Model Predictive Control of Residential Central Heating using Economic and Weather-based Data

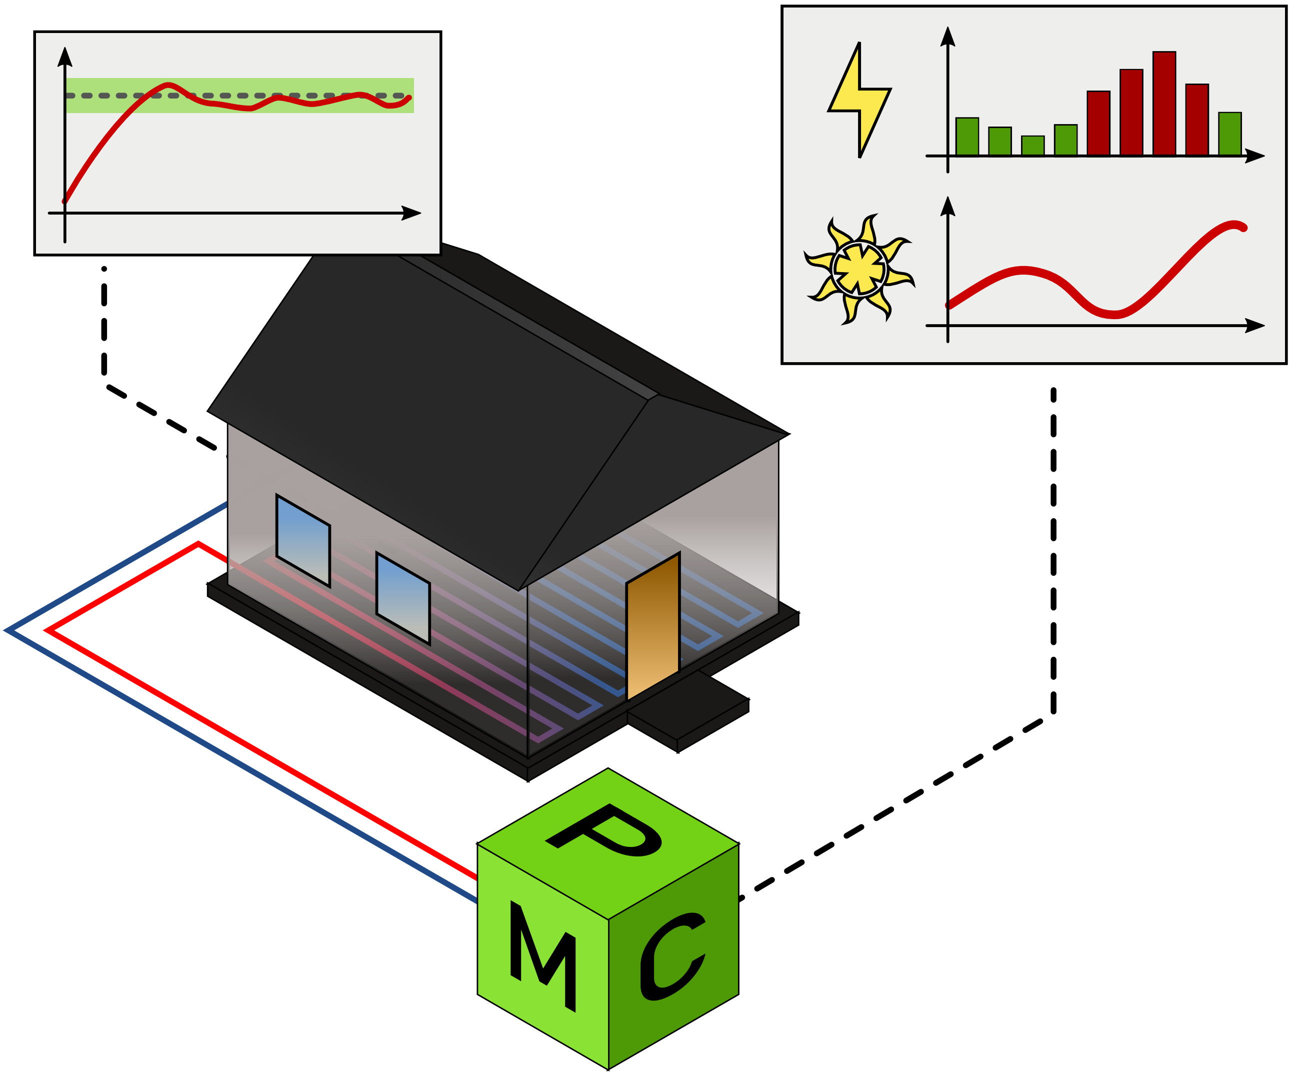

Control and Automation Master-Thesis

Andreas Just Poulsen & Christian Schneider Pedersen

Aalborg University

4^{th} Semester Control and Automation Master Programme

Date of Completion: June 1, 2023

AALBORG UNIVERSITY
STUDENT REPORT

Study Board of Electronics and IT
Aalborg University
http://www.aau.dk

Title:
Model Predictive Control of Residential Central Heating using Economic and Weather-based Data

Theme:
Master's Thesis

Project Period:
February 2023 - June 2023

Project Group:
CA10 1033

Participant(s):
Andreas Just Poulsen
Christian Schneider Pedersen

Supervisor(s):
Jan Dimon Bendtsen
John-Josef Leth

Copies: 0

Page Numbers: 79

Date of Completion:
June 1, 2023

Abstract:

This report investigates the design of a model predictive controller (MPC) for reducing operation costs of a central heating system in a single family residential building. A resistor-capacitor (RC) equivalent model for the residence is fitted to measured data using a subspace method, producing a discrete state space system with realistic dynamics and a small constant mean squared error. Using this model an MPC is designed to continuously solve a quadratic programming (QP) problem to obtain optimising plant inputs wrt. the designed cost based on several factors including Luenberger observer estimated states, forecast electricity price, and forecast ambient temperature. The output of the MPC is constrained to a chosen rate of change and range dependent on the ambient temperature disturbance signal. A dynamic prediction horizon is implemented corresponding to the Danish electricity price forecast schedule. The QP problem is revised as a direct term is needed to incorporate the electricity price in the reference tracking weight of the MPC, the value of which is found systematically by simulation. The fitted model and MPC are simulated and compared to a baseline MPC with no price weights. For the period 2022 to 2023, the simulation shows a 7.3 % reduction in the total cost of operation. The total savings of other simulated periods are generally less, likely due to a lack of price normalisation.

Contents

Preface

This Master's Thesis is written on the 4th semester of Control and Automation Master Programme at AAU in collaboration with Bitzer Electronics. The contents of this report are freely available.

The authors would like to thank Bitzer Electronics and especially the representative Kresten for the collaboration, providing data, and providing access to the house actuators. Special thanks to Simon Thorsteinsson for guidance and discussions related to the modelling of the provided residence.

Andreas Just Poulsen
<ajpo18@student.aau.dk>

Christian Schneider Pedersen
<cspe18@student.aau.dk>

1 Introduction

Residential buildings are one of the largest consumers of energy, mainly due to the heating required for comfortable living. In Denmark alone, residential buildings were responsible for 31 % of energy consumption in 2021, of which 81 % was used for residential heating [1]. Even small reductions or shifts in consumption locally may improve the cost-effectiveness of the energy grid as a whole, and are also of interest for the end-consumer if the costs are not too great. As a majority of residential occupants own their respective housing, they must themselves pay the steep up-front cost of improvement in energy effectiveness [2]. Therefore, that improvement has to be rather cost-effective to have any impact at all. One cost-effective strategy is shifting some energy consumption from time slots with a high energy price to time slots with a low energy price. This is further incentivised by the new tariff increases in Denmark and overall energy price inflation experienced in Europe in 2022 [3] [4].

A method of shifting the energy consumption while keeping thermal comfort of the building, is by utilising Thermal Energy Storage (TES) such as a designated water tank or thermal mass of the building. Residential thermal mass consists of floors, walls, roof, and items inside the building, all of which can be warmed up in advance, such that they release their heat during costly time slots and thus lowering the heating costs in this period. Especially floors and walls tend to have a large thermal capacity, allowing a great storage of energy at the cost of slow temperature changes.

In systems with central heating, sensors and actuators are likely to exist, which are used for local control. In conjunction with a model for the heating system and residential thermal mass, these sensors and actuators allow for the design of more advanced controllers, which are capable of shifting the energy consumption according to the forecast energy price. A first-principles construction of a residential thermal model can be onerous and expensive, and therefore an empiric grey-box model is often sufficient and preferable. However, grey-box modelling can be difficult when only sensors in the heating system are available, as they provide limited to no information of the conditions inside the rooms. In turn, this makes ensuring thermal comfort difficult as the observability of indoor temperatures becomes poor or impossible.

This project is a case-study of a single-family household in collaboration with the company Bitzer. A representative of Bitzer, currently living in the house, has provided access to the buildings sensors, actuators, and a database of sensor data over a time span of approximately one year. The house in question is heated centrally by an in-house heat pump, which draws geothermal heat from a connected ground pipe. This household has the aforementioned problem of scarce information, as sensors are mounted only on the end-points of the heating system. This is common problem in the field in which Bitzer operates, as most buildings they control are akin to the provided residence wrt. placement and type of sensors. The provided residence is further formalised in the following section.

1.1 Formalisation of Provided Residence

In order to construct a model of the provided residence, the layout and confines are established. This formalisation and the available sensors are presented in this section. The floor plan of the house is seen in Figure 1.1, which shows the ground level of the provided two-story building. Part of the building is not residence and this is marked by a grey background. Floor heating is installed in parallel for each room in the residence, resulting in a 7-way manifold at the hot side of the central heating system. The central heating system is a heat pump, marked **HP** on the figure, which draws energy from the brine of ground source heat pipes totalling a length of 450 m. Furthermore, the house has a wood-fired boiler marked **B**, and a hot-water tank marked **HW**, the latter providing hot water for showers, sink etc. The wood-fired boiler is disregarded for the rest of the project as it is not in use, and as such is considered disconnected. The blue lines indicate which rooms contain water sinks and/or sources. The dimensions of the building and rooms are known but omitted for clarity.

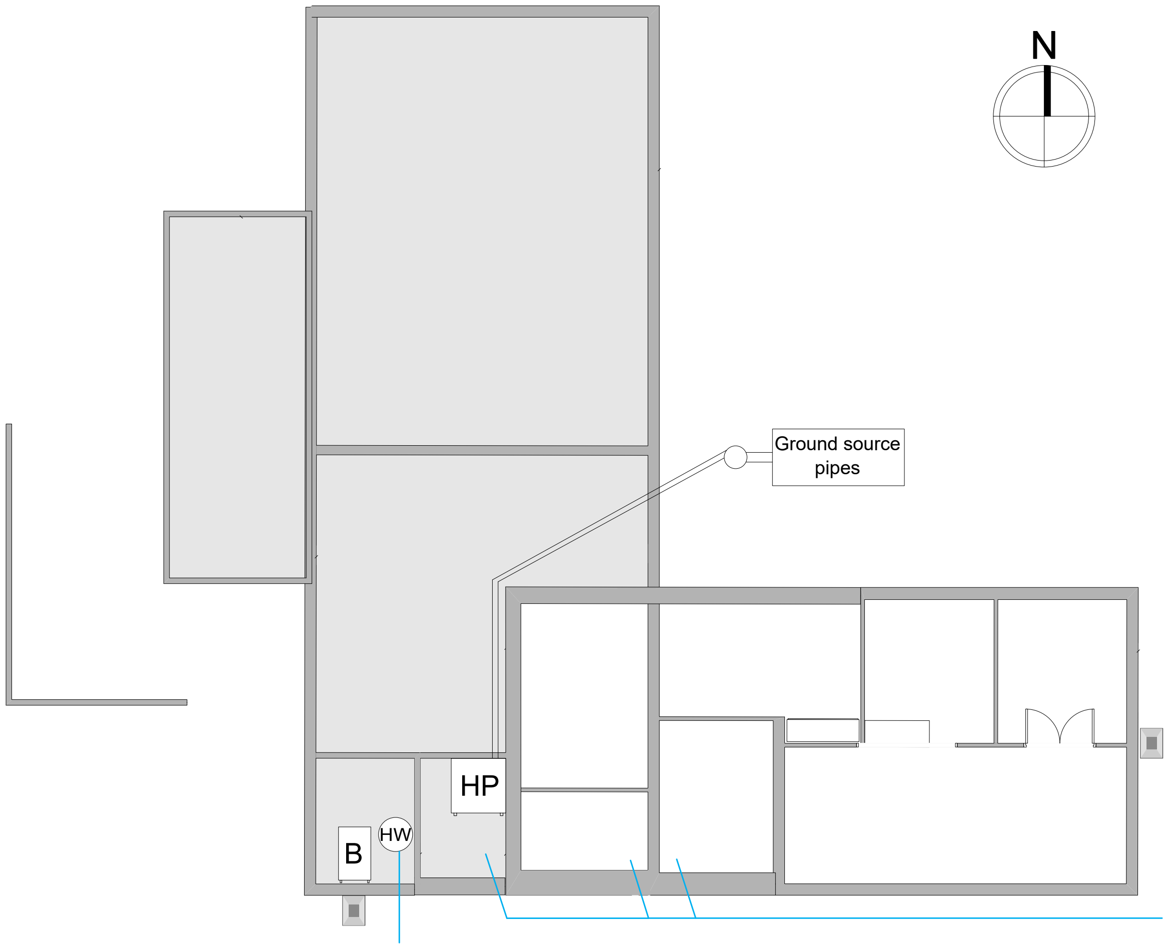

Figure 1.1: Ground floor plan of the provided residence. The boiler, hot water tank and heat pump are marked by **B**, **HW** and **HP** respectively.

Not pictured in Figure 1.1 is the upper floor, whose layout is different from that of the ground floor but not specified. On the upper floor exists a single radiator, which is in parallel with the floor heating below. It is dimensioned in such a way, that allows it to utilise the same input temperature as the floor heating while providing a similar comfort level.

The central heating system is further detailed in Figure 1.2, which illustrates the the three heating/-cooling cycles in the provided residence. These are:

1. The brine cycle seen to the left delivers heat from the ground pipe brine to the heat pump evaporator.
2. The refrigerant cycle seen in the middle transfers the absorbed heat by a refrigeration cycle to the water cycle.
3. The water cycle seen to the right delivers heat to the floor heating and radiator or the hot water tank.

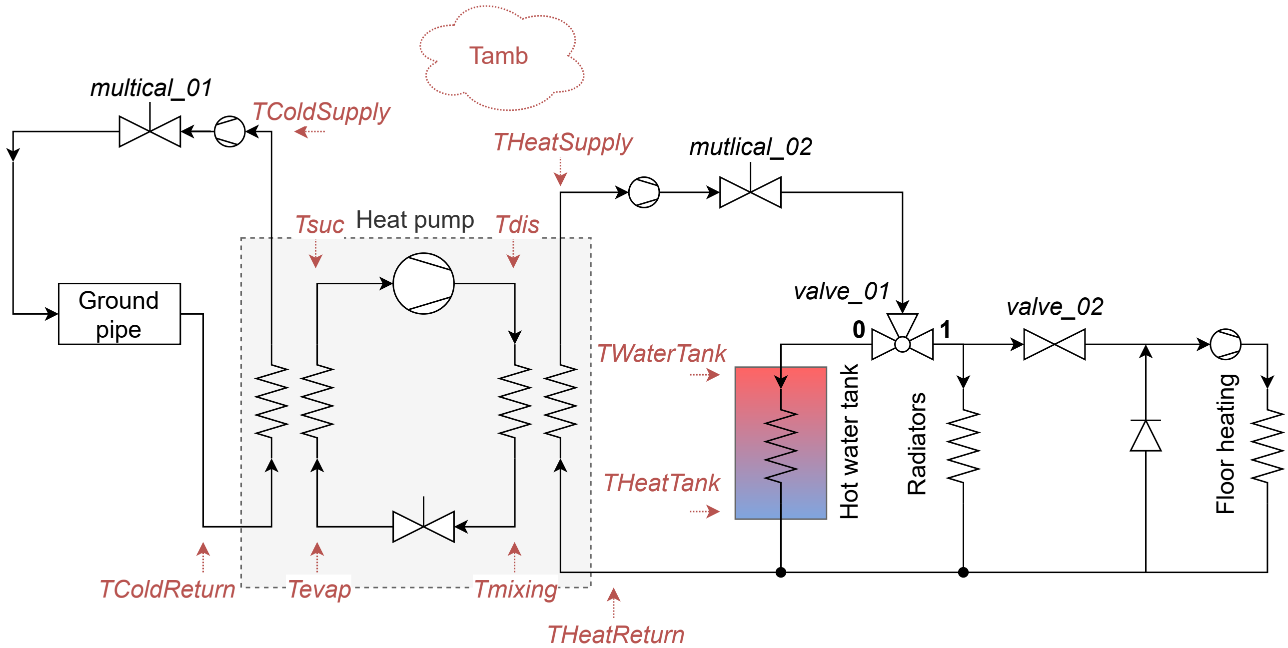

Figure 1.2: Diagram of the three heating cycles present in the provided residence.

Various temperature sensors are installed in the heating system, which are currently used for local control. These available temperature sensors are marked by red striped arrows in Figure 1.2. The ambient temperature *Tamb* is available from weather forecasts. Furthermore, two flow sensors are available and marked *multical_01* and *multical_02*. These flow sensors are capable of determining the energy flow in the cooling and heating cycles by measuring the temperature and mass flow.

During normal operation, a reference for the supply temperature is set, such that the inhabitants are comfortable. A local controller calculates and adjusts the speed of the compressor, such that this supply temperature is achieved. The valve *valve_01* is set to 1 such that the supply water is routed through the floor heating and radiator. This mode of operation is referred to as Floor Heating (FH) mode. During warm periods, the requested compressor speed may not be possible due to hardware limitations, at which point the local controller turns the compressor on and off repeatedly to same effect. This is referred to as on-off compressor operation. During even warmer periods, the compressor is turned off completely.

Another mode of operation occurs when the residents wish to use hot water, referred to as Hot Water Production (HWP) mode. In this mode, the local controller sets *valve_01* to 0 and turns up the compressor on the heat pump until the water tank temperature *TWaterTank* is at the desired temperature. The production of hot water is considered a priority.

The actuators in this system mainly reside within the heat pump. In fact, the main control variable for the heat pump is the speed of the compressor which controls the pressure at the condenser. It is also possible to utilise the pre-existing local heat pump controller and treat THeatSupply as an input. This may be favourable, as that very controller also redirects hot water to the water tank on the occasions it is needed. This means that the project may focus solely on control during the FH mode. This approach is chosen, the disadvantages of which are discussed in Section 5.3. Table 1.1 details each sensor available in the provided residence. Some sensor values are attributed mathematical symbols if they are utilised in this project.

Table 1.1: Available sensors in the provided residence, their attributed mathematical symbols, descriptions and units.

Sensor	Sym	Description	Unit
TColdSupply	-	Brine supply temperature at the cold side of the heat pump	°C
TColdReturn	-	Brine return temperature at the cold side of the heat pump	°C
multical_01	-	Brine mass flow at the cold side of the heat pump	kg/s
mutlical_01	-	Brine energy flow at the cold side of the heat pump	W
Tsuc	-	Refrigerant temperature at the inlet of the compressor	°C
Psuc	-	Pressure at the inlet of the compressor	bar
Tdis	-	Refrigerant temperature at the outlet of the compressor	°C
Pdis	-	Pressure at the outlet of the compressor	bar
Tmixing	-	Refrigerant temperature at the inlet of the expansion valve	°C
Tevap	-	Refrigerant temperature at the outlet of the expansion valve	°C
EM340	-	Power sensor for the heat pump	W
THeatSupply	$T_{w,s}$	Water supply temperature at the hot side of the heat pump	°C
THeatReturn	$T_{w,r}$	Water return temperature at the hot side of the heat pump	°C
multical_02	-	Water mass flow at the hot side of the heat pump	kg/s
multical_02	-	Water energy flow at the hot side of the heat pump	W
TWaterTank	T_{hw}	Water temperature at the top of the hot water tan	°C
THeatTank	-	Water temperature at the bottom of the hot water tank	°C
Tamb	T_a	Ambient temperature	°C

1.2 Problem formulation

In light of the problems highlighted earlier, it is of interest to investigate the possibility of using the thermal mass of a house to act as a TES, such that the end-consumer cost may be lowered, without compromising the occupants comfort. This is achieved by constructing a controller, which considers the building model, outside temperature etc. while minimising the aggregated costs of heating. One such controller strategy is Model Predictive Control (MPC), which is an increasingly popular area of study in the realm of residential heating [5]. Specifically, MPC of a building with scarce indoor information is stated as a point of interest for the collaborating company Bitzer Electronics. This leads to the following problem formulation:

> How can a model predictive controller be designed which optimises the end-consumer cost of residential heating by utilising residential thermal masses and considering forecast ambient temperature and electricity prices?

To do this, a grey-box model of the system is built by constructing a simplified resistance-capacitance model of the thermal system and fitting the state space entries using a subspace method. This resistance-capacitance model is constructed iteratively by addition of new features and states until a desirable input-output characteristic is found. This model is tested with the various provided data sets to validate its realism.

The fitted models form the basis for the prediction model and cost of two model predictive controllers, the first of which acts as a baseline. This baseline controller aims to maximise the comfort of the residents by minimising the error between indoor room temperature and the desired reference thereof. The general theory for predictive control is presented and the controller structure is introduced. Furthermore, the cost of operating the central heating system is estimated, which is used to compare both controllers alongside the relative comfort. A second model predictive controller is designed, which utilises the residential heating mass while considering the comfort level and forecast electricity prices. This controller is compared to the baseline controller by their simulated responses and cost of operation.

Finally, it is discussed whether a model of the disturbances may improve the stability of the MPC, how the electricity price can be normalised to account for general price evolution and whether a model of the central heating heat pump would improve the MPC response. A reflection on the feasibility of integration in the provided residence is also made and the overall functionality of both controllers is discussed.

2 Modelling of the Provided House

To enable an MPC to utilise the thermal storage of the provided residence, a model of the residence along with the heating system is constructed. This model is not all-encompassing as it is not necessary for the main objective of investigating MPC design for residential heating. The success criteria of this model is purely based on the realism of its output wrt. the input, and realistic behaviour of internal states or values is not required. As such, a simplified thermal model is constructed, and the provided data is used to determine the coefficients. This is henceforth referred to as grey-box modelling. Three approaches are investigated, the last of which is employed for the remainder of this report.

All three approaches are based on simplified resistance-capacitance (RC) models of the thermal system. In such models, temperature is analogous to voltage while heat flow is analogous to current. By extension, degrees Celsius (°C) is analogous to volts (V) while joules per second (J/s) is analogous to ampere (A). States are the nodes within the system which have dynamics, that is nodes with connected capacitors. Inputs to the system are modelled as voltage sources and outputs are modelled as some linear combination of the states. This method is also used in other literature [6]. The method is further exemplified by Figure 2.1 which shows a first order RC-circuit. Kirchhoff's current law manifests a dynamic equation as seen in Equation (2.1) from which a state space model or an ordinary differential equation (ODE) can be constructed.

$$0 = T_{in}\frac{1}{R} - \frac{d}{dt}T_1 C, \tag{2.1}$$

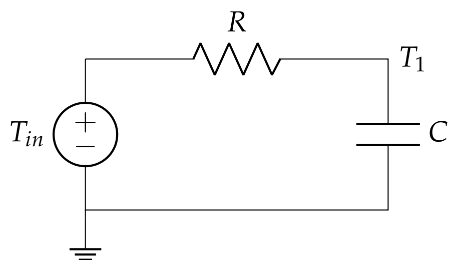

Figure 2.1: Example of an RC circuit equivalent of a thermal system.

2.1 Linear State Space Approach

An initial RC equivalent model is built with six states, three pipe temperature states $T_{w,i}$, an output temperature state T_{out}, a floor temperature state T_f and a room temperature state T_r. Only the output temperature state T_{out} is considered measurable. The input temperature T_{in}, which represents the supply water temperature $T_{w,s}$, and ambient temperature T_a are considered inputs. The RC equivalent circuit diagram for this model is shown in Figure 2.2 and the equivalent state space equations are seen in Equations (2.2).

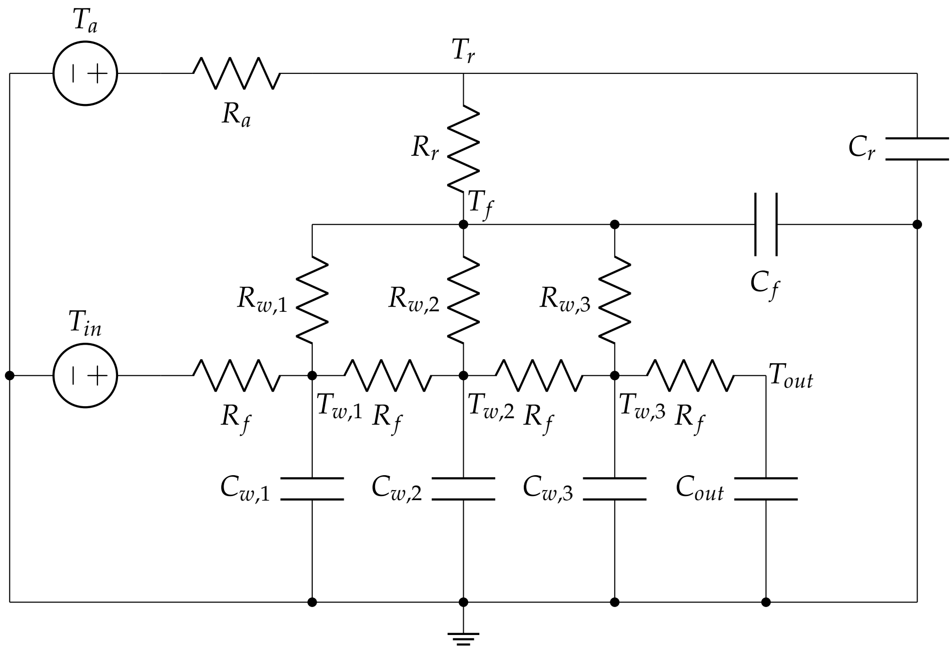

Figure 2.2: Initial 6th order RC model for the linear state space approach.

$$\dot{x} = Ax + Bu = \Gamma^{-1}(\tilde{A}x + \tilde{B}u) \tag{2.2a}$$

$$y = T_{w,out} = Cx + Du \tag{2.2b}$$

$$x = \begin{bmatrix} T_{w,1} & T_{w,2} & T_{w,3} & T_{out} & T_f & T_r \end{bmatrix}^T \tag{2.2c}$$

$$u = \begin{bmatrix} T_{in} & T_a \end{bmatrix}^T \tag{2.2d}$$

$$\Gamma = \mathrm{diag}(\begin{bmatrix} C_{w,1} & C_{w,2} & C_{w,3} & C_{out} & C_f & C_r \end{bmatrix}) \tag{2.2e}$$

$$\tilde{A} = \begin{bmatrix} -2\frac{1}{R_f} - \frac{1}{R_{w,1}} & \frac{1}{R_f} & 0 & 0 & \frac{1}{R_{w_1}} & 0 \\ \frac{1}{R_f} & -2\frac{1}{R_f} - \frac{1}{R_{w,2}} & \frac{1}{R_f} & 0 & \frac{1}{R_{w,2}} & 0 \\ 0 & \frac{1}{R_f} & -2\frac{1}{R_f} - \frac{1}{R_{w,3}} & \frac{1}{R_f} & \frac{1}{R_{w,3}} & 0 \\ 0 & 0 & \frac{1}{R_f} & -\frac{1}{R_f} & 0 & 0 \\ \frac{1}{R_{w,1}} & \frac{1}{R_{w,2}} & \frac{1}{R_{w,3}} & 0 & -\frac{1}{R_r} - \frac{1}{R_{w,1}} - \frac{1}{R_{w,2}} - \frac{1}{R_{w,3}} & \frac{1}{R_r} \\ 0 & 0 & 0 & 0 & \frac{1}{R_r} & -\frac{1}{R_r} - \frac{1}{R_a} \end{bmatrix} \tag{2.2f}$$

$$\tilde{B} = \begin{bmatrix} \frac{1}{R_f} & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & \frac{1}{R_a} \end{bmatrix} \tag{2.2g}$$

$$C = \begin{bmatrix} 0 & 0 & 0 & \alpha_1 & 0 & 0 \end{bmatrix} \tag{2.2h}$$

$$D = \begin{bmatrix} 0 & 0 \end{bmatrix} \tag{2.2i}$$

where:

T_{in}	is the water supply temperature	[°C]
$T_{w,i}$	is the water temperature in the i$^{\text{th}}$ segment of pipe	[°C]
T_f	is the temperature of the flooring	[°C]
T_r	is the temperature of the room	[°C]
R_f	is the thermal resistance between pipe segments	[°C · s/J]
$R_{w,i}$	is the thermal resistance between pipe segment i and flooring	[°C · s/J]
R_r	is the thermal resistance between flooring and room	[°C · s/J]
R_a	is the thermal resistance between room and outside	[°C · s/J]
$C_{w,i}$	is the thermal capacitance of the i$^{\text{th}}$ pipe segment	[J/°C]
C_f	is the thermal capacitance of the flooring	[J/°C]
C_r	is the thermal capacitance of the room	[J/°C]

The resistance between each pipe segment is modelled to be the reciprocal of the energy flow due to mass flow, that is:

$$R_f = (q \cdot c_w)^{-1} \tag{2.3}$$

where:

q	is the mass flow of water	[kg/s]
c_w	is the specific heat capacity of water	[J/kg/°C]

The initial RC equivalent model is observable, which is checked symbolically in MATLAB® with finite nonzero RC values by examining the rank of the observability matrix. In a similar fashion it is confirmed that the system is controllable. However, to make fitting of the parameters easier, it is decided to combine the floor temperature and room temperature into one, and discarding the output temperature state and measuring the last pipe temperature instead.

2.1.1 Fitting to Outlier Replaced Data

As mentioned previously, a large amount of data is available for fitting the coefficients of this RC-model. However, in the provided data the ambient temperature is constantly −70 °C for a large majority of the data sets, which is why it is replaced with historical weather data from the Open-Meteo service[1]. A method for fitting state space models based on subspace methods is detailed in Ljung [10] and available in MATLAB® as the function `ssest()`. The function works by accepting an *identifiable* state space object and providing an *input-output data* object to which the coefficients are then fitted. In particular, the identifiable state space object allows for some structuring such as constraining certain entries in the state space matrices, which will be relevant later. The main problem of this approach is the available data not fitting the dynamics of the designed RC equivalent. This is mainly due to the HWP which happens fairly often and can be seen as the large spikes in the supply temperature in Figure 2.3.

[1]Open-Meteo data is available at `www.open-meteo.com` and is based on the ERA5 [7] ERA5-land [8] and CERRA [9] projects.

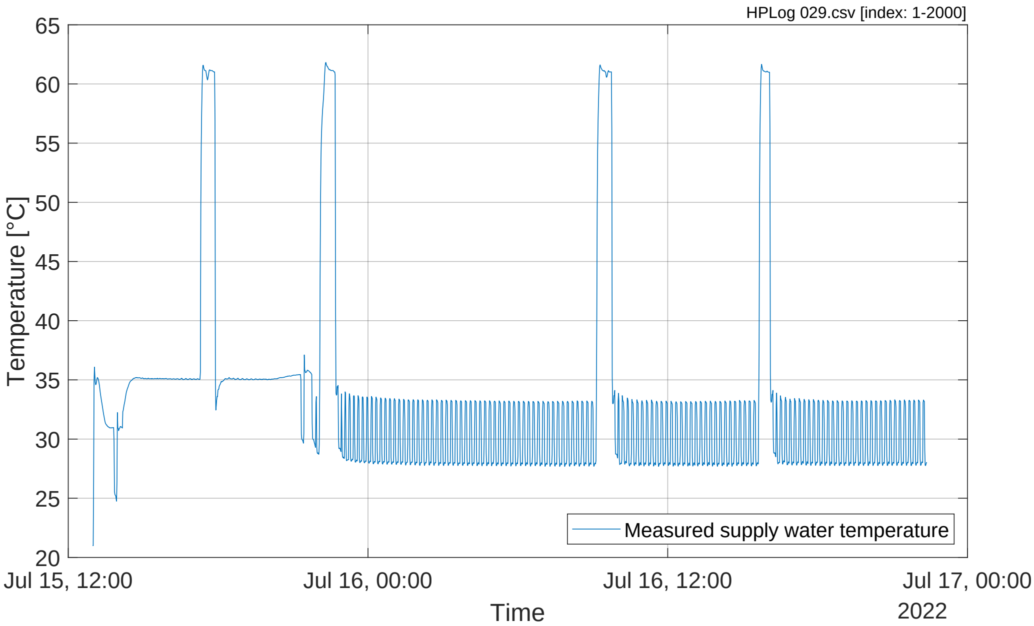

Figure 2.3: Water supply temperature for a time span of the provided data. The label in the top right indicates the used data set and how it is indexed.

The model is designed only with regards to FH operation, and when exposed to the input, which includes the HWP mode spikes, oscillations occur in the output signal as seen in Figure 2.4. As `ssest()` tries to minimise the error between model and measurement, the model must follow the HWP spikes. This requires fast dynamics wrt. the input which then oscillates due to the on-off compressor behaviour described in Section 1.1, finally producing unwanted oscillations in the modelled return temperature. As a consequence of these poor results, this method is disregarded.

