CHALMERS
UNIVERSITY OF TECHNOLOGY

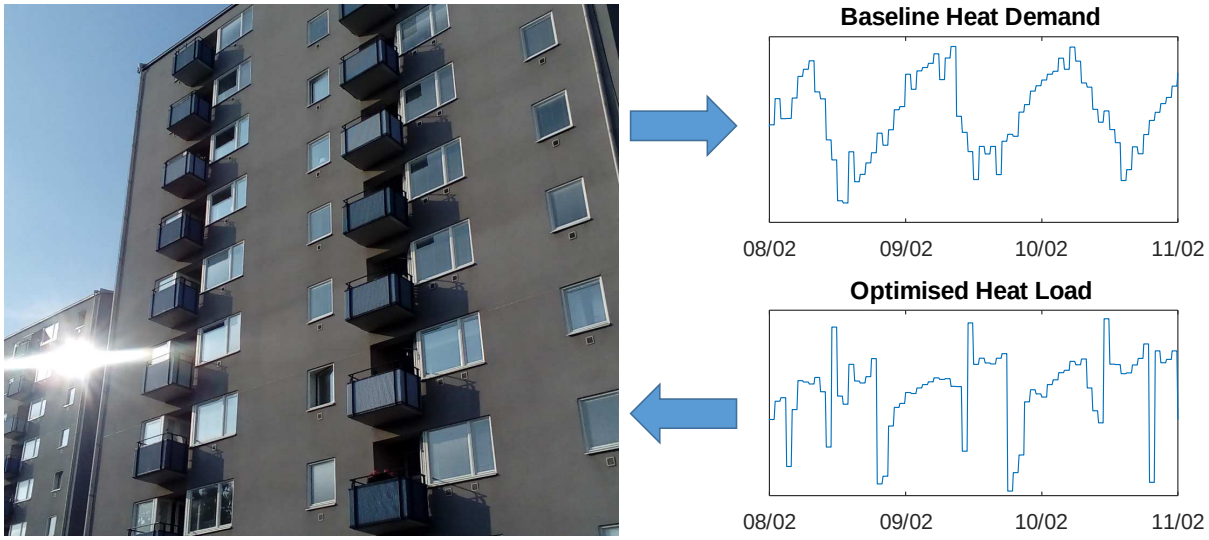

Marginal Price Control of Buildings Utilised as Thermal Energy Storage

Optimising the heating cost of a modelled residential building with respect to the district heating network marginal generation cost

Master thesis in Complex Adaptive Systems

Jens Carlsson

Department of Civil and Environmental Engineering
Division of Building Services Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2016

Master's thesis BOMX02-16-89

Marginal Price Control of Buildings Utilised as Thermal Eneergy Storage

Optimising the heating cost of a modelled residential building with respect to the district heating network marginal generation cost

JENS CARLSSON

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Civil and Environmental Engineering
Division of Building Services Engineering
Chalmers University of Technology
Gothenburg, Sweden 2016

Marginal Price Control of Buildings Utilised as Thermal Energy Storage
Optimising the heating cost of a modelled residential building with respect to the district heating network marginal generation cost
JENS CARLSSON

Supervisor: Johan Kensby, Department of Civil and Environmental Engineering
Examiner: Anders Truschel, Department of Civil and Environmental Engineering

Master's Thesis 2016:89
Department of Civil and Environmental Engineering
Division of Building Service Engineering
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Qualitative illustration of how instantaneous heat load can vary from the nominal heat demand without negatively affecting thermal comfort in a building utilised as thermal energy storage.

Typeset in LaTeX
Gothenburg, Sweden 2016

Marginal Price Control of Buildings Utilised as Thermal Energy Storage
Optimising the heating cost of a modelled residential building with respect to the district heating network marginal generation cost
Jens Carlsson
Department of Civil and Environmental Engineering
Division of Building Services Engineering
Chalmers University of Technology

Abstract

District heating networks often experience significant variations in heat load on an hourly timescale, with one or two noticeable peaks per day. These peaks are associated with increased marginal heat generation cost and environmental impact, and it would therefore be beneficial to reduce the daily variations. It has been previously shown that heavy residential buildings, thanks to their high thermal inertia, can be used as short-term thermal energy storages for the purpose of load shifting. In this thesis project the economic viability of implementing load-shifting control systems was investigated. A dynamical building model was created based on data collected by Göteborg Energi. Control systems were designed that minimised the overall building heating cost under the assumption that district heating prices were proportional to the marginal generation cost. It was found that, when only focusing on temperature stabilisation, a 10% reduction in heat use leads to less than 7% reduced heating costs due to naturally high indoor temperatures coinciding with low marginal costs. When applying a combination of temperature stabilisation and load shifting however, savings above 13% were achieved without negatively impacting thermal comfort.

Keywords: District heating, marginal price, control, thermal energy storage, load shifting.

Contents

Nomenclature

Abbreviation	Name
DHN	District Heating Network
DH	District Heat
TES	Thermal Energy Storage
HOB	Heat Only Boiler
CHP	Combined Heat and Power
MPC	Model Predictive Control
HX	Heat Exchanger
MLR	Multiple Linear Regression
ODE	Orginary Differential Equation
DDE	Delay-Differential Equation
RKN	Runge-Kutta method of order N
PID	Proportional-Integrating-Differentiating controller
PS	Proportional-Summation controller
HVAC	Heat, Ventilation and Air-Conditioning
PSO	Particle Swarm Optimisation
HL	Heat Load
DHW	Domestic Hot Water
cs_i	Control system i
T-module	Part of control system attempting to maintain a stable indoor temperature
P-module	Part of control system attempting to shift heat load to hours with lower marginal cost

Symbol	Unit	Name/Description
G_d	%	Relative daily variation
ΔT	$°C$	Control Signal
T_{nom}	$°C$	Nominal indoor temperature, $22°C$
$T_{outside}$	$°C$	Outdoor temperature
$T_{adjusted}$	$°C$	Adjusted outdoor temperature, $T_{nom} + \Delta T$
τ	h	Building response time constant
T_{ap}	$°C$	Apartment and shallow storage temperature
T_{deep}	$°C$	Deep thermal storage temperature
HWVp1	Unitless, $\in [0, 1]$	DHW system valve position 1
$c_{p,eff}$	$\frac{MWh}{K}$	Effective heat capacity
HL_{DHW}	MW	DHW subsystem Heat Load
HL_{rad}	MW	Radiator subsystem Heat Load
$T_{rad,setpoint}$	$°C$	Radiator supply setpoint temperature
$T_{rad,return}$	$°C$	Radiator return temperature
F^{-1}	Unitless	Inverse Fitness, or cost
x_i	MWh	Model energy storage i
u_i	$°C$	Model input signal i
r_i	MW	Model transfer rate i
c_i	Various	Model parameter i
d_i	Various	Control system parameter i
ΔT_T	$°C$	T-module control suggestion
ΔT_P	$°C$	P-module control suggestion
ΔT_{lim}	$°C$	Maximum allowed control magnitude
ΔT_{varlim}	$°C$	Maximum rate of change of control magnitude
$T_{min,n}$	$°C$	Lowest recorded T_{ap}, last n days
T_{min}	$°C$	Lower indoor temperature limit
$f_{sub,n}$	Unitless, $\in [0, 1]$	Fraction of time last n days with $T_{ap} < T_{min}$
$f_{sub,lim}$	Unitless, $\in [0, 1]$	Highest acceptable values of f_{sub}
T_{var}	$°C$	Largest T_{ap} difference in last 24 h, negative if cold peak most recent
$T_{var,lim}$	$°C$	Highest acceptable absolute value of T_{var}

1

Introduction

In light of the current growing environmental concerns, serious efforts have in recent years been put into reducing nationwide energy consumption, both related to electricity and heat use. The energy source can however in many cases be as important as the amount, a fact only partly reflected in todays market. Some electricity providers now allow end consumers to purchase electricity at prices varying by the hour, giving them incentives to spread their usage in a more system-friendly manner. The same method should in principle be applicable to district heating networks (DHN), as many of those also draw energy from multiple sources with widely varying production characteristics. According to Svenska Energimyndigheten[1], 12% of Swedens total energy demand in 2014 was heating of apartments and office spaces with district heating (DH). Thus, the successful implementation of a system that maximises the use of preferential district heating sources could potentially have a significant positive environmental impact.

In Gothenburg, whose DHN comprises many sources [21] with widely differing characteristics, the shift of heat load from less to more environmentally friendly sources has an especially large potential. The largest, in terms of total annual energy delivered, and most environmentally friendly sources are excess heat from industries and garbage incineration. On the other end of the spectrum lies heat only boilers run with fossil fuels, whose annual active duration period should be limited if possible. See Figure 1.1 for a qualitative illustration of how the energy mix can vary with varying heat load. “Marginal” cost, economic or environmental, is the energy generation cost of the plant whose output would be lowered first if the demand was reduced, which usually is the most expensive one. As heat demand varies with the daily cycle, it is not uncommon for the marginal price to vary with a factor two during a single day.

A promising method for reducing these variations is short term thermal energy storage (TES). A TES can absorb heat during off-peak hours and release it back into the system as demand increases. The most straightforward implementation is a hot water tank situated near a major plant, as investigated by Verda and Colella [20]. While in principle an elegant solution, the tank itself can be difficult to implement for a number of practical reasons. This is especially true in a distributed network where no single plant provides heat for a majority of the system. An alternative method is to utilise the thermal inertia of buildings already connected to the grid, and make the building cores act as temporary storage units.

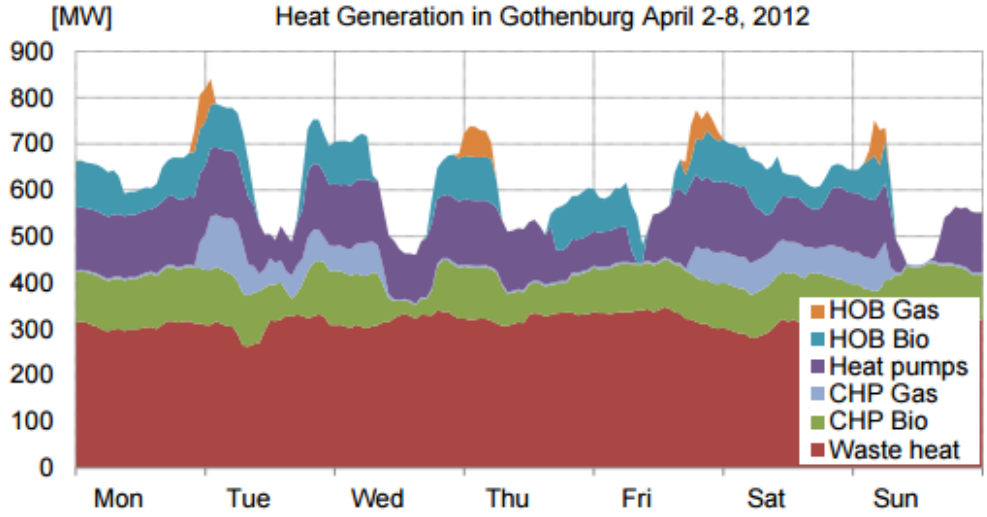

Figure 1.1: Total heat load and energy source breakdown for the Gothenburg DHN during a week in April, 2012. Image source: `http://publications.lib.chalmers.se/records/fulltext/216470/216470.pdf`

A pilot test (referred to as "*the pilot test*" throughout the rest of this report) regarding the potential of this method has been conducted by Göteborg Energi in collaboration with Chalmers University of Technology, the main findings of which were summarised by Johan Kensby [12]. It was found that many buildings with a concrete core in Gothenburg are well suited for this purpose, and that the method is economically preferable to the installation of a hot water tank if the system installation cost is kept below €3000 per substation. However, making efficient use of this type of storage is more difficult than optimally utilising a hot water tank. Thus, this project aims to evaluate how well relatively simple control systems can be expected to perform at this task.

1.1 Aims and limitations

This project will expand on the findings of [12] and evaluate how well the predictions regarding the potential of building TES are affected by taking building dynamics and control difficulties into account. A control system which controls building heat load from the current and future marginal price will be developed and evaluated. To accomplish this, a building model will be created based on data from *the pilot project* that can predict how a building's heat demand and indoor temperature will vary when the signal from its outdoor temperature sensor is adjusted. Control systems will then be evaluated on these model buildings for the purpose of minimising economic cost while maintaining thermal comfort, with building heating cost based on the marginal cost of heat generation. Control methods will be evaluated on multiple buildings over two years, for which climate and marginal cost data was available. Building characteristics will be identified and compared to the conclusions drawn from the pilot project, to validate results of the project and draw further con-

clusions regarding how the building would react when subjected to other types of control cycles.

Speed of development, simulation and evaluation will be prioritised throughout the project. Thus the intention is not to create an especially detailed building model, but instead one which can be quickly modified to simulate a large number of buildings with varying characteristics. Then, control algorithms will be selected based on their performance on a number of buildings. While the best possible control method likely is a form of model predictive control (MPC) working from an on-line adaptive grey-box model, the implementation of this was deemed outside the scope of the project. Therefore this project did not aim to create the most advanced possible control system but rather show what emergent difficulties one can expect, give an indication of how well a more advanced control system could perform and provide a potential platform for the development of a more advanced control system.

2

Background

2.1 Thermal energy storage

Utilising short-term thermal energy storage in buildings is not a new concept. In 2015 Heier et al [10] published a comprehensive review of options for thermal energy storage in buildings, focusing on systems explicitly installed to increase the thermal mass of the building. They noted that high thermal inertia can increase thermal comfort while potentially reducing heating or cooling demands depending on the season, as well as facilitate load shifting. In Sweden, there are many residential buildings with a concrete core whos naturally high thermal mass could be utilised in a similar manner without the need for additional installations. In the final report of project Cerbof 2:2[21] the authors note that a heavier building core can, but does not necessarily, lead to lower heat demands. In order to utilise the possibilities associated with larger thermal inertia the heating control system has to be designed around it, and in a complementary argumentative article [18] three of the authors recommend that existing heavy residential buildings are retrofitted with control systems that take this into account.

The recent rapid increase in availability of small general purpose computers and network solutions has inspired a lot of research regarding advanced control algorithms for buildings. A comprehensive survey regarding recent developments was done in 2014 by Afram et al [6]. They concluded that systems making use of model predictive control (MPC) typically outperform control schemes that don't, both in economy and thermal comfort, although the concept of MPC is very broad and there are many ways it can be implemented into or combined with other control methods. MPC requires a dynamical model of the building that can be evaluated quickly in real-time by the control system. This can be a physical model built from detailed knowledge of the building, as done by Schirrer et al [19], or a grey- or black box model with on-line parameter updating.

