Distributed Heating Networks

Rijad Alisic

LUND
UNIVERSITY

Department of Automatic Control

MSc Thesis
TFRT-6063
ISSN 0280-5316

Department of Automatic Control
Lund University
Box 118
SE-221 00 LUND
Sweden

Printed in Sweden by Tryckeriet i E-huset
Lund 2018

Abstract

There is a shortage of models and analysis methods of fourth generation district heating networks, which are capable of both extracting and depositing heat energy to some thermal network grid. This thesis fills that gap by combining mathematical models of components into a network that is capable of sending heat energy between its nodes. Questions regarding good heating strategies for controlling the nodes were posed, and based on these, some simplifications were made to produce simpler systems to work with. Near-optimal distributed control strategies were produced and tested on simulations of the full mathematical models. For comparison, tuned P- and PI-controllers were also simulated on the full system. The results showed that the optimal controllers induce less oscillations and had less stationary error, however this caused larger control signals and more grid-wise interaction, causing neighboring users to feel more of the impact when a single node changed its operating point. This effect can be suppressed if a heating battery is connected to the system.

Acknowledgements

I am eternally grateful to my loving family, Malin, Fedra, Smail and Neo. You guys have always been there for me and pushed me to be the best that I could ever be. This would never have been possible without you!

I'd would like to express my deepest gratitude to my supervisor, Richard Pates, and my examiner, Anders Rantzer. Your inputs at our weekly meetings have been essential for keeping the momentum of this thesis up. There has not been a single moment during this thesis where there has been a shortage of ideas to try out.

Finally, I'd also like to thank all the people at the department of Automatic Control. You have always made me feel like a part of the crew for these past months. A special shout-out goes to everyone that was stationed in the Master Thesis room. You guys were the first I turned to when I ran into difficulties with this thesis.

Contents

1

Introduction

In 2015, the Swedish Government announced its Fossil Free Sweden initiative, with the specific goal of becoming a country free of fossil fuels before 2045. The initiative provides a direct link of discussions between industry and the government on how to provide legislation not only to ease this transition, but also to make sure that it is driven by economic growth. Currently, no scalable infrastructure replacement for fossil fuels exists in the market, due to the high fluctuations of ideal operating conditions that renewable sources need to have to be able to supply energy for the large demand. This is especially true in the Nordic countries, where direct sunlight or wind might be absent for often large, unknown periods of time.

1.1 Fourth Generation District Heating

In [Lund et al., 2014], the concept of a fourth generation district heating system is defined, where the renewables are not only optimally combined to provide energy for heating a district, but also cooling is integrated in the thermal grid (along with other necessary things, such as improved insulation). Here, a historical review is shown of the previous three generations of district heating, where the network heat is ever decreasing as progress is being made. A projection is made for the supply and return temperatures in the next generation heating system, lowering it to around 65 ^{0}C on the supply and roughly 20-50 ^{0}C on the return side.

Even if one were to successfully integrate renewable energy sources for the purpose of heating and cooling buildings, there is another development lurking just beyond the horizon. Following the large foothold Tesla has gained in the vehicle market, especially in Norway, several key people in the field are worried of how the massive load increase of the power networks will affect power stability. There are a few articles treating this for small, local scale systems, see [Clement-Nyns et al., 2010]. Due to this massive new energy demand, it is clear that changes solely on the producer side will not come close to solving the issue. Because of this, large efforts have been put to make the consumer change its behavior as well. For quite some time now, there have been incentives put forward on researchers and developers to

lower the energy either by increase of energy efficiency of products or by urging consumers to use less and to save more.

This way of approaching the problem looks towards decreasing the energy demands proactively and in some cases, especially in district heating situations, requires full knowledge of future consumer behavior to be effective, which quite often is impossible to determine. On top of that, it might also lead to a trade-off between energy efficiency and comfort in the sense that people might be required to lower their indoor temperatures during some periods of time, most likely right after they arrive to a room and right before they leave it. One may instead look towards acting retroactively and try to recapture residual heat of a space that is no longer used, which is seen as a major source of energy spillage to the environment.

A calculation put forth by EON shows that there might be an energy reduction of up to 70 percent in densely populated urban environments by recapturing residual heat and reusing it in other parts of a heating network [E.ON, 2018]. This changes the consumers role to become a quasi-prosumer, where they do not fully generate or buy their own heat energy, but rather resell all bought energy that has not been utilized, which would otherwise be lost to the environment if left in a room for quite some time. In addition, one may also connect cooling equipment to this exchange market, which often generates large thermal flows. This creates a space for exchanging heating energy and the lack of it (in the form of very cold water), so that refrigerators do not need to use all of their electricity to inefficiently cool down its interior. Instead, it can purchase cold medium from this exchange market in order to use it to partly cool down its interior.

This is, in essence, what EON aspires towards with their new ectogrid technology. By combining heating pumps and cooling machines to work in an integrated network of heating flows, where each building works as one of these quasi-prosumers, it is possible to send energy between different actors in the network. One may think of the network being one giant thermal battery, although one may also add components that more specifically work as batteries. Only when there is no more surplus energy, does additional energy become bought or generated from another local ectogrid network or centralized power plant. Furthermore, EON claims that the grid operates at a very low temperature, the same as the surrounding environment which not only implies that there are extremely low energy losses in due to heat dissipation, but also that it is possible to use cheaper, less insulating material in the pipes.

1.2 Purpose

Even though this idea of sending packets of heating energy is in itself a large boost for minimizing energy generation, one could greatly improve upon the efficiency of transfer between the different prosumers, minimizing losses and the electrical bill. While a small grid might be easy to control through information exchange with a

central plant that uses this information to dictate control signals, the true reduction in energy comes from the ability of having lots of prosumers in the same grid, in order to, for the most time, be able to have someone to send the energy to.

For large networks that communicate through some cloud service, reliability and bandwidth of that communication line becomes of vital importance. However, this also implies that there are a number of security risks that come with using such a service. Will the heating network still work in case of a power outage? What happens when the internet connection is lost? This could potentially leave several people stranded without any means of heating.

Finally, it is of interest to discuss some privacy concerns through the use of a central controller. In theory, the cloud could be used for spying on its users, creating profiles that could tell them things ranging from the desired temperature of a room, where in the building they currently are, to how many are there in the building and when the users leave their homes.

1.3 Problem Formulation

This thesis aims to look at fourth generation district heating networks, of which ectogrid is an example, in order to analyze the entire network dynamics and control the behavior, preferably with optimal controllers that only use local information. District heating networks, along with all of its components and use of liquid or gas as an energy transportation medium, are inherently nonlinear systems coupled with several delays of variable lengths. It is a key element of this thesis to find and motivate simplifications both to the model and to the control strategy.