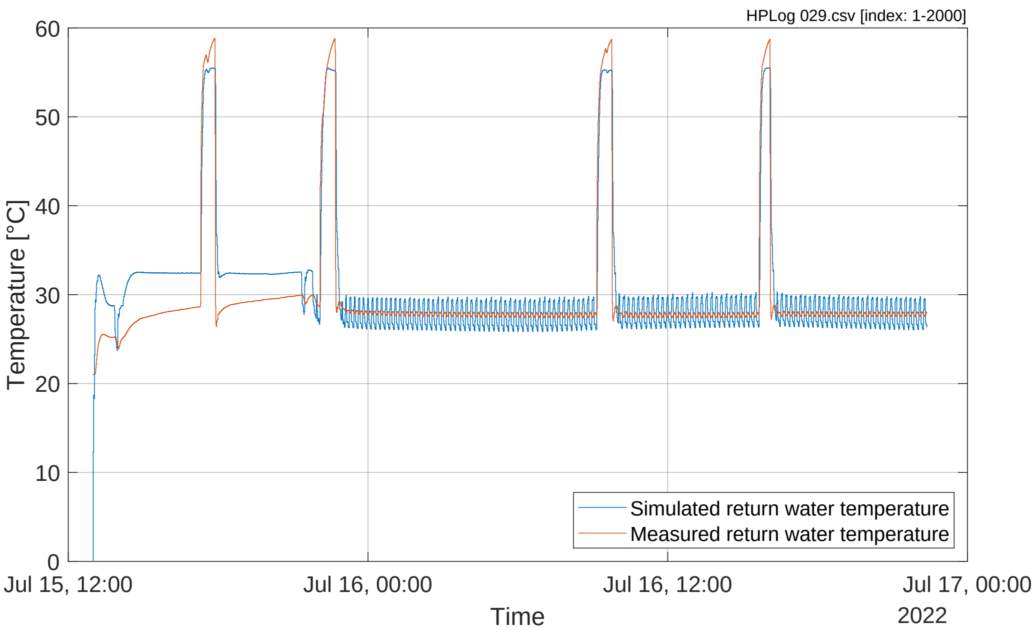

Figure 2.4: Return water temperature measurement and simulated response for the 5$^{\text{th}}$ order fitted model.

2.1.2 Fitting to Separate Experiments

The periods in which the system operates in HWP mode is considered a disturbance, in which the system is controlled by a different controller. The goal of this project is controlling the system outside of this region, and therefore time spans without hot water production are mainly of interest. One way of solving the problem with unmodelled HWP dynamics in the underlying fitted model is by removing them entirely. Two methods for removing hot water production spikes are attempted. The first is done by detecting outliers by deviation from the median by some percentile and replacing affected values with the nearest non-outlier value. Additional data clean-up is then performed by low-pass filtering the inputs and outputs in order to reduce the high frequency dynamics visible in Figure 2.3. Replacing the outliers leads to discontinuities when the remaining time spans are stitched together. in turn leading to unrealistic dynamics and poor fits.

The second attempt at removing hot water spikes mitigates the problem of discontinuity by fitting to various *experiments* rather than the whole data set as before. An experiment, in this context, is a time span of data between two HWP spikes. This time span is found with a non-causal moving average filter, seen in Equation (2.4).

$$y[k+1] = \frac{1}{N}\left(x[k] + \sum_{i=1}^{N/2} x[k-i] + \sum_{i=1}^{N/2} x[k+i]\right) \tag{2.4}$$

where:

N	is the window size
$y[k]$	is the moving average process
$x[k]$	is the current evaluated sample

While a moving average process of the input, denoted $\hat{T}_{in,ma}$, is above some threshold, denoted $\hat{T}_{in,tr}$, it is assumed that the system produces hot water. That is:

- $\hat{T}_{in,ma} < \hat{T}_{in,tr}$ means that no hot water is being produced
- $\hat{T}_{in,ma} \geq \hat{T}_{in,tr}$ means that hot water is being produced

For the sake of clarity, the time spans in which hot water is being produced is denoted a HWP region and time spans in between are separate experiments in the FH regions. The experiments start the sample after HWP ends and stop the sample before HWP starts.

The threshold $T_{in,tr}$ cannot be constant, as several data sets include steps and have different working temperature areas. Therefore, the threshold must vary along with the input, and it is chosen to construct the threshold as a moving average process with a greater window size than the previous. Specifically, good results have been achieved with a window size of 75 for $\hat{T}_{in,ma}$ and 2000 for $\hat{T}_{in,tr}$. An example of the input T_{in} and the corresponding moving average and threshold is seen in Figure 2.5. Remark that detected HWP regions are denoted by the marked area.

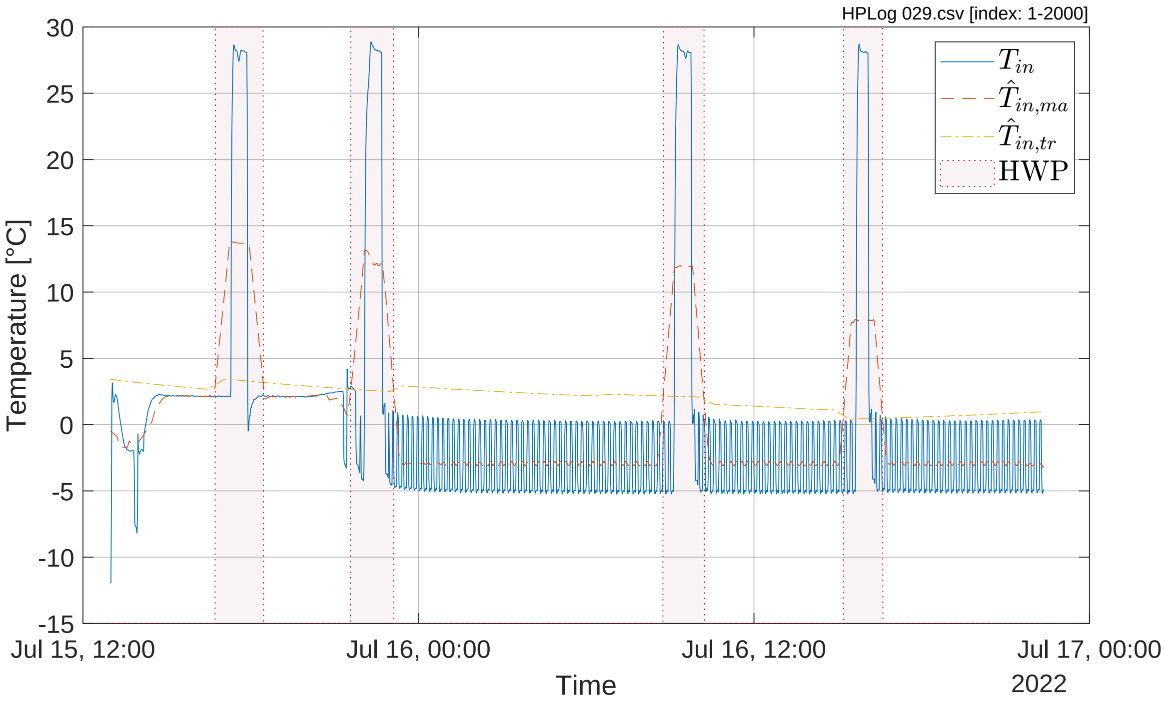

Figure 2.5: The supply water temperature T_{in}, the moving average temperature $\hat{T}_{in,ma}$, and the moving average threshold $\hat{T}_{in,tr}$. The temperatures have been normalised to the median. The area denoted by HWP is the hot-water production region.

The model introduced earlier is fitted on multiple experiments found by the aforementioned method. Each experiment is therefore just a small set of sample points in between the spikes, typically between 1000 and 3000 samples with a sample period $\tau_s = 60\,\text{s}$. This, however, also yields poor results. This may be due to the initial states supplied to the subspace method causing the modelled response to fail converging in the short available time. In general, this RC equivalent model does not provide sufficient results, mainly due to problems fitting unmodelled dynamics and difficulty in separating modelled from unmodelled dynamics prior to fitting. A new RC model is constructed to account for the HWP dynamics.

2.2 One Segment Switched-Input State Space Approach

During HWP, the dynamics of the system change drastically due to the valve bypassing the FH. This switched nature of the system can be incorporated by adding a switch in the RC equivalent circuit from before, which transports water to a new pipe segment which in turn, exchanges heat with the water in the hot water storage tank. This gives one new input and one new state as seen in the circuit diagram in Figure 2.6. The Boolean switch $S \in 0, 1$ is 1 during FH operation and 0 during HWP operation.

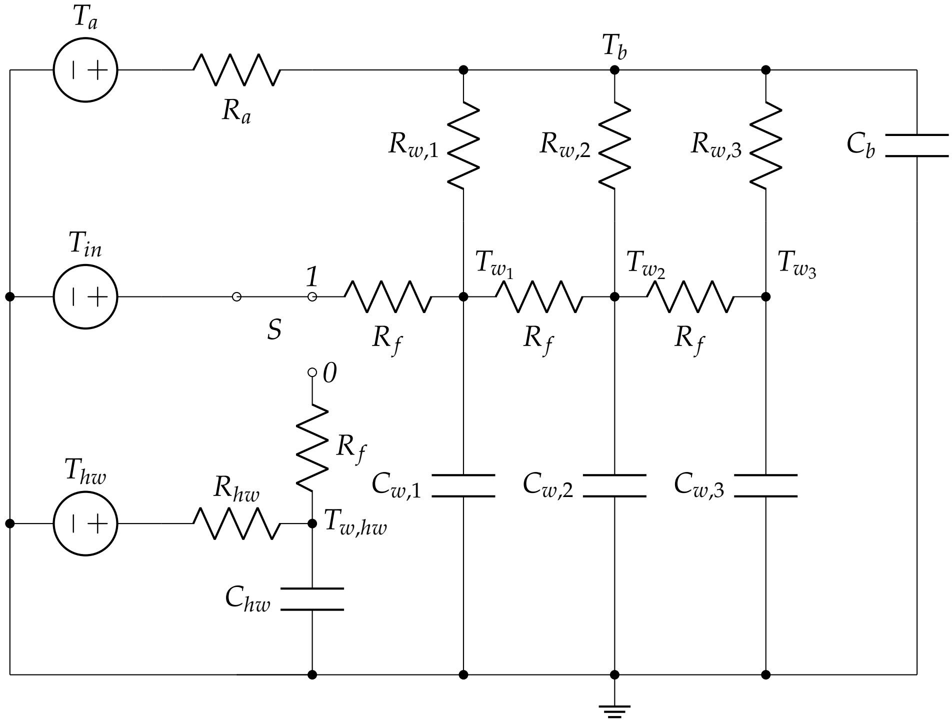

Figure 2.6: 5th order switched RC model for the one segment switched-input state space approach.

where:

T_{hw}	is the temperature of the hot water tank	[°C]
T_b	is the joint temperature of the building	[°C]
$T_{w,hw}$	is the water temperature of the pipe segment near the hot water tank	[°C]
C_{hw}	is the thermal capacitance of the pipe segment near the hot water tank	[J/°C]
C_b	is the thermal capacitance of the building	[J/°C]

Notice that the switch is positioned right after T_{in} and this allows for the switched input to be treated as *two* non-switched inputs. The principle of this transformation is shown in Figure 2.7. In essence, this allows for a single state space representation of Figure 2.6 with two inputs ST_{in} and $\tilde{S}T_{in}$, explicitly stated in Equation (2.5).

$$ST_{in} = \begin{cases} T_{in}, & S = 1 \\ 0, & S = 0 \end{cases} \qquad \tilde{S}T_{in} = \begin{cases} 0, & S = 1 \\ T_{in}, & S = 0 \end{cases} \tag{2.5}$$

This state space is essentially two state space representations in one, one part belonging to the system for $S = 1$ and one belonging to the system for $S = 0$. The states used for $S = 1$ have no effect on $T_{w,hw}$ and vice versa.

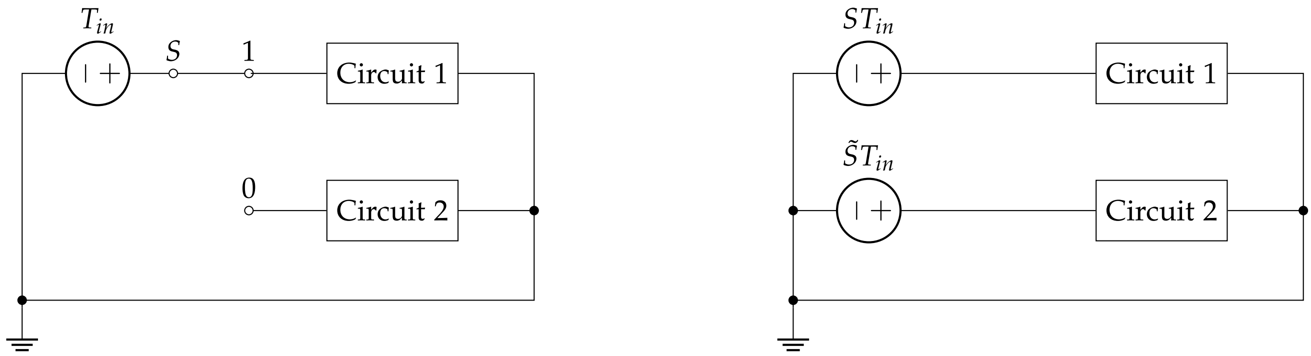

(a) A switched single-input system. S is a Boolean switch.

(b) A transformation of Figure 2.7a in which the switched input is split into two inputs. $\tilde{S}$ is the logical NOT of S.

Figure 2.7: Transformation between a single switched input and two individually switched inputs.

Figure 2.8 shows an example of how this switched input is practically produced. The switch to hot water production is detected non-causally using the moving average filter as described in Section 2.1.2. Remark that a low-pass filter has been applied to the inputs before splitting. Furthermore, remark that the final non-zero $\tilde{S} \cdot T_{in,(lp)}$ is a negligible one sample in length and exists due to the algorithm used.

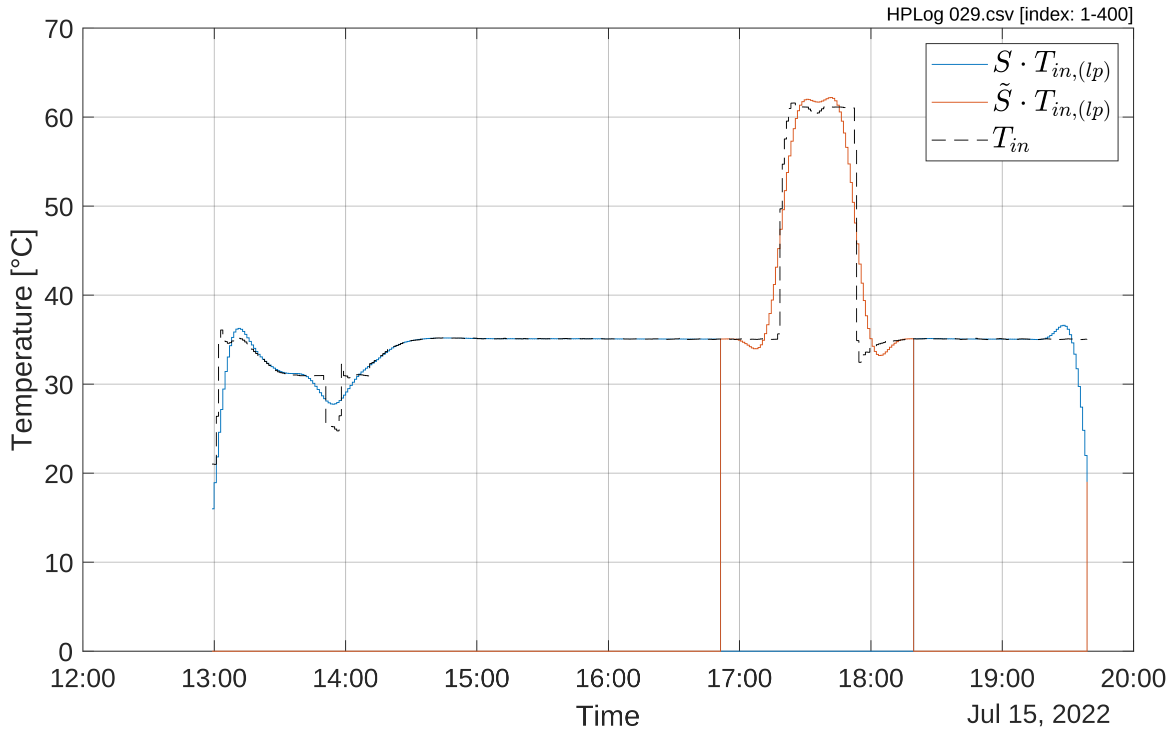

Figure 2.8: Water supply temperature split into two low-pass-filtered inputs depending on HWP. HWP is detected by a non-causal moving average filter described in Section 2.1.2.

2.2.1 State Space Representation

The state space equation with switched inputs of the RC equivalent in Figure 2.6 is seen in Equations (2.6).

$$\dot{x} = Ax + Bu = \Gamma^{-1}(\tilde{A}x + \tilde{B}u) \tag{2.6a}$$

$$y = T_{w,r} = Cx + Du \tag{2.6b}$$

$$x = \begin{bmatrix} T_{w,1} & T_{w,2} & T_{w,3} & T_b & T_{w,hw} \end{bmatrix}^T \tag{2.6c}$$

$$u = \begin{bmatrix} ST_{in} & T_a & \tilde{S}T_{in} & T_{hw} \end{bmatrix}^T \tag{2.6d}$$

$$\Gamma = \mathrm{diag}(\begin{bmatrix} C_{w,1} & C_{w,2} & C_{w,3} & C_b & C_{hw} \end{bmatrix}) \tag{2.6e}$$

$$\tilde{A} = \begin{bmatrix} -2\frac{1}{R_f} - \frac{1}{R_{w,1}} & \frac{1}{R_f} & 0 & \frac{1}{R_{w,1}} & 0 \\ \frac{1}{R_f} & -2\frac{1}{R_f} - \frac{1}{R_{w,2}} & \frac{1}{R_f} & \frac{1}{R_{w,2}} & 0 \\ 0 & \frac{1}{R_f} & -\frac{1}{R_f} - \frac{1}{R_{w,3}} & \frac{1}{R_{w,3}} & 0 \\ \frac{1}{R_{w,1}} & \frac{1}{R_{w,2}} & \frac{1}{R_{w,3}} & -\frac{1}{R_{w,1}} - \frac{1}{R_{w,2}} - \frac{1}{R_{w,3}} - \frac{1}{R_a} & 0 \\ 0 & 0 & 0 & 0 & -\frac{1}{R_f} - \frac{1}{R_{hw}} \end{bmatrix} \tag{2.6f}$$

$$\tilde{B} = \begin{bmatrix} \frac{1}{R_f} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & \frac{1}{R_a} & 0 & 0 \\ 0 & 0 & \frac{1}{R_f} & \frac{1}{R_{hw}} \end{bmatrix} \tag{2.6g}$$

$$C = \begin{bmatrix} 0 & 0 & \alpha_1 & 0 & \alpha_2 \end{bmatrix} \tag{2.6h}$$

$$D = \begin{bmatrix} 0 & 0 & 0 & 0 \end{bmatrix} \tag{2.6i}$$

It is chosen to fit the model in Equations (2.6) to a single data set using the `ssest()` algorithm described earlier. The data set `HPLog 029.csv` is chosen as it is sufficiently long, has several HWP spikes, and different steady states. The initial parameters are chosen by hand and the system is discretized using zero-order hold in MATLAB®. Additionally, the matrices in the identifiable state space object are constrained, such that fitting may only change the non-zero values of A, B, C, and D in Equations (2.6). Remark that the `ssest()` algorithm is, in this case, extremely sensitive to initial parameters and initial states. Therefore, obtaining a good result requires many attempts at manually selecting parameters and initial states, as no suitable alternative was found. The initial parameters used in the `ssest()` model are seen in Table 2.1.

Table 2.1: Initial parameters and states used for fitting by `ssest()` to the 5^{th} order hybrid model in Equations (2.6).

Parameter	$C_{w,1}$	$C_{w,2}$	$C_{w,3}$	C_b	C_{hw}	-
Value [J/°C]	$1 \cdot 10^{-9}$	$1 \cdot 10^{-9}$	$1 \cdot 10^{-9}$	$50 \cdot 10^{-9}$	$1 \cdot 10^{-9}$	-
Parameter	R_f	$R_{w,1}$	$R_{w,2}$	$R_{w,3}$	R_a	R_{hw}
Value [s°C/J]	50	200	200	200	$1 \cdot 10^3$	200
Parameter	α_1	α_2	-	-	-	-
Value [·]	1	1	-	-	-	-
Initial state	$T_{w,1}$	$T_{w,2}$	$T_{w,3}$	T_b	$T_{w,hw}$	-
Value [°C]	10	10	20	10	20	-

The fit is performed and the coefficients are noted in Equations (2.7). Remark that the matrices are not necessarily constrained to the form in Equations (2.6). This is due to the discretization method used. This discretization maintains the disconnect between the FH and HWP subsystems.

$$x[k+1] = Ax[k] + Bu[k] \tag{2.7a}$$

$$y[k] = Cx[k] + Du[k] \tag{2.7b}$$

$$A = \begin{bmatrix} -0.329 & -0.319 & 0 & 0.394 & 0 \\ -0.083 & -0.041 & -0.288 & -0.006 & 0 \\ 0 & -0.581 & 0.009 & 0.092 & 0 \\ -0.293 & 0.038 & -0.191 & 0.252 & 0 \\ 0 & 0 & 0 & 0 & -0.483 \end{bmatrix} \tag{2.7c}$$

$$B = \begin{bmatrix} 1.212 & 0.000 & 0 & 0 \\ 1.000 & 0.000 & 0 & 0 \\ 1.000 & 0.000 & 0 & 0 \\ 0.937 & 0.175 & 0 & 0 \\ 0 & 0 & 1.593 & 0.098 \end{bmatrix} \tag{2.7d}$$

$$C = \begin{bmatrix} 0 & 0 & 1.236 & 0 & 0.782 \end{bmatrix} \tag{2.7e}$$

$$D = \begin{bmatrix} 0 & 0 & 0 & 0 \end{bmatrix} \tag{2.7f}$$

The observability and controllability of the model in Equations (2.7) is investigated using the same method utilised earlier. The fitted models controllability and observability matrices both have full rank, their conditioning numbers being 34.7 and 223.1 respectively. As this model essentially contains two subsystems, investigating the subsystems controllability and observerability is also of interest. For the FH system, corresponding to the first 4 states, the controllability and observability matrix conditioning numbers are 34.7 and 36.1 respectively. For the HWP system, corresponding to the 5$^{\text{th}}$ state, the condition numbers are 1 and 1 respectively.

2.2.2 Fitting Results

Figure 2.9 shows the output response of the fit compared to the measured data. The fitted response has a mean square error (MSE) of 3.0058 °C^2 which is due to the visible offset. Specifically, the fitted model reacts faster to changes in input, resulting in slightly unrealistic steps. This is apparent in Figure 2.10 which shows the same fit, zoomed in on the step occurring around the 30$^{\text{th}}$ of July 2022, also denoted 30-07-2022.

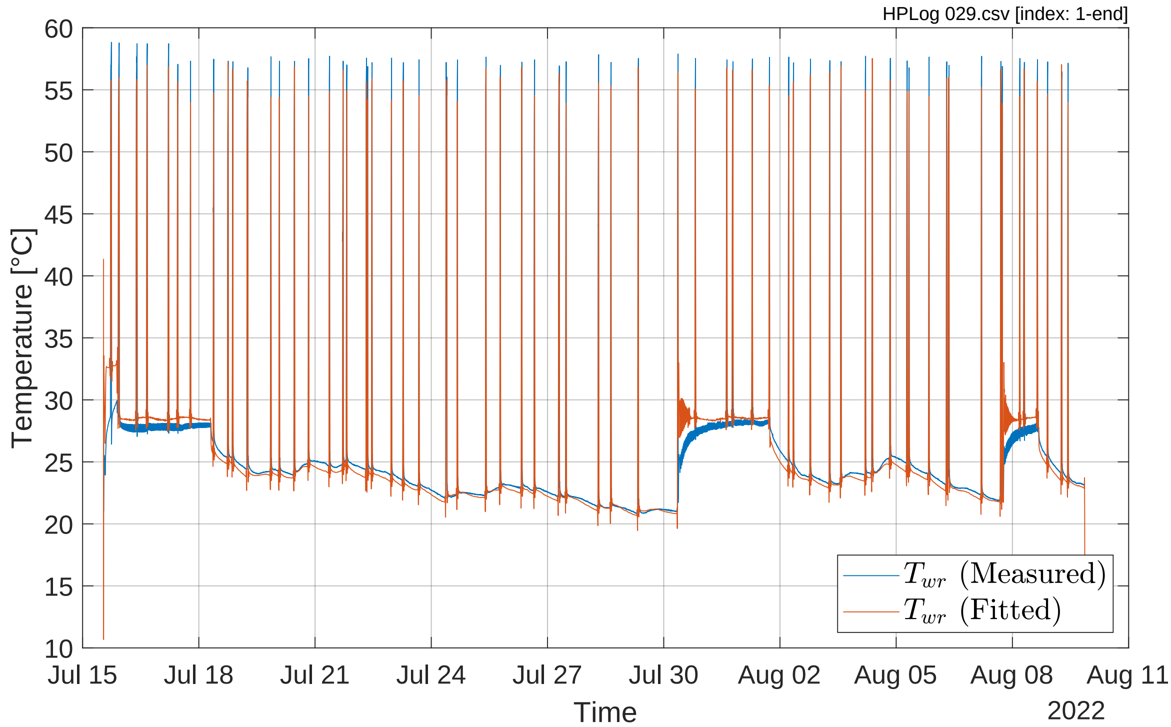

Figure 2.9: Output response of model in Equations (2.6) fitted to the data set `HPLog 029.csv` and compared with the data set `HPLog 029.csv`.

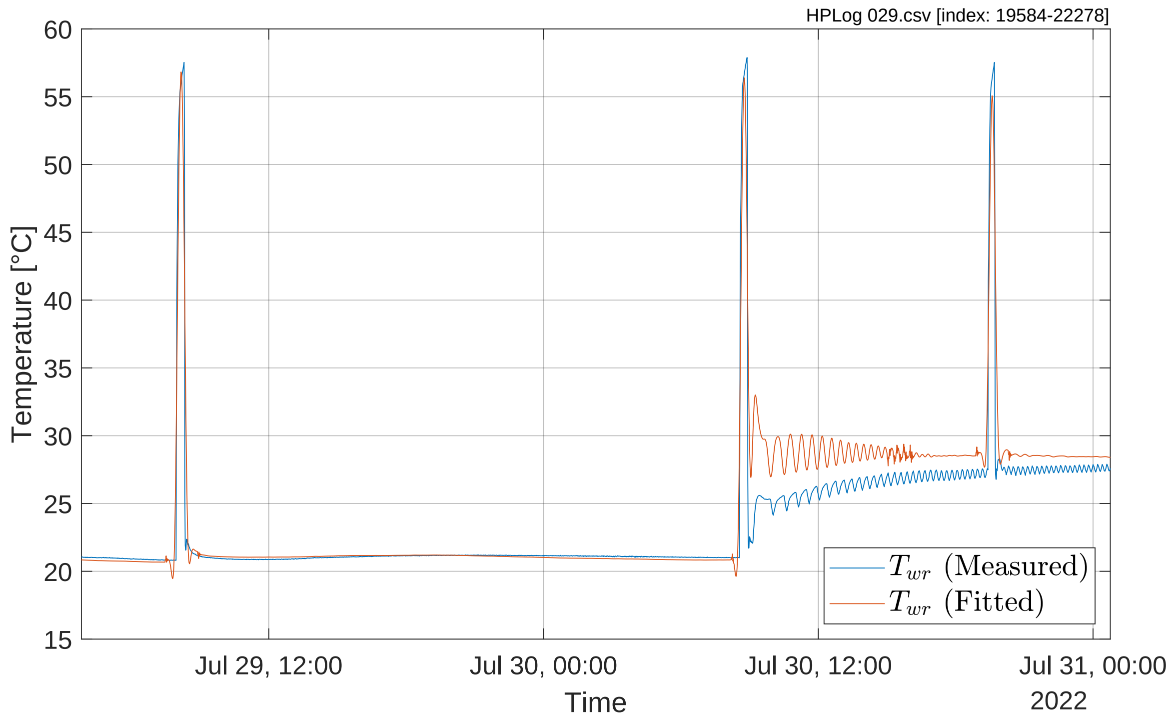

Figure 2.10: Output response of model in Equations (2.6) fitted to the data set `HPLog 029.csv` and compared with the data set `HPLog 029.csv`. Zoomed in around 30-07-2022.

In order to further confirm that the fitted model behaves sufficiently realistic, it is compared against different data sets. To do this validation, the fitted model of data set `HPLog 029.csv` is utilised with the input/output data of every other data set. For each simulation of the model, the result is compared to the measurements and an MSE is calculated. Table 2.2 shows the best three and the worst three results. The average MSE for all data sets is 6.42 °C^2 for each time step and 8.08 °C^2 per data set. Remark that some data sets were excluded due to being too short, that is, less than 120 minutes. This is the case for 12 of the 43 data sets available, but since they are so short, almost no data is lost. Remark also that the heat pump was turned off for the entirety of data set HPLog 026.csv which is the likely cause of the small MSE.

Table 2.2: Three smallest and largest mean square errors (MSE) between simulation using the fitted model in Equations (2.7) and measured data sets.

Category	MSE	Data set
Smallest error	1.27 °C^2	HPLog 028.csv
	2.57 °C^2	HPLog 026.csv
	2.68 °C^2	HPLog 015.csv
Largest error	34.40 °C^2	HPLog 004.csv
	17.94 °C^2	HPLog 034.csv
	17.92 °C^2	HPLog 043.csv

Figure 2.11 and Figure 2.12 show the simulated and measured data for the data sets corresponding to the smallest and largest MSE respectively. Incorrectly estimated initial values are one of the main causes of larger MSE.

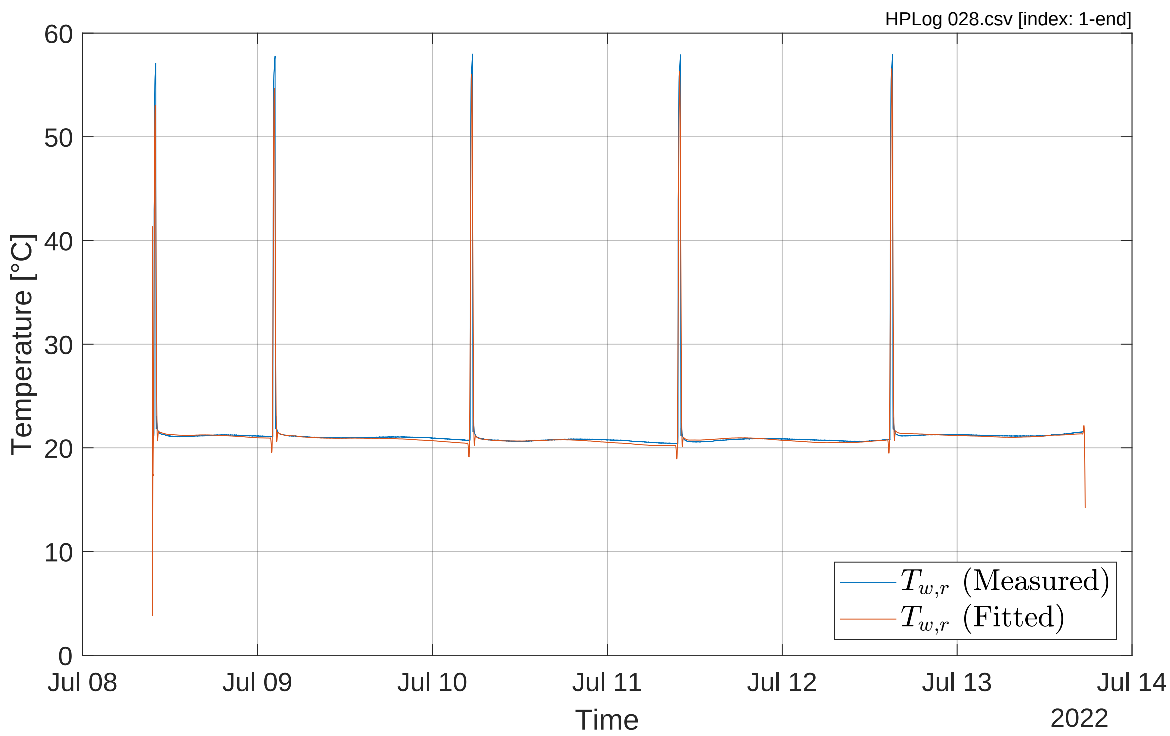

Figure 2.11: Output response of model in Equations (2.6) fitted to the data set `HPLog 029.csv` and compared with the data set `HPLog 028.csv`.

