Master's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Engineering Cybernetics

Axel Tveiten Bech

Distributed MPC for Peak Power Reduction in Smart Homes

Master's thesis in Cybernetics and Robotics
Supervisor: Sebastien Gros
June 2022

NTNU
Norwegian University of
Science and Technology

Axel Tveiten Bech

Distributed MPC for Peak Power Reduction in Smart Homes

Master's thesis in Cybernetics and Robotics
Supervisor: Sebastien Gros
June 2022

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Engineering Cybernetics

NTNU
Norwegian University of
Science and Technology

Abstract

In the coming years, electricity demand and non-dispatchable power production are projected to increase. Improved consumer flexibility is needed to address the increase in production volatility. A peak power tariff will be introduced for residential electricity consumers in Norway in 2022. This tariff creates an economic incentive for residential consumers to reduce their peak power demand.

Research in model predictive control (MPC) methods for residential building power use mostly focus on reductions in total power use and cost. This thesis investigates model predictive control methods for reducing the combined peak power consumption in groups of houses by controlling heat-pump use. A distributed MPC approach, based on dual decomposition, allows for both parallel computation and coordination between houses to reduce joint power peaks while mostly preserving the privacy and independence of each house.

Simulation of the control schemes show that the distributed MPC managed to reduce momentary and hourly peak power, but at the expense of comfort in the houses. Further research should refine the peak reduction formulations, implement a more realistic room temperature model, and investigate the feasibility of residential power coalitions.

Sammendrag

I de kommende årene anslås elektrisitetsetterspørselen og andelen ikke-kontrollerbar kraftproduksjon å øke. Forbedret forbrukerfleksibilitet er nødvendig for å imøtekomme økt produksjonsvolatilitet. I Norge i 2022 skal et nytt 'kapasitetsledd' erstatte det tidligere fastleddet i nettleia. Dette gir husstander et økonomisk insentiv til å redusere toppstrømbehovet deres.

Forskning i modellprediktiv kontroll (MPC) av energibruk i boliger fokuserer i hovedsak på reduksjoner i totalt energibruk og kostnader. Denne oppgaven undersøker modellprediktive kontrollmetoder for å redusere det kombinerte maksimale strømforbruket i grupper av hus ved å kontrollere varmepumpebruken. En distribuert MPC-tilnærming, basert på 'dual decomposition', tillater både parallelle utregninger og koordinering mellom hus for å redusere felles effekttopper, samtidig som privatvern og uavhengigheten til hvert hus tas hensyn til.

Simulering av kontrollmetodene viser at den distribuerte MPC klarte å redusere momentant og timesmessing toppforbruk, men på bekostning av komforten i husene. Videre forskning bør videreutvikle toppreduksjonsformuleringene, implementere en mer realistisk romtemperaturmodell og undersøke gjennomførbarheten av kombinerte strømavtaler for flere boliger.

Preface

This master's thesis in the course TTK4900 Engineering Cybernetics was written in the spring semester of 2022 by Axel Tveiten Bech. It is the final course in the 5-year master's degree programme, *Cybernetics and Robotics*, at the Norwegian University of Science and Technology (NTNU).

This thesis is related to a larger project, POWIOT. The project uses commercial smart-house technology to implement advanced control algorithms for a house in Trondheim. One of the aims of the project is to use smart control of heating to reduce costs and address current issues in power distribution. The project is led by Prof. Sebastien Gros, who is also the supervisor for this thesis.

I am grateful to Sebastien Gros for providing invaluable feedback, support, and discussions. The data collection software he has running on his house provided much of the external data used in the simulations in this thesis. Additionally, the control formulations that are explored in this thesis, are based on his original house control formulations.

I would also like to thank Kang Qui for providing feedback on the report.

Contents

Figures

Tables

Chapter 1

Introduction

1.1 Problem Description

This thesis investigates control methods to reduce peak power consumption in residential power coalitions. The control methodologies employed are centralized-, decentralized-, and distributed model predictive control, where the distributed control formulation is obtained using dual decomposition. A residential power coalition is an imagined scenario where multiple residential homes band together to negotiate a common agreement with a power provider. Crucially, the peak power cost would be calculated from the highest combined power consumption in the coalition.

1.2 POWIOT

This thesis expands upon the control aspect of the 'POWIOT' project, led by Prof. Sebastien Gros. The aim of the POWIOT project is to implement a model predictive control (MPC) algorithm to control home heating in a way that reduces power consumption and power costs, feeding the MPC with data collected by a Raspberry Pi using smart-home IoT technology. Information about the daily electricity spot market and a temperature model of the house allows the algorithm to exploit the inherent thermal capacity of the house to e.g. pre-heat the home prior to a spike in electricity prices. Large-scale adoption of such a solution might help spread out total energy consumption close to spikes since heating accounts for a significant portion of energy usage.

Figure 1.1 shows an overview of the system currently in place in Trondheim, in the professor's house. A Raspberry Pi collects information about power consumption, spot prices, indoor temperature, heat pump settings, and weather forecasts. This data is accessed by another computer running a moving horizon estimator (MHE) and the MPC. The computer, as of now the professor's laptop, sends commands to the heat pump. Access to the heat pumps, which also includes indoor temperature measurements, is provided by two Sensibo Sky gate-

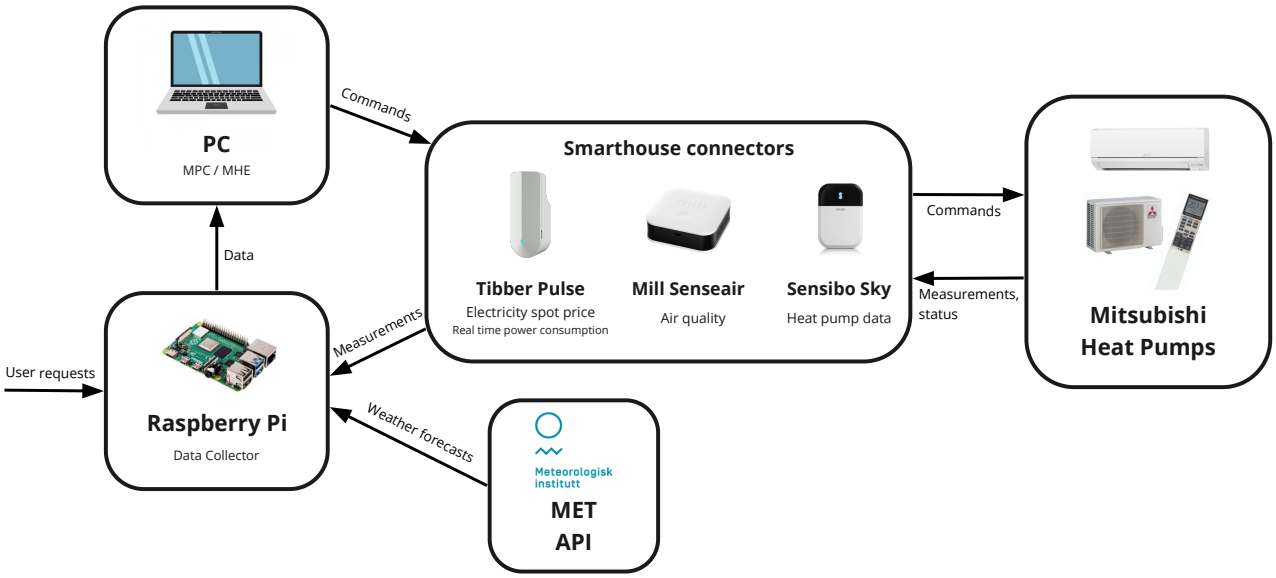

Figure 1.1: Hardware overview of the system

ways. A Tibber Pulse gateway makes the real-time power consumption available through an application programmable interface (API). The same API is used to access daily electricity prices but does not depend on the Tibber Pulse device. The MET API provides weather forecasts from The Norwegian Meteorological Institute. The model used in the MPC has been generated using system identification methods (SYSID).

1.3 Motivation

1.3.1 Challenges in Power Markets

The electrical grid consists of two parts. The transmission grid is where the power is produced and transported. The voltage is higher because for transporting the same amount of power, a higher electrical current gives more losses than high voltage. The power distribution grid is where the electricity is distributed to consumers. It has a lower voltage, closer to that of most electrical appliances. In Norway, the power consumption of especially residential consumers varies much with both time of year and time of day [1]. If consumer demands are not met during peak power consumption, brown-outs or black-outs can occur. Because of this, the electrical grid is typically oversized compared to the average consumption.

Buildings account for 40 % of energy consumption in the EU [2]. Household

energy demands are likely to increase with more electric cars and increased use of electric heating solutions like heat pumps. Energy demands in industry and the transportation sector are also projected to increase [3]. Additionally, power production in most countries is shifting towards renewable energy sources. Some of these, like solar and wind power, can have a fluctuating output that depends on weather conditions. The fluctuations in consumption and power production make the need for smart solutions to ensure energy demands do not exceed production.

A Statnett report from 2018 details some predicted developments in the Nordic power market [4]. The widespread adoption of advanced metering systems (AMS), increased automation, and more aggressive price signaling will make it more attractive for power consumers to adapt to the production of power instead of the other way round. At the moment, however, consumer flexibility is underutilized as a solution to discrepancies between power production and consumption. They also note that space-heating accounts for 64 % of household power consumption in Norway, which in addition to the many available options for thermal energy storage, makes heating an ideal source of household consumer flexibility.

1.3.2 New Capacity Cost

Transmission system operators (TSOs) are responsible for transmission grids in a region of a country. These companies are different from electricity providers, which users pay for the electricity itself. In Norway, users pay a 'nettleie', grid cost, to use the power infrastructure their local TSO is responsible for. The grid cost consists of a fixed cost (fastleddet), as well as a cost based on consumption (energileddet). The consumption cost is fixed per kWh, and independent of the spot-price-based cost paid to an electricity provider.

From July 2022, Norwegian transmission system operators will change their price models. The previous fixed cost will be replaced by a 'kapasitetsledd', a capacity cost [5]. Figure 1.2 shows an example of the new pricing model. The energy cost stays mostly the same in the new pricing model but is now cheaper at night. The new day-night shift is not included in the models in this thesis. The capacity cost is calculated each month, based on the hour the household consumed the most energy. That is, momentarily high power consumption may not put a household in a higher price group, but maintaining high consumption over the course of an hour will. In the example, consuming between 5-10 kWh in an hour will put a household in capacity price group 2, which is 115 kr more expensive per month than group 1. Crucially, this new model gives users an economic incentive to reduce peak power consumption.

1.3.3 Residential Power Coalitions

In many countries, large-scale power consumers have power agreements in which distribution costs are based on their maximum peak load. This incentivizes peak shaving, in which power peaks are reduced. This means less money has to be spent on expanding the power distribution network to accommodate peak demand. The

Energiledd	Energiledd øre/kWh eks avgifter	Energiledd øre/kWh (jan-mars) inkl. avgifter	Energiledd øre/kWh (april-des) inkl. avgifter
Dag (0600-2200)	18	34,89	43,01
Natt (2200-0600)	10	24,89	33,01
Kapasitetsledd	**kW**	**Kr/mnd (inkl. mva)**	**Kr/år (inkl. mva)**
Trinn 1	0-5	175	2100
Trinn 2	5-10	290	3480
Trinn 3	10-15	405	4860
Trinn 4	15-20	520	6240
Trinn 5	20-25	635	7620

Figure 1.2: New power price model [6]

Statnett report [4] predicts smaller power consumers may join together in aggregates to improve sway over markets. In a scenario in which a coalition of multiple homes, or an apartment complex, can prove they have the same control over total power consumption as a typical large industrial consumer, they may be granted an agreement with power distribution networks similar to that of e.g. large industrial consumers. In that case, they could employ similar peak shaving strategies to reduce distribution costs. Their aggregated peak cost may be similar to the capacity tariff being introduced to Norwegian homes in 2022, where cost is based on the highest hourly power consumption(s). Or it may be similar to how larger power customers pay for their peaks, which considers momentary peaks.

Chapter 2

Background

2.1 Review

2.1.1 Model Predictive Control in Home Energy Management

Early research on home energy control employed simple control algorithms like proportional-integral-derivative- (PID) and on-off controllers. On-off controllers are still very common in thermostats. They deploy temperature hysteresis, where they don't switch to on or off until the temperature has changed a little past the set point. The oscillations this causes means that the temperature will overshoot the desired temperature, and will not stabilize completely [7]. PID controllers are better, but still have difficulties handling the non-linear and noisy behaviour of indoor temperature [7]. They also do not consider desired comfort levels in their environments, and possess no information about the system they control.

More recent projects typically favour model predictive control (MPC) algorithms [7]. Using a plant model that can be obtained through e.g. system identification (SYSID), MPC algorithms iteratively solve a constrained finite-horizon optimization problem based on the system model. This allows it to provide optimal control for the current time step while taking into account the future behaviour of the system. The working principles of model predictive control is described in more detail in 2.2.1. Following its introduction in the 1980s, model predictive control was mostly used in petro-chemical processes. This was because the processes ideally operated close to the systems limits, or constraints, which MPC handles well [8]. Figure 2.1 shows the working principle of the algorithm.

Model predictive control is very well suited for building heating control [9]. Multiple demands, like comfort level, operating costs, and peak demand can be accounted for as optimization criteria or constraints. Additionally, its predictive nature allows it to take into consideration the inherent thermal storage properties of buildings, allowing it to plan hours or even days ahead. Advances in the availability of both processing power and information collected from buildings further compound the usability of the solution. However, one risks disenfranchising the inhabitants of a building if they are not given a say in how the system should

operate, or informed what it is planning [9].

Building management systems, which automate many facets of a building, including heating, ventilation, and air conditioning (HVAC) systems, are mostly deployed in large buildings. The research on MPC for HVAC reflects this, but there is increasing research on deploying the system in residential smart buildings [10]. The same review also states that the performance of the MPC can be improved by better estimation and prediction of disturbances such as humidity and solar irradiation. In one Italian study not dissimilar to the POWIOT project, modern IoT and ICT solutions were used to deploy an MPC on the HVAC systems in a smart building [11]. This included using a cloud based architecture for the system design, as well as taking advantage of external application programmable interfaces (APIs) to collect e.g. weather data. This allows the system to be accessed from anywhere.

2.1.2 Distributed Model Predictive Control

The centralized, decentralized, and distributed MPC approaches are presented and discussed in [12]. The traditional, centralized approach includes all information and model parameters in one place. This can be problematic if the formulation includes an increasing number of subsystems, or agents, as the optimization problem can become computationally infeasible. The original problem may in some cases be split into several sub-problems pertaining to each agent, and solved independently by local controllers. This is known as decentralized MPC. The decentralized approach may be less optimal than the centralized approach since the independent agents do not cooperate or take into account their effect on the larger system. Distributed MPC is a middle ground between the centralized and decentralized approach, where there is some limited communication between agents. The information exchange may give better global performance, at best the same as the centralized approach.

A recent paper investigates these MPC strategies in the context of energy management in buildings and energy hubs [13], with a particular focus on the distributed and centralized approach. The privacy of the building occupants as well as computational feasibility and robustness are their main motivations for a distributed approach over a centralized one. In the distributed approach, the local controllers are derived from performing dual decomposition on the original problem, where a shared variable couples the systems. The shared variables are the amount of heat taken from a shared water storage tank. In terms of temperature violation and heating efficiency, the distributed and centralized approach have nearly identical performance. The decentralized approach performs significantly worse in those regards. The distributed approach is only more computationally efficient than the centralized approach when the number of agents is large, around 50 or more.

2.1.3 Peak Power Reduction

One paper incorporates incentives to make homes consume power in a grid-friendly way [14]. This is achieved by designing and introducing a special tariff that reflects the cost of providing the end-user with electricity at any given time. Additionally, a MPC scheme is introduced so the climate control takes the new tariffs into account. The resulting scheme reduces peak power demand primarily by shifting electricity demands associated with the thermal load. One disadvantage of their approach is that the total electricity consumed increases due to the load shifting. They also note that some consumption cannot realistically be shifted, like lighting. This limits the efficacy of the approach.

Another paper proposes optimization models to manage energy load for single- and multi-house cases, while taking into account distributed energy sources and batteries [15]. The single-house optimization model introduces a fixed limit on the amount of energy purchased at each time slot. The multi-house optimization model introduces a similar constraint, but the limit applies to the sum of purchased energy by the houses at each time slot. Test show that the multi-user optimization model produces much lower peak electricity demand compared to using the single-user model on the same number of homes.

2.2 Theory

2.2.1 Model Predictive Control

Model predictive control (MPC) is a control method that employs a model of the target system to decide the best course of action. MPC algorithms solve a constrained finite-horizon optimization problem based on the system model. At each time step, the solution to the optimization problem will be a sequence of inputs, as well as the system states that would result from that sequence of inputs. The first input/action in the solution sequence is applied to the real system, and the process repeats itself at the next time step. This allows the MPC to provide optimal control for the current time step while taking into account the future behaviour of the system. Figure 2.1 shows the working principle of the algorithm. As time progresses, the optimization window moves forward so that it starts at the current time step. The optimization problem the MPC solves at each new optimization window can be written as follows:

$$\min_{\boldsymbol{x},\boldsymbol{u}} \quad J(\boldsymbol{x},\boldsymbol{u}) = \sum_{i=0}^{N-1} l(\boldsymbol{x}_i,\boldsymbol{u}_i) + V(\boldsymbol{x}_N) \tag{2.1a}$$

$$\text{subject to} \quad \boldsymbol{x}_{i+1} = \boldsymbol{f}\left(\boldsymbol{x}_i,\boldsymbol{u}_i\right) \qquad i = 0,1,...N-1 \tag{2.1b}$$

$$\boldsymbol{x}_0 = \bar{\boldsymbol{x}}_0 \tag{2.1c}$$

$$g(\boldsymbol{x}_i,\boldsymbol{u}_i) \leq 0 \qquad i = 0,1,...N-1 \tag{2.1d}$$

$$h(\boldsymbol{x}_i,\boldsymbol{u}_i) = 0 \qquad i = 0,1,...N-1 \tag{2.1e}$$

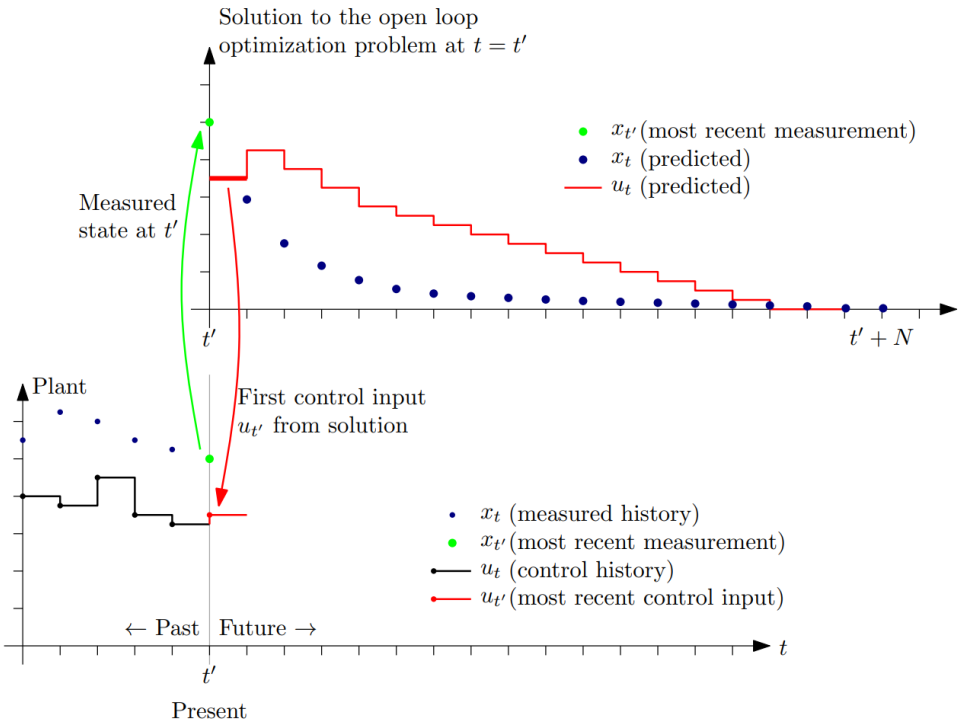

Figure 2.1: An illustration showing the working principle of model predictive control [16]

Equation (2.1) is a generic MPC formulation. The subscript i represents the prediction steps for the current optimization horizon, with a prediction horizon of N steps.

