{t_f} |\delta_z(t)| dt, \tag{5a}$$

$$C(t_0, t_f) = \sum_{i \in \xi} \int_{t_0}^{t_f} p_i(t) P_i(t) dt \tag{5b}$$

D is the total discomfort in units of Kh; $\mathcal{Z}$ is the set of zones in the multi-zone building; $\delta_z(t)$ is the slack or deviation of the temperature of zone z at time instant t; C is the total operational cost; ξ is the set of equipment in the system with an associated energy use; $p_i(t)$ is the price for energy use of equipment i at time t; and $P_i(t)$ is the instantaneous power use of equipment i.
In our particular case, the two pieces of equipment using energy are a thermal gas boiler and a hydraulic circulation pump of the heating distribution circuit. The values of the discomfort and the total operational cost are computed between the initial t_0 and final t_f simulation times. The quantitative results are shown in Figure 10 with circled markers.
The decentralized approach is the only one that presents higher discomfort than the baseline controller for both validation months. This indicates that single-zone and centralized multi-zone models are more robust in the sense that they are unlikely to incorporate dangling zones with critical discomfort for not being properly modelled. The reason is that, in the centralized and single-zone models, all building states are interconnected ensuring more likely variations among the foreseen zone temperatures.
The centralized and single-zone MPC controllers show a better performance than the baseline controller even though this building case is not particularly interesting for MPC as it does not present large thermal inertia and is not exposed to dynamic

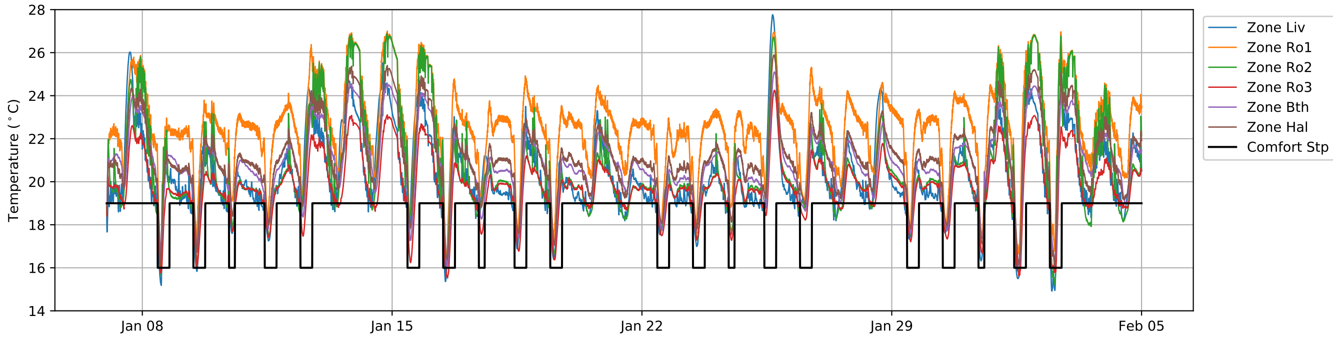

(a) Centralized multi-zone grey-box model into a model predictive controller.

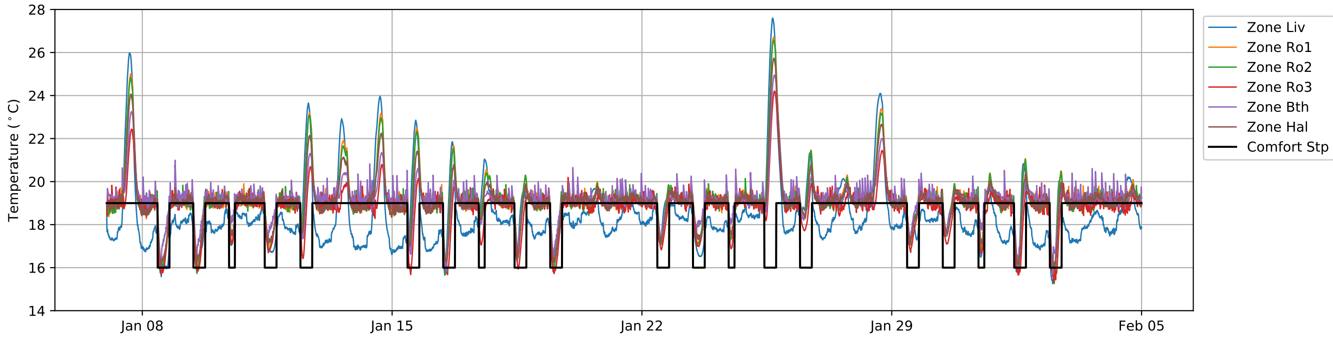

(b) Decentralized multi-zone grey-box model into a model predictive controller.