2

Methods

Due to the fact that fourth generation district heating systems is a relatively new concept, where the main idea is that there should be a way to store thermal energy for later use, there is a lot of research effort that has yet to be put into this area. For example, there does not exist any publicly known model for a system of this type. In order to be able to analyze control schemes of a fourth generation district heating system, such a model needs to be created. It is on the basis of that model, that different control schemes will be created.

2.1 Modeling

The modeling of the network will take on a bottom-up approach, by creating mathematical models for various components that will be a part of the network. The mathematical description of the components will start from a physical description, where relevant effects will be incorporated into the model. This will lead to a model that is highly nonlinear with time-delays. Therefore, it is of interest to find ways to simplify the model. For some simplification, this means ordinary linearization around some operating point in order to synthesize a local controller. Other types of approximations will require that the signals in the system are limited in the frequency domain.

2.2 Benchmarking of Control Laws

Several types of control will be created in this thesis, depending on what behavior is desirable. For the most part, focus will lie on optimal controllers whose main job is to minimize certain non-desirable effects, such as energy loss. These control schemes will be synthesized from Pontryagin's maximum principle using various simplifications of the model. It is of major interest to generate decentralized controllers that only use local information to create a feedback controller, so this feature will also be emphasized.

For comparison, they will also be tested against local P- and PI-controllers that will be tuned so that the gain removes as much stationary error as possible and the integrator has a low gain while still managing to remove the stationary error within a reasonable time scale. It will be shown that if a these PI controllers are used as the primary control method, it is possible to minimize the energy usage by the proper choice of temperature in the medium flowing in the pipes of the thermal grid.

2.3 Simulations

The derived control schemes will then be tested in the full, non-simplified mathematical model. The various model components will be created in block functions using Simulink, a modeling and simulation environment tightly integrated with the MATLAB platform. The components will be put together as described in Section 3.4, and the control modules will be implemented in their own blocks.

Simulations will mostly focus around a set of operating points. For example, the environment temperature will always be set to 288K, and the goal of the prosumers will be to successfully control energy flows so that the two rooms will have a temperature of 293K and 283K, respectively. The mass flows in the network will be fixed and have the same values throughout all the simulations. The control laws that are derived will be created in their own separate blocks, so that it is easy to swap out only the controller part, minimizing the risk of benchmarking against entirely different mathematical models.

3

Modeling an Energy Flow Network

Current district heating networks operate from a central heating plant that not only has to make sure that the hot water needed to heat the buildings has the correct temperature (around 95 ^{0}C), but that the pressure across the network is maintained so that the water can flow with sufficient velocity even for the consumer that is the farthest away from them. As shown in Section 3.3, a fast mass flow implies better heating for all members in the grid. Additionally, the closer one is to the power plant, the better the heating due to larger pressure differences and temperatures. Another thing that has to be dealt with in the current heating network models is that Reynolds number is subject to change, as the enthalpy of the medium drops across the network, leading to a different mixture of gas bubbles and consequently, different mass flows and convex transfer coefficients. In this thermal heating system, these effects will be ignored, due to the lower temperatures of the water in the pipes and due to much lower mass flows.

3.1 Models of Fourth Generation District Heating Networks

According to the ectogrid website, in EONs version of the fourth generation heating system, users can tap into the grid in any way they want [E.ON, 2018]. Should they only want to extract heat from the grid or use it to keep the temperature low in a cold storage or both, they will be able to do so. Essentially this means that there needs to be different models for the different ways the consumer extracts or deposits heat into the network. On top of that, there are also several different economic options, such as the consumer directly pays for the net heating energy that they use, similar to contemporary practices, or they could become prosumers, where they own their own equipment and not only use it to extract heat from the grid, but also to generate their own heat from other sources, and then sell this excess heat when it is no longer

needed. The third option is to own a part of the grid itself, allowing for, for example, free energy transfer within the grid before interacting with the "outside world". This last option is more viable for organizations or cooperations owning multiple buildings in the same vicinity.

Another complication is that there are a couple of different ways to store the intermittent energy. One could store it directly in the network, by letting the entire network temperature rise during times of excess heat depositing into the network and cool during times of high heating demands, or one could dedicate components that in effect work as thermal batteries for the system, such as a large water storage somewhere in the network that mathematically works as an integrator, keeping the network temperature at a desired level. The former is a short term approach since it quite likely will move the average temperature of the grid away from the surrounding environment temperature, creating unnecessary losses to the environment. An example is that during the winter, there will be a surplus of heat demand and with all those buildings continuously extracting heat from the grid, its average temperature will become much lower than the environment, creating losses due to the large heat gradient. The large heat gradient between the grid and the spaces that are being heated will also cause the heat pumps to work extra hard, due to the fact that their efficiency is inversely proportional to the temperature difference, see Section 3.3. By introducing a thermal storage, it has the potential to bring the grid temperature back to an environment average, removing much of the losses.

3.2 Network Restrictions

This thesis focused on the quasi-prosumers case, where the user was assumed to be a consumer buying the energy required to heat its premises and then selling it back to the company that is controlling the grid, allowing them to move the very same energy to another part of the grid. In other words, the consumer was not explicitly able to generate its own heat and sell it to the company. Another simplification was that every user has precisely one space to heat or cool and that the heating/cooling happens everywhere in the room at the same time. An extension to the model where one for example splits up the air closest to the radiator and the room average is not difficult to do, but it was not considered here for simplicity.

The assumed network topology was that of a straight line, see Figure 3.1, where each user had access to two pipes, one flowing away from the heat storage and production facility and the other flowing towards it. This was done because most of the worlds current district heating systems have this topology, and that they are easy to combine into larger systems. However, one should keep in mind that other topologies are also possible and could, in principle, be better. Another simplification of the topology was that every building was assumed to be on the same height above sea level, in other words, the pressure would not change due to the different elevations of the various users.

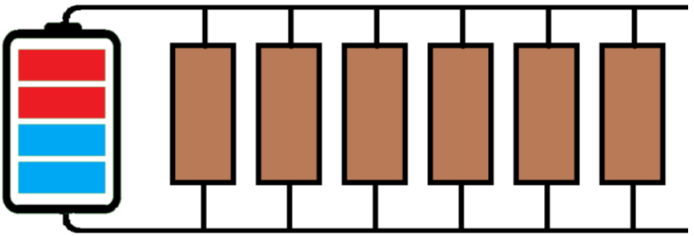

Figure 3.1 An example of how the fourth generation heating networks could look like. It is this topology of the network that this thesis will work with.

