A modelling workflow for predictive control in residential buildings

E. O'Dwyer, E. Atam, P. Falugi, E. C. Kerrigan, M. A. Zagorowska and N. Shah

5.1 Introduction

With electrification of heating and transport, as well as an increased reliance on varying renewable generation sources, buildings will be required to act in a flexible manner, allowing for the shifting of demands and coordination of local low-carbon generation [1]. Rather than simply acting as a rigid demand that must be satisfied by the grid, buildings can act as useful system of components by offering the opportunity for flexibility provision [2]. By exploiting the energy storage capacity of a building's fabric and using information about occupancy and comfort requirements, the heat/cold supplied to a building can be shifted away from times of high demand [3]. Furthermore, storage devices including batteries [4], thermal stores [5] and electric vehicles [6], along with renewable generation sources (particularly so-

E. O'Dwyer
Department of Chemical Engineering, Imperial College London, e-mail: e.odwyer@imperial.ac.uk

E. Atam
Department of Electrical and Electronic Engineering, Imperial College London, e-mail: ercan.atam@gmail.com

P. Falugi
Department of Electrical and Electronic Engineering, Imperial College London, e-mail: p.falugi@imperial.ac.uk

E. C. Kerrigan
Department of Electrical and Electronic Engineering, Department of Aeronautics, Imperial College London, e-mail: e.kerrigan@imperial.ac.uk

M. A. Zagorowska
Department of Electrical and Electronic Engineering, Imperial College London, e-mail: m.zagorowska@imperial.ac.uk

N. Shah
Department of Chemical Engineering, Imperial College London, e-mail: n.shah@imperial.ac.uk

This research was supported by the EPSRC under the grant EP/V012053/1.

lar generation), can be managed appropriately to allow power to be absorbed from the grid and supplied back to the grid as needed. Aside from the physical storage and generation assets, communication technology and computational techniques are needed to determine optimal control decisions for the various system components in a timely fashion.

A technique commonly proposed for the task of optimal real-time decision-making in the building sector is Model Predictive Control (MPC) [7]. An MPC-based strategy uses a predictive model of a system's dynamics and constraints to form an optimisation problem, which is solved to determine input trajectories over a receding horizon. The ability of MPC to incorporate future behaviour in the decision-making process makes it particularly suited to the problem of optimally time-shifting energy demands. Furthermore, the manner in which system constraints can be explicitly applied makes MPC a natural fit for the role of building energy management [8]. Despite the promising potential of MPC to enable demand-side flexibility, the uptake of the technology in the sector remains limited. This is particularly challenging for the domestic sector, as a high-level of adoption would be required to ensure sufficient *shiftable* demand is available to encourage worthwhile interaction with the grid.

A commonly cited obstacle to its wider roll-out is the modelling challenge [9, 10]. Buildings can vary greatly in their design and are inherently formed of multiple inputs and outputs. Furthermore, the way in which a building is used (particularly in a domestic setting) is highly dependent on the individual occupants and one-size-fits-all solutions are unlikely to provide adequate performance. Modelling techniques that are not contingent on large time investments from building modelling experts would greatly facilitate the deployment of MPC in the sector, thus expediting the shift towards more flexible operation and a lower-carbon energy sector.

This chapter establishes a systematic workflow, from detailed building modelling based on standard tools, to MPC design. This workflow begins with detailed physics-based modelling followed by low-order control-oriented modelling and heating system component modelling. The workflow can be summarised as follows:

1. The building shape is drawn and building materials are defined.
2. A detailed building simulation model is created.
3. The underlying mathematical representation of the model is extracted.
4. Option A: The model complexity is reduced using truncation techniques for control implementation
5. Option B: Low-complexity models are derived from data.
6. Design and analysis of control strategy is carried out using detailed simulation modelling platform.

A case study is used to illustrate the application of the proposed methods, whereby the performance of MPC strategy for managing a heating system is analysed and the ability to interact with an external price signal is demonstrated. The contribution of the work is primarily as a guide for practitioners applying MPC to residential buildings, by describing and illustrating, in detail, modelling principles that can underpin a successful strategy implementation. Furthermore, some of the remaining open

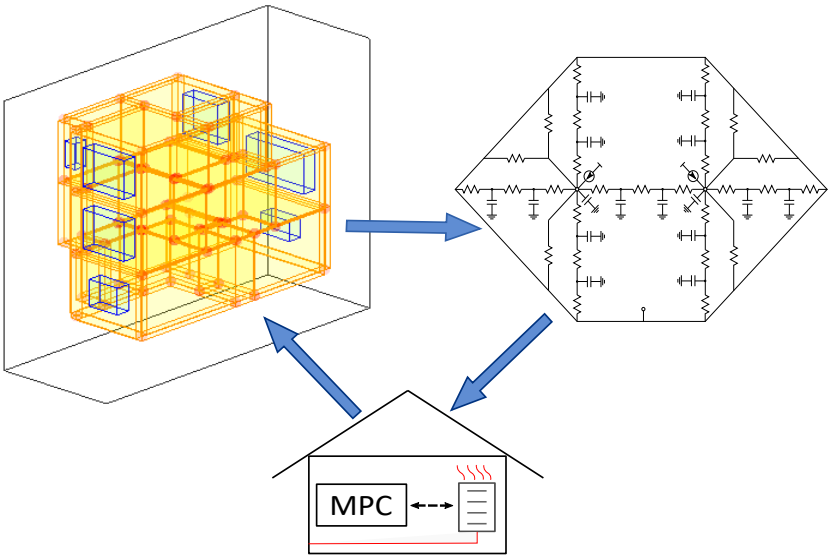

Fig. 5.1: Modelling workflow for MPC design in buildings

challenges and key design questions are emphasised, encouraging future research in the area.

In the second section, the development of detailed physics-based modelling is considered both from a theoretical and software perspective. By commencing with such an approach, comprehensive analyses and design evaluations can be carried out, while the use of industry-standard software reduces the need for specialist expertise. The third section examines the generation of low-order, low-complexity representations of building models to better fit the optimisation-based focus of an MPC strategy. Low-order modelling techniques are outlined from two perspectives. Firstly, a model-based approach is summarised whereby a detailed physics-based model is truncated to emphasise the dynamics most relevant to the timescales of an MPC controller. Following this, data-centric approaches based on system identification are considered. Finally, the outlined techniques are combined in the design of both a centralised and a decentralised MPC strategy, which minimises energy cost while maintaining user comfort. By outlining suitable tools, models and design decisions, the chapter is intended to provide a systematic methodology for application of MPC in the sector.

The chapter is structured as follows: Section 5.2 briefly describes the current state and future outlook of the field, Section 5.3 outlines the production of a detailed simulation environment, Section 5.4 describes methods for deriving low-complexity control-oriented models and Section 5.5 introduces suitable MPC formulations for a residential case study. The methods are illustrated through the use of a case study in

which an MPC strategy for a heating system in a three-storey dwelling is developed and evaluated.

5.2 Status Quo, Challenges and Outlook

The need for new tools and methods to tackle the modelling challenge faced in the building energy domain has been recognised, leading to the development of interoperable tools that can bridge the gap between building simulation software and control design. The Building Controls Virtual Test Bed (BCVTB) [11], for example, is an environment that enables co-simulation of different building modelling tools and programming languages (e.g. EnergyPlus, Modelica, Matlab) for the purpose of control design, within the open-source Ptolemy II framework. The Building Resistance-Capacitance Modelling (BRCM) toolbox [12] is focussed on MPC for buildings and can interface with EnergyPlus (as done in this chapter). Aside from software tools, modelling techniques for building simulation come in different forms, from the physics-based focus of white-box techniques [13], to black-box or machine learning techniques that seek to model building behaviour without in-depth knowledge of the building structure [14]. Furthermore, grey-box techniques that use data to parameterise models are often proposed with the advantage of maintaining a desirable model structure without needing detailed knowledge of the building fabric, with a review of methods for matching simulation models to data presented in [15]. An overall review of building modelling methods and tools can be found in [16].

Despite progress in these areas, challenges remain. Data-driven techniques can be used to improve model accuracy and reduce modelling effort, but care must be taken with regard to system excitation. As noted in [17], closed loop data may not provide sufficient richness for successful parameter identification. Carrying out forced-response tests on a building can be a challenge if the building is occupied and would add to the implementation cost, which is already a hindrance to wider scale-up [18]. The challenge for researchers at present is not to demonstrate that MPC can lead to improved performance in a building, but to demonstrate that it can be implemented at scale without excessive cost. Nonetheless, recent development in data-driven predictive control approaches (e.g. [19, 20] may be promising as they can provide low implementation effort while maintaining certain robustness guarantees. A review on data-driven building control methods can be found in [21]. Furthermore, as simulation tools become more common for building design, it has been recognised that such tools can also be useful for operational analysis. In this manner, the digital twin concept (see for example the Gemini principles [22]) fits with these goals. In [23], the BCVTB environment is deployed as a digital twin for a distributed set of building energy and smart city assets. Using a digital representation of a system enables control strategies to be analysed in-silico prior to deployment, while running a digital twin in parallel allows for operational insights to be attained that may otherwise have been missed.