$\boldsymbol{x}$ is the state vector for the system. It describes the state of the system the MPC is trying to control. $\boldsymbol{u}$ is the input vector, which contains the inputs the MPC has to control the system. Equation (2.1a) shows the cost function to be minimized. The function $l(\boldsymbol{x}_i, \boldsymbol{u}_i)$ is known as the stage cost and $V(\boldsymbol{x}_N)$ is known as the terminal cost, which applies to the last state in the optimization horizon. Together these define the cost function $J(\boldsymbol{x}, \boldsymbol{u})$. Equation (2.1b) are the dynamics constraints of the system. The dynamics constraints ensure continuity between steps $i+1$ and i in terms of the system model. Equation (2.1c) is a constraint on the initial state of the system, which is set to be equal to an estimate of the initial state $\bar{\boldsymbol{x}}_0$. Any additional inequality or equality constraints in the system are included in (2.1d) and (2.1e).

(2.1) is a so-called optimal control problem (OCP), 'optimal' because the solution provides an (usually) optimal way to control the system with respect to the provided cost function and constraints. This generalized form may represent non-linear functions. Model predictive control with non-linear constraints or cost functions is known as non-linear model predictive control (NMPC).

2.2.2 Numerical Optimal Control

CasADi [17] is a symbolic framework for optimal control, among other things. For nonlinear programming (NLP), it solves problems of the following form.

$$\begin{aligned} &\min_{z} \quad f(z,p) \\ &z_{lb} \leq z \leq z_{ub} \\ &g_{lb} \leq g(z,p) \leq g_{ub} \end{aligned} \tag{2.2}$$

Where z is the decision variable and p is a known parameter vector. z_{lb} and z_{ub} are bounds on the decision variable. $g(z,p)$ contains the constraint functions for the problem, and may encompass both equality and inequality constraints. g_{lb} and g_{ub} are the bounds on the constraint functions, and are equal when $g(z,p)$ is an equality constraint. To utilize CasADi to solve the MPC problem, (2.1) must be reformulated to fit the NLP solver in CasADi (see (2.2)). First, the constraints in (2.1) may be reformulated as follows:

$$\min_{\boldsymbol{x},\boldsymbol{u}} \quad J(\boldsymbol{x},\boldsymbol{u}) \tag{2.3a}$$

$$\text{subject to} \quad 0 \leq \boldsymbol{x}_{i+1} - \boldsymbol{f}(\boldsymbol{x}_i, \boldsymbol{u}_i) \leq 0 \quad i = 0,1,...N-1 \tag{2.3b}$$

$$\bar{\boldsymbol{x}}_0 \leq \boldsymbol{x}_0 \leq \bar{\boldsymbol{x}}_0 \tag{2.3c}$$

$$-\infty \leq g(\boldsymbol{x}_i, \boldsymbol{u}_i) \leq 0 \quad i = 0,1,...N-1 \tag{2.3d}$$

$$0 \leq h(\boldsymbol{x}_i, \boldsymbol{u}_i) \leq 0 \quad i = 0,1,...N-1 \tag{2.3e}$$

Equation (2.3b) is the rewritten form of the dynamics constraints (2.1b). The upper and lower bounds are the same because it is an equality constraint. The constraint (2.3c) is based on the initial state equality constraint (2.1c). The initial state estimate $\bar{\boldsymbol{x}}_0$ is set as the upper and lower bound on the initial state $\boldsymbol{x}_0$. The equality constraint (2.1e) is rewritten as (2.3e), where once again the upper and lower bounds are the same. (2.1d) is an inequality constraint, and as such it is written as (2.3d). Since it only has an upper bound of 0, its 'lower bound' is set to negative infinity to fit with the NLP solver.

In addition to rewriting the constraints as (2.3b)-(2.3e), the states $\mathbf{x}$ and inputs $\mathbf{u}$ from may be both be included as decision variables. Then the MPC formulation (2.1) may be written in the same form as the CasADi NLP, (2.2):

$$\boldsymbol{y} = \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{u} \end{bmatrix} \tag{2.4a}$$

$$\min_{\boldsymbol{y}} \; J(\boldsymbol{y}, \boldsymbol{p}) \tag{2.4b}$$

$$\boldsymbol{y}_{lb} \leq \boldsymbol{y} \leq \boldsymbol{y}_{ub} \tag{2.4c}$$

$$\boldsymbol{g}_{lb} \leq \boldsymbol{g}(\boldsymbol{y}, \boldsymbol{p}) \leq \boldsymbol{g}_{ub} \tag{2.4d}$$

The cost function (2.4a) equates to the MPC cost function (2.1a). In the context of (2.1), contents of $\boldsymbol{p}$ may be weights for the cost function, model parameters, external data, or other information. The dynamics constraints (2.3b), inequality constraints (2.3d), and equality constraints (2.3e) are all part of the constraint functions (2.4d). The initial state constraint is a part of the decision variable bounds (2.4c).

2.2.3 Dual Decomposition

Decomposition methods break up large problems into smaller ones to be solved separately. This has the advantage of reducing problem complexity as well as making it possible to compute the problems in parallel. In the context of model predictive control, the problems may be solved by separate controllers at different locations.

A problem may be separable, in which case the subproblems are independent of each other and the solution to the original problem can be re-assembled from the solutions to the subproblems [18]. The following is an example of a separable problem:

$$\begin{aligned} &\min_{x_1, x_2} \; f_1(x_1) + f_2(x_2) \\ &x_1 \in C_1 \\ &x_2 \in C_2 \end{aligned} \tag{2.5}$$

Equation (2.5) is a two-variable example of a separable problem. Both the objective function and constraints are decoupled, and the problem may be solved

separately. The function f_1 and the constraints C_1 constitute the subproblem pertaining to x_1. f_2 and the constraints C_2 constitute the subproblem pertaining to x_2:

$$\begin{aligned} &\min_{x_1} \ f_1(x_1) \\ &x_1 \in C_1 \end{aligned} \tag{2.6}$$

$$\begin{aligned} &\min_{x_2} \ f_1(x_2) \\ &x_2 \in C_2 \end{aligned} \tag{2.7}$$

The solutions to problem (2.6) and (2.7) can be combined for the solution to the original problem. Decomposition methods come into play when problems are not separable:

$$\begin{aligned} &\min_{x_1,x_2} \ f_1(x_1) + f_2(x_2) \\ &x_1 \in C_1 \\ &x_2 \in C_2 \\ &h_1(x_1) + h_2(x_2) \leq h_0 \end{aligned} \tag{2.8}$$

The problem in (2.8) has a coupling constraint $h_1(x_1) + h_2(x_2) \leq h_0$, and cannot be trivially separated into subproblems without decomposition. To perform dual decomposition on this example, the first step is forming the Lagrangian function:

$$L(x_1, x_2, \mu) = f_1(x_1) + f_2(x_2) + \mu\,(h_1(x_1) + h_2(x_2) - h_0) \tag{2.9}$$

As shown in (2.9) the Lagrangian relaxation entails penalizing violation of the inequality constraint in the objective function, while removing the inequality as a hard constraint. The Lagrange multiplier μ dictates how expensive constraint violations should be.

Using (2.9), the dual function $d(\mu)$ for (2.8) is

$$\begin{aligned} d(\mu) = \min_{x_1,x_2} \ & L(x_1, x_2, \mu) \\ & x_1 \in C_1 \\ & x_2 \in C_2 \end{aligned} \tag{2.10}$$

The primal problem (2.8) is a minimization problem, and as such its dual problem is a maximization problem:

$$\begin{aligned} \max_{\mu} \quad & d(\mu) \\ s.t. \quad & \mu \geq 0 \end{aligned} \tag{2.11}$$

(2.11) is known as the dual problem. The weak duality theorem states that an optimal solution to the dual problem provides a lower bound on the optimal value

of the original primal problem [19], in this case, (2.8). If the original problem is non-convex, the solution provides a local (or possibly global) maximizer of the dual problem.

The difference between the optimal value of the original and dual problem is known as the *duality gap*. If the primal problem is convex and satisfies a constraint qualification, the duality gap is zero. A duality gap of zero is known as *strong duality*, where the optimal value of the dual problem is the same as the optimal value of the primal problem. The next step is to rewrite the Lagrangian.

$$\begin{aligned} &L(x_1, x_2, \mu) \\ &= f_1(x_1) + f_2(x_2) + \mu \left(h_1(x_1) + h_2(x_2) - h_0 \right) \\ &= \left(f_1(x_1) + \mu\, h_1(x_1) \right) + \left(f_2(x_2) + \mu\, h_2(x_2) \right) - \mu\, h_0 \end{aligned} \tag{2.12}$$

Equation (2.12) is separable in x_1 and x_2, and by fixing μ a subproblem for each variable is obtained:

$$\begin{aligned} d_1(\mu) = \min_{x_1} \; & f_1(x_1) + \mu\, h_1(x_1) \\ & x_1 \in C_1 \end{aligned} \tag{2.13}$$

$$\begin{aligned} d_2(\mu) = \min_{x_2} \; & f_2(x_2) + \mu\, h_2(x_2) \\ & x_2 \in C_2 \end{aligned} \tag{2.14}$$

$$d(\mu) \;=\; g_1(\mu) + g_2(\mu) - \mu h_0 \tag{2.15}$$

Where (2.13) and (2.14) are the dual problems for x_1 and x_2, referred to as $d_1(\mu)$ and $d_2(\mu)$. The function (2.15) is the dual function that is maximized in the dual problem (2.11).

The dual problem (2.11) can be solved with a projected subgradient method [20] if the dual function $d(\mu)$ is not differentiable. If the dual function is differentiable, a projected gradient method suffices. Starting from some initial point μ_0, at each iteration the method updates the dual variable μ in the direction of the (sub)gradient at the current point. This update is then projected to stay feasible:

$$\mu_{k+1} = P\Big(\mu_k + \alpha_k \, \nabla d(\mu_k) \Big) \tag{2.16}$$

(2.16) shows the gradient method. $\nabla d(\mu_k)$ is the gradient of the dual function at μ_k. Since the dual problem (2.11) is a maximization problem, the direction of ascent is used ($+\nabla d(\mu_k)$). P is a projection function for the dual variable μ to ensure iterates are feasible. α_k is the step size at iteration k. If we consider the gradient method (2.16) on the dual problem (2.11) (with the dual function defined as (2.15), the dual variable μ will be updated as follows:

$$\begin{aligned}\mu_{k+1} &= \Big(\mu_k + \alpha_k \nabla d(\mu_k)\Big)_+ \\ &= \Big(\mu_k + \alpha_k \nabla\big(d_1(\mu_k) + d_2(\mu_k) - h_0\,\mu_k\big)\Big)_+ \\ &= \Big(\mu_k + \alpha_k\big(h_1({x_1}^*) + h_2({x_2}^*) - h_0\big)\Big)_+\end{aligned} \tag{2.17}$$

(2.17) shows how the dual variables are updated at each iteration. $h_1(x_1^*)$ and $h_2(x_2^*)$ are the values of the constraint functions at the solution of the minimization problems (2.13) and (2.14). The projection function $()_+$ ensures the dual variable iterates stay positive, in line with the positivity constraint $\mu \geq 0$ in (2.11).

Chapter 3

Method

In the POWIOT project, heat pumps are used to control the temperature in three rooms in a house in Trondheim. The power use of the heat pumps is directed by a controller in a way that saves money in electricity costs and considers preferences for room temperature.

A simplified house control formulation, based on the one currently in use in Trondheim, is developed in 3.1. This simplified formulation is expanded upon in 3.3 to include peak power reduction. Based on those peak power reduction formulations, three formulations are presented in 3.5. In chapter 4 those three formulations are tested in a variety of simulation scenarios to determine their performance.

Data collection software installed in the house in Trondheim collects data used by the control algorithm that controls its heat pumps. Much of this same data is used in the simulations in this thesis. An overview of the external data used in the simulations can be found in section 3.2.

3.1 House Control Formulation

An MPC formulation that controls the room temperature in a house needs a temperature model. Figure 3.1 shows how the temperature model for a single room in a building. It is a simplified version of the room temperature model used by the MPC that controls the house in Trondheim. It is the basis of the dynamics constraints (3.3a) and (3.3b).

$T_{outdoor}$ is the outdoor temperature, T_{wall} is the temperature in the wall, and T_{room} is the temperature in the room. P_{HP} is the power output from the heat pump in the room, and COP is the coefficient of power for that heat pump. The COP tells how much heating power is produced per power provided to the heat pump. If the COP is 3.5, then if the heat pump is provided with 1 kW of power, it will produce 3.5 kW of heating. The coefficient of power depends on operating conditions, especially outdoor temperature. ρ_{out} and ρ_{in} are parameters that represent how well heat travels between the wall and the outside, and between the wall and the room.

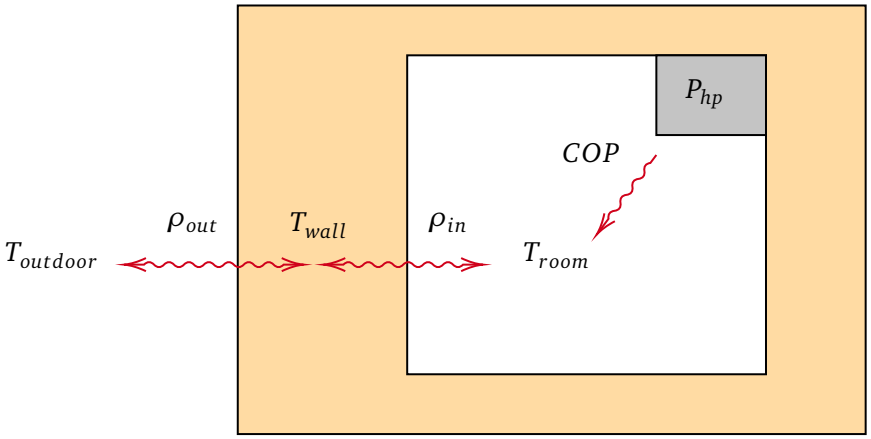

Figure 3.1: Temperature model of room used in control algorithm. Red lines represent energy transfer.

The MPC formulation in use on the house in Trondheim, including its temperature model, is developed by Prof. Sebastien Gros. It considers multiple heat pumps installed in multiple rooms. In this thesis, a simplified version of that MPC formulation is developed, where each house is assumed to have one room with a single heat pump.

The input $\boldsymbol{u}$ of the system is the heat pump power, and the state $\boldsymbol{x}$ of the system is the wall- and room temperature. In the real system the heat pump measure room temperature with a sampling time of 5 minutes. As such the length between each prediction step i is 5 minutes. The total prediction length, represented by N, is 24 hours. The states and inputs are defined as follows:

$$\boldsymbol{x}_i^T = \left[T_{room,i}, T_{wall,i}, T_{slack,i}\right] \qquad i = 0, ..., N \tag{3.1a}$$

$$\boldsymbol{u}_i = \left[P_{HP,i}\right] \qquad i = 0, ..., N-1 \tag{3.1b}$$

The state vector for each step i, $\boldsymbol{x}_i$, contains a slack temperature $T_{slack,i}$ that helps keep the room temperature above some limit T_{min}. The following is the simplified MPC formulation:

$$\begin{aligned} &\min_{\boldsymbol{x},\boldsymbol{u}} \ \Phi(\boldsymbol{x},\boldsymbol{u}) \\ &= \min_{\boldsymbol{x},\boldsymbol{u}} \ \sum_{i=1}^{N-1} l_i(\boldsymbol{x}_i,\boldsymbol{u}_i) + V(\boldsymbol{x}_N) \end{aligned} \tag{3.2a}$$

$$\begin{aligned} l_i(\boldsymbol{x}_i,\boldsymbol{u}_i) = &\, w_{power} \cdot C_{cost,i} \cdot P(\boldsymbol{u}_i) \\ &+ w_{comfort} \cdot C_{comfort,i+1}^2 + w_{slack} \cdot T_{slack,i+1}^2 \end{aligned} \tag{3.2b}$$

$$V(\boldsymbol{x}_N) = w_{comfort} \cdot C_{comfort,N}^2 + w_{slack} \cdot T_{slack,N}^2 \tag{3.2c}$$

subject to the constraints

$$T_{room,i+1} = T_{room,i} + \left(\rho_{in} \cdot (T_{wall,i} - T_{room,i}) + COP \cdot P_{HP,i}\right) \tag{3.3a}$$

$$T_{wall,i+1} = T_{wall,i} + \left(\rho_{in} \cdot (T_{room,i} - T_{wall,i}) + \rho_{out} \cdot (T_{outdoor,i} - T_{wall,i})\right) \tag{3.3b}$$

$$\boldsymbol{x}_0 = \bar{\boldsymbol{x}} \tag{3.3c}$$

$$T_{min,i} \leq T_{room,i} + T_{slack,i} \tag{3.3d}$$

$$0 \leq T_{slack,i} \tag{3.3e}$$

$$0 \leq P_{HP,i} \leq P_{max} \tag{3.3f}$$

where

$$\begin{aligned} C_{cost,i} &= \text{Spot market price}_i \cdot \text{VAT} \\ C_{comfort,i} &= T_{room,i} - T_{reference,i} \\ P(u_i) &= P_{HP,i} + P_{ext,i} \end{aligned} \tag{3.4}$$

Table 3.1: House MPC model parameters

Parameter	Value
COP	2.50
ρ_{in}	0.37
ρ_{out}	0.018

Equation (3.2) shows the cost function for the house control problem. (3.2b) is the stage cost for each prediction step i, and (3.2c) is the terminal cost associated with the last state $\boldsymbol{x}_N$. C_{cost} is the cost associated with power consumption. It only considers spot price, omitting any fixed grid costs. $C_{comfort}$ is a term that represents the cost of deviations from the reference temperature $T_{reference}$. T_{min} is the minimum room temperature. The minimum room temperature is not enforced through a hard constraint, and is instead strongly penalized using the slack variable T_{slack}. Equations (3.3d) and (3.3e) are the constraints on T_{slack}.

$P(u_i)$ is the power consumption as a function of the system input $\boldsymbol{u}$. It is the sum of the heat pump power P_{HP} and 'external' power use P_{ext}, like lighting or appliances. The external power is not controllable by the MPC. P_{HP} is the only element of the input $\boldsymbol{u}$. Direct control over the heat pump power P_{HP}, stated in (3.1b), is another simplification. In reality, the heat pumps take a temperature reference as input, and the resulting power output must be estimated with some power model. Equation (3.3f) shows the constraints on the heat pump power, where P_{max} is the maximum power output of a heat pump, in our case 1.5 kW. Since thesis only considers houses with a single heat pump and a single room, P_{ext} (obtained from a large house) is scaled down to emulate a smaller house in the tests in chapter 4.

The constraint (3.3c) states that the start state $\boldsymbol{x}_0$ must be equal to some current state estimate $\bar{\boldsymbol{x}}$. Typically an estimator is employed to produce a state estimate. For this thesis the prediction models are assumed to be exact, and the estimate $\bar{\boldsymbol{x}}$ is simply set to the second optimal state $\boldsymbol{x}_1^*$ from the optimization results of the previous time step $t-1$.

The parameter values for the room temperature model in table 3.1 were obtained using system identification methods in a previous POWIOT project. They are provided to the dynamics constraints (3.3a) and (3.3b). The heat pump coefficient of power (COP) varies depending on the outdoor temperature. In this case, it is simplified and given a relatively low, fixed value, representative of the low outside temperature considered in the test case. ρ_{in} and ρ_{out} are parameters that model the effects of thermal conduction through the building envelope. In other words, they determine how fast temperature propagates between the room, the wall, and the outside. The discrete temperature model (see (3.3a) and (3.3b) is based on the 5-minute sampling time of the heat pumps in the POWIOT project. As such, each prediction horizon step i represents a 5-minute window. The horizon length, represented by N, is 24 hours. The same temperature model is used for all the houses.

By gathering the equality and inequality constraints we formulate the following optimization problem:

$$\min_{\boldsymbol{x},\boldsymbol{u}} \ \Phi(\boldsymbol{x},\boldsymbol{u}) \tag{3.5a}$$
$$\text{s.t. } \boldsymbol{x}_{i+1} = f(\boldsymbol{x}_i,\boldsymbol{u}_i) \tag{3.5b}$$
$$\boldsymbol{x}_0 = \bar{\boldsymbol{x}} \tag{3.5c}$$
$$T_{min,i} - T_{slack,i} \leq T_{room,i} \tag{3.5d}$$
$$0 \leq T_{slack,i} \tag{3.5e}$$
$$0 \leq P_{HP,i} \leq P_{max} \tag{3.5f}$$

where

$$\boldsymbol{x}_i^T = \left[T_{room,i}, T_{wall,i}, T_{slack,i}\right] \qquad i = 0,...,N \tag{3.6a}$$
$$\boldsymbol{u}_i = \left[P_{HP,i}\right] \qquad i = 0,...,N-1 \tag{3.6b}$$

The constraint (3.5b) contains the dynamics constraints (3.3a) and (3.3b).