Commercially, multiple companies and solutions already exist with varying implementation methods, two examples of which will be given here. Kabona Ecopilot®is a centralised control system that monitors both heating, cooling and ventilation systems. It utilises zone control, weather forecasts and locally produced excess heat to minimise unnecessary heating or cooling of offices and other locales with large open spaces. While many of their customers report savings of 20-40% total heating and cooling costs [4], many of the energy saving measures might not

be applicable to residential buildings. Additionally its installation might require a significant overhaul of the current heating and ventilation system for proper use as well as extensive expert knowledge both for installation and fine tuning. Thus its applicability for mass installation on residential buildings might be limited. A somewhat easier to implement solution is provided by NODA Intelligent Systems. Their Smart Heat Building Solutions^{TM}[5] system utilises a continuously updated mathematical model of the building combined with weather reports to improve indoor temperature stability and reduce unnecessary heating, with reported average energy savings of about 10%. Indoor temperature and humidity sensors are required, but the control system is easily installed on top of the existing building substation and control scheme adaptation is automatic, which should make it more suitable for large scale implementation.

2.2 District heating building subsystems

District heating is based on the principle large centralised energy plants often are more efficient than many small distributed ones. By organising acquisition of heat on a central level many sources – such as industrial excess heat, combined heat and power installations and geothermal energy – can be incorporated into the energy mix that for market reasons would be difficult to utilise if every household was responsible for its own heating. As discussed in section 1, the energy mix of a given DHN usually becomes less economical and environmentally friendly the higher the instantaneous heat load. This means the average network efficiency for a given daily heat load is maximised when short term heat load variations are minimised, which also minimises costs related to startup and shutdown of plants. A measurement for quantifying the severity of short-term variations was suggested by Gadd and Werner [8]. They define *relative daily variation* (G_d) as

$$G_d = \frac{\sum_{h=1}^{24} |P_h - P_d|}{2 \cdot P_a \cdot 24} \cdot 100(\%), \tag{2.1}$$

where P_h is the heat load over hour h, P_d the total heat load over that day and P_a the annual average daily heat load. They analysed 20 Swedish DHNs and found that $4\% < G_d < 5.5\%$ for 16 of those.

The connection point between the DHN and a buildings internal energy system is referred to as a substation, of which there are many different configurations. One of the simplest solutions is illustrated in Figure 2.1. Here, hot water is tapped from the DHN and split between two heat exchangers (HXs). One exchanges heat with the radiator system, the other with the domestic hot water (DHW) system of the building. After the heat exchangers, the two cold streams are merged and the water is sent to the DHN return stream. The building owner is billed based on the measured flow rate and heat difference between the supply and return temperature of the DHN water. Flow rates on both sides of the heat exchangers are controlled locally. Cold water is periodically added to the DHW circuit as this is not a closed loop, so the heat load over the DHW HX is naturally very uneven. Over the radiator HX, flows are adjusted so that the radiator supply water temperature matches

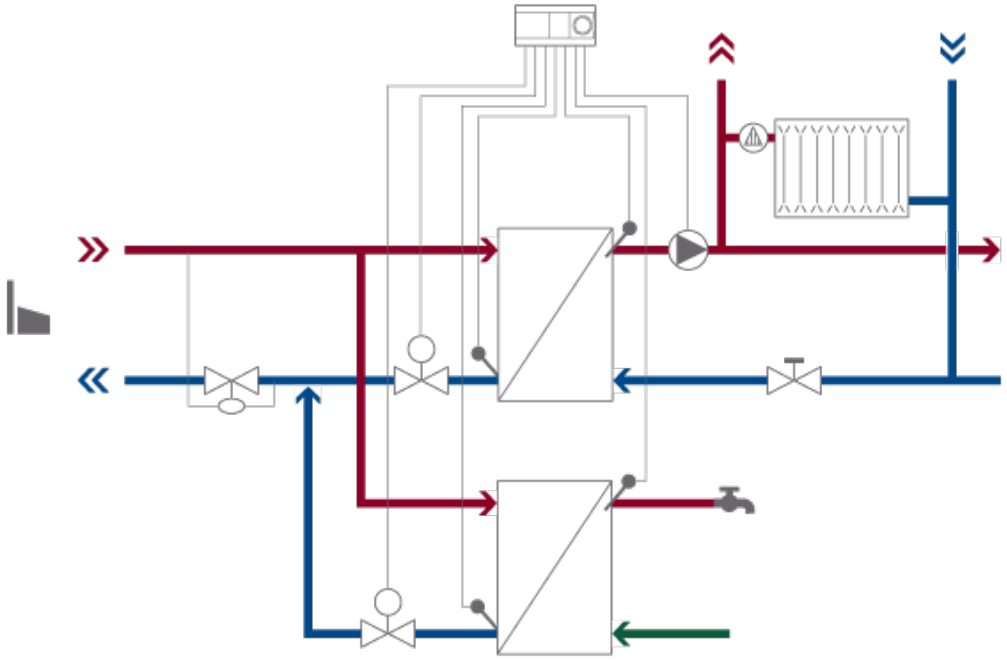

Figure 2.1: Schematic of residential building substation connected to district heating network. Image retrieved from `http://dh-applications.danfoss.com/application411.html?build=mfc&app=com`.

some setpoint given by the building's internal control system. There are alternative implementations, for example with partial series connections between the two HXs, which from the perspective of the systems discussed in this report are equivalent.

Traditionally, the radiator setpoint is determined by the building's setpoint curve, which often is a piecewise linear function of outdoor temperature. The apartment radiators are then simply dimensioned such that, during nominal conditions, a comfortable indoor climate is achieved without the need for system feedback. As discussed in section 2.1 both thermal comfort and energy use can be improved by making the setpoint also depend on indoor temperature, weather conditions, current apartment activity and similar, but there are many possibilities for incorporating this setpoint control into the existing system. A simple implementation method was used in the pilot project between Chalmers university of Technology and Göteborg Energi [12]. Instead of controlling the radiator setpoint directly, a control signal ΔT was added to the outdoor temperature reading. This adjusted signal, from here on referred to as $T_{adjusted}$, was then sent to the built-in control system instead of the actual outdoor temperature. See Figure 2.2 for a schematic illustration.

This implementation method is easy to apply to existing buildings, and so is suitable for retrofit purposes. It is also easy to interpret, since if a constant adjustment signal is maintained for long periods of time the indoor climate should eventually adjust towards $T_{nom} - \Delta T$, where T_{nom} is the nominal temperature experienced without any external manipulation. In reality, this statement only holds if ΔT is small enough that no secondary effects in the building activate, such as

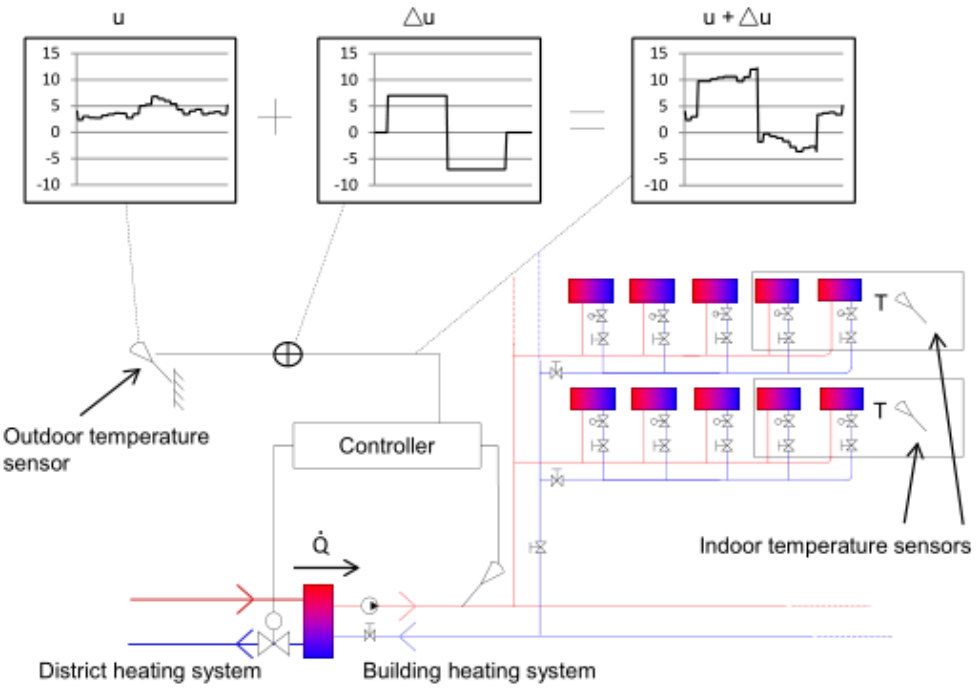

Figure 2.2: Schematic of pilot test setup. Image retrieved from `http://publications.lib.chalmers.se/records/fulltext/216470/216470.pdf`

apartment thermostats throttling radiator flows or tenants opening windows. Attempts were made to describe the studied buildings in terms of a time constant τ, by assuming that indoor temperature T_{ap} upon setpoint adjustment behaves according to

$$T_{ap}(t, \Delta T) = T_{ap}(0) - \Delta T \left(1 - e^{\frac{-t}{\tau}}\right) \tag{2.2}$$

This model was found to not be sufficient to characterise the buildings, as different values of τ were obtained depending on measurement time and magnitude of ΔT. Nevertheless, equation (2.2) can be useful for comparing the responsiveness of different buildings, as long as the experiment parameters are kept consistent.

2.3 Modelling and Simulation

A simulation model is created primarily as a faster and more convenient alternative to running real-world experiments. A model can either be created before its real world equivalent is constructed, or created to mimic some physical system. In the latter case, the model is often fine-tuned by simulating it with the same inputs as the real system experienced and modified until it produces a similar output. Once this is achieved, one can insert new inputs into the model and acquire a prediction of how the real world system would react to that scenario. When used in conjunction with an optimisation routine, this is done repeatedly until the combination of inputs that produce the most favourable outcome is found.

Many comprehensive tools for simulating the energy flows in buildings to

great detail exist. Examples include open-source projects such as EnergyPlus [2] and commercial products such as IDA ICE [3]. These tools allow for very high fidelity models and accurate representation, but both creating and simulating models was deemed too time consuming to be within the scope of this project. Simpler, faster-to-iterate, mathematical models can be built by hand relatively quickly. To find relationships between data, Multiple Linear Regression (MLR) can be used. Relatively simple systems of ordinary differential equations can also describe many dynamical systems to reasonable degree.

2.3.1 Regression modelling

MLR is a statistical tool for finding correlations between data [17]. One assumes that a recorded output response y_r can be estimated by a linear combination of state variables $\mathbf{x} = [x_1, \ldots, x_k]^T$. The linear weight coefficients $\mathbf{b} = [b_i \ldots, b_k]^T$ are chosen to minimise the sum of squared errors over N datapoints, that is:

$$\mathbf{b} = \text{argmin}_{\mathbf{b}} \sum_{j=1}^{N} \left(y_{r,j} - \sum_{i=1}^{k} b_i x_{i,j} \right)^2$$

$\mathbf{x}$ is often made to contain a vector with constant terms of magnitude 1. The corresponding weight is then referred to as the *model bias*. After the weights have been found, the system responses y_p for any possible future state of $\mathbf{x}$ can be estimated as

$$y_p = \sum_{i=1}^{k} b_i x_i$$

The algorithms for finding these coefficients are deterministic, making this a very quick and reliable method for data analysis. In this project the "regress" function, provided with the MATLAB®Statistics and Machine Learning Toolbox, was used.

2.3.2 Dynamical models

In this report, the term "Dynamical Model" refers to a model consisting of a set of differential equations, specifically systems of restricted ordinary differential equations (ODEs). A system of ODEs consists of some time dependent inputs $\mathbf{u} = [u_1, \ldots, u_m]^T$, a set of state variables $\mathbf{x} = [x_1, \ldots, x_k]^T$ and their respective first order derivative functions $\mathbf{f}(\mathbf{x}, \mathbf{u}) = [f_1(\mathbf{x}, \mathbf{u}), \ldots, f_k(\mathbf{x}, \mathbf{u})]^T$, such that

$$\begin{cases} \frac{dx_1}{dt} = \dot{x_1} = f_1(x_1, \ldots, x_k, u_1, \ldots, u_m) \\ \vdots \\ \frac{dx_k}{dt} = \dot{x_k} = f_k(x_1, \ldots, x_k, u_1, \ldots, u_m) \end{cases} \tag{2.3}$$

If the equations are linear with respect to $\mathbf{x}$ and $\mathbf{u}$, i.e $f_i(\mathbf{x}) = \sum_{i=1}^{k}(a_i x_i) + \sum_{j=1}^{m}(b_j u_j)$, and there are no restrictions on either $\mathbf{x}$ or $\frac{d\mathbf{x}}{dt}$, the system is referred to as a state-space model. In that case equation (2.3) can be rewritten as

$$\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u}$$

where A and B are constant matrices. Due to their linear and unrestricted nature, many analysis tools and solving techniques can be used that make state-space models faster and easier to work with than general systems [14].

In some special cases, a system of ODEs can be solved analytically, giving exact solutions. When this is not possible, an approximate numerical solver has to be used. The most common techniques involve approximating the system rate of change, $\mathbf{g}$, as piecewise constant. The simulation period is split into many smaller intervals, and within each interval $\{t, t+h\}$ one performs the calculation

$$\mathbf{x}(t+h) = \mathbf{x}(t) + h \cdot \mathbf{g}(t)$$

The characteristics of the solver – such as speed, accuracy and stability – are determined by what step size h is used and how $\mathbf{g}$ is calculated from (2.3). The simplest method, known as the *Forward Euler* method, is to let $\mathbf{g}(t) = \mathbf{f}(t)$ for the entire interval. This method is incredibly easy to implement, but requires a very small time step h to ensure good accuracy and stability, which potentially makes it computationally expensive. The FE method can be regarded as a special case of a larger family of solvers known as the *Runge-Kutta* (RK) methods. In an explicit RK method of order q, the step slope is calculated as

$$\mathbf{g}(t) = \sum_{i=1}^{q} b_i \mathbf{k}_i,$$

where

$$\begin{aligned} \mathbf{k}_1 &= \mathbf{f}(t, \mathbf{x}), \\ \mathbf{k}_k &= \mathbf{f}(t + c_k h, \mathbf{x} + \textstyle\sum_{j=1}^{k-1} a_{kj} \mathbf{k}_j). \end{aligned}$$

Effectively, this method starts with a *Forward Euler* estimate ($\mathbf{k}_1$), creates a second estimate $\mathbf{k}_2$ based on $\mathbf{k}_1$, creates a third estimate from $\mathbf{k}_1$ and $\mathbf{k}_2$ and so on, until finally letting $\mathbf{g}$ be a weighted average of all these partial estimates. There are criteria that dictate how to choose method parameters a, b and c to achieve optimal results, for details see for example [9].