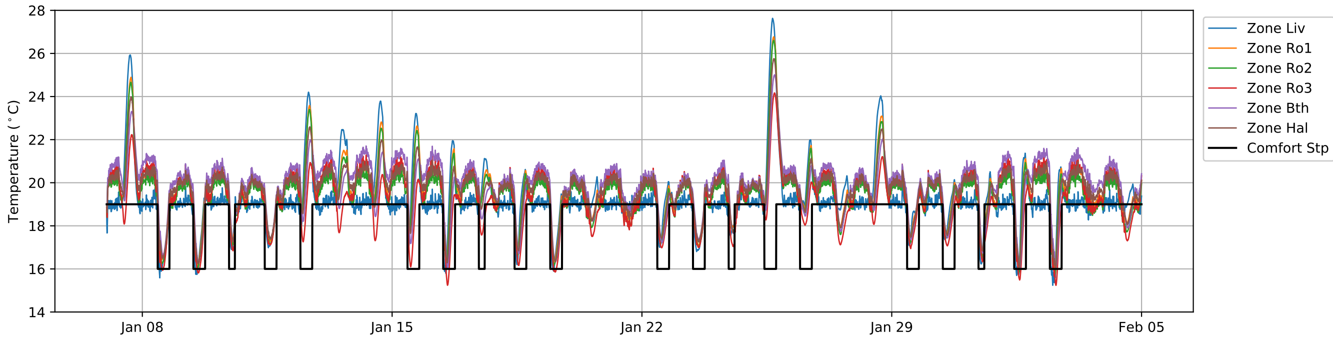

(c) Single-zone grey-box model into a model predictive controller.

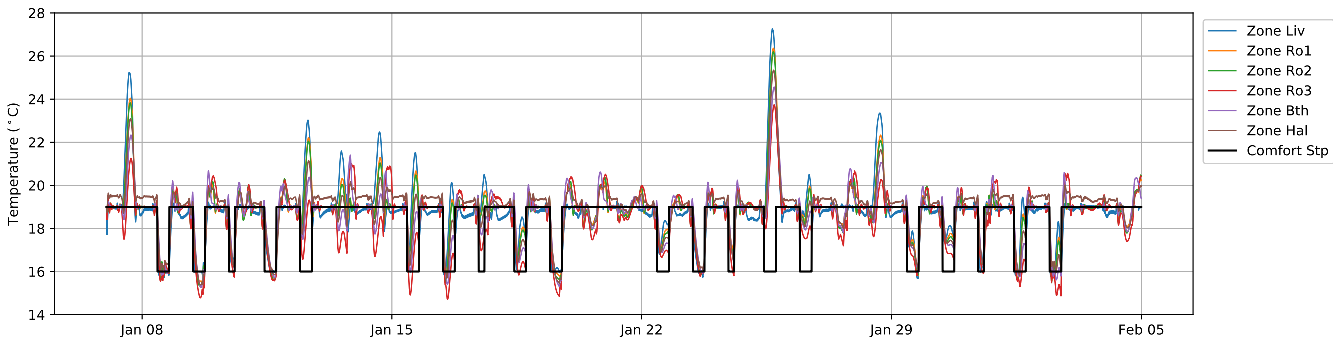

(d) Baseline controller

Figure 8: Temperature evolution in each of the six heated zones for different control strategies.