3.1.1 Centralized MPC

The formulation (3.5) formulation considers a single home. To consider multiple houses, as in the proposed residential power coalitions, we introduce a subscript h for each house in the coalition, with H total houses. By including all houses in the same optimization problem, we formulate a centralized optimization problem. In

this centralized formulation, the cost functions and constraints from each house are considered together:

$$\min_{\boldsymbol{x},\boldsymbol{u}} \sum_{h=1}^{H} \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \tag{3.7a}$$

$$\text{s.t. } \boldsymbol{x}_{i+1,h} = f_h(\boldsymbol{x}_{i,h}, \boldsymbol{u}_{i,h}) \tag{3.7b}$$

$$\boldsymbol{x}_{0,h} = \bar{\boldsymbol{x}}_h \tag{3.7c}$$

$$T_{min,i,h} - T_{slack,i,h} \leq T_{room,i,h} \tag{3.7d}$$

$$0 \leq T_{slack,i,h} \tag{3.7e}$$

$$0 \leq P_{HP,i,h} \leq P_{max} \tag{3.7f}$$

The state and input vectors for each house are defined as:

$$\boldsymbol{x}_{i,h}^T = \left[T_{room,i,h}, T_{wall,i,h}, T_{slack,i,h}\right] \qquad i = 0,...,N \quad h = 1,...,H \tag{3.8a}$$

$$\boldsymbol{u}_{i,h}^T = \left[P_{HP,i,h}\right] \qquad i = 0,...,N-1 \quad h = 1,...,H \tag{3.8b}$$

such that the total state vector x and input vector u consist of the state $\boldsymbol{x}_h$ and input $\boldsymbol{u}_h$ vectors from each house:

$$\boldsymbol{x}^T = \left[\boldsymbol{x}_1^T, ..., \boldsymbol{x}_H^T\right] \tag{3.9a}$$

$$\boldsymbol{u}^T = \left[\boldsymbol{u}_1^T, ..., \boldsymbol{u}_H^T\right] \tag{3.9b}$$

The problem (3.7) is the centralized optimization problem for the power coalition as a whole. This centralized formulation does not have any coupling between the houses, and can therefore be trivially decomposed into a problem pertaining to each house (see the decentralized formulation in 3.1.2). However, this centralized formulation is used as the basis of the formulations in 3.3, which explores formulations to reduce the peak total power among all the houses. In those formulations, there are coupling constraints between the houses, and they can therefore not be trivially decomposed to a problem for each house.

Figure 3.2 shows how the centralized formulation could be implemented in practice. A centralized controller receives state measurements $\boldsymbol{x}_h$ and external parameters $\boldsymbol{p}_h$ from each house, and in turn gives commands to the heat pumps, $\boldsymbol{u}_h$.

3.1.2 Decentralized MPC

Since there are no coupling constraints or states between the houses, the problem (3.7) may be decomposed and solved using separate controllers on a per-house basis:

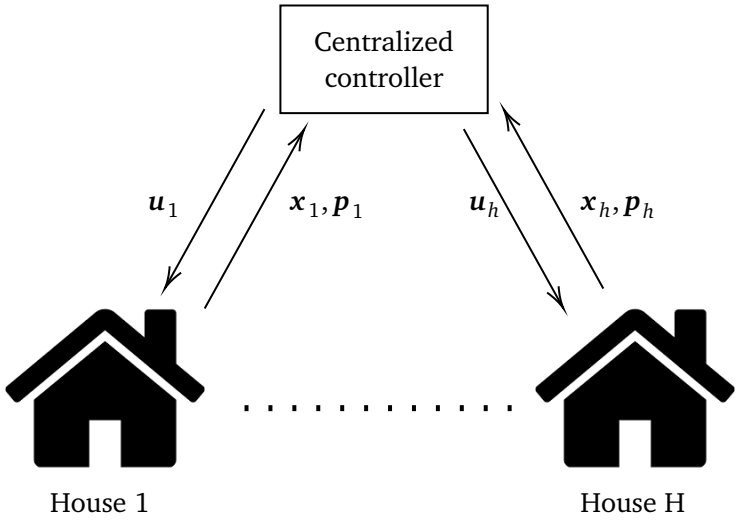

Figure 3.2: Principle of centralized MPC for house control

$$\min_{\boldsymbol{x}_h, \boldsymbol{u}_h} \ \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \tag{3.10a}$$

$$\text{s.t. } \boldsymbol{x}_{i+1,h} = f_h(\boldsymbol{x}_{i,h}, \boldsymbol{u}_{i,h}) \tag{3.10b}$$

$$\boldsymbol{x}_{0,h} = \bar{\boldsymbol{x}}_h \tag{3.10c}$$

$$T_{min,i,h} - T_{slack,i,h} \leq T_{room,i,h} \tag{3.10d}$$

$$0 \leq T_{slack,i,h} \tag{3.10e}$$

$$0 \leq P_{HP,i,h} \leq P_{max} \tag{3.10f}$$

Note that this is just the single-home formulation (3.5) with the added subscript h to indicate which house it is. This formulation is the decentralized version of (3.7), and solving the sub-problems for each house h will yield the same solution as its centralized equivalent. Figure 3.3 shows how this formulation could be implemented in practice. A controller in each house receives state measurements and external parameters from their house, and in turn gives commands to the heat pumps. Compare this to a centralized approach, shown in figure 3.2, where a single controller collects data from and sends commands to all houses in the system.

The decentralized approach is based on the formulation (3.10)**. The decentralized approach is used as a benchmark in chapter 4 which the various centralized and distributed formulations for peak power reduction from section 3.5 are compared against. It emulates a scenario where each home has a local controller for heating control, and the controllers do not coordinate with each other.**

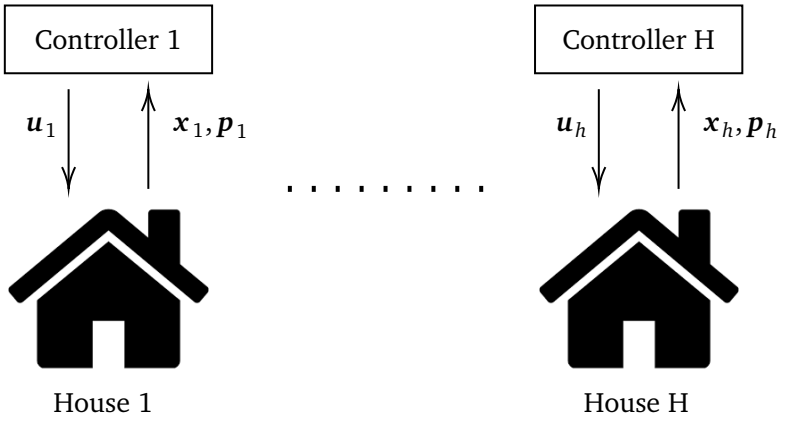

Figure 3.3: Principle of decentralized MPC for house control

3.2 External Factors

The data collection software that is installed in the house in Trondheim collects information that is relevant to the estimation and control of power and indoor temperature. The collected data includes power consumption, weather forecasts for the area, and spot prices for the Trondheim day ahead market. Data has been collected for a total time span of roughly a year and is the source of spot price, outdoor temperature, and external power use data used in the simulations in this thesis.

Figure 3.4 shows the spot prices in Trondheim on November 29 and November 30, 2021. The spot prices reached their highest points during these days in 2021, with significant fluctuations from hour to hour. They represent a sort of worst-case scenario for power prices. The prices are retrieved from Nord Pool [21], a European power exchange. Figure 3.5 shows the outdoor temperature in Trondheim from November 29 through November 30, 2021. The data has an hourly resolution and is retrieved from the Norwegian Meteorological Institute [22].

Figure 3.6 shows the power consumption of the house in Trondheim, excluding the heat pumps, from November 29 through November 30, 2021. Power use in the mornings and evenings is especially high. The power data is collected using a Tibber Pulse device [23], which connects to smart meters in homes and tracks power consumption in real-time. Unlike spot prices and weather, there are no existing forecasts for external power consumption. It can be assumed, however, that power consumption profiles from previous days provide reasonable predictions for the following days. Figure 3.7 shows a smoothed average of non-heat pump power consumption from 18 different days.

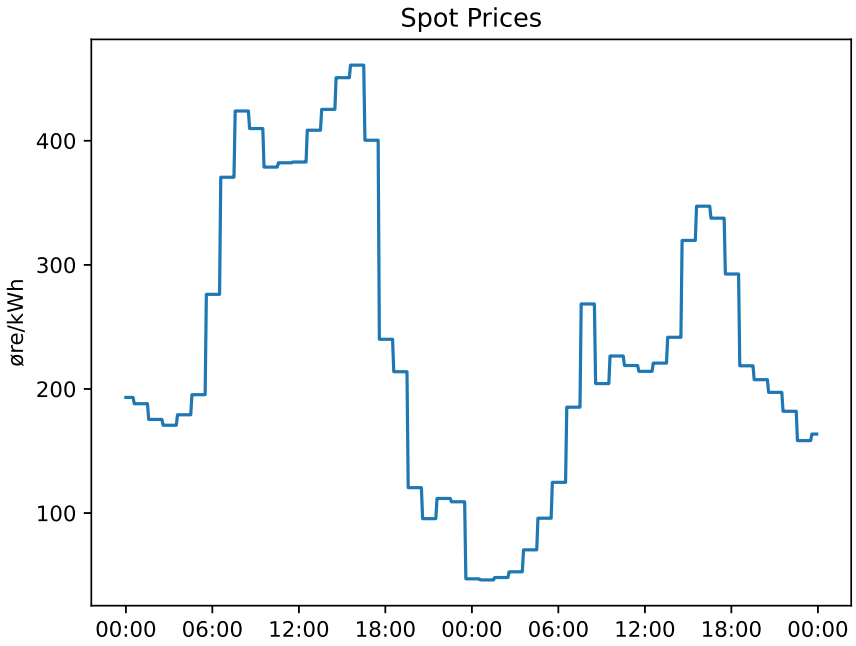

Figure 3.4: Spot price in Trondheim November 29 and November 30, 2021

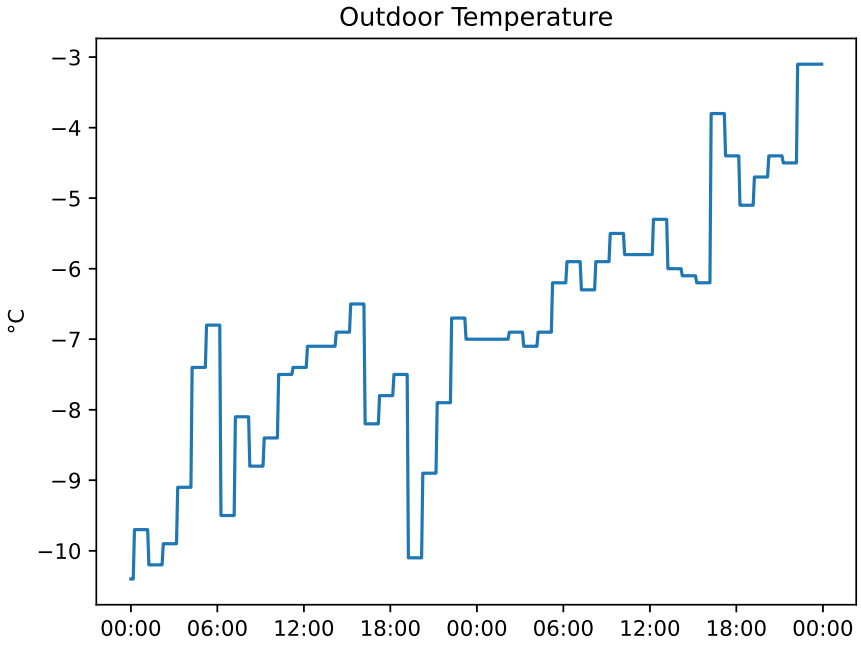

Figure 3.5: Outdoor temperature in Trondheim November 29 and November 30, 2021

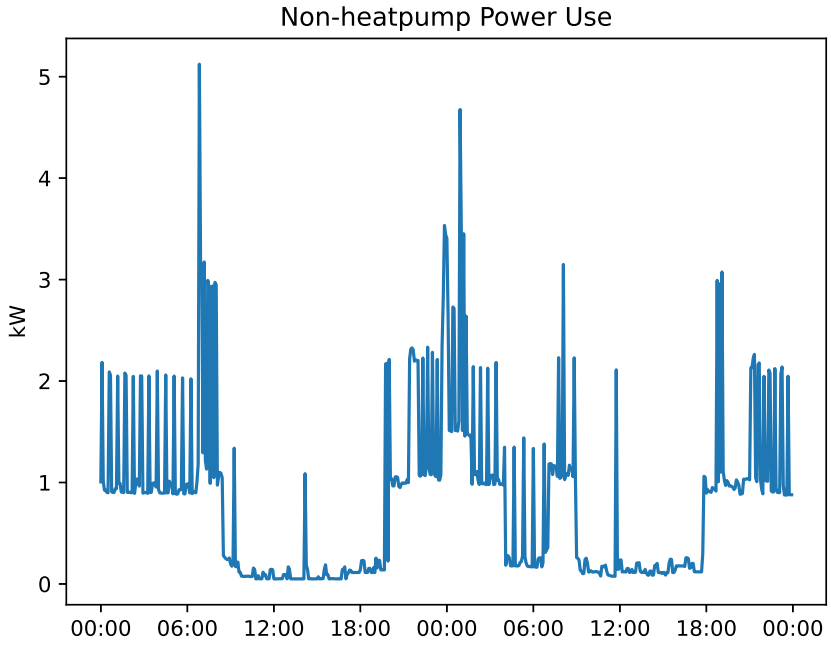

Figure 3.6: Non-heat pump power use in house November 29 and November 30, 2021

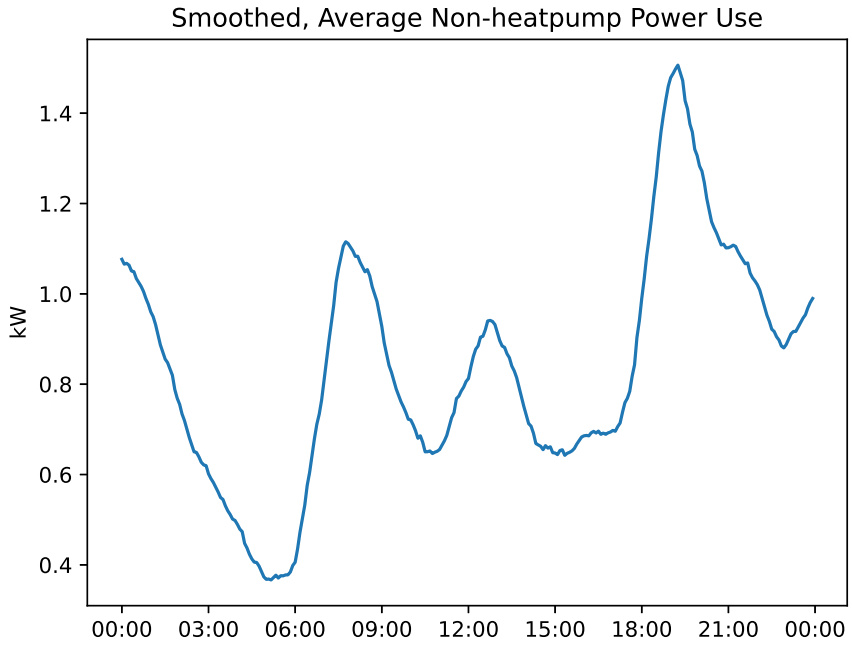

Figure 3.7: Smoothed average of 18 days of non-heat pump power use

3.3 Peak Power Reduction Formulations

The house control problem (3.5) is designed to take the electricity spot price into account when deciding the heat pump power usage. It is not designed with peak power reduction in mind. Therefore, this thesis extends the current control algorithm to systematically handle peak power consumption. This section proposes several formulations to implement this feature. It investigates whether the formulations are fit for dual decomposition, and consequently distributed control. To keep the formulations short, we rewrite the constraints (3.7b)-(3.7f) for the following centralized problem:

$$\min_{\boldsymbol{x},\boldsymbol{u}} \sum_{h=1}^{H} \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \tag{3.11a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.11b}$$

$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.11c}$$

The equality constraints (3.11b) contain the dynamics constraints (3.7b) and initial state constraints (3.7c). The inequality constraints (3.11c) contain the slack state constraints (3.7d) and (3.7e), and heat pump power constraints (3.7f).

3.3.1 Fixed Power Limit

The most simple approach to reducing peak power is to simply introduce a constraint on the sum of power consumption for each time step:

$$\sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}) \leq P_{lim} \tag{3.12}$$

Note that if the sum of uncontrollable powers is larger than the total power limit P_{lim}, the constraint (3.12) cannot be satisfied and the problem becomes infeasible. We may remedy this by including a slack variable for the power consumption:

$$\sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}) \leq P_{lim} + P_{slack,i} \tag{3.13a}$$

$$P_{slack,i} \geq 0 \tag{3.13b}$$

By adding this constraint to (3.11) and adding a cost associated with the new slack variable P_{slack}, we obtain a centralized formulation with a fixed limit on the total power consumed by the houses:

$$\min_{\boldsymbol{x},\boldsymbol{u}} \sum_{h=1}^{H} \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) + \sum_{i=0}^{N-1} w_{P_{slack}} P_{slack,i} \tag{3.14a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.14b}$$
$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.14c}$$
$$\sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}) \leq P_{lim} + P_{slack,i} \tag{3.14d}$$
$$P_{slack,i} \geq 0 \tag{3.14e}$$

The slack variable P_{slack} ensures that the constraint (3.13) is met, even if the sum of uncontrollable powers is larger than the limit P_{lim}. The weight on the power slack $w_{P_{slack}}$ is set very high so that the fixed power constraint is always enforced if it is feasible to do so. We may transform the slack variable P_{slack} into a 'localized' version for each house:

$$P_{slack,i} = \sum_{h=1}^{H} P_{slack,i,h} \tag{3.15}$$

Then the power slack constraints (3.13) become

$$-P_{lim} + \sum_{h=1}^{H} \Big(P_h(\boldsymbol{u}_{i,h}) - P_{slack,i,h} \Big) \leq 0 \tag{3.16a}$$
$$P_{slack,i,h} \geq 0 \tag{3.16b}$$

We can 'dualize' (move into the cost function) constraint (3.16a). This is the first step of the dual decomposition described in section 2.2.3. The problem (3.14) may then be rewritten as

$$\begin{aligned} d(\boldsymbol{\mu}) = \min_{\boldsymbol{x},\boldsymbol{u}} \; & \sum_{i=0}^{N-1} \mu_i \left(-P_{lim} + \sum_{h=1}^{H} \Big(P_h(\boldsymbol{u}_{i,h}) - P_{slack,h} \Big) \right) \\ & + \sum_{h=1}^{H} \left(\Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) + \sum_{i=0}^{N-1} w_{P_{slack}} P_{slack,i,h} \right) \end{aligned} \tag{3.17a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.17b}$$
$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.17c}$$
$$P_{slack,i,h} \geq 0 \tag{3.17d}$$

Where $\boldsymbol{\mu}$ is the dual variable vector associated with constraint (3.16a). We seek the solution to the dual problem

$$\max_{\boldsymbol{\mu}} \; d(\boldsymbol{\mu}) \tag{3.18a}$$
$$\text{s.t. } \boldsymbol{\mu} \geq 0 \tag{3.18b}$$

We may decompose (3.17) to be solved on a per-house basis. This is the next step of the dual decomposition:

$$\begin{aligned} d_h(\boldsymbol{\mu}) = \min_{\boldsymbol{x}_h, \boldsymbol{u}_h} \; & \sum_{i=0}^{N-1} \mu_i \Big(P_h(\boldsymbol{u}_{i,h}) - P_{slack,i,h} \Big) \\ & + \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) + \sum_{i=0}^{N-1} w_{P_{slack}} P_{slack,i,h} \end{aligned} \tag{3.19a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.19b}$$

$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.19c}$$

$$P_{slack,i,h} \geq 0 \tag{3.19d}$$

The cost of each slack variable $P_{slack,i,h}$ in the minimization problem in (3.19a) is

$$w_{P_{slack}} P_{slack,i,h} - \mu_i P_{slack,i,h} \tag{3.20}$$

Since there is no upper bound on the slack variables $P_{slack,i,h}$, the problem (3.19) may be unbounded if the dual variables are larger than the slack weight $w_{P_{slack}}$.

$$d_h(\boldsymbol{\mu}) = -\infty \quad \text{for} \quad \mu_i > w_{P_{slack}} \tag{3.21}$$

With $P_{slack} = \infty$. One the other hand, if the weight is larger than the dual variable, the minimization problem will produce an optimal slack variable of zero, $P_{slack,i,h} = 0$.