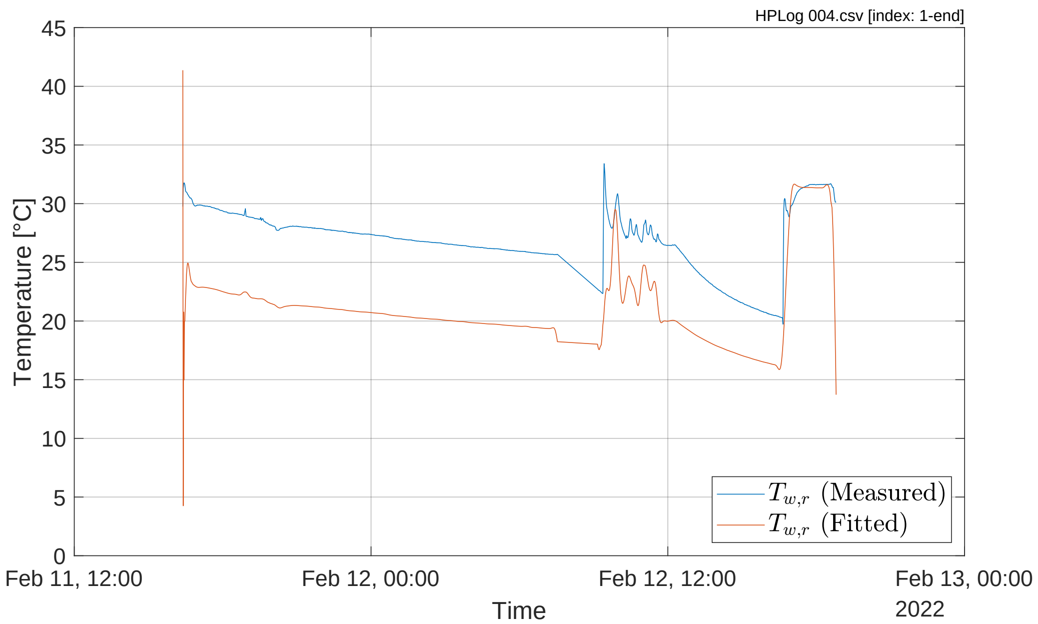

Figure 2.12: Output response of model in Equations (2.6) fitted to the data set `HPLog 029.csv` and compared with the data set `HPLog 004.csv`.

2.3 Two Segment Switched-Input State Space Approach

As explained in Chapter 1, the provided house does not generally have a temperature sensor in inhabited spaces. However, on request of the authors a temperature sensor has been installed, and a small data set with said sensor is available. The one segment model from before only has a joint building temperature T_b which has quite odd dynamics as seen in the model response in Figure 2.13. This is not surprising, as this state can represent many things as a result of `ssest()` and is not fitted to any real measurement. This is in of itself not problem, since the only success criteria of the model is the input-output characteristic, which was proved to be sufficiently realistic as per the previous section.

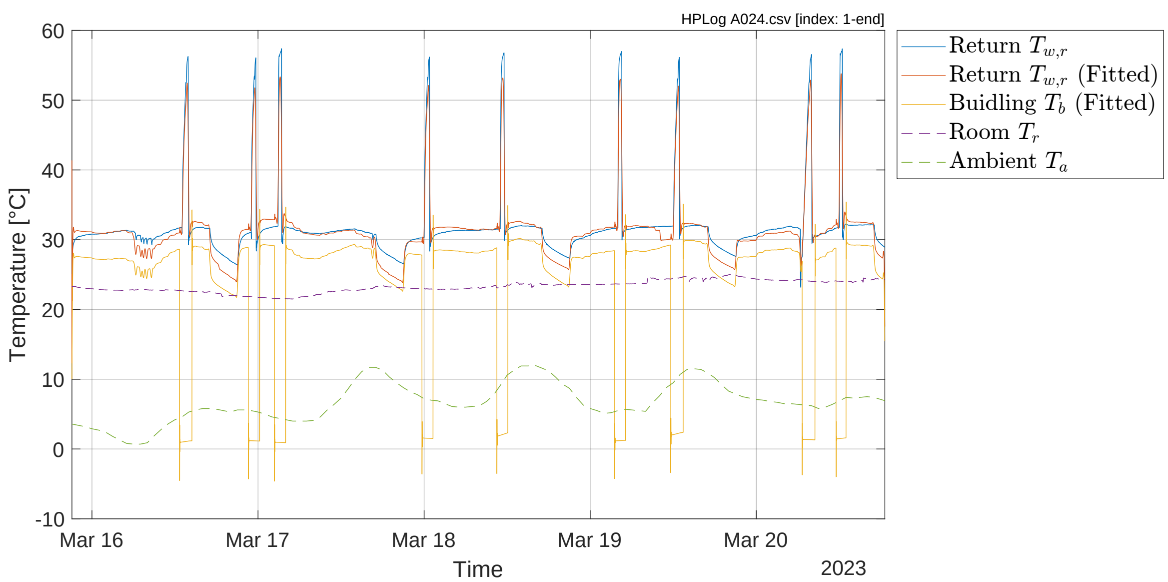

Figure 2.13: Simulated return water temperature and building temperature alongside measured return temperature, room temperature and ambient temperature for a period of provided data.

The main purpose of the heating system is to provide a comfortably heated environment for its users. Therefore, a cost should be attached to the inside temperature, such that large deviations from the reference are penalised. This requires the indoor temperature being either measurable or observable, and therefore the one segment approach from before is insufficient as it has no such state or and because T_b does not seem to accurately represent or follow the room temperature, as expected of a state not tied to any measurements. A new model is constructed in order to account for the room dynamics. The idea behind this new model and approach is utilising the provided room temperature as an output for the model while fitting, hoping to capture more realistic dynamics of the buildings thermal capacitance. To do this, the model is altered such that the one segment state of the building T_b is once again split into a floor temperature T_f and a room temperature T_r as was the case in the model in Figure 2.2. Figure 2.14 shows the RC-equivalent of this model and Equations (2.8) shows the corresponding state space model.

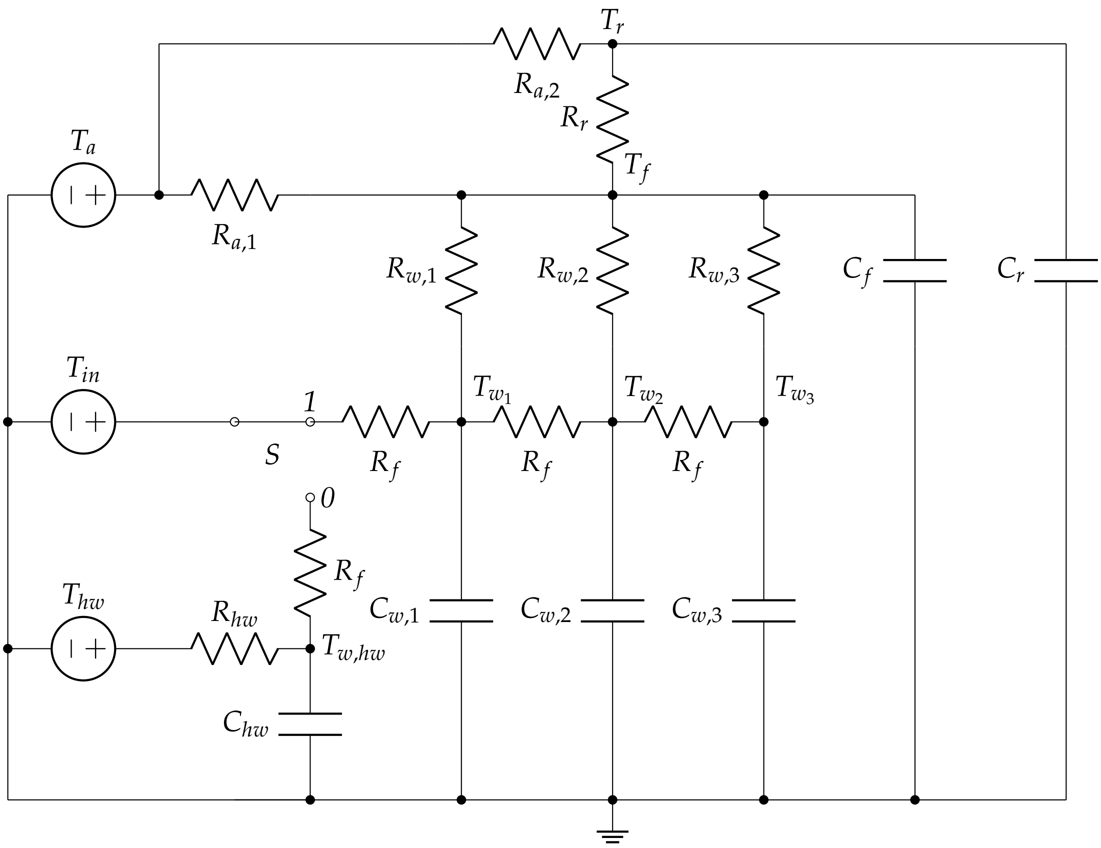

Figure 2.14: 6th order RC model for the two segment switched-input state space approach.

where:

$R_{a,1}$	is the thermal resistance between floor and outside	$[°C \cdot s/J]$
$R_{a,2}$	is the thermal resistance between room and outside	$[°C \cdot s/J]$

$$\dot{x} = Ax + Bu = \Gamma^{-1}(\tilde{A}x + \tilde{B}u) \tag{2.8a}$$

$$y = \begin{bmatrix} \alpha_1 T_{w,3} + \alpha_2 T_{w,hw} \\ T_r \end{bmatrix} = Cx + Du \tag{2.8b}$$

$$x = \begin{bmatrix} T_{w,1} & T_{w,2} & T_{w,3} & T_f & T_r & T_{w,hw} \end{bmatrix}^T \tag{2.8c}$$

$$u = \begin{bmatrix} ST_{in} & T_a & \tilde{S}T_in & T_{hw} \end{bmatrix}^T \tag{2.8d}$$

$$\Gamma = \mathrm{diag}(\begin{bmatrix} C_{w,1} & C_{w,2} & C_{w,3} & C_f & C_r & C_{hw} \end{bmatrix}) \tag{2.8e}$$

$$\tilde{A} = \begin{bmatrix} -2\frac{1}{R_f} - \frac{1}{R_{w,1}} & \frac{1}{R_f} & 0 & \frac{1}{R_{w_1}} & 0 & 0 \\ \frac{1}{R_f} & -2\frac{1}{R_f} - \frac{1}{R_{w,2}} & \frac{1}{R_f} & \frac{1}{R_{w,2}} & 0 & 0 \\ 0 & \frac{1}{R_f} & -\frac{1}{R_f} - \frac{1}{R_{w,3}} & \frac{1}{R_{w,3}} & 0 & 0 \\ \frac{1}{R_{w,1}} & \frac{1}{R_{w,2}} & \frac{1}{R_{w,3}} & -\frac{1}{R_r} - \frac{1}{R_{a,1}} - \sum_{i=1}^{3} \frac{1}{R_{w,i}} & \frac{1}{R_r} & 0 \\ 0 & 0 & 0 & \frac{1}{R_r} & -\frac{1}{R_r} - \frac{1}{R_{a,2}} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{R_f} - \frac{1}{R_{hw}} \end{bmatrix} \tag{2.8f}$$

$$\tilde{B} = \begin{bmatrix} \frac{1}{R_f} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & \frac{1}{R_{a,1}} & 0 & 0 \\ 0 & \frac{1}{R_{a,2}} & 0 & 0 \\ 0 & 0 & -\frac{1}{R_f} & \frac{1}{R_{hw}} \end{bmatrix} \tag{2.8g}$$

$$C = \begin{bmatrix} C_1 \\ C_2 \end{bmatrix} = \begin{bmatrix} 0 & 0 & \alpha_1 & 0 & \alpha_2 \\ 0 & 0 & 0 & \alpha_3 & 0 \end{bmatrix} \tag{2.8h}$$

$$D = \begin{bmatrix} D_1 \\ D_2 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{2.8i}$$

Using the data sets with indoor temperature measurements, the fitting is performed with `ssest()` as described earlier. Unlike the previous approach, the parameters are not chosen by hand, as this gave poor results. Instead, the space of initial parameters are explored using Monte Carlo methods. In particular, for 1000 iterations, the parameters are drawn from a uniform distribution between 0 and 1000 and the fit is performed. The best case fit has an average MSE of 2.62 °C^2 and is achieved with the parameters seen in Table 2.3.

Table 2.3: Initial Monte Carlo parameters and states for the best fit by `ssest()` for the two segment 6th order hybrid model in Equations (2.8).

Parameter	$C_{w,1}$	$C_{w,2}$	$C_{w,3}$	C_f	C_r	C_{hw}	-	-
Value [J/°C]	134.1	670.6	86.1	965.7	377.0	568.8	-	-
Parameter	R_f	$R_{w,1}$	$R_{w,2}$	$R_{w,3}$	R_r	$R_{a,1}$	$R_{a,2}$	R_{hw}
Value [s°C/J]	556.1	607.4	289.7	843.7	966.6	483.1	494.8	197.8
Parameter	α_1	α_2	α_3	-	-	-	-	-
Value [$\cdot$]	1	1	1	-	-	-	-	-
Initial state	$T_{w,1}$	$T_{w,2}$	$T_{w,3}$	T_f	T_r	$T_{w,hw}$	-	-
Value [°C]	10	10	20	10	20	20	-	-

By using the parameters in Table 2.3 the state space system in Equations (2.8) can be written as

Equations (2.9). Remark that a value of 0.000 indicates a rounded non-zero value below 0.0004 while a 0 represents strictly 0.

$$A = \begin{bmatrix} -0.021 & -0.026 & 0.000 & -0.001 & 0.000 & 0 \\ -0.038 & -0.006 & -0.023 & -0.018 & 0.000 & 0 \\ 0.000 & -0.024 & -0.025 & 0.001 & 0.000 & 0 \\ -0.013 & -0.003 & 0.016 & 0.998 & -0.001 & 0 \\ 0.000 & 0.000 & 0.000 & 0.001 & 1.000 & 0 \\ 0 & 0 & 0 & 0 & 0 & -0.025 \end{bmatrix} \tag{2.9a}$$

$$B = \begin{bmatrix} 0.999 & 0.000 & 0 & 0 \\ 1.000 & 0.000 & 0 & 0 \\ 1.000 & 0.000 & 0 & 0 \\ 0.001 & 0.000 & 0 & 0 \\ 0.000 & 0.000 & 0 & 0 \\ 0 & 0 & 0.979 & -0.008 \end{bmatrix} \tag{2.9b}$$

$$C = \begin{bmatrix} 0 & 0 & 0.973 & 0 & 0 & 0.979 \\ 0 & 0 & 0 & 0 & 0.992 & 0 \end{bmatrix} \tag{2.9c}$$

$$D = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{2.9d}$$

The observability and controllability of the model in Equations (2.9) is investigated using the same method utilised earlier. The fitted models controllability and observability matrices both have full rank, their conditioning numbers being $2.9060 \cdot 10^4$ and $3.8860 \cdot 10^5$ respectively. As this model also contains two subsystems, investigating the subsystems controllability and observerability is also of interest. For the FH system, corresponding to the first 5 states, the controllability and observability matrix conditioning numbers are $3.0660 \cdot 10^4$ and $1.1383 \cdot 10^6$ respectively. For the HWP system, corresponding to the 6^{th}, the condition numbers are 1 and 1 respectively. The conditioning numbers for the first subsystem, which is the one intended to be controlled in this project, is rather large. Through direct experimentation it is, however, deemed not to be a problem.

2.3.1 Fitting Results

Figure 2.15 shows the output response of the fit compared to the measured data. The fitted response has an average MSE of 2.62 $°\text{C}^2$. Remark the small difference between the fitted room temperature and measured room temperature, which was the goal of this model design.

An MSE is also calculated between the measured and fitted response for every other data set as well. This MSE does not take into account the room temperature residual as that cannot be generated due to the lack of room temperature data generally. In essence, the MSE is only calculated from output C_1 in Equations (2.8). Table 2.4 shows the MSE between simulated return water temperature and fitted return water temperature. The average MSE is 5.41 $°\text{C}^2$ for each time step and 6.78 $°\text{C}^2$ per data set. This is a 15.7 % and 20.5 % respective reduction in MSE compared to the one segment model in Section 2.2. Notice, that both `HPLog A023.csv` and `HPLog A024.csv` have been used for fitting. `HPLog 026.csv`, which is the largest data set, is the third best fitting data set with an MSE of 2.57 $°\text{C}^2$, which was likewise the second best fit in the previous fitting with the same MSE of 2.57 $°\text{C}^2$. This alludes to a similar good fit. Similarly the worst fitting of `HPLog 006.csv` is a rather short data set of about a day, akin to the worst fitting of the previous model `HPLog 004.csv` which is about the same length.

Figure 2.15: Output response of model in Equations (2.8) fitted to the data sets `HPLog A023.csv`,`HPLog A024.csv`, `HPLog A025.csv`, and `HPLog A026.csv`, which is compared with the data set `HPLog A023.csv`.

Table 2.4: Three smallest and largest mean square errors (MSE) between simulation using the fitted model in Equations (2.8) and measured data sets.

Category	MSE	Data set
Smallest error	2.26 °C^2	HPLog A023.csv
	2.43 °C^2	HPLog A024.csv
	2.57 °C^2	HPLog 026.csv
Largest error	27.25 °C^2	HPLog 006.csv
	16.78 °C^2	HPLog 034.csv
	16.29 °C^2	HPLog 043.csv

The room temperature measurement C_2 cannot generally be verified, except for the few sets of data onto which the model is fitted. This leads to a further investigation of the simulated responses to see if the room temperature acts sufficiently realistic. Using the model, a simulation is made in which the water supply temperature is raised and then lowered. With `HPLog 026.csv` as a starting point, the water supply temperature is split into 4 parts. In the second and fourth quarter of the data, the supply temperature is set to 5 °C and 50 °C respectively. Figure 2.16 shows the response of the fitted two segment switched input model with this modified input.

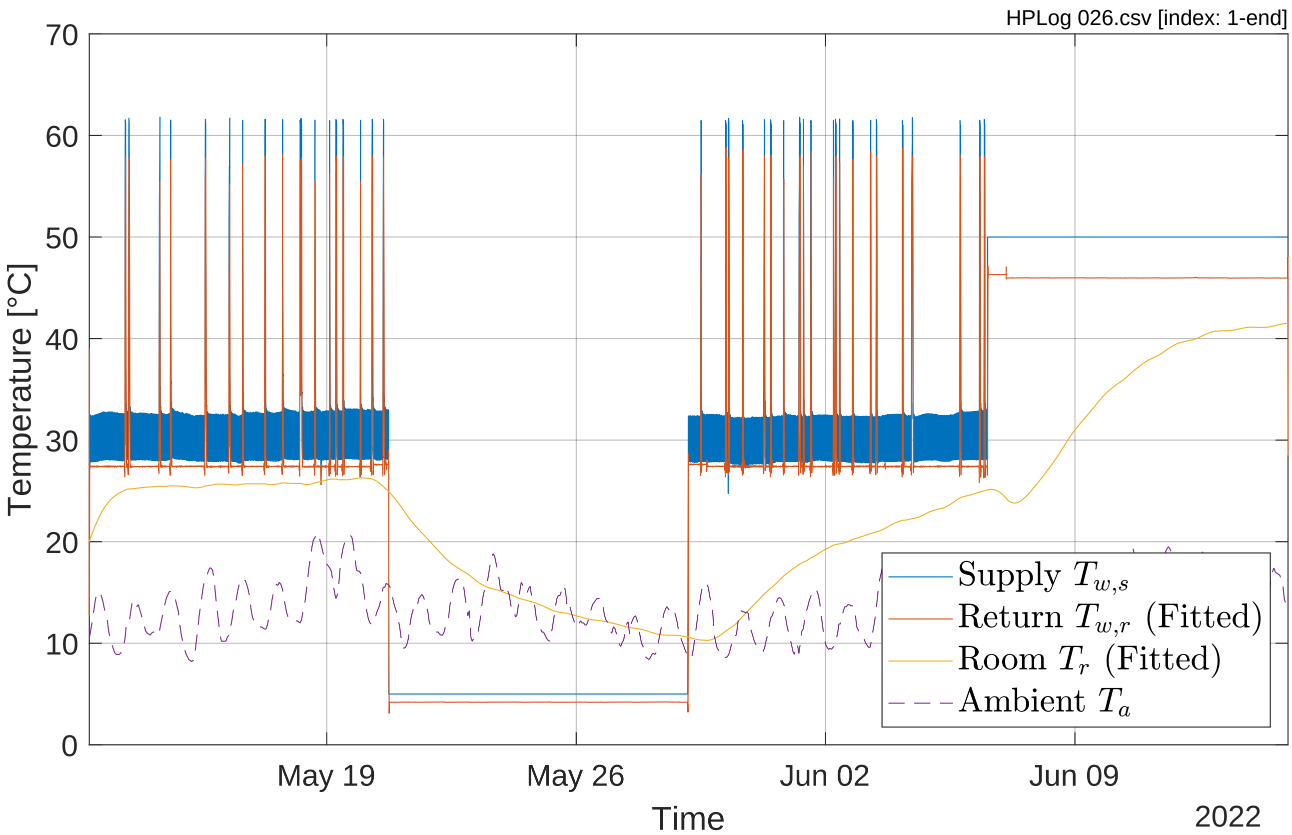

Figure 2.16: Two segment switched input model room temperature response to changes in the supply temperature.

As seen on Figure 2.16 when the supply temperature is set to 5 °C in the second quarter, the room temperature slowly falls and approaches the ambient temperature. When the supply temperature reverts to the data set inputs, the room temperature slowly starts to rise and climbs higher when set to 50 °C. At the ends of each quarter it takes some time for the system to start rising or falling. This is caused by the peak detection algorithm which detects the sudden increase in input as HWP peaks thus applying the step as an HWP input, which is why the room is not heated. The input is detected as a HWP peak as the moving average process is lower than the moving average threshold process as described in Section 2.1.2. The lengths of these HWP peaks are approximately 1 day.

As mentioned in the beginning of this chapter, the model aims to behave realistically wrt. the output with a given input. Since these states are not fitted to any output, their behaviours cannot generally be relied on, which is evident when looking at the state behaviours in Figure 2.17, which corresponds to the simulation response shown in Figure 2.16. As is seen, the floor temperature is usually sub-zero, while the hot-water tank temperature and water pipe temperatures dip into the sub-zero range depending on operation mode. These behaviours are the results of the grey-box estimation method used, and, while illogical, provide good results on the room temperature.

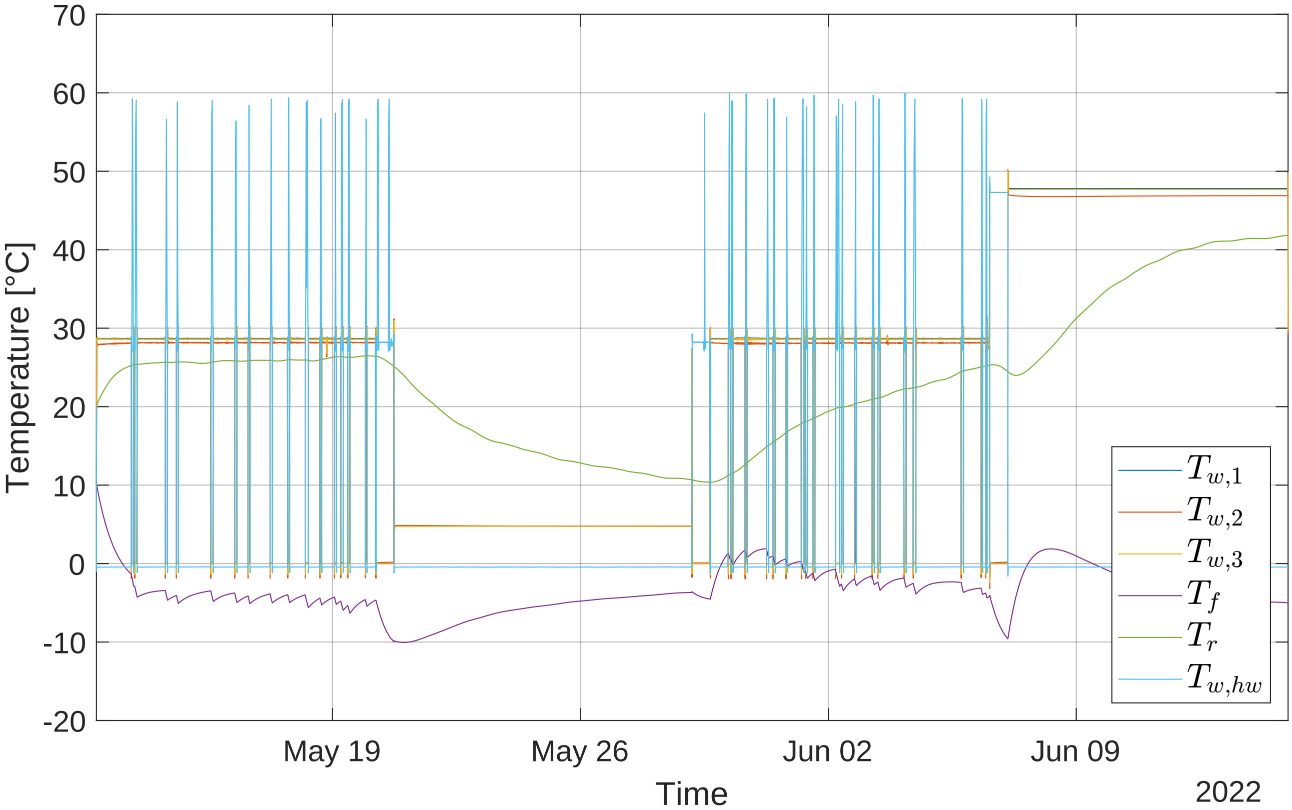

Figure 2.17: Two segment switched input model room temperature states corresponding to the response in Figure 2.16.

Generally, this model is deemed appropriate for initial controller design, as the return temperature and room temperature have realistic behaviours and follow the given data well. This is despite the odd state behaviour which is disregarded as only the input-output relation is of interest. Still, this model lacks inputs such as solar radiation and heat radiation from inhabitants, and large variations in these may cause the fitted model to act unrealistically. Furthermore, the ill-conditioned observability and controlability matrices suggest difficulty in controller design, but due to the slow nature of the system, this may not become a problem. These problems are set aside for now, as a model predictive controller is designed, which maximises the comfort of the inhabitants.

3 Comfort Model Predictive Controller

With the model behaving sufficiently realistic, the comfort model predictive controller (MPC) of the provided residence can be designed. The problem of controlling central heating is well-suited for predictive control for the following reasons:

- Actuator limits such as range and rate of change can be incorporated.
- Disturbances known in advance such as ambient temperature and electricity prices can be taken into account.
- The slow dynamics of the system allow for slow control rates, which may be necessary as MPC operation can be computationally costly.

These reasons for selecting predictive control schemes are not unique to this system, and are generally the reasons why one would chose to construct an MPC [11]. This chapter details the cost function and general structure of MPC, how the problem is constrained and how it is implemented in simulation. Lastly, the controller is downsampled and simulated results are presented.

3.1 MPC Theory

MPC is a control scheme which calculates an input based on an optimisation problem, such that a cost function J is minimised [11]. This cost function J typically includes reference error and is typically a function of the change in input [11]. In essence, the MPC produces the optimal changes to the input, denoted $\Delta\hat{\mathcal{U}}(k)$, by the argument which minimises J as seen in Equation (3.1).

$$\Delta\hat{\mathcal{U}}(k) = \begin{bmatrix} \Delta\hat{u}(k) \\ \Delta\hat{u}(k+1) \\ \vdots \\ \Delta\hat{u}(k+H_u) \end{bmatrix} = \underset{\Delta\mathcal{U}(k)}{\operatorname{argmin}} \left(J(\Delta\mathcal{U}(k), r(k), d(k), \hat{z}(k), H_p, H_u) \right) \tag{3.1}$$

where:

$\Delta\mathcal{U}(k)$	is the vector of changes in input from time k to $k+H_u$
$\Delta\hat{\mathcal{U}}(k)$	is the optimising change in input at time k to $k+H_u$
$r(k)$	is the reference signal from k to $k+H_p$
$d(k)$	is the disturbance signal from k to $k+H_p$
$\hat{z}(k)$	is the predicted output to be controlled
H_p	is the prediction horizon
H_u	is the control horizon

The input to the system at time k, $u(k)$ is then chosen as the last input plus the change, i.e:

$$u(k) = u(k-1) + \Delta u(k) \tag{3.2}$$

The future references and disturbances are used together with an internal model in the optimisation step to calculate the tracking error. The prediction horizon, H_p, determines how far in time the internal model is advanced while the control horizon, H_u, determines how far in time the optimising input is considered. That is, the MPC examines the model H_p steps into the future and selects H_u inputs which minimise the cost function.

For clarity, the general system structure used in this chapter is seen in Figure 3.1, in which a controller supplies an input to the plant which provides a measured output. Additionally, a disturbance is also supplied to the plant. The output of the plant is supplied to an observer, which estimates the states and the output to be controlled. The controlled output and states are fed back into the controller, closing the loop.

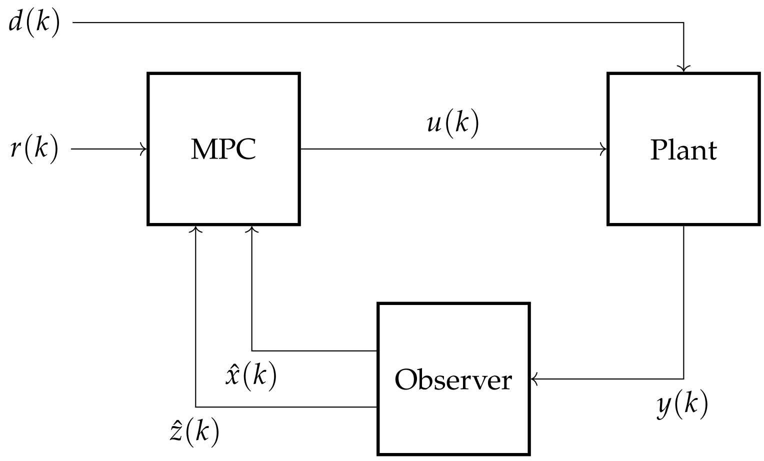

Figure 3.1: The system structure used in this chapter consisting of the MPC, the plant, the observer, and the signals in between.

As mentioned previously, the room temperature can, under normal circumstances, only be estimated why an observer is necessary. This estimation enables the incorporation of room temperature in the cost function within the MPC, and is in fact, for now, the only cost factor. The cost function J is implemented specifically as a quadratic cost function and penalises deviations from the reference signal. The quadratic cost is seen in Equation (3.3) [11].

$$J(k) = \sum_{i=H_w}^{H_p} \|\hat{z}(k+i) - r(k+i)\|^2_{Q(i)} + \sum_{i=0}^{H_u-1} \|\Delta\hat{u}(k+i)\|^2_{R(i)} \tag{3.3}$$

where:

H_w	is prediction window
$Q(i)$	is the tracking error weight
$R(i)$	is the input rate weight

The quadratic cost function predicts the tracking error from the prediction window H_w to the prediction horizon H_p and weighs the samples with the iteration dependent tracking error weight $Q(i)$. The cost function also weighs the changes in the input $\Delta\hat{u}$ with the corresponding weight $R(i)$. Remark that $Q(i)$ and $R(i)$ are expressed in quadratic form, that is:

$$\|x\|^2_Q \triangleq x^T Q x \tag{3.4}$$

The weights Q and R must both be positive semi-definite matrices, where Q has the dimensions of $H_p \times H_p$ and R has the dimensions of $H_u \times H_u$ [11]. Notice that the weights may be dependent on the sample, hence the index i. The cost function is split into the prediction term, which puts a cost on predicted tracking error and the control effort term, which puts a cost on change in input. As mentioned previously, the internal model is used to predict $\hat{z}$ while optimising $\Delta\hat{u}$. This internal model is predicted from the prediction window H_w to H_p steps forward in time. The optimising $\Delta\mathcal{U}(k)$ is H_u in length and is shorter than the prediction horizon, i.e. $H_u < H_p$. Beyond $H_u - 1$ the change in input is considered 0, that is $\Delta\hat{u}(k+i) = 0$ for $i \geq H_u$.