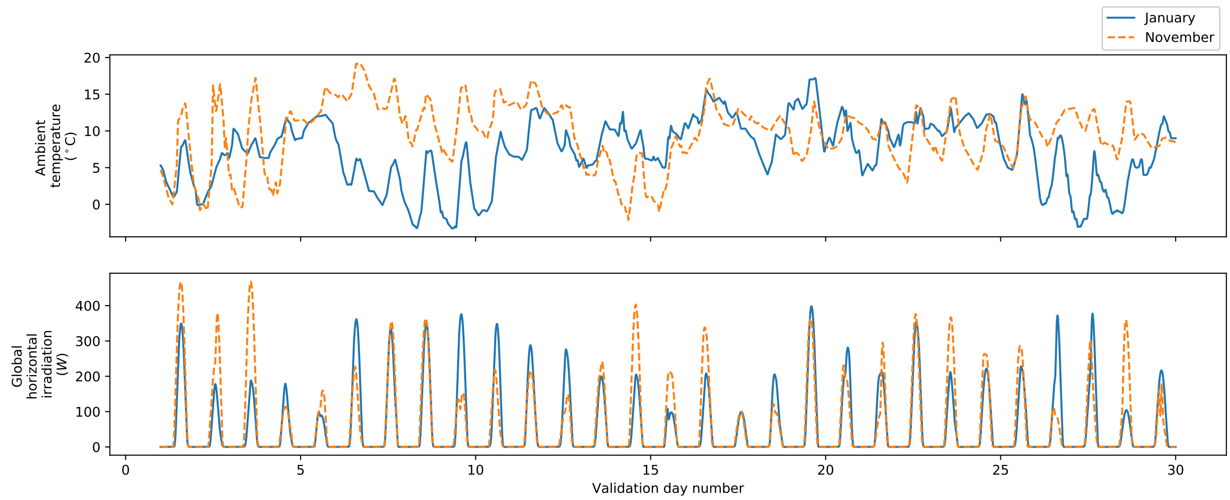

Figure 9: Ambient temperature and solar irradiation for the two envisaged cross-validation periods.

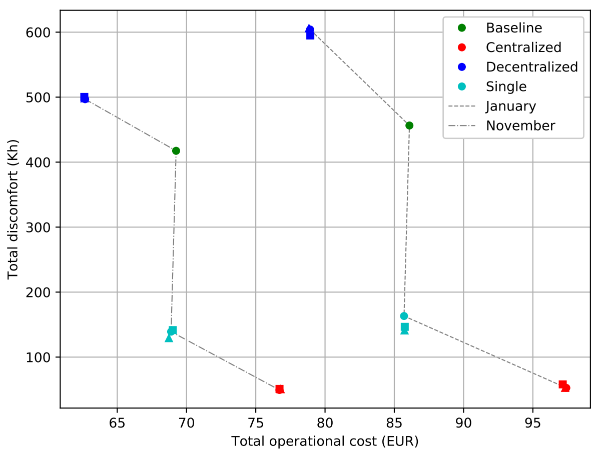

Figure 10: Comparison of control performance.

pricing. Particularly, the MPC using the centralized model in January makes a 88.47% reduction on the thermal discomfort when compared to the baseline controller by increasing the cost only by 13.15%. The validation in November shows similar results with a thermal discomfort reduction of 88.17% at an extra cost of 10.77%. The operational cost of the MPC with the single-zone grey-box model is approximately the same as the baseline controller and the thermal discomfort is reduced by 64.24% in January and by 66.67% in November. From the optimization perspective, the results are favorable to the multi-zone centralized model as it achieves a lower discomfort which is more heavily penalized in the objective function through the weight w at Equation 4a.

4.4. Online updating of parameters

In the previous sections the performance of static controller models has been studied. The parameters were trained once and remained constant over the full control period. This could deteriorate the control performance if the parameters do not suit to the continuously changing boundary conditions. In this scenario erratic conclusions could be obtained on the grey-box building model structures investigated in this paper.

To ensure that this is not the case, the models are made adaptive over time. The objective is twofold: first, to guarantee that the achieved conclusions are not dependent on the training period but only on the model structure. Second, to study the potential benefit of making these models adaptive. Two cases are envisaged: one where the objective function for parameter estimation remains unmodified compared to 1, and another where the objective function is modified to include aspects of the system dynamics.