If the 'correct' value of $P_{slack,i,h}$ (solution of original problem) is not 0, the dual problem will consequently produce an incorrect result. As such (3.14) is unfit for dual decomposition in its current form.

3.3.2 Momentary Peak State

The last approach formulated a constraint on the sum of power consumption between the homes. A less direct approach to try to reduce peak power is to introduce new states that in some way represent the peak power consumption and penalize them. For this formulation, we consider a single peak state s, that represents the highest total power consumption between the homes at any given prediction step. To enforce this representation, we introduce the constraint:

$$\sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}) \leq s \tag{3.22}$$

The constraint (3.22) ensures that the peak state s is larger than the highest sum of powers for all prediction steps i. In other words, the peak state s represents

the peak sum of powers in the current prediction horizon. Including (3.22) in (3.11) with a cost on the peak state s yields:

$$\min_{\boldsymbol{x},\boldsymbol{u},s} \ \varphi(s) + \sum_{h=1}^{H} \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \tag{3.23a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.23b}$$

$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.23c}$$

$$\sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}) \leq s \tag{3.23d}$$

Where $\varphi(s)$ is the cost function associated with the peak state s. Because peak power consumption is penalized but not constrained, a controller that considers (3.23) always has the flexibility to let the peak total power consumption increase if other demands (such as a freezing room temperature) are more pressing.

Dualizing (3.22) in (3.23) yields:

$$\begin{aligned} d(\boldsymbol{\mu}) \ = \ \min_{\boldsymbol{x},\boldsymbol{u},s} \ & \varphi(s) - \sum_{i=0}^{N-1} \mu_i \cdot s \\ & + \sum_{h=1}^{H} \Big(\Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) + \sum_{i=0}^{N-1} \mu_i \cdot P_h(\boldsymbol{u}_{i,h}) \Big) \end{aligned} \tag{3.24a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.24b}$$

$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.24c}$$

We may decompose (3.24) into a house- and peak state component:

$$d_h(\boldsymbol{\mu}) = \min_{\boldsymbol{x}_h,\boldsymbol{u}_h} \ \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) + \sum_{i=0}^{N-1} \mu_i \cdot P_h(\boldsymbol{u}_{i,h}) \tag{3.25a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.25b}$$

$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.25c}$$

and

$$d_s(\boldsymbol{\mu}) = \min_{s} \ \varphi(s) - \left(\sum_{i=0}^{N-1} \mu_i \right) \cdot s \tag{3.26}$$

Equation (3.26) solved for s by satisfying:

$$\varphi'(s) - \sum_{i=0}^{N-1} \mu_i = 0 \tag{3.27}$$

Note that for a linear peak cost $\varphi(s)$, (3.26) is unbounded or ill-posed. Indeed, for

$$\varphi(s) = w \cdot s \tag{3.28}$$

problem (3.26) yields

$$d_s(\boldsymbol{\mu}) = -\infty \quad \text{for} \quad w - \sum_{i=0}^{N-1} \mu_i \neq 0 \tag{3.29}$$

with $s = \infty$ for $w < \sum_{i=0}^{N-1} \mu_i$ and $s = -\infty$ for $w > \sum_{i=0}^{N-1} \mu_i$. We have

$$d_s(\boldsymbol{\mu}) = 0 \quad \text{for} \quad w - \sum_{i=0}^{N-1} \mu_i = 0 \tag{3.30}$$

with s undefined. The function $d_s(\boldsymbol{\mu})$ is therefore discontinuous, but requires that

$$w - \sum_{i=0}^{N-1} \mu_i = 0 \tag{3.31}$$

in order to be maximized. The dual problem

$$\begin{aligned} \max_{\boldsymbol{\mu}} \quad & d_s(\boldsymbol{\mu}) + \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \\ s.t. \quad & \boldsymbol{\mu} \geq 0 \end{aligned} \tag{3.32}$$

then translates into

$$\max_{\boldsymbol{\mu}} \quad \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \tag{3.33a}$$

$$s.t. \quad w = \sum_{i=0}^{N-1} \mu_i \tag{3.33b}$$

$$\boldsymbol{\mu} \geq 0 \tag{3.33c}$$

The constraint (3.33b) is analogous to marginal-cost pricing in economics, which is the practice of setting the price of a product to equal the extra cost of producing an extra unit of that product. The peak state s may be viewed as the product, w the cost of producing an extra unit of s (the marginal cost), and $\sum_{i=0}^{N-1} \mu_i$ the 'price' of s.

The (total) derivatives of $d_h(\boldsymbol{\mu})$ are given by:

$$\frac{d}{d\boldsymbol{\mu}} d_h(\boldsymbol{\mu}) = \frac{\partial \mathcal{L}_h}{\partial \boldsymbol{\mu}} \tag{3.34}$$

where $\mathcal{L}_h$ is the Lagrange function associated with (3.25), evaluated at the primal-dual solutions of (3.25) and given by:

$$\begin{aligned} \mathcal{L}_h \; &= \lambda_0 \cdot (\boldsymbol{x}_{0,h} - \bar{\boldsymbol{x}}_h) \; + \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \\ &+ \sum_{i=0}^{N-1} \Big(\mu_i \cdot P_h(\boldsymbol{u}_{i,h}) + \lambda_{i+1} \big(\boldsymbol{x}_{i+1,h} - \boldsymbol{f}_h(\boldsymbol{x}_{i,h}, \boldsymbol{u}_{i,h}) \big) \Big) \end{aligned} \tag{3.35}$$

Then its partial derivative w.r.t. $\boldsymbol{\mu}$ reads as:

$$\frac{\partial \mathcal{L}_h}{\partial \boldsymbol{\mu}} = \begin{bmatrix} P_h(\boldsymbol{u}_{0,h}) \\ \vdots \\ P_h(\boldsymbol{u}_{N,h}) \end{bmatrix}^\top \tag{3.36}$$

Hence the gradient of the cost (3.33a) reads as:

$$\nabla_{\mu_i} \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) = \sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}) \tag{3.37}$$

Problem (3.33) is well suited for a proximal gradient method, where the gradient step on the cost is projected in the Euclidian sense onto the feasible domain of the problem. In the case of (3.33), this projection reads as:

$$\min_{\boldsymbol{\mu}_\perp} \quad \frac{1}{2} ||\boldsymbol{\mu}_\perp - \boldsymbol{\mu}_+||^2 \tag{3.38a}$$

$$\begin{aligned} s.t. \quad & w - \mathbf{1}^\top \boldsymbol{\mu}_\perp = 0 \\ & \boldsymbol{\mu}_\perp \geq 0 \end{aligned} \tag{3.38b}$$

where $\boldsymbol{\mu}_+$ is the "free" gradient step:

$$\boldsymbol{\mu}_+ = \boldsymbol{\mu} + \alpha \cdot \nabla_{\mu_i} \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \tag{3.39}$$

and $\boldsymbol{\mu}_\perp$ is the projected gradient step.

3.3.3 Quadratic Peak Cost

To formulate the peak state s without the advanced projection of 3.3.2, we again consider (3.23), but with $\varphi(s)$ as a quadratic function:

$$\varphi(s) \; = \; w s^2 \tag{3.40}$$

and

$$d_s(\boldsymbol{\mu}) = \min_s \; w s^2 - \left(\sum_{i=0}^{N-1} \mu_i \right) \cdot s \tag{3.41}$$

Then (3.27) becomes:

$$2 w s - \sum_{i=0}^{N-1} \mu_i = 0 \tag{3.42}$$

Which can be solved for s. The dual problem (3.32) can then be solved without transforming the problem $d_s(\boldsymbol{\mu})$ to a constraint. The gradient of (3.32) is then defined as:

$$\nabla_{\mu_i} d(\boldsymbol{\mu}) = \nabla_{\mu_i} \left(d_s(\boldsymbol{\mu}) + \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \right) = -s + \sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}) \tag{3.43}$$

3.3.4 Hourly Peak State

The peak state s penalizes momentary peak states. The new capacity cost, outlined in section 1.3.2 penalizes high hourly consumption. Although this capacity cost will be applied to homes individually, it is still interesting to consider in the context of the residential power coalitions presented in section 1.3.3.

Unlike momentary power use, which is used to reduce peak power in 3.3.2, we do not directly control hourly power. The controller needs a way to 'see' hourly power consumption, and reduce its highest value if need be. In line with the hourly consumption-based capacity cost, we formulate a new centralized optimization problem with a single peak hourly power consumption state E:

$$\min_{\boldsymbol{x},\boldsymbol{u},E} \; \varphi(E) + \sum_{h=1}^{H} \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \tag{3.44a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.44b}$$

$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.44c}$$

$$E \geq e_j \quad \forall j \tag{3.44d}$$

The constraint (3.44d) states that E must be larger than the hourly power consumption e_j for all hours j. In other words, the peak state E represents the highest hourly consumption for the group of houses. We define the power consumption at hour j as follows:

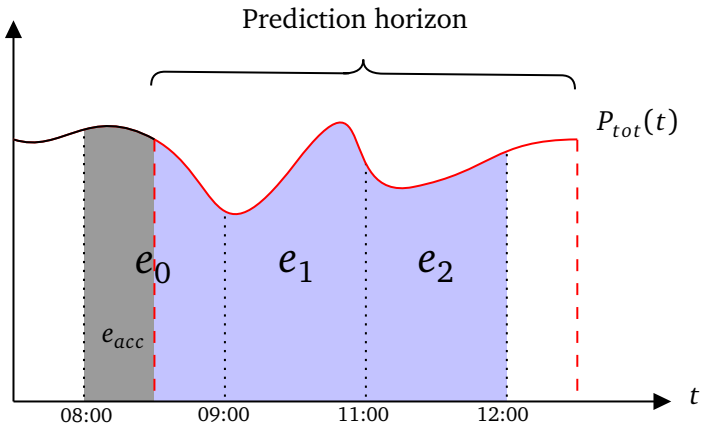

Figure 3.8: Principle of hourly power consumptions e_j and accumulated power for the first hour e_{acc}

$$\begin{aligned} e_j &= \sum_{h=1}^{H} \sum_{i \in \Theta_j} P_h(\boldsymbol{u}_{i,h}) \quad \forall j \neq 1 \\ e_0 &= e_{acc} + \sum_{h=1}^{H} \sum_{i \in \Theta_0} P_h(\boldsymbol{u}_{i,h}) \end{aligned} \tag{3.45}$$

The index j represents every whole hour in the horizon N, including the current hour. Θ_j is the set of all prediction steps i that fall into the hour j. Since the prediction horizon moves one 5-minute step forward at each iteration, some prediction steps i will fall into the next whole hour at the new iteration. As such, each set Θ_j will change at each new iteration of the MPC.

The variable e_{acc} is the total power consumed so far in the current hour. It must be added to the first total hourly power consumption e_0 so that it actually represents the total power consumption for that hour, including power consumption from earlier that hour. Figure 3.8 shows how each variable e_j represents the power consumed for each whole hour, and how e_{acc} represents the power consumption for the first hour that falls outside the prediction horizon. e_{acc} is updated by the following rules:

$$\begin{aligned} e_{acc} &= e_{acc} + \sum_{h=1}^{H} P(\boldsymbol{u}_{0,h}) \quad \text{after each time step} \\ e_{acc} &= 0 \quad \text{when a new hour begins} \end{aligned} \tag{3.46}$$

In other words the first computed power use is added to e_{acc} at each time step until a new hour is reached, at which point it is set back to zero. Since the energy consumed in the current hour includes previous consumption, the accumulated

power consumption for the current hour e_{acc} is added to e_0. By dualizing the constraint (3.44d) we obtain:

$$\begin{aligned} d(\boldsymbol{\mu}) = \min_{\boldsymbol{x},\boldsymbol{u},E} \ & \varphi(E) + \sum_{h=1}^{H} \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \\ & + \sum_j \mu_j \Big(e_j - E \Big) \end{aligned} \tag{3.47a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.47b}$$

$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.47c}$$

We seek the solution of

$$\max_{\mu} \ d(\mu) \tag{3.48a}$$

$$\text{s.t. } \mu \geq 0 \tag{3.48b}$$

We want to decompose the dual problem into multiple sub-problems:

$$\max_{\boldsymbol{\mu}} \ d_E(\boldsymbol{\mu}) + \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \tag{3.49a}$$

$$\text{s.t. } \boldsymbol{\mu} \geq 0 \tag{3.49b}$$

Where $d_E(\boldsymbol{\mu})$ is the problem associated with the hourly peak state, and $d_h(\boldsymbol{\mu})$ is the problem associated with house h. To do this, we first consider the hourly power consumption e_j on a per-house basis:

$$\begin{aligned} e_j &= \sum_{h=1}^{H} e_{j,h} \quad \forall j \\ e_{j,h} &= \sum_{i \in \Theta_j} P_h(\boldsymbol{u}_{i,h}) \quad \forall j \neq 1 \\ e_{0,h} &= e_{acc,h} + \sum_{i \in \Theta_0} P_h(\boldsymbol{u}_{i,h}) \\ e_{acc,h} &= e_{acc,h} + P_h(\boldsymbol{u}_{0,h}) \quad \text{after each time step} \\ e_{acc,h} &= 0 \quad \text{at each new hour} \end{aligned} \tag{3.50}$$

Where $e_{acc,h}$ is the accumulated hourly power consumption for house h. Using (3.50) we may consider (3.47) on a per-house basis:

$$d_h(\boldsymbol{\mu}) = \min_{\boldsymbol{x},\boldsymbol{u}} \ \Phi(\boldsymbol{x}_h, \boldsymbol{u}_h) + \sum_j \mu_j e_{j,h} \tag{3.51a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.51b}$$
$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.51c}$$

and a problem for the peak hourly power consumption E

$$d_E(\boldsymbol{\mu}) = \min_E \ \varphi(E) - \sum_j \mu_j \cdot E \tag{3.52}$$

Where the gradient of the dual function is given by

$$\nabla_{\mu_j} d(\boldsymbol{\mu}) = \nabla_{\mu_j} \Big(d_E(\boldsymbol{\mu}) + \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \Big) = -E + \sum_{h=1}^{H} e_{j,h} \tag{3.53}$$

Judging from the pitfalls detailed in 3.3.2, a linear cost may not be the best choice for $\varphi(E)$. A quadratic peak cost similar to 3.3.3 is likely a better candidate:

$$\varphi(E) = E^2 \tag{3.54}$$

In which case (3.52) is a well-posed optimization problem, and does not require any additional modification to work with a distributed approach.

3.4 Distributed MPC

Figure 3.9 shows how a distributed model predictive control approach might be implemented in the case of house control. Note the similarity with figure 3.3, which shows a decentralized house control approach. In both cases there is a controller for each house, so a failure in a controller only affects one house. The houses do not need to share all their information with some centralized controller, as for the centralized approach, shown in figure 3.2.

Unlike the decentralized approach, the distributed approach uses a coordinator to update the dual variables $\boldsymbol{\mu}$ associated with the coupling constraints between the houses. Also note that in addition to the controllers for each house, the final formulations described in 3.5.2 and 3.5.2 require a separate controller for the peak states s and E.

3.4.1 DMPC Algorithm

Algorithm 1 shows the algorithmic formulation of the DMPC approach used in this thesis. The algorithm is designed to perform dual decomposition on a Model Predictive Control problem with a coupling constraint. P_{ext} is the external power use, which we do not control and only know for the current time step. T_{out} is the outdoor temperature. T_{ref} is the temperature reference.

w_{lb} and w_{ub} are the lower and upper bounds on the optimization state w. T is the total simulation time, and t is the current time step. k is the current dual decomposition step, f_{diff} and μ_{diff} are the changes in the sum of cost functions

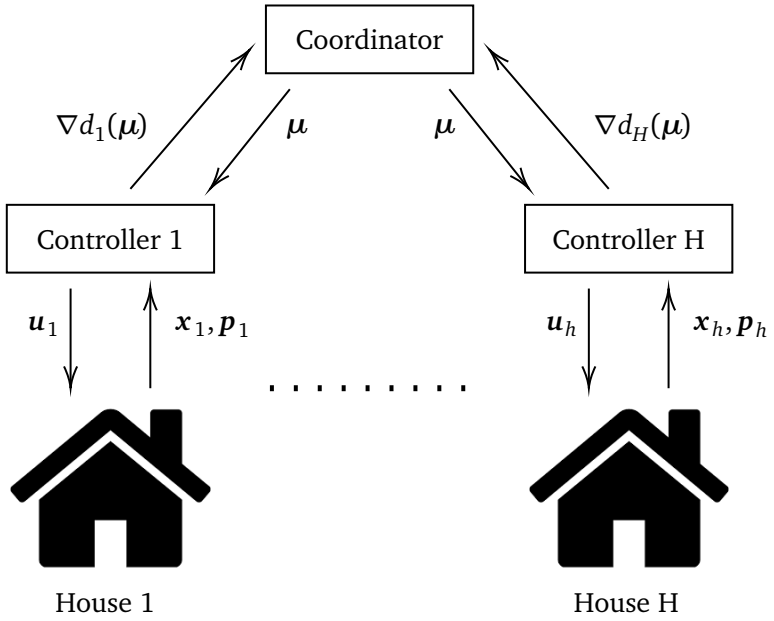

Figure 3.9: Principle of distributed MPC for house control

and dual variables, respectively. f_{tol}, μ_{tol} and k_{max} form the termination criteria for the dual decomposition.

μ contains the dual variables. f_i^* and w_i^* is the local optimal value and optimal state for controller i. c_μ is a constant based on any constant values in the coupling constraint, e.g. h_0 in (2.8). $\mu_i^+(w_i^*)$ is the gradient update contribution to μ from controller i, computed using the optimal local state w_i^*. α is some step size for the gradient method, and P is a function for projecting the updated μ onto its feasible region. This will typically be a function to set all negative values to zero, associated with the constraint $\mu \geq 0$, though it may be some other function (see section 3.3.2).

Algorithm 1 Distributed model predictive control

$t = 0$
while $t < T$ **do**
  Controllers prepare P_{ext}, T_{out}, T_{ref}
  Controllers prepare state constraints w_{lb}, w_{ub}
  $k = 0$
  while $\left(f_{diff} > f_{tol} \text{ or } \mu_{diff} > \mu_{tol}\right)$ and $k < k_{max}$ **do**
    Submit dual variable μ to controllers
    Controllers compute local solutions w_i^*, f_i^* given μ
    $\mu^+ = c_\mu + \sum_i \mu_i^+(w_i^*)$ ▷ Dual gradient from local solutions
    $\mu = P\left(\mu + \alpha\mu^+\right)$ ▷ Projected gradient method
    $f_{sum} = \sum_i f_i^*$
    $\mu_{diff} = ||\mu - \mu_{last}||, \quad f_{diff} = |f_{sum} - f_{sum,last}|$
    $\mu_{last} = \mu, \quad f_{sum,last} = f_{sum}, \quad k = k + 1$
  end while
  Controllers apply solution and move to next time step
  $t = t + 1$
end while

3.5 Final Formulations

These are the formulations that are implemented. They are based on the peak reduction formulations in 3.3. Each formulation has a centralized version, and a corresponding distributed version obtained from dual decomposition. Both versions are simulated in chapter 4.