5.3 Physics-based modelling of the building

There are different reasons for developing a detailed physics-based model of a building and its components. Many methods and international standards have been developed to predict the expected energy performance of a building, thus allowing the evaluation of design decisions and retrofit options. Such methods tend to be focussed on long-term steady-state performance, mostly independent of the control approach taken (the Tabula webtool [24] is an excellent example of a database of such models based on thermal properties and heat balances). The approaches discussed in this section are distinct from these insofar as they are specifically focussed on the impact of control techniques and, as such, require the transient behaviour over a wider range of timescales to be captured. The thermal response of the air in the building (which has a relatively low heat capacity) must be captured as well as the thermal response of larger concrete slabs (which have relatively high heat capacities). With such models, it should be possible to carry out detailed analysis and evaluation in-silico prior to implementation of a strategy. Simulation allows for control techniques to be compared without being impacted by changeable external influences that cannot otherwise be controlled (most notably the weather). The commonly used resistance-capacitance (RC) methodology is detailed here along with software packages that can be used for implementation.

5.3.1 Modelling background

The equations and concepts underpinning the development of an appropriate simulation model are first described. Suitable methods are outlined for modelling the thermal fabric, the heating system and the occupants in the building. These components can be then brought together in a single simulation model.

5.3.1.1 Building fabric: the resistance-capacitance model

When some knowledge of the building fabric dimensions and materials are known, a thermal model of the building can be formed as structure analogous to an electrical circuit composed of resistors and capacitors [25, 26]. In such an RC-network, each component of the physical structure can be represented as a system of resistors and capacitors with parameters corresponding to the thermal properties of the structure, while each room or airspace can be represented by a single capacitor. Heat flows are represented in such a circuit as current, while temperature differences are represented as voltages. To illustrate the method, a model of a single wall connecting a zone of temperature T_z to the external air at temperature T_e is shown in Fig. 5.2. In this diagram, the zone capacitance is denoted C_1, while the three wall resistances are denoted R_1, R_2 and R_3. Two intermediate wall temperatures are defined (T_{W1} and

T_{W2}), with the wall capacitances given as C_2 and C_3. The window is represented by a single resistance, given as R_4. The heat supply to the zone given by Q_{in}.

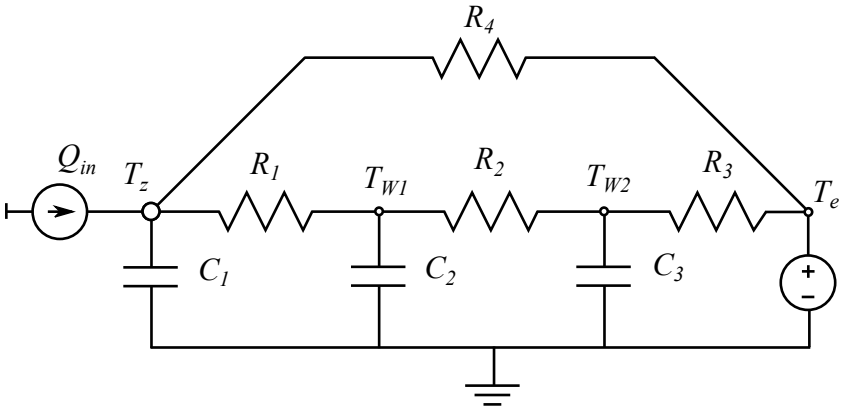

Fig. 5.2: RC-network representation of single external wall with window

A key feature of this form is that the full building can be represented as a linear dissipative state-space system. In a such a system, the internal states correspond to wall and air temperatures, heat supplies to the zones can be regarded as controlled inputs while external influences, such as the ambient temperature or the solar radiation incident on the walls, can be categorised as disturbances to the model. With this approach, the wall of Fig. 5.2 can be represented in a standard form by the following set of differential equations:

$$\begin{bmatrix} \dot{T}_z(t) \\ \dot{T}_{W1}(t) \\ \dot{T}_{W2}(t) \end{bmatrix} = \begin{bmatrix} \frac{-1}{C_1}\left(\frac{1}{R_1}+\frac{1}{R_4}\right) & \frac{1}{C_1R_1} & 0 \\ \frac{1}{C_2R_1} & \frac{-1}{C_2}\left(\frac{1}{R_1}+\frac{1}{R_2}\right) & \frac{1}{C_2R_2} \\ 0 & \frac{1}{C_3R_2} & \frac{-1}{C_3}\left(\frac{1}{R_2}+\frac{1}{R_3}\right) \end{bmatrix} \begin{bmatrix} T_z(t) \\ T_{W1}(t) \\ T_{W2}(t) \end{bmatrix} + \begin{bmatrix} \frac{1}{C_1} \\ 0 \\ 0 \end{bmatrix} Q_{in}(t) + \begin{bmatrix} \frac{1}{C_1R_4} \\ 0 \\ \frac{1}{C_3R_3} \end{bmatrix} T_e(t) \tag{1}$$

$$T_z(t) = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} T_z(t) \\ T_{W1}(t) \\ T_{W2}(t) \end{bmatrix}. \tag{2}$$

By representing all walls, floors and ceilings in a similar manner, the full building can be given a general form, where the wall and zone air temperature states are given by the vector $\mathbf{x}(t)$ and the measured zone air temperatures are given by the vector $\mathbf{y}(t)$. Controlled inputs are given by $\mathbf{u}(t)$, disturbances (e.g. the external temperature,

ground temperature, solar incident radiation and internal gains) given here as $\mathbf{d}(t)$, C is a matrix mapping states to measurements and A, B and E are parameter matrices:

$$\dot{\mathbf{x}}(t) = A\mathbf{x}(t) + B\mathbf{u}(t) + E\mathbf{d}(t) \tag{3}$$

$$\mathbf{y}(t) = C\mathbf{x}(t). \tag{4}$$

This is a standard linear state-space representation. Generally (4) can also contain an additional term denoted $D\mathbf{u}(t)$ representing the direct feedthrough from input to output. This is omitted as no such feedthrough is present in the model.

While (3) describes the flow of heat through the building fabric, the heating or cooling systems associated with the inputs can vary in form and appropriate modelling methods will depend on the specific systems used. For example, the control variables themselves may be system flow rates (in a radiator-based system) or air-flows (in an air-based system) and the heating or cooling supply in a particular zone can become a function of the product of the controlled flow rate and the air temperature of the zone, which is a model state. Thus, the model of (3) would become bilinear. Furthermore, the system may include boilers, chillers or heat pumps, which may introduce nonlinearities in the overall system behaviour as discussed in the next section.

5.3.1.2 Heating system: heat pump modelling

As the set of potential heating or cooling system components is open ended, an exhaustive list of different modelling techniques for different technologies is not presented here. Nonetheless, to demonstrate the use of the techniques and the modelling toolchain established here via case studies, it is assumed that a heating system is present in the building, with heat supplied by an electric heat pump. The use of predictive control for energy management of a heat pump-based system is sensible, since the lower-system temperatures preferable for heat pump utilisation can lead to longer heat-up times, thus making predictive strategies favourable. Furthermore, electrification of heating is likely to become ever more prevalent, leading to an increased use of heat pumps, and by choosing a heat pump as the heat source, the interaction between the electrical grid and the building can be better illustrated. Advanced control strategies such as MPC are needed to manage this interaction, once again making a heat pump the most suitable choice for illustrating an MPC-based modelling workflow. A suitable heat pump model is shown in this section.

Whether using an air or ground source heat pump, the relationship between electrical power consumed (given here as P_e) and the heat supplied to the heating system (given here as Q_h) is defined by its coefficient of performance (COP) as:

$$COP = \frac{Q_h}{P_e}. \tag{5}$$

The COP is often assumed to be a constant for modelling simplicity, however more detailed representations consider COP to be a function of a number of operational variables, of which perhaps the most important is the *lift*. In the context of a heat pump, lift is the temperature difference between the source temperature (ground or ambient air) and the sink temperature (the heating system flow temperature). Lower lift values tend to correspond to higher COP values. This makes lower temperature heating systems favourable, while also implying that an efficiency reduction can be expected during cold weather. While manufacturers may supply spec sheets with COPs included, field tests of heat pumps tend to show a wide variation in performance depending on context. To account for this, an extensive review of domestic heat pumps was carried out by [27], whereby COP values taken from manufacturers data, field trials and experimental studies were plotted as a function of lift. The resulting empirical relationships derived for air-source (ASHP) and ground-source (GSHP) heat pumps are as follows (where lift is denoted ΔT):

$$COP_{ASHP} = 6.81 - 0.121\Delta T + 0.00063\Delta T^2, 15 \leq \Delta T \leq 60, \tag{6}$$

$$COP_{GSHP} = 8.77 - 0.15\Delta T + 0.000734\Delta T^2, 12 \leq \Delta T \leq 60. \tag{7}$$

These equations are suitable for the purpose of a simulation model that can capture the nonlinear relationship between system efficiency and lift, while not requiring detailed descriptions of the inner workings of the heat pumps themselves.