4.4.1. Unmodified objective function

This case uses the same objective function as in Equation 1 to re-train the model parameters every day at time 00:00 using the previous day as training data. Notice that the model structure remains always the same and that the optimized parameters from the previous day are used as initial guess for every new re-training.

The adaptation is made for all models of Section 4.3 that were implemented into an MPC. Exactly the same control parameters and period are used as in Section 4.3. The performance results are shown in Figure 10 in the same color for each corresponding control case but with a square-shaped marker. No significant performance differences can be highlighted from each control case when making the models adaptive. Even though the parameters undergo substantial changes over the control period, the performance associated to each model structure remains almost the same, confirming the conclusions made in Section 4.3. The controller in January with the single-zone model experiences the largest improvement in the comparison against the baseline controller, with an extra reduction of 3.65% in thermal discomfort compared to the single-zone static model.

To demonstrate the large changes in the parameter values, Figure 11 presents the evolution of the adaptive thermal capaci-

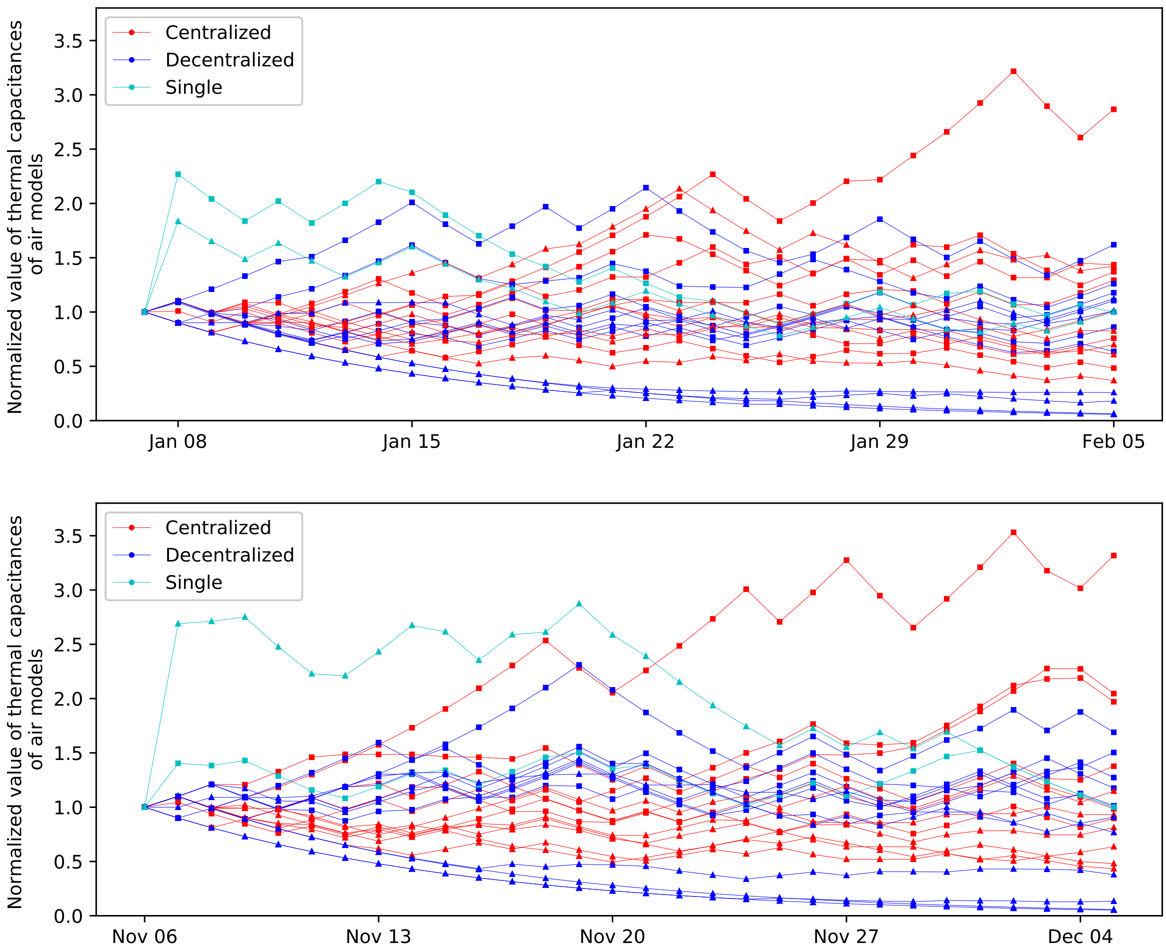

Figure 11: Evolution of the normalized thermal capacitance values for the air nodes when making an online updating of parameters for both validation months: January (top) and November (bottom). Two types of results are shown for each of the three model structures: the results with a squared marker indicate the use of an unmodified objective function; the results with a triangular marker show the use of a modified objective function introducing a penalty for derivative deviations.