While it is usually possible to estimate beforehand what step size and method order will be required to achieve good results within a given system, in the interest of generality many commonly used solvers are built on the principle of adaptive step size. An example is the MATLAB®solver ode45, which is based on the work of Dormand and Prince [7]. This method calculates two simultaneous solutions using Runge-Kutta methods of orders 4 and 5, then compares these solutions to estimate current numerical error and determine optimal step size.

2.4 Control Systems

An "optimal" control system uses an optimisation algorithm to periodically find the optimal control decisions, often using a predictive model of the system. While these are potentially very powerful, their implementation was deemed beyond the scope of this project and so they will not be considered further. A more commonly used

control method is the classical controller. These units operate on the principle of being fed an error value e – which is the difference between a measured quantity and its preferred *setpoint* value – and outputting a control signal based on this error. The most general is the Proportional-Integrating-Differentiating (PID) controller, which has three parameters that can be tuned to balance speed and stability [13]. Often a PI controller – which omits the differentiating term – is used instead, as it is less sensitive to measurement noise and less prone to cause system instability. A discretised PI controller, whose integrating term often is replaced with simple summation, will at time step n send a control signal $u(t_n)$ according to

$$u(t_n) = K_p e(t_n) + K_I \sum_{i=0}^{n} e(t_i) \tag{2.4}$$

where K_p is the proportional controller magnitude and K_I the integrating controller magnitude. An alternative, but equivalent, implementation of the summation part is to store a summation control output term u_s which at every time step is updated according to the current error. Then, equation (2.4) can be rewritten as

$$\begin{cases} u_s(t_n) = u_s(t_{n-1}) + K_I e(t_n) \\ u(t_n) = K_p e(t_n) + u_i(t_n) \end{cases} .$$

The applications of a PI-controller can be extended by varying the control magnitudes or how the error is calculated depending on the current system state. One can also implement multiple concurrent controllers in a hierarchical configuration, where each unit sets some parameters – usually the setpoint – to the next unit. This is known as cascade control and is useful when there are relevant dynamics acting on separable time scales. An example of this is found in building substations, where the radiator supply temperature setpoint is chosen by a mathematical function and fed to a PI or PID controller which attempts to minimise the error between this setpoint and the actual supply temperature. This mathematical function can be another PID controller, a simple setpoint curve or a more advanced decision system, such as a neural network or a model predictive control (MPC) algorithm [6].

A control method which requires less computing power than optimal MPC, less training data than a neural network but still has more utility than a classical PID controller is a goal oriented decision making algorithm. An novel application of this on a system with a limited energy storage and conflicting goals has been presented by Kaiser et al [11]. They formulated four goals with different weights, whose utilities from different operational strategies could be calculated relatively quickly. Three of these were compounded into a total utility value, and at regular intervals operating conditions were set to maximise this value subject to limitations set by the first goal. This implementation method was found to significantly increase lifecycle savings of the installed storage system.

2.5 Parameter Optimisation

Optimisation is an indirect process of finding a set of input parameters $\mathbf{x}$ that yield the optimal value of some objective function $f(\mathbf{x})$. Numerical iterative optimisation is necessary when this cannot be done exactly, for example if there are constraints on the input variables which make a direct search difficult or the objective function has to be solved numerically. Many factors can affect what results are achieved from the process, including how the objective function is chosen and what optimisation method is used.

The goal of the objective function is to quantify the suitability, represented as a *fitness value* (F) or *inverse fitness* (F^{-1}), of a combination of input parameters, so that many such combinations can be easily compared. When optimising a physical model to gathered data, it is common to attempt to minimise the sum of squares of errors between simulated and measured data. If multiple quantities are to be matched simultaneously they have to be normalised so they can be meaningfully compared. The standard procedure on a given data vector is to subtract its mean, then divide by its standard deviation [16]. The same operations are performed on the corresponding simulated vector. Let $y_{j,i,d}$ be the i'th measured value of normalised data vector $\mathbf{y_j}$, $y_{j,i,s}(\mathbf{x})$ the simulated value for a given input vector ($\mathbf{x}$) and w_j the error weight for that output. One possible inverse fitness function, which is to be minimised, is then

$$F^{-1} = f(\mathbf{x}) = \frac{1}{n\sum_{j=1}^{m}(w_j)} \sum_{j=1}^{m} w_j \sum_{i=1}^{n} \left(y_{j,i,s}(\mathbf{x}) - y_{j,i,d}\right)^2 . \tag{2.5}$$

This formulation is also beneficial in that the cost function in principle are independent of series length (although this fails if there are long-term seasonal changes), allowing comparison between different time series. When a control system is optimised, the objective function has to be designed to encourage the intended behaviour of the system. An example of this is given in [15], where a control system split between two hierarchical layers is judged on running cost, consistency between control layers and tendency to maintain a preferable system state.

The goal of the optimisation method is to effectively search the parameter space and find an optimum input configuration using as little computing time as possible. There are many possible choices, each with their strengths and weaknesses. A comprehensive, albeit somewhat out of date, review of optimisation algorithms used for control of HVAC systems was published in 2007 by Wang and Ma [23]. Their review stresses the nonlinear cause-and-effect relationships often present within HVAC systems, meaning that many local optima are expected within the parameter space creating potential difficulties for local search algorithms.

One method not presented in [23] is the particle swarm optimisation (PSO) algorithm. It is a general stochastic search algorithm, where the tradeoff between the rate of convergence and risk of converging to local optima can be managed by modification of a few parameters. The algorithm will be qualitatively described here, see [22] for a more detailed description. In PSO, the parameter combinations

to be evaluated are represented by a swarm of particles moving through parameter space. Initially, each particle $\mathbf{x}_i$ is placed randomly within some provided upper and lower bounds $[\mathbf{x}_{min}, \mathbf{x}_{max}]$ and is given a random initial velocity. Each iteration step then follows the following procedure:

1. The objective function is evaluated once for each particle, and the resulting fitness or cost value is attached to the respective particle.
2. Each particle updates its personal best evaluated value, and its associated location $\mathbf{x}_i^{pb}$. The historically best global value and position $\mathbf{x}^{sb}$ is also noted.
3. Each particle updates its velocity, by adding to its current velocity a linear combination (with partly random weights) of the direction vectors from its current location $\mathbf{x}_i$ towards $\mathbf{x}_i^{pb}$ and $\mathbf{x}^{sb}$.
4. The swarm is simulated for one time step, during which all particles move according to their current velocity vector.

Many additions can be made to this basic algorithm to adjust performance. In this work, a decaying inertia weight which increases the rate at which particles can turn over time is added, and particles will be randomly "pushed" once their speed goes below a certain threshold to avoid stagnation.

3

Methodology

3.1 Model Design

Before a control system can be evaluated, a dynamical building model that reacts realistically to control signals is required. For ease of implementation and speed of evaluation, a restricted system of first order ordinary differential equations (ODEs) was used. All data manipulation and simulations was done using MATLAB®, unless otherwise specified. Throughout this section, considerations regarding data handling and model design are presented.

3.1.1 Data Preprocessing

A constructed model mimicking a real system will have to be validated against known and reliable data. Therefore all data first has to be organised and evaluated. On the district heating side, network-wide heat generation and normalised marginal and average heat cost was available at hourly resolution for 2013 and 2014. Their respective relative daily variation values (see Equation (2.1)) are displayed in Table 3.1, which highlights how the marginal cost fluctuates much more strongly than the heat load. On the building side, available data consisted of continuous readings from various sensors installed on a single building from 2010-02-01 through 2011-03-15. Most quantities were logged every 10 minutes, one exception being the accumulated total building heat load (HL) which was measured every hour. The registered value at every hour from this sensor corresponds to the cumulative measurement at the end of that hour, with the exception of the first value each day corresponding to the measured value at the end of that day. This means the actual value at 01:00 is unknown, and had to be estimated by assuming a constant load between 00:00 and 02:00. An hourly average instantaneous heat load could then be created as the difference of cumulative heat load between two measurements.

$\mathbf{G_d}[\%]$	**Heat load**	**Average cost**	**Marginal cost**
2013	3.94	3.53	11.11
2014	3.57	2.81	9.50

Table 3.1: Relative daily variation of network wide data

The total building HL is shared between two subsystems, the radiator system and the domestic hot water (DHW) system. The proposed control systems will not be able to affect DHW usage, and so an estimate of the split between radiator heat load (HL_{rad}) and DHW heat load (HL_{DHW}) was required. This split can be estimated by isolating data from a time periods during which one of the systems is known to dominate the total HL and calculating this load from available sensor measurements. Data from all hours for which the radiator return temperature ($T_{rad,return}$) was higher than the setpoint ($T_{rad,setpoint}$) was extracted, as $HL_{DHW} \approx HL$ was assumed during those hours. Initially attempts were made to fit nonlinear equations based on assumed relevant parameters, but no consistently good fit was found and so MLR was used to determine which parameters were relevant. Initially all available temperatures and valve positions were used, and the ones whose resulting regression parameters had small relative values and large confidence intervals were removed. Eventually only the constant bias term and DHW HX valve position 1 ($HWVp1$) were kept, the latter represented three times with exponential terms 0.5, 1 and 2, to account for the non-linear nature of valves. This equation performed slightly better than the expression $HL_{DHW} = a_1 \cdot HWVp1^{a_2}$.

Indoor temperature was measured in two apartments at all times. There are naturally variations in indoor temperature that can not be predicted by the available data due to tenant behaviour, which should be removed from the data if at all possible. This was done using a custom function for identifying unrealistic measurements based on the rate of change compared to surrounding readings. Identified outliers were replaced by linear interpolation between the closest non-outlier on either side. It was also found that both sensors started reporting continuously unrealistic data around 75 days into the experiment, with values either rapidly oscillating between two or three values, or remaining completely constant for unrealistically long periods of time. Thus, reliable indoor temperature data during the heating season was only available for two and a half months during spring, 2010.

3.1.2 Dynamical Building Model

The model assumes the building can be split into three energy storages with internal homogeneity. Energy flow in and out of each of these storages is described by a linear first order differential equation, with some restrictions regarding flow directions and magnitudes. The variables and parameters used in the model are summarised in Table 3.2. Note that the model temperature of a given energy storage is found by multiplying the energy level with its corresponding heat sensitivity coefficient, so for example $T_{ap} = T_2 = c_2 x_2$. Radiator return temperature ($c_1 x_4$), rather than total radiator system energy, is used to estimate the flow rate from the DHN to the radiator system (r_1). There is a natural delay between radiator setpoint (u_1) being changed and return temperature reacting, which is included in the model by letting radiator return energy (x_4) be a state variable which continuously moves towards an algebraically estimated return energy (x_4^{alg}). This method was used because the more intuitive formulation, of using a pure delay, lead to an oscillatory solution behaviour that was not reflected in the data. A qualitative illustration of the model

Table 3.2: Summary of variables and parameters used in the building model. For applications, see Figure 3.1 and equations (3.1) and (3.2).

Letter	Description	Unit
x_1	Radiator system energy	MWh
x_2	Shallow storage energy	MWh
x_3	Deep storage energy	MWh
x_4	Radiator return energy	MWh
x_4^{alg}	Algebraic radiator return energy	MWh
u_1	Radiator system setpoint	$°C$
u_2	Outdoor temperature	$°C$
r_1	Transfer rate DHN to radiator system	MW
r_2	Transfer rate radiator system to shallow storage	MW
r_3	Transfer rate deep storage to shallow storage	MW
r_4	Transfer rate shallow storage to outside	MW
c_1	Heat sensitivity of radiator subsystem, x_1 and x_4	$\frac{K}{MWh}$
c_2	Heat sensitivity of shallow storage, x_2	$\frac{K}{MWh}$
c_3	Heat sensitivity of deep storage	$\frac{K}{MWh}$
c_4	r_1 transfer coefficient	$\frac{MW}{K}$
c_5	r_2 transfer coefficient	$\frac{MW}{K}$
c_6	r_3 transfer coefficient	$\frac{MW}{K}$
c_7	r_4 transfer coefficient	$\frac{MW}{K}$
c_8	Stable temperature difference	$°C$
c_9	Radiator setpoint factor	Unitless
c_{10}	Average circulation time	h
c_{11}	Radiator return diffusion coefficient	Unitless

is shown in Figure 3.1. The mathematical equations governing the model are written explicitly in equation (3.1), with restrictions summarised in equation (3.2).

$$\begin{cases} \frac{dx_1}{dt} = & r_1 - r_2 \\ \frac{dx_2}{dt} = & r_2 + r_3 - r_4 \\ \frac{dx_3}{dt} = & -r_3 \\ \frac{dx_4}{dt} = & \frac{c_{11}}{c_{10}}(x_4^{alg} - x_4) \end{cases} \quad \text{where} \quad \begin{cases} r_1 = & c_4(c_9u_1 - c_1x_4) \\ r_2 = & c_5(c_1x_1 - c_2x_2) \\ r_3 = & c_6(c_3x_3 - c_2x_2) \\ r_4 = & c_7(c_2x_2 - (u_2 - c_8)) \\ x_4^{alg} = & x_1 - \frac{1}{2}c_{10}r_2 \end{cases} \tag{3.1}$$

$$\begin{cases} \{r_1, r_2, r_4\} \geq 0 \\ r_1 \leq 0 \text{ if } c_1x_1 \geq u_1 \\ c_1x_4^{\text{alg}} \geq c_2x_2 \end{cases} \tag{3.2}$$

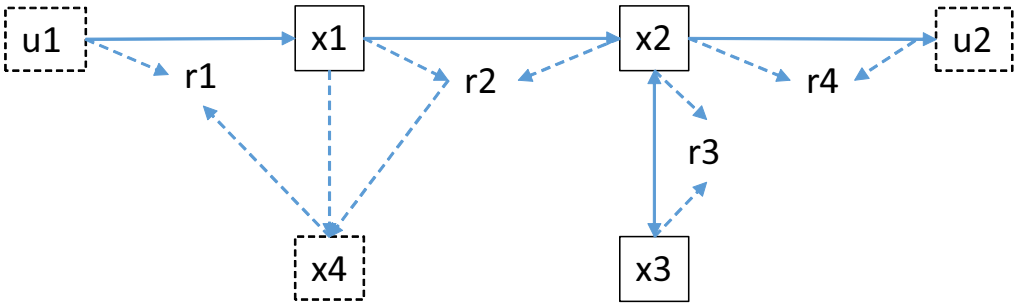

Figure 3.1: Block diagram illustration of the building model. Solid boxes and arrows represent storages and flows of thermal energy, dashed boxes and arrows represent information. u1 and u2 are external signals, of which u1 can be controlled. x4 is the energy level corresponding to the estimated radiator return temperature, which affects flow rate r1 but does not count towards the total amount of stored thermal energy.