3.1.1 Posing the Quadratic Programming Problem

In mathematical optimisation theory, an optimisation problem is called a Quadratic Programming (QP) problem if it is on the form in Equations (3.5) [11].

$$\min \frac{1}{2} x^T H x + f^T x \tag{3.5a}$$

$$\text{subject to:} \quad Ax \leq b \tag{3.5b}$$

Algorithms and software for solving QP problems are readily available. In order to obtain a problem on the form in Equations (3.5), the cost function in Equation (3.3) is written as a sum of all terms in a matrix notation. This revision procedure is henceforth referred to as *lifting* and the matrices that result are referred to as *lifted* matrices. The procedure begins by recursively expressing the predicted state as a function of the change in input as seen in Equations (3.6) [11].

$$\hat{x}(k+1) = A\hat{x}(k) + Bu(k) \tag{3.6a}$$

$$\hat{x}(k+1) = A\hat{x}(k) + B\left(\Delta u(k) + u(k-1)\right) \tag{3.6b}$$

$$\hat{x}(k+2) = A\hat{x}(k+1) + B\left(\Delta u(k+1) + \Delta u(k) + u(k-1)\right) \tag{3.6c}$$

$$\hat{x}(k+2) = A\left[A\hat{x}(k) + B\left(\Delta u(k) + u(k-1)\right)\right] + B\left[\Delta u(k+1) + \Delta u(k) + u(k-1)\right] \tag{3.6d}$$

$$\hat{x}(k+2) = A^2\hat{x}(k) + (A+I)B\Delta\hat{u}(k) + B\Delta\hat{u}(k+1) + (A+I)Bu(k-1) \tag{3.6e}$$

$$\begin{aligned} \hat{x}(k+H_u) =& A^{H_u}\hat{x}(k) + (A^{H_u-1} + \cdots + A + I)B\Delta\hat{u}(k) + \cdots + B\Delta\hat{u}(k+H_u-1) + \\ & (A^{H_u-1} + \cdots + A + I)Bu(k-1) \end{aligned} \tag{3.6f}$$

Beyond the control horizon $H_u - 1$, the change in input is assumed zero and new input change terms stop appearing. That is at state $k+i$, all the propagations from previous input changes still remain and are still propagated, those being the terms with the $B\Delta\hat{u}(k+j-1)$ factors where j is less than i. However, new input change terms do not appear which they did previously on the form $B\Delta\hat{u}(k+j-1)$ where $j = i$. This is seen in Equations (3.7).

$$\hat{x}(k+H_u+1) = A\hat{x}(k+H_u) + Bu(k+H_u) = A\hat{x}(k+H_u) + 0 \tag{3.7a}$$

$$\begin{aligned} \hat{x}(k+H_u+1) =& A^{H_u+1}\hat{x}(k) + (A^{H_u} + \cdots + A + I)B\Delta\hat{u}(k) + \cdots + (A+I)B\Delta\hat{u}(k+H_u-1) + \\ & (A^{H_u} + \cdots + A + I)Bu(k-1) \end{aligned} \tag{3.7b}$$

$$\begin{aligned} \hat{x}(k+H_p) =& A^{H_p}\hat{x}(k) + (A^{H_p-1} + \cdots + A + I)B\Delta\hat{u}(k) + \cdots + (A+I)B\Delta\hat{u}(k+H_u-1) + \\ & (A^{H_p-1} + \cdots + A + I)Bu(k-1) \end{aligned} \tag{3.7c}$$

These predicted states can be collected and turned into a lifted representation of each predicted step as seen in Equation (3.8).

$$\mathcal{X} = \mathcal{A}\hat{x}(k) + \mathcal{B}_u u(k-1) + \mathcal{B}_{\Delta u}\Delta\mathcal{U}(k). \tag{3.8}$$

The matrices of the lifted representation in Equation (3.8) are shown below. The lifted $\mathcal{X}$ in Equation (3.9) consists of all the future states up to the prediction horizon. The dimensions of this matrix is $H_p m \times 1$ where m is the number of states.

$$\mathcal{X} = \begin{bmatrix} \hat{x}(k+1) \\ \hat{x}(k+2) \\ \vdots \\ \hat{x}(k+H_u) \\ \vdots \\ \hat{x}(k+H_p) \end{bmatrix} \tag{3.9}$$

Likewise $\mathcal{A}$, seen in Equation (3.10), also has the dimensions of $H_p m \times m$ as it consists of all the future A matrices.

$$\mathcal{A} = \begin{bmatrix} A \\ \vdots \\ A^{H_u} \\ \vdots \\ A^{H_p} \end{bmatrix} \tag{3.10}$$

The lifted $\mathcal{B}_u$ relates the previous input to the states and has dimensions $H_p m \times n$, where n is the number of inputs. It is seen in Equation (3.11).

$$\mathcal{B}_u = \begin{bmatrix} B \\ AB + B \\ \vdots \\ \sum_{i=0}^{H_u} A^i B \\ \vdots \\ \sum_{i=0}^{H_p-1} A^i B \end{bmatrix} \tag{3.11}$$

The lifted $\mathcal{B}_{\Delta u}$ relates the predicted future input changes to the states and has dimensions $H_p m \times H_u n$ and is seen in Equation (3.12).

$$\mathcal{B}_{\Delta u} = \begin{bmatrix} B & 0 & 0 & \cdots & 0 \\ AB + B & B & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & & \vdots \\ \sum_{i=0}^{H_u-1} A^i B & & \cdots & & B \\ \sum_{i=0}^{H_u} A^i B & & \cdots & & AB + B \\ \vdots & & \vdots & & \vdots \\ \sum_{i=0}^{H_p-1} A^i B & & \cdots & & \sum_{i=0}^{H_p-H_u} A^i B \end{bmatrix} \tag{3.12}$$

The predicted output for each step can be lifted in the same fashion as seen in Equations (3.13). The direct term D_z is neglected for now, as it is zero in the model.

$$\hat{z}(k+1) = C_z\hat{x}(k+1) \tag{3.13a}$$
$$\hat{z}(k+2) = C_z\hat{x}(k+2) \tag{3.13b}$$
$$\vdots$$
$$\hat{z}(k+H_p) = C_z\hat{x}(k+H_p) \tag{3.13c}$$

where:

C_z is the output matrix for the output to be controlled

This results in the lifted matrices of Equations (3.14).

$$\mathcal{Z}(k) = \mathcal{C}\mathcal{X} = \mathcal{C}\left(\mathcal{A}\hat{x}(k) + \mathcal{B}_u u(k-1) + \mathcal{B}_{\Delta u}\Delta\mathcal{U}(k)\right) \tag{3.14a}$$
$$= \Psi\hat{x}(k) + \Upsilon u(k-1) + \Theta\Delta\mathcal{U}(k) \tag{3.14b}$$

where $\mathcal{C}$ is a diagonal matrix of size $H_p p \times H_p m$ with C_z as block entries and where p is the number of controlled outputs. The matrices in Equations (3.14) are not shown as they are trivial to construct.

Now, the tracking error is defined in Equation (3.15) as the difference between predicted system response and free response of the system if no changes had been made, i.e. $\Delta\mathcal{U}(k) = 0$ [11].

$$\mathcal{E}(k) = \mathcal{T}(k) - \Psi\hat{x}(k) - \Upsilon U(k-1) \tag{3.15}$$

where:

$\mathcal{E}(k)$ is the tracking error of the free response at time k
$\mathcal{T}(k)$ is the vector of reference signals from k to $k+H_p$

The cost from Equation (3.3) can now be reformulated as a QP problem expressed in terms of $\Delta\mathcal{U}(k)$ using the tracking error as seen in Equations (3.16) [11].

$$J(k) = \|\mathcal{Z}(k) - \mathcal{T}(k)\|^2_{\mathcal{Q}} + \|\Delta\mathcal{U}(k)\|^2_{\mathcal{R}} \tag{3.16a}$$
$$= \|\Theta\Delta\mathcal{U}(k) + \Psi\hat{x}(k) + \Upsilon u(k-1) - \mathcal{T}(k)\|^2_{\mathcal{Q}} + \|\Delta\mathcal{U}(k)\|^2_{\mathcal{R}} \tag{3.16b}$$
$$= \|\Theta\Delta\mathcal{U}(k) - \mathcal{E}(k)\|^2_{\mathcal{Q}} + \|\Delta\mathcal{U}(k)\|^2_{\mathcal{R}} \tag{3.16c}$$
$$= \left[\Delta\mathcal{U}(k)^T\Theta^T - \mathcal{E}(k)^T\right]\mathcal{Q}\left[\Theta\Delta\mathcal{U}(k) - \mathcal{E}(k)\right] + \Delta\mathcal{U}(k)^T\mathcal{R}\Delta\mathcal{U}(k) \tag{3.16d}$$
$$= \mathcal{E}(k)^T\mathcal{Q}\mathcal{E}(k) - 2\Delta\mathcal{U}(k)^T\Theta^T\mathcal{Q}\mathcal{E}(k) + \Delta\mathcal{U}(k)^T\left[\Theta^T\mathcal{Q}\Theta + \mathcal{R}\right]\Delta\mathcal{U}(k) \tag{3.16e}$$

where:

$\mathcal{Q}$ is a diagonal matrix whose i'th entry is the tracking error weight $Q(i)$
$\mathcal{R}$ is a diagonal matrix whose i'th entry is input change weight $R(i)$

The terms are collected into a constant term ι a linear term $\mathcal{G}$ and a quadratic term $\mathcal{H}$ as seen in Equations (3.17).

$$J = \iota - \Delta\mathcal{U}(k)^T\mathcal{G} + \Delta\mathcal{U}(k)^T\mathcal{H}\Delta\mathcal{U}(k) \tag{3.17a}$$
$$\iota = \mathcal{E}(k)^T\mathcal{Q}\mathcal{E}(k) \tag{3.17b}$$
$$\mathcal{G} = 2\Theta^T\mathcal{Q}\mathcal{E}(k) \tag{3.17c}$$
$$\mathcal{H} = \Theta^T\mathcal{Q}\Theta + \mathcal{R} \tag{3.17d}$$

The cost J is now clearly on the QP form as specified in Equations (3.5), and the constant term ι can be disregarded as it does not affect the optimising $\Delta\mathcal{U}$.

3.1.2 Incorporating Constraints

As the central heating system is subject to the laws of physics, constraints can be a useful and necessary for the MPC to act within the bounds of reality. Since it is chosen to formulate the MPC as a QP problem the constraints must be on the form of linear inequalities as seen in Equations (3.5). The linear inequality is split into three parts, those being input rate constraints, input constraints and output constraints [11]. Each constraint category is modelled mathematically in Equations (3.18) as a linear combination of either input rates, inputs, or outputs being less than or equal to zero.

$$E\begin{bmatrix}\Delta\mathcal{U}(k)\\1\end{bmatrix} \leq 0 \tag{3.18a}$$
$$F\begin{bmatrix}\mathcal{U}(k)\\1\end{bmatrix} \leq 0 \tag{3.18b}$$
$$G\begin{bmatrix}\mathcal{Z}(k)\\1\end{bmatrix} \leq 0 \tag{3.18c}$$

where:

E is the input rate constraint matrix
F is the input constraint matrix
G is the output constraint matrix

To include these constraints in the QP problem, they must be dependent only on the same variable as the cost function, that is $\Delta\mathcal{U}(k)$. The constraint E is already expressed in terms of $\Delta\mathcal{U}(k)$ so no reformulation is required. However, the other two constraints must be lifted in order to express them in terms of $\Delta\mathcal{U}(k)$ [11]. Starting with the input constraint F, it is split into columns pertaining to each input and lastly the constant:

$$F = \begin{bmatrix}F_1 & F_2 & \dots & F_{H_u} & f\end{bmatrix} \tag{3.19}$$

The constraint is rewritten as a sum and $u(k+i)$ is written as the sum of changes $\Delta u(k)$ up until that point:

$$\sum_{i=1}^{H_u} F_i \hat{u}(k+i-1) + f \leq 0 \tag{3.20a}$$

$$\sum_{i=1}^{H_u} F_i \left[\hat{u}(k-1) + \sum_{j=0}^{i-1} \Delta \hat{u}(k+j) \right] + f \leq 0 \tag{3.20b}$$

$$\sum_{i=1}^{H_u} F_i \left[\hat{u}(k-1) + \sum_{j=1}^{i} \Delta \hat{u}(k+j-1) \right] + f \leq 0 \tag{3.20c}$$

$$\sum_{i=1}^{H_u} F_i \hat{u}(k-1) + \sum_{i=1}^{H_u} F_i \sum_{j=1}^{i} \Delta \hat{u}(k+j-1) + f \leq 0 \tag{3.20d}$$

The summation identity in Equation (3.21) is used to reformulate Equations (3.20) as seen in Equations (3.22).

$$\sum_{i=k}^{n} \sum_{j=k}^{i} a_{i,j} = \sum_{j=k}^{n} \sum_{i=j}^{n} a_{i,j} \tag{3.21}$$

$$\sum_{i=1}^{H_u} F_i u(k-1) + \sum_{j=1}^{H_u} \sum_{i=j}^{H_u} F_i \Delta \hat{u}(k+j-1) + f \leq 0 \tag{3.22a}$$

$$\sum_{i=1}^{H_u} F_i u(k-1) + \sum_{i=1}^{H_u} F_i \Delta \hat{u}(k) + \sum_{i=2}^{H_u} F_i \Delta \hat{u}(k+1) + \cdots + \sum_{i=H_u}^{H_u} F_i \Delta \hat{u}(k+H_u-1) + f \leq 0 \tag{3.22b}$$

$$\mathbf{F_1} u(k-1) + \mathbf{F_1} \Delta \hat{u}(k) + \mathbf{F_2} \Delta \hat{u}(k+1) + \cdots + \mathbf{F_{H_u}} \Delta \hat{u}(k+H_u-1) + f \leq 0 \tag{3.22c}$$

This can be collected into Equations (3.23) using the input change vector $\Delta\mathcal{U}(k)$.

$$\mathbf{F} \Delta \mathcal{U} \leq -\mathbf{F_1} u(k-1) - f \tag{3.23a}$$

$$\mathbf{F} = \begin{bmatrix} \mathbf{F_1} & \mathbf{F_2} & \cdots & \mathbf{F_{H_u}} \end{bmatrix} \tag{3.23b}$$

$$\mathbf{F_j} = \sum_{i=j}^{H_u} F_i \tag{3.23c}$$

Likewise, the input rate constraint and output constraint matrices are lifted, albeit significantly easier, as seen in Equation (3.24) and Equations (3.25) respectively.

$$E \begin{bmatrix} \Delta \mathcal{U}(k) \\ 1 \end{bmatrix} \leq 0 \quad \Leftrightarrow \quad \begin{bmatrix} W & -w \end{bmatrix} \begin{bmatrix} \Delta \mathcal{U}(k) \\ 1 \end{bmatrix} \leq 0 \quad \Leftrightarrow \quad W \Delta \mathcal{U}(k) \leq w \tag{3.24}$$

$$G \begin{bmatrix} \mathcal{Z}(k) \\ 1 \end{bmatrix} \leq 0 \quad \Leftrightarrow \quad \begin{bmatrix} \Gamma & \gamma \end{bmatrix} \begin{bmatrix} \mathcal{Z}(k) \\ 1 \end{bmatrix} \leq 0 \quad \Leftrightarrow \quad \Gamma \mathcal{Z}(k) \leq -\gamma \tag{3.25a}$$

$$\Leftrightarrow \quad \Gamma \left(\Psi \hat{x}(k) + \Upsilon u(k-1) + \Theta \Delta \mathcal{U}(k) \right) \leq -\gamma \tag{3.25b}$$

$$\Leftrightarrow \quad \Gamma \Theta \Delta \mathcal{U}(k) \leq -\Gamma \left[\Psi \hat{x}(k) + \Upsilon u(k-1) \right] - \gamma \tag{3.25c}$$

Collecting Equations (3.23), Equation (3.24) and Equations (3.25) into one linear inequality gives the result in Equation (3.26).

$$\begin{bmatrix} \mathbf{F} \\ \Gamma\Theta \\ W \end{bmatrix} \Delta\mathcal{U}(k) \leq \begin{bmatrix} -\mathbf{F}_1 u(k-1) - f \\ -\Gamma[\Psi x(k) + \Upsilon u(k-1)] - \gamma \\ w \end{bmatrix} \tag{3.26}$$

Now the QP-problem has been posed where the quadratic function in Equations (3.17) is minimised subject to the constraints Equation (3.26). The full QP-problem is restated in Equations (3.27) for clarity.

$$\min_{\Delta U(k)} J(k) = -\Delta\mathcal{U}(k)^T \mathcal{G} + \Delta\mathcal{U}(k)^T \mathcal{H} \Delta\mathcal{U}(k) \tag{3.27a}$$

subject to:

$$\begin{bmatrix} \mathbf{F} \\ \Gamma\Theta \\ W \end{bmatrix} \Delta\mathcal{U}(k) \leq \begin{bmatrix} -\mathbf{F}_1 u(k-1) - f \\ -\Gamma[\Psi x(k) + \Upsilon u(k-1)] - \gamma \\ w \end{bmatrix} \tag{3.27b}$$

Equations (3.27) is clearly on the quadratic form stated previously in Equations (3.5) and can, as such, be solved by available algorithms and tools.

3.2 Controller Structure

The general structure in Figure 3.1 is implemented in practice. As mentioned, the system uses a simple Luenberger observer to estimate the states $\hat{x}(k)$, which in turn are used to produce the estimated output to be controlled $\hat{z}(k)$. For clarity, the system equations are presented below:

$$x(k+1) = Ax(k) + Bu(k) + B_d d(k) \tag{3.28}$$

$$y(k) = C_y x(k) + D_y u(k) + D_{dy} d(k) \tag{3.29}$$

The Luenberger observer is then characterised by the following [11]:

$$\hat{y}(k) = \hat{C}_y \hat{x}(k) + \hat{D}_y u(k) + \hat{D}_{dy} d(k) \tag{3.30}$$

$$\hat{z}(k) = \hat{C}_z \hat{x}(k) + \hat{D}_z u(k) + \hat{D}_{dz} d(k) \tag{3.31}$$

$$\hat{x}(k+1) = \hat{A}\hat{x}(k) + L(y(k) - \hat{y}(k)) + \hat{B}u(k) + \hat{B}_d d(k) \tag{3.32}$$

where:

$\hat{C}_y, \hat{D}_y, \hat{D}_{dy}, \hat{A}, \hat{B}, \hat{B}_d$ are the modelled state space matrices.

In practice, $\hat{z}(k)$ is not calculated as it is implicitly predicted in the MPC so long as the lifted matrices are constructed using the matrices pertaining to $\hat{z}(k)$, i.e. $\hat{C}_z, \hat{D}_z, \hat{D}_{dz}$. The observer feedback L is chosen such that estimation error is stable and converges to 0. This is the case if the system is observable and the eigenvalues of $A - LC_y$ are within the unit circle [11]. The eigenvalues are placed using the MATLAB `place()` command in a small region around 0.

The model of the provided house has one input and one disturbance, those being the supply temperature and ambient temperature respectively. Remark that the direct terms for both input and disturbance are 0 in the fitted model (see Equations (2.9)) and are therefore disregarded going forward. The ambient temperature is not a controllable input but rather a measured disturbance which has to be accounted for in the controller. A method to handle disturbances is not present in the MPC structure formulated in Section 3.1 and must be added. This is done by lifting the matrices pertaining to the disturbance $d(k)$, that being B_d, in the same fashion as in Section 3.1. Since the disturbance matrix does not depend on $\Delta\mathcal{U}(k)$ it is lifted and premultiplied by $\mathcal{C}$ as well as added to the lifted output in Equations (3.14) as seen in Equations (3.33) [11].

$$\mathcal{Z}(k) = \Psi\hat{x}(k) + \Upsilon u(k-1) + \Theta\Delta\mathcal{U}(k) + \Xi\mathcal{D}(k) \tag{3.33a}$$

$$\mathcal{D} = \begin{bmatrix} d(k) \\ \hat{d}(k+1) \\ \vdots \\ \hat{d}(k+H_p-1) \end{bmatrix} \tag{3.33b}$$

$$\Xi = \begin{bmatrix} C_zB_d & 0 & \cdots & 0 \\ C_zAB_d & C_zB_d & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ C_zA^{H_p-1}B_d & C_zA^{H_p-2}B_d & \cdots & C_zB_d \end{bmatrix} \tag{3.33c}$$

where:

$\hat{d}(k)$	is the predicted disturbance at time k
$\mathcal{D}$	is the vector of disturbances from k until $k+H_p-1$
Ξ	is the lifted disturbance matrix

This changes the tracking error $\mathcal{E}(k)$ but otherwise, the cost function in Equations (3.16) remains the same:

$$\mathcal{E}(k) = \mathcal{T} - \Psi\hat{x}(k) - \Upsilon u(k-1) - \Xi\mathcal{D}(k) \tag{3.34}$$

The output matrix for the output to be controlled C_z is chosen such that the room temperature is the only measurement. In accordance with the MPC theory from Section 3.1, the cost is then purely based on how far the predicted room temperature is from the reference. As the room temperature T_r is available as the last state in the state vector, the output matrix simply becomes:

$$C_z = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \end{bmatrix} \tag{3.35}$$

The constraints of the system are set such that the system behaves realistically wrt. the repsonse of the pre-existing heat pump controller. The input constraints on the heat pump controller reference, i.e. the supply water temperature, is chosen heuristically as:

1. The supplied FH water cannot be cooler than $\mu_1 = 1\,°C$
2. The supplied FH water cannot be warmer than $\mu_2 = 45\,°C$
3. The supplied FH water cannot be cooled, and as such, cannot generally be colder than at least the ambient temperature. It is chosen to be at least $\mu_3 = 2\,°C$ warmer than the ambient temperature.

4. The heating system cannot change the supplied water temperature more than $\mu_4 = 0.1\,°C$ per minute.

These constraints are sketched in Figure 3.2 which also shows the set of infeasible inputs due to these constraints.

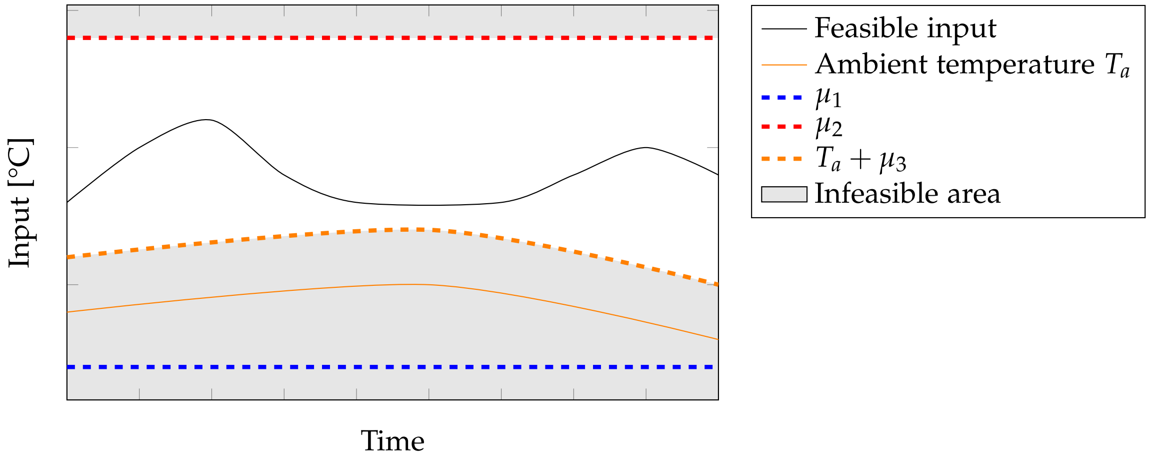

Figure 3.2: Sketch of input constraints and the infeasible set of inputs they produce. The actual values of μ_1, μ_2, μ_3 are exaggerated for demonstrational purposes.

These constraints can be expressed in terms of the predicted inputs $\mathcal{U}(k)$ and the change in input $\Delta\mathcal{U}(k)$ as seen in Equations (3.36).

$$-\mathcal{U}(k) + \mu_1 \leq 0 \tag{3.36a}$$
$$\mathcal{U}(k) - \mu_2 \leq 0 \tag{3.36b}$$
$$-\mathcal{U}(k) + \mu_3 + \mathcal{D}(k) \leq 0 \tag{3.36c}$$
$$\Delta\mathcal{U}(k) - \mu_4\kappa \leq 0 \tag{3.36d}$$
$$-\Delta\mathcal{U}(k) - \mu_4\kappa \leq 0 \tag{3.36e}$$

where:

κ is the scale factor between system sample time and controller sample time (see Section 3.2.1).

The downsampling factor κ is required in the constraint, in order to scale the chosen $0.1\,°C$ per minute to match the sample time of controller. These constraints are put on the form in Equations (3.18) as seen in Equation (3.37).

$$E \begin{bmatrix} \Delta\mathcal{U}(k) \\ 1 \end{bmatrix} \leq 0 \tag{3.37}$$

$$F \begin{bmatrix} \mathcal{U}(k) \\ 1 \end{bmatrix} \leq 0 \tag{3.38}$$

$$E = \begin{bmatrix} +1 & 0 & 0 & \cdots & 0 & -\mu_4\kappa \\ -1 & 0 & 0 & \cdots & 0 & -\mu_4\kappa \\ 0 & +1 & 0 & \cdots & 0 & -\mu_4\kappa \\ 0 & -1 & 0 & \cdots & 0 & -\mu_4\kappa \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & +1 & -\mu_4\kappa \\ 0 & 0 & 0 & \cdots & -1 & -\mu_4\kappa \end{bmatrix} \tag{3.39}$$

$$F = \begin{bmatrix} -1 & 0 & 0 & \cdots & 0 & \mu_1 \\ +1 & 0 & 0 & \cdots & 0 & -\mu_2 \\ -1 & 0 & 0 & \cdots & 0 & \mu_3 + d(1) \\ 0 & -1 & 0 & \cdots & 0 & \mu_1 \\ 0 & +1 & 0 & \cdots & 0 & -\mu_2 \\ 0 & -1 & 0 & \cdots & 0 & \mu_3 + d(2) \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -1 & \mu_1 \\ 0 & 0 & 0 & \cdots & +1 & -\mu_2 \\ 0 & 0 & 0 & \cdots & -1 & \mu_3 + d(H_u) \end{bmatrix} \tag{3.40}$$

These constraints are lifted as detailed in Equation (3.26) such that they are dependent only on $\Delta\mathcal{U}(k)$. This matrix is quite extensive and not of particular interest and is thus not shown.

3.2.1 Downsampling of the Controller

Depending on the prediction horizon and control horizon the lifted matrix can potentially become huge, leading to slow computation times. This is especially the case if the controller is expected to react to changes in weather, as this changes very slowly compared to the designed system sample time of $\tau_s = 60\,\text{s}$. While it is possible to downsample the entire system and the controller, this is not necessary as the observer and control update can easily be run within τ_s. Therefore the control loop is changed, such that the MPC control loop is only run every $\kappa \cdot \tau_s$ where κ is some positive integer. This leads to a few considerations:

1. The model in the MPC, and only the MPC, must be downsampled to $\kappa\tau_s$.
2. The input is only available every $\kappa\tau_s$ but is used every τ_s, thus requiring interpolation and a scheme for selecting $u(k)$.

The first consideration is easily handled by converting the system to continuous time and downsampling it with a sample time of $\kappa\tau_s$. This is done by the zero-order hold method both ways using the MATLAB® command `d2d()`. The second consideration is addressed by interpolating the optimising input sequence produced by the MPC. At time k, if κ is a divisor of k, the MPC produces inputs for $k, k+\kappa, \ldots, k+\kappa H_u$ which is interpolated to $k, k+1, \ldots, k+\kappa H_u$. In practice, this input is generated such that a vector is indexed by $k - k_{mpc}$ where k_{mpc} is the k at which the MPC output was generated. This results in the first index of the vector corresponding to k being accessed first, followed by the second index corresponding to $k+1$ and so forth. This is illustrated in Figure 3.3. This interpolation scheme ensures that the system always has access to the correct input at time k.

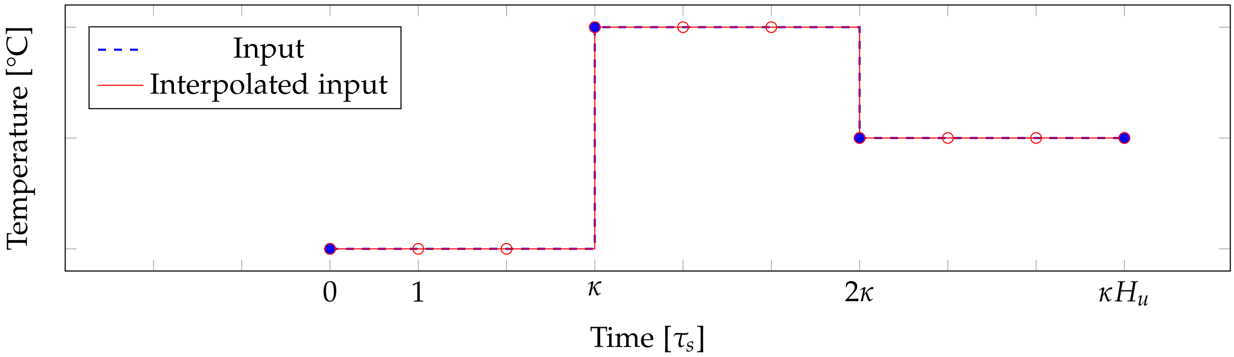

Figure 3.3: Illustration of the interpolation of the input produced by the MPC.

In this implementation, only the first κ inputs will be used, as the MPC loop is run every κ samples, producing new inputs from that point on. However, using this implementation, it is in fact possible to run the controller at any positive integer sample interval $\bar{\kappa}$ which is smaller than or equal to κH_u thus utilising more of the optimising input before calculating a new one. The plant input must be calculated within one sample of length τ_s, and as such, κ and $\bar{\kappa}$ have no effect on the computation time during online operation. However, a smaller κ and $\bar{\kappa}$ result in the controller being run more often, significantly increasing the simulation time. A balanced approach is taken in which $\bar{\kappa} = \kappa = 30$ as this provides decent results with a moderate simulation time.

An algorithmic description is constructed and shown in Algorithm 1. The pseudocode presented in this algorithm defines the initial variables, and repeats the control loop of acquiring a measurement, updating the observer, and choosing an input. Evidently, the plant input vector is only updated every κ sample, upon which the lifted system is obtained. As is it is not necessary to construct these lifted matrices more than once, they are stored and re-used. The output is chosen by subtracting the running variable k by the value of k when the input was generated, i.e. k_{mpc} and using that as an index. For an overview of the routines `LIFT_SYSTEM()` and `LIFT_CONSTRAINTS()` refer to Appendix A.

Thus the controller described in this section is capable of being simulated, such that its input, states and corresponding output may be examined. While such a simulation allows for the comparison against another controllers via their responses, it provides only one metric for success. Therefore, the other metric for comparison is now constructed in the form of an estimated cost of operation of the central heating system.

Algorithm 1 Implementation of the algorithm for comfort control presented in this chapter. Refer to Appendix A for an overview of the non-trivial internal functions used in this algorithm.

Require: $1 \leq H_p, \quad 1 \leq H_u \leq H_p$
Require: $\kappa \in \mathbb{Z}, \quad 1 \leq \kappa$ ▷ κ must be a positive integer
Require: $\mathcal{D} \in \mathbb{R}^{\kappa H_p}, \quad \mathcal{T} \in \mathbb{R}^{\kappa H_p}$
$k = 0$
$k_{mpc} = 1$
$\hat{x} = \hat{x}_0 \in \mathbb{R}^m$
$u = u_0 \in \mathbb{R}^1$
repeat
  $k = k + 1$
  $y =$ GET_MEASUREMENT ▷ Measurement is performed
  if $(k-1) \bmod \kappa = 0$ **then** ▷ Controller loop is downsampled by κ
    $k_{mpc} = k$
    $(\Psi, \Upsilon, \Theta, \Xi) =$ LIFT_SYSTEM(H_p, H_u)
    $\mathcal{H} = \Theta^T \mathcal{Q} \Theta + \mathcal{R}$
    $\mathcal{E} = \mathcal{T} - \Psi x - \Upsilon u - \Xi d$
    $\mathcal{G} = 2\Theta^T \mathcal{Q} \mathcal{E}$
    $(A_{con}, b_{con}) =$ LIFT_CONSTRAINTS$(H_p, H_u, u, \mathcal{D}, \mathcal{T}, \Upsilon, \Theta, \Xi)$
    $\Delta\mathcal{U} =$ QUADPROG$(2\mathcal{H}, -G^T, A_{con}, b_{con})$ ▷ Constrained QP problem solved
    $\mathcal{U} = \begin{bmatrix} u + \Delta\mathcal{U}(1) & u + \Delta\mathcal{U}(2) & \dots & u + \Delta\mathcal{U}(\kappa H_u) \end{bmatrix}$
  end if
  $\hat{y} = C_y \hat{x} + D_y u + D_{dy} d$ ▷ Estimates are produced
  $\hat{x} = A\hat{x} + L(y - \hat{y}) + Bu + B_d d$
  $u = \mathcal{U}(k - k_{mpc} + 1)$
until $k > k_{final}$

3.3 Cost of Operation

As the goal of the project is to determine the possibility of constructing a cost-optimising controller, an important metric in determining success of a controller is the price of operation. The cost of operation can, in practice, be calculated as the power usage, which is measured, multiplied by the electricity price at that time, which is given by electricity suppliers and known beforehand in some horizon. The electricity price varies as the demand and supply on the electricity grid changes and is therefore considered as an outside disturbance to the system.

Problems arise when calculating power usage in simulation since this variable is not part of the state space, and is not directly derived therefrom. It is hypothesised that the power usage is close to proportional to the difference between supply and return temperature in the FH system. This is because the assumption that the primary energy consumption of the central heating system happens as the return water temperature is raised, by the central heating system, to that of the supply water temperature. The difference in supply and return temperature $\Delta T_{r,s}$ is defined as follows:

$$\Delta T_{r,s} = T_{w,s} - T_{w,r} = T_{in} - T_{w,r} \tag{3.41}$$

This can be incorporated in the state space by new rows in the C_z and D_z matrices respectively. This is not necessary for now, as the cost is computed offline after simulation by means of the saved states and inputs. The variable $\Delta T_{r,s}$ is constructed from a real set of data using measured supply and return temperature. This is then approximately analogous to the input T_{in} and the state $T_{w,r}$ they are the supply and return temperatures to which the model is fitted. Additional power consumers are assumed negligible or constant and can thus be disregarded in this cost.

The temperature difference $\Delta T_{r,s}$ is plotted against the measured power usage P_c, and a linear relationship is found as seen in Figure 3.4. The points in Figure 3.4 have a clear tendency why some points are defined outliers, while the rest are denoted inlier points. The inlier points lie within the inlier zone, which is determined heuristically. These points generally correspond to the HWP periods as seen in Figure 3.5 which shows points considered outliers and points considered inliers by this method. During fitting, the outlier points are removed and the fit is performed on the remaining points.