5.3.1.3 Renewable generation

Photovoltaic (PV) generation can be included to the simulation model in the following manner. The power generated can be calculated as a nonlinear function of the solar irradiance hitting the panel and the ambient temperature, following the approach described by Pepe et al. [28]. Where the power generated is given as P_{pv}, the ambient temperature is T_a and θ_1, θ_2 and θ_3 are parameters, the model can be represented as:

$$P_{pv} = \theta_1 \left(1 + \theta_2 I_{panel} + \theta_3 T_a\right) I_{panel} \tag{8}$$

These parameters can be derived from data using simple regression techniques. In the absence of data, nominal values are provided in [28], with θ_2 and θ_3 falling in the following ranges:

$$\theta_2 \in [-2.5 \times 10^{-4}, -1.9 \times 10^{-5}] \tag{9}$$

$$\theta_3 \in [-4.8 \times 10^{-3}, -1.7 \times 10^{-3}] \tag{10}$$

Similarly, solar thermal generation can be readily included. Solar thermal panels transfer heat to a fluid (water for the purposes of this framework) which can be used as a heat source for a heat pump or a thermal store. The model suggested here is the widely adopted Hottel-Whillier-Bliss model [29]. The useful energy removed by the panel (Q_u) can be calculated as a function of the collector area (A_{Sol}), the overall loss coefficient (U_L), the heat removal factor (F_R), the inlet flow temperature (T_i) and the ambient temperature (T_a) as follows:

$$Q_u = A_{Sol} F_R \left(I_{panel} - U_L (T_i - T_a) \right) \tag{11}$$

Parameter values (F_R and U_L) can be found from literature, for example, using values from the experimental set up of Anderson et al. [30]. Once again, these values could be calibrated to specific system with measured data.

Hybrid technologies that combine PV and solar thermal could also be considered. The efficiency of a PV panel tends to decrease with increasing temperature (note that the parameter θ_3 in Eq.[8] is negative). As such, PV efficiency improvements can be achieved by removing heat from the panel. This heat can be used in the manner of a solar thermal collector, resulting in a PV-thermal (PVT) hybrid.

To represent this, the Hottel-Whillier-Bliss model of Eq.[11] as modified by Florschuetz [31] to account for the electrical efficiency drop due to temperature increase can be used. The modified thermal efficiency calculation can be represented as:

$$\eta_t = F_{R_{PVT}} \left[(\tau\alpha)_e (1 - \eta_e) - U_L \left(\frac{T_{i_{PVT}} - Ta}{I_{panel}} \right) \right] \tag{12}$$

where η_e is the electrical efficiency at ambient and $(\tau\alpha)_e$ is the effective transmittance. The full area of a PVT collector will not be covered in PV cells, and as such, a term representing packing factor is included to relate the electrical cell efficiency to an efficiency per panel area as follows [32]:

$$\eta_e = \frac{A_{cell}\eta_{cell}}{A_{Sol}} = \beta_{pack}\eta_{cell} \tag{13}$$

where β_{pack} is the packing factor, A_{cell} and A_{panel} are the cell and panel areas and ηe and η_{cell} are the cell and panel efficiencies. Parameters can once again be obtained from Anderson et al. [30] in the absence of data.

5.3.1.4 Demand-side: occupancy modelling

While the previous sections illustrated techniques for capturing the heat flows, the demand for heat is dependent on the occupants of the building. Appropriate methods for modelling different user types are needed here and are of particular relevance to a domestic setting. The building control literature is dominated by non-domestic buildings (such as office buildings), in which usage schedules are regular and predictable

(e.g. office hours). This is not the case in a domestic setting whereby hours of absence and activity can vary depending on the types of occupant and their lifestyles. Despite this, uniform occupancy schedules (e.g. based on national averages) are prevalent in the literature. From the perspective of a useful building modelling framework, the use of different occupancy profiles should then be considered.

Given the wide range of potential occupant types and activities, defining realistic occupancy profiles can be a challenge. The occupancy-integrated archetype approach of [33] seeks to overcome the complexity present by using a data-driven approach to categorise the user behaviour recorded in an extensive national Time Use Survey (TUS). With this method, activities were mapped to one of three states: absent from the dwelling, present but inactive or present and active. The authors then used a clustering approach to create five representative occupant activity types. In the modelling framework described here, a modelled dwelling is assigned a user-type corresponding to one of these archetypes. From these, the heating schedule of the building is determined.

5.3.1.5 Demand-side: domestic hot water modelling

Aside from space-heating, the domestic hot water (DHW) demand can make up a significant proportion of the overall heat requirement of any building. In a well-insulated building, this hot water demand can exceed the space-heating demand. When considering predictive strategies that rely on forecasts of the heating requirement, this requires consideration (though in many cases, a separate heat source may be used to satisfy this demand). From a modelling perspective, DHW demand is usually characterised by demand spikes representing usage of taps, showers and baths. Such events could be modelled as stochastic processes whereby the likelihood of a DHW event varies through the day according to some average usage profile, such as that of Fig. 5.3, which is taken from a report by the Energy Savings Trust [34] for average usage in the UK. The profile can be modified to ensure that events only occur during occupied periods and scaled depending on the number of dwelling occupants following the guidelines indicated in [35]. This relationship between occupant numbers (denoted N_{occ}) and water usage can be summarised as

$$DHW_{litres} = 38 + 25N_{occ}. \tag{14}$$

The energy consumption associated with each DHW event is related to the volume of water consumed and the temperature rise required to increase the cold inlet water temperature (given here as T_c) to a suitable outlet temperature (given here as T_h). Poisson distributions that describe typical volume requirements for different event types can be found as part of the Centre for Renewable Energy Systems Technology (CREST) model [36]. Each DHW event can be assigned a required volume of water and the energy needed can then be calculated as:

$$DHW_{kWh} = 1000 DHW_{litres} \rho_w C_p \left(T_h - T_c\right), \tag{15}$$

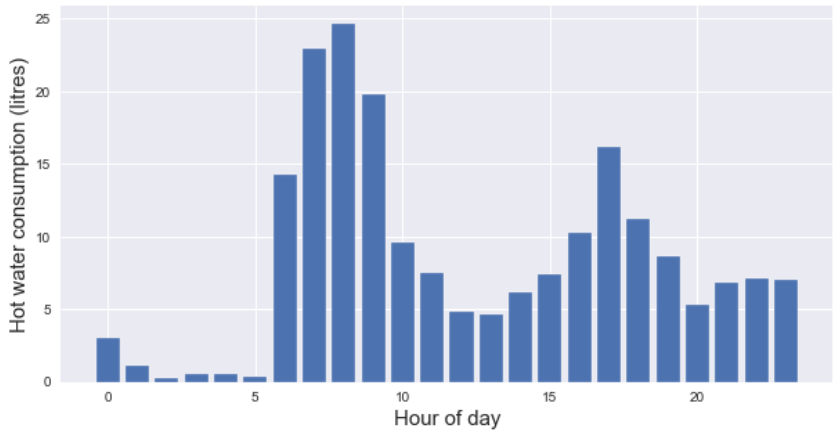

Fig. 5.3: Average daily DHW consumption

where ρ_w and C_p are the density and specific heat capacity of water, respectively. Suitable values for T_c and T_h can be found in [34].

5.3.1.6 Demand-side: Auxiliary electrical loads

Auxiliary electrical demands relating to lighting and other appliances can also be modelled. One simple method approach is to use two uniform daily profiles representing workdays and holidays/weekends. These can be taken from a report [37] in which measurements from 251 households were obtained and characterised in terms of dwelling and occupant type, then scaled based on the floor area.

Using an averaged approach such as this neglects the peaks and troughs that naturally occur during normal daily operation. Other approaches consider various appliance models and usage patterns to generate more realistic demand profiles (an excellent example can be found in the CREST model [36]). Here however, these auxiliary loads are considered non-controllable from the perspective of the thermally-focussed controllers discussed here and as such, a more detailed approach is not elaborated further.

5.3.2 Implementation of methods

Using the methodology developed in the previous sections, this section describes the creation and simulation of a model of a three-storey dwelling supplied by an air-source heat pump. The tools used are first described, followed by the outputs of the model.

5.3.2.1 Software tools

To put the above methods into practice, different software tools are available with different purposes. A state-of-the-art review of commonly used software and tools for simulation/controller design in energy-efficient buildings and comparing their capabilities is given in [38]. One such tool, particularly useful in control design, is the BRCM toolbox for Matlab developed at ETH Zurich [12]. The toolbox allows for the generation of RC type thermal building models. Since it is Matlab-based, additional system models and statistical demand models, such as those described in Sections 5.3.1.2-5.3.1.5, can be incorporated into a full system model. Furthermore, the functionality is provided in the BRCM models to import models created using Energyplus, a widely used tool for building energy performance evaluation [39]. Further tools have been developed to assist a user in the definition of an Energyplus model, notably Openstudio, which is used here. The full toolchain is then as follows:

1. The building shape is drawn and building materials are defined using Openstudio.
2. The Openstudio model is then exported as an Energyplus model.
3. The BRCM toolbox is used to extract the model dynamics from the Energyplus model.
4. The model matrices are used in Matlab along with a statistical DHW model and a heat pump model to simulate the heat demand of the building for a given set of occupants.

5.3.2.2 Case study: three-story dwelling

To illustrate the implementation of the modelling methods described in the previous section, a simulation model of a three-storey residential terraced house with 14 rooms was created and used to simulate thermal behaviour. The building model is based on a modern dwelling in the Trent Basin development of Nottingham UK with building fabric materials chosen to ensure the structure corresponds to the Energy Performance Certificate (EPC) of the building. It should be noted that the choice of building and location used here is arbitrary. The workflow is not restricted to a specific set of building properties or location characteristics. The resulting thermal characteristics of the building are summarised in Table 5.1.