tance parameters of each of the zones normalized with their initial value. This initial value is the one already used statically in Section 4.3. Only the air thermal capacitances are presented. The rest of the parameters do not have a physical meaning and experience major changes, which makes the comparison more difficult. The case with unmodified objective is again shown by a squared marker. It is observed that there are parameters that experience important changes, ranging from 5% up to 300% of their original value. These large changes contrast with the small performance variations seen in Figure 10, indicating that the control performance is more intensively influenced by the model structure than by the model parameter values.

4.4.2. Modified objective function

In this case the same daily updating scheme is maintained yet modifying the objective function for parameter estimation to introduce a penalty for derivative deviations. The intention is to reward the models that best follow the system dynamics rather than only get their outputs close to the absolute zone temperatures. The resulting objective function is shown in Equation 6

$$J' = \sum_{z=1}^{\mathcal{Z}} \int_{t_0}^{t_f} e_z(t)^2 + W \left(\frac{dy_z(t)}{dt} - \frac{dm_z(t)}{dt} \right)^2 dt \qquad (6)$$

In this objective function, an additional term accounts for the derivative deviations of the zone model outcome $y_z(t)$ with respect to the zone measured evolution $m_z(z)$. W is a weight to balance both, the absolute and derivative terms of the objective. The resulting optimized parameter values of the model thermal capacitances are represented by the triangle markers in Figure 11. This graph shows that there is a tendency of the parameters obtained with the modified objective function to reduce the thermal capacity of the modelled zones as this leads to faster model reactions and therefore better fitness with the data derivatives. However, this still does not result in a significant improvement in control performance when compared to the online updating of parameters with the unmodified objective function as illustrated in Figure 10, where the results of the modified objective function are again shown with a triangle marker.

5. Conclusion

Grey-box models have a large potential for optimal control, which is reflected in their wide use. The use of physics and data brings advantages from both sides, in the same way it brings challenges. This paper proposes a methodology to facilitate the identification of higher-order grey-box models for multi-zone buildings following a centralized approach. The methodology is implemented for an emulator building of the BOPTEST framework and compared against a decentralized and a single-zone model in both simulation and control performance.

The results show a relevant impact of the used training data length. One week of data is enough to identify the multi-zone building model. When comparing the models in simulation performance, the centralized model slightly outperforms the decentralized model and shows similar accuracy as the single-zone model.

In control performance the differences are more significant: the decentralized model is the one exhibiting worst comfort whereas the centralized model does not overestimate the temperature in any zone which leads to the minimum possible comfort violations. The multi-zone centralized model also outperforms the single-zone model by achieving lower discomfort that is more heavily penalized than the cost in the objective function. However, the single-zone model shows a surprisingly good performance, although a perfect hydraulic balance is assumed. These results suggest that the thermal interactions among zones should be modelled for multi-zone buildings, and that single-zone models are suitable as well if the heat distribution to the zones is properly balanced.

To test the generality of these findings, a daily online updating of the parameters is carried out using both, an unmodified objective function and an objective function that includes penalization for derivative deviations. The small variation in control performance for all cases endorses the previous conclusions by proving that the differences come from the model structures, and not from the parameter values that significantly vary with the training data.

Acknowledgement

The authors gratefully acknowledge the financial support of the European Commission in the H2020 program under Grant Agreement no. 731231 FHP. 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. The work of Javier Arroyo is financed by VITO through a PhD Fellowship under the grant no. 1710754. Finally, the authors wish to thank to David Blum for his greatly appreciated assistance, to Valentin Gavan for providing the simulation model, as well as to Ján Drgoňa and Filip Jorissen, for kindly reviewing the paper.