3.3 Components

As was briefly touched upon in the beginning of this chapter, the network was composed of several different components, which together made up the network. Some simplifications have already been made on the network side, but there were a few more whose purpose also was to simplify the network model, although they may not be completely obvious on the individual component level. Naturally, these simplifications imposed some restrictions on the operating procedure, often leading to some sub-optimal behavior or the loss of some effects.

Pipe Model

The most important component was the pipes, since they are the backbone of the network which connect the different users to each other. On the construction level, there are several degrees of freedom that are possible to choose from, which greatly impact network dynamics, efficiency and cost. This includes a choice of pipe material, pipe dimensions and the medium used as the heat transport vessel.

This model assumed that the temperature of the medium propagated throughout the pipe as if the flow was laminar, in other words, the cross-section temperature profile that entered a section of the pipe remained the same upon exit. Also, it was assumed that the fluid in the pipes was incompressible, such that the mass that entered a pipe had to immediately exit it on the other side. Considering Bernoulli's principle, the mass velocity, V, is governed by the pressure difference, $\Delta p = p_1 - p_2$, the inlet and outlet,

$$\frac{V_1^2}{2}\rho + \rho g h_1 + p_1 = \frac{V_2^2}{2}\rho + \rho g h_2 + p_2 \tag{3.1}$$

where ρ is the fluid density, g is the gravitational constant and h is the elevation above a reference plane.

Bernoulli's principle imposes that the pressure drops linearly across the pipe, and in extension, throughout the network. A large part of this pressure drop is due to the friction occurring between the walls of the pipe and the fluid. By assuming that the fluid density is constant throughout the pipe and using the continuity equation, it became easy to show that this interaction could be represented as heat generation of the fluid in the pipe. Initially, if one starts by considering the linear thermal expansion of a fluid, $\Delta V = \beta \Delta T$, then by simply taking the reciprocal of these relations, it was possible to obtain the density increase,

$$\rho_1 = \frac{m}{V_1} = \frac{m}{V_0(1+\beta(T_1 - T_0))} = \frac{\rho_0}{(1+\beta(T_1 - T_0))} \tag{3.2}$$

where β is the volumetric temperature coefficient of the fluid.

Similarly, one may consider the volumetric expansion due to a change in pressure and arrive at a similar result. In order to consider both the effects at the same time, one could just multiply the factors to each other, giving the relation,

$$\rho_1 = \rho_0 \frac{1}{(1+\beta(T_1 - T_0))} \frac{1}{(1 - \frac{p_1 - p_0}{E})} \tag{3.3}$$

where E is the bulk modulus of the fluid. By setting the density to be roughly the same and applying the Darcy-Weishbach equation for the major head loss, $\Delta P = \frac{F_f \dot{m}^2}{2\rho A^2 d_i} x$, gave the final form of the temperature increment in a infinitesimal section of the pipe,

$$\mathrm{d}(T_1 - T_0) = \frac{E}{\beta} \frac{1}{(E + \Delta P)^2} \frac{\mathrm{d}P}{\mathrm{d}x} \mathrm{d}x \approx \frac{1}{E\beta} \frac{F_f \dot{m}^2}{2\rho A^2 d_i} \mathrm{d}x \tag{3.4}$$

where F_f is the friction factor that is extracted from a Moody chart, $\dot{m}$ is the mass flow in the pipe, A is the cross-sectional area of the pipe and d_i is the internal diameter of the pipe.

Finally, convective heat transfer that occurred between the fluid and the outside environment through the pipe walls was considered. [Spakovszky, 2018] states that the heat transfer is in an infinitesimal section of a pipe is given by,

$$\mathrm{d}(T_1 - T_0) = \frac{h\pi d_i}{\dot{m} c_p} (T_e - T_1) \mathrm{d}x \tag{3.5}$$

where h is the convective heat transfer coefficient through the pipe and c_p is the specific heat capacity of the fluid. T_e is the temperature outside the pipe.

By combining this with equation (3.4), it was possible to derive a relation between the inlet temperature and the outlet temperature,

$$\mathrm{d}(T_1 - T_0) = \left(\frac{h\pi d_i}{\dot{m} c_p} (T_e - T_1) + \frac{1}{E\beta} \frac{F_f \dot{m}^2}{2\rho A^2 d_i} \right) \mathrm{d}x \Rightarrow \frac{T_e^* - T_1}{T_e^* - T_2} = \mathrm{e}^{-\frac{h\pi d_i}{\dot{m} c_p} x} \tag{3.6}$$

where the second term in equation (3.6), which arose from the Darcy-Weishbach pressure drop, has been embedded in T_e^*. If the mass flow is very small, which is common for the pipes entering and exiting the buildings in current heating systems, the Darcy-Weisbach pressure-drop may be ignored, prompting $T_e^* \approx T_e$. However, for large networks, it is an effect that one needs to take into consideration, since the plant has to generate a massive mass flow in order to provide heating for all its customers.

A final assumption of the pipe model was that it is completely a one-phase flow, meaning that there are was no gaseous flows within the pipes, which is particularly valid due to the low operating temperature and the assumption that the pumps generate a small pressure difference with respect to the atmosphere.

A major effect that pipes bring with them are time delays. This effect in the Laplace domain, could easily be represented as an exponential function multiplied with the gain of the pipe that was derived in equation (3.6). By, using the signals as temperature differences with T_e^*,

$$y(t) = \mathrm{e}^{-\frac{h\pi d_i}{\dot{m}c_p}x} u(t-t_0) \Rightarrow Y(s) = \mathrm{e}^{-\frac{h\pi d_i}{\dot{m}c_p}x} \mathrm{e}^{-st_0} U(s) \tag{3.7}$$

The exponential function, is however very difficult to work with, especially with a networked system with several delays present in different places throughout the network. To accurately describe a single time delay block, one would need an infinite amount of states. However, one could try to limit the amount of these states by the use of some approximation on the transfer function, however some additional constraints need to be imposed in order to maintain physical validity. Consider for example the following Padé approximation of the time delay

$$\mathrm{e}^{-st_0} \approx \frac{2-st_0}{2+st_0} = \frac{4}{2+st_0} - 1 \tag{3.8}$$

One sees that the approximation is composed of a simple first order system and a direct term. However, the direct term implies that what happens at the outlet is instantaneously affected by the inlet temperature, which is not a valid physical property. A comparison of the Bode diagram of these two expressions can be seen in Figure 3.2. One might then instead resort to a Taylor expansion of the form,

$$\mathrm{e}^{-st_0} \approx \frac{1}{1+st_0} \tag{3.9}$$

This approximation is valid under the condition that the input temperature spectral density is $st_0 < 1$, where the time it takes to traverse the pipe is directly proportional to the mass flow and the length of the pipe. This limitation can be eased by splitting the pipe up into multiple segments, allowing the traversal time to be lowered and increasing the model validity for higher frequencies of the input signal.