The linear state-space model representing the building consists of 223 states, 14 controlled inputs (heat supply to each zone), 14 outputs (zone air temperatures) and 20 disturbance channels. The disturbances include internal gains for each zone, the ground temperature, the air temperature and solar gains corresponding to each orientation. An air-source heat pump is used to supply heat via a buffer tank which controlled to maintain a temperature of 60°C at all times. A diagram of the building structure is shown in Fig. 5.4.

Prior to development of the MPC strategy (which is detailed in Section 5.5), an on/off control strategy is applied whereby the controller of each radiator follows a hysteresis loop. The maximum heat supply is applied when the temperature falls below a specified low threshold and remains in this state until exceeding a specified

Table 5.1: Three-storey dwelling construction details

	Thermal transmittance W/m^2K	Heat capacitance (per m^2 surface area) kJ/m^2K	Solar Factor (G-value)	Area m^2
Roof	0.120	70.69	-	-
Floor	0.155	305.1	-	-
Ext. walls	0.128	187.5	-	-
Internal walls	2.581	33.1	-	-
Glazing	1	-	0.5	20.1

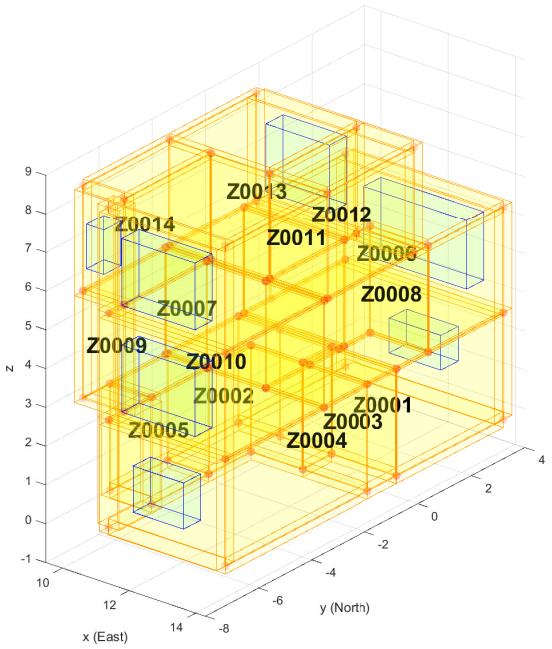

Fig. 5.4: Trent Basin terraced dwelling (floor area = $129.1m^2$)

high threshold at which point the heat supply is switched off. An occupancy schedule is chosen such that heating is required for a short period in the morning followed by a longer period in the evening each weekday. At weekends, a single longer occupied period is assumed. At night and during unoccupied periods, the desired temperature set-point is 16°C while during occupied periods it is raised to 21°C with a 1°C comfort band allowed for the on/off hysteresis loop.

This set-up is simulated for a period of seven days with weather data taken from January 2019 in London. The 14 zone temperatures are plotted in Fig. 5.5 with the room set-point temperature plotted as the thick black line. Since the strategy is purely reactive, it can be seen that when the lower set-point increases, a certain amount of time is needed before the zone temperatures begin to approach the desired comfort level. The influence of solar gains can be seen in the temperature spikes that occur in the third and sixth day.

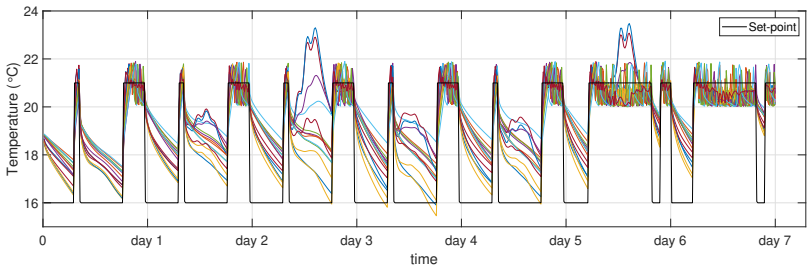

Fig. 5.5: Simulated zone temperatures for seven-day simulation with on/off radiator control strategy

The heat demand from the building for this 7-day scenario is shown in Fig. 5.6. This includes both the space heating demand and the DHW heat demand.

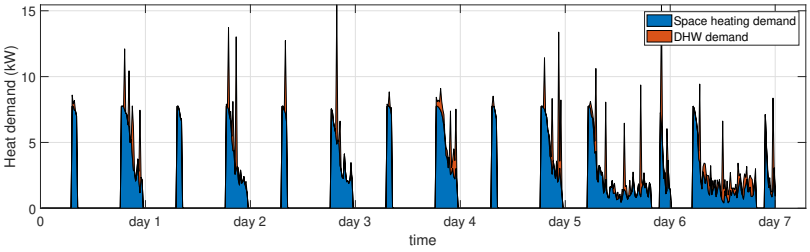

Fig. 5.6: Simulated heat demand (both space heating and hot water) for one week

5.4 Low-order modelling for MPC applications

With a simulation environment developed as outlined in the previous sections, the development of an MPC formulation is made more straightforward. In this section, modelling methods that can be used for MPC development are outlined. An opti-

misation problem lies at the heart of MPC-based strategies, whereby the behaviour of the building must be predicted with control choices made to optimise some predefined criteria. A model of the system behaviour is needed to carry out such an optimisation. Suitable models for MPC implementation differ from those used for simulation, however, in that the high degrees of complexity may lead to optimisation formulations that are not tractable in real-time. With this in mind, strategies for developing low-complexity models in the context of building energy systems are discussed here. The methods are implemented using the simulation environment developed in Section 5.3. By analysing the performance of the models in a simulated setting, the methods can be assessed prior to implementation in a real building.

5.4.1 A model-based pathway

Using the methods outlined in Section 5.3, the production of a detailed simulation model can be achieved without significant modelling effort. The resulting model can contain a large number of equations and parameters, due to the need to characterise each individual structural component of the building fabric. Such a model may be unmanageable from an optimisation perspective. Nonetheless, model-reduction techniques can be applied to the detailed model to derive lower-complexity representations. Different model reduction approaches exist (see for example [40, 41]), with the approach chosen here being balanced truncation.

5.4.1.1 Balanced truncation

Using a balanced truncation approach, a linear transformation is carried out on the full state-space model and the system's Hankel singular values (HSVs) are identified. The HSVs of a system can be defined as the square-root of the eigenvalues of the product of the system's observability and controllability gramians. A transformation chosen to lead to equivalent, diagonal gramians results in what is referred to as a *balanced* model and the HSVs can be found as the main diagonal of either gramian. Relatively smaller HSVs imply lower energy states, the removal of which will have less effect on the more dominant system dynamics. As such, an appropriate number of states can be retained to ensure sufficient predictive performance with greatly reduced model-order [42].

The general theory is briefly described here. A transformation $\tilde{x} = Tx$ can be applied to the model of (3)-(4):

$$\dot{\tilde{\boldsymbol{x}}}(t) = \tilde{A}\tilde{\boldsymbol{x}}(t) + \tilde{B}\boldsymbol{u}(t) \tag{16}$$

$$\tilde{\boldsymbol{y}}(t) = \tilde{C}\tilde{\boldsymbol{x}}(t), \tag{17}$$

whereby

$$\left(\tilde{A}, \tilde{B}, \tilde{C}\right) = \left(TAT^{-1}, TB, CT^{-1}\right). \tag{18}$$

The controllability gramian W_c and observability gramian W_o of this new system can be defined as:

$$\tilde{W}_c = TW_cT^T \tag{19}$$

$$\tilde{W}_o = \left(T^{-1}\right)^T W_0 T^{-1}. \tag{20}$$

This system can be easily transformed to a balanced system with equal, diagonal gramians given by:

$$\tilde{W}_c = \tilde{W}_c = diag(\sigma_{hsv_1}, \sigma_{hsv_2} ... \sigma_{hsv_{172}}), \tag{21}$$

arranged such that $\sigma_{hsv_1} \geq \sigma_{hsv_2} \geq \cdots \geq \sigma_{hsv_{172}} \geq 0$. σ_{hsv_p} denotes the p^{th} HSV. In this transformed system, the states that correspond to smaller HSVs have less impact on system behaviour and can be removed.

5.4.1.2 Implementation in the simulation environment

Applying balanced truncation to the simulation model developed in Section 5.3 leads to a model with reduced order. If each zone is viewed as a separate entity, the states that best describe the interaction between the input and output of that zone can be found by looking at the HSVs. For the three-storey building model, the HSVs (found using Matlab's *hsvd* function) for zone 1 are shown in Fig. 5.7 in descending order of state energy. States that don't contribute greatly to the system behaviour have a low energy and can be discarded without significantly impacting the performance of the model.

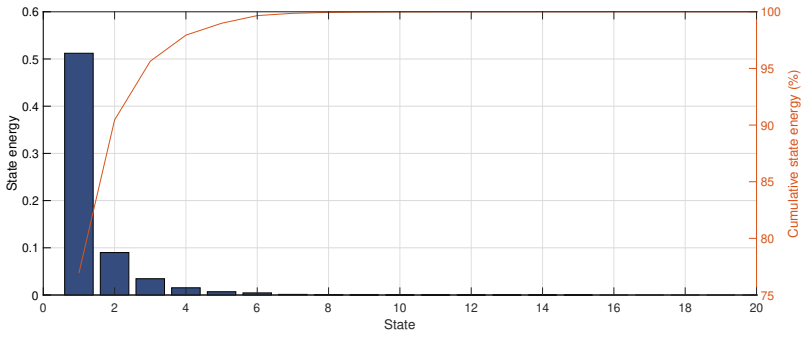

Fig. 5.7: Hankel singular values associated with the first zone of the simulation model

It can be seen from this figure that the system behaviour can be mostly captured by a low number of states. As such, the full modelled system of Section 5.3, which

was previously made up of 233 state equations, can now be replaced by zone models with as few as 2-5 states. A comparison of different model orders is shown in Fig. 5.8 where the same inputs and disturbances are applied to each. The difference between the full-order representation (the black dashed line) and the reduced-order representation increases with lower model-orders, though the 3^{rd} and 5^{th} order models appear to capture the behaviour quite well.