3.1.3 Numerical Solver

Beyond accuracy, two requirements dictated the choice of numerical solver: evaluation speed and easy access to previously simulated values during simulation. For these reasons, an in-line solver was written that utilised a fourth order Runge-Kutta method with a fixed step size of 10 minutes. External data provided with a 10 minute resolution is assumed constant throughout these 10 minutes, so no interpolation is needed within each step. The results from this method were compared to simulations of the same system with the built-in solver ode45. Both function evaluation counts and results were very similar, and so the fixed-step method was deemed accurate enough for this case.

3.2 Control System

The control system can be regarded as consisting of two modules that give individual suggestions, with the final control signal being the addition of these suggestions subjected to some constraints. Restrictions related to the minimum indoor temperature as well as maximum daily variation of indoor temperature are imposed. The first control module, referred to as the T-module as it is concerned with indoor temperature, consists of a P-controller with a function that adjusts the setpoint over time depending on how well the indoor temperature restrictions are observed.

The second module, called the P-module, considers current and future marginal heat prices. Relative magnitudes of prices at different times are compared. The resulting expression is "squashed" in a symmetric double sinusoid function Φ_2, see equation (3.4). This equation was chosen above the simpler single sinusoid Φ_1 (3.3), because it provides an interval close to the origin where the two modules can compete, while exhibiting a very sharp response beyond this. The two functions are illustrated in Figure 3.2. Because the utility of function parameters s_1 and s_2 partly

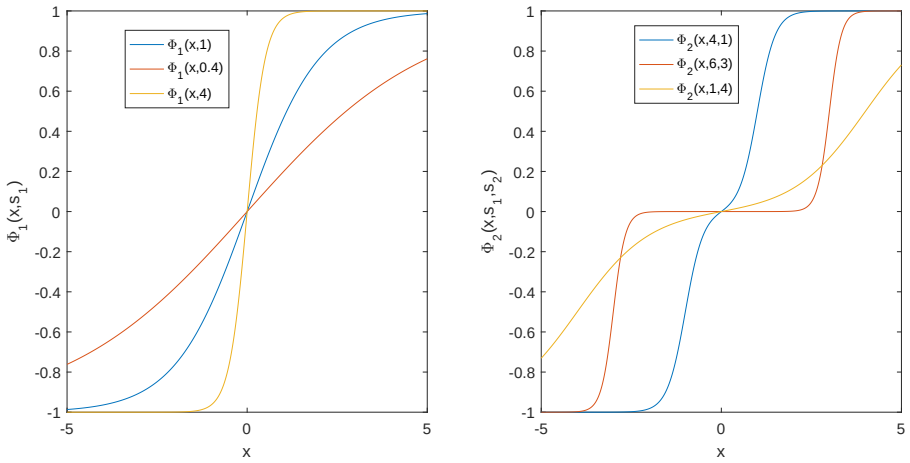

Figure 3.2: Illustration of the squash functions Φ_1 (3.3) and Φ_2 (3.4). Three sets of function parameters are displayed, to illustrate the range of possible function behaviours

overlap with those of the P-module, these were fixed during parameter optimisation to limit the number of active parameters. All parameters and variables relevant to the control system are displayed in Table 3.3, with the mathematical equations describing the system summarised in Table 3.4.

$$\Phi_1(x, s_1) = 2 \cdot \left(\frac{1}{1 + \exp(-s_1 x)} - 0.5 \right) \tag{3.3}$$

$$\Phi_2(x, s_1, s_2) = \left(\frac{1}{1 + \exp(s_1(s_2 - x))} + \frac{1}{1 + \exp(s_1(-s_2 - x))} - 1 \right) \tag{3.4}$$

No hard values have been given for the restrictions discussed throughout this section, because these can be tuned to reach a satisfactory compromise between potential savings and thermal comfort. One aspect of this project was to investigate how varying these will affect the savings potential. Four cases with different restriction variables were created, as displayed in Table 3.5. The primary interest is Case 1, which allows for a control signal magnitude of 7 $^\circ C$ and a daily variation of 1 $^\circ C$.

Restriction Case Nr	T_{min} $[^\circ C]$	$f_{sub,lim}$	$T_{var,lim}$ $[^\circ C]$	ΔT_{lim} $[^\circ C]$	ΔT_{varlim} $[^\circ C]$
1	21	0.1	1	7	3.5
2	21	0.1	1	14	7
3	21	0.1	2	7	3.5
4	21	0.1	2	14	7

Table 3.5: Collections of restriction parameters, referred to as *restriction cases*, which control systems were evaluated against

Table 3.3: Control parameters and states observed by the control system

Variable	Description
d_1	T-module setpoint, relative to T_{min}
d_2	T-module control magnitude
d_3	P-module maximum magnitude
d_4	P-module overall sensitivity
d_5	P-module consideration of price 12-36 hours forward
d_6	P-module consideration of price 36-48 hours forward
ΔT_T	Control suggestion from T-module
ΔT_P	Control suggestion from P-module
ΔT_{lim}	Maximum allowed absolute control magnitude
ΔT_{varlim}	Maximum rate of change of control magnitude
ΔT	Final control signal
T_{ap}	Indoor apartment temperature
$T_{min,n}$	Lowest recorded indoor temperature, last n days
T_{min}	Lower indoor temperature limit
$f_{sub,n}$	Fraction of time indoor temperature was below T_{min} last n days
$f_{sub,lim}$	Highest acceptable value of $f_{sub,n}$
T_{var}	Largest T difference in last 24 hours, negative if cold extremum most recent
$T_{var,lim}$	Highest acceptable absolute value of T_{var}
P	Current marginal production price of district heat
P_{12}	Mean price over coming 12 hours
P_{24}	Mean projected price between 12 and 24 hours forward
P_{48}	Mean projected price between 24 and 48 hours forward

Table 3.4: List of equations governing the control system. For parameter and variable descriptions, see Table 3.3. For the definition of Φ_2, see equation (3.4).

T-module, updates every 10 minutes
$\Delta T_T = d_2 \cdot \Delta T_{lim}(T_{ap} - (T_{min} + d_1))$
P-module, updates hourly
$\Delta T_P = d_3 \cdot T_{varlim} \cdot \Delta T_{lim} \cdot \Phi_2(P_\Sigma, 4, 1)$, $P_\Sigma = d_4 \left(\frac{P - P_{12}}{\min(P, P_{12})} + d_5 \cdot \frac{P_{12} - P_{24}}{\min(P_{12}, P_{24})} + d_6 \cdot \frac{P_{12} - P_{48}}{\min(p_{12}, p_{48})} \right)$
T-module setpoint, updates daily
if $f_{sub,1} > f_{sub,lim}$, $\quad d_1 \leftarrow d_1 + 0.5 \cdot (T_{min} - T_{min,1})$ if $f_{sub,7} < \frac{1}{7} f_{sub,lim}$, $\quad d_1 \leftarrow d_1 \cdot \left(1 - \frac{0.5}{7}\right)$
Output signal, updates every 10 minutes
The assignment $\Delta T_s = \Delta T_T + \Delta T_P$ is attempted, with the following restictions: if $T < (T_{min} + d_1)$, $\Delta T_s \leq \Delta T_T$ if $T < T_{min}$, $\Delta T_s \leq 0$ if $T_{var} > T_{var,lim}$, $\Delta T_s \geq 0$ if $T_{var} < -T_{var,lim}$, $\Delta T_s \leq 0$ $\lvert \Delta T_s - (\Delta T - d_1) \rvert < \Delta T_{varlim}$ After which the control signal is assigned as $\Delta T = \Delta T_s + d_1$

3.3 Parameter Fitting

When adapting parameters of a dynamical model one attempts to find the parameters that, upon simulation of the model, minimises some cost function that quantifies the discrepancy between simulation and experimental data. Equation (2.5) was used as cost function, with outputs and weights described in Table 3.6. Varying what properties were included or the respective property weights yields different sets of optimal parameters. The setup displayed in Table 3.6 represents a cost function that was found to provide a good compromise between minimising heat load prediction error while encouraging a reasonable physical behaviour. More weight was put upon the variability of indoor temperature than average temperature, because short-term behaviour related to control of the building was deemed more important to capture than seasonal behaviour and long-term variations. Alternative parameter sets for the other buildings used for control system evaluation were chosen by modifying and locking some parameters, then optimising the remaining such that the simulated apartment temperature duration curve in the absence of active control shared key characteristics with that of a modern, uncontrolled building. The resulting model parameters and average uncontrolled indoor temperature are displayed in Table A.1 in appendix A. Data for this uncontrolled building was provided by Göteborg Energi, the average indoor temperature for January through April in that building was $22.7°C$.

Table 3.6: Example of cost function properties that were used for dynamical model parameter fitting

Property	Example weight w_j
Energy transfer rate from DHN to radiator system	1
Radiator system return temperature	0.2
Mean difference between current indoor temperature and average temperature last 24 hours	0.4
Mean indoor temperature last 24 hours	0.1

Several of the parameters could be algebraically estimated directly from data, but for the final parameter search particle swarm optimisation was used. The swarm was made up of particles moving in a normalised k-dimensional cuboid with sides of length 1, where k is the number of parameters to be optimised. Real-world upper and lower bounds (uB and lB) were provided for each parameter, and these were used to map particle locations between the normalised optimisation space and the real parameter space. This was done to ensure that the relative magnitudes of the parameters did not affect how they were treated during optimisation, as well as simplifying the setting of optimisation parameters such as maximum particle velocity.

For parameters that were strictly positive and had a large search span, i.e 0 <lB≪uB, a nonlinear function was used to map between the real and normalised space. This mapping ensured that equal attention was given above and below the geometric mean, as opposed to the arithmetic mean, of the upper and lower bounds. Let x_R denote a parameter value in real space and x_N its corresponding normalised location, and the mapping was done according to equation (3.5). An illustration of this mapping, with $lB = 1$ and $uB = 20$ is displayed in Figure 3.3. For parameters where this mapping was either not possible or undesirable, β was set to 1 to obtain a linear mapping [lB,uB] ↔ [0,1].

$$x_N = \left(\frac{x_R - lB}{uB - lB}\right)^{\beta} \quad , \quad x_R = x_N^{1/\beta} \cdot (uB - lB) + lB \ , \qquad \beta = \frac{\log_2(1/2)}{\log_2\left(\frac{\sqrt{lB \cdot uB}}{uB - lB}\right)} \tag{3.5}$$

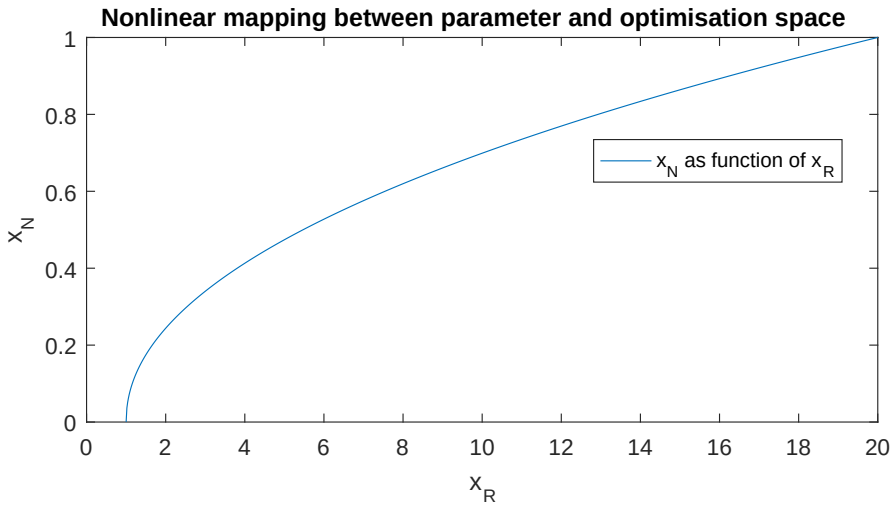

Figure 3.3: Illustration of the nonlinear mapping between real parameter space and normalised optimisation space, according to equation (3.5), when $lB = 1$ and $uB = 20$.

4

Results

4.1 Model Performance

Attempting to estimate the domestic hot water heat load (HL_{DHW}) from other known data using linear regression led to an expression on the form of equation (4.1), where $HWVp1$ refers to the normalised position of the first hot water valve. b_1 and b_3 are negative.

$$HL_{DHW} = b_1 + b_2\sqrt{\text{HWVp1}} + b_3 \cdot \text{HWVp1} + b_4(\text{HWVp1})^2 \tag{4.1}$$

Some limitations regarding this expression should be noted. First off, as can be seen in Figure 4.1, the estimated HL_{DHW} never drops below $0.0035MW$ even during night when hot water load is expected to be very close to 0. This is a result of the nonlinearity of the system combined with the relatively low time resolution, which results in the model typically underestimating the peaks and overestimating the throughs of HL_{DHW}. When extrapolated, equation (4.1) yielded negative values for less than 0.1% of all data points, which were simply set to 0. $HL_{DHW} > HL$ for 18.9% of all datapoints, 100.0% of which occurred while $(T_{rad,return}+2) > T_{rad,setpoint}$. Simply setting HL_{rad} to 0 for these datapoints is not necessarily problematic, as both radiator heat load and marginal price typically are very small in these scenarios, but more problematic is the fact that $HL_{DHW} < HL$ for 38.7% of all data points where $T_{rad,return} > T_{rad,setpoint}$. During these moments no radiator heat load is possible according to the model formulation, yet one is expected from the data. No adjustments were made to the data in these cases, instead model fitting was confined to the colder pats of the year where these effects did not occur.