3.5.1 Formulation 1

Formulation 1 is based on the formulation detailed in section 3.3.2, a **momentary peak state** s **with a linear cost**. The centralized formulation is as follows:

$$\min_{\boldsymbol{x},\boldsymbol{u},s} \quad w_{lin}\, s + \sum_{h=1}^{H} \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \tag{3.55a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.55b}$$

$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.55c}$$

$$\sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}) \leq s \tag{3.55d}$$

And the dual decomposition-based distributed formulation is

$$\max_{\boldsymbol{\mu}} \quad \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \tag{3.56a}$$

$$s.t. \quad w_{lin} - \sum_{i=0}^{N-1} \mu_i = 0 \tag{3.56b}$$

$$\boldsymbol{\mu} \geq 0 \tag{3.56c}$$

where

$$d_h(\boldsymbol{\mu}) = \min_{\boldsymbol{x}_h, \boldsymbol{u}_h} \; \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) + \sum_{i=0}^{N-1} \mu_i \cdot P_h(\boldsymbol{u}_{i,h}) \tag{3.57a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.57b}$$

$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.57c}$$

At each new iteration of a gradient method used to maximize $\boldsymbol{\mu}$, the constraints (3.56b) and (3.56c) are enforced using a proximal gradient method:

$$\min_{\boldsymbol{\mu}_\perp} \quad \frac{1}{2} ||\boldsymbol{\mu}_\perp - \boldsymbol{\mu}_+||^2 \tag{3.58a}$$

$$\begin{aligned} s.t. \quad & w - \mathbf{1}^\top \boldsymbol{\mu}_\perp = 0 \\ & \boldsymbol{\mu}_\perp \geq 0 \end{aligned} \tag{3.58b}$$

where $\boldsymbol{\mu}_\perp$ is the projected gradient step. In the context of the distributed model predictive control algorithm presented in section 3.4.1, the dual variable projection step, represented with the function P, amounts to solving this optimization problem. $\boldsymbol{\mu}_+$ is the "free", unprojected, gradient step:

$$\boldsymbol{\mu}_+ = \boldsymbol{\mu} + \alpha \cdot \nabla_\mu \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \tag{3.59}$$

The gradient of (3.56a), which is used in a gradient method to obtain $\boldsymbol{\mu}_+$, is given by

$$\nabla_{\mu_i} \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) = \sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}^*) \tag{3.60}$$

Where $P_h(\boldsymbol{u}_{i,h}^*)$ is the optimal power output at time step i for the controller for house h, given a fixed dual variable $\boldsymbol{\mu}$. It is obtained from the solution of the minimization problem in (3.57).

3.5.2 Formulation 2

Formulation 2 is based on the formulation detailed in section 3.3.3, a **momentary peak state s with a quadratic cost**. The centralized formulation is as follows:

$$\min_{\boldsymbol{x}, \boldsymbol{u}, s} \; w_{quad}\, s^2 + \sum_{h=1}^{H} \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \tag{3.61a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.61b}$$
$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.61c}$$
$$\sum_{h=1}^{H} P_h(\boldsymbol{u}_{i,h}) \leq s \tag{3.61d}$$

And the distributed formulation is

$$\begin{aligned} \max_{\boldsymbol{\mu}} \quad & d_s(\boldsymbol{\mu}) + \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \\ s.t. \quad & \boldsymbol{\mu} \geq 0 \end{aligned} \tag{3.62}$$

where

$$d_s(\boldsymbol{\mu}) = \min_{s} \; w_{quad}\, s^2 - \sum_{i=0}^{N-1} \mu_i\, s \tag{3.63}$$

and

$$d_h(\boldsymbol{\mu}) = \min_{\boldsymbol{x}_h, \boldsymbol{u}_h} \; \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) + \sum_{i=0}^{N-1} \mu_i \cdot P_h(\boldsymbol{u}_{i,h}) \tag{3.64a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.64b}$$
$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.64c}$$

Updating the dual variable $\boldsymbol{\mu}$ is done as follows:

$$\boldsymbol{\mu} = \Big[\boldsymbol{\mu} + \alpha\, \nabla_{\mu} d(\boldsymbol{\mu}) \Big]_{+} \tag{3.65}$$

Where $[\,]_+$ is the dual variable projection that ensures the positivity constraint $\boldsymbol{\mu} \geq 0$ is satisfied, and the gradient of the dual function is given by

$$\nabla_{\mu_i} d(\boldsymbol{\mu}) = \nabla_{\mu_i} \Big(d_s(\boldsymbol{\mu}) + \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \Big) = -s^* + \sum_{h=1}^{H} P_h(\boldsymbol{u}^*_{i,h}) \tag{3.66}$$

Where $P_h(\boldsymbol{u}^*_{i,h})$ is obtained from each house-controller solving (3.64), and s^* is obtained through solving (3.63).

3.5.3 Formulation 3

Formulation 3 is based on the formulation detailed in section 3.3.4, an **hourly peak state E with a quadratic cost**. The centralized formulation is as follows:

$$\min_{\boldsymbol{x}, \boldsymbol{u}, E} \; w_{hourly}\, E^2 + \sum_{h=1}^{H} \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \tag{3.67a}$$

$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.67b}$$
$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.67c}$$
$$E \geq e_j \quad \forall j \tag{3.67d}$$

The hourly peak state E and the variables representing the power consumption for each whole hour, e_j, as well as its decomposition, are explained in more detail in 3.3.4. The distributed formulation is

$$\max_{\boldsymbol{\mu}} \; d_E(\boldsymbol{\mu}) + \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \tag{3.68a}$$
$$\text{s.t. } \boldsymbol{\mu} \geq 0 \tag{3.68b}$$

where

$$d_E(\boldsymbol{\mu}) = \min_{E} \; w_{hourly}\, E^2 - \sum_j \mu_j \cdot E \tag{3.69}$$

and

$$d_h(\boldsymbol{\mu}) = \min_{\boldsymbol{x}_h, \boldsymbol{u}_h} \; \Phi_h(\boldsymbol{x}_h, \boldsymbol{u}_h) + \sum_j \mu_j e_{j,h} \tag{3.70a}$$
$$\text{s.t. } h_h(\boldsymbol{x}_h, \boldsymbol{u}_h) = 0 \tag{3.70b}$$
$$g_h(\boldsymbol{x}_h, \boldsymbol{u}_h) \leq 0 \tag{3.70c}$$

Similarly to 3.5.2, updating the dual variable $\boldsymbol{\mu}$ is done as follows:

$$\boldsymbol{\mu} = \Big[\boldsymbol{\mu} + \alpha \, \nabla_{\mu} d(\boldsymbol{\mu}) \Big]_+ \tag{3.71}$$

The gradient of the dual function is given by

$$\nabla_{\mu_j} d(\boldsymbol{\mu}) = \nabla_{\mu_j} \Big(d_E(\boldsymbol{\mu}) + \sum_{h=1}^{H} d_h(\boldsymbol{\mu}) \Big) = -E^* + \sum_{h=1}^{H} e_{j,h}^* \tag{3.72}$$

Chapter 4

Results

The results were generated on an Intel i7-9700k 8-core desktop processor. Parallel computation is necessary to properly evaluate the computational performance of the decentralized and distributed approaches since parallelism is a key feature (see 2.1.2). Parallel computation was achieved using the multiprocessing module in Python.

The proposed centralized and distributed peak reduction formulations are validated in two scenarios: one considering a coalition of 2 houses and one considering a coalition of 8 houses, respectively. The scenarios simulate 24 hours at a resolution of 5 minutes. The MPCs have a prediction horizon of 24 hours. The high spot prices and low outdoor temperature provided to the simulations represent a challenging validation scenario.

Each house is assumed to have one room with one heat pump. We assume that the houses register the same outdoor temperature, spot price, reference temperature, and predicted external power use. Each house will differ in its real external power use. The real external power is provided based on real data collected from the house in Trondheim at various two-day intervals. To fit the validation scenario, the external power use is scaled down to account for the smaller size of the house model considered in the simulations compared to the large residential home from which the external power use is gathered. It is scaled by 0.33.

This thesis proposes a decentralized approach with MPC formulation (3.10), which has no formulations for peak reduction (details are outlined in section 3.1.2). For validation purposes, the centralized and distributed MPC formulations are compared to the same decentralized MPC formulation.

For each scenario, the formulations and MPC approaches are validated on their peak reduction capabilities, their ability to keep the room temperature at desired levels and keep power costs down, as well as their computational performance. The performance measures for peak reduction are the highest aggregated momentary and hourly power consumption. Performance with respect to comfort is measured in total deviations in room temperature from the reference temperature. Power use performance is measured in both total power use as well as total power costs. Computational performance is measured in average MPC iteration

time.

The results presented in this chapter are discussed in chapter 5

4.1 Formulation 1

This section shows the results of implementing the formulation in 3.5.1, which contains a single peak state s that represents the highest momentary power consumption in a prediction horizon. This peak state s is penalized with a linear cost.

4.1.1 2 Houses

Table 4.1 shows the highest hourly total power consumption in the centralized, distributed, and decentralized approaches, with two houses. The centralized and distributed approaches have similar results, slightly lower than the decentralized approach.

Table 4.1: Max hourly power consumption, formulation 1, 2 houses

	Max hourly power consumption [kWh]
Centralized MPC	1.564
Decentralized MPC	1.793
Distributed MPC	1.580

Table 4.2 shows the highest momentary total power consumption in the centralized, distributed, and decentralized approaches, with two houses. The centralized and distributed approaches have slightly lower values than the decentralized approach.

Table 4.2: Max power consumption, formulation 1, 2 houses

	Max power consumption [kW]
Centralized MPC	2.070
Decentralized MPC	2.367
Distributed MPC	2.070

Table 4.3 shows the total temperature deviations in the same scenario. The decentralized approach has by far the lowest temperature deviations, and the distributed approach has slightly higher temperature deviations than the centralized approach.

Table 4.4 shows the total cost from heat pump power use. The results are quite similar for all three approaches.

Table 4.5 shows the total heat pump power use. The decentralized approach uses a little more power in total than the centralized and distributed approach.

Figure 4.1 shows the power use, including external power, for the three approaches. The distributed approach follows the centralized approach well, and

Table 4.3: Room temperature deviation, formulation 1, 2 houses

	Room temperature deviation [°Ch]
Centralized MPC	25.573
Decentralized MPC	16.464
Distributed MPC	24.281

Table 4.4: Heat pump power consumption cost, formulation 1, 2 houses

	Heat pump power consumption cost [kr]
Centralized MPC	26.669
Decentralized MPC	26.492
Distributed MPC	26.697

Table 4.5: Total heat pump power consumption, formulation 1, 2 houses

	Total heat pump power consumption [kWh]
Centralized MPC	9.813
Decentralized MPC	10.108
Distributed MPC	9.833

certain spikes in power use that are present in the decentralized approach are slightly reduced in the centralized and distributed approaches.

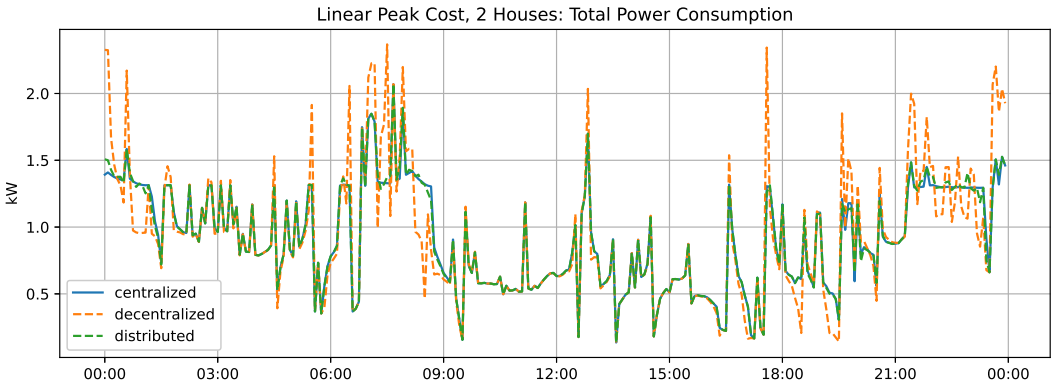

Figure 4.1: Comparison of total power use, formulation 1, 2 houses

Figure 4.2 shows the room temperature in the 2 houses for the three approaches. The distributed approach mostly follows the centralized approach, with some discrepancies around the time 18-20. These approaches are also further from the temperature reference than the centralized approach at certain points of time.

Figure 4.3 shows the energy consumption at each hour of the day. All three approaches are largely similar, but the highest value at 6-7 in the morning is smaller for the centralized and decentralized approaches.

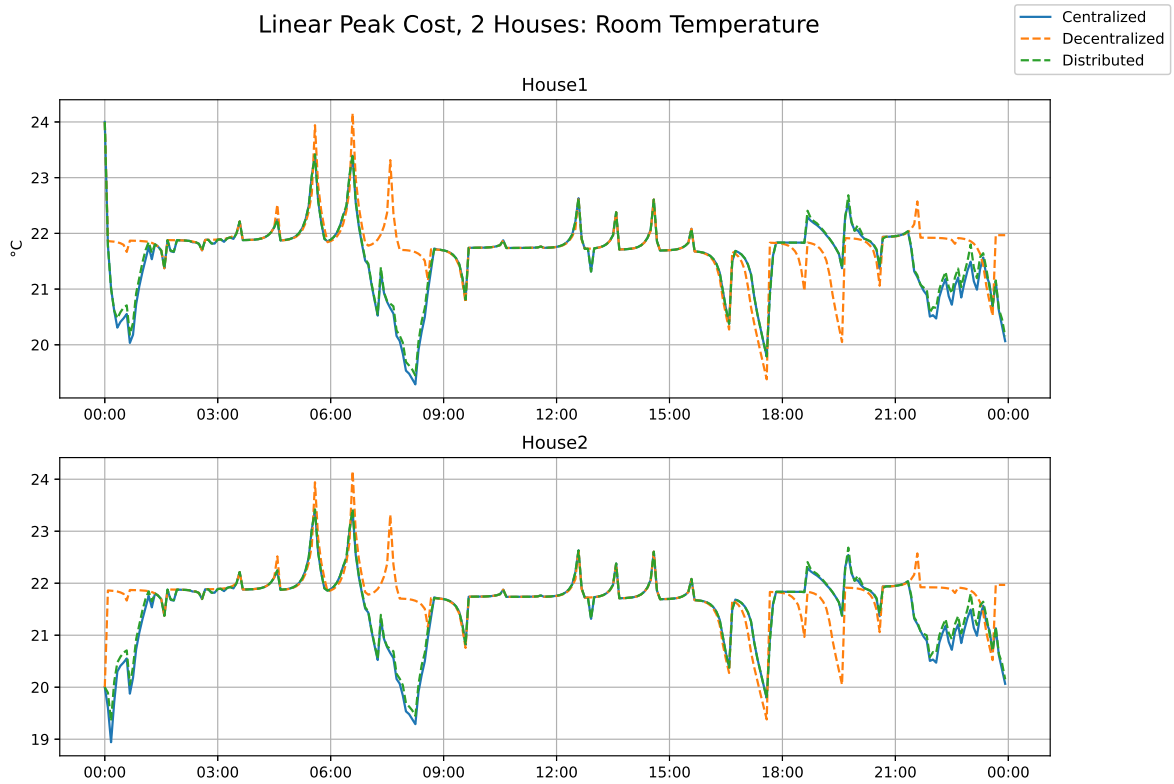

Figure 4.2: Room temperature comparison, formulation 1, 2 houses

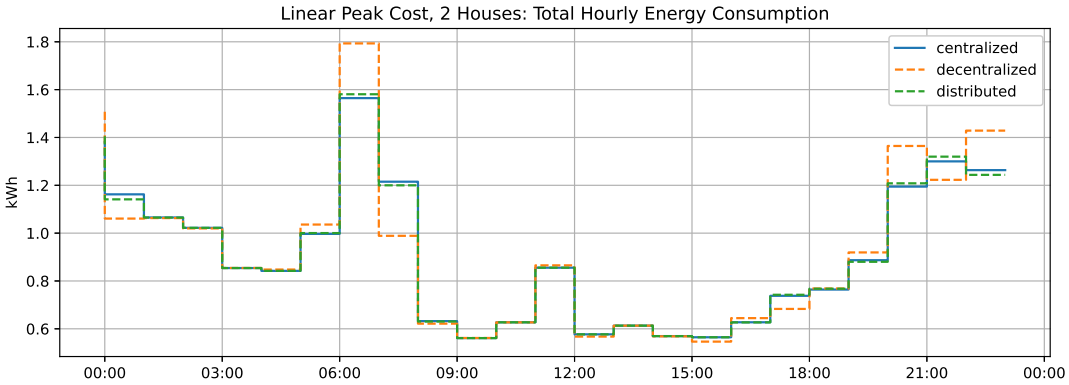

Figure 4.3: Hourly total power consumption, formulation 1, 2 houses

4.1.2 8 Houses

These are the results using formulation 1, simulating on 8 houses. Table 4.6 shows the highest hourly total power consumption in the centralized, distributed, and

decentralized approaches, with two houses. The centralized and distributed approaches have nearly identical results, slightly lower than the decentralized approach.

Table 4.6: Max hourly power consumption, formulation 1, 8 houses

	Max hourly power consumption [kWh]
Centralized MPC	5.266
Decentralized MPC	5.739
Distributed MPC	5.265

Table 4.7 shows the highest momentary total power consumption in the centralized, distributed, and decentralized approaches, with two houses. The centralized and distributed approaches have around half the max power consumption of the decentralized approach, with the distributed approach having a little higher max power consumption than the centralized approach.

Table 4.7: Max power consumption, formulation 1, 8 houses

	Max power consumption [kW]
Centralized MPC	5.357
Decentralized MPC	11.232
Distributed MPC	5.390

Table 4.8 shows the total temperature deviations in the same scenario. The decentralized approach has lower temperature deviations than the centralized and distributed approaches.

Table 4.8: Room temperature deviation, formulation 1, 8 houses

	Room temperature deviation [°Ch]
Centralized MPC	78.008
Decentralized MPC	65.857
Distributed MPC	77.360

Table 4.9 shows the total cost from heat pump power use. The results are nearly identical for all three approaches, with slightly lower costs in the decentralized approach.

Table 4.10 shows the total heat pump power use. All three approaches are nearly identical.

Figure 4.4 shows the power use, including external power, for the three approaches. The distributed approach follows the centralized approach nearly perfectly, and they both have lower power peaks than the decentralized approach.

Figure 4.5 shows the energy consumption at each hour of the day. Again the distributed approach follows the centralized approach nearly perfectly. The three

Table 4.9: Heat pump power consumption cost, formulation 1, 8 houses

	Heat pump power consumption cost [kr]
Centralized MPC	105.399
Decentralized MPC	104.657
Distributed MPC	105.445

Table 4.10: Total heat pump power consumption, formulation 1, 8 houses

	Total heat pump power consumption [kWh]
Centralized MPC	39.316
Decentralized MPC	39.755
Distributed MPC	39.364

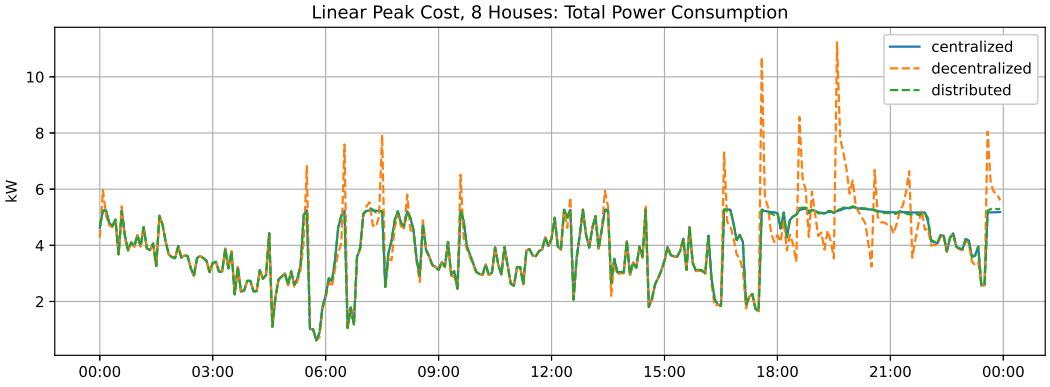

Figure 4.4: Comparison of total power use, formulation 1, 8 houses

approaches are nearly identical, with the exception of 18-19 in the evening when the decentralized approach has higher hourly energy consumption.