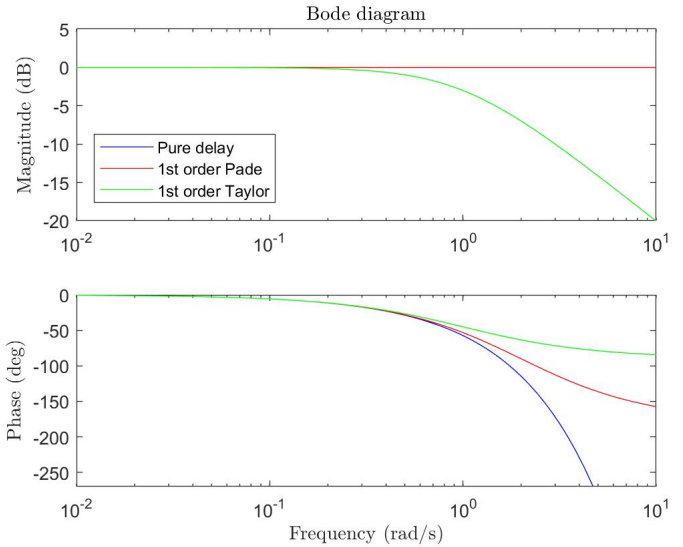

Figure 3.2 Bode diagram of two approximations to time delay, together with the pure delay.

$$e^{-st_0} \approx \left(\frac{1}{1 + s\frac{t_0}{k}} \right)^k \tag{3.10}$$

which is valid under the frequency range $s\frac{t_0}{k} < 1$.

Heat Exchangers

There are many different types of heat exchanger constructions differing by their directionality (counter- or parallel flow) and interaction area. See Figure 3.3 for an example of the type of heat exchanger that was used here, called the counter-flow heat exchanger. These heat exchangers are useful from a theoretical point of view, since the temperature relations between the two inlet and two outlet flows can be derived without any assumptions made on how it is constructed. Following the derivations found in [Spakovszky, 2018], it is possible to eliminate two temperatures flowing through the heat exchanger. In this thesis, they will almost exclusively be the outlet temperatures on both sides of the exchanger. So, given that the inlet flows are known, one of the outlet flows is given by the relationship

$$T_1^{in} - T_1^{out} = \mu(T_1^{in} - T_2^{in}) \tag{3.11}$$

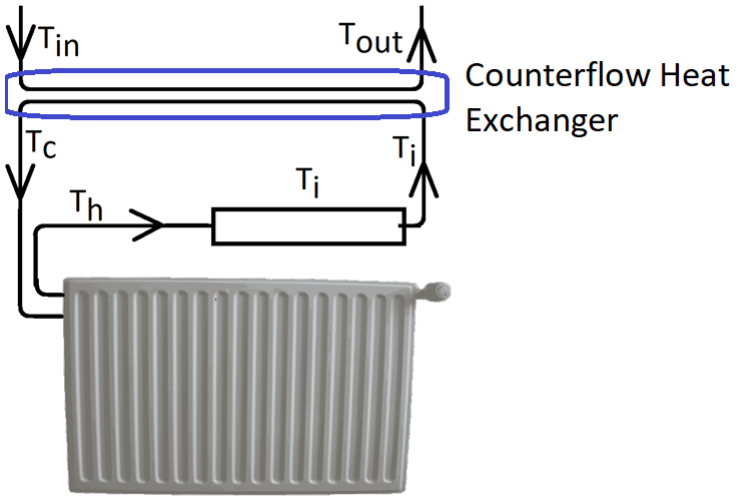

Figure 3.3 A schematic of a counter-flow heat exchanger that exchanges energy with an external pipe. The heated medium is then cooled down by interacting with the room through a radiator which also makes it a heat exchanger.

where μ is an efficiency factor that was assumed to be constant for a relatively high mass flow. It is governed by the following relation assuming that the same medium is used in both sides,

$$\mu = \frac{1 - \mathrm{e}^{-\alpha}}{1 - \frac{\dot{m}_{in}}{\dot{m}_{out}} \mathrm{e}^{-\alpha}} \tag{3.12}$$

$$\alpha = h\pi DL \left(\frac{1}{\dot{m}_{in}} - \frac{1}{\dot{m}_{out}} \right) \frac{1}{c_p} \tag{3.13}$$

where D and L are the diameter and the length of the interaction pipes

For rough values of the other physical constants, one can see that this expression is approximately constant for a range of mass flows that was used here, namely that the internal mass flow of each heat pump was smaller than the one out in the heating grid pipes, see Figure 3.4.

The equation for the second outlet temperature is similar, with only a change in the μ factor,

$$T_2^{out} - T_2^{in} = \mu'(T_1^{in} - T_2^{in}) \tag{3.14}$$

$$\mu' = \frac{\dot{m}_{in}}{\dot{m}_{out}} \mu \tag{3.15}$$

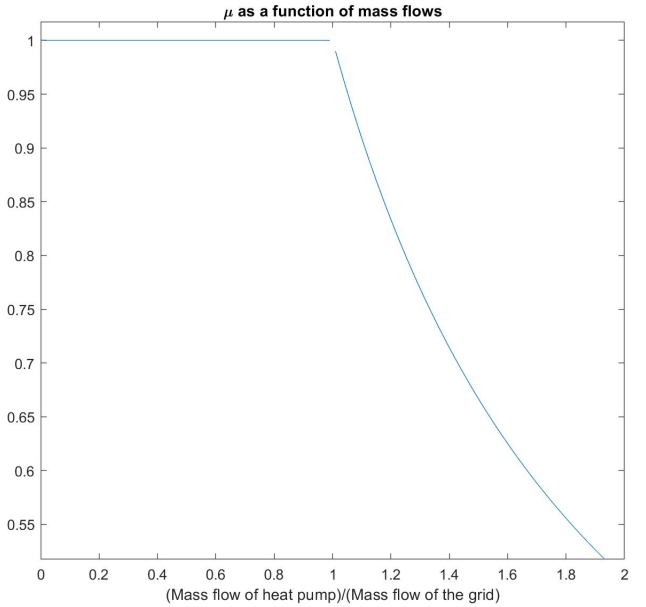

Figure 3.4 The graph shows in what ranges the approximation of μ being a constant is valid using some representative values of the heat physical parameters.