The dynamical model presented in Section 3.1 was optimised towards given data between 01/02-2010 and 10/04-2010, using the inverse fitness function (2.5) and parameter setup as displayed in Table 3.6. Note that no feedback or corrections were given to the model throughout the simulation period. The resulting model parameter set, referred to as "building A1", is displayed together with all other buildings used for control system evaluation in Table A.1 in Appendix A. Typical resulting behaviour of the system compared to given data is illustrated in Figures A.1 and 4.2. Figure 4.3 illustrates how the simulated total thermal energy of the building varies with temperature. By comparing the TES curves to the second graph in Figure 4.2 it is possible to estimate the effective heat capacity $c_{p,eff}$ of the

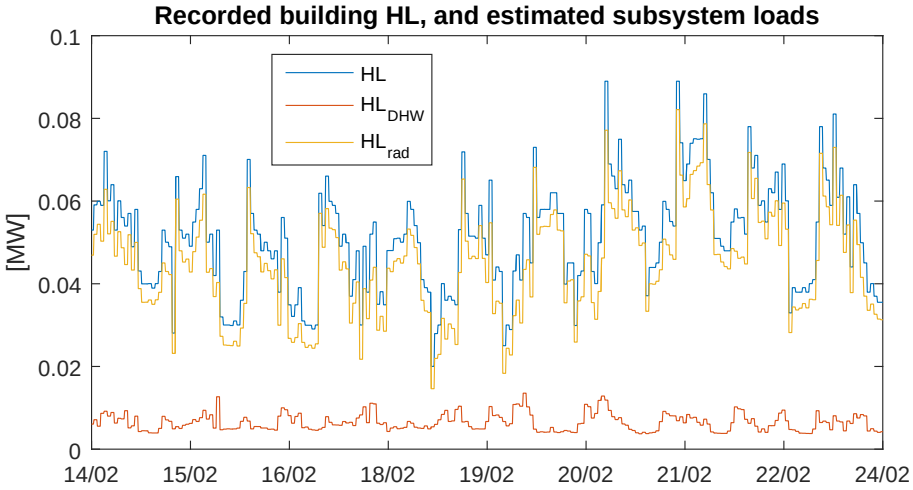

Figure 4.1: Illustration of the DHW system HL from measured valve positions, according to equation (4.1)

building, as calculated by the difference in internal energy divided by the difference of indoor temperature, over different time scales. Over individual cycles, an apparent heat capacity of $0.13 < c_{p,eff} < 0.14 \frac{MWh}{K}$ is observed, similarly to what was found by [12]. Performing the same calculation between the first and last local minimum of the series yields $c_{p,eff} \approx 0.34 \frac{MWh}{K}$, which is much closer to the absolute heat capacity achieved by adding the inverse of the two heat sensitivity parameters c_2 and c_3.

There were no other sufficiently long periods with reliable temperature data available for proper model validation. Instead the model (with parameters A1) was evaluated only on radiator heat load against the span 28/8-2010 through 15/3-2011. An error of 0.162 was achieved, where 0 is a perfect fit and 1 means the error has the same average magnitude as the data's standard deviation. An illustration of typical behaviour is displayed in Figure 4.4.

4.2 Control System Performance

Control system parameters were optimised towards minimising heating cost off building A1, then run on the remaining buildings to verify expected performance. The optimisation was ran multiple times under different restrictions on the control parameters, to achieve several characteristically different control schemes. Five such schemes will be presented and discussed here, displayed in Table 4.1. Control system 1 (cs_1) represents the hands-off case, with no control. cs_2 only contains the T-module, while cs_3 only utilises the P-module. cs_4 makes use of both modules, but only had price forecasts available up to 12 hours forward. cs_5 used price forecasts up to 48 hours forward, but its optimisation ran with stricter restrictions regarding amplitude parameters d_2 and d_3.

An illustration of how the different control systems tend to affect indoor

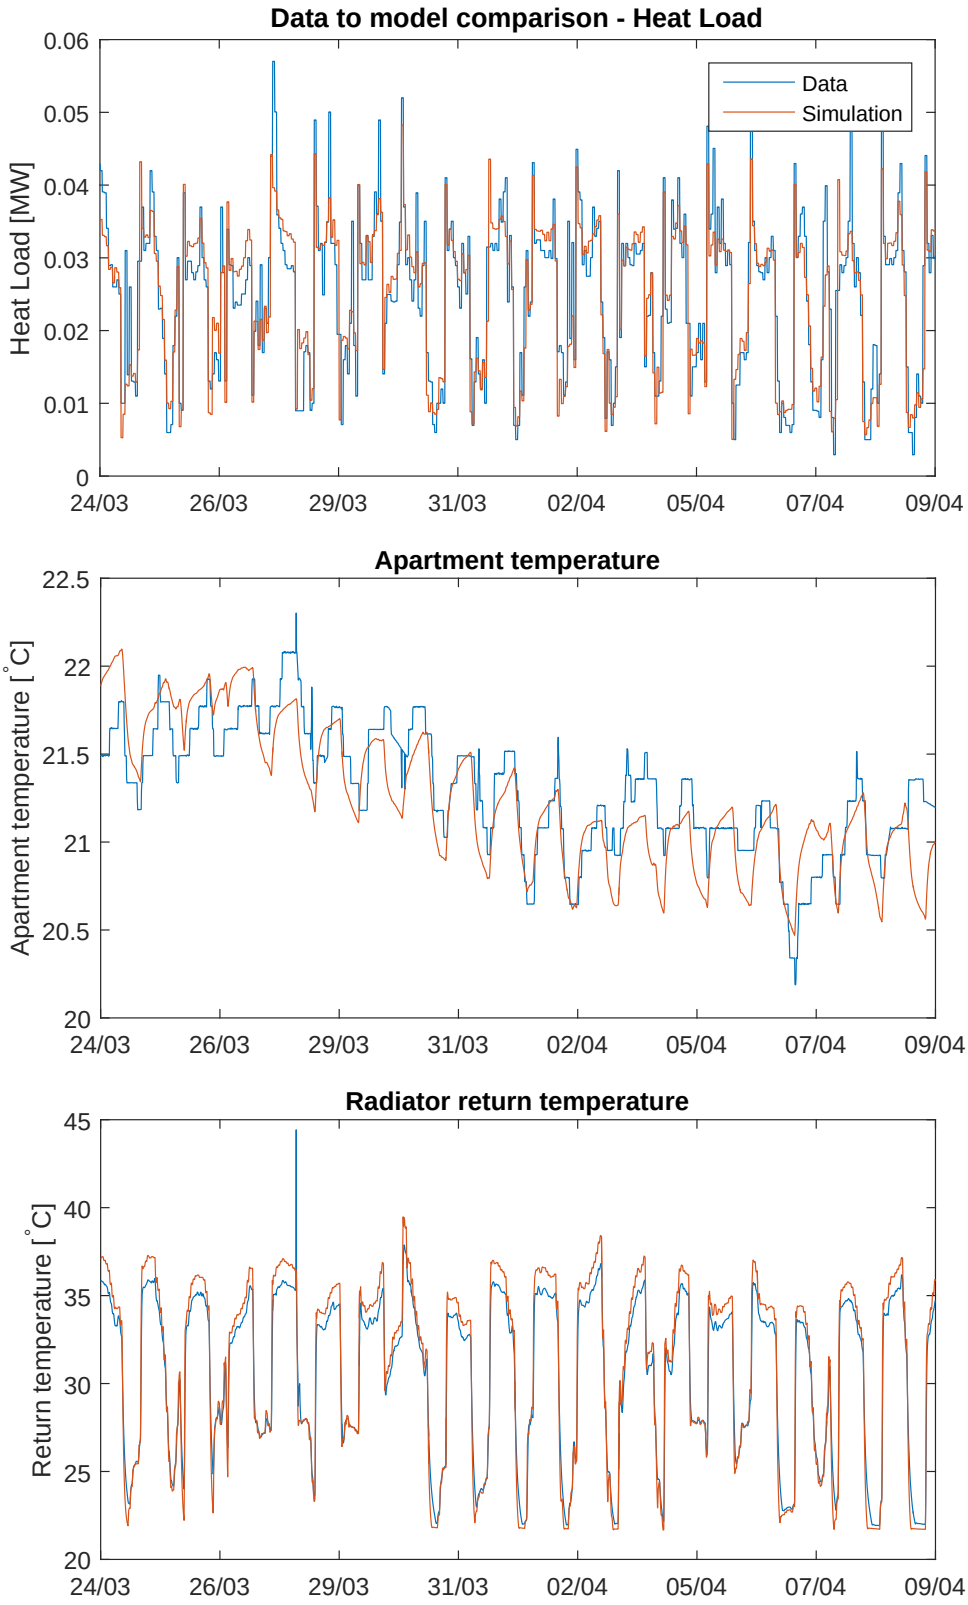

Figure 4.2: Comparison between model behaviour and recorded data

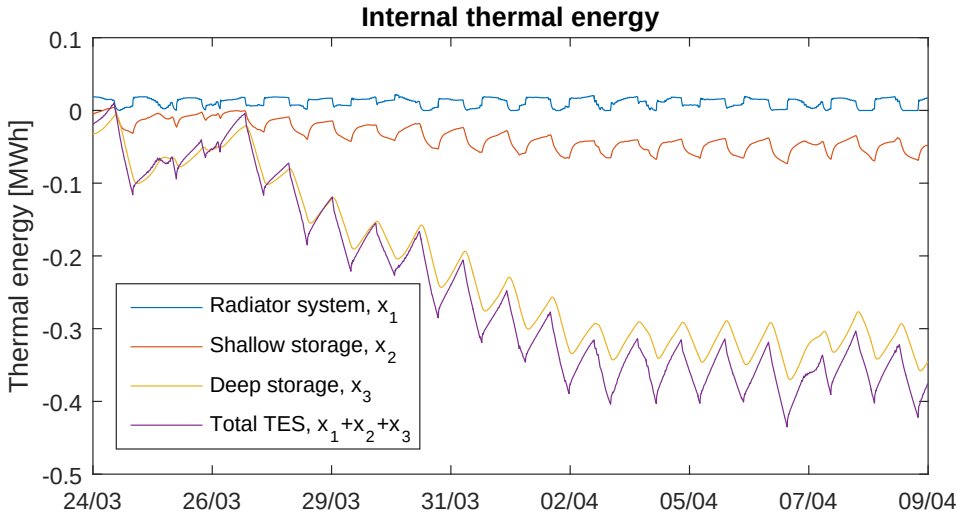

Figure 4.3: Simulated internal thermal energy of building A1, same dataset and timespan as displayed in Figure 4.2

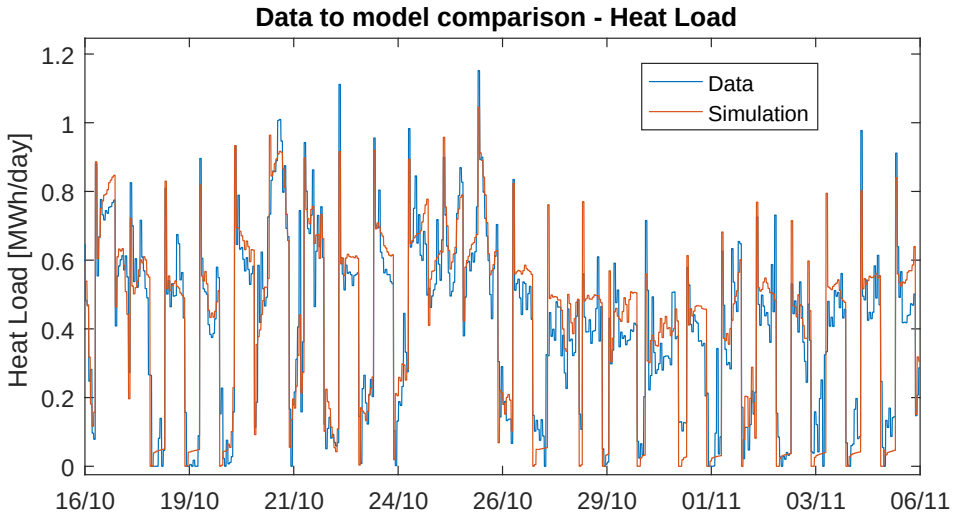

Figure 4.4: Comparison between recorded data and heat load prediction of model building A1, on a dataset the model was not trained on

Control Scheme	d_1	d_2	d_3	d_4	d_5	d_6
$\mathbf{cs_1}$: No control	0	0	0	0	0	0
$\mathbf{cs_2}$: T-module only	0.1	3.8571	0	0	0	0
$\mathbf{cs_3}$: P-module only	0	0	4.285	5	0.1758	0.2007
$\mathbf{cs_4}$: Both modules, $12h$ forecast	0.1	3.5714	5.7143	3.0304	0	0
$\mathbf{cs_5}$: Both modules, $48h$ forecast	0.1	1.6	2	4.9241	0.0584	0.4625

Table 4.1: Optimised parameters for control schemes

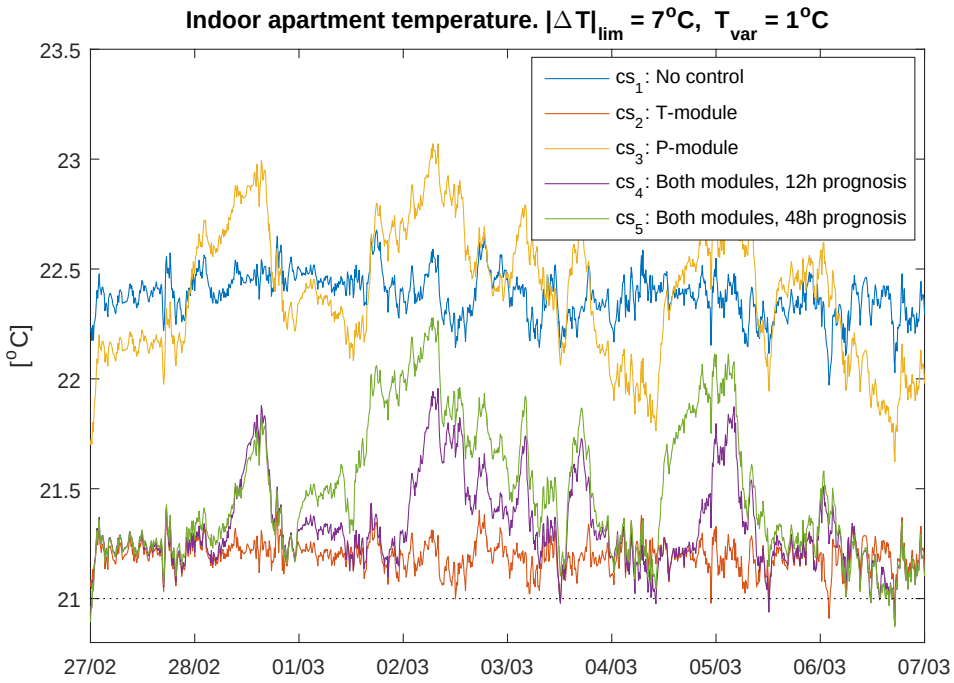

Figure 4.5: Illustration of how indoor temperature tends to vary with five different control systems

temperature is displayed in Figure 4.5. When only controlling based on the marginal price, the temperature oscillates around the simulated uncontrolled temperature. With only the T-module active, it tends to lie between 21 and 21.5°C. With both modules active the temperature behaviour shifts between staying relatively stable and quickly increasing or decreasing depending on price variations.