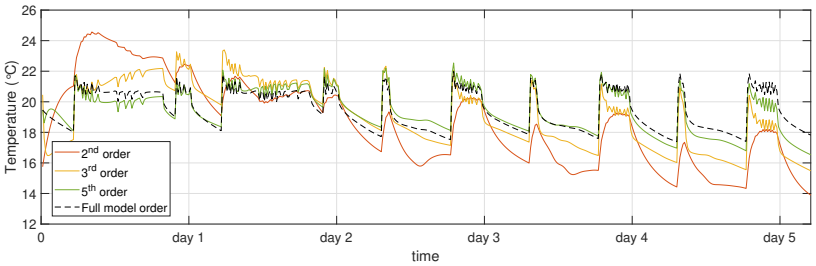

Fig. 5.8: Comparison of outputs for different model orders derived using truncation

While methods exist for quantification of the difference between models, for example the gap metric approach [46], a true assessment of the usefulness of the model can be achieved by observing the performance of the model within an MPC formulation which will be carried out in Section 5.5.

5.4.2 A data-driven pathway

Two disadvantages of the model-based pathway are that a detailed model must be available, and this detailed model must be a good representation of the real building. While having a detailed model is useful from an analysis perspective, from a scale-up perspective it may be unrealistic to assume that such a modelling effort would be applied to every building to which MPC is deployed. By directly deriving model parameters using data measurements taken from the building however, the need for a detailed simulation model is removed. The main prerequisite in this case is for sensors to be deployed appropriately in the building along with data handling and storage capabilities, which would most likely be required for MPC in any case. Typically measurements of the air temperature in the temperature-controlled rooms of the building would be used along with the heat flux into the rooms, or a proxy thereof. Examples would be the flow rate and temperature of the water flowing through a radiator-based heating system, or the temperature and airflow through an air-handling unit (AHU).

5.4.2.1 Parameter identification

The general concept is to derive the relationships between the controlled inputs, measured or estimated disturbances and measured outputs of the building, but specific requirements depend on the control approach used. As to be expected, there are a large number of methods investigated in the literature for data-driven building model identification, covering a variety of machine-learning and system identification techniques [21, 47]. For the purposes of MPC, the model structure is paramount as it must be suited to a receding horizon optimisation-based framework. While this narrows the range of possible choices, there remains no single accepted technique and the field is likely to evolve in the future. In the approach focussed on here, the model structure is chosen to be the same as for the model-based approach (i.e. the thermal dynamics of the building fabric are represented in a linear state-space form). As such, we refer to it as a grey-box strategy insofar as the model structure is assumed (unlike a black-box approach) but parameters are empirically derived (unlike a white-box approach).

5.4.2.2 Implementation in the simulation environment

Applying this technique using data obtained from the detailed simulation model, models for each zone can again be derived for different model orders. Inputs are the heat flows into each room/zone, outputs are the room/zone temperatures and external temperature, while solar radiation and internal gains can be included as disturbances. Inclusion of disturbance channels in a real strategy would depend on the availability of certain measurements. The model-order is user-specified, and the model parameters are identified using the N4SID function of Matlab's System Identification toolbox [48].

Appropriate training data is crucial in the modelling process and the collection of such data is a challenging feature of the building energy application. Building heating schedules tend to be quite repetitive, leading to a low degree of excitation in the training data. Carrying out functional testing (e.g. pseudo-random-binary-sequence (PRBS) events) can be difficult if the building is occupied without violating comfort criteria. Furthermore, external uncontrollable disturbances (particularly the weather) play a large role in the thermal behaviour of the building. Long timescales may be needed to capture a wide range of weather events. In the case presented here, it is assumed that no functional testing is possible and training data is generated using the on/off strategy of Fig. 5.5.

In Fig. 5.9 the outputs generated by different model orders are plotted against the full-order model output where the same inputs have been applied to all. The data-set used for training the models was different to the data-set shown here. It can be seen that in all cases the behaviour has been captured quite well. Note however, no noise is present in the full-order system or in the measurements taken. In reality, noise and other data issues can degrade the performance. The development of data-driven

methods that can provide performance guarantees in the face of uncertainty, without requiring intrusive training-data collection periods, remains an open challenge.

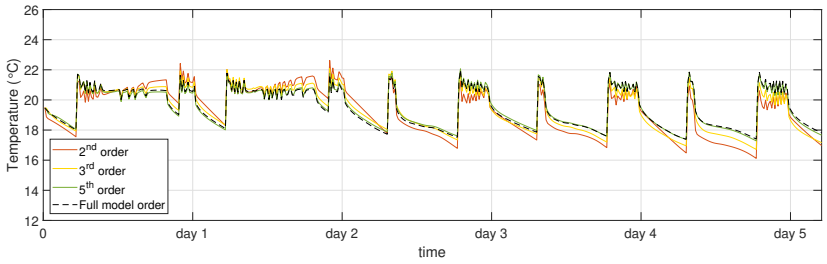

Fig. 5.9: Comparison of outputs for different model orders derived using system identification

To illustrate the use of the modelling methods outlined, their implementation in an MPC strategy is required. Since methods are independent of the control architecture, two different control architectures, centralised and decentralised, are developed in the next section. The strategies will incorporate the low-complexity models of Section 5.4 with simulation and evaluation carried out in the simulation environment of Section 5.3.

5.5 Building models within an MPC framework

With a suitable modelling framework established along with suitable low-order models, the MPC strategy can now be developed. In this section, a suitable formulation is outlined and demonstrated using the simulation model of Section 5.3. The aim of the MPC formulation is to maintain the temperatures of air in each zone within a satisfactory comfort range while minimising energy consumption, subject to constraints in the system. At each sample, an optimisation problem is solved to determine such a set of inputs that satisfy the specified criteria over a given prediction horizon. The first elements of the optimal input trajectory are applied, and the optimisation is repeated at the next step in a receding horizon manner. The models developed in Section 5.4 form the basis of the optimisation problem.

Two architectures are considered here to highlight different possible modelling structures: centralised and decentralised. In both cases the MPC seeks to find optimal heat inputs to each zone, while the heating system itself (e.g. the heat pumps and/or thermal stores) is not considered to keep the focus solely on the building models. Centralised (CMPC) and decentralised (DMPC) approaches have different advantages and disadvantages with suitability depending on context. CMPC can require a larger computation time, potentially exceeding the sample-time, whereas DMPC

formulations, though more computationally lightweight, may neglect interaction between zones, potentially degrading performance [42]. Selection of a strategy requires careful consideration. This is a pressing challenge in the residential sector, where formulations that span multiple zones and multiple buildings are required to allow for aggregation of flexibility provision. A detailed comparison between centralised and decentralised approaches can be found in [43], whereby the influence of the partition insulation level is considered. Examples of distributed strategies that seek to find a compromise between the two through the use of minimal information flows between otherwise decentralised controllers can be found in [44] and [45].

5.5.1 Centralised MPC: zone level control

In the centralised formulation, all zones are considered within the same optimisation problem. The optimisation formulation takes measurements of the zone temperatures at each time step along with weather forecasts to determine optimal radiator heat supply settings for all zones using the low complexity models for prediction. The schematic of this is shown in Fig. 5.10. A key point to note is that in the centralised strategy, all zones are considered in the same model as well as the interactions between them. This necessitates greater complexity but may result in improved performance.