A fit is performed on the inlier points to the relationship in Equation (3.42).

$$P_c(k) = \alpha \Delta T_{r,s}(k) \tag{3.42}$$

where:

P_c	is the total power consumption of the central heating system	[kW]
α	is a proportionality constant	[kW/°C]

The relationship is fitted using least squares with the result being $\alpha = 326.6\,\text{kW}/°\text{C}$.

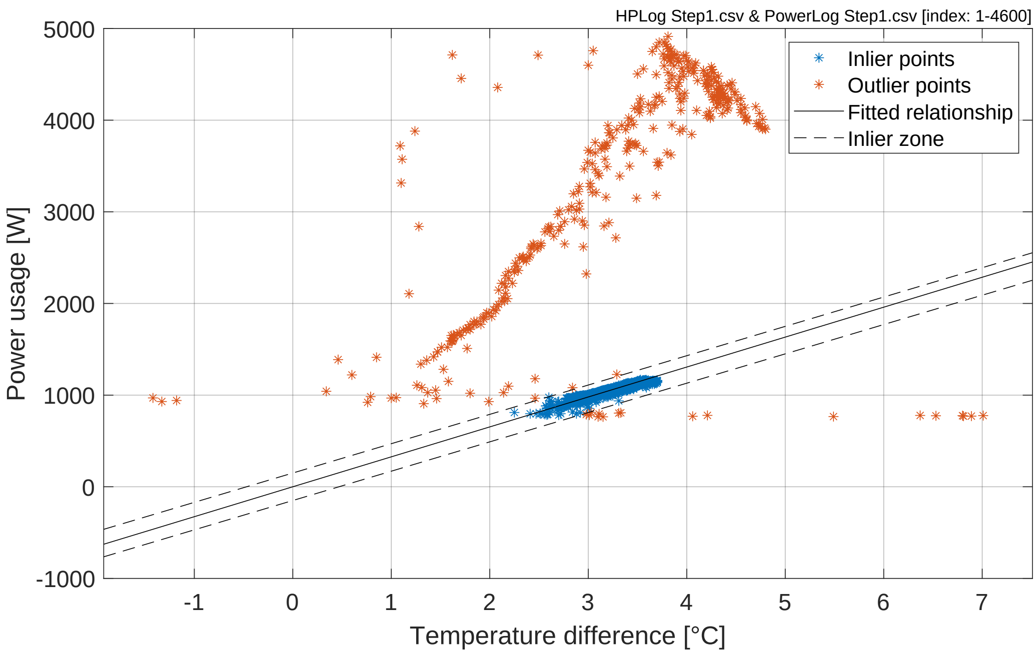

Figure 3.4: Measured power consumption as a function of measured supply and return water temperature difference with outlier classification and fit onto inlier points.

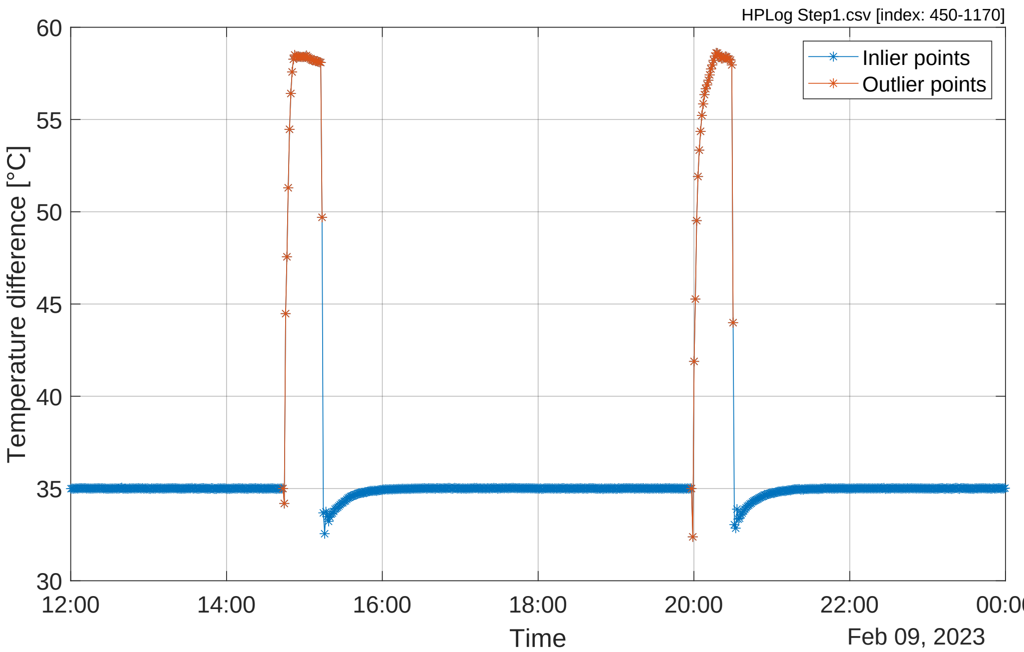

Figure 3.5: Supply water temperature as a function of time. Outlier and inlier points correspond to the points in Figure 3.4.

To obtain the actual cost of operation the electricity price is multiplied onto this power consumption as seen in Equation (3.43).

$$V(k) = P_c(k) \cdot \frac{\tau_s}{60^2} \cdot m_e(k) \tag{3.43}$$

where:

$V(k)$	is the cost of operation for a single step of length τ_s at time k	[DKK]
$m_e(k)$	is the cost of electricity at time k	[DKK/(kWh)]

The factor 60^{-2} exists to align the units since $\tau_s = 60\,\text{s}$. The total cost of operation across some time span from 0 to k is simply the sum, denoted the cumulative cost of operations $\bar{V}(k)$:

$$\bar{V}(k) = \sum_{i=0}^{k} V(k) \tag{3.44}$$

The cost of electricity is, as mentioned previously, retrieved from the supplier. This retrieved price is dynamic in nature, but also has dynamics tariffs which makes the price especially volatile. The electricity supplier of the provided residence is the company N1, which has the following additional relevant tariffs defined [12]:

- Periods from October (01-10) to March (31-03) between the hours 17:00-20:00 inclusive are denoted peak consumption periods. These periods have an additional tariff of 0.8409 DKK/(kWh).
- Any period which is not a peak consumption period is denoted a low consumption period. These periods have an additional tariff of 0.3286 DKK/(kWh).

The price periods are retrieved from Energi Data Service which provides electricity prices from any period ranging back to 2004 as well as forecast electricity prices. Figure 3.6 shows a period of time with and without peak consumption periods before and after the addition of these periodic tariffs.

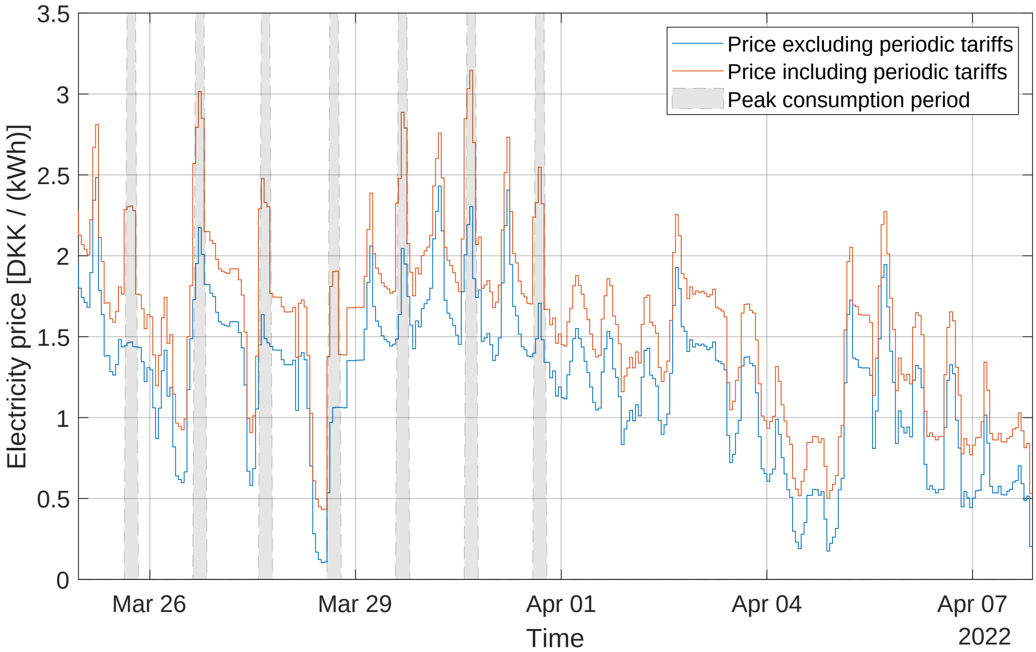

Figure 3.6: Electricity prices with and without periodic tariffs from the period 25-03-2022 to 08-04-2022. The grey zone indicates a peak consumption period where the tariffs are increased.

This cost including tariffs does not necessarily cover subscription costs, additional transport costs etc. which originate from the specific electricity provider of the provided residence. However, these additional costs are assumed either negligible or constant, in which case they play no role in the optimising input. This is also the reason why the proportionality constant in Equation (3.42) need not be precise as a change in the proportionality does not affect the final optimising input. Both unmodelled costs and imprecise power calculations do affect the final cumulative cost why this metric is not necessarily accurate in reality but merely used as a comparison criteria in simulation.

3.4 Comfort Controller Simulation Results

In this section, the simulated response of the comfort controller and costs of operation are presented. A year is chosen as the length of the simulation, as this allows for investigation of the controllers response to seasonal changes. The period 1st of January 2022 to the 1st of January 2023 is chosen as the simulation period. The system is simulated by the model equations in Section 2.3, observed and controlled using the comfort MPC designed in this chapter. The weather is obtained from the Open-Meteo weather service as explained in Section 2.1.1 while the electricity data including tariffs is retrieved from Energi Data Service as explained in Section 3.3. For a full list of parameters used for this simulation refer to Appendix C. Figure 3.7 shows the response of the output to be controlled, i.e. the room temperature alongside its reference, the input and ambient temperature.

As can be seen in Figure 3.7, the room temperature stays close to the reference except for a region between June and September. This is likely due to the ambient temperature raising the room temperature above the reference. As the input is constrained to never go below ambient temperature, it is not possible to cool the room, so large ambient temperatures are detrimental to the performance. Handling of such a scenario is outside the central heating system and FH circuit, and is therefore outside the scope of this project. Nevertheless, this will affect the cost of operation due to the way it is modelled. Remark the clear counter action of the input temperature to the ambient temperature in the cooler parts of the year, in which the input roughly mirrors the ambient temperature.

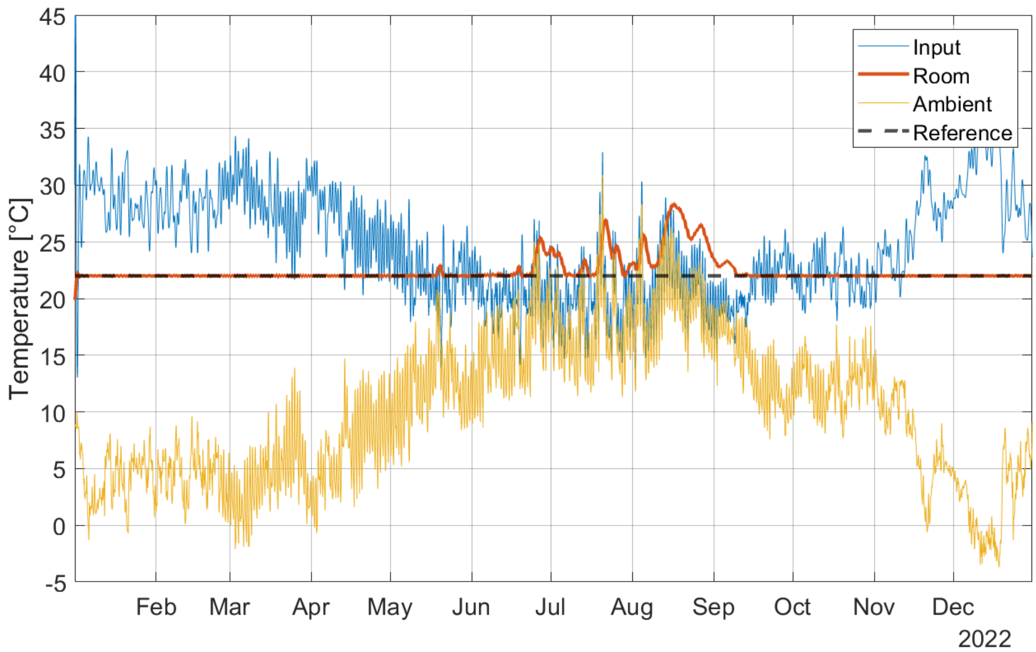

Figure 3.7: Simulation of the comfort controller input-output response to a reference of 22 °C and ambient temperature data from 1st of January 2022 to 1st of January 2023. The ambient temperature data is retrieved from the Open-Meteo weather service and used as disturbance in the simulated model.

A zoomed in view of the beginning of the input-output response, i.e. the room temperature, can be seen in Figure 3.8. The initial room temperature state is chosen as 20 °C and the reference temperature of 22 °C is reached within a day. However, there is an overshoot present which peaks at 22.3 °C. This is assumed to be negligible for human comfort. Be that as it may, the controller does reach a steady state at the reference with only minor oscillations. A region is denoted a steady state region if the input does not change significantly over time despite the disturbances the system faces in the form of ambient temperature changes.

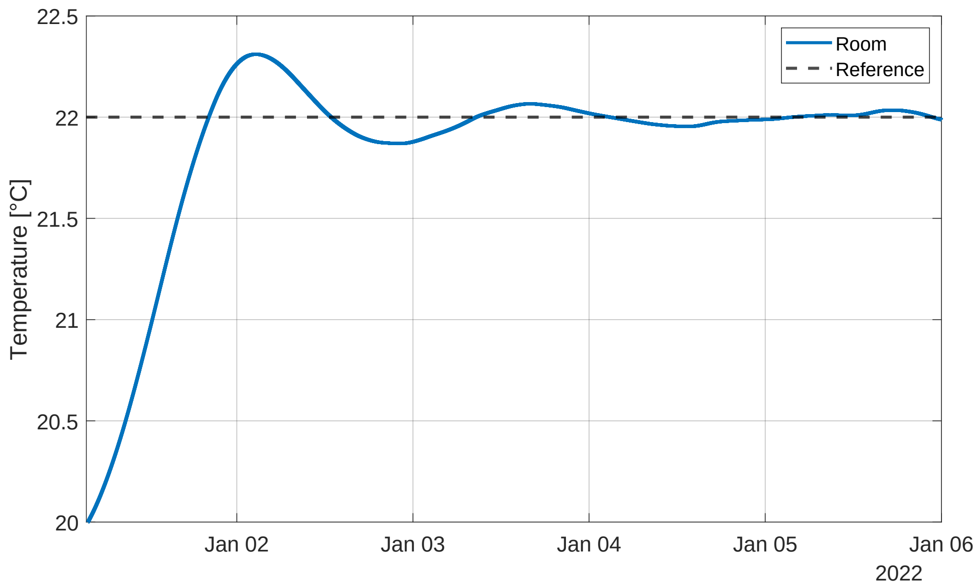

Figure 3.8: Initial part of the simulated output response in Figure 3.7.

Examining the steady state behaviour, seen in Figure 3.9, it becomes clear that minor temperature changes in the room temperature do occur. However, these changes are generally less than 0.1 °C. It is assumed that the residents cannot discern between the reference temperature and one that is 0.1 °C colder or warmer, and as such it is deemed acceptable.

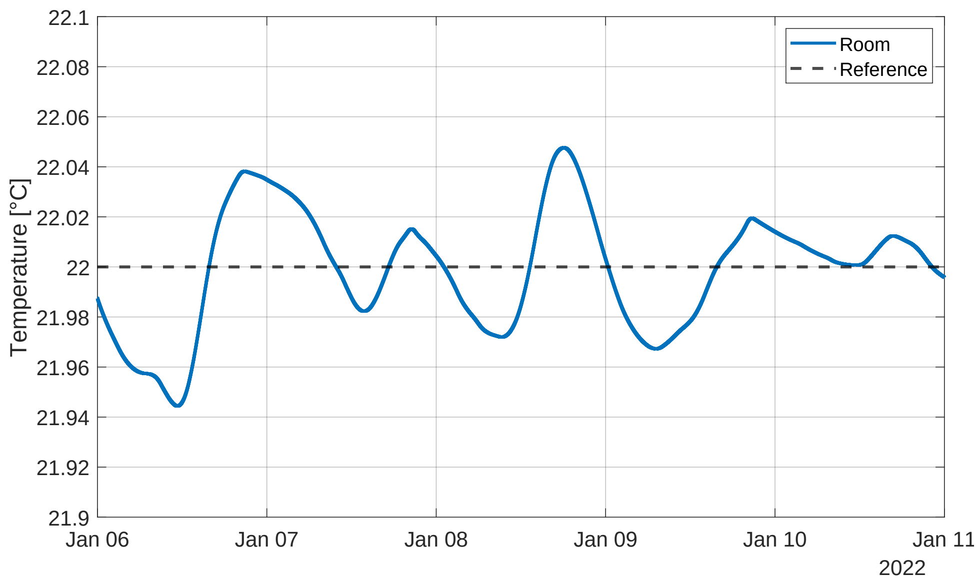

Figure 3.9: Zoomed in view of the simulated room temperature in Figure 3.7 at a chosen steady state region.

Another point of interest is the behaviour of the input temperature in the very beginning of the response, which can be seen in Figure 3.10. It can be seen that the input temperature is correctly limited by the controllers constraints to a maximum heating of 45 °C. Another observation is, that the controller decides to reduce the input temperature as the room temperature approaches the reference to counteract the potential overshoot. In this case, the reduction is possible since the controller sees a relatively low ambient temperature. In summer, or other warm periods, the ambient temperature may not allow the controller to reduce the input temperature in anticipation of reaching the reference which could result in higher overshoot than present on Figure 3.8.

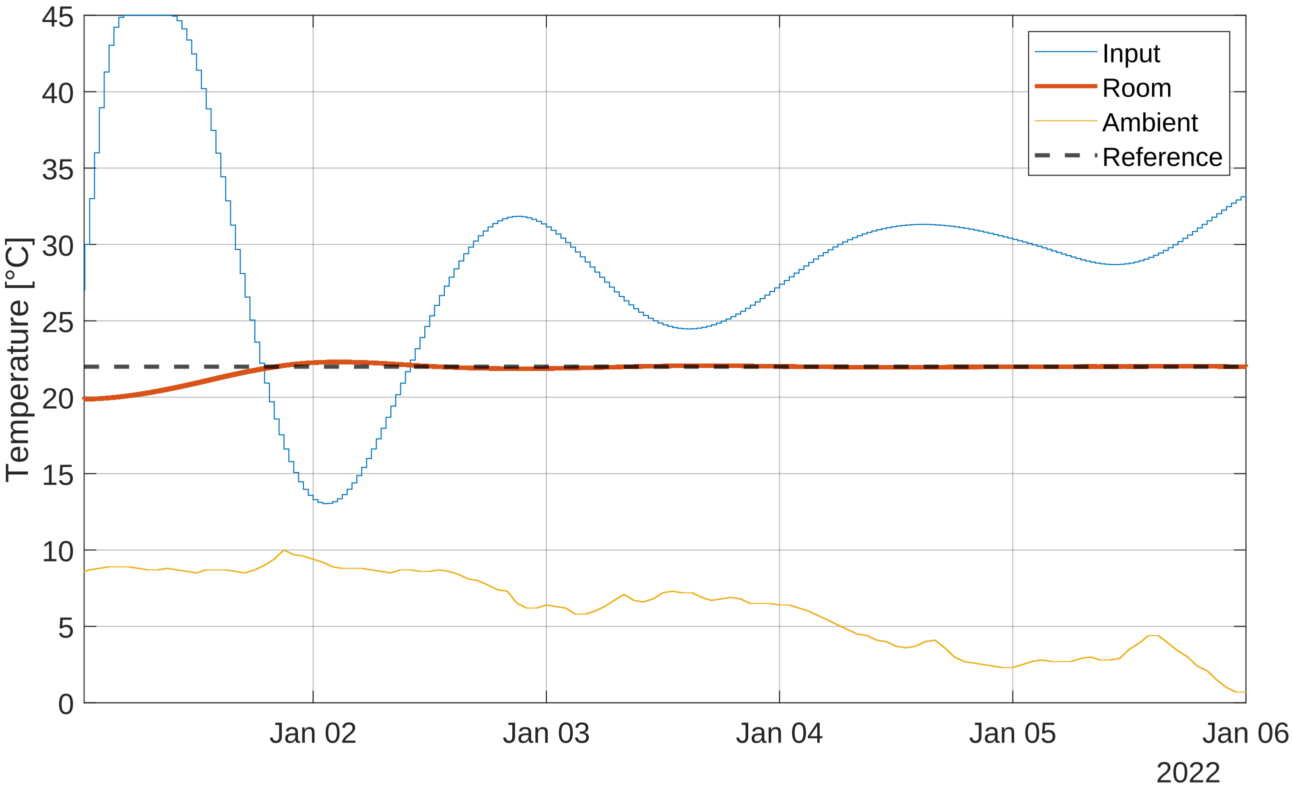

Figure 3.10: The beginning of the output response of the room temperature in Figure 3.7.

The economical cost of running this controller is of interest for comparison reasons. Figure 3.11 shows the historical Danish electricity price $m_e(k)$ for the entirety of year 2022, with the electricity price spiking somewhere above 6 DKK. This price includes both the actual cost of the electricity and the tariffs as mention in Section 3.3. $P_c(k)$ is the simulated power consumption of the simulated system throughout 2022. As mentioned in Section 3.3, the difference in input and output temperature is utilised to get the hypothesised primary energy consumption. $\bar{V}(k)$ is the cumulative sum of the simulated total price over the entire year, ending at a cost of 6463.56 DKK. Remark, that this is not necessarily accurate wrt. the cost of operation in reality.

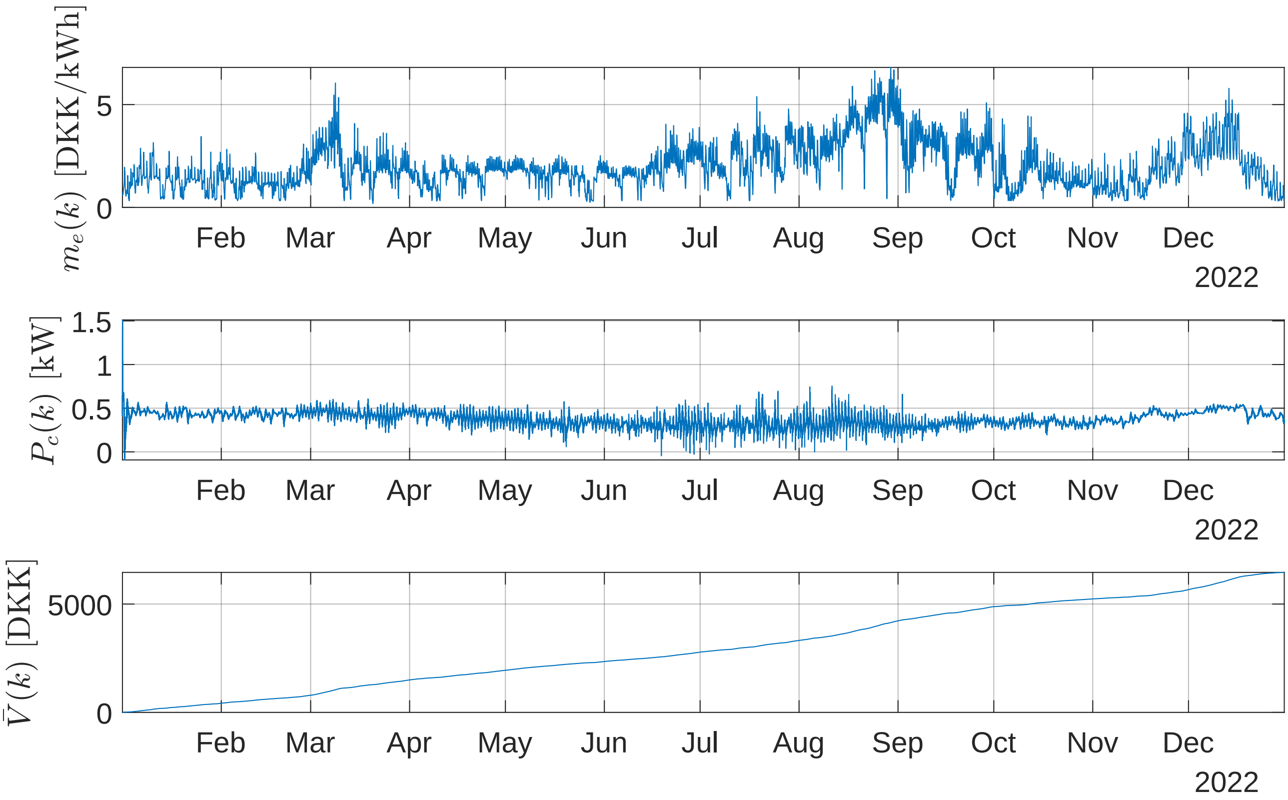

Figure 3.11: Graph of price data, simulated power usage from the state response, and calculated total price spanning the year 2022. The signal $m_e(k)$ is the electricity price, $P_c(k)$ is the simulated power usage, and $\bar{V}(k)$ is the cumulative sum of the simulated total price.

In general, the behaviour of the system is as expected; the input temperature is raised quickly to compensate for the 2 °C tracking error that is immediately present, and the temperature is reduced to compensate for overshoot. As time goes on, the tracking error is lowered and at steady state, after a couple of days, it is negligible. However, ambient temperature changes may force great tracking error as the room temperature is raised beyond reference. Since the plant input is constrained to stay above ambient temperature the controller is unable to lower the room temperature. However, this scenario is deemed outside the scope of this project. With the baseline comfort controller performance deemed acceptable, the design of the economic cost MPC can begin.

4 Economic Cost Controller

In this chapter the comfort controller will be extended to the economic cost controller such that it incorporates the power consumption and the price of said power consumption. The price is sourced from Energi Data Service as mentioned previously. As the controller is designed to be run online, the horizon of electricity price data must be considered. This, along with a new measurement incorporating the cost is introduced in this chapter. The inherent cost of power consumption is tuned, such that the controller is capable of utilising low-cost periods while avoiding high-cost periods while still maintaining a comfortable indoor temperature. Furthermore, a scheme is chosen to limit and alter the electricity price, such that the controller is justifiably always within some temperature range even in the face of extreme price periods. Akin to Section 3.4, this chapter presents results of a simulation of the designed controller and model with real weather and electricity price data for one year.

4.1 Introducing Electricity Prices

If the controller is to respond to electricity prices, they must, in some way, factor into the cost function. In Section 3.3, a new variable $\Delta T_{r,s}$ is constructed as the difference between supply and return temperature after which it is proportionally fitted to the measured power consumption. This new variable was stated in Equation (3.41) and re-stated in Equation (4.1).

$$\Delta T_{r,s} = T_{in} - T_{w,r} \tag{4.1}$$

Furthermore, Section 3.3 theorises that the one step power consumption is proportional to this variable with the factor α. This relationship is re-stated in Equations (4.2).

$$P_c(k) = \alpha \Delta T_{r,s}(k) \tag{4.2a}$$

$$V(k) = P_c(k) \cdot \frac{\tau_s}{60^2} \cdot m_e(k) \tag{4.2b}$$

where:

$P_c(k)$	is the power consumption of the central heating system	[kW]
$V(k)$	is the cost of operation for a single step of length τ_s at time k	[DKK]
$m_e(k)$	is the cost of electricity at time k	[DKK/(kWh)]

The proportionality coefficient α and the coefficient due to the changing units is combined into one:

$$\alpha_V = \alpha \cdot \frac{\tau_s}{60^2} \tag{4.3}$$

Both $P_c(k)$ and $m_e(k)$ are non-constant, and their product $V(k)$ cannot be represented by the state space representation used to lift the QP-problem. As such it is chosen to incorporate $P_c(k)$ as a measurement in the state space, and use the electricity price $m_e(k)$ in the cost weight.

The power consumption $P_c(k)$ is added as a measurement in $z(k)$ and the lifting performed as in Section 3.1. However, as the dependent variable $\Delta T_{r,s}$ is dependent on the input, it cannot be directly incorporated into the lifted matrices presented in Section 3.1, as this section assumes empty direct terms. A direct term is therefore added such that the input temperature is fed directly to the output enabling the constructing of the desired measurement in Equation (4.1), giving the measurement vector $z(k)$ stated in full in Equations (4.4).

$$z(k) = \begin{bmatrix} T_r \\ P_c \end{bmatrix} = \begin{bmatrix} & & C_z & & \\ 0 & 0 & -\alpha_V & 0 & 0 \end{bmatrix} x(k) + \begin{bmatrix} D_z \\ \alpha_V \end{bmatrix} u(k) + \begin{bmatrix} D_{dz} \\ 0 \end{bmatrix} d(k) \tag{4.4a}$$

where

$$x(k) = \begin{bmatrix} T_{w,1} & T_{w,2} & T_{w,3} & T_f & T_r \end{bmatrix}^T \tag{4.4b}$$

$$u(k) = T_{in} \tag{4.4c}$$

$$d(k) = T_a \tag{4.4d}$$

and $T_{w,3} = T_{w,r}$ as per the fitting in Section 2.3. These new matrices replace C_z, D_z and D_{dz} from this point on.

The direct term is lifted similarly to the previous methods, i.e. by expressing the output as a function of change in input before collecting into matrices. This procedure is performed in Equations (4.5). Remark that the change in input beyond $k + H_u - 1$ is considered 0, why the rows of the lifted matrices are the same from that point onward.

$$z(k) = C_z\hat{x}(k) + D_z u(k) \tag{4.5a}$$

$$z(k) = C_z\hat{x}(k) + D_z(\Delta u(k) + u(k-1)) \tag{4.5b}$$

$$z(k+1) = C_z\hat{x}(k+1) + D_z(\Delta u(k+1) + \Delta u(k) + u(k-1)) \tag{4.5c}$$

$$\vdots$$

$$z(k+H_u-1) = C_z\hat{x}(k+H_u-1) + D_z\left(\sum_{i=0}^{H_u-1} \Delta u(k+i) + u(k-1)\right) \tag{4.5d}$$

$$z(k+H_u) = C_z\hat{x}(k+H_u) + D_z\left(\sum_{i=0}^{H_u-1} \Delta u(k+i) + u(k-1)\right) \tag{4.5e}$$

$$\vdots$$

$$z(k+H_p) = C_z\hat{x}(k+H_p) + D_z\left(\sum_{i=0}^{H_u-1} \Delta u(k+i) + u(k-1)\right) \tag{4.5f}$$

$$\Rightarrow \mathcal{Z} = \mathcal{C}\mathcal{X} + \Lambda u(k-1) + \Lambda_\Delta \Delta\mathcal{U}(k) \tag{4.5g}$$

where

$$\Lambda = \begin{bmatrix} D_z \\ D_z \\ \vdots \\ D_z \end{bmatrix} \qquad \Lambda_\Delta = \begin{bmatrix} D_z & 0 & 0 & \cdots & 0 \\ D_z & D_z & 0 & \cdots & 0 \\ & \vdots & & \ddots & \\ D_z & D_z & D_z & \cdots & D_z \\ & & \vdots & & \\ D_z & D_z & D_z & \cdots & D_z \end{bmatrix}$$

The dimensions of Λ is $pH_p \times n$ and Λ_Δ is $pH_p \times H_u n$ where p is the number of outputs and n is the number of inputs. By considering the direct terms as altering the lifted matrices, they are easily incorporated in the QP-problem. This is done by, once again, expanding the cost function as seen in Equations (4.6).

$$J(k) = \|\mathcal{Z}(k) - \mathcal{E}(k)\|^2_{\mathcal{Q}} + \|\Delta\mathcal{U}(k)\|^2_{\mathcal{R}} \tag{4.6a}$$

$$\mathcal{Z}(k) = \|\Theta\Delta\mathcal{U}(k) + \Psi\hat{x}(k) + \Upsilon u(k-1) + \Xi d(k) + \Lambda u(k-1) + \Lambda_\Delta\Delta\mathcal{U}(k) - \mathcal{E}(k)\|^2_{\mathcal{Q}} \tag{4.6b}$$

$$= \|(\Theta + \Lambda_\Delta)\Delta\mathcal{U}(k) + \Psi\hat{x}(k) + (\Upsilon + \Lambda)u(k-1) + \Xi d(k) - \mathcal{E}(k)\|^2_{\mathcal{Q}} \tag{4.6c}$$

$$= \|\Theta_D\Delta\mathcal{U}(k) + \Psi\hat{x}(k) + \Upsilon_D u(k-1) + \Xi d(k) - \mathcal{E}(k)\|^2_{\mathcal{Q}} \tag{4.6d}$$

This changes the QP-problem in Equations (3.27) including the constraints, as stated in Equations (4.7).

$$\min_{\Delta U(k)} J(k) = -\Delta\mathcal{U}(k)^T\mathcal{G} + \Delta\mathcal{U}(k)^T\mathcal{H}\Delta\mathcal{U}(k) \tag{4.7a}$$

subject to:

$$\begin{bmatrix} \mathbf{F} \\ \Gamma\Theta_D \\ W \end{bmatrix} \Delta\mathcal{U}(k) \leq \begin{bmatrix} -\mathbf{F}_1 u(k-1) - f \\ -\Gamma[\Psi x(k) + \Upsilon_D u(k-1)] - \gamma \\ w \end{bmatrix} \tag{4.7b}$$

where

$$\mathcal{H} = \Theta_D^T\mathcal{Q}\Theta_D + \mathcal{R} \tag{4.7c}$$

$$\mathcal{E}(k) = \mathcal{T} - (\Psi x(k) + \Upsilon_D u(k) + \Xi d(k)) \tag{4.7d}$$

$$\mathcal{G} = 2\Theta_D^T\mathcal{Q}\mathcal{E}(k) \tag{4.7e}$$

As mentioned, it is chosen that the electricity price $m_e(k)$ enters at the cost weight. While it is possible to insert the electricity price elsewhere in the MPC, such as inside the state transition matrices, this is deemed more intuitive. Recall that the cost matrix $Q(i)$ is the quadratic cost weight on reference tracking, which means that a reference of 0 results in a positive cost for any non-zero output. Since the state space now has two measurements, the reference is a vector consisting of the room temperature reference and power consumption reference:

$$r(k) = \begin{bmatrix} r_{T_r} \\ r_{P_c} \end{bmatrix} = \begin{bmatrix} r_{T_r} \\ 0 \end{bmatrix} \tag{4.8}$$

The reference for the power consumption is set to 0, as it is desired to be as low as possible. In turn, the cost matrix becomes a 2×2 matrix which is chosen in Equation (4.9).

$$Q(i) = \begin{bmatrix} \alpha_Q & 0 \\ 0 & \beta_Q(\alpha_V m_e(i))^2 \end{bmatrix} \tag{4.9}$$

This results in a cost function $J(k)$ which can be seen in Equations (4.10).

$$J(k) = \sum_{i=1}^{H_p} \|\hat{z}(k+1) - r(k+1)\|^2_{Q(i)} + \sum_{i=0}^{H_u-1} \|\Delta\hat{u}(k+i)\|^2_{R(i)} \tag{4.10a}$$

$$= \sum_{i=1}^{H_p} \alpha_Q (T_r - r_{T_r})^2 + \beta_Q(\alpha_V m_e(k) P_c(k))^2 + \sum_{i=0}^{H_u-1} R(i)\Delta T_{in}^2 \tag{4.10b}$$

$$= \sum_{i=1}^{H_p} \alpha_Q (T_r - r_{T_r})^2 + \beta_Q V(k)^2 + \sum_{i=0}^{H_u-1} R(i)\Delta T_{in}^2 \tag{4.10c}$$

Now, the cost of operation $V(k)$ is directly included in the cost function. This design of $Q(i)$ is favourable as it poses the following distinguishable design parameters:

- The coefficient α_Q, which determines the cost of reference tracking error on the room temperature.
- The coefficient β_Q, which determines the cost of the squared cost of operation, that being the product of the power consumption and the electricity price.
- The scalar weight $R(i)$, which determines the cost of a change in input.