A summary of control system performance, characterised by Funds saved, heat saved and average indoor temperature above T_{min}, is displayed in Table 4.2. Equivalent data, with values taken as the average performance over all simulated buildings, is displayed in Table A.2 in Appendix A. First we note that T-module only control saves $\approx$ 10% heat, but only 6.5% Funds. This is because high indoor temperatures naturally coincide with lower heat costs. Meanwhile the total heat use increases when only controlling against marginal price (cs_3), but over 6% Funds are still saved. Overall the T- and P-modules independently perform relatively similar, especially in the most realistically implementable Case 1.

The effect of allowing both modules to act simultaneously is a slightly lower heat saving than when only controlling temperature, but significantly increased monetary savings to only having one module active. The two modules compete somewhat, with the price control leading to increased heat use while the temperature control sometimes inhibits load shifting, which can be seen in that the cumulative savings from having both modules active is smaller than the sum of each individually. The 48 hour forecast of cs_5 leads to slightly increased heat use and monetary savings compared to the 12 hour forecast of cs_4, but the difference is less than 10%

Specifier	Funds saved [%]	Heat saved [%]	$\mathbf{T_{mean} - T_{min}}[^\circ\mathbf{C}]$
Per control system (restriction case 1 and full dataset)			
$\mathbf{cs_1}$ (No control)	0	0	1.19
$\mathbf{cs_2}$ (T-module only)	6.5	9.9	0.22
$\mathbf{cs_3}$ (P-module only)	6.2	-3.7	1.63
$\mathbf{cs_4}$ ($12h$ forecast)	12.6	8.0	0.43
$\mathbf{cs_5}$ ($48h$ forecast)	13.5	7.3	0.49
Per restriction case (control system 5 and full dataset)			
Case 1 (Default restrictions)	13.5	7.3	0.49
Case 2 (Higher ΔT_{lim})	15.6	7.2	0.50
Case 3 (Higher $T_{var,lim}$)	14.1	5.2	0.70
Case 4 (Both higher)	19.3	4.1	0.83
Per dataset (control system 5 and restriction case 1)			
Spring 2013	5.7	1.2	0.41
Winter 2013 – 2014	12.4	8.5	0.40
Full 2013 – 2014	13.5	7.3	0.49

Table 4.2: Primary performance results, building A. $\mathbf{T_{mean} - T_{min}}$ refers to the mean indoor apartment temperature above the minimum acceptable temperature, in this instance $21^\circ C$.

either way.

Comparing different restriction cases, it is notable that increasing the control magnitude limit (ΔT_{lim}) alone leads to greater savings than only increasing the daily allowed variation ($T_{var,lim}$). The monetary benefit of increasing both simultaneously is greater than the sum of the benefits of increasing one at a time. This indicates that in case 2, the faster response is bottlenecked by strict limits while in case 3 the system is too slow to fully utilise the entire allowed temperature span. Comparing performance of control system 5 over different datasets it's clear just how much external factors, like weather conditions and marginal price profile, affect the control system performance. In order for the system to be effective, there must be sufficient price variations for the load shifting mechanism to offset the increased heat consumption. As a curio, Table A.3 in Appendix A shows how each control scheme performs when controlled against the average, instead of the marginal, heat generation price. In this case price control alone saves less than 1% Funds, while all systems utilising temperature control display very similar performance.

The best performing scheme, cs_5, was chosen for further study. General behaviour of the control signal and resulting apartment temperature for restriction

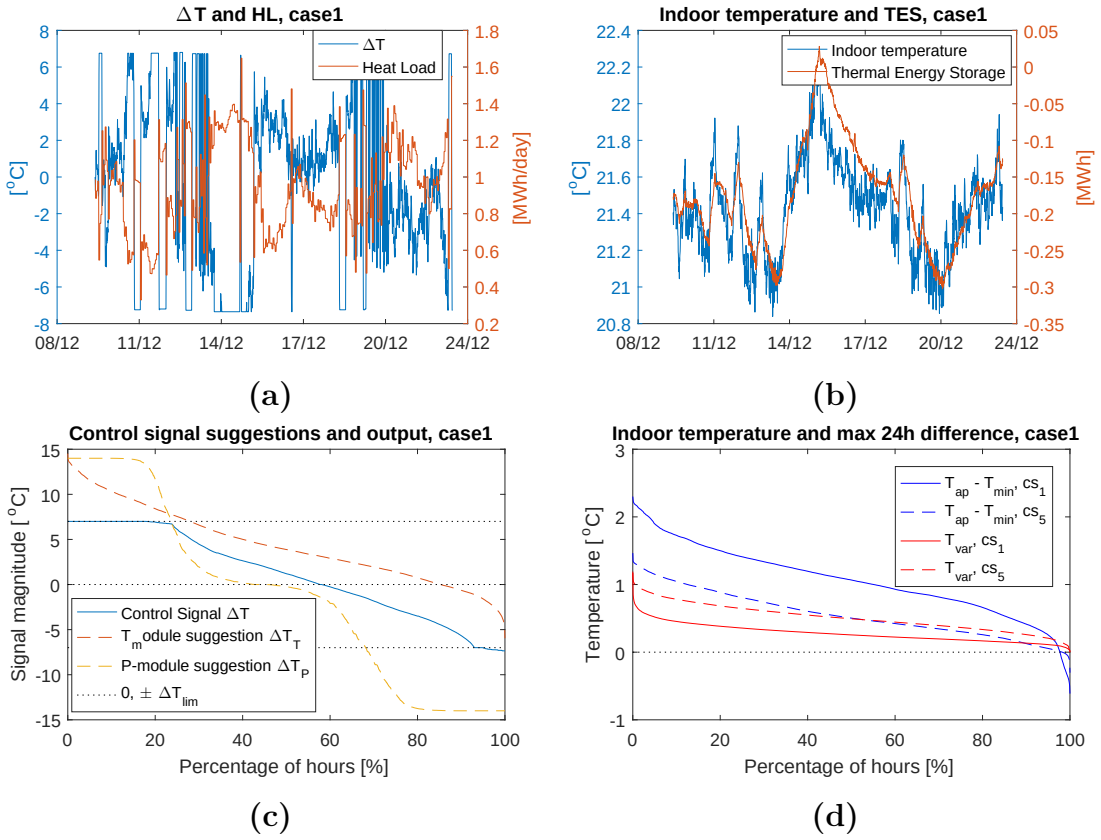

Figure 4.6: Illustration of behaviour of control system cs_5, restriction case 1. (a) shows a sample of how ΔT and HL tend to vary with time, (b) the resulting T_{ap} and simulated TES. Duration curves from the whole simulation period of ΔT and its module components are displayed in (c), (d) displays duration curves of the controlled building, compared to those from a simulation of the building with no control (cs_1).

cases 1 and 4 is presented in Figures 4.6 and 4.7. In case 1, T_{ap} stays within 21 to 22°C for 90% of the time with significantly lower overall temperature than the uncontrolled case despite higher daily variations. Stored TES follows T_{ap} relatively closely, but is somewhat delayed and occasionally continues to increase after T_{ap} peaks, indicating that the building's potential for thermal storage is not fully utilised during these timescales. In case 4 the control appears more deliberate, with temperature either remaining relatively constant or sharply increasing and decreasing in well defined cycles. This is related to the fact that $\Delta T = \pm\Delta T_{lim}$ 33% of the time for case 1, 48% of the time during case 4. This change is not unexpected, since the maximum magnitude of ΔT is increased by a factor 2 between the two cases, while that of ΔT_P is increased by a factor of 4.

The T- and P-modules are mostly in competition, they only share sign 36% of the time in case 1 and 34% in case 4. The duration curve of ΔT_P (P-module control suggestion) somewhat favours charging over discharging, possibly because of the formulation of P_Σ, while ΔT_T predominantly is positive. In both cases, the highest noted indoor temperature is the one which leads to $\Delta T_T \cdot \Delta T_{varlim} \approx \Delta T_{P,max}$.

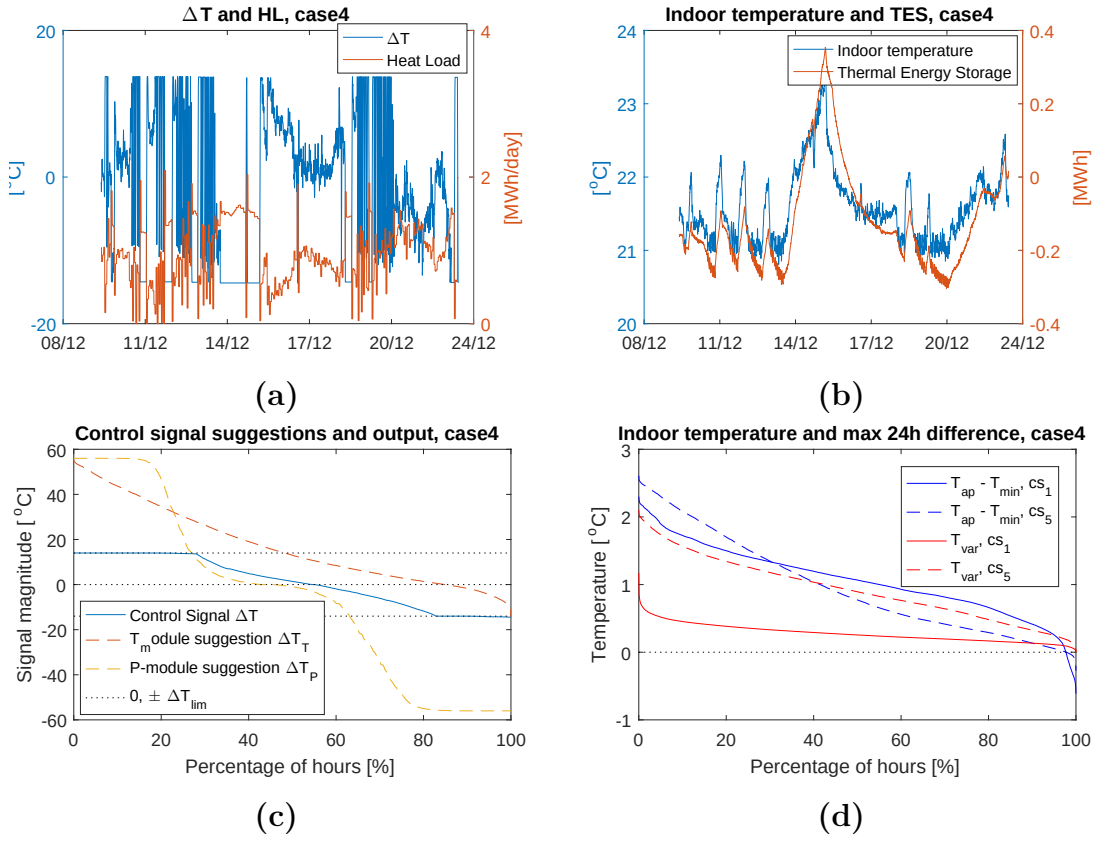

Figure 4.7: Illustration of behaviour of control system cs_5, restriction case 4.

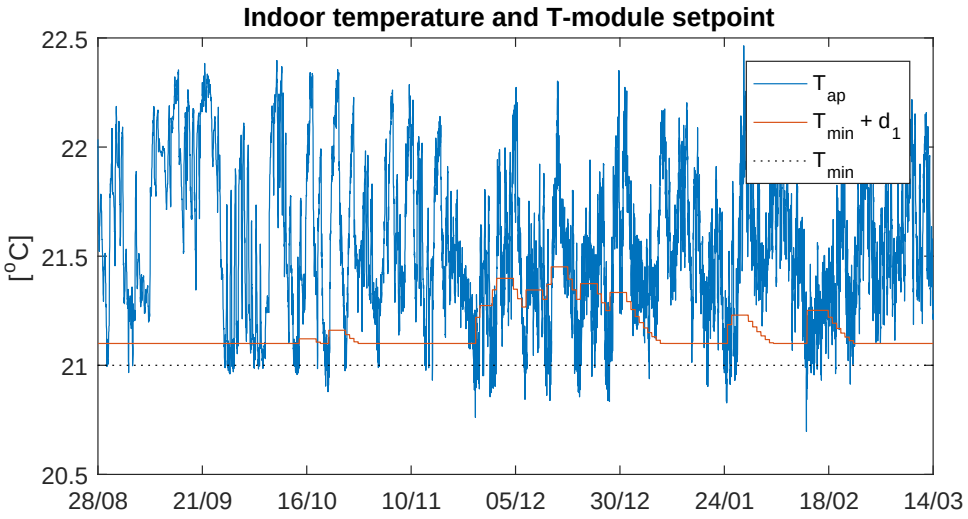

Figure 4.8: Illustration of how the T-module setpoint $T_{min} + d_1$ reacts to achieved indoor temperature. Note how the setpoint increases whenever T_{ap} stays below T_{min} for too long, and starts to decrease after a number of days with no temperature violations. This simulation was done with cs_5, Case 1.

With this information one could estimate the highest occurring indoor temperature from d_2 and d_3, and potentially create a formulation of the control system which remains within some desired bounds without the need for an explicit limit $T_{var,lim}$. Specifically, if $\Phi_2 = -1$, then $\Delta T_P = -\Delta T_T$ when $(T_{ap} - T_{min}) = \frac{d_3 \cdot T_{varlim}}{d_2}$.

Typical behaviour of the moving T-module setpoint $(T_{min} + d_1)$ is shown in Figure 4.8. This behaviour changes surprisingly little between restriction cases and control systems, and d_1 rarely exceeds 0.5.

4.3 Building response analysis

The simulation data was also used to make attempts at more generally characterising buildings used for thermal energy storage. Charge and discharge spans were identified from the simulation with cs_5 running on building A1 over the full dataset 2013-2014 with restriction case 1. A charge/discharge span was defined by ΔT being at least 80% of its maximum or minimum values continuously for a number of hours. Discharges proved difficult to analyse, since $T_{adjusted}$ in these cases often was above $15°C$ where the building no longer reacts to control input due to not providing any space heating, and so focus was put on charge cycles ($\Delta T < 0$). Characteristics across these spans were then compiled, specifically: Effective heat capacity $c_{p,eff}$ across the span, or change in total stored energy divided by change in indoor temperature; Time constant τ, as calculated from equation (2.2); Charge rate, or change in stored energy over time. For an ideal energy storage, $c_{p,eff}$ and charge rate is constant.