The heat exchanger was assumed to have no temporal delays when the medium flowed through it, since the system was relatively small in comparison with the time it took to change the temperature of the entire grid in comparison to the pipes, which were in lengths of tens of meters.

In third generation district heating, the radiators utilize the pressure difference between the hot and cold pipes to force the medium to run through the radiator, heating the room. This requires, however, that the temperature is very high on the hot side, unless some local pressure mechanism compresses the flow going through the radiator, increasing its temperature (which would imply that this is no longer a heat exchanger, but a heat pump). The outlet temperature, on the other hand, becomes very high (about 55-60 ^{0}C) which is not only not suitable to use for cooling machines, but also susceptible of heat loss to the environment.

The idea was to instead, say during heating operation, only draw heat from the hot side of the ectogrid and not touch the cold side. A way to model this was by having a heat exchanger extract heat from the hot side, and use this to heat up the room, see Figure 3.1, where the top pipe is the hotter one. By varying the mass flow of this internal heat exchanger, one also gained control of the degree of heating of the room, namely one could set the equilibrium exchange of heat energy. However,

it ass not possible to heat the room to any temperature higher than the temperature in the pipes.

Thus, only using this mechanism ass not scalable indefinitely, especially when it was claimed that the grid should be able to operate with the same low temperatures as the surrounding environment. The temperature of the hot side could very well drop to exactly the environment temperature, making all subsequent users connected to the hot side unable to further extract heat. One could make the same argument for the cooling mechanism by extracting cold water from the cooler pipes.

Heat Pump

The discussion in Section 3.3 showed that there is a need of being able to transport heat energy in the opposite direction of the temperature gradient, which is precisely what a heat pump does. Essentially, by running a Carnot engine in reverse, investing work instead of harvesting it, one is able to cause the heat to flow in the opposite direction. The property of the theoretical framework for how much heat is being extracted in the Carnot engine, and in particular the heat pump, does not depend on the precise construction, making it very easy to use for calculations. The central term in a heat pump is the Coefficient of Performance (COP), defined as the useful work W divided by the power input, P,

$$COP = \frac{W}{P} \tag{3.16}$$

According to the second law law of thermodynamics, total entropy can never be reduced in an isolated system. Using this, one is able to extract the largest possible COP,

$$COP \leq \frac{T_h}{T_h - T_c} \tag{3.17}$$

where subscript h denotes the hotter side and c denotes the cooler.

However, this maximum COP is only reached when the process is run very slowly, too slow for any practical uses. Because of this, in normal operating procedures, the physicist approximates any losses to the system by an efficiency factor, ν,

$$COP = \nu \frac{T_h}{T_h - T_c} \tag{3.18}$$

In general however, this factor is not only also dependent on the temperatures on both sides, but also of the frequency the pump engine is currently working with, which is why on most heat pumps, the manufacturers give a polynomial fit of the COP as a function of these three parameters. A simplification was made here, namely that the frequency ass eliminated as done in [Verhelst et al., 2012], keeping the definition of COP to be the ratio between extracted heat and the power input solely dependent on the temperatures on each side.

$$COP = a_0 + a_1 T_h + a_2 T_c + a_3 T_h^2 + a_4 T_c^2 + a_5 T_h T_c \tag{3.19}$$

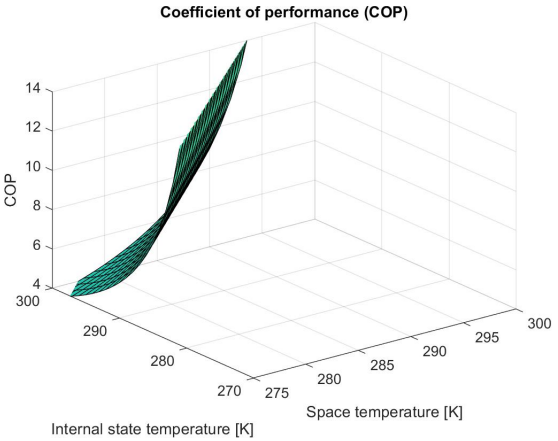

Figure 3.5 The diagram shows how the Carnot COP depends on the temperature of either side of its heat baths.

A side by side comparison of these two COP can be seen in Figures 3.5 and 3.6, where a ν of 0.3 has been used and the coefficients for (3.19) have been taken from [Verhelst et al., 2012]. One can see that there are some differences between these two images mainly due to the efficiency difference as the total temperature of both sides is raised, but the overall shape is roughly the same. According to [Verhelst et al., 2012], the largest error in the area of (3.19) is a few percent due to the elimination of the frequency terms. In this model, equation (3.18) was used as the COP, with the efficiency factor chosen so that it matches the size of equation (3.19) as closely as possible in the operating range.

If one were to blindly connect the heat pump to a radiator and a pipe in the thermal grid, some natural questions would arise. Such as what is the exact temperature of the pipe needed to calculate the COP. A simple answer such as using the inlet temperature or calculating the COP as some average between the inlet and outlet temperature, however, this does not truly capture the dynamics of the system. Instead, it was modelled as small fluid storage tank with continuous in- and outflows. The in- and outflows were connected to a heat exchanger, which extracted or deposited energy into the grid in a manner that was proportional to the surface area of interaction. Under the assumption that the temperature in this tank is raised (lowered) everywhere simultaneously due to the hotter (colder) medium constantly flowing in, it became easier to calculate the temperature in the tank, while at the same time providing a good approximation of the continuous temperature drop in the pipe that was part of the thermal grid. The heat pump then extracted (deposited) the heat energy to (from) the room. See Figure 3.7 for a overview of how this system model looks like.

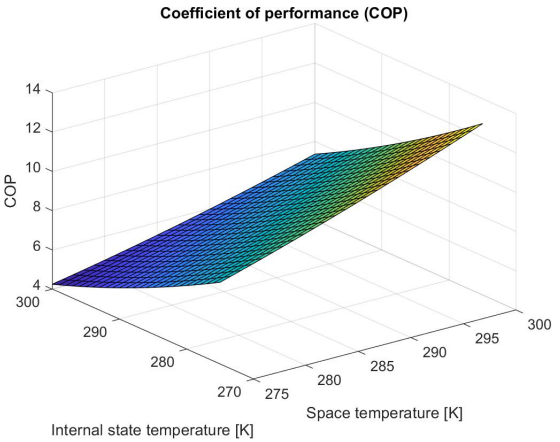

Figure 3.6 The diagram shows how a least squared polynomial fit of the COP depends on the temperature of either side of its heat baths.