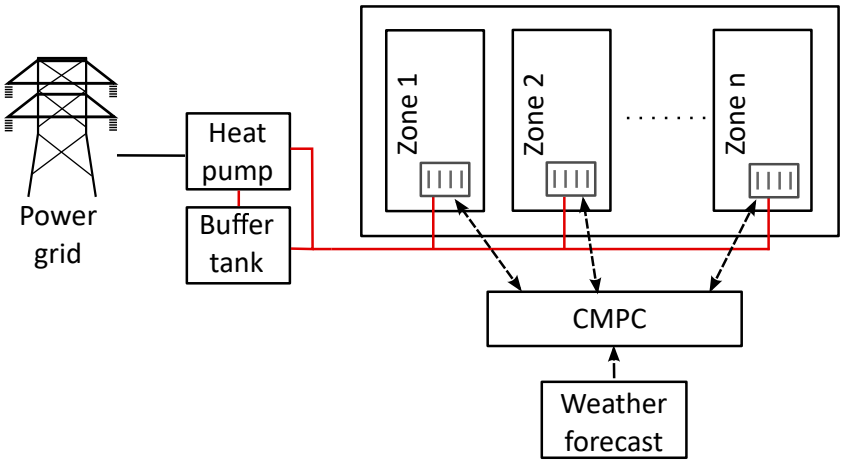

Fig. 5.10: Centralised building control strategy

The low-complexity model matrices are given for this centralised strategy as A, B, C and E, with $\boldsymbol{y}(k)$ representing the vector of zone temperatures at time-step k, $\boldsymbol{x}(k)$ representing the vector of internal model states, $\boldsymbol{u}(k)$ representing the inputs to each zone and $\boldsymbol{v}(k)$ representing the disturbances. Note that this is a discrete time

model, as opposed to the continuous time representation of (3) with k representing the discrete time-steps. The deviation outside the comfort set-point band in each zone is captured by the vector $\boldsymbol{\varepsilon}(k)$. Instead of applying a hard constraint on comfort (by using a constraint that states that the zone temperatures must stay within the specified comfort bounds), a soft constraint is used whereby comfort deviations are allowed, but are heavily penalised with a quadratic cost. This ensures that a feasible solution can be achieved, and larger deviations are penalised more heavily than smaller deviations. The comfort band is bounded by an upper set-point given as $T_{sp_{hi}}$ and a lower set-point $T_{sp_{lo}}$. The vector of maximum radiator heat supplies (one for each zone) is given as $\boldsymbol{q}^{max}$. Where R and S are weighting matrices balancing the energy minimisation and set-point tracking objectives, the objective and constraints of the centralised MPC problem are then given as:

$$\boldsymbol{u}^* = \arg\min \sum_{k=0}^{N-1} \left(\boldsymbol{\varepsilon}(k)^T R \boldsymbol{\varepsilon}(k) + S \left[\boldsymbol{u}(k)^T, \boldsymbol{\varepsilon}(k)^T \right]^T \right) \tag{22}$$

$$s.t. \quad \boldsymbol{x}(k+1) = A\boldsymbol{x}(k) + B\boldsymbol{u}(k) + E\boldsymbol{v}(k) \tag{23}$$

$$\boldsymbol{y}(k) = C\boldsymbol{x}(k) \tag{24}$$

$$0 \leq \boldsymbol{u}(k) \leq \boldsymbol{q}^{max} \tag{25}$$

$$T_{sp_{hi}}(k) - \boldsymbol{\varepsilon}(k) \leq \boldsymbol{y}(k) \leq T_{sp_{hi}}(k) + \boldsymbol{\varepsilon}(k) \tag{26}$$

$$0 \leq \boldsymbol{\varepsilon}(k) \tag{27}$$

The model-based or the data-driven pathway can be used to derive the A, B, C and E matrices. Using such a formulation, it is possible to influence the behaviour of the heating system by updating the objective and constraints accordingly. Perhaps the most common way in which this is achieved is by incorporating a time-varying energy cost in the weighting matrix S. Energy use at times of high cost is penalised more heavily than at times of low cost, resulting in solutions that shift the timing of the energy use while maintaining constraint satisfaction. Alternatively, input constraints can be tightened as required, to force curtailment of the energy supply. The quadratic form of the objective and the linear constraints ensure solutions can be obtained with standard quadratic programming solvers.

In either centralised or decentralised form, an estimation of the model state vector $\boldsymbol{x}$ at time k is needed to initialise the optimisation model at time k. These states are not representative of physical quantities (unlike in the full-order physics-based representation of the building, in which all states represent wall, air, floor or ceiling temperatures) unless a first-order model is used, in which case, the state is equivalent to the measured air temperature times a constant C. In the modelling framework, state estimation is carried whenever an MPC strategy is deployed by use of Kalman filtering.

5.5.2 Decentralised MPC: system level control

In the decentralised formulation, a separate MPC problem is solved for each individual zone. Each zone has its own model that relates the zone's radiator heat input and any relevant disturbances to the zone air temperature. Thermal interactions between different zones are not modelled. This is shown in the schematic of Fig. 5.11.

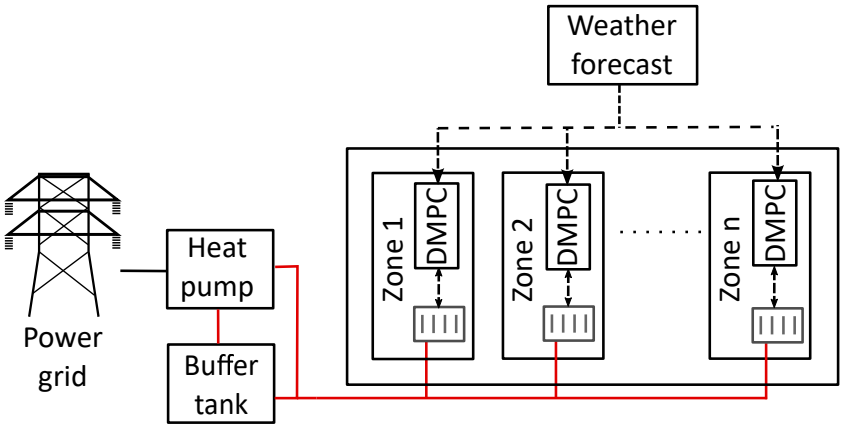

Fig. 5.11: Decentralised building control strategy

For a building with J zones, model matrices are given for zone $j \in J$ as A_j^z, B_j^z, C_j^z and E_j^z, with $y_j(k)$ representing the zone temperature at time k, $\boldsymbol{x}_j(k)$ representing the vector of internal model states, $u_j(k)$ representing the heat input to the zone and $\boldsymbol{v}_j(k)$ representing the zone-relevant disturbances. Once again, the deviation outside the comfort set-point band (bounded by $T_{sp_{lo}}$ and $T_{sp_{hi}}$) is penalised, with deviations denoted by $\varepsilon_j(k)$. The maximum radiator heat supply is given as q_j^{max}. The objective weighting matrices are given as R_j and S_j for the zone, leading to the following formulation:

$$\boldsymbol{u_j^*} = \arg\min \sum_{k=0}^{N-1} \left(\varepsilon_j^2(k) R + S \left[u_j(k), \varepsilon_j(k) \right]^T \right) \tag{28}$$

$$s.t. \quad \boldsymbol{x}_j(k+1) = A_j^z \boldsymbol{x}_j(k) + B_j^z u_j(k) + E_j^z v_j(k) \tag{29}$$

$$y_j(k) = C_j^z \boldsymbol{x}_j(k) \tag{30}$$

$$0 \leq u_j(k) \leq q_j^{max} \tag{31}$$

$$T_{sp_{lo},j}(k) - \varepsilon_j(k) \leq y_j(k) \leq T_{sp_{hi},j}(k) + \varepsilon_j(k) \tag{32}$$

$$0 \leq \varepsilon_j(k) \tag{33}$$

The purpose of including different control structures is to show the adaptability of the proposed strategies to fit in multi-input-multi-output (MIMO) and single-input-single-output (SISO) contexts, rather than presupposing the existence of a one-size-fits-all solution.

5.5.3 Implementation in the simulation model

The behaviour of the control approach is demonstrated here by implementation of the decentralised formulation with 5^{th}-order zone models derived using the model truncation method. A 15-minute sample-time and a 20-step prediction horizon (5 hours) were applied. The weighting matrices R_j and S_j were chosen to ensure that comfort violation is heavily penalised relative to financial cost. In this way the controller attempts to keep the zone temperatures within or as close as possible to the comfort bounds. Of course, by increasing the relative weight on the financial penalty, larger deviations from the comfort bounds would be allowed for the sake of financial savings. Finding an appropriate balance for a specific user is part of the design process. Once again, by having a detailed simulation model available, this potentially abstract balance can be analysed in terms of specific financial savings and tangible comfort deviations.

The resulting simulated zone temperatures for the building over a 2-weekday simulation are shown in Fig. 5.12. The dark black lines indicate the designated comfortable temperature range for the building. It can be seen that the strategy successfully maintains the temperatures within these bounds without unnecessary overheating. The ability of the decentralised approach to maintain internal comfort indicates that the cross-zone interactions do not significantly impact the performance, which is unsurprising, since the temperature difference between zones is low so heat flow between zones can be assumed to be small.

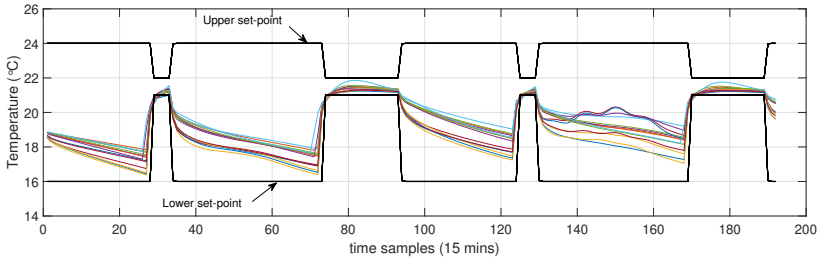

Fig. 5.12: Zone temperatures simulated in the building for a two-weekday simulation using a decentralised MPC control strategy

A comparison of the MPC performance with that of the On/Off strategy from Section 5.3 is shown in Table 5.2. To quantify the level of comfort satisfaction,

Table 5.2: Space heating performance of predictive and non-predictive strategies

	Energy demand kWh	Comfort deviation K.hr
On/off strategy	70.1	3.8
MPC strategy	72.4	0.7

discomfort is measured here in K.hr, whereby a deviation outside the comfort bounds of 1K for one hour corresponds to 1K.hr. The reactive nature of the On/Off strategy leads to 3.8 K.hr deviation from the comfort bounds, compared to 0.7 K.hr for the MPC strategy. Interestingly, the energy requirement is greater for the MPC strategy, which reflects the conflicting nature of the comfort and energy goals. More energy was required to ensure the comfort goals were achieved. This can be seen clearly in Fig. 5.13 in which the average zone temperature of the buildings is plotted for the MPC and On/Off scenarios. To achieve comfort, the zones are maintained at a higher temperature when MPC is used.