Identified charge rate, effective heat capacity $c_{p,eff}$ and time constant τ for

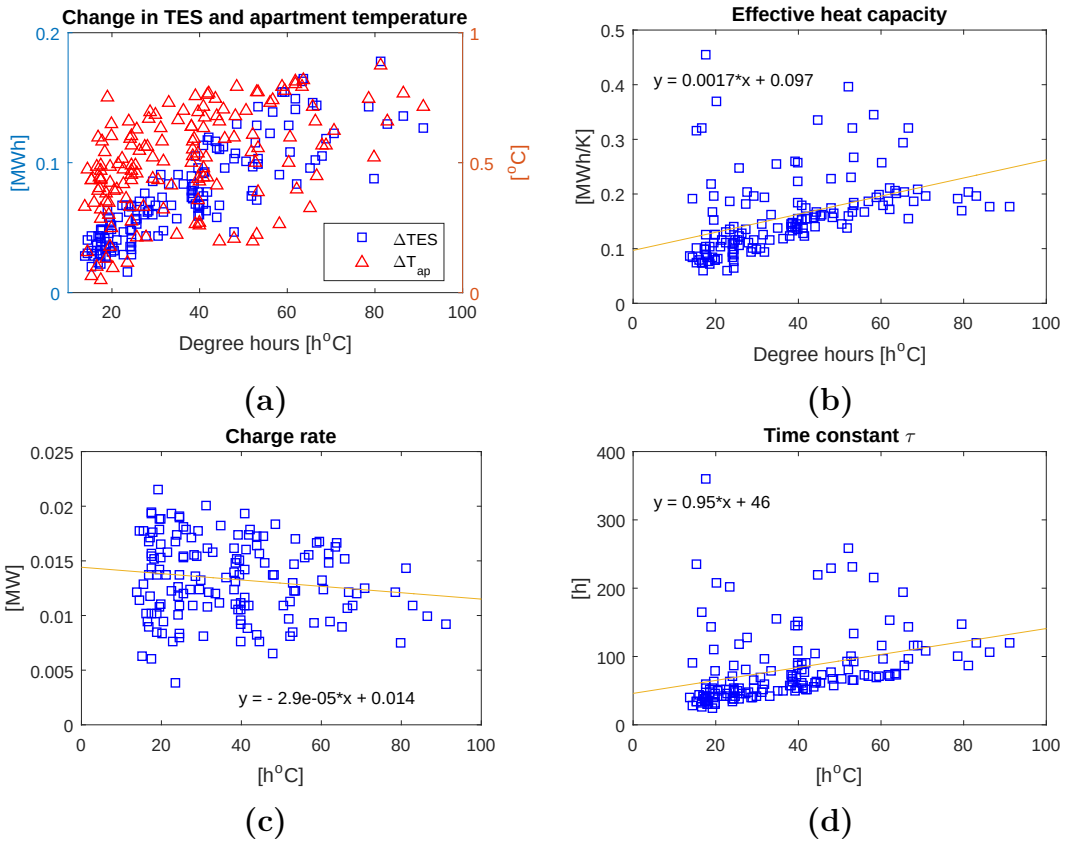

Figure 4.9: Scatter plot of apparent characteristics during long consecutive charge periods. Against charge duration in either hours or degree-hours the figures show: (a) Distribution of change in indoor temperature and internal energy during charge spans (b) Effective heat capacity (c) Rate of change of internal energy (charge rate) over the entire span (d) Time constant τ, as defined by equation (2.2)

each charge cycles are displayed in Figure 4.9. The horizontal axis, *Degree hours*, represents the charge duration (in hours) multiplied with the average absolute control signal magnitude. There are large variations in these characteristics, which decrease with cycle time. Only charge spans of at least 4 hours are displayed for clarity. The straight lines are linear least square fits. While it seems clear that both $c_{p,eff}$ and τ increase during longer charge cycles, there are too many exceptions to classify them all as outliers. Additionally the correlation between charge rate and span length is very weak despite significant variations. This all indicates that these characteristics depend on more than cycle duration. This is illustrated in Figure 4.10. $c_{p,eff}$ is plotted against the difference between the shallow and deep storage temperatures, T_{ap} and T_{deep}, at the start of the charge span and a strong correlation is found. Charge rate is plotted against outdoor temperature, and shows a curious behaviour. Above an outdoor temperature of $15°C$ charge rate starts to decrease dramatically because the setpoint curve reaches its minimum at that outdoor temperature, meaning its dependence on ΔT vanishes as $T_{outside}$ approaches $15 + \Delta T^{\circ}_{lim} C$. Conversely, charge rate decreases with decreasing outdoor temperatures primarily because the setpoint curve slope is less steep below $0°C$.

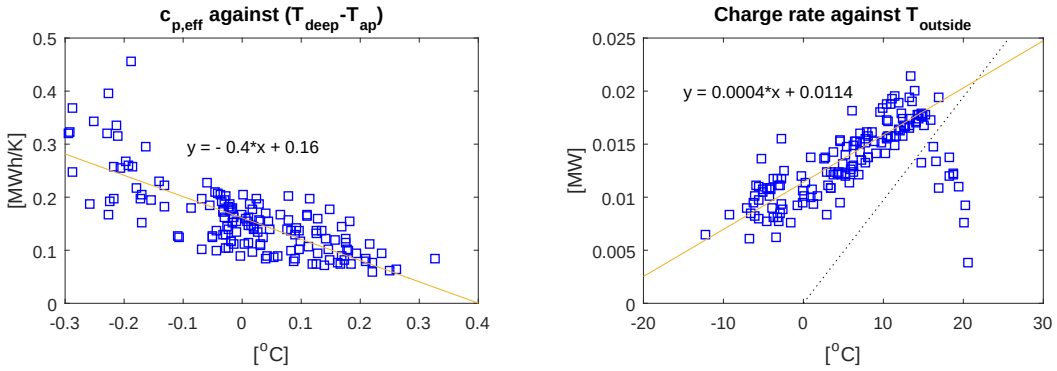

(a) (b)

Figure 4.10: Scatter plot showing: (a) Effective heat capacity over initial temperature difference between shallow and deep thermal storage (b) Charge rate over average outdoor temperature during charge span. The solid line is the linear regression line for all data points, excluding the ones below the dotted line.

	$\mathbf{D_{span}}$	$\mathbf{\Delta T_{avg}}$	$\mathbf{T_{deep} - T_{ap}}$	$\mathbf{T_{outside,avg}}$	$\mathbf{T_1 - T_{setpoint}}$
$\mathbf{\Delta T_{ap}}$	**0.607**	**0.369**	**0.413**	**0.288**	-0.110
ΔTES	**0.894**	**0.214**	-0.049	0.195	-0.131
$\mathbf{c_{p,eff}}$	0.132	-0.127	**-0.253**	0.108	0.098
τ	0.065	-0.105	-0.092	0.129	0.165
Charge rate	-0.179	**0.832**	-0.089	**0.374**	-0.134

Table 4.3: Regression coefficients indicating to what extent building characteristics depend on external factors during a charge. Coefficients with a magnitude above 0.2 are bold for highlighting purposes.

A more systematic attempt to quantify charge and building characteristics was made using MLR. Five parameters were chosen: Duration of charge (D_{span}), average signal strength during charge (ΔT_{avg}), difference between T_{ap} and T_{deep} at the start of the charge, average outdoor temperature during charge ($T_{outside,avg}$) and difference between average radiator temperature T_1 and radiator setpoint $T_{setpoint}$ at the start of the charge. Five characteristics were investigated: Change in T_{ap} during charge (ΔT_{ap}), change in TES during charge (ΔTES), $c_{p,eff}$, τ and rate of charge ($\Delta TES/D_{span}$). Data was collected for all spans over 1 hour with $|\Delta T| > 0.5 \cdot \Delta T_{lim}$. All data columns were centred and normalised to have a mean of 0 and standard deviation of 1, to facilitate simple comparison between regression coefficient magnitudes. The resulting coefficients are displayed in Table 4.3.

From Table 4.3 a few observations can be made.

- As long as a constant signal direction is maintained, duration of charge is more important than average charge strength to achieve a change in stored energy.
- Overall charge rate decreases somewhat for longer spans.
- Maximum effective heat capacity is achieved with long charge periods at small

control magnitudes, with a low initial deep storage temperature.

- Initial internal energy difference ($T_{deep} - T_{ap}$) has a significant impact on temperature response, and therefore also effective heat capacity, but barely impacts how much energy actually is stored during the charge or effective charge rate.
- Outdoor temperature has a noticeable impact on how the building responds to control.
- The initial radiator circuit state is often relevant, but never dominating. Interestingly, the relevance of initial radiator state does not increase if only shorter spans are considered.

5

Discussion & Conclusions

5.1 Discussion

The scope of this project was limited far beyond what would be required to create a realistic candidate for a control system that optimally utilises the thermal inertia of buildings to minimise their heating cost when considering the marginal generation cost of heat. Instead the aims were to predict how beneficial such a system can be, as well as provide useful pointers and highlight difficulties that can be encountered when designing such a system. There are also some inherent benefits to simple systems, meaning that expected performance of a rather minimal control system such as the one presented here can be interesting in themselves. A discussion based on these themes will be presented below.

5.1.1 Results

While the two goals of the control system – to reduce overall heat use and to temporarily raise indoor temperature in order to shift load from peak hours – seem to conflict with one another, the benefits from having all control parameters active is still very significant. One reason is that the two goals partly complement each other, since having a lower average indoor temperature means stored energy dissipates slower, increasing the utility of the storage. Increasing the available forecast from 12 to 48 hours does lead to increased savings, the gains are relatively small and might therefore in reality not be worth the risks associated with longer, less reliable, forecasts.

There are clear interactions between the control restrictions, meaning that these should be considered in tandem to maximise storage utility without compromising thermal comfort. Under a constant limit of control signal magnitude (ΔT_{lim}), there will be diminishing returns if the allowed daily variation (T_{varlim}) is increased to the point where the control system is too slow to make full use of its allowed temperature interval. If ΔT_{lim} is increased instead, the indoor temperature variations will be exceeded very quickly. There are still benefits to a faster system however, since the deep storage charge rate increases with apartment temperature. Therefore hitting a maximum indoor temperature quickly and remaining there for a number of hours will result in somewhat increased storage utility compared to slowly reaching that indoor temperature over the same number of hours. It therefore appears that

a higher control signal magnitude always is beneficial, but in reality there will be an upper practical limit to this parameter due to factors not considered here.

The actual savings by temperature-control alone is somewhat lower than what is expected from the literature and industrial implementations, even though a comparable amount of energy was saved in simulations. This can be explained by that the indoor apartment temperature naturally is higher during warmer parts of the year and day, coinciding with lower marginal prices in the DHN. It might be reasonable to vary the target indoor temperature with time of year to combat this, but primarily this result highlights that pure energy use minimisation is not the optimal choice for environmental or monetary savings.

5.1.2 Limitations

Several limitations regarding both available data, the chosen model and allowances of the control system that have to be taken into account when considering the achieved results. Among the most obvious are the limitations of the available data. Heat load was only available at an hourly scale, with two hours every day having to be estimated. Additionally no satisfactory estimate was found of the split of this load between the radiator and DHW subsystem. This, as well as the low amount of reliable indoor temperature readings, lowered the quality of the data available for model parameter fitting and meant that it often was difficult to determine whether simulation discrepancies were due to issues with the model or the data. The most significant resulting uncertainty from this is the deep storage heat capacity, since determining that parameter requires reliable temperature-response curves during control cycles with multiple characteristic time scales.

Internal homogeneity is assumed for the three energy storages of the building, which has a unique implication for each storage. For the radiator system, radiator return and supply temperatures are arguably more important than the average temperature, and so an estimate of the return temperature was used in addition to the supply temperature in this model. While this estimate was shown to give an acceptable prediction of the real return temperature against available data, there is no guarantee it remains realistic for all possible control behaviours. A 1D periodic model would not only be more predictable in this regard but also make it easier to utilise the logarithmic, as opposed to the herein used arithmetic, mean temperature difference for the radiator-to-room temperature exchange. The shallow energy storage currently represents both the indoor air, which absorbs radiator heat by convection, and the outermost wall layer, which absorbs radiation energy. By splitting these into two layers with individual transfer rates from the radiators, it is possible that nonlinear radiator transfer effects could be observed properly. Finally, using a single container for the building core severely limits how long and short term storage dynamics can differ. A more detailed construction could include multiple deep storage modules connected in series, each representing sections of the wall at different depths.

No tenant behaviour, secondary building effects – such as ventilation system

responses or thermostat throttling – or weather effects beyond outdoor temperature were taken into account. This means a lot of opportunities considered by modern advanced control systems, such as predicting solar radiation or utilising zone control, were not included. It also completely ignores the difficulties that arise from controlling an unpredictable syste, such as tenants opening windows or using heat-generating utensils for long periods of time. The control system is quite minimal, and is not expected to perform as well on a real building as on this model system. Several studies, summarised in [6], indicate that classical controllers not necessarily are sufficient for maintaining thermal comfort in buildings. Given these studies it seems likely that the herein used P-controller with a gradually adjusting setpoint would be too slow to control around erratic tenant behaviour. Additionally, the P-module consists of four parameters and was trained to make reasonable choices most of the time, but inspection of price-control curves reveals it can make seemingly poor decisions when the price curve exhibits certain specific behaviours.

However, these limitations and uncertainties do not necessarily have a major impact on the overall findings. In a realistic scenario there are more complications a control system has to take into account, but there are also more opportunities to make use of excess heat from for example sunlight or food preparation. It is difficult to say if a more accomplished control system in a more complicated environment would perform slightly better or worse than a simple system on a simple model, but the trade off between maintaining a lower indoor temperature and shifting the buildings heat load should remain fairly consistent. Thus the main conclusion from this work, that a system which takes marginal heat price into account will achieve a lower overall heating cost than a system which simply minimises total heat use, is not invalidated.