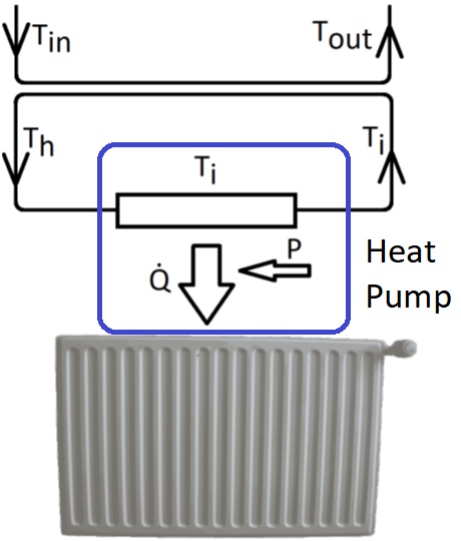

Figure 3.7 The schematic shows the connection between the heat pump, radiator and the internal state, which in turn interacts through a heat exchanger with the heating grid pipes.

Battery and Production Plant

A dynamical model of this component mash-up was not used explicitly in this thesis. However, due to the assumption that the plant was very large and could control large amounts of flows and heat them to very high temperatures, the plant was assumed to simply have some gain that increased the temperatures. This was used in the dynamical derivations of controllers. For the P- and PI-controllers, it was assumed that the plant could set any temperature of the water that is sent into the grid. It could do this perhaps by preheating vast amounts of energy and storing it in a battery or by simply collecting all the residual energy created by the heating grid, storing it and then using it when needed. A similar model is used in [Scholten et al., 2017].

3.4 Creating the States - a Simple Network

Now that the components have been modelled, it was time to combine them into a grid. Since the components had been modelled separately from each other, and from the ground up, it was possible to combine them into a general network. In this case, the network topology was assumed to be linear, with each building having a pipe to both the hot and the cold side. Each user that connected to the thermal grid was assumed to precisely have one space they wished to control the temperature of, be it a room, boiler or freezer. The state equations of the space temperature also assumed that there were some heat losses with the ambient environment, proportional to the surface area separating the two. This, together with the increase in temperature due to a heat exchanger or heat pump, was governed by the following equation with $X_k = T_k - T_e$, for some label k,

$$\dot{X}_r = -aX_r + g(T_r, T_i, T_e)P \tag{3.20}$$

where P is the input power variable, g is a nonlinear function describing the heat transport between the space with temperature T_r and the internal state, T_i and a is a collection of a number of physical constants.

The second term, g, increased the temperature due to a heat pump or a heat exchanger as it transported energy between the internal state to the space. What g that was used depended on the temperatures of the room, environment and the internal tank state. This was due to the way the different machines alter the temperature in the state after energy was supplied to them. A heat pump pushes the temperatures in opposite directions, causing the heat to flow in opposite directions while the heat exchanger pushed them closer to each other as more power was supplied to it. Under the assumption that heat loss to the environment was very minimal, it was possible to reach almost any state with the proper combination of heat machine, power and grid temperature. An example of a situation where a specific desired temperature was not possible to reach is shown in Figure 3.8. Here it was not possible to drive the space temperature to the same temperature as the grid pipe, since an infinite flow (infinite input to the machine pumping the water), was required to achieve it.

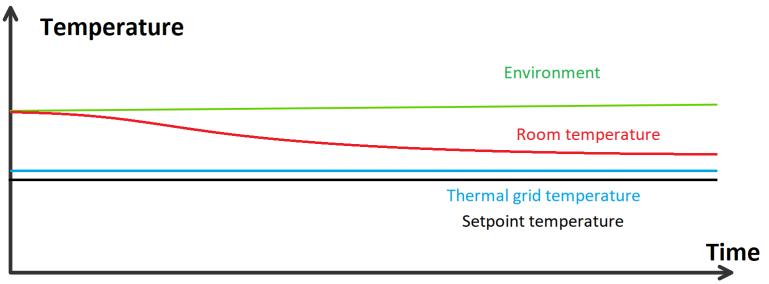

Figure 3.8 A special example of what could go wrong when a room changes it desired setpoint to be on the "other side" of the grid temperature. In this case, the room temperature will only ever be between the environment temperature and the thermal grid temperature. In order to make the room temperature cross to the other side of the thermal temperature, a different cooling machine than the ones mentioned here would need to be used - one that is not as efficient as a heat pump or heat exchanger. This thesis will ignore this case, since a simple solution could be to get as close as possible, before temporarily switching to a less efficient way of removing the final heat before the heat pump kicks in.

Therefore, a way to get around this problem could be by changing the grid temperature. The different machines worked, in some sense, only effectively under their operating conditions.

It was quite natural that as the heat machine deposited or extracted heat from the internal state, its temperature, T_i, should have changed accordingly. However, not all of the heat that was delivered to the space was always extracted from the internal state, at least not when a heat pump was used. The electric work that was generated by the pump eventually got converted into heat and was supplied to the hotter side. This meant that the opposite side was not cooled as much as the other side was heated. Using the COP notation, one could mathematically write it as

$$COP_c^{max} = 1 - COP_h^{max} \leq \frac{T_c}{T_h - T_c} \tag{3.21}$$

Similarly to the heating case, the efficiency of the engines was incorporated into the expression by multiplying the COP with a factor ν'. Here, it was assumed that $\nu' = \nu$, since the same pump that was extracting the heat from the cold state, was used for the delivering it to the hot state.

The internal state was used as an intermediate state, where it was possible to control how much energy one would like to extract from the grid. Its connection to the thermal grid was modelled by a counter flow heat exchanger, whose flows got heated or cooled by extracting or depositing heat to the grid and in turn, the grid changed its temperature based on the amount of surface area interaction. Referring to Figure 3.7, T_{in} was defined as the temperature of a pipe in the thermal grid that

was going through the building and into the heat exchanger. On the other side of the heat exchanger, a pipe with water flowing from the internal state entered, T_i. At the respective end of the two sides, the two pipes have changed their temperatures, where T_h flowed into the internal state and T_{out}, was the temperature of the water in the thermal grid, which flowed to the next unit that was connected to the grid. Using the equation (3.11) given in Section 3.3 and assuming that the heat pump was used as the heat transport between internal and space state, the time evolution of the internal state temperature was obtained as,

$$\dot{X}_i = \dot{m}_i c(X_{in} - X_i) \mp d\nu \frac{T_i}{X_r - X_i} P \tag{3.22}$$

The first term in equation (3.22) described the thermal flow between the grid and the internal state with $T_{in} = X_{in} + T_e$ being the temperature of the pipe in the thermal grid. With the mass flow factor, $\dot{m}$, it was possible to determine the degree of interaction with the grid. The second term described the energy interaction between the space and internal state using a heat pump. Also, precisely as was before, c and d are collections of physical constants and efficiency factors.