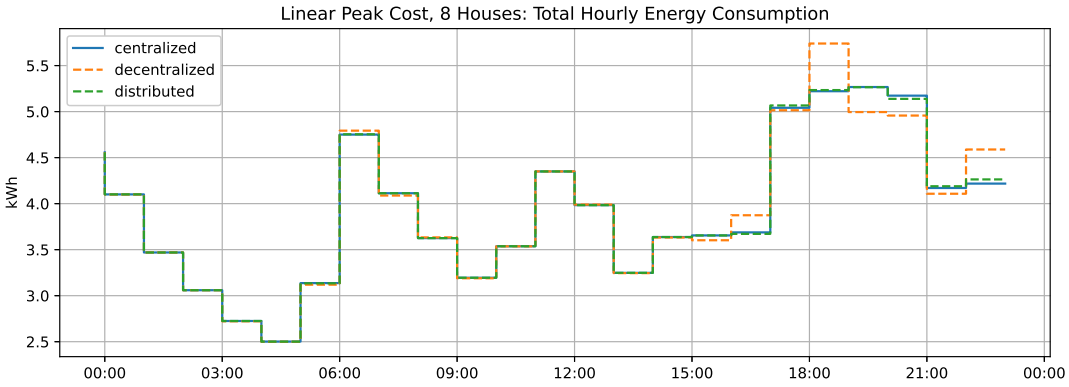

Figure 4.5: Hourly total power consumption, formulation 1, 8 houses

4.2 Formulation 2

This section shows the results of implementing the formulation in 3.5.2, which contains a single peak state s that represents the highest momentary power consumption in a prediction horizon. This peak state s is penalized with a quadratic cost.

4.2.1 2 Houses

Table 4.11 shows the highest hourly total power consumption in the centralized, distributed, and decentralized approaches, with two houses. The centralized and distributed approaches have similar results, slightly lower than the decentralized approach.

Table 4.11: Max hourly power consumption, formulation 2, 2 houses

	Max hourly power consumption [kWh]
Centralized MPC	1.670
Decentralized MPC	1.793
Distributed MPC	1.690

Table 4.12 shows the highest momentary total power consumption in the centralized, distributed, and decentralized approaches, with two houses. The centralized and distributed approaches have slightly lower values than the decentralized approach. These results are identical to formulation 1.

Table 4.13 shows the total temperature deviations in the same scenario. The decentralized approach has the lowest temperature deviations, with the centralized and distributed approaches a few °Ch higher.

Table 4.14 shows the total cost from heat pump power use. The results are nearly identical for all three approaches.

Table 4.12: Max power consumption, formulation 2, 2 houses

	Max power consumption [kW]
Centralized MPC	2.070
Decentralized MPC	2.367
Distributed MPC	2.070

Table 4.13: Room temperature deviation, formulation 2, 2 houses

	Room temperature deviation [°Ch]
Centralized MPC	19.820
Decentralized MPC	16.464
Distributed MPC	18.579

Table 4.15 shows the total heat pump power use. The decentralized approach uses a little more power in total than the centralized and distributed approach, which are nearly identical.

Figure 4.6 shows the power use, including external power, for the three approaches. The distributed approach follows the centralized approach well, and many spikes in power use that are present in the decentralized approach are slightly reduced in the centralized and distributed approaches.

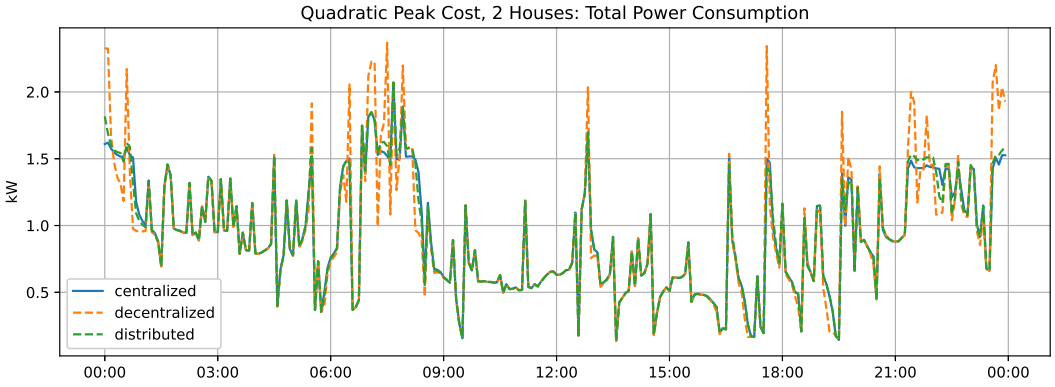

Figure 4.6: Comparison of total power use , formulation 2, 2 houses

Figure 4.7 shows the room temperature in the 2 houses for the three approaches. The room temperature is very similar for all three approaches, with the notable exception that the centralized and distributed approaches have lower room temperatures around 7-8 in the morning.

Figure 4.8 shows the energy consumption at each hour of the day. All three approaches are largely similar, but the highest value at 6-7 in the morning is smaller for the centralized and decentralized approaches.

Table 4.14: Heat pump power consumption cost, formulation 2, 2 houses

	Heat pump power consumption cost [kr]
Centralized MPC	26.564
Decentralized MPC	26.492
Distributed MPC	26.582

Table 4.15: Total heat pump power consumption, formulation 2, 2 houses

	Total heat pump power consumption [kWh]
Centralized MPC	9.878
Decentralized MPC	10.108
Distributed MPC	9.906

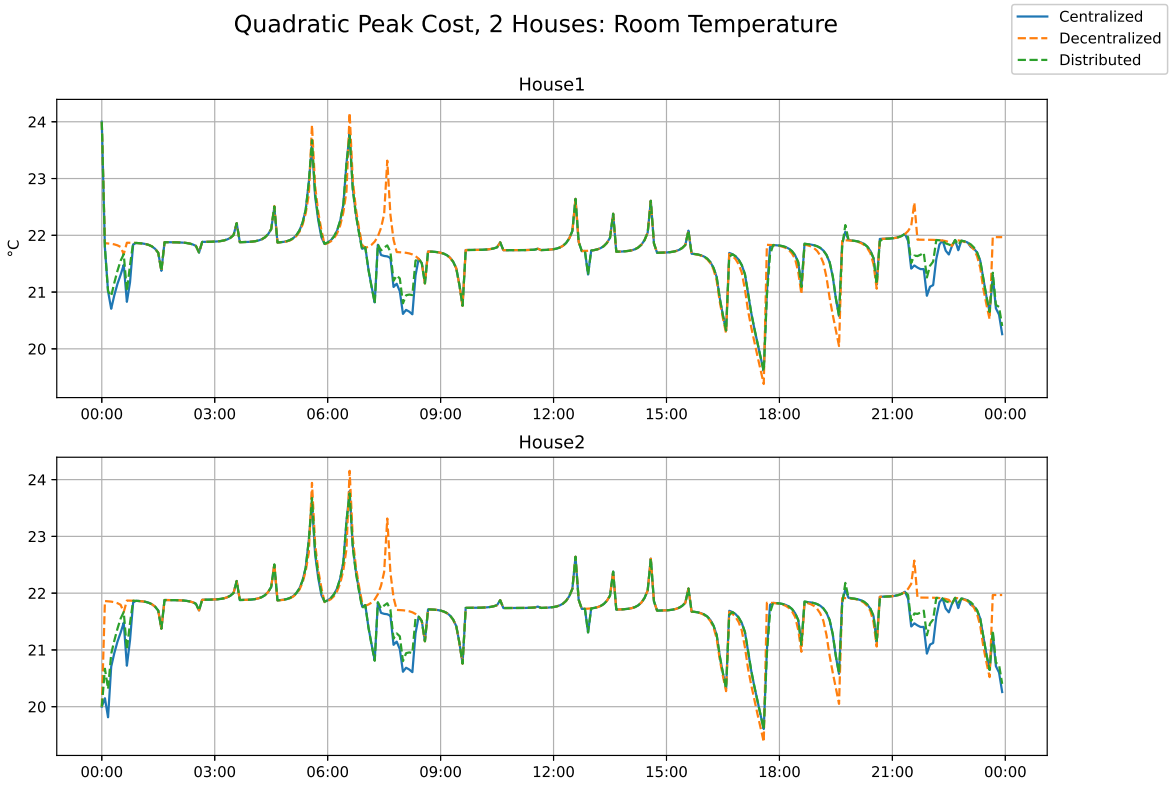

Figure 4.7: Room temperature comparison, formulation 2, 2 houses

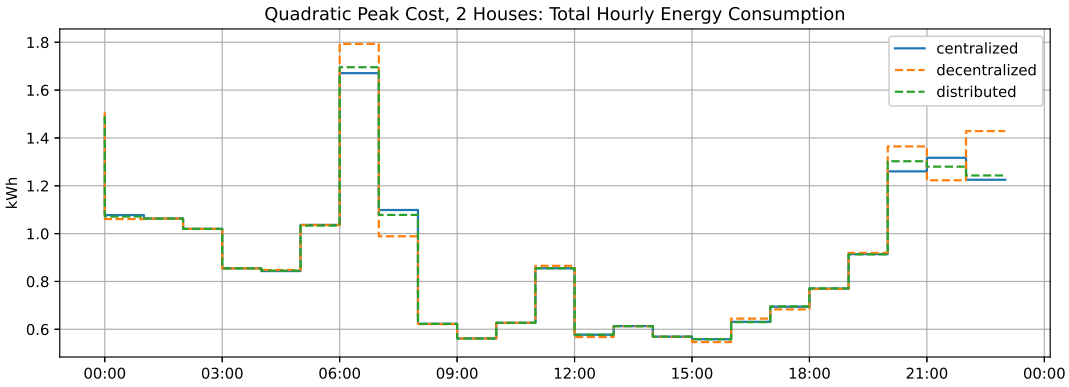

Figure 4.8: Hourly total power consumption, formulation 2, 2 houses

4.2.2 8 Houses

These are the results using formulation 2, simulating on 8 houses. Table 4.16 shows the highest hourly total power consumption in the centralized, distributed, and decentralized approaches, with eight houses. The centralized and distributed approaches have similar results, slightly lower than the decentralized approach.

Table 4.16: Max hourly power consumption, formulation 2, 8 houses

	Max hourly power consumption [kWh]
Centralized MPC	5.257
Decentralized MPC	5.739
Distributed MPC	5.255

Table 4.17 shows the highest momentary total power consumption in the centralized, distributed, and decentralized approaches, with eight houses. The centralized and distributed approaches have around half the max power consumption of the decentralized approach, with the distributed approach having a little higher max power consumption than the centralized approach.

Table 4.17: Max power consumption, formulation 2, 8 houses

	Max power consumption [kW]
Centralized MPC	5.340
Decentralized MPC	11.232
Distributed MPC	5.474

Table 4.18 shows the total temperature deviations in the same scenario. The decentralized approach has lower temperature deviations than the centralized and distributed approaches.

Table 4.18: Room temperature deviation, formulation 2, 8 houses

	Room temperature deviation [°Ch]
Centralized MPC	78.501
Decentralized MPC	65.857
Distributed MPC	75.571

Table 4.19 shows the total cost from heat pump power use. The results are nearly identical for all three approaches, with slightly lower costs in the decentralized approach.

Table 4.20 shows the total heat pump power use. All three approaches are nearly identical.

Figure 4.9 shows the power use, including external power, for the three approaches. Similar to formulation 1, the distributed approach follows the centralized approach nearly perfectly, and they both have lower power peaks than the decentralized approach.

Table 4.19: Heat pump power consumption cost, formulation 2, 8 houses

	Heat pump power consumption cost [kr]
Centralized MPC	105.409
Decentralized MPC	104.657
Distributed MPC	105.541

Table 4.20: Total heat pump power consumption, formulation 2, 8 houses

	Total heat pump power consumption [kWh]
Centralized MPC	39.308
Decentralized MPC	39.755
Distributed MPC	39.417

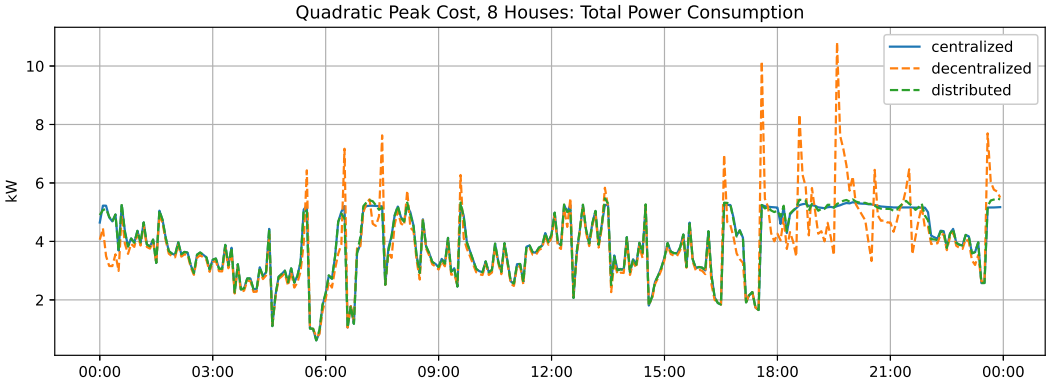

Figure 4.9: Comparison of total power use, formulation 2, 8 houses

Figure 4.10 shows the energy consumption at each hour of the day. Once again the results are very similar to formulation 1. The three approaches are nearly identical, with the exception of 18-19 in the evening when the decentralized approach has higher hourly energy consumption.

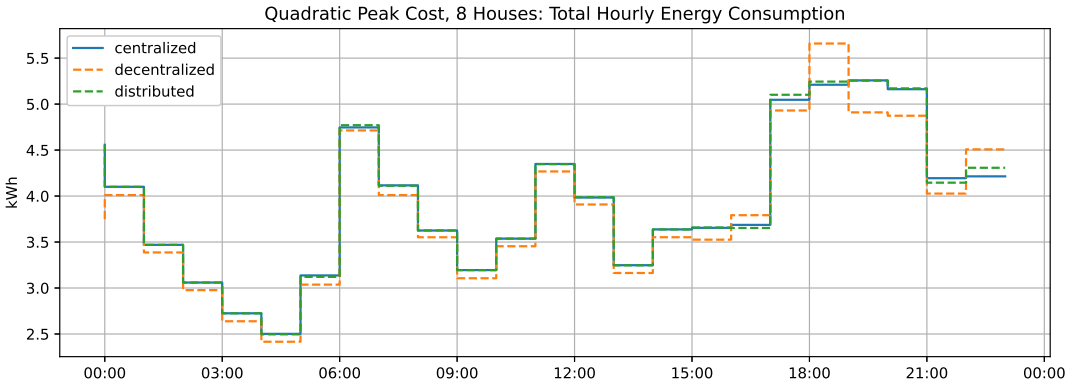

Figure 4.10: Hourly total power consumption, formulation 2, 8 houses

4.3 Formulation 3

This section shows the results of implementing the formulation in 3.5.3, which contains a single peak state E that represents the highest hourly power consumption in a prediction horizon. This peak state E is penalized with a quadratic cost.

4.3.1 2 Houses

Table 4.21 shows the highest hourly total power consumption in the centralized, distributed, and decentralized approaches, with two houses. The centralized and decentralized approaches have a significantly lower hourly peak.

Table 4.21: Max hourly power consumption, formulation 3, 2 houses

	Max hourly power consumption [kWh]
Centralized MPC	1.472
Decentralized MPC	1.793
Distributed MPC	1.481

Table 4.22 shows the highest momentary total power consumption in the centralized, distributed, and decentralized approaches, with two houses. The centralized and distributed approaches have a much higher maximum power consumption.

Table 4.22: Max power consumption, formulation 3, 2 houses

	Max power consumption [kW]
Centralized MPC	3.734
Decentralized MPC	2.367
Distributed MPC	3.734

Table 4.23 shows the total temperature deviations in the same scenario. The centralized and distributed approaches have significantly higher temperature deviations than the decentralized approach.

Table 4.23: Room temperature deviation, formulation 3, 2 houses

	Room temperature deviation [°Ch]
Centralized MPC	41.105
Decentralized MPC	16.464
Distributed MPC	41.543

Table 4.24 shows the total cost from heat pump power use. The results nearly identical for all three approaches.

Table 4.24: Heat pump power consumption cost, formulation 3, 2 houses

	Heat pump power consumption cost [kr]
Centralized MPC	27.235
Decentralized MPC	26.492
Distributed MPC	27.235

Table 4.25 shows the total heat pump power use. The decentralized approach uses a little more power in total than the centralized and distributed approach, which are nearly identical.

Table 4.25: Total heat pump power consumption, formulation 3, 2 houses

	Total heat pump power consumption [kWh]
Centralized MPC	9.813
Decentralized MPC	10.108
Distributed MPC	9.838

Figure 4.11 shows the power use, including external power, for the three approaches. The distributed approach follows the centralized approach well, but they both have more power spikes than the decentralized approach and seem to oscillate every hour.

Figure 4.12 shows the room temperature in the 2 houses for the three approaches. The room temperature in the centralized and distributed approach is very different from the room temperature in the decentralized approach at many points, especially between 18-21 in the evening, where the temperature is as high as 5 degrees higher than the reference temperature.

Figure 4.13 shows the energy consumption at each hour of the day. The three approaches are somewhat similar, but the centralized and distributed approaches have a much smaller hourly consumption than the decentralized approach at 6-7 in the morning.

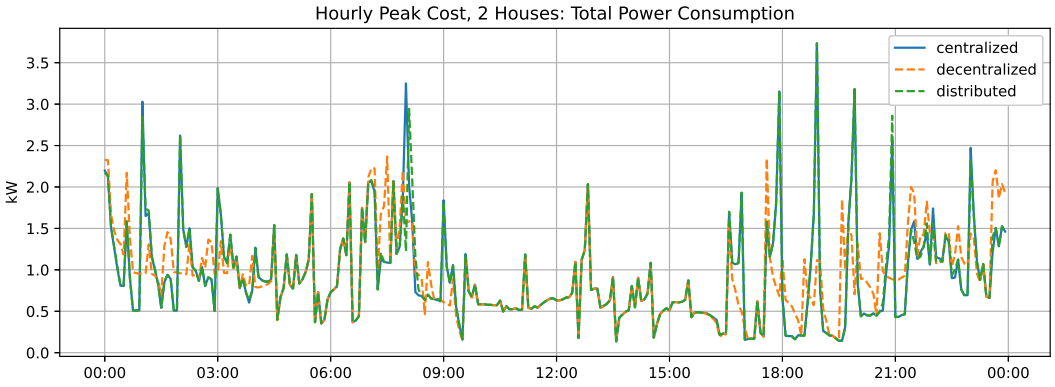

Figure 4.11: Comparison of total power use, formulation 3, 2 houses

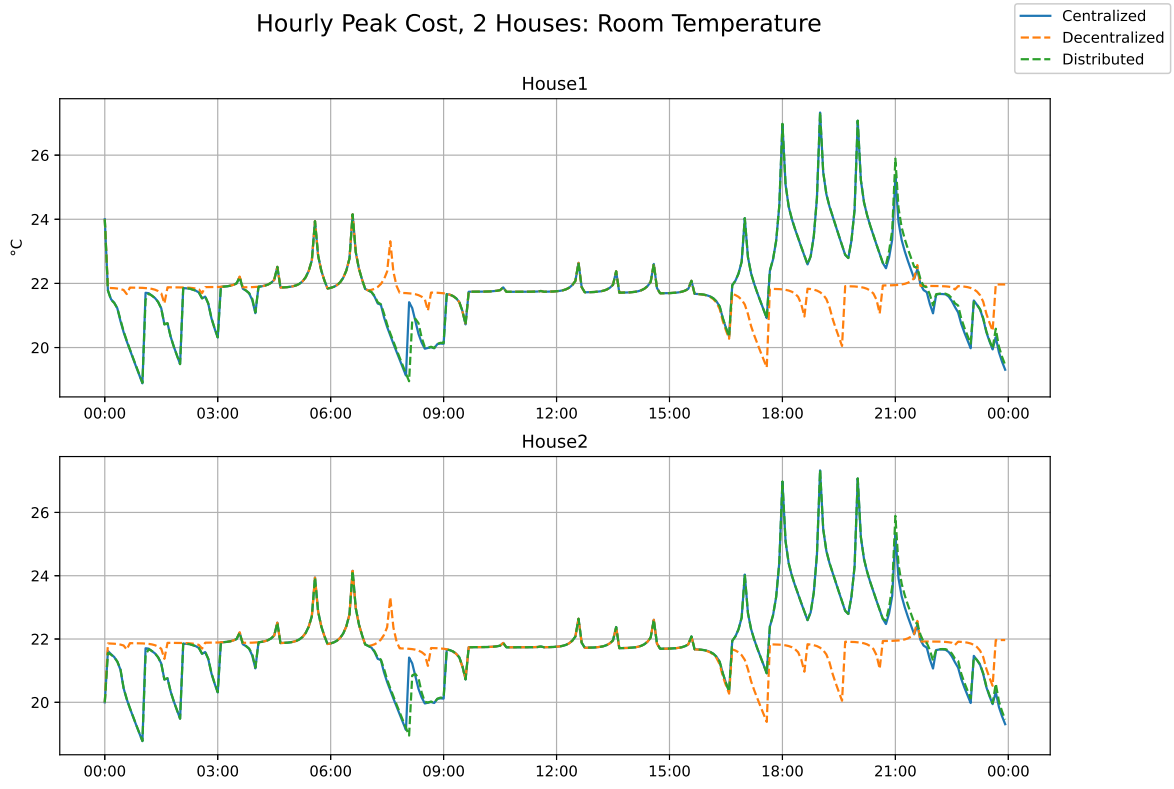

Figure 4.12: Room temperature comparison, formulation 3, 2 houses

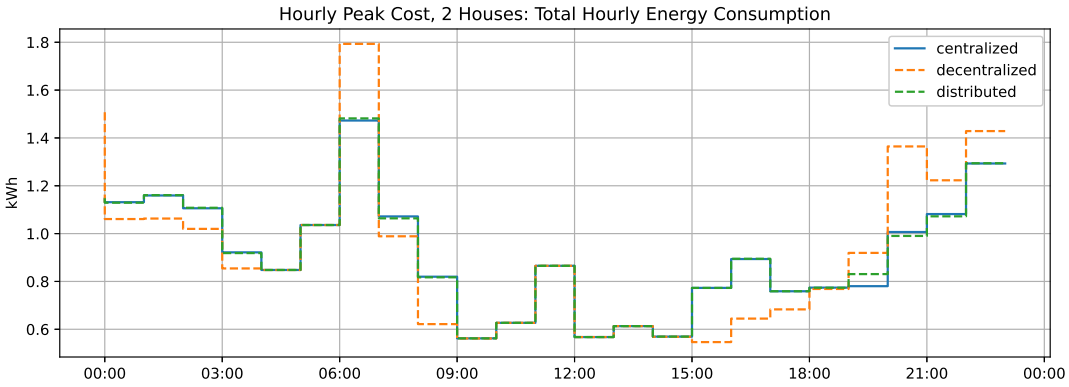

Figure 4.13: Hourly total power consumption, formulation 3, 2 houses

4.3.2 8 Houses

These are the results using formulation 3, simulating on 8 houses. Table 4.26 shows the highest hourly total power consumption. The centralized and distributed approaches are very similar, both lower than the decentralized approach.