References

[1] International Energy Agency and International Partnership for Energy Efficiency Cooperation. Building Energy Performance Metrics. Supporting Energy Efficiency Progress in Major Economies. Technical report, OECD/IEA, April 2015.

[2] International Energy Agency. Market Report Series: Energy Efficiency. Technical report, 2018.

[3] K. W. Roth, D. Westphalen, J. Dieckmann, S. D. Hamilton, and W. Goetzler. Energy Consumption Characteristics of Commercial Building HVAC Systems - Volume III: Energy Savings Potential. Technical report, TIAX LLC, 2002.

[4] D. Gyalistras, M. Gwerder, F. Schildbach, C.N. Jones, M. Morari, B. Lehmann, K. Wirth, and V. Stauch. Analysis of Energy Savings Potentials for Integrated Room Automation. In *Clima - RHEVA World Congress*, Antalya, Turkey, May 2010.

[5] R. De Coninck and L. Helsen. Quantification of flexibility in buildings by cost curves – methodology and application. *Applied Energy*, 162:653–665, 01 2016.

[6] J. Arroyo, S. Gowri, F. De Ridder, and L. Helsen. Flexibility quantification in the context of flexible heat and power for buildings. In *Proc. of the Federation of European Heating, Ventilation and Air Conditioning associations conference in Brussels*. REHVA - Federation of European Heating, Ventilation and Air Conditioning Associations, 2017.

[7] M. Sourbron, C. Verhelst, and L. Helsen. Building models for model pre-

dictive control of office buildings with concrete core activation. *Journal of Building Performance Simulation*, 6(3):175–198, 2013b.

[8] X. Li and J. Wen. Review of building energy modeling for control and operation. *Renewable and Sustainable Energy Reviews*, 37:517–537, 2014.

[9] S. Prívara, J. Cigler, Z. Váňa, F. Oldewurtel, C. Sagerschnig, and E. Žáčeková. Building modeling as a crucial part for building predictive control. *Energy & Buildings*, 56:8–22, January 2013.

[10] P. Bacher and H. Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings*, 43(7):1511–1522, 2011.

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

[12] R. De Coninck. *Grey-Box Based Optimal Control for Thermal Systems in Buildings-Unlocking Energy Efficiency and Flexibility*. PhD, KU Leuven, Leuven, Belgium, 2015.

[13] R. De Coninck, F. Magnusson, J. Å kesson, and L. Helsen. Grey-box building models for model order reduction and control. In *Proceedings of the 10th International Modelica Conference*, pages 657–666, Lund, Sweden, 03 2014.

[14] R. De Coninck and L. Helsen. Practical implementation and evaluation of model predictive control for an office building in brussels. *Energy and Buildings*, 111:290 – 298, 2016.

[15] M. Vande Cavey, R. De Coninck, and L. Helsen. Setting up a framework for model predictive control with moving horizon state estimation using JModelica. In *10th International Modelica Conference, March 10-12, 2014, Lund, Sweden*, 2014.

[16] J. Arroyo, B. van der Heijde, F. Spiessens, and L. Helsen. A Python-Based Toolbox for Model Predictive Control Applied to Buildings. In *Proc. of the Intelligent Building Operations (IBO) Workshop in Purdue*, Purdue University, West Lafayette, Indiana, July 2018.

[17] I. Hazyuk, C. Ghiaus, and D. Penhouet. Optimal temperature control of intermittently heated buildings using model predictive control: Part i – building modeling. *Building and Environment*, 51:379 – 387, 2012.

[18] E. O'Dwyer, L. De Tommasi, K. Kouramas, M. Cychowski, and G. Lightbody. Modelling and disturbance estimation for model predictive control in building heating systems. *Energy and Buildings*, 130:532 – 545, 2016.

[19] T. Bohlin and S.F. Graebe. Issues in nonlinear stochastic grey-box identification. *IFAC Proceedings Volumes*, 27(8):1175 – 1180, 1994. IFAC Symposium on System Identification (SYSID'94), Copenhagen, Denmark, 4-6 July.