The design parameters are difficult to choose, as no performance metrics can be calculated without running the controller. Therefore, choosing $\alpha_Q, \beta_Q, R(i)$ requires a large number of simulations in order to search the parameter space for a sufficient solution. Luckily, only the ratio between these coefficients matter, as the same optimising input is produced for any coefficients as long as the ratio is maintained. This search is done later in this chapter, as a few issues need to be addressed first, including the variable length of the price signal $m_e(k)$.

4.2 Dynamic Electricity Price Horizon

The horizon of the electricity price gathered from Energi Data Service changes, which causes problems in the MPC if the prediction horizon sees beyond the last known point. The electricity price signal varies in length as the electricity price forecast is published at 13:00 at which point the electricity price is available until 23:59 the day after. This means, that at 13:00 the electricity price is available for the next $l_h = 35$ whole hours. The hour before, at 12:00, the electricity price is available until 23:59, that is, the price is available for only $l_h = 12$ whole hours. This is illustrated in Figure 4.1.

As the electricity price signal is interpolated to match the sample interval of the system, the actual length l of the disturbance becomes far greater. More precisely:

$$l = 60^2 \cdot \frac{l_h}{\tau_s} = 60 l_h \tag{4.11}$$

Due to the downsampling performed in Section 3.2.1, the controller is capable of seeing H_p samples with an interval of κ up until the last sample at κH_p. The lifted matrices in Section 3.1 are designed for this, and deviation from this input size makes the matrix calculations incompatible. Depending on the chosen design parameters, it is quite likely that κH_p exceeds the number of samples provided by the electricity price forecast, which is in the interval $l_h \in [12, 35]$ corresponding to samples $l \in [720, 2100]$ of sample time $\tau_s = 60\,\text{s}$. Therefore, a scheme is chosen to allow the controller to continue running, even if it cannot see electricity prices for the entire duration κH_p.

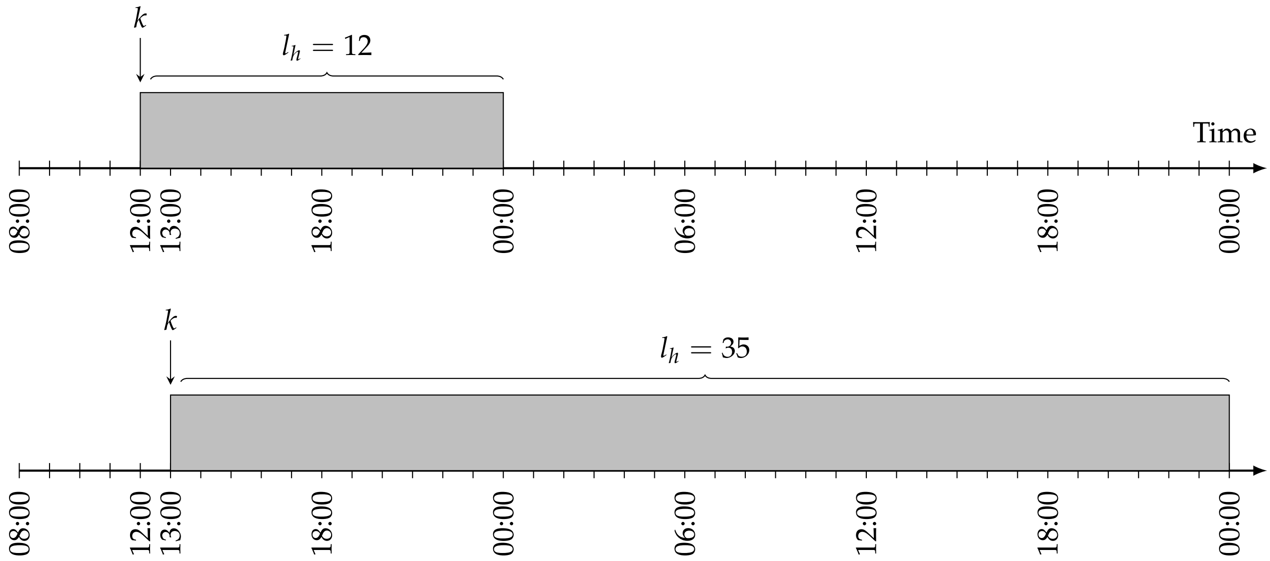

Figure 4.1: Illustration of a changing horizon due to the schedule in which forecast electricity prices are published at 13:00 until midnight the following day.

While there are several ways of doing this, the implementation of a variable prediction horizon in the MPC is chosen. This requires altering the lifted matrices from Section 3.1 dynamically. Constructing these matrices from scratch is mathematically trivial but computationally significant, as the horizon may change rapidly. Therefore the variable price length l is used to limit the variable prediction horizon and crop the lifted matrices. The idea is that a variable prediction horizon $\bar{H}_p$ is constructed as the minimum of the chosen upper bound H_p and the disturbance length l normalised to κ and rounded down to nearest integer value. This variable prediction horizon is seen in Equation (4.12).

$$\bar{H}_p = \min\left(H_p, \left\lfloor \frac{l}{\kappa} \right\rfloor\right) \tag{4.12}$$

where:

$\lfloor x \rfloor$ is x rounded down to the nearest integer value (floor function)

Equation (4.12) has the fixed length H_p as the upper limit and the variable length l/κ as the lower limit. This results in $\bar{H}_p$ being mostly H_p except when the last sample κH_p exceeds l, in which case $\bar{H}_p$ becomes the largest it can be, that being the floored l/κ. A similar function is constructed for the control horizon to ensure it will never eclipse the prediction horizon, which can be seen in Equation (4.13).

$$\bar{H}_u = \min(H_u, \bar{H}_p) \tag{4.13}$$

The behaviour of Equation (4.13) is very similar to the behaviour of $\bar{H}_p$, instead having H_u as the upper limit and $\bar{H}_p$ as the lower. Figure 4.2 shows an illustration of the variable electricity price length l and the corresponding prediction and control horizons in system sample periods, that is, $\kappa\bar{H}_p$ and $\kappa\bar{H}_u$. When constrained by the forecast horizon limit, the actual value of $\bar{H}_p$ is the largest integer such that $\kappa\bar{H}_p$ is less than or equal to l and likewise for κH_u. While Figure 4.2 shows $\kappa\bar{H}_p$ and l coinciding, this is only the case if κ is a divisor of 60. In every case, $\kappa\bar{H}_p$ is always less than or equal to l at every controller instant.

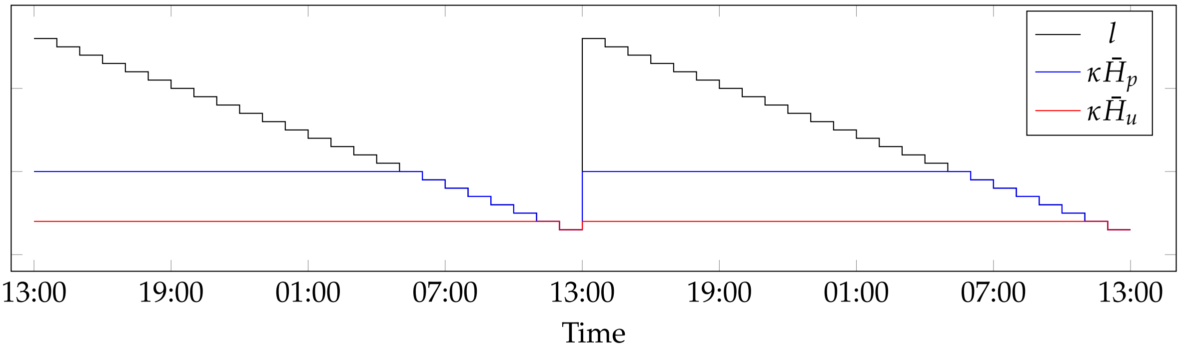

Figure 4.2: Illustration of the desired behaviour of the dynamic horizons $\bar{H}_p$ and $\bar{H}_u$ as the length l of the electricity price forecast changes.

The reference signal and ambient temperature are indexed to be the same length as $\bar{H}_p$ to ensure consistency in the controller. The lifted matrices, Ψ, Υ, Θ, and Ξ are cropped such that the matrix equations are compatible. If these lifted matrices are cropped correctly, it is equivalent to redoing the lifting with a smaller H_p and H_u because of their structure (see Section 3.1). The notation in Equation (4.14) is used to indicate a matrix A which is cropped to from entries $a_{i1,j1}$ to $a_{i2,j2}$.

$$A_{>(i1,j1),(i2,j2)<} = \begin{bmatrix} a_{i1,j1} & a_{i1+1,j1} & \cdots & a_{i2,j1} \\ a_{i1,j1+1} & a_{i1+1,j1+1} & \cdots & a_{i2,j1+1} \\ \vdots & \vdots & \ddots & \vdots \\ a_{i1,j2} & a_{i1+1,j2} & \cdots & a_{i2,j2} \end{bmatrix} \tag{4.14}$$

As an example, $\bar{\Psi}$ is constructed by cropping Ψ using $\bar{H}_p$ as seen in Equation (4.15).

$$\Psi = \begin{bmatrix} CA \\ CA^2 \\ \vdots \\ CA^{H_p} \end{bmatrix} \qquad \bar{\Psi} = \begin{bmatrix} CA \\ CA^2 \\ \vdots \\ CA^{\bar{H}_p} \end{bmatrix} = \Psi_{>(1,1),(\bar{H}_p m,m)<} \tag{4.15}$$

where:

m is the number of states

The other lifted matrices are systematically cropped in the same fashion and presented in Equations (4.16).

$$\bar{\Psi} = \Psi_{>(1,1),(\bar{H}_p m,m)<} \tag{4.16a}$$

$$\bar{\Upsilon} = \Upsilon_{>(1,1),(\bar{H}_p p,n)<} \tag{4.16b}$$

$$\bar{\Theta} = \Theta_{>(1,1),(\bar{H}_p p,\bar{H}_u n)<} \tag{4.16c}$$

$$\bar{\Xi} = \Xi_{>(1,1),(\bar{H}_p p,H_p j)<} \tag{4.16d}$$

The variable control horizon also affects the constraint matrices, and as such need to be resized for compatibility. The constraint matrices are constructed from scratch, as they require the disturbance in the form of the ambient temperature. Therefore, it does not need to be cropped, as the variable

$\bar{H}_u$ is accounted for during construction.

Finally, akin to the algorithm for the comfort controller presented in Algorithm 1, a pseudocode algorithm is written for the economic controller, to provide an overview of its functionality. This algorithm is presented in Algorithm 2 and works similarly to before, updating the observer with measured values and only running the controller every κ sample. Additionally, the economic MPC algorithm incorporates the dynamic prediction and control horizons $\bar{H}_p$ and $\bar{H}_u$ and uses them to crop disturbance, reference, weights, and the lifted system. For a full overview of the algorithms used in this report, refer to Appendix A.

4.3 Ensuring a Region of Comfort

The cost now includes two different objectives, those being the tracking of the room temperature reference and the lowering of the cost of operation. These objectives pull the temperature in two different directions simultaneously, resulting in the objective with the highest cost dominating the temperature. At the extremes, weighing reference tracking significantly higher than operation cost approximately leads to the same response as the comfort controller. The opposite weights result in the room temperature approaching ambient, as the input temperature is lowered to the ambient temperature. This is exemplified in Figure 4.3 which shows the MPC response of one year for two extreme ratios between α_Q and β_Q. No price weight corresponds to $\beta_Q = 0$ while excessive price weight corresponds to β_Q being far greater than α_Q.

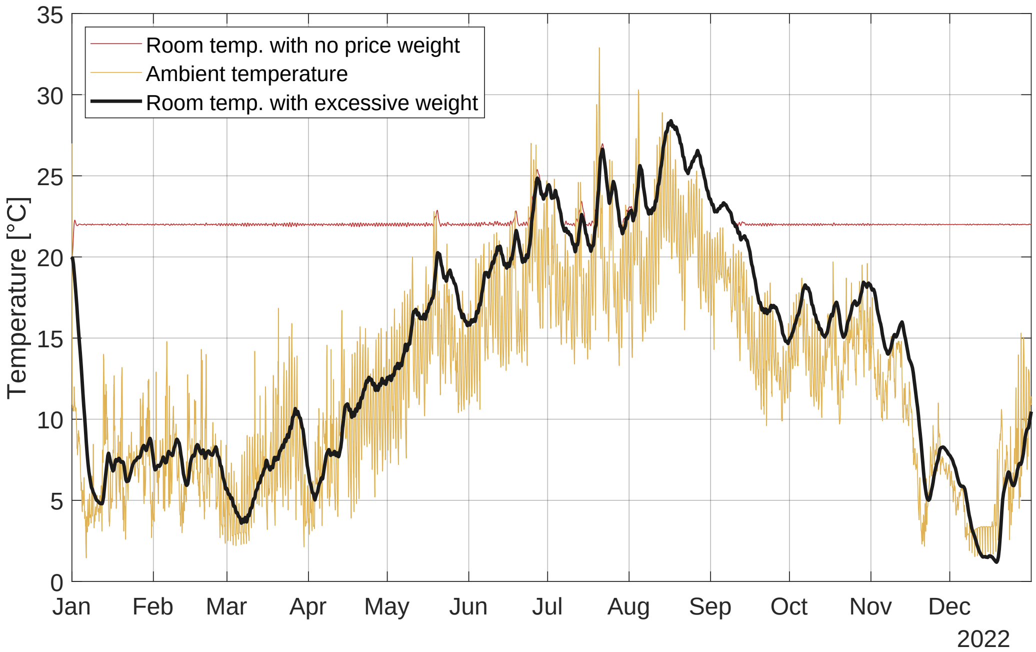

Figure 4.3: The simulated output response of the MPC for the period 01-01-2022 to 01-01-2023 for two extreme weight ratios.

For this reason a satisfactory medium, where the end user of the provided house can both have a sufficiently comfortable room temperature while reducing the cost of heating, is necessary for this controller to achieve success.

Algorithm 2 Implementation of the economic MPC with downsampling, dynamic prediction horizon and cropping. Refer to Appendix A for an overview of the non-trivial internal functions used.

Require: $A \in \mathbb{R}^{m \times m}, B \in \mathbb{R}^{m \times n}, B_d \in \mathbb{R}^{m \times j}$ ▷ System equations and their dimensions
Require: $C_z \in \mathbb{R}^{p \times m}, D_z \in \mathbb{R}^{p \times n}, D_{dz} \in \mathbb{R}^{p \times j}$
Require: $\mathcal{D} \in \mathbb{R}^{j_d}, \quad \mathcal{M}_e \in \mathbb{R}^{j_m}, \quad \mathcal{T} \in \mathbb{R}^{j_t}, j_t > \max(j_d, j_m)$
Require: $1 \le H_p, \quad 1 \le H_u \le H_p$
Require: $\kappa \in \mathbb{Z}, \quad 1 \le \kappa \le \min(j_d, j_m, j_t)$ ▷ κ must be integer. Ensure at least one sample point.
$k = 0$
$k_{mpc} = 1$
$\hat{x} = \hat{x}_0 \in \mathbb{R}^n$
$u = u_0 \in \mathbb{R}^1$
repeat
  $k = k + 1$
  $t = \text{GET\_TIME}(k)$
  $y = \text{GET\_MEASUREMENT}$ ▷ Measurement is performed
  if $(k-1) \bmod \kappa = 0$ **then** ▷ Controller loop is downsampled by κ
    $k_{mpc} = k$
    $\bar{H}_p = \left\lfloor \min\left(H_p, \frac{j_d}{\kappa}, \frac{j_m}{\kappa}\right) \right\rfloor$ ▷ Prediction horizon limited to disturbance length
    $\bar{H}_u = \min(H_u, \bar{H}_p)$
    $\bar{d} = \begin{bmatrix} \mathcal{D}(k) & \mathcal{D}(k+\kappa) & \dots & \mathcal{D}(\bar{H}_p \cdot \kappa) \end{bmatrix}^T$ ▷ Disturbance & reference picked accordingly
    $\bar{r} = \begin{bmatrix} \mathcal{T}(k) & \mathcal{T}(k+\kappa) & \dots & \mathcal{T}(\bar{H}_p \cdot \kappa) \end{bmatrix}^T$
    $\bar{m}_e = \begin{bmatrix} \mathcal{M}_e(k) & \mathcal{M}_e(k+\kappa) & \dots & \mathcal{M}_e(\bar{H}_p \cdot \kappa) \end{bmatrix}^T$
    $\mathcal{Q} = \text{WEIGHT\_PRICE}(\bar{m}_e)$ ▷ Produced as explained in Section 4.1
    $\bar{\mathcal{Q}} = \mathcal{Q}_{>(1,1),(p\bar{H}_p, n\bar{H}_u)<}$
    $\bar{\mathcal{R}} = \mathcal{R}_{>(1,1),(p\bar{H}_p, n\bar{H}_u)<}$
    $(\bar{\Psi}, \bar{\Upsilon}, \bar{\Theta}, \bar{\Xi}, \bar{\Lambda}, \bar{\Lambda}_D) = \text{LIFT\_SYSTEM}(\bar{H}_p, \bar{H}_u)$
    $\mathcal{H} = (\Theta + \Lambda_D)^T \mathcal{Q} (\Theta + \Lambda_D) + \mathcal{R}$
    $\mathcal{E} = \bar{r} - \bar{\Psi} x - (\Upsilon + \Lambda) u - \Xi \bar{d}$
    $\mathcal{G} = 2(\Theta + \Lambda_D)^T \mathcal{Q} \mathcal{E}$
    $(A_{con}, b_{con}) = \text{LIFT\_CONSTRAINTS}(\bar{H}_p, \bar{H}_u, u, \bar{d}, \bar{r}, \bar{\Upsilon}, \bar{\Theta}, \bar{\Xi}, \bar{\Lambda}, \bar{\Lambda}_D)$
    $\Delta\mathcal{U} = \text{QUADPROG}(2\mathcal{H}, -G^T, A_{con}, b_{con})$ ▷ Constrained QP problem solved
    $\mathcal{U} = \begin{bmatrix} u + \Delta\mathcal{U}(k) & u + \Delta\mathcal{U}(k+1) & \dots & u + \Delta\mathcal{U}(k + \bar{\kappa} H_p) \end{bmatrix}$
  end if
  $\hat{y} = C_y \hat{x} + D_y u + D_{dy} d$ ▷ Estimates are produced
  $\hat{x} = A\hat{x} + L(y - \hat{y}) + Bu + B_d d$
  $u = \mathcal{U}(k - k_{mpc} + 1)$
until $k > k_{final}$

4.3.1 Electricity Price Limits

An important consideration when attempting to ensure a region of comfort, is that the electricity price disturbance is dynamic and, in theory, limitless. Therefore, large price spikes may occur, which without some kind of restriction, can lead to a significant drop in the room temperature and thus the residents' relative comfort will suffer. A too significant drop is unacceptable for the controller, so a solution must be found. Multiple solutions to this problem are proposed:

- Constraining the MPC output, such that room temperature is always within some acceptable range.
- Limiting the price disturbance by linearly mapping minimum and maximum electricity price of the current period to some lower and upper bound respectively.
- Making the mean of the electricity price signal of the given period zero.
- Limiting the price disturbance by an upper and a lower bound.

The three approaches which alter the electricity price are illustrated in Figure 4.4 alongside the original electricity price, for a high-cost period with much variation, denoted *turbulent*, and a low-cost period with less variation, denoted *calm*.

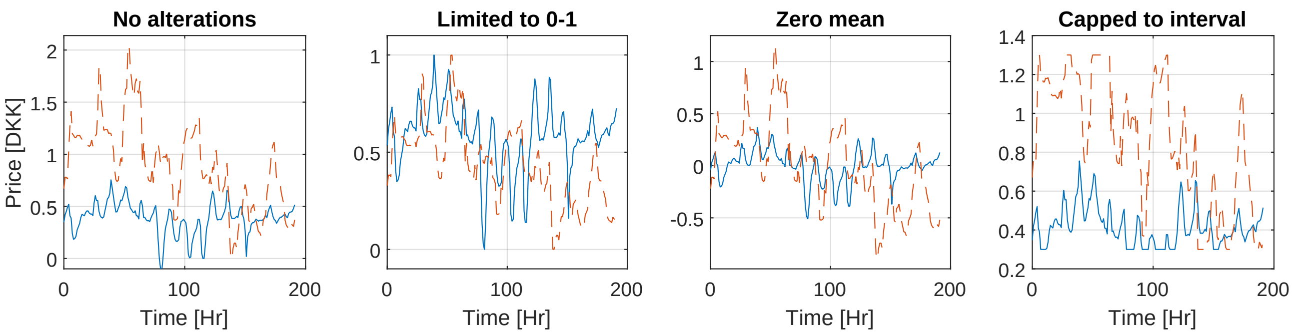

Figure 4.4: The electricity price $m_e(k)$ and three electricity prices with applied schemes for limiting the effect of extreme price values. The blue solid line marks a calm period from 01-05-2021 to 08-05-2021, while the dashed orange line marks a turbulent period from 01-01-2022 to 08-01-2022.

The advantages and disadvantages of the four approaches are now explored. Output constraints are the easiest to design but have a major drawback. The output constraints ensure that the room temperature, at least in theory, never exceeds or falls below a given temperature and can thus guarantee comfort of the room. The drawback is, that the system can enter an infeasible area in which this constraint is violated and be unable to leave. This is especially prevalent in the start-up phase of the MPC and upon major changes in the ambient temperature disturbance. As an example, the results of the comfort controller in Figure 3.7 shows an ambient temperature so high that the room temperature is raised way beyond the reference. A realistic output constraint would certainly be violated in this case. Methods for handling infeasibility do exist, such as introducing slack variables which allow for slack in the constraint at some cost in the cost function [11]. However, constraints are usually tied to physical constraints within the system which this comfort output constraint is not. Therefore, this approach is not chosen.

Another possible solution is limiting the price disturbance between 0 and 1, which, in this context, means finding the minimum and maximum of the predicted price disturbance and normalising it to 0 and 1 respectively. This normalised factor is then weighted with β_Q. This approach has the advantage of a known maximum and minimum, which aids in establishing a region of comfort, while also keeping the peaks and valleys, i.e. the high and low cost regions of the price signal, intact. It is hypothesised that this solution performs well at turbulent price periods, where there are significant differences in price peaks and valleys. Since this method forces the price disturbance within the interval 0 to 1, the comfort of the house can be relatively guaranteed with the right choice of β_Q. Au contraire, for periods with relatively low difference between peaks and valleys, this method will treat even minor spikes as large disturbances, which will elicit an overreaction by the MPC. This results in the controller being unable to truly utilise low price periods, and as such, this solution is also not chosen.

Another attempted solution is making the price period mean zero. This is done by subtracting the sample mean of the electricity prices in a given region from the electricity price signal. The idea is, that prolonged periods of high cost electricity do not cause the room temperature to go too low, while still enabling the MPC to act upon price peaks and valleys. This method entails a price signal which is often negative, which poses a problem in the cost function as it weights the cost of operation squared, leading to high costs for low energy prices. To combat this, the zero mean is simply lifted. There are several problems with this method the least of which being finding a new mean that ensures a positive price in the face of an unbounded disturbance. In spite of the mean being set to some value, the peaks and valleys may still reach extreme values, in turn sacrificing comfort for the residents. Consequently, this method is discarded as well.

Lastly, the conceptually easiest method is to limit the price signal within some interval. This means that an electricity price above some upper bound is set to the upper bound and an electricity price below some lower bound is set to the lower bound. This is referred to as *capping* the signal. If done correctly, this capping prevents extreme price peaks or valleys from causing too much change in room temperature, as is the case when normalising the mean. Furthermore, as long as the price is not extreme, the price dynamic is not altered, unlike the limit proposed earlier. Lastly, this method also avoids the possible infeasabilities caused by an output constraint. For these reasons, this method is chosen to justifiably ascertain a region of comfort for the economic cost controller. The choice of price normalisation is further discussed in Section 5.2.

Finding the limiting interval is not a trivial task. If the interval is too narrow, prices are capped needlessly causing the MPC to only respond mildly to the dynamics. On the contrary, an interval too wide approaches the non-limited price signal, which features the problems of extreme electricity prices discussed earlier. The bounds of the interval could be part of the parameter search for α_Q and β_Q, but this increases computation time. Instead, a qualatitive method is used to find the interval, in which electricity prices from January 2018 to January 2023 are examined. The lower threshold is chosen as the 5^{th} percentile of these electricity prices. Likewise for the upper threshold, the 95^{th} percentile is utilised as the upper limit. This will ensure the largest of price spikes do not have a too significant influence on the system. The capping interval obtained from these percentiles is $0.05\,\mathrm{DKK}/(\mathrm{kWh})$ to $2.2\,\mathrm{DKK}/(\mathrm{kWh})$. Capping the price disturbance seems to be the most viable solution to the dynamic disturbance problem. However, tuning of β_Q is still required.

4.3.2 Choice of Economic Weights

There are three weights that should be tuned such that the MPC provides a reasonable comfort level while responding to the dynamically changing electricity price. These three weights are the room temperature reference tracking weight α_Q, the electricity price weight factor β_Q and the input rate weight $R(i)$. As mentioned previously, only the ratio between α_Q, β_Q, and $R(i)$ is of importance. The following procedure is performed to select a suitable set of parameters:

1. The input rate weight $R(i)$ is fastened to some constant.
2. The reference tracking weight α_Q is fastened to some constant.
3. A lower and an upper limit of the region of comfort is chosen. These are denoted $T_{r,min}$ and $T_{r,max}$ respectively.
4. A price weight β_Q is found heuristically, such that a simulation of the room temperature response over one year exceeds either $T_{r,min}$ or $T_{r,max}$. This is denoted $\beta_{Q,max}$ and must be greater than zero.
5. A total of n simulations is made, in which $\beta_{Q,i}$ iterates over n equally spaced real values between 0 and $\beta_{Q,max}$. These simulations should cover one year or more.
6. The largest $\beta_{Q,i}$ for which the simulated room temperature never exceeds the interval $[T_{r,min}, T_{r,max}]$ is chosen as the electricity price weight factor β_Q.

This procedure allows for choosing the parameters based on the requirements of the region of comfort. It is chosen that $R(i) = 0.1$ and $\alpha_Q = 1$. While the change in input T_{in} is constrained, a nonzero $R(i)$ is preferred as this tends to give a slower response acting as a buffer, in case the input rate constraint is modelled incorrectly. Still, the reference tracking is much more important, and as such α_Q is a factor 10 larger. Heuristically it is found that $\beta_{Q,max} = 3$, why a an iterative simulation is performed for $\beta_Q = 0, 0.125, 0.25, \ldots, 3$ for a total of $n = 25$ points. The sample size n is limited by the computation time of the simulations, but nevertheless provides valuable results as seen in Figure 4.5, which shows the simulated room temperature for a small region of the simulated period from 01-01-2022 to 01-01-2023. The comfort region chosen is the reference $\pm 2\,°\mathrm{C}$. The full list of parameters used in the search for β_Q is seen in Equations (4.17). Additional parameters used for simulation are found in Appendix C.

$$R(i) = 0.1 \tag{4.17a}$$
$$\alpha_Q = 1 \tag{4.17b}$$
$$[T_{r,min}, T_{r,max}] = [20, 24]\,°\mathrm{C} \tag{4.17c}$$
$$\beta_{Q,max} = 3 \tag{4.17d}$$
$$n = 25 \tag{4.17e}$$

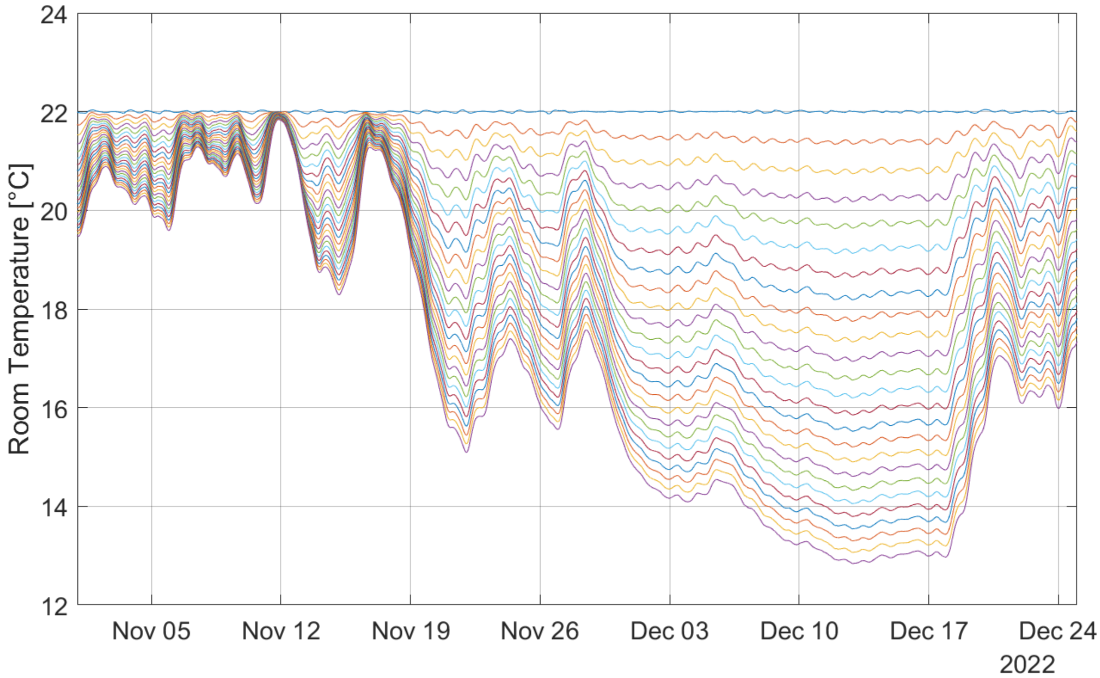

Figure 4.5: Simulated room temperature response for various β_Q values. The purple line fourth from the top is the largest $\beta_{Q,i}$ for which the chosen room temperature interval is never exceeded.

As seen in Figure 4.5, the responses converge to the same value at certain peaks, while diverging elsewhere. These peaks, in which the responses converge, are points in time with low electricity prices. As the electricity price $m_e(k)$ tends to zero, the corresponding weight tends to zero as well, regardless of β_Q. This results in α_Q being the dominant weight, in turn heavily prioritising the reduction of room temperature tracking error. When the price is high, β_Q plays a larger role resulting in lower room temperatures. The region shown in Figure 4.5 includes the minimum value for every iteration, and as such, an informed choice of β_Q can be made merely from this. However, to better showcase the relationship between β_Q and the minimum simulated room temperature, the relationship is plotted in Figure 4.6.