Table 4.26: Max hourly power consumption, formulation 3, 8 houses

	Max hourly power consumption [kWh]
Centralized MPC	4.328
Decentralized MPC	5.739
Distributed MPC	4.327

Table 4.27 shows the highest momentary total power consumption in the centralized, distributed, and decentralized approaches, with eight houses.

Table 4.27: Max power consumption, formulation 3, 8 houses

	Max power consumption [kW]
Centralized MPC	13.991
Decentralized MPC	11.232
Distributed MPC	13.883

Table 4.28 shows the total temperature deviations in the same scenario. The temperature deviations are much lower in the decentralized approach compared to the centralized and distributed approaches.

Table 4.28: Room temperature deviation, formulation 3, 8 houses

	Room temperature deviation [°Ch]
Centralized MPC	209.440
Decentralized MPC	65.857
Distributed MPC	212.319

Table 4.29 shows the total cost from heat pump power use. The results are nearly identical for all three approaches.

Table 4.30 shows the total heat pump power use. The decentralized approach uses slightly more power than the centralized and distributed approaches.

Figure 4.14 shows the power use, including external power, for the three approaches. The hourly oscillations present in the 2-house scenario for the centralized and distributed approaches are also present when simulating with 8 houses.

Figure 4.15 shows the energy consumption at each hour of the day. The centralized and distributed approaches have a much more even hourly energy consumption than the decentralized approach.

Table 4.29: Heat pump power consumption cost, formulation 3, 8 houses

	Heat pump power consumption cost [kr]
Centralized MPC	104.189
Decentralized MPC	104.657
Distributed MPC	103.755

Table 4.30: Total heat pump power consumption, formulation 3, 8 houses

	Total heat pump power consumption [kWh]
Centralized MPC	36.343
Decentralized MPC	39.755
Distributed MPC	37.060

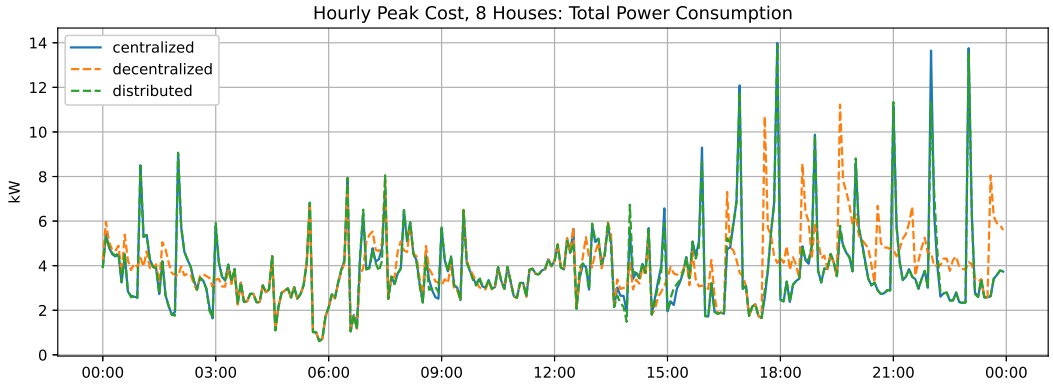

Figure 4.14: Comparison of total power use, formulation 3, 8 houses

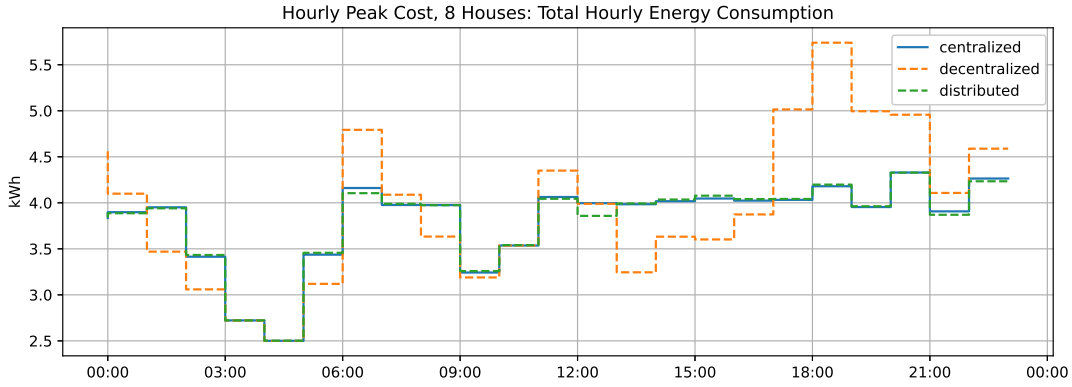

Figure 4.15: Hourly total power consumption, formulation 3, 8 houses

4.4 Computational Performance

To determine the computational performance of the formulations and MPC approaches, their average iteration time during the simulations is considered. The average iteration time is calculated by dividing the total simulation time by the number of time steps. It expresses how much time, on average, the various approaches needed to complete one time step.

The centralized and distributed approach are compared for all the simulation scenarios; formulations 1, 2, and 3 for 2 and 8 houses. The performance of the decentralized approach is also calculated.

Figure 4.16 shows the average iteration time using formulation 1, for the centralized and distributed approaches. The centralized approach is around 4 times slower for 8 houses compared to 2 houses. The distributed approach is slightly faster for 8 houses compared to 2 houses. The distributed approach is slower overall, but the difference between the approaches is far larger for 2 houses than 8 houses.

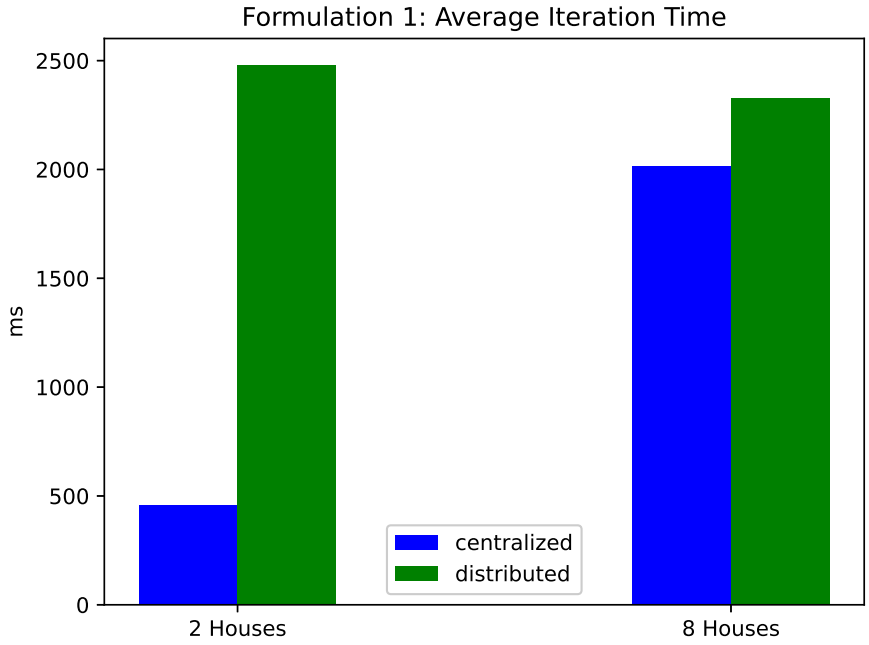

Figure 4.16: Average iteration time in milliseconds, formulation 1

Figure 4.17 shows the average iteration time using formulation 2, for the centralized and distributed approaches. The centralized approach is again around 4 times slower for 8 houses compared to 2 houses. The distributed approach is slightly slower for 8 houses compared to 2 houses. The distributed approach is much slower than the centralized approach for 2 houses but slightly faster for 8

houses.

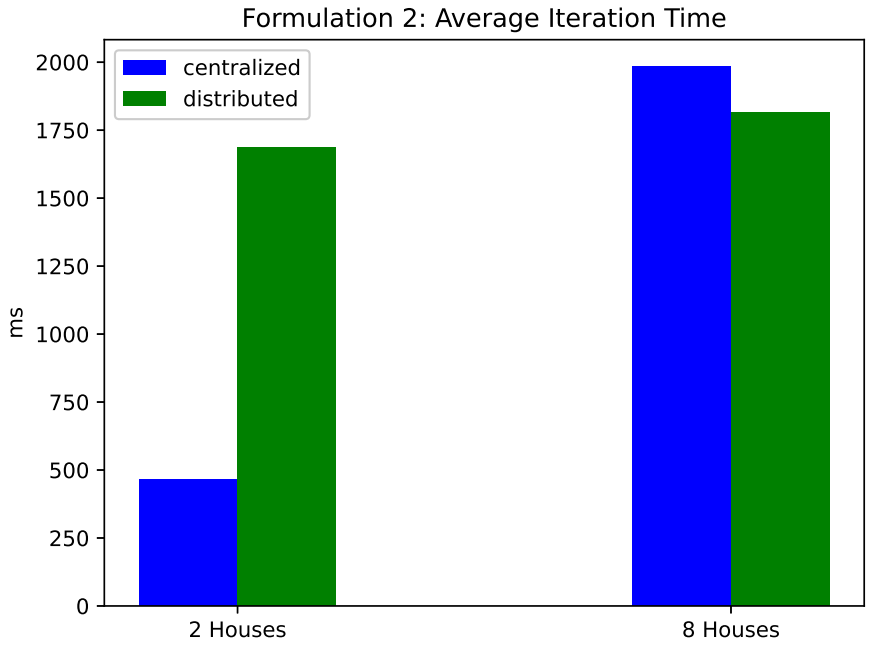

Figure 4.17: Average iteration time in milliseconds, formulation 2

Figure 4.18 shows the average iteration time using formulation 3, for the centralized and distributed approaches. The centralized approach is around 4 times slower for 8 houses compared to 2 houses. The centralized performance is very similar for all three approaches. The distributed approach is nearly identical for 8 houses compared to 2 houses. For both 2 and 8 houses, the distributed performance is far worse than the centralized performance, but whereas the centralized performance decreases as the number of houses increases, the distributed performance stays largely the same when the number of houses increases.

Figure 4.19 shows the average iteration time for the decentralized approach. The decentralized approach (see 3.1.2) has far better performance than the centralized and distributed approaches,

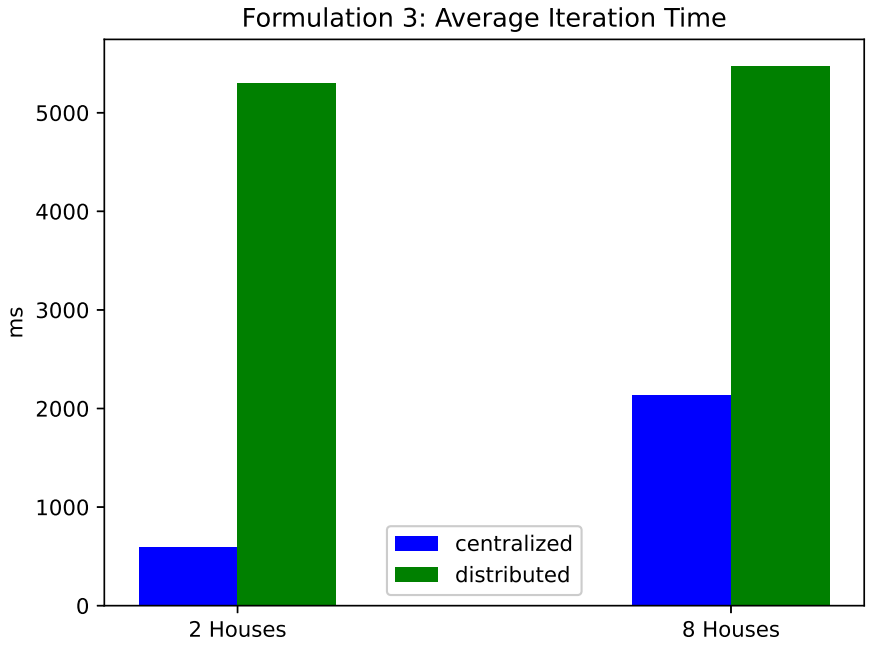

Figure 4.18: Average iteration time in milliseconds, formulation 3

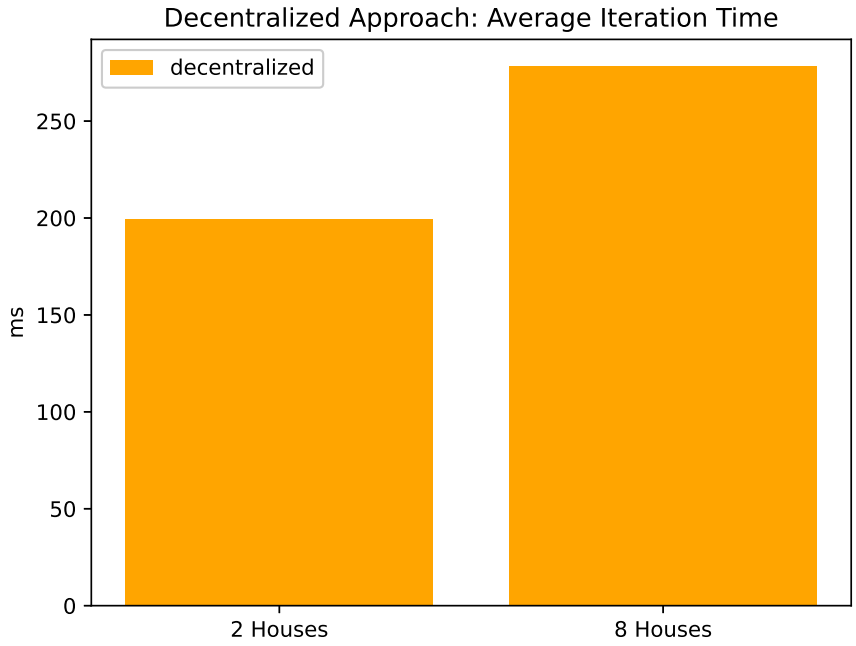

Figure 4.19: Average iteration time in milliseconds, decentralized approach

Chapter 5

Discussion

5.1 Limitations of Formulations

Section 3.1 mentions many of the simplifications in the house control formulation, chief of which is only considering one heat pump and one room for each house. In a real implementation, a house may have multiple heat pumps and several rooms that need heating. In addition to different external power use, larger houses also need different models. The current model considers the heat conduction between one room and the outdoors, as well as the heat generated by a heat pump. In a larger house, the heat convection between rooms would need to be considered; the temperature in one room might affect the temperature in another room. The effect of solar irradiation on room temperature is also not considered.

Additionally, the MPC approaches consider the predictions from the models to be perfect estimates. Though the discrepancy between model and real value may be small in some systems, it is very hard to model the temperatures of a house correctly. Consequently, in a real house control system, the decisions a controller makes will likely be based on estimates and predictions which are incorrect.

This, along with the various other simplifications, limits the validity of the results. For this reason, expanding the models and formulations used in the simulation would be crucial in any future work.

5.2 Residential Power Coalitions

Residential power coalitions would require multiple households/homes to agree to share a joint cost agreement. How exactly the cost is then divided can be determined according to the individual consumption of each household. Additionally, each home would have to agree to have their heating controlled by one of the control approaches mentioned in this thesis, instead of manual heating control.

For a centralized approach, this is a hard sell, since the users would have to accept that some centralized controller elsewhere, collects all their data and simply sends commands to their local heating control.

A decentralized approach is easier to justify to a house owner since the controller operates locally, but the lack of coordination between houses detracts from the justification for a power coalition in the first place. If the coalition cannot reduce their joint peak power, the electricity provider has less incentive to give them a better deal.

In a distributed approach, users would be able to consider the virtual 'price' of pushing on some peak power constraint, in this case represented by μ, and then decide if they want to use more power regardless. In this scenario, it would be practical to devise some scheme to associate a real cost with the virtual one, to indirectly incentivize coordination for peak power reduction.

5.3 Simulation Results

5.3.1 Hourly Peak Power Reduction

The new capacity based cost introduced in 1.3.2 is based on the hour with the highest consumption. As such, peaks in hourly consumption is a natural way to judge the peak power reduction abilities of the different approaches.

Simulating with two houses, figures 4.3 and 4.8 show the hourly power consumption for formulation 1, which has a linear momentary peak cost, and formulation 2, which has a quadratic momentary peak cost. Table 4.1 and table 4.11 show the corresponding maximum hourly power consumption. Formulation 1 performs better than formulation 2 in this regard, with a decrease of about 0.2 kWh compared to 0.1 kWh for formulation 2. It is unclear why formulation 1 performs better than formulation 2 in terms of maximum hourly consumption, as it does not so for the 8-house scenario when considering the same metric. In both cases, the results in the distributed approach are nearly identical to the centralized approach.

Figure 4.13 shows the hourly power consumption for formulation 3, which has a quadratic hourly peak cost. Table 4.21 shows the corresponding maximum hourly power consumption. The distributed approach is nearly identical to the centralized approach. The hourly peak approach had an improvement over the decentralized approach of 0.3 kWh for the highest hourly consumption. Slightly better peak power reduction is achieved compared to formulations 1 and 2, which consider momentary peak electricity.

When simulating with 8 houses, the differences between the formulations in terms of hourly peak reduction are larger. Figure 4.15 shows the hourly power consumption for formulation 3. The figure shows that the hourly consumption never increases much over 4 kWh. Compare that with 4.5 and 4.10 for the momentary peak formulations 1 and 2. Their hourly consumption is above 5 kWh from 17:00-21:00, which is similar to the results of the decentralized approach.

From table 4.6 and table 4.16, the max hourly power consumption for formulation 1 and 2, were about 0.5 kWh lower than the decentralized approach. The max hourly power consumption for formulation 3, found in table 4.26, was

around 1.4 kWh lower than the decentralized approach.

Formulation 3 performs well in terms of hourly peak reduction. This can be attributed to its direct penalization of the peak hourly consumption. In this regard, formulation 3 works as intended. Formulations 1 and 2 exhibit limited hourly peak reduction capabilities. Their penalization of momentary power peaks evidently does not correlate to significant reductions in peak hourly consumption.

5.3.2 Momentary Peak Power Reduction

When considering a larger amount of consumers together, in this case multiple houses in a power coalition, their combined momentary peaks in demand may have an impact on the performance of their local distribution grid. If power demand suddenly peaks, the supply of power must be increased to avoid brownouts or blackouts. As such these momentary peaks are another way to judge the peak power reduction abilities of the different approaches.