Referring back to equation (3.14) derived in Section 3.3, one could use it to write the output temperature of the fluid going on the next user in the grid as,

$$Y = X_{out} = \mu X_i + (1 - \mu) X_{in} \tag{3.23}$$

It was now possible to write a good representation of the heat flows between two states connected by a heat pump and whether one needed it for heating or cooling, it was possible to just switch the sign of the COP dependent on what was used.

$$\dot{X}_r = -aX_r \pm b\nu \frac{T_r}{X_r - X_i} P \tag{3.24}$$

One could also replace the heat pump for a heat exchanger in some dynamic cases, such as when there was a of need cooling/heating a room in order to quickly save energy without investing too much work.

With the model including how the states affect the thermal grid, it was time to turn the attention to the grid itself. The inlet to each user could be considered to come from an outlet of a different user, minus the effects that happened in the pipes as discussed in Section 3.3. As shown in this section, it was possible to treat the inlet temperature as one of the inputs to a heat exchanger, whereas the other input to the heat exchanger was the temperature of the internal state, T_i, see Figure 3.7. Given each subset of nodes consisting of a space to be heated and a pipe supplying the heat to the space, it was written in the following coupled state differential equation

$$\dot{X}_r = -aX_r \pm b\nu \frac{T_r}{X_r - X_i} P \tag{3.25}$$

$$\dot{X}_i = \dot{m}_i c(X_{in} - X_i) \mp d\nu \frac{T_i}{X_r - X_i} P \tag{3.26}$$

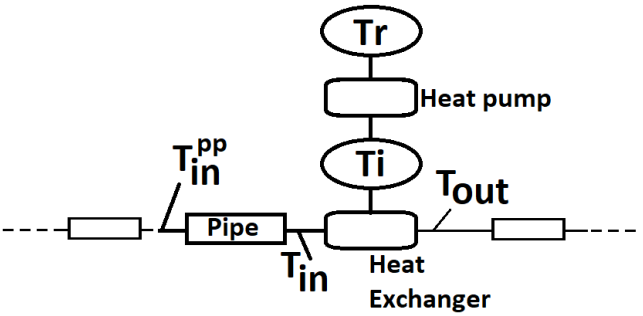

Figure 3.9 A schematic layout of the prosumer (the vertical nodes above the horizontal pipes). Here one can see how the states are connected with each other while using a heat pump to extract energy from the internal state.

$$X_{in}(t) = X_{in}^{pp}(t - t_0)\mathrm{e}^{-\frac{h\pi d_i}{\dot{m}c_p}L} \tag{3.27}$$

$$Y = \mu X_i(t) + (1-\mu)X_{in}(t) \tag{3.28}$$

where X_{in}^{pp} is the fluid temperature immediately after it has left the previous house, μ is the heat exchanger efficiency factor, L is the pipe length and t_0 is the time it takes for the fluid to travel through the pipe. A schematic view of how the nodes were connected to each other can be seen in Figure 3.9.

Example: Single Prosumer

Consider a single building that is connected to a fourth generation heating grid according to our model. Assume that the building temperature initially is at a bit above 288K and the temperature in the grid is 288K, see Figure 3.10. After some time, the people in that building find it to be too cold, so they increase the power input to the heat pump which is extracting heat from the grid. One can then see the effect of that change at $3.8 * 10^4$ seconds in Figure 3.10. In that Figure, it is possible to see how the difference between the internal state and the room temperature becomes larger, this is because the heat pump has started to extract more energy from the internal state compared to the rate of which the internal state extracts energy from the thermal grid. That more energy is extracted from the grid could be seen from the outlet temperature Figure 3.11, where it has dropped about 0.4K.

Full Model Description

Since a linear topology was considered with a simple battery/producer being a part of it, it was assumed that the pressure pumps were located there, creating a pressure difference pushing the liquid to flow throughout the network. However, the actual

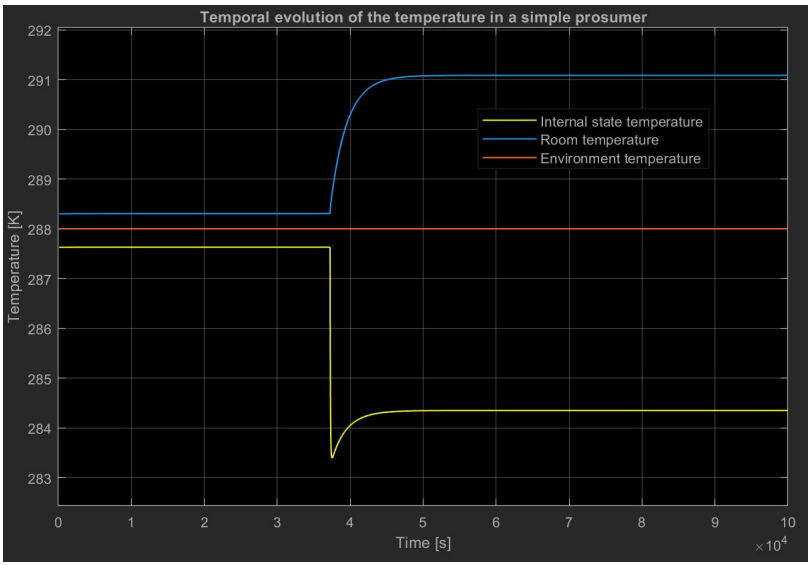

Figure 3.10 The image shows the temperature evolution of a simple house that is connected to the thermal network.

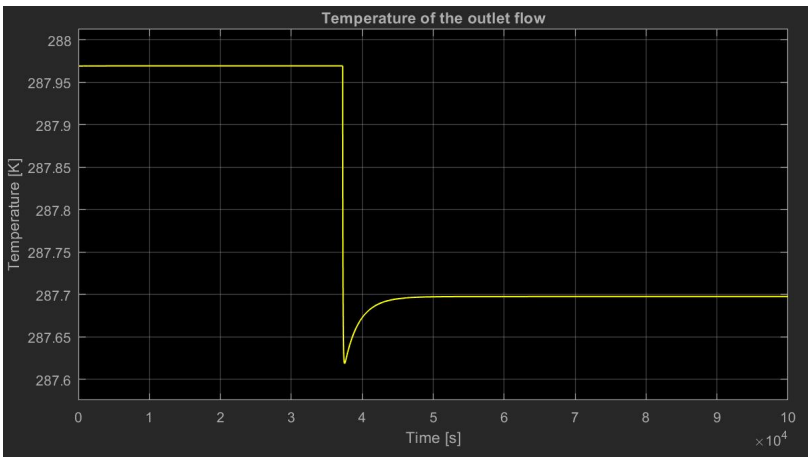

Figure 3.11 The graph shows the outlet return temperature of a simple house. The temperature should be lower than the inlet temperature, 288K, indicating that energy has been extracted from the water flowing through the building.

purpose of these pumps was to balance flows in the subsequent net and also to partly regulate the temperature of the water in the pipes. It was formulated as the control input that defines the amount of energy that was added and extracted from the outside "world". The input flow was defined from the individual flows in the network, where by individual flows it was meant that the buildings themselves in Figure 3.1 determined the mass flows along the vertical pipes. The job of the plant became, more precisely, to balance the flows in the first connection nodes. Another consequence of the network that was used is that the flow exiting the network must be equal the flow that enters it, since a vital assumption was that the fluid used in the pipes is incompressible.