5.1.3 Comments Regarding Future Implementation

The main conclusion that can be drawn from the results presented in section 4.3 is that instantaneous apartment temperature alone is not a sufficient indicator for how a building will react when its temperature sensor is adjusted. One of the biggest reasons is that the overall building internal energy varies much slower than indoor temperature and so should be regarded separately, potentially estimated from the apartment temperature history. Since the building core slowly absorbs energy based on the internal temperature difference, maintaining a certain indoor temperature for a duration of time is just as important as reaching that temperature. Secondly, the varying slope of the internal setpoint curve means that the same control signal will elicit a different response depending on current outdoor temperature. While the setpoint curve non-linearity does improve indoor climate overall in uncontrolled buildings, it does not necessarily aid a control system which controls towards a below-nominal indoor temperature. Especially if the control system takes weather effects into account, giving it direct radiator setpoint control will likely be preferable to the herein assumed method. The most significant addition to the large-scale analysis by [12] is that their estimated storage capacity limitation likely is a lower estimate due to their test cycle not allowing the deep building storage to fully charge

in either direction. It is possible that a sufficiently advanced control system could improve the effective storage capacity of buildings compared to their estimate.

Optimising control according to marginal heat price is not a trivial task. For example, if a long discharge period is expected it is important not only that a high apartment temperature is reached beforehand, but that it has been held for some time to ensure the deep storage temperature also is high. On the other hand, maintaining high indoor temperature for too long beforehand will increase the overall heat use and is therefore not advisable either. Additionally, when considering a real building with secondary effects and uneven tenant behaviour, reaching a desired temperature with minimum energy expenditure is a complicated problem to solve in itself. Since there already are commercially successful products that tackle the second problem, an elegant implementation of a total cost-optimising control routine would be a two stage system, where the first calculates the indoor temperature which would make optimal use of the buildings thermal inertia, and the second attempts to reach that temperature while minimising overall energy use.

5.2 Conclusion

The complexity of the dynamical building model and control system used in this project was somewhat restricted by various constraints, but the end result is still a valuable demonstration of the potential of utilising a building's thermal inertia for load shifting within a district heating network. When purchasing heat at marginal generation costs, a relatively simple control system with a 48 hour price forecast was able to save over 13% monetary units using a combination of load shifting and temperature stabilisation with minimal impact on thermal comfort. The effective storage capacity limitations previously estimated by [12] were shown to be safe-side estimates, but a well considered control system will be required to fully utilise the potential of the buildings. The estimated power limitation will likely vary with outdoor temperature, thus an alternative implementation of the control signal might be advisable.

Even when not considering load shifting, it is not uncommon that buildings are fitted with monitoring systems that allow for indoor temperature feedback control for improved thermal comfort and energy savings. Because of the rapid technological advance within the area, a number of new control methods are not unlikely to appear in the near future. Therefore the author strongly recommends that all planned building control units are designed to be reprogrammable and capable of receiving generic signals, to facilitate simple implementation of any future control systems.

Bibliography

[1] Energiläget i siffror 2016. http://www.energimyndigheten.se/globalassets/statistik/overgripande-rapporter/energilaget-i-siffror-2016_160218.xls. Accessed 2016-05-13.

[2] Energyplus. https://energyplus.net. Accessed 2016-05-09.

[3] Ida indoor climate and energy. www.equa.se/en/ida-ice. Accessed 2016-05-09.

[4] Kabona ecopilot public website. http://www.ecopilot.com/references/. Accessed 2016-05-12.

[5] Smart heat building, a technical overview. https://www.noda.se/static/pdf/NODA_smartheatbuildning_en.pdf. Accessed 2016-05-11.

[6] Abdul Afram and Farrokh Janabi-Sharifi. Theory and applications of hvac control systems – a review of model predictive control (mpc). *Building and Environment*, 72:343 – 355, 2014.

[7] J.R. Dormand and P.J. Prince. A family of embedded runge-kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1):19 – 26, 1980.

[8] Henrik Gadd and Sven Werner. Daily heat load variations in swedish district heating systems. *Applied Energy*, 106:47 – 55, 2013.

[9] David F. Griffiths and D. J. Higham. *Numerical methods for ordinary differential equations: initial value problems*. Springer, New York;London;, 2010.

[10] Johan Heier, Chris Bales, and Viktoria Martin. Combining thermal energy storage with buildings – a review. *Renewable and Sustainable Energy Reviews*, 42:1305 – 1325, 2015.

[11] R. Kaiser, D.U. Sauer, A. Armbruster, G. Bopp, and H.G. Puls. New concepts for system design and operation control of photovoltaic systems. In *14th European Photovoltaic Solar Energy Conference*, July 1997.

[12] Johan Kensby. *Buildings as thermal energy storage*. PhD thesis, Chalmers University of Technology, 2015.

[13] Bengt Lennartson. *Reglerteknikens grunder*. Studentlitteratur, Lund, 4. uppl. edition, 2002.

[14] Lennart Ljung and Torkel Glad. *Modeling of dynamic systems*. Prentice Hall, Englewood Cliffs, 1994.

[15] Barbara Mayer, Michaela Killian, and Martin Kozek. Management of hybrid energy supply systems in buildings using mixed-integer model predictive control. *Energy Conversion and Management*, 98:470 – 483, 2015.

[16] Tahir Mehmood and Bilal Ahmed. The diversity in the applications of partial least squares: an overview. *Journal of Chemometrics*, 30(1):4–17, 1 2016.

[17] John A. Rice. *Mathematical statistics and data analysis*. Thomson Brooks/Cole, Belmont, Calif, 3. edition, 2007.

[18] Anders Rönneblad, Jan Forslund, and Ronny Andersson. Byt styrsystem i miljonprogramsbostäder. *Samhällsbyggaren*, (5):13 – 17, 2011.

[19] Alexander Schirrer, Markus Brandstetter, Ines Leobner, Stefan Hauer, and Martin Kozek. Nonlinear model predictive control for a heating and cooling system of a low-energy office building. *Energy and Buildings*, 125:86 – 98, 2016.

[20] Vittorio Verda and Francesco Colella. Primary energy savings through thermal storage in district heating networks. *Energy*, 36(7):4278 – 4286, 2011.

[21] Lars Wadsö, Jonathan Karlsson, Anders Rönneblad, Ronny Andersson, Eva-Lotta Kurkinen, Mats Emborg, Mats Öberg, and Ulf Ohlsson. Energy saving through the utilization of the thermal behaviour of heavy buildings, based on new materials, building frameworks and heat storage systems. Technical report, Cerbof - Centrum för Energi- och Resurseffektivitet i Byggande och Förvaltning, 2012. Accessed 2016-05-08 from `http://www.byggnadsmaterial.lth.se/forskning/cerbof_projekt/`.

[22] M Wahde. *Biologically inspired optimization methods: An introduction*, chapter 5. WIT Press, 2008.

[23] Shengwei Wang and Zhenjun Ma. Supervisory and optimal control of building hvac systems: A review. *HVAC&R Research*, 14(1):3–32, 2008.

A

Appendix 1

	A1	**A2**	**A3**	**A4**	**A5**
$\mathbf{c_1}$	963.58	1126.5	1284.8	800	722.685
$\mathbf{c_2}$	20.958	13.743	15.718	20	15.718
$\mathbf{c_3}$	3.317	2.5214	2.4878	3.5	4.4227
$\mathbf{c_4}$	0.1728	0.21317	0.21656	0.1385	0.0392
$\mathbf{c_5}$	0.0368	0.040045	0.039548	0.0688	0.1238
$\mathbf{c_6}$	0.74	0.70387	0.67393	0.148	0.3427
$\mathbf{c_7}$	0.0477	0.051137	0.047993	0.04	0.0429
$\mathbf{c_8}$	5	5	5	5	5
$\mathbf{c_9}$	1	1	1	0.6	1
$\mathbf{c_{10}}$	1.25	1.613	1.6667	1.25	1.096
$\mathbf{c_{11}}$	130	130	130	130	130
$\mathbf{T_{ap,mean} - T_{min}}[°\mathbf{C}]$	1.07	1.32	1.22	1.75	1.32

Table A.1: Model parameters of buildings used during control system evaluation, and average indoor temperature during simulation period 2 with no temperature control. For parameter descriptions, see Table 3.2. Note that transfer rate coefficients c_4 through c_7 are given in $\frac{MWh}{day,K}$ instead of $\frac{MW}{K}$ in this table, and c_{11} in *day* instead of *h*.

	Funds saved [%]	**Heat saved** [%]	$\mathbf{T_{mean} - T_{min}}[°\mathbf{C}]$
Per restriction case			
$\mathbf{cs_1}$	0	0	1.4031
$\mathbf{cs_2}$	8.5141	11.7504	0.2455
$\mathbf{cs_3}$	5.6843	-3.1843	1.8031
$\mathbf{cs_4}$	13.7279	9.7787	0.4651
$\mathbf{cs_5}$	14.5701	9.1659	0.5148
Per restriction case			
Case 1	14.5701	9.1659	0.5148
Case 2	16.6826	9.2843	0.5101
Case 3	15.2758	6.8638	0.7546
Case 4	20.1892	6.1422	0.8406
Per restriction case			
Spring 2013	6.8432	3.3032	0.4709
Winter 2013 – 2014	13.5031	10.3420	0.4434
Full 2013 – 2014	14.5701	9.1659	0.5148

Table A.2: Primary performance results, all buildings

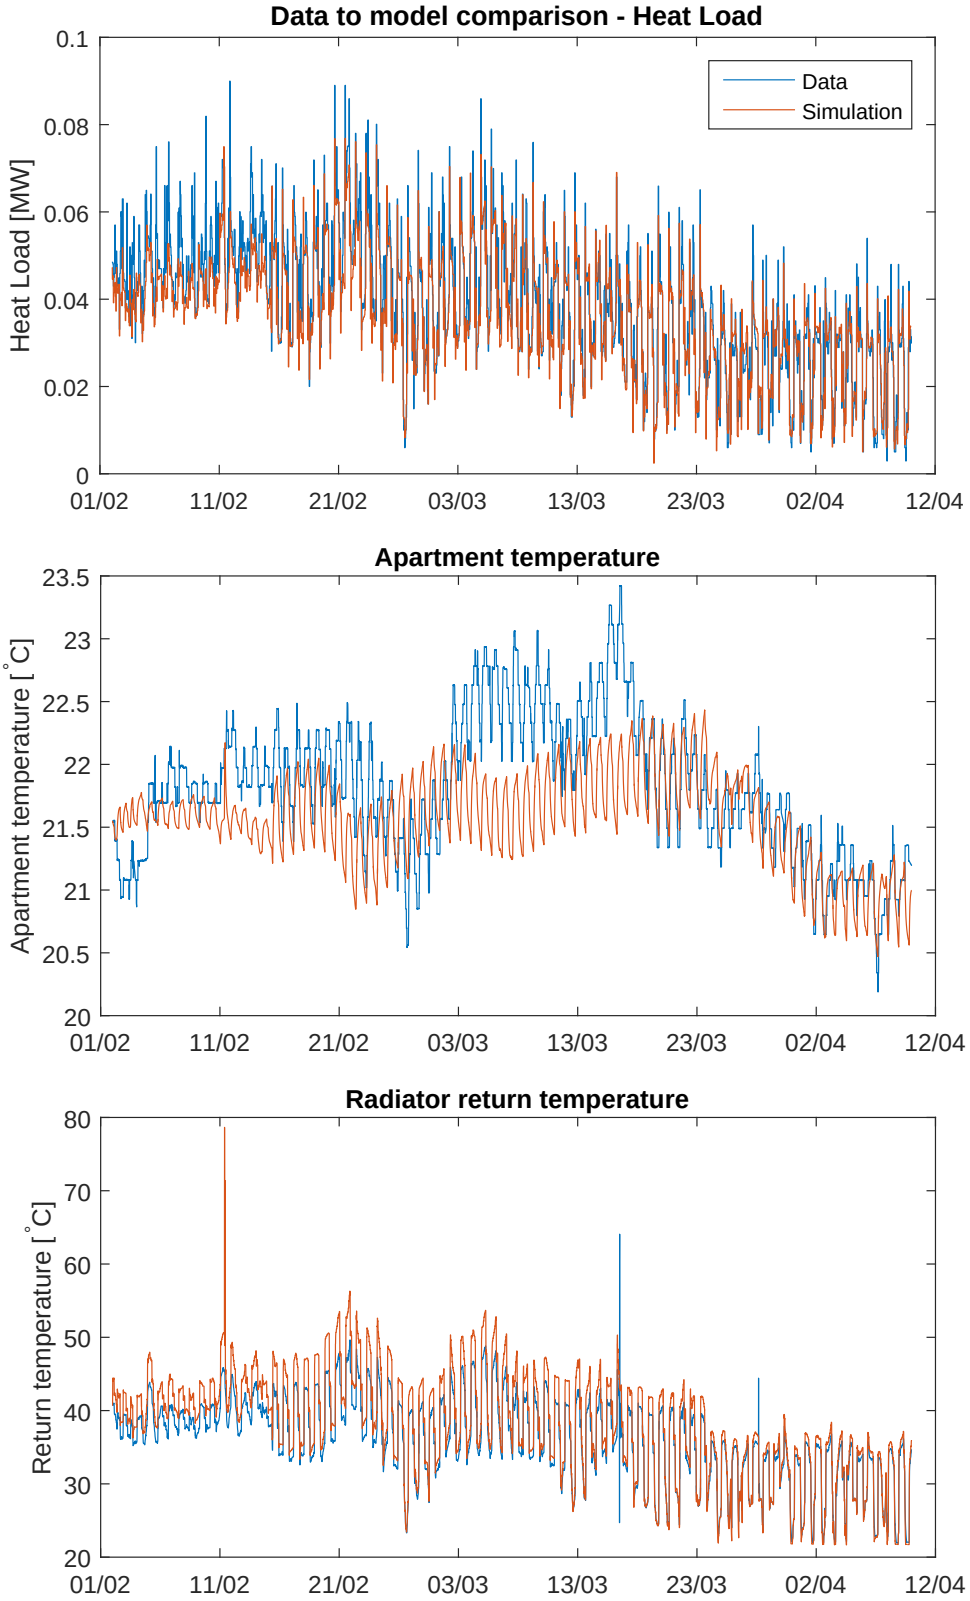

Figure A.1: Comparison between model behaviour and recorded data, entire simulation period

Funds saved [%]	**Case 1**	**Case 2**	**Case 3**	**Case 4**
$\mathbf{cs_1}$	0	0	0	0
$\mathbf{cs_2}$	10.0229	10.4496	10.0229	10.4496
$\mathbf{cs_3}$	0.8899	0.1055	0.9816	0.8392
$\mathbf{cs_4}$	10.9203	11.5089	11.0824	11.7680
$\mathbf{cs_5}$	10.3312	11.3044	10.4819	11.6309

Table A.3: Performance of control systems during 2013/2014, when controlling against average price instead of marginal