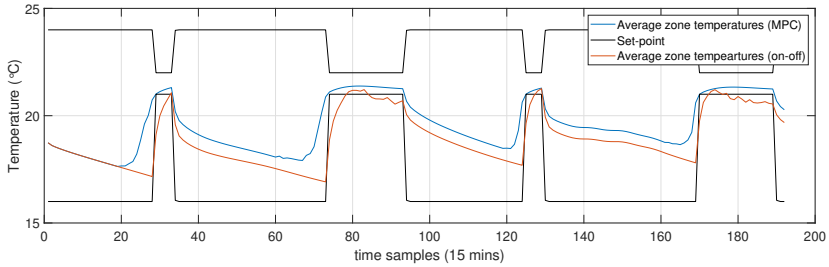

Fig. 5.13: Average zone temperatures for MPC (decentralised) and On/Off strategies for a two-weekday simulation

The handling of multiple objectives (in this case, energy and comfort), is not necessarily trivial. The need to handle the objectives of different users across financial, environmental and resilience vectors, without the need for excessive design input, remains an open challenge, particularly given the diversity of stakeholder requirements in the residential sector.

5.5.3.1 Including price-awareness

Aside from a better thermal comfort performance, one of the key advantages of MPC when compared to traditional strategies is the ability to incorporate external factors in the control decisions, thus enabling the shifting of demands in a manner that benefits the power grid. A typical mechanism for achieving this is to introduce

a variable electricity price, such as the *Agile* tariff used by Octopus Energy [49]. At times of high demand, the price can be increased to disincentivize energy use, encouraging a shift towards times of lower price. To show this using the modelling framework, the objective of the MPC strategy was adjusted to minimise electricity consumption cost instead of the energy consumption using a variable electricity tariff. The behaviour of this *price-aware* strategy is shown in Fig. 5.14 in which the average zone temperatures are plotted, as well as the electricity price. Two large price spikes can be seen in the two-weekday simulation. Prior to these spikes, an increase in the zone temperatures is visible. By shifting the energy use, less energy is consumed during the high-price period.

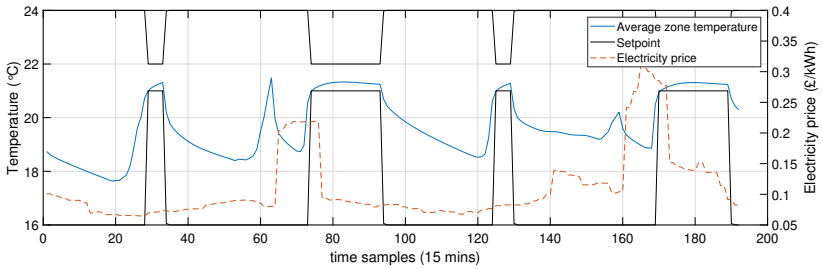

Fig. 5.14: Average zone temperature simulated in the building with price-aware MPC strategy for two-weekday simulation with decentralised MPC control strategy. The red, dashed line indicates the electricity price

This behaviour can be seen more clearly by observing the heat supplied to the building, as shown in Fig. 5.15 in which the electricity price is also plotted. During the two price spikes, the heat supply drops to zero, with a certain amount of pre-heating carried out before the price-spikes occur.

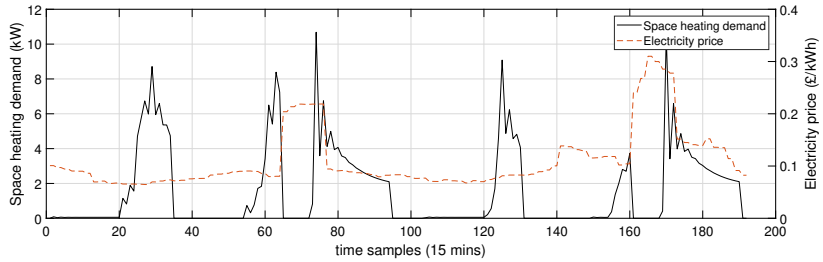

Fig. 5.15: Space heating demand and electricity price for 2-day simulation period

This behaviour results in a lower electricity cost (the total amount paid for electricity for the simulation length) when compared to a non-price aware strategy with

Table 5.3: Space heating performance of MPC with a time-varying price tariff strategy

	Electricity cost £	Energy demand kWh	Comfort deviation K.hr
Price-aware MPC	8.77	73.9	0.7
Non-price-aware MPC	9.82	72.4	0.8

little change in the comfort satisfaction, achieved by shifting energy use away from high-tariff times. By supplying heat earlier the total heat requirement increases slightly. The amount of additional heat lost due to this shift in demand will depend on the level of insulation of the building. While it's generally understood that well-insulated buildings are more suited to this type of flexibility provision, the modelling framework can explicitly capture the trade-off. Simple methods for quantifying this trade-off without significant analysis remain a challenge.

5.5.3.2 Parameter selection considerations

The control design process can be streamlined using the methods and the workflow presented here, however the design process is not completely eliminated. Any given choice of parameters and structures may not lead to a controller that behaves in a similar manner across all buildings and contexts. To show this, the impact of some of these parameters on the computation time is examined here.

Choices such as model-order and control structure will affect the resulting performance. For example, the impact of model order, prediction horizon and structure selection on the time taken to carry out a single optimisation using a 3.2GHz 4-core processor is shown in the box-plots of Fig. 5.16-5.17. The use of box plots here is to capture the variation found in the optimisation times. In Fig. 5.16 it can be seen that although there is an improvement in using a 30^{th}-order model in comparison to the full-order RC model, the significant time improvement is achieved when the low-order decentralised models are used.

In Fig. 5.16, a prediction horizon of 20 steps is used. The increased relevance of model order when longer horizons are used is shown in Fig. 5.17, with box plots once again used to compare the time taken to complete an optimisation, this time comparing a 20-step (5 hours) and a 96-step (24 hours) horizon. It can be seen that the time taken to carry out an optimisation using the full-order model becomes prohibitively large when the 1-day horizon is used (in the context of a 15-minute sample time). For the low-order models, even though the shorter time-horizon again produces a much faster solution, the average time taken for the longer horizon is approximately 0.1s (depending on the computational power available).

The purpose of this is to illustrate that the choices of prediction horizon, model order and control structure are all connected from the perspective of problem complexity and the correct choice of parameters may not be obvious but will depend

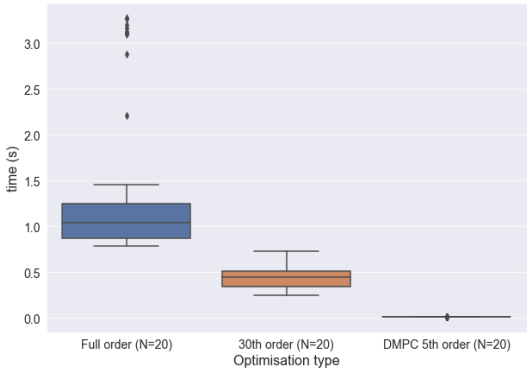

Fig. 5.16: Time taken to solve MPC optimisation formulations at each time-step for different model orders

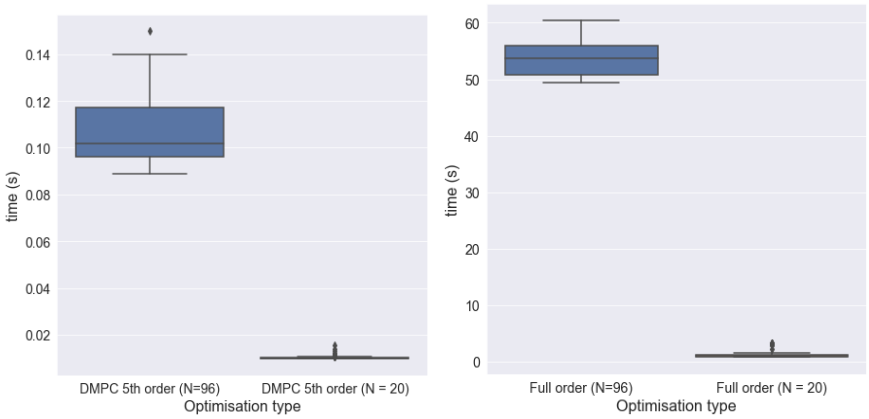

(a) (b)

Fig. 5.17: Time taken to solve MPC optimisation formulations once for different model types and prediction horizons

on the context. High-order models may provide additional accuracy that is needed is some contexts, but not in others. Similarly, longer prediction horizons may or may not be required. The trade-off that stems from these decisions is all part of the design challenge as it stands. While the availability of a simulation environment can aid selection of such parameters, methods for automating the design process further could offer significant benefits.