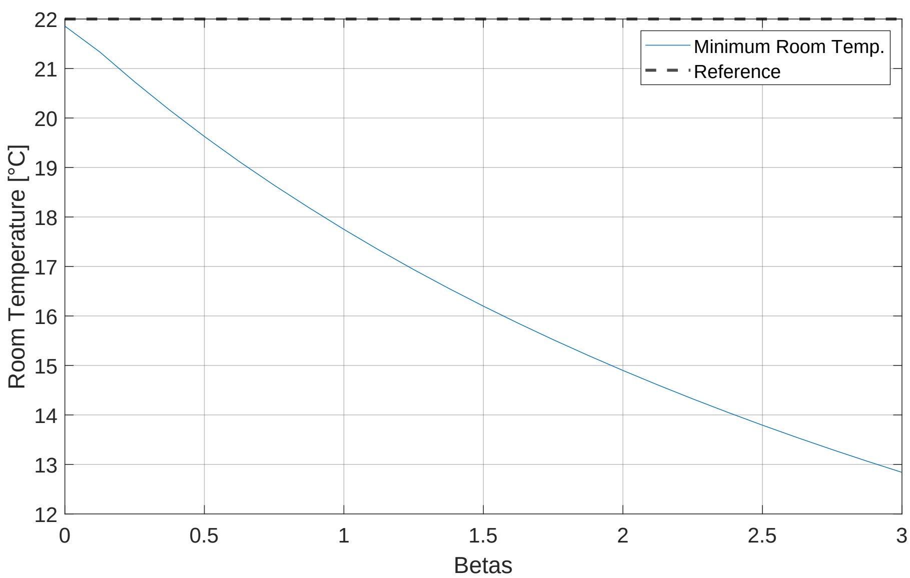

Figure 4.6: The relationship between values of β_Q and the minimum of the simulated room temperature.

It is now possible to choose a β_Q such that indoor room temperature justifiably stays within the chosen comfort interval. The weight β_Q is chosen to be 0.375, since it is the largest β_Q with a minimum above $T_{r,min} = 20\,°C$, that minimum being exactly 20.16 °C. With the parameters tuned, it is now possible to perform the simulation, in which the results are compared to the baseline comfort controller.

4.4 Economic Controller Simulation Results

In this section, the simulated response of the economic cost controller and costs of operation are presented. The same simulation period is chosen as for the comfort controller, that period being from 01-01-2022 to 01-01-2023, as this allows for a comparison of the two controllers. The system is still simulated with the model equations produced in Section 2.3 and utilises the same observer. The weather and economic data is retrieved from the same providers to ensure consistency in the two simulations. The full list of parameters used is seen in Appendix C. Figure 4.7 shows the response of the output to be controlled, i.e. the room temperature alongside its reference, the input, and ambient temperature.

It can be seen on Figure 4.7 that the behaviour of the controller from July to September is similar to Figure 3.7. This is to be expected since the system is unable to cool the room, why it applies the smallest input it can in this period. The interesting parts of Figure 4.7 are the areas where it diverges from the comfort controller. Specifically, major divergence occurs at the start of March and almost all of December, where the lowest room temperature is experienced. These time periods are also some of the most expensive price-wise, as seen in Figure 4.12. The hour with the most expensive electricity price is experienced at the start of September, but the MPC does not respond thereto per the problems regarding large ambient temperatures. Notice, that the room temperature rarely reaches the reference. This is likely due to it not being financially viable to store excess energy in the room with the chosen design parameters.

Even if the controller does not actively charge beyond the reference, it does indeed respond to the electricity price as expected. Examining Figure 4.8 it is seen that the controller tends to increase the room temperature during low-cost periods while letting it decrease during high-cost periods. Charging of the room requires larger inputs which is directly tied to the electricity price as per the differential temperature measurement produced in Section 3.3.

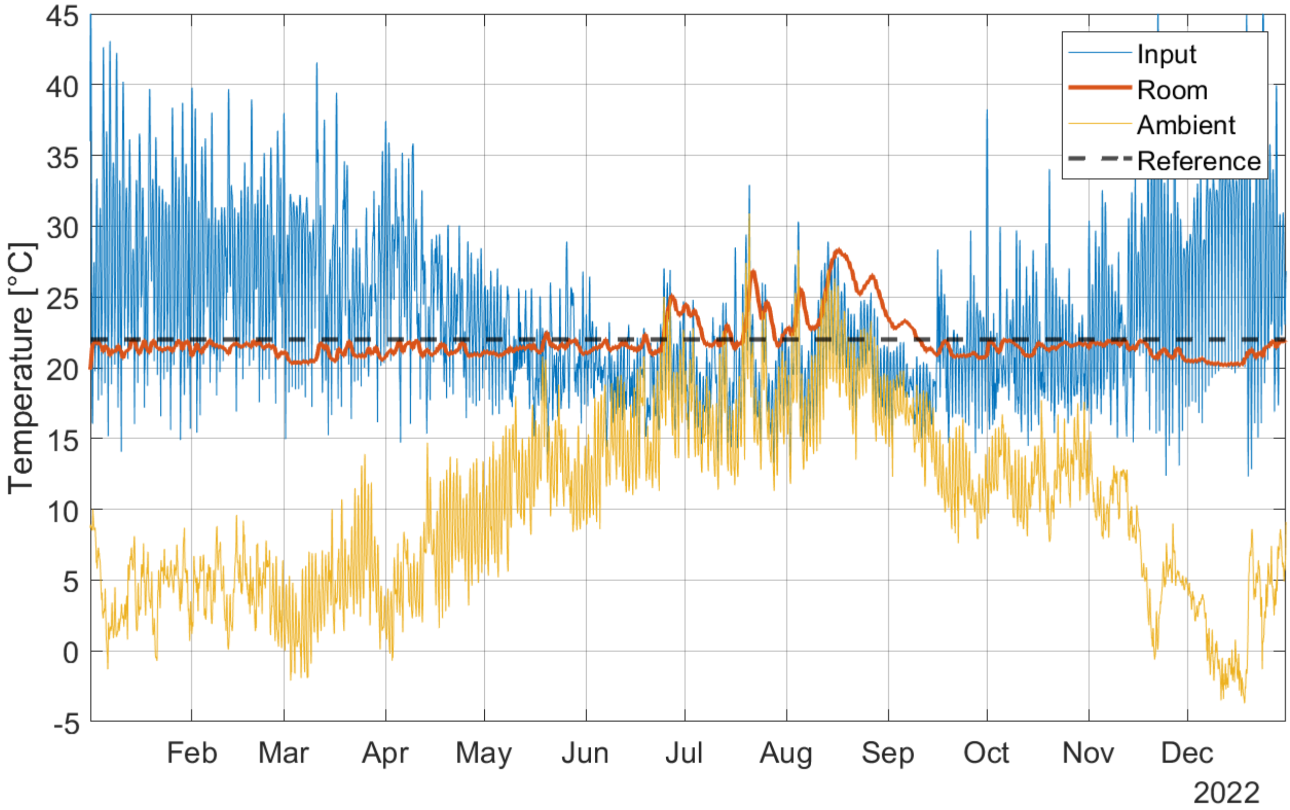

Figure 4.7: Simulation of the economic cost controller output response to a reference of 22 °C and ambient temperature data from 1st of January 2022 to 1st of January 2023. The ambient temperature data is retrieved from the Open-Meteo weather service and used as disturbance in the simulated model. Likewise, the economical data is sourced from Energi Data Service and utilised in the controller weights.

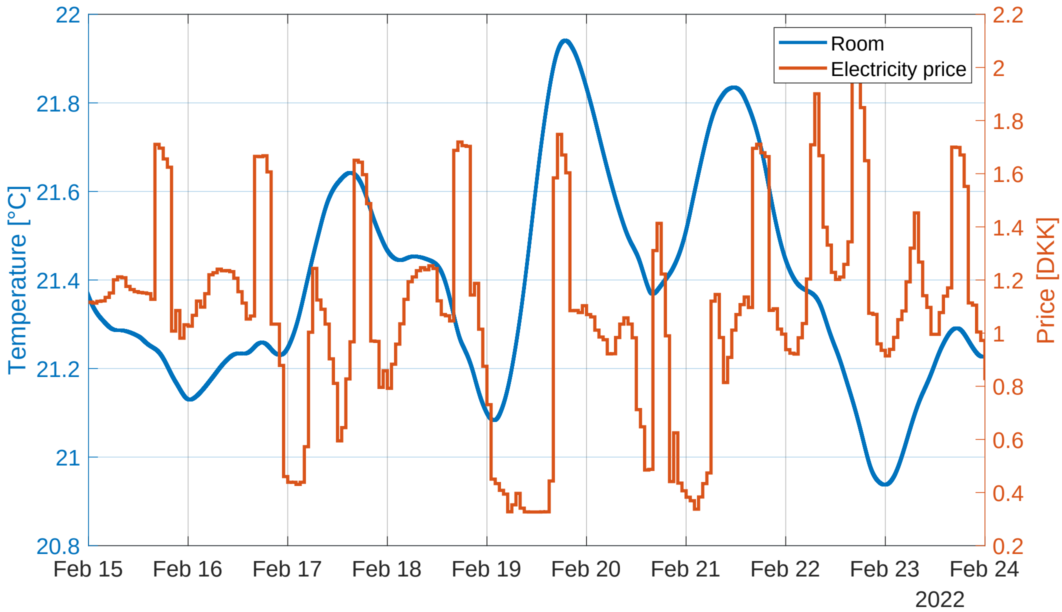

Figure 4.8: Zoomed in view the of a period of the room temperature response in Figure 4.7 alongside the electricity price in that period.

Figure 4.9 shows the same zoomed in view as seen in Figure 3.8, in which the initial part of the response is seen. Precisely like the comfort controller, the initial room temperature value is 20 °C with a reference of 22 °C. As mentioned earlier, the reference is rarely reached and the start of the simulation is not such an occasion. Obviously, this means no overshoot is present unlike Figure 3.8. Similarly, no true steady state region exists due to the effects of the electricity price disturbance.

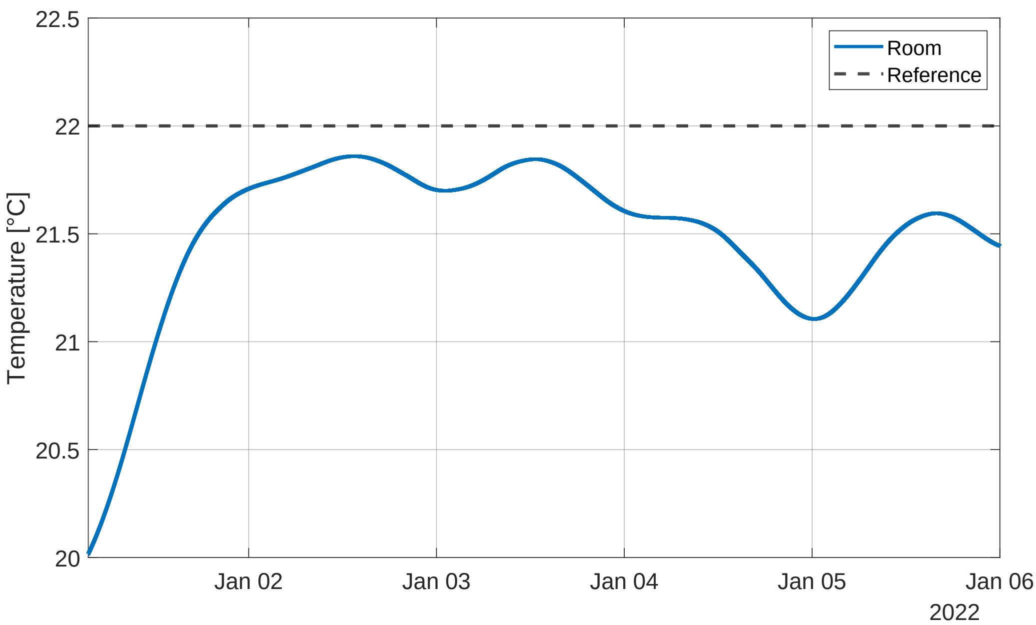

Figure 4.9: Initial part of the simulated output response in Figure 4.7.

The absence of a steady state truly becomes evident when examining the same time period as Figure 3.9, which showed a almost constant tracking error of less than 0.1 °C. As can be seen on Figure 4.10 there are some variations which is largely dominated by the price fluctuations seen in Figure 4.12. However, this is not a problem but the intended way for the controller to operate.

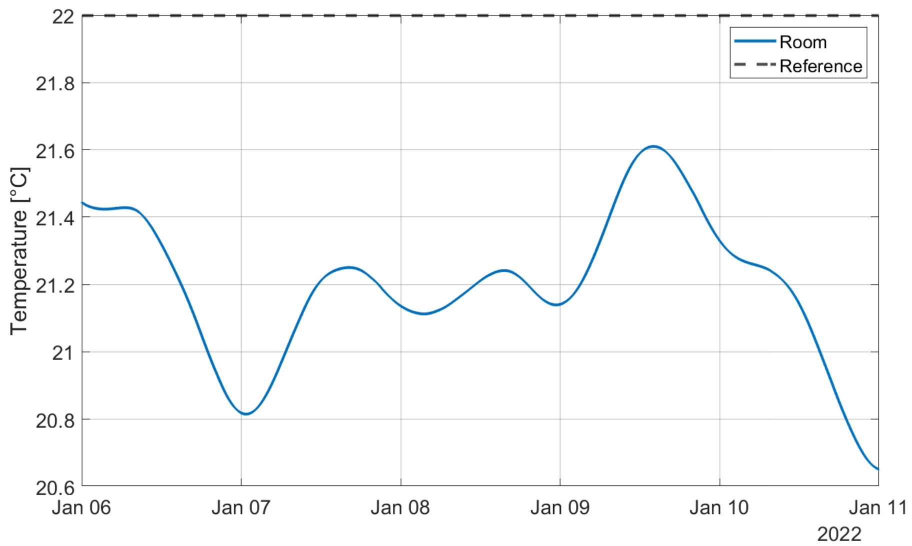

Figure 4.10: Zoomed in view of the simulated room temperature in Figure 4.7 at the same time span as for the comfort controller seen in Figure 3.9.

Figure 4.11 also utilises the same time period as Figure 3.10, where it can be seen that the input behaves similarly for the majority of the 1st of January. The difference is that the input temperature falls below the room temperature faster in Figure 4.11 but does not go below 15 °C. This reduced initial heating period is the reason for not reaching the reference value. Again, this is simply how the controller is designed and it appears to be working as intended.

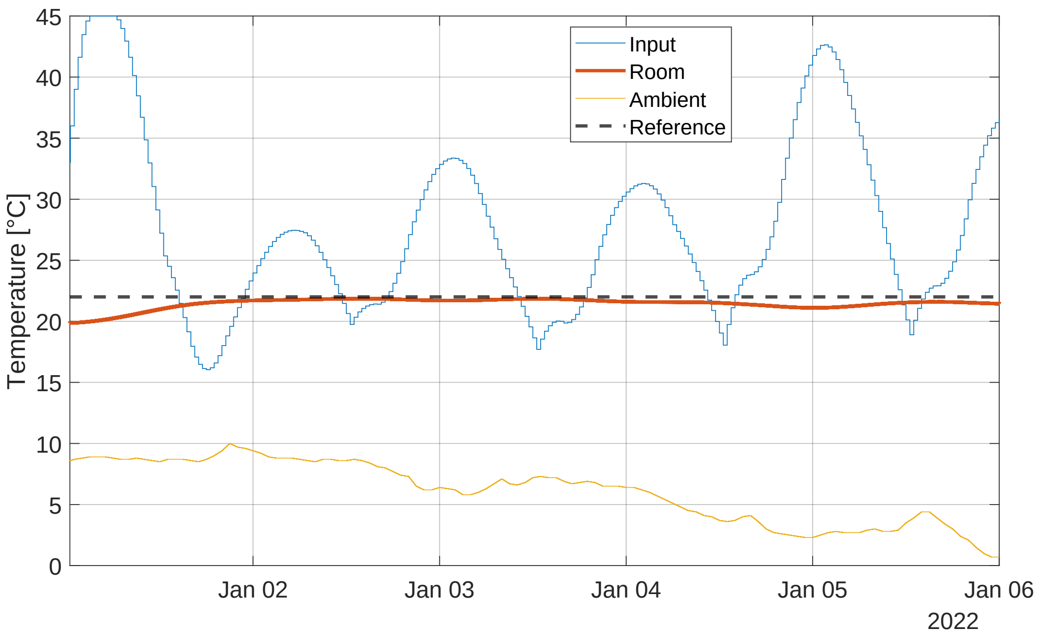

Figure 4.11: The beginning of the output response of the simulated room temperature in Figure 4.7.

Previously in this chapter, a power consumption measurement was introduced in the state space. The product of this power consumption and the electricity price for the simulated period is calculated and shown in Figure 4.12 alongside the electricity price. As can be seen, the power usage is generally less in the summer and more in the winter. Furthermore, the cumulative sum of total simulated price is 6124.37 DKK.

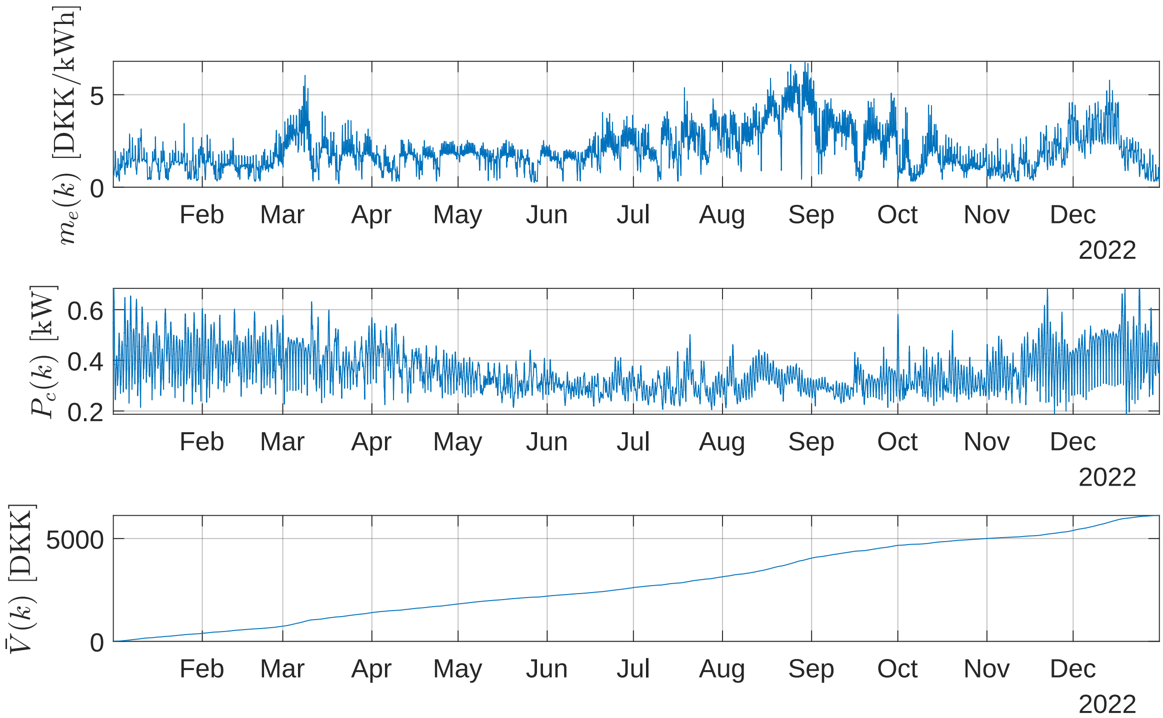

Figure 4.12: Electricity price, simulated power usage, and calculated total price spanning the year 2022. The signal $m_e(k)$ is the electricity price, $P_c(k)$ is the simulated power usage, and $\bar{V}(k)$ is the cumulative sum of the simulated total price.

Comparing the cumulative total price of this controller at 6124.37 DKK to that of the comfort controller at 6463.56 DKK for the same period, it is seen that the economic controller does indeed reduce overall cost of operation. Percentagewise, this reduction equals 5.25 %.

Remark, however that the controller runs and calculates a power consumption for the hot ambient temperature region from July to September as seen in Figure 4.7. In this region, the economic and comfort controller act alike as they both reduce the input to the minimum, just barely above ambient temperature. In reality, the controller is turned off in times like these, and this should also be accounted for in the total price calculation. To do this, every point in which the ambient temperature is above 15 °C is removed, and the price is re-calculated. This amounts to 4239.7 DKK.

In general, it is seen that the controller is capable of utilising the known electricity prices to alter its input, such that the overall cost of operation is reduced when compared to the baseline comfort controller. The controller response shows that the MPC tends to heat the room in low-cost periods and avoids doing so in high-cost periods. On the contrary, the economic MPC does not provide a room temperature tracking error as low as the comfort controller, which is also to be expected. However, the economic MPC is designed to stay within some comfort region, which it does uphold. To further explore the response of the economic MPC, it is further compared to the baseline comfort controller.

4.5 Comparison to Comfort Controller

The comfort controller designed in Chapter 3, merely minimises the room temperature tracking error at all costs. Meanwhile, the economic controller designed in this chapter considers the same objective, while also minimising the cost of operation. This is two conflicting objectives, which certainly gives a poorer tracking error. Nonetheless, it is of interest to compare the response of these two controllers directly, in order to see similarities and differences.

The simulation procedure presented in Section 4.4 is repeated for the economic MPC and the comfort MPC each year from 2018 to 2023. The minimum simulated temperature and total cost with and without points where the ambient temperature is above 15 °C is seen in Table 4.1.

Table 4.1: Simulated minimum temperature and cost for both the comfort and economic controllers for the periods between 2018 and 2023. Every period starts at the 1$^{\text{th}}$ of January and ends at 1$^{\text{th}}$ of January in the corresponding years. The masked cost indicates the total cost of operation, where points with ambient temperatures greater than 15 °C are omitted.

Period	Controller	Min. temp.	Total cost	Masked cost	Reduction
2022-2023	Comfort	21.9 °C	6465.1 DKK	4571.7 DKK	-
	Economic	20.2 °C	6124.4 DKK	4239.7 DKK	5.3 % / 7.3 %
2021-2022	Comfort	21.8 °C	3440.1 DKK	2766.8 DKK	-
	Economic	20.2 °C	3282.0 DKK	2611.2 DKK	4.6 % / 5.6 %
2020-2021	Comfort	21.8 °C	1817.1 DKK	1443.5 DKK	-
	Economic	21.3 °C	1783.2 DKK	1409.1 DKK	1.9 % / 2.2 %
2019-2020	Comfort	21.8 °C	2179.5 DKK	1778.7 DKK	-
	Economic	21.4 °C	2140.7 DKK	1739.7 DKK	1.8 % / 2.1 %
2018-2019	Comfort	21.8 °C	2335.1 DKK	1706.4 DKK	-
	Economic	21.4 °C	2294.3 DKK	1665.8 DKK	1.7 % / 2.3 %

It appears that the chosen design parameters do not properly account for the general change in electricity price from 2018 to 2023, as the minimum temperature for some years is close to the reference. This is likely because the parameters were tuned specifically for 2022-2023, a year in which the electricity prices were generally greater than the previous years. This means that the fitted electricity price weight factor β_Q is smaller to lessen the impact of these large price regions, in turn causing even less impact of lower price regions. Table 4.1 indicates this problem, but it is further showcased by a plot of the room temperature in the yearly period beginning in 2019 and 2022, as seen in Figure 4.13.

As showcased, the economic MPC approaches the comfort MPC for lower energy prices, which is undesirable. As the MPC is allowed to vary the room temperature within the chosen comfort region, it should do so to reduce the overall cost of operation. This could be fixed by raising the electricity price weight β_Q, which would result in the comfort region being exceeded during high cost periods, such as parts of the period from 2022-2023. This was the original problem hypothesised in Section 4.3, which was why the electricity price cap was implemented. As the choice of β_Q in conjunction with the capping scheme has not provided adequate results for all examined price periods, a new scheme should be considered. This is further discussed in Section 5.2.

In conclusion, the MPC designed in this chapter does, in simulation, optimise the end-consumer cost of the modelled residential heating system, thereby indicating the success of such a design which answers the problem formulation in Section 1.2. However, lack of real life data makes it difficult to judge the design, as unforeseen dynamics or problems may occur during implementation, lessening or counteracting the cost reductions shown in simulation. This topic, amongst others, is discussed in the following chapter.

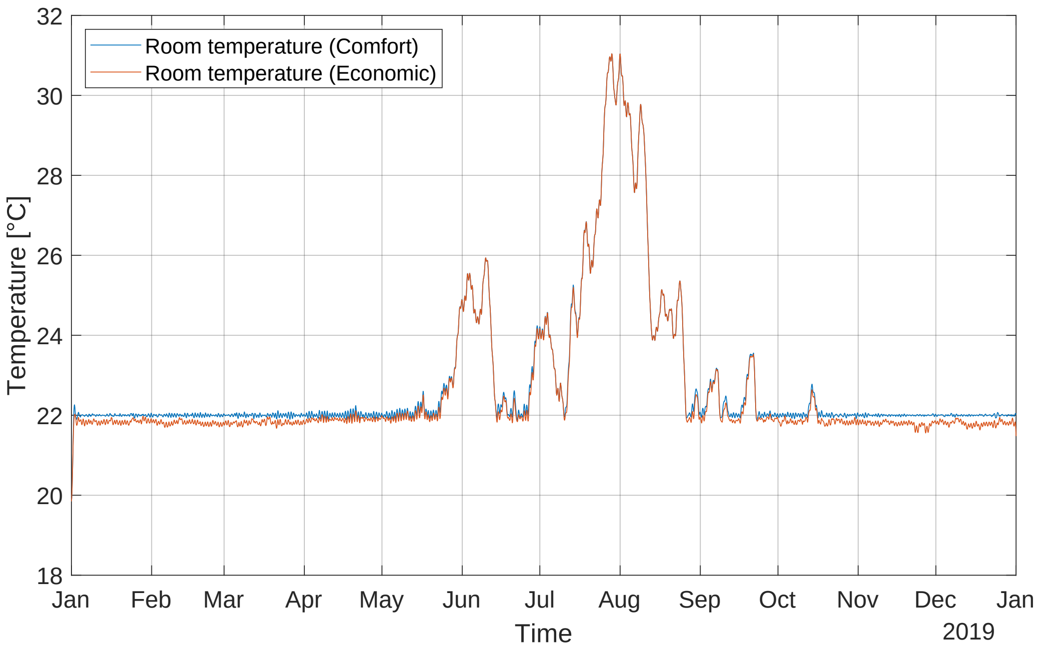

(a) Response in the period 2019-2020.

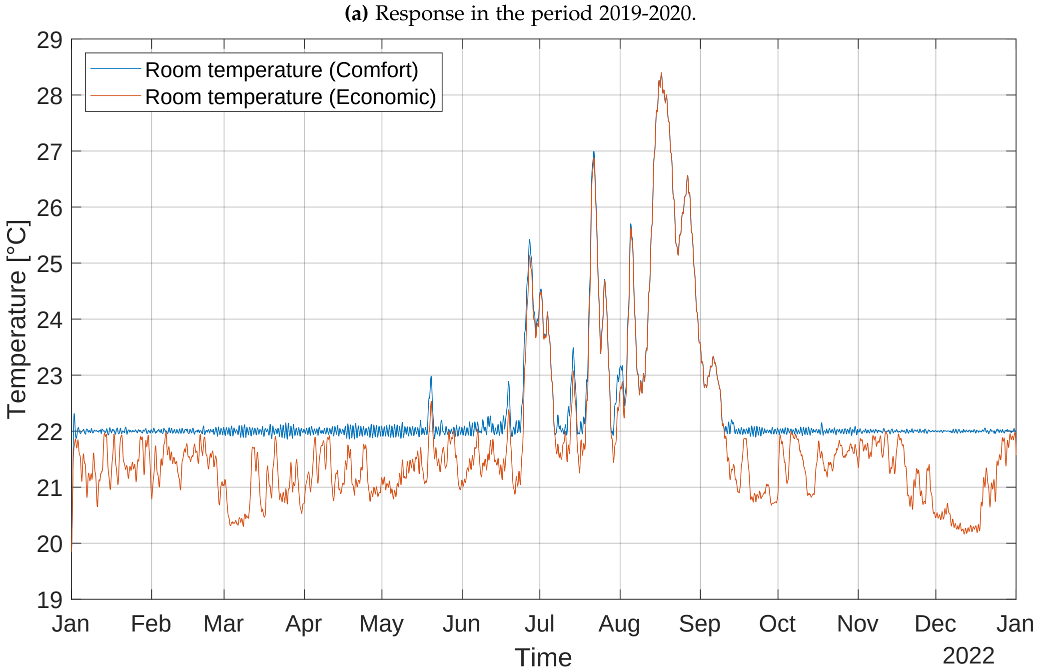

(b) Response in the period 2022-2023.

Figure 4.13: Simulated room temperature response for two separate periods for both the comfort and economic MPC. Each period begins on the 1st of January and ends on the 1st of January.

5 Discussion

Although an economic MPC capable of shifting power consumption from high-cost periods to low-cost periods have been proven, in theory, to function, there is still room for improvement and features which have not been approached due to either the limited project scope or time constraints. Furthermore, the implications of some design choices have not been fully explored. These new concepts are discussed in this chapter.

5.1 Modelling of Disturbances and Stability

Both the weather and electricity prices are considered disturbances in the system, and are, in view of the system, entirely stochastic as they are unmodelled and unbounded. While the disturbances are limitless in theory, it is trivial to establish realistic and conservative limits, such that the disturbances always lie within. In Section 4.3.1, the electricity price bounds were established by forcing the price outside some region to be within. These bounds can be used to establish a worst-case scenario, which likely occurs when the electricity price is high while the ambient temperature is low.

A theoretical worst case scenario can be used to verify that the region of comfort is never exceeded. This worst case scenario is most likely approached using Monte Carlo simulation as traditional stability analysis fails in this regard. Specifically, infinite horizon MPC requires a constant tracking error weight $Q(i) = Q$ preventing the price weighting scheme defined in Section 3.3 [11]. Furthermore, the problem is not suited for infinite horizon methods because of the nature of both weather and electricity price forecasts. This may be mitigated by the construction of conservative and realistic models for both disturbances, but this is not a trivial task.

Even if infinite horizon stability cannot be guaranteed, these models can be used in the MPC to predict further into the future. Specifically, an electricity price model is favourable as the 12 to 35 hour periods of available data is currently the limiting factor in terms of the prediction horizon $\bar{H}_p$. An electricity price model would allow for $\bar{H}_p$ to reach the limit of the weather service that being 7 days into the future, but this comes with the drawback of possibly being incorrect, which is quite likely as the electricity price has shown to be quite volatile. An incorrectly predicted price may cause the MPC to react poorly, e.g. charging prior to a modelled electricity peak only to discover that none ever appeared once the according forecast is published.

5.2 Accounting for Electricity Price Tendencies

In Section 4.3.1, limits were placed on the electricity price to prevent extremities from causing an unacceptable amount of discomfort. After choosing these limits, the electricity price weight factor β_Q was tuned to give satisfactory results for the period 01-01-2022 to 01-01-2023. As was shown in Section 4.4 the MPC response for a simulation of previous years proved to be inadequate as the controller did not properly react to to the dynamics of the electricity price, since it was generally lower. In fact, the economic MPC design presented in this report does not take into account the evolution of the electricity price average which occurs due to the changing electricity production landscape, inflation and other external factors. For this reason, the MPC is not future proof.

One way of addressing this issue was explained in Section 4.3.1, in which either the mean or entirety of an electricity price period is normalised to some value. Naively, the scheme was only used for one period at a time, which resulted in radically changing dynamics due to periods with either large peaks, no peaks, low means or high means. The issue of normalising wrt. a single period is further explained in Section 4.3.1. Normalising wrt. multiple periods lessens the effect of such extremities and is therefore a viable option. The window of normalisation must not be too small, as this presents the same issues outlined in Section 4.3.1 and cannot be too large as that prevents capturing the general tendency of the electricity price evolution. This general tendency is exemplified in Figure 5.1, which shows the moving average of the electricity price from 2018 to 2023 for a window size of one hour (akin to Section 4.3.1), 24 hours and 30 days. Remark that the moving average at sample k is the average of sample k and the last $w-1$ samples where w is the window length.