[20] E. Žáčeková, S. Prívara, and Z. Váňa. Model predictive control relevant identification using partial least squares for building modeling. In *2011 Australian Control Conference*, pages 422–427, Nov 2011.

[21] S. Prívara, J. Široký, L. Ferkl, and J. Cigler. Model predictive control of a building heating system: The first experience. *Energy and Buildings*, 43(2):564–572, 2011.

[22] C. Verhelst. *Model predictive control of ground coupled heat pump systems for office buildings*. PhD, KU Leuven, Leuven, Belgium, 2012.

[23] D. Picard, J. Drgoňa, M. Kvasnica, and L. Helsen. Impact of the controller model complexity on model predictive control performance for buildings. *Energy and Buildings*, 152:739 – 751, 2017.

[24] D. Picard, M. Sourbron, F. Jorissen, Z. Váňa, J Cigler, L. Ferkl, and L. Helsen. Comparison of model predictive control performance using grey-box and white-box controller models of a multi-zone office building. In *4th International High Performance Buildings Conference*, Purdue, 2016.

[25] K. Arendt, M. Jradi, H. R. Shaker, and C. Veje. Comparative analysis of white-, gray- and black-box models for thermal simulation of indoor environment: Teaching building case study. In *Proceedings of the 2018 Building Performance Modeling Conference and SimBuild co-organized by ASHRAE and IBPSA-USA*, pages 173–180. ASHRAE, 9 2018.

[26] T. Berthou, P. Stabat, Salvazet R., and Marchio D. Development and validation of a gray box model to predict thermal behavior of occupied office buildings. *Energy and Buildings*, 74:91 – 100, 2014.

[27] R. De Coninck, F. Magnusson, J. Åkesson, and L. Helsen. Toolbox for development and validation of grey-box building models for forecasting and control. *Journal of Building Performance Simulation*, 9(3):288–303, 2016.

[28] C. Agbi, Z. Song, and B. Krogh. Parameter identifiability for multi-zone building models. In *Proceedings of the IEEE Conference on Decision and Control*, pages 6951–6956, 2012.

[29] J. Cai. A practical and scalable inverse modeling approach for multi-zone buildings. In *9th International Conference on System Simulation in Buildings, Liege, December 10-12, 2014*, Liege, 2014.

[30] E. Atam and L. Helsen. Control-oriented thermal modeling of multizone buildings: Methods and issues: Intelligent control of a building system. *IEEE Control Systems Magazine*, 36(3):86–111, June 2016.

[31] Modelon AB. Jmodelica.org user guide version 2.2. Technical report, 2018.

[32] M. Wetter, W. Zuo, T. Nouidui, and X. Pang. Modelica buildings library. *Journal of Building Performance Simulation*, 7(4):253–270, 2014.

[33] R. Baetens, R. De Coninck, F. Jorissen, D. Picard, L. Helsen, and D. Saelens. OpenIDEAS - An open framework for integrated district energy simulations. In *Proceedings of Building simulation 2015*, Hyderabad, India, 2015.

[34] J. Andersson, J. Gillis, G. Horn, J. Rawlings, and M. Diehl. CasADi – a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019.

[35] Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, may 2006.

[36] I. S. Duff. MA57 - A code for the solution of sparse symmetric definite and indefinite systems. *ACM Transactions on Mathematical Software*, 30(2):118–144, jun 2004.

[37] HSL. HSL - a collection of fortran codes for large-scale scientific computation, 2011.

[38] Byron Tasseff, Carleton Coffrin, Andreas Wachter, and Carl D. Laird. Exploring benefits of linear solver parallelism on modern nonlinear optimization applications. In *arXiv:1909.08104 [math.OC]*, September 2019.

[39] A. Kelman. *Utilizing Problem Structure in Optimization Algorithms for Model Predictive Control*. PhD thesis.

[40] D. Blum, F. Jorissen, S. Huang, Y. Chen, J. Arroyo, K. Benne, Y. Li, V. Gaven, L. Rivalin, L. Helsen, D. Vrabie, M. Wetter, and M. Sofos. Prototyping the BOPTEST framework for simulation-based testing of advanced control strategies in buildings. In *Proceedings of the 16th IBPSA Conference*, Rome, 2019.