5.6 Conclusions

This chapter outlines a workflow for designing and evaluating MPC strategies in the residential building sector to act as a guide for practitioners in the area. With electrification of the heating and transport sectors well underway, the need for intelligent control approaches that can be applied to the domestic building sector to enable effective demand-side management is urgent. The development of appropriate models for this purpose acts as a significant barrier to the scale-up of such methods. The workflow introduced here incorporates the establishment of a detailed physics-based simulation environment, the development of low-complexity control-oriented models, the formulation of a suitable optimisation strategy and the evaluation of the strategy using the simulation environment. Suitable methods and tools for simulating the thermal behaviour of a building were first detailed, with the production of a detailed physics-based model of a three-storey building carried out. Model-based and data-driven methods that can be used to capture the building's thermal behaviour in a low-complexity form were then presented, followed by the formulation of an MPC strategy that seeks to ensure occupant comfort with minimal energy use. Using the developed simulation environment, it was shown that the MPC strategy can significantly reduce violations of the comfort criteria when compared to a more traditional reactive control approach. The possibility of the MPC strategy to incorporate price information was also demonstrated. Using a price-aware strategy, the energy cost for space-heating reduced by approximately 11% though the overall energy consumption slightly increased. With the workflow established here, such strategic trade-offs can be properly analysed prior to implementation. By illustrating the process in detail in the manner presented, in addition to highlighting the modelling principles required in the sector, some of the challenges that remain are emphasised, including parameter selection, the impact of which was shown in terms of computational performance.

Acknowledgements This work has received funding from the EPSRC (Engineering and Physical Sciences) under the Active Building Centre project (reference number: EP/V012053/1).

References

1. Kathirgamanathan A, De Rosa M, Mangina E, Finn DP (2021) Data-driven predictive control for unlocking building energy flexibility: A review. Renew Sustain Energy Rev 135:110120.
2. Greater London Authority (2018) Zero carbon London: A 1.5oC compatible plan.
3. Gonzato S, Chimento J, O'Dwyer E, et al (2019) Hierarchical price coordination of heat pumps in a building network controlled using model predictive control. Energy Build 202:.
4. Al Essa MJM (2019) Home energy management of thermostatically controlled loads and photovoltaic-battery systems. Energy 176:742–752.
5. O'Dwyer E, Pan I, Charlesworth R, et al (2020) Integration of an Energy Management Tool and Digital Twin for Coordination and Control of Multi-vector Smart Energy Systems. Sustain Cities Soc 62:102412.

6. Barone G, Buonomano A, Calise F, et al (2019) Building to vehicle to building concept toward a novel zero energy paradigm: Modelling and case studies. Renew Sustain Energy Rev 101:625–648.
7. Shaikh PH, Nor NBM, Nallagownden P, et al (2014) A review on optimized control systems for building energy and comfort management of smart sustainable buildings. Renew Sustain Energy Rev 34:409–429.
8. Drgoňa J, Arroyo J, Cupeiro Figueroa I, et al (2020) All you need to know about model predictive control for buildings. Annu Rev Control.
9. Killian M, Kozek M (2016) Ten questions concerning model predictive control for energy efficient buildings. Build Environ 105:403–412.
10. Atam E, Helsen L(2016) Control-oriented thermal modeling of multizone buildings: methods and issues: intelligent control of a building system. IEEE Control Systems Magazine 36(3):86–111.
11. Simulationresearch (2016) FrontPage - bcvtb. https://simulationresearch.lbl.gov/bcvtb. Accessed 13 May 2018.
12. Sturzenegger D, Gyalistras D, Semeraro V, et al (2014) BRCM Matlab Toolbox: Model generation for model predictive building control. Proc Am Control Conf 1063–1069.
13. Zwickel P, Engelmann A, Groll L, et al (2019) A Comparison of Economic MPC Formulations for Thermal Building Control. In: 2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe). IEEE, pp 1–5.
14. Rätz M, Javadi AP, Baranski M, et al (2019) Automated data-driven modeling of building energy systems via machine learning algorithms. Energy and Buildings 202.
15. Coakley D, Raftery P, Keane M (2014) A review of methods to match building energy simulation models to measured data. Renewable and Sustainable Energy Reviews 37:123–141.
16. Harish VSKV, Kumar A (2016) A review on modeling and simulation of building energy systems. Renewable and Sustainable Energy Reviews 56:1272–1292.
17. Lin Y, Middelkoop T, Barooah P (2012) Issues in identification of control-oriented thermal models of zones in multi-zone buildings. In: Proceedings of the IEEE Conference on Decision and Control. Ieee, pp 6932–6937.
18. Sturzenegger D, Gyalistras D, Morari M, Smith RS (2016) Model Predictive Climate Control of a Swiss Office Building: Implementation, Results, and Cost–Benefit Analysis. IEEE Trans Control Syst Technol 24:1–12. https://doi.org/10.1109/TCST.2015.2415411
19. Bünning F, Huber B, Heer P, et al (2020) Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings. Energy Build 211:.
20. Coulson J, Lygeros J, Dörfler F (2020) Distributionally Robust Chance Constrained Data-enabled Predictive Control. 1–14
21. Maddalena ET, Lian Y, Jones CN (2020) Data-driven methods for building control — A review and promising future directions. Control Eng Pract 95:104211.
22. CDBB (2018) The Gemini Priciples. Univ Cambridge, UK 2018 15.
23. O'Dwyer E, Pan I, Acha S, et al (2019) Modelling and Evaluation of Multi-Vector Energy Networks in Smart Cities. In: International Conference on Smart Infrastructure and Construction 2019 (ICSIC). ICE Publishing, pp 161–168.
24. EASME TABULA webtool. http://webtool.building-typology.eu/#bm
25. Ma Y, Kelman A, Daly A, Borelli F (2012) Predictive control for energy efficient buildings with thermal storage. IEEE Control Syst Mag 32:44–64.
26. Oliveira Panão MJN, Mateus NM, Carrilho da Graça G (2019) Measured and modeled performance of internal mass as a thermal energy battery for energy flexible residential buildings. Appl Energy 239:252–267.
27. Staffell I, Brett D, Brandon N, Hawkes A (2012) A review of domestic heat pumps. Energy Environ Sci 5:9291.
28. Pepe D, Bianchini G, Vicino A (2018) Estimating PV forecasting models from power data. In: 2018 IEEE International Energy Conference, ENERGYCON 2018. IEEE, pp 1–6.
29. Duffie JA, Beckman WA (2013) Solar Engineering of Thermal. John Wiley & Sons, Inc., Hoboken, NJ, USA.

30. Anderson TN, Duke M, Morrison GL, Carson JK (2009) Performance of a building integrated photovoltaic/thermal (BIPVT) solar collector. Solar Energy 83:445–455.
31. Florschuetz LW (1976) Extension of the Hottel-Whillier-Bliss Model To the Analysis of Combined Photovoltaic/Thermal Flat Plate Collectors. Solar Energy 22:361–366.
32. Chow TT (2010) A review on photovoltaic/thermal hybrid solar technology. Appl Energy 87:365–379.
33. Buttitta G, Turner WJN, Neu O, Finn DP (2019) Development of occupancy-integrated archetypes: Use of data mining clustering techniques to embed occupant behaviour profiles in archetypes. Energy Build 198:84–99.
34. Energy Savings Trust (2008) Measurement of domestic hot water consumption in dwellings. Energy Savings Trust 1–62.
35. Henderson J, Hart J (2015) BREDEM 2012: A technical description of the BRE Domestic Energy Model - Version 1.1.
36. McKenna E, Thomson M (2016) High-resolution stochastic integrated thermal-electrical domestic demand model. Applied Energy.
37. Zimmermann J-P, Evans M, Lineham T, et al (2012) Household Electricity Survey: A study of domestic electrical product usage. Intertek, 1-600.
38. Atam E (2017) Current software barriers to advanced model-based control design for energy-efficient buildings. Renew Sustain Energy Rev, 73:1031-1040.
39. DOE (2017) EnergyPlus | EnergyPlus. In: U.S. Dep. Energy's. https://www.energyplus.net/. Accessed 10 May 2018.
40. Astolfi A (2010) Model reduction by moment matching for linear and nonlinear systems. IEEE Transactions Automatic Control.
41. Deng K, Goyal S, Barooah P, Mehta PG (2014) Structure-preserving model reduction of nonlinear building thermal models. Automatica 50:1188–1195.
42. Lyons B, O'Dwyer E, Shah N (2020) Model reduction for Model Predictive Control of district and communal heating systems within cooperative energy systems. Energy 197:117178.
43. Pedersen TH, Hedegaard RE, Knudsen MD, Petersen S (2017) Comparison of centralized and decentralized model predictive control in a building retrofit scenario. In: Energy Procedia.
44. Morosan P-D, Bourdais R, Dumur D, Buisson J (2010) Distributed model predictive control for building temperature regulation. American Control Conference (ACC), 2010 3174–3179.
45. Worthmann K, Kellett CM, Braun P, et al (2015) Distributed and Decentralized Control of Residential Energy Systems Incorporating Battery Storage. IEEE Transactions on Smart Grid.
46. Georgiou TT (1988) On the computation of the gap metric. In: Proceedings of the IEEE Conference on Decision and Control.
47. Li X, Wen J (2014) Review of building energy modeling for control and operation. Renew Sustain Energy Rev 37:517–537.
48. Ljung L (1988) System identification toolbox. Matlab user's Guide.
49. Octopus Energy Agile pricing explained. https://octopus.energy/blog/agile-pricing-explained/. Accessed 26 Oct 2020