Simulating with two houses, figure 4.1 and 4.6 show the total power consumption for formulation 1 and formulation 2. Table 4.2 and table 4.12 show the corresponding maximum momentary power consumption. The results are nearly identical, with both centralized and distributed momentary peak formulations giving a decrease of about 15 % in maximum momentary power consumption compared to the decentralized approach.

On the other hand, figure 4.11 shows the total power consumption for formulation 3. Table 4.22 shows the corresponding maximum momentary power consumption. For this scenario, the results are very different from formulations 1 and 2, with the centralized and distributed approaches using 60 % more power than the decentralized approach.

The superior performance of the momentary peak formulations (formulation 1 and 2) compared to the hourly peak formulation (formulation 3) for reducing momentary power peaks are much more prevalent in the 8-house simulations. Figure 4.4 and figure 4.9 show the total power consumption between the 8 houses for formulation 1 and formulation 2. Comparing them to the results from formulation 3, figure 4.14, the momentary peak approaches have much lower peaks than formulation 3. From table 4.7, table 4.17, and table 4.27, we see that the highest momentary peak for formulation 3 (around 13.9 kW) is almost three times as large as for formulation 1 and 2 (both around 5.4 kW).

While the penalization of momentary power peaks in Formulations 1 and 2 does little to decrease hourly peaks, the momentary peaks are greatly reduced. This seems to indicate that the added peak state s does a good job of representing the peak in momentary power consumption since its penalization leads to a reduction in momentary peaks.

It is not clear why formulation 3 performs worse in terms of momentary peaks than the decentralized approach. The highest momentary peaks produced by formulation 3 appear around whole hours. Whole hours are the start and end points of the sets Θ_j, which define which hour momentary power use falls into. As such

the high momentary peaks produced by formulation 3 may be the byproduct of some intangible feature of the hourly peak formulation. If formulation 3 is to be considered in future work, its issues with producing high momentary peaks would need to be remedied.

5.3.3 Temperature Deviations and Power Consumption

Figure 4.2 shows the room temperature for formulation 1, and figure 4.7 for formulation 2, simulated on 2 houses. In both cases, the distributed approach has nearly identical results to the centralized approach. Formulation 1 produces a large dip in room temperature for both houses around 07:00-09:00, which is much less prevalent for formulation 2. This dip is not present in the decentralized approach. In this regard, formulations 1 and 2 yield more different results than their fairly similar power use.

Figure 4.12 shows the room temperature for formulation 3, simulated on 2 houses. The distributed and centralized approaches have very similar results, both very different from the decentralized approach. The deviations from the reference temperature for formulation 3 are problematically large, with room temperatures as high as 27 degrees. It is not clear why formulation 3 produces such high temperatures. It should incur a high cost for both temperature deviations and power use. Such results would not be acceptable in a real scenario.

Table 4.3 and table 4.8, the total temperature deviations using formulation 1, and table 4.13 and table 4.18, the total temperature deviations using formulation 2, are very similar. The similarity is likely due to temperature deviations being penalized the same for all formulations, and formulations 1 and 2 both penalizing momentary peak consumption (albeit with different costs). Granted, the temperature trajectories produced by formulations 1 and 2 are more different than the similar total temperature deviations indicate.

Considering the results, both room temperature trajectories and total temperature deviations, the decentralized approach performs better than all formulations in terms of tracking the room temperature reference. All formulations had the same cost associated with temperature deviations and power use, but the decentralized approach does not have any additional peak state formulations that might further penalize power use at the expense of temperature reference tracking.

Indeed, compared to the decentralized approach, all formulations had lower total power use. Table 4.5, 4.15, and 4.25 show the total power for all three formulations, simulating with 2 houses. Table 4.10, 4.20, and 4.30 show the total power for all three formulations, simulating with 8 houses. In all 6 cases, both centralized and distributed approaches had less total power use than the decentralized approach.

Despite this, the monetary cost associated with power consumption was higher for all formulations compared to the decentralized approach. Table 4.4, 4.14, and 4.24 show the total power cost for all three formulations, simulating with 2 houses. Table 4.9, 4.19, and 4.29 show the total power cost for all three formu-

lations, simulating with 8 houses. In all 6 cases, both centralized and distributed approaches had higher power use costs than the decentralized approach. While the peak state formulations may lead to reduced total power use, they also seem to make it harder for the controller to schedule power use based on spot prices, increasing total costs in that regard.

5.4 Comparing MPC Methods

5.4.1 Computational Performance

Section 4.4 shows the computational performance of the centralized and distributed MPC approaches for formulations 1, 2, 3, as well as the decentralized approach, during simulations.

Figure 4.16 shows the average iteration time for 2 and 8 houses, using formulation 1, a momentary peak state with a linear cost. The centralized approach has a longer average iteration time as the number of houses increases due to increased problem complexity. Interestingly, the average iteration time for the distributed approach is lower for 8 houses than for 2 houses. This is most likely due to different termination criteria for the dual decomposition (see f_{tol} and μ_{tol} in algorithm 1), which may cause it to terminate earlier and move to the next iteration, resulting in a lower average iteration time.

Figure 4.17 shows the average iteration time for 2 and 8 houses, using formulation 2, a momentary peak state with a quadratic cost. The 8-house simulation scenario using formulation 2 is the only case where the distributed approach is faster than the centralized approach. Formulation 2 has the best distributed performance of all three formulations. There may be multiple reasons for this. The distributed formulation 2 has a much simpler projection $P(\mu)$ than formulation 1, simply ensuring positivity of the dual variables $\boldsymbol{\mu}$, which means less overhead at each dual decomposition step.

Figure 4.18 shows the average iteration time for 2 and 8 houses, using formulation 3, an hourly peak state with a quadratic cost. Formulation 3 has a similar centralized performance to formulations 1 and 2. However, it has the worst distributed performance of all three formulations, around twice as slow. The likely culprit is the computation of the dual function gradient for each house $\nabla d_h(\boldsymbol{\mu})$. This computation is done at each dual decomposition step and requires the projection of each prediction step i to their corresponding hour j. This calculation is not complicated but has a lot of overhead. Fortunately, the average iteration time does not increase much from 2 houses to 8 houses.

Parallelism

The decentralized computational performance, shown in figure 4.19, is much better than the centralized and distributed performance. The average iteration time

also does not increase much from 2 to 8 houses. This is because the decentralized problems pertaining to each house can be run in parallel.

The distributed approach can also be run in parallel, but for the dual decomposition to converge, the parallel controllers may need to repeatedly solve the problem for the current prediction horizon. As such the average iteration time for the distributed approach is around 8-20 times higher than for the decentralized approach. That being said, a crucial property shared by the decentralized and distributed approaches is that the average iteration time does not change much as the number of houses increases. Although the simulations in this thesis only consider 2 and 8 houses, it is a fair assumption that the average iteration time would not increase by much if the number of houses was increased (see the review of [13] in section 2.1.2). This is a big advantage for a residential power coalition with many homes.

5.4.2 Distributed Convergence to Centralized Approach

The results in chapter 4 show that for all three formulations, there is little difference between the distributed and centralized approaches. This is the desired outcome, as the dualized problems the distributed approach solves, which are a result of the dual decomposition, are ideally identical to the primal problems which they are based upon. These primal problems are the ones that are solved by the centralized MPC approaches.

The implementation and performance of the distributed MPC approach is a central aspect of the thesis. Section 3.3 outlines several formulations for peak reduction, but only those where applying dual decomposition yielded a reasonable result work for the distributed approach. With that being said, it could be interesting to include and test formulations that would not work with a distributed MPC approach. In that case, only the centralized MPC approach would be used in simulations.

5.4.3 Privacy Concerns

One of the main motivations for choosing a distributed MPC approach over a centralized MPC approach is privacy. In the context of house heating, each subsystem represents not only an individual control scenario but a home whose inhabitants require comfort and privacy. Though these concerns are harder to price than power use, it is obvious that they must be considered when designing any control scheme that affects them. Privacy becomes an important design point if a centralized controller for multiple homes is considered.

In the centralized formulations, a single central controller needs to know not only the user preferences of all the houses but also their power consumption profiles, external and heat-pump based, as well as room temperatures. Though such information is not as sensitive as, say, location information collected from a smartphone, a power coalition between houses would require all its members to accept sharing this data with some central entity.

In a decentralized control approach, an in-house controller, physically accessible by its inhabitants, may be used to host both control algorithms and data collection. In this approach, information collected for use in the controller stays in the house it is installed in (technically this is not true when the IoT solutions used to collect data from smart devices are based on external databases, but in practice, it is nearly impossible for someone else to gain access to the information).

In the distributed approach, some information is shared from each house to coordinate the coupling constraints. In formulations 1 and 2, this information is the optimal power use for the current prediction horizon, $P(\boldsymbol{u}_h)$. If gathered over multiple time steps, this information gives the power use of a home with a resolution of 5 minutes. In formulation 3, the shared information is the optimal hourly power use for the current prediction horizon. $e_{j,h}$ for each hour j. This information can be used to create the power use of a house with a resolution of one hour. In both cases, the power use information does not distinguish between external and heat pump power use.

Earlier formulations for peak power reduction, since omitted from 3.3, included a formulation where the peak state itself was decomposed on a per house basis, in which case a separate peak state problem $d_s(\boldsymbol{\mu})$ was not needed. For that approach, the information needed to coordinate the coupling constraint was the localized peak state itself, which is less invasive than power use. The formulation was not used due to infeasibility in the resulting distributed formulation. Formulations where the users would not have to share their power use would be preferable, and should be considered in future work.

5.4.4 Improvements for Decentralized Approach

The decentralized approach, based on the formulation in 3.1.2, is what the three peak power reduction formulations described in 3.5 are compared against. Unlike those formulations, the decentralized formulation does not contain any elements that specifically target peak power use.

An alternative decentralized formulation could contain some of the peak power reduction formulations described in 3.3 but on a completely localized level. Decentralized formulations that penalize momentary or hourly power peaks might have provided a better benchmark to compare the centralized and distributed approaches against.

Chapter 6

Conclusion

From the simulations, formulations 1 and 2, which penalize momentary power peaks, work well for reducing momentary peaks but have a smaller impact on the highest hourly consumption. Formulation 3, which penalizes the highest hourly power consumption, is far better at reducing the highest hourly power consumption but has higher momentary peaks than the decentralized approach, which does not penalize peaks at all. Additionally, formulation 3 produced very high deviations from the reference room temperature. For all scenarios, the distributed MPC approach performed very similarly to the centralized MPC approach, which means the decomposed variants of the formulations were reasonable. Additionally, the distributed MPC approach scaled better than the centralized MPC approach in terms of computational performance when the number of houses was larger.

Future work should refine the formulations to improve their peak reduction capabilities, both momentary and hourly, while maintaining an acceptable room temperature and power costs. A more realistic temperature model and power model should be implemented. A real-life implementation could be considered if simulation results are promising enough. The idea of a residential power coalition should be investigated further.

Bibliography

[1] Energifaktanorge, *Electricity production*, Last accessed 12 November 2021, 2021. [Online]. Available: https://energifaktanorge.no/en/norsk-energiforsyning/kraftproduksjon/.

[2] U. Berardi, 'Building energy consumption in us, eu, and bric countries,' *Procedia Engineering*, vol. 118, pp. 128–136, 2015. DOI: https://doi.org/10.1016/j.proeng.2015.08.411. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877705815020664.

[3] Norwegian Water Resources and Energy Directorate, *Electricity consumption in norway towards 2030*, Last accessed 19 November 2021, 2021. [Online]. Available: https://www.nve.no/energy-consumption-and-efficiency/energy-consumption-in-norway/electricity-consumption-in-norway-towards-2030/.

[4] Statnett, 'Fleksibilitet i det nordiske kraftmarkedet,' 2018. [Online]. Available: https://www.statnett.no/globalassets/for-aktorer-i-kraftsystemet/planer-og-analyser/2018-Fleksibilitet-i-det-nordiske-kraftmarkedet-2018-2040.

[5] Linda Ørstavik Öberg, *Regjeringen har bestemt seg: Slik blir din nettleie fra 1. juli*, 2022. [Online]. Available: https://www.huseierne.no/nyheter/regjeringen-har-bestemt-seg-slik-blir-din-nettleie-fra-1.-juli-2022/.

[6] Tibber, *Ny nettleie del 1: Hvordan beregnes den?* 2022. [Online]. Available: https://tibber.com/no/magazine/power-hacks/ny-nettleie-del-1.

[7] P. H. Shaikh, N. B. M. Nor, P. Nallagownden, I. Elamvazuthi and T. Ibrahim, 'A review on optimized control systems for building energy and comfort management of smart sustainable buildings,' *Renewable and Sustainable Energy Reviews*, vol. 34, pp. 409–429, 2014, ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2014.03.027. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032114001889.

[8] D. Mayne, J. Rawlings, C. Rao and P. Scokaert, 'Constrained model predictive control: Stability and optimality,' *Automatica*, vol. 36, no. 6, pp. 789–814, 2000, ISSN: 0005-1098. DOI: https://doi.org/10.1016/S0005-

1098(99)00214-9. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0005109899002149.

[9] G. P. Henze, *Model predictive control for buildings: A quantum leap?* 2013.

[10] Y. Yao and D. K. Shekhar, 'State of the art review on model predictive control (mpc) in heating ventilation and air-conditioning (hvac) field,' *Building and Environment*, vol. 200, p. 107 952, 2021, ISSN: 0360-1323. DOI: https://doi.org/10.1016/j.buildenv.2021.107952. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360132321003565.

[11] R. Carli, G. Cavone, S. Ben Othman and M. Dotoli, 'Iot based architecture for model predictive control of hvac systems in smart buildings,' *Sensors*, vol. 20, no. 3, 2020, ISSN: 1424-8220. DOI: 10.3390/s20030781. [Online]. Available: https://www.mdpi.com/1424-8220/20/3/781.

[12] P. D. Christofides, R. Scattolini, D. Muñoz de la Peña and J. Liu, 'Distributed model predictive control: A tutorial review and future research directions,' *Computers & Chemical Engineering*, vol. 51, pp. 21–41, 2013, CPC VIII, ISSN: 0098-1354. DOI: https://doi.org/10.1016/j.compchemeng.2012.05.011. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0098135412001573.

[13] N. Lefebure, M. Khosravi, M. Hudobade Badyn, F. Bünning, J. Lygeros, C. Jones and R. S. Smith, 'Distributed model predictive control of buildings and energy hubs,' *Energy and Buildings*, vol. 259, p. 111 806, 2022, ISSN: 0378-7788. DOI: https://doi.org/10.1016/j.enbuild.2021.111806. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378778821010902.

[14] F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson and M. Morari, 'Reducing peak electricity demand in building climate control using real-time pricing and model predictive control,' in *49th IEEE Conference on Decision and Control (CDC)*, 2010, pp. 1927–1932. DOI: 10.1109/CDC.2010.5717458.

[15] A. Barbato, A. Capone, G. Carello, M. Delfanti, M. Merlo and A. Zaminga, 'House energy demand optimization in single and multi-user scenarios,' in *2011 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2011, pp. 345–350. DOI: 10.1109/SmartGridComm.2011.6102345.

[16] B. Foss and A. N. Heirung, *Merging optimization and control*, 2016.

[17] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings and M. Diehl, 'CasADi – A software framework for nonlinear optimization and optimal control,' *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019. DOI: 10.1007/s12532-018-0139-4.

[18] S. Boyd, L. Xiao, A. Mutapcic and J. Mattingley, *Notes on decomposition methods*, 2015.

[19] M. Bierlaire, *Optimization: Principles and Algorithms*, 2nd. Lausanne: EPFL Press, 2018, ISBN: 9782940222780.

[20] D. P. Palomar and Y. C. Eldar, *Convex Optimization in Signal Processing and Communications*. Cambridge University Press, 2009. DOI: 10.1017/CBO9780511804458.

[21] Nord Pool, *Nord pool*, 2022. [Online]. Available: https://www.nordpoolgroup.com/en/.

[22] The Norwegian Meteorological Institute, *Norwegian meteorological institute*, 2022. [Online]. Available: https://www.met.no/en.

[23] Tibber, *Tibber pulse*, 2022. [Online]. Available: https://tibber.com/no/store/produkt/pulse.

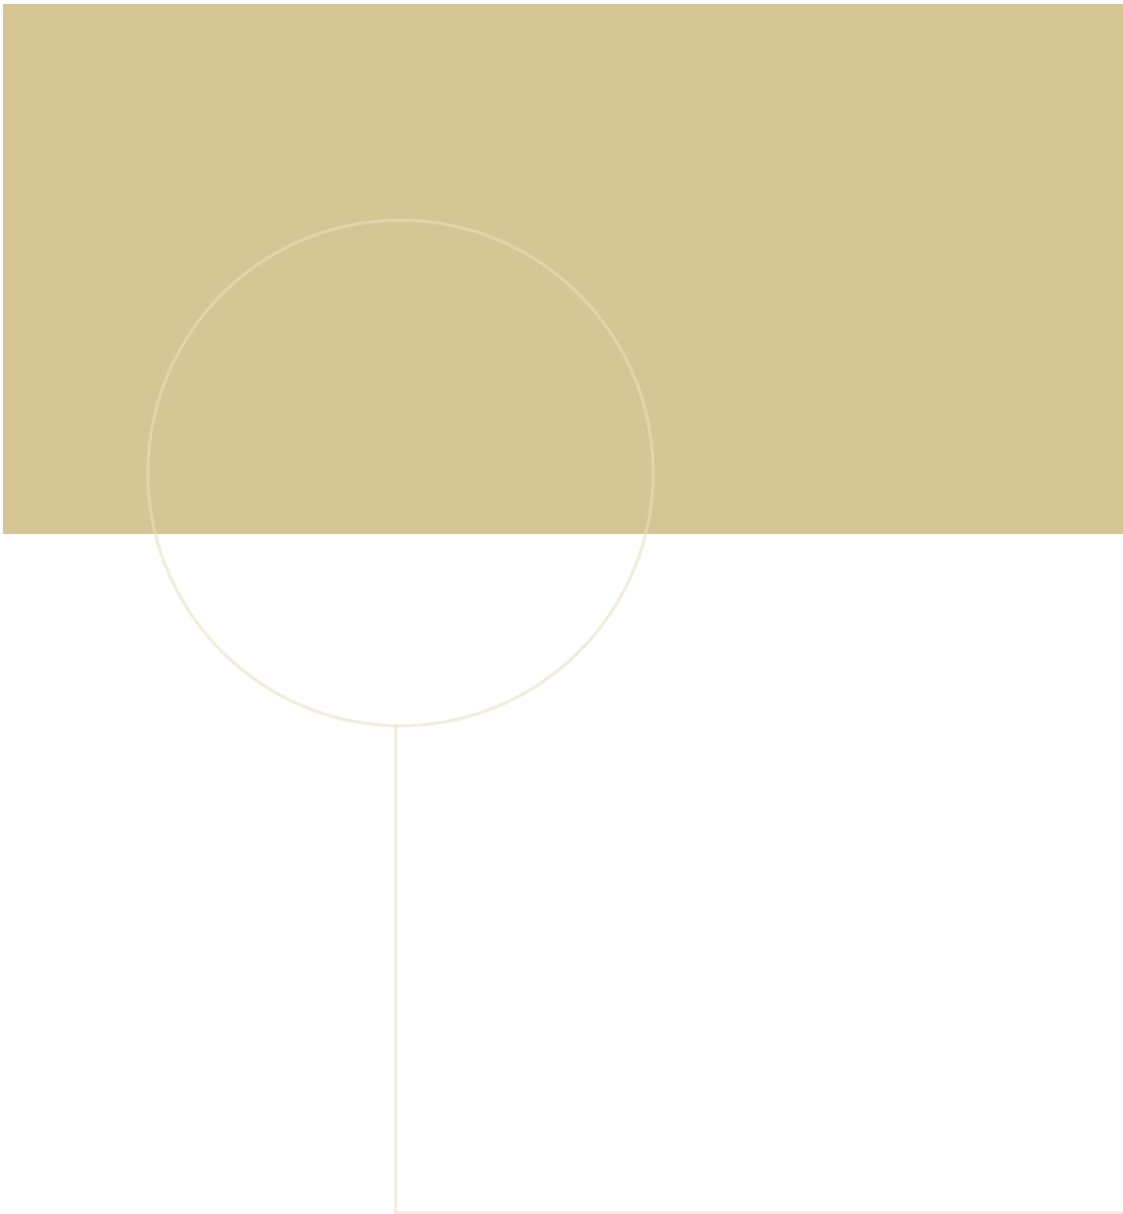

NTNU

Norwegian University of
Science and Technology