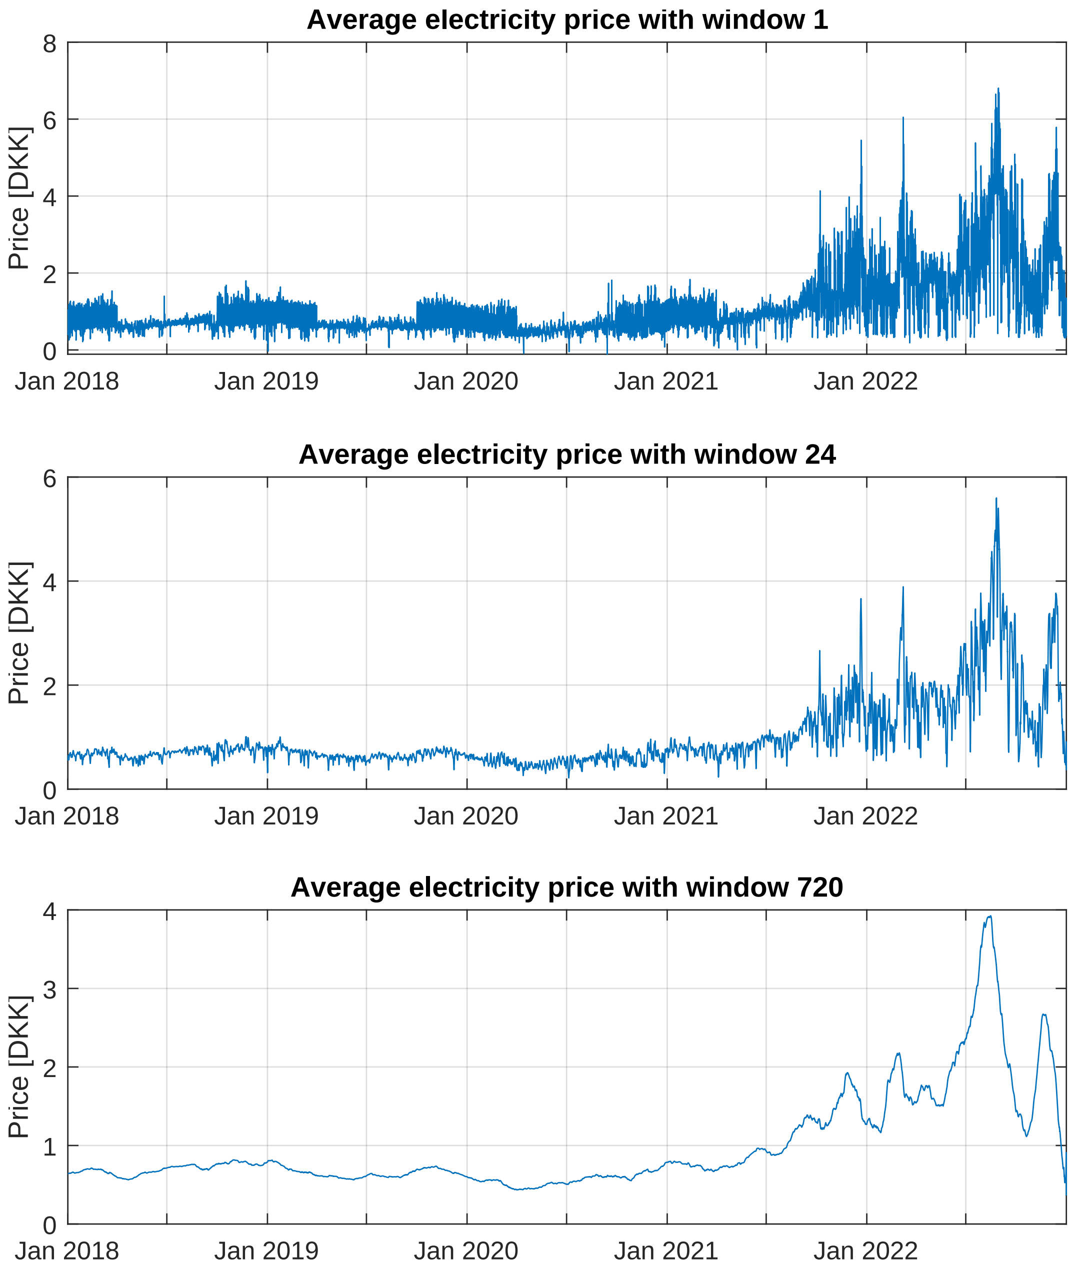

Figure 5.1: Moving average of the electricity price for various windows. The data is retrieved from Energi Data Service and pertains to the period 01-01-2018 to 01-01-2023.

When the price normalisation is performed wrt. the moving average price signal, the dynamics between forecast price periods does not suddenly change. As such, the price dynamics are kept consistent while accounting for the change in mean. Once the change in mean has been accounted for, for example by subtracting the moving average price mean from the period, the price weight factor β_Q can be tuned. Tuning using the entire set of data in Figure 5.1 is not realistic due to computation times, but as per the goal of price normalisation, that should not matter. Since the price is normalised to some degree, the MPC will most likely perform better in price periods it is not tuned upon, thereby making the MPC more future proof. While a normalisation using multiple price periods is not difficult to implement, it was not done due to time constraints as new simulations would be required for the results in Section 4.4.

5.3 Modelling of the Heat Pump

As explained in Section 1.1, the heat pump is responsible for drawing energy from the ground source heat pipes and transferring it to the water in the FH water cycle. In reality, the compressor of the heat pump is controlled in such a way, that the water supplied to the FH tends to the reference. This reference is then controlled by the comfort and economic controllers. It is furthermore assumed, that the supply water temperature is equal to the reference the MPC sets, which it is not as that would require a perfect controller on the heat pump. To circumvent this, constraints were placed on the input in the MPC, which limits the range and rate of change to match the expected performance of the heat pump controller. This was all done to prevent modelling of the heat pump, as it was deemed outside the scope of the project, but some advantages could indeed be gained in the implementation of the MPC by making such a model.

Modelling of the heat pump could be done using first principles or using a grey-box method similarly to that in Section 2.3. Incorporation of this model into the MPC, could potentially allow for even greater optimisation of the cost of operation at the cost of increased complexity. For example, it is known that heat pumps are more efficient at certain operating temperatures, which could be taken into account using such a model. Furthermore, the model would allow the constraints, which currently are set rather arbitrarily based on expected heat pump controller response, to be based on physical limitations of the heat pump such as the temperature, pressure, flow range, and rate of change. Therefore, a more thorough model is required in order to establish true constraints to the physical system. This thorough model is not a trivial task to construct, since the heat pump is a complex system with potential subsystem that need modelling and not just a temperature provider, which it is currently treated as. Because of these reasons a thorough model is not developed in this report, but is nevertheless a point of improvement, which could be investigated in future work on this project.

5.4 Controller Implementation

The designed controllers have two modes, an offline and an online mode. The offline mode is used during simulation and utilises historical data for both the temperature and electricity price. The online mode is designed to be run in real-time on a machine mounted in the provided residence and utilises forecast data from the same APIs as the historical data. In the online mode, the measurement is gathered from the actual system in the provided residence, which is not a problem since the existing soft- and hardware supports MATLAB® in which the controller is written. The offline mode uses the model for its measurement and is advanced alongside the observer. The comfort controller was indeed implemented in the provided residence, however no results will be shown in this report. This is due to a detrimental error present in the observer implementation which was found relatively late in the time span of the project, causing a unexpected and sub-par controller response. Although the error is fixed, due to the time constraints only a small amount of actual data can be gathered. This is compounded by the comparison of the two controllers, since they would have to run consecutively but with relatively similar prerequisites for a fair comparison.

To accurately compare the controllers, the circumstances of each run should be similar, those being the ambient temperature, electricity price and other unused but possibly significant factors such as residence behaviour, wind etc. To do this, similar days could be categorised, which may require a complex scheme and an indeterminably large amount of data. The advantage of such a scheme would be the possibility of investigating the cost reduction between the two designed controllers and the pre-existing controller. Furthermore, the advantages of having the controllers running in the provided residence, is that the accuracy of the model is tested and the addition of unmodelled dynamics such as human behaviour is investigated.

5.5 Controller Comparison and Conflicting Objectives

The cost of the economic controllers will always be less than or equal to that of the comfort controller. This is due to the conflicting costs of the economic controller pulling the room temperature in different directions. In essence, the comfort controller is an extreme of the economic controller, in which the weight on the electricity price tends to 0. Because of these conflicting objectives, a direct comparison between the two controllers can be difficult.

As mentioned in Section 4.5, the chosen electricity price weight β_Q, along with the other chosen design parameters, have resulted in a controller which appears to be too conservative in some periods. With the chosen design parameters, it is, however, likely that the changes in room temperature between 20 °C and 22 °C seen in the results might be unnoticeable for the residents. If it is noticeable, a new comfort region may be chosen, and the design procedure in Section 4.3.2 allows for a relatively quick redesign at the cost of saving less money on electricity. Be that as it may, a model predictive controller method which utilises both the predicted price and weather data and a design procedure has been constructed, as was the goal of the project as explained in Section 1.2. This controller has been shown to, at least in theory, optimise for the end-consumers comfort needs while optimising the cost of operation to some degree.

6 Conclusion

The large economical costs of residential heating can be reduced by shifting the consumption of individual houses to periods with low electricity prices. This can be achieved by a model predictive controller (MPC), the design of which is investigated in this report.

A model is designed of a residence provided by the company Bitzer Electronics, currently occupied by a single family. The provided residence is a two-story building with multiple rooms and a central heating system consisting of a heat pump which transfers energy from ground source heat pipes to the floor heating circuit. The model is constructed as an resistor-capacitor (RC) equivalent and acts purely as an input-output model. The model is fitted to data acquired from sensors in the provided residence by a subspace method resulting in a discrete state space model. The RC equivalent is found by an iterative design process and has three pipe segments, room temperature, and floor temperature as states with supply temperature as controlled input and ambient temperature as disturbance. The model is a switched input model encompassing a hot water production (HWP) and a floor heating (FH) subsystem. The switch is determined by a non-causal moving average filtering of the data as the fit is performed, after which the HWP subsystem is removed from the state space. The model is validated against different sets of data, and generally has satisfactory input-output characteristics. For most of the provided data, the mean squared error (MSE) between simulated and measured output, i.e. the return temperature, is less than $6.5\,°\mathrm{C}^2$.

A baseline comfort MPC is constructed with the goal of guiding the room temperature to be as close as possible to some reference. The mathematical theory behind MPC in terms of a quadratic programming (QP) problem is presented. The MPC output is a supply temperature reference to a pre-existing heat pump controller, which is constrained to match the response thereof by rate of change and range dependent on the ambient temperature. The MPC optimises its output wrt. a designed quadratic cost function, the constraints, and the system states estimated by a Luenberger observer. The controller is downsampled to decrease computation time for large prediction horizons. The simulated fitted model is controlled by the comfort MPC which results in a room temperature deviating less than 0.1 °C from the reference after converging.

The baseline controller counterpart is the economic MPC with the goal of minimising the total cost of operation while keeping the room temperature within a chosen region of comfort. A differential temperature measurement is constructed, empirically proportional to the power consumption. Once an electricity price signal extended with appropriate Danish tariffs is included in the MPC tracking error weight, the cost of operation is produced. The new measurement introduces a direct state space term which is incorporated into the QP problem. The prediction horizon is made dynamic lasting between 12 and 35 hours corresponding to the Danish electricity price forecast schedule. Compatibility of the QP problem is ensured by cropping system matrices and indexing signals accordingly. A tuning procedure is proposed in which the parameters are systematically altered until the simulated room temperature response never exceeds the chosen region of comfort. Under identical conditions as the comfort controller simulations, the simulated model is controlled by the economic controller for a simulation period of one year. Generally, the room temperature is below reference but within the chosen comfort region and reacts to the electricity price signal by heating more in low-cost periods and less in high-cost periods. The total cost of operation for the economic controller is 7.3 % less than that of the comfort controller in the simulated period 2022 to 2023. Both controllers are further simulated in other periods one year in length, but the cost reduction between the economic and comfort MPC reaches as low as 2.1 %. This is likely due to the tuned price weight being too conservative, as it is fitted to the highly volatile and generally costly electricity price signal of the period 2022 to 2023.

A way of ensuring more consistent results across different time periods is discussed, and it is theorised that normalisation of the electricity price mean wrt. a month-long moving average signal may suffice. Furthermore, the electricity price is limited to a range determined empirically, such that extreme electricity prices do not cause a violation of the comfort region. While modelling of disturbances do not help ensuring stability, they can be used to increase the theoretical upper limit of the dynamic prediction horizon. Furthermore, a heat pump model could improve the MPC response and aid in the construction of realistic constraints. While the differing objectives of the two controllers make them difficult to compare fairly, it is still shown in this report, how an MPC can be designed for a central heating system in a residential building. This economic MPC has proven, at least in simulation, capable of utilising the residential thermal mass in conjunction with the forecast ambient temperature and electricity price, to overall reduce the cost of operation while keeping the indoor temperature comfortable. Future work on this economic MPC includes electricity price normalisation and implementation on real hardware, the latter of which is within reach since the controller is designed with hardware implementation in mind.

Bibliography

[1] Energistyrelsen, *Energistatistik 2021*, Dec. 2022.

[2] Danmarks Statistik. “Boligbestanden”. (2020), [Online]. Available: `https://www.dst.dk/da/Statistik/emner/borgere/boligforhold/boligbestanden`.

[3] Energinet. “Energinet justerer eltarifferne for 2023”. (Sep. 2022), [Online]. Available: `https://energinet.dk/om-nyheder/nyheder/2022/09/06/energinet-justerer-tarif/`.

[4] O. Fletcher. “European energy bills hit record despite government support”. (Nov. 2022), [Online]. Available: `https://www.bloomberg.com/news/articles/2022-11-07/european-energy-bills-hit-record-despite-government-support?leadSource=uverify%5C%20wall`.

[5] T. Hilliard, M. Kavgic, and L. Swan, “Model predictive control for commercial buildings: Trends and opportunities”, eng, *Advances in building energy research*, vol. 10, no. 2, pp. 172–190, 2016, ISSN: 1751-2549.

[6] S. Thorsteinsson, H. Lund Stærmose, and J. D. Bendtsen, “Transient analysis of individual return temperatures in hydronic floor heating systems”, *CLIMA 2022 conference*, May 2022. DOI: `10.34641/clima.2022.140`. [Online]. Available: `https://proceedings.open.tudelft.nl/clima2022/article/view/140`.

[7] H. Hersbach, B. Bell, P. Berrisford, *et al.*, “Era5 hourly data on single levels from 1959 to present.”, Copernicus Climate Change Service (C3S) Climate Data Store (CDS), Tech. Rep., 2018.

[8] Muñoz Sabater, J. and Copernicus Climate Change Service, “Era5-land hourly data from 2001 to present”, Tech. Rep., 2019. DOI: `10.24381/CDS.E2161BAC`. [Online]. Available: `https://cds.climate.copernicus.eu/doi/10.24381/cds.e2161bac`.

[9] S. S., R. M., L. M. P., *et al.*, “Cerra sub-daily regional reanalysis data for europe on single levels from 1984 to present”, Tech. Rep., 2021. DOI: `10.24381/CDS.622A565A`. [Online]. Available: `https://cds.climate.copernicus.eu/doi/10.24381/cds.622a565a`.

[10] L. Ljung, *System identification: Theory for the user* (Prentice Hall information and system sciences series), en, 2nd ed. Philadelphia, PA: Prentice Hall, Dec. 1998.

[11] J. M. Maciejowski, *Predictive control: with Constraints.* eng. Harlow: Prentice Hall, 2002, ISBN: 9780201398236.

[12] N1. “Priser og vilkår”. (May 2023), [Online]. Available: `https://n1.dk/priser-og-vilkaar`.

Appendix

A Controller Algorithms

This appendix details the main algorithms which are implemented and used within the comfort controller from Chapter 3 and economic controller from Chapter 4. For an actual implementation of the algorithms, examine the code referenced in Appendix B. Both controllers utilise the lifting of the system equations and constraints presented in Section 3.1, yet it is only shown for the economic controller, as this is of most interest. The lifting for the comfort controller is similar and simply lacks some of the variables pertaining only to the economic controller. The implemented algorithms for the comfort and economic MPC are seen in Algorithm 3 and Algorithm 4 respectively. The procedures used to lift the system and constraints are seen in Algorithm 5 and Algorithm 6 respectively.

The notation $I_{n\times m}$ is used to indicate an identity matrix with n rows and m columns. Likewise, $\mathbb{1}$ is a matrix consisting of all ones while O is a matrix consisting of all zeroes. The symbol ($\otimes$) is a *Kronecker product* which produces the block matrix consisting of the right-side matrix in each entry multiplied with the corresponding left-side entry. The symbol ($\odot$) is an *element-wise product* between a matrix and a scalar which produces the matrix with every entry multiplied by the scalar.

Algorithm 3 An implementation of the Comfort MPC with downsampling and constraints described in Chapter 3.

Require: $1 \leq H_p, \quad 1 \leq H_u \leq H_p$
Require: $\kappa \in \mathbb{Z}, \quad 1 \leq \kappa$ ▷ κ must be a positive integer
Require: $\mathcal{D} \in \mathbb{R}^{\kappa H_p}, \quad \mathcal{T} \in \mathbb{R}^{\kappa H_p}$
$k = 0$
$k_{mpc} = 1$
$\hat{x} = \hat{x}_0 \in \mathbb{R}^m$
$u = u_0 \in \mathbb{R}^1$
repeat
  $k = k + 1$
  $y =$ GET_MEASUREMENT ▷ Measurement is performed
  if $(k-1) \bmod \kappa = 0$ **then** ▷ Controller loop is downsampled by κ
    $k_{mpc} = k$
    $(\Psi, \Upsilon, \Theta, \Xi) =$ LIFT_SYSTEM(H_p, H_u)
    $\mathcal{H} = \Theta^T \mathcal{Q} \Theta + \mathcal{R}$
    $\mathcal{E} = \mathcal{T} - \Psi x - \Upsilon u - \Xi d$
    $\mathcal{G} = 2\Theta^T \mathcal{Q} \mathcal{E}$
    $(A_{con}, b_{con}) =$ LIFT_CONSTRAINTS$(H_p, H_u, u, \mathcal{D}, \mathcal{T}, \Upsilon, \Theta, \Xi)$
    $\Delta\mathcal{U} =$ QUADPROG$(2\mathcal{H},\ -G^T,\ A_{con},\ b_{con})$ ▷ Constrained QP problem solved
    $\mathcal{U} = \begin{bmatrix} u + \Delta\mathcal{U}(1) & u + \Delta\mathcal{U}(2) & \dots & u + \Delta\mathcal{U}(\kappa H_u) \end{bmatrix}$
  end if
  $\hat{y} = C_y \hat{x} + D_y u + D_{dy} d$ ▷ Estimates are produced
  $\hat{x} = A\hat{x} + L(y - \hat{y}) + Bu + B_d d$
  $u = \mathcal{U}(k - k_{mpc} + 1)$
until $k > k_{final}$

Algorithm 4 An implementation of the Economic MPC with downsampling, constraints and dynamic prediction horizon described in Chapter 4.

Require: $A \in \mathbb{R}^{m\times m}, B \in \mathbb{R}^{m\times n}, B_d \in \mathbb{R}^{m\times j}$ ▷ System equations and their dimensions
Require: $C_z \in \mathbb{R}^{p\times m}, D_z \in \mathbb{R}^{p\times n}, D_{dz} \in \mathbb{R}^{p\times j}$
Require: $\mathcal{D} \in \mathbb{R}^{j_d}, \quad \mathcal{M}_e \in \mathbb{R}^{j_m}, \quad \mathcal{T} \in \mathbb{R}^{j_t}, j_t > \max(j_d, j_m)$
Require: $1 \le H_p, \quad 1 \le H_u \le H_p$
Require: $\kappa \in \mathbb{Z}, \quad 1 \le \kappa \le \min(j_d, j_m, j_t)$ ▷ κ must be integer. Ensure at least one sample point.

$k = 0$
$k_{mpc} = 1$
$\hat{x} = \hat{x}_0 \in \mathbb{R}^n$
$u = u_0 \in \mathbb{R}^1$
repeat
    $k = k + 1$
    $t = $ GET_TIME(k)
    $y = $ GET_MEASUREMENT ▷ Measurement is performed
    if $(k-1) \bmod \kappa = 0$ **then** ▷ Controller loop is downsampled by κ
        $k_{mpc} = k$
        $\bar{H}_p = \left\lfloor \min\left(H_p, \frac{j_d}{\kappa}, \frac{j_m}{\kappa}\right)\right\rfloor$ ▷ Prediction horizon limited to disturbance length
        $\bar{H}_u = \min(H_u, \bar{H}_p)$
        $\bar{d} = \begin{bmatrix}\mathcal{D}(k) & \mathcal{D}(k+\kappa) & \dots & \mathcal{D}(\bar{H}_p \cdot \kappa)\end{bmatrix}^T$ ▷ Disturbance & reference picked accordingly
        $\bar{r} = \begin{bmatrix}\mathcal{T}(k) & \mathcal{T}(k+\kappa) & \dots & \mathcal{T}(\bar{H}_p \cdot \kappa)\end{bmatrix}^T$
        $\bar{m}_e = \begin{bmatrix}\mathcal{M}_e(k) & \mathcal{M}_e(k+\kappa) & \dots & \mathcal{M}_e(\bar{H}_p \cdot \kappa)\end{bmatrix}^T$
        $\mathcal{Q} = $ WEIGHT_PRICE($\bar{m}_e$) ▷ Produced as explained in Section 4.1
        $\bar{\mathcal{Q}} = \mathcal{Q}_{>(1,1),(p\bar{H}_p, n\bar{H}_u)<}$
        $\bar{\mathcal{R}} = \mathcal{R}_{>(1,1),(p\bar{H}_p, n\bar{H}_u)<}$
        $(\bar{\Psi}, \bar{\Upsilon}, \bar{\Theta}, \bar{\Xi}, \bar{\Lambda}, \bar{\Lambda}_D) = $ LIFT_SYSTEM($\bar{H}_p, \bar{H}_u$)
        $\mathcal{H} = (\Theta + \Lambda_D)^T \mathcal{Q} (\Theta + \Lambda_D) + \mathcal{R}$
        $\mathcal{E} = \bar{r} - \bar{\Psi} x - (\Upsilon + \Lambda) u - \Xi \bar{d}$
        $\mathcal{G} = 2(\Theta + \Lambda_D)^T \mathcal{Q} \mathcal{E}$
        $(A_{con}, b_{con}) = $ LIFT_CONSTRAINTS($\bar{H}_p, \bar{H}_u, u, \bar{d}, \bar{r}, \bar{\Upsilon}, \bar{\Theta}, \bar{\Xi}, \bar{\Lambda}, \bar{\Lambda}_D$)
        $\Delta\mathcal{U} = $ QUADPROG($2\mathcal{H}, -G^T, A_{con}, b_{con}$) ▷ Constrained QP problem solved
        $\mathcal{U} = \begin{bmatrix}u + \Delta\mathcal{U}(k) & u + \Delta\mathcal{U}(k+1) & \dots & u + \Delta\mathcal{U}(k + \bar{\kappa} H_p)\end{bmatrix}$
    end if
    $\hat{y} = C_y \hat{x} + D_y u + D_{dy} d$ ▷ Estimates are produced
    $\hat{x} = A\hat{x} + L(y - \hat{y}) + Bu + B_d d$
    $u = \mathcal{U}(k - k_{mpc} + 1)$
until $k > k_{final}$

Algorithm 5 Procedure for lifting the state space equations to a QP problem described in Section 3.1.

Require: $A \in \mathbb{R}^{m\times m}, B \in \mathbb{R}^{m\times n}, B_d \in \mathbb{R}^{m\times j}$ ▷ System equations and their dimensions
Require: $C \in \mathbb{R}^{p\times m}, D \in \mathbb{R}^{p\times n}, D_d \in \mathbb{R}^{p\times j}$
Require: $1 \leq \bar{H}_p, \quad 1 \leq \bar{H}_u \leq \bar{H}_p$

function LIFT_SYSTEM($\bar{H}_p, \bar{H}_u$)

$\mathcal{A} = \begin{bmatrix} A & A^2 & A^3 & \cdots & A^{H_p} \end{bmatrix}^T$

$\mathbf{B}_0 = B$

for $i = 1, \bar{H}_p - 1$ **do** ▷ Recursive lifting of B

$\mathbf{B}_i = \mathbf{B}_{i-1} + A^i \cdot B$

end for

$\mathcal{B} = \begin{bmatrix} \mathbf{B}_0 & \mathbf{B}_1 & \cdots & \mathbf{B}_{H_p-1} \end{bmatrix}^T$

$\mathbf{B}_{\Delta,0} = \mathcal{B}$

for $i = 1, \bar{H}_u - 1$ **do**

$\mathbf{B}_{\Delta,i-1} = \begin{bmatrix} B_{\Delta,1} & B_{\Delta,2} & \cdots & B_{\Delta,\bar{H}_p} \end{bmatrix}^T$ ▷ Each block is named for the next step

$\mathbf{B}_{\Delta,i} = \begin{bmatrix} 0 & B_{\Delta,1} & B_{\Delta,2} & \cdots & B_{\Delta,\bar{H}_p-1} \end{bmatrix}^T$ ▷ Shift by discarding last entry & prepending 0

end for

$\mathcal{B}_{\Delta u} = \begin{bmatrix} \mathbf{B}_{\Delta,0} & \mathbf{B}_{\Delta,1} & \cdots & \mathbf{B}_{\Delta,\bar{H}_u-1} \end{bmatrix}^T$

$\mathbf{X}_0 = \begin{bmatrix} B_d & AB_d & \cdots & A^{H_p-1}B_d \end{bmatrix}^T$ ▷ In reality produced recursively like $\mathbf{B}_i$

for $i = 1, \bar{H}_p - 1$ **do** ▷ Similar shifting is performed

$\mathbf{X}_{i-1} = \begin{bmatrix} X_1 & X_2 & \cdots & X_{H_p} \end{bmatrix}^T$

$\mathbf{X}_i = \begin{bmatrix} 0 & X_1 & X_2 & \cdots & X_{H_p-1} \end{bmatrix}^T$

end for

$\mathcal{X} = \begin{bmatrix} \mathbf{X}_0 & \mathbf{X}_1 & \cdots & \mathbf{X}_{H_p-1} \end{bmatrix}^T$ ▷ The coefficients are collected

$\mathcal{C} = I_{\bar{H}_p \times \bar{H}_p} \otimes C$

$\bar{\Psi} = \mathcal{C}\mathcal{A}$

$\bar{\Upsilon} = \mathcal{C}\mathcal{B}$

$\bar{\Theta} = \mathcal{C}\mathcal{B}_{\Delta u}$

$\bar{\Xi} = \mathcal{C}\mathcal{X}$

$\Lambda = \mathbb{1}_{\bar{H}_p \times 1} \otimes D$

$$\Lambda_D = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 1 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{bmatrix} \otimes D$$ ▷ Remark $\Lambda_D \in \mathbb{R}^{\bar{H}_p \cdot p \times \bar{H}_u \cdot n}$

end function

Algorithm 6 Procedure for lifting the constraints to a QP problem described in Section 3.1.

Require: $n_f \geq 1, \quad n_u \geq 1$ ▷ Positive number of inputs and constraints
Require: $1 \leq \bar{H}_p, \quad 1 \leq \bar{H}_u \leq \bar{H}_p$
Require: $u \in \mathbb{R}^1, \quad d \in \mathbb{R}^{\bar{H}_p}$
function LIFT_CONSTRAINTS($\bar{H}_p, \bar{H}_u, u, d$)
  $F = I_{\bar{H}_u \times \bar{H}_u} \otimes \begin{bmatrix} -1 & 1 & -1 \end{bmatrix}^T$ ▷ Input constraints with minimum & maximum limit
  $f_1 = \mathbb{1}_{\bar{H}_u \times 1} \otimes \begin{bmatrix} \mu_1 & -\mu_2 & \mu_3 \end{bmatrix}^T$
  $f_2 = \left(I_{\bar{H}_u \times \bar{H}_u} \otimes \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T \right) d$ ▷ Disturbance dependent minimum input limit
  $F = \begin{bmatrix} F & f_1 + f_2 \end{bmatrix} = \begin{bmatrix} F_1 & F_2 & \cdots & F_{H_u} & f \end{bmatrix}$
  $W = \begin{bmatrix} I_{\bar{H}_u \times \bar{H}_u} & -I_{\bar{H}_u \times \bar{H}_u} \end{bmatrix}^T$ ▷ Input rate constraints
  $w = \begin{bmatrix} \mathbb{1}_{\bar{H}_u \times 1} & -\mathbb{1}_{\bar{H}_u \times 1} \end{bmatrix}^T \odot (\mu_4 \cdot \kappa)$
  for $i = 1, \bar{H}_u$ **do** ▷ Lifting of input constraint F
    $\mathbf{F_i} = O_{\bar{H}_u \cdot n_f \times n_u}$
    for $j = i, \bar{H}_u$ **do**
      $\mathbf{F_i} = \mathbf{F_i} + F_j$
    end for
  end for
  $\mathcal{F} = \begin{bmatrix} \mathbf{F}_1 & \mathbf{F}_2 & \ldots & \mathbf{F}_{H_u} \end{bmatrix}$
  $A_{con} = \begin{bmatrix} \mathcal{F} & W \end{bmatrix}^T$ ▷ Lifted constraint matrices used in `QUADPROG()`
  $b_{con} = \begin{bmatrix} -\mathbf{F_1} \cdot u - f & w \end{bmatrix}^T$
end function

B Overview of Implemented Code

The economic MPC is implemented in MATLAB® and available on GitHub at the following internet address:

`https://github.com/joulsen/CA10-1033-MPC`

Remark that only one controller is found in this repository, as either the comfort controller or economic controller can be implemented by setting the correct options (see Appendix C). This appendix details what the important files of the directories do. Some files are only used when run in real-time on the machine, denoted *online mode*, while some are run unthrottled for simulation purposes, denoted *offline mode*. At the moment, the main file `mpc_run.m` is not functional due to a redaction of personal information relating to the provided residence.

At the time of writing (commit `a88ab1f`) 15 files exist in the repository. All files, except two trivial, are presented in Table B.1 alongside a description of their functionality.

Table B.1: Files present in the directory containing the implemented MPC and their descriptions.

Filename	Description
`MultiYear.m`	Script used to determine the cost reductions of various yearly periods in Table 4.1 and the responses seen in Figure 4.13.
`get_forecast.m`	Function for retrieving weather forecast in online mode.
`get_forecast_historical.m`	Function for retrieving weather forecast in offline mode.
`get_plant.m`	Function producing the plant model by properly resampling and dissecting the fitted model in `model.mat`.
`get_price.m`	Function for retrieving forecasted electricity prices in online mode.
`get_price_historical.m`	Function for retrieving historical electricity prices as if they were forecast, abiding by the Danish spot price forecast schedule. Used in the controller in offline mode.
`get_price_historical_all.m`	Function for retrieving historical electricity prices with no limits or scheduling. Is not used in either online or offline mode.
`model.mat`	Binary file containing data for the fitted two segment model presented in Section 2.3.
`mpc_base.m`	Function implementing the majority of the MPC controller loop except for the lifting and QP solving step specifically which is done by `mpc_iteration.m`.
`mpc_beta_search.m`	Script used to tune the price factor weight β_Q as seen in Section 4.3.
`mpc_initialize.m`	Function producing the system matrices, downsampled system matrices, Luenberger observer and lifted matrices as seen in Section 3.1.
`mpc_iteration.m`	Function lifting the constraints and solving the QP problem posed in Section 3.1. This function also crops the lifted system equations according to the dynamic prediction horizon presented in Section 4.2
`mpc_run.m`	Main script for running the MPC in either offline or online mode. Utilises most other files in this directory in some capacity.

C Simulation Parameters

This chapter houses the parameters used during simulation of the comfort MPC and economic MPC. The parameters directly relate to the MPC implementation and may be set in any script running the MPC implementation, including but not limited to the file `mpc_run.m` (refer to Appendix B). As described in Appendix B, the same implementation is used for both and this is mainly controlled by setting values in the *option* variable. Parameters and options not noted in this appendix are either never changed or easily determined from the simulation results in Section 3.4 and Section 4.4 such as the simulation dates. Both controller simulations share many parameters and options, most notably the initial state x_0, the initial input u_0 and the Luenberger observer poles P seen below:

$$x_0 = \begin{bmatrix} 35 & 30 & 25 & -2 & 20 \end{bmatrix}^T \tag{C.1a}$$

$$u_0 = 30 \tag{C.1b}$$

$$P = \begin{bmatrix} 0.01 & 0.02 & 0.03 & 0.04 & 0.05 \end{bmatrix}^T \tag{C.1c}$$

Table C.1 and Table C.2 show the parameters and options used during simulation of the comfort controller in Section 3.4. Table C.3 and Table C.4 show the parameters and options used during simulation of the economic controller in Section 4.4.

Table C.1: Parameters used in the simulation of the comfort controller in Section 3.4.

Parameter	Value	Description
H_u	20	Control horizon
H_p	25	Prediction horizon
μ_1	1	Minimum temperature constraint
μ_2	45	Maximum temperature constraint
μ_3	2	Minimum difference between input and ambient temperature constraint
μ_4	0.1	Maximum rate of change. Scaled with κ in implementation
κ	30	Controller downsampling factor

Table C.2: Options used in the simulation of the comfort controller in Section 3.4.

Option	Value	Description
PriceOn	False	Utilise electricity price in the controller

Table C.3: Parameters used in the simulation of the economic controller in Section 4.4.

Parameter	Value	Description
H_u	20	Control horizon
H_p	35	Prediction horizon
μ_1	1	Minimum temperature constraint
μ_2	45	Maximum temperature constraint
μ_3	2	Minimum difference between input and ambient temperature constraint
μ_4	0.1	Maximum rate of change. Scaled with κ in implementation
κ	30	Controller downsampling factor

Table C.4: Options used in the simulation of the economic controller in Section 4.4.

Option	Value	Description
PriceOn	True	Utilise electricity price in the controller
PriceNormalisation	"cap"	Type of price normalisation used (see Section 4.3.1).
NormMin	0.05	Minimum of the capped electricity price.
NormMax	2.2	Maximum of the capped electricity price.