The necessity of external energy can be explained in the following way; suppose that the network is stationary at some operating point and that there is no external deposition or extraction of energy thus the buildings in the network successfully cancel out each others energy usage. Suppose now that one of the buildings wishes to extract more and more energy from the grid, due to some change in the reference temperature or perhaps that someone opens a couple of windows in that building to get some fresh air. In order to maintain the room temperature, more energy will need to be extracted from the grid, thus lowering the temperature in the grid. Difficulties arise when the temperature becomes so low that the medium in the pipes surpasses its own melting point. Although one could alter the pressure of the medium to lower the melting point, however the condensation point or, at the very least, the density of the medium is also altered, creating a more complex phenomena that would require more complex models.

The network model that was examined in this thesis however, is that of a two-building system, making it possible to capture certain critical aspects that could be of interest when implementing control of a larger system. The battery/production facility was assumed to be able to immediately change the temperature of their outlet, which physically can be implemented with a correct mixing of the medium exiting the network and the stored medium of the battery/production facility. By setting up the network in this way, one can draw some parallels with current literature of battery/production plant systems, where in for example [Scholten et al., 2017], a control scheme is synthesized to combine the battery and production facility to optimize over some time dependent restraints on the battery that could be linked to operating cost or safety margins for unforeseen effects. Although that result would have to be generalized to incorporate time delays in the system. In this model, minimizing the difference between the inlet and outlet to the producer/battery block was a part of the objective, in order to not blow of all of the stored energy immediately or to minimize the amount of energy that was needed to be produced at the production facility.

The first prosumer in the network, immediately to the right of the battery in Figure 3.12, in this model was assumed to have the objective of having a temperature higher than the environment temperature, it wants to heat it's space. It was therefore be called the heated building. In order to achieve this, it took the medium

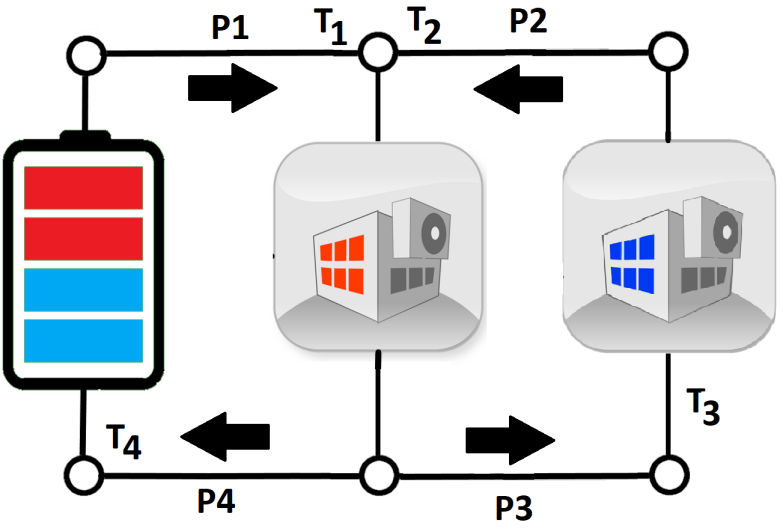

Figure 3.12 A scheme of the full model used that is used in order to derive decentralized controllers. Notice that the building that is extracting heat, here called the "heated building" is marked in red while the cooled building, marked in blue, deposits heat energy into the grid.

from the hot side and extracted the heat energy from it. The cooled medium was then deposited on the cold side of the grid to be used either by a cooling machine or to be heated up again by the battery/producer. The second prosumer was thus assumed to have an objective that is lower than the environment temperature, the objective is to cool its premises and will thus be called the cooled building. It thus took the medium on the cold side and deposited energy to it before it released it to the hot side. Under some circumstances it was possible to completely cut of the battery/producer and let the energy flow between the two prosumers. It was, however, not always possible to be able to cancel out each others energy needs and for the most part, it was very costly to maintain large temperature differences between the buildings.

Using the equations derived in sections leading up to 3.4, with reference to Figure 3.12, it was possible to build the full model. From the first input of the network, i.e., the inlet temperature, it was possible to create a state, x_1, of the temperature that exited the pipe labeled as P_1 in Figure 3.12,

$$x_1 = u_0(t-t_0)\mathrm{e}^{-\beta_1} \tag{3.29}$$

Similarly, it was possible to create the states x_2, x_3 and x_4 using the exit temperature of pipes P_2, P_3 and P_4, respectively, using their physical constants in the equations. However, the input to these pipes were simply the outputs of other components as seen in in Figure 3.12. For all three remaining pipes, their inputs were defined as the exit flows of the prosumers. As seen in Section 3.4, for state x_2, it was straight forward to calculate the input using the equations there

$$x_2 = (x_i^c(t-t_0)\mu + (1-\mu)x_3(t-t_0))\mathrm{e}^{-\beta_1} \tag{3.30}$$

$$\dot{x}_i^c = \dot{m}_i^c c^c (x_3 - x_i^c) \mp d^c v^c \frac{T_i^c}{x_r^c - x_i^c} P^c \tag{3.31}$$

$$\dot{x}_r^c = -a^c x_r^c \pm b^c v^c \frac{T_r^c}{x_r^c - x_i^c} P^c \tag{3.32}$$

Special attention was given to the intersection of pipe P_1 and P_2, it is clear that some sort of mixing occurs here between the medium exiting both pipes. It was assumed that this mixing occurs quickly, and that the resulting temperature was a weighted average of the temperatures, where the weights are the medium mass flow in the two respective pipes. This average was then used as an input to the heated prosumer which was governed by the equations derived in Section 3.4.

$$x_{in} = \frac{\dot{m}_0 x_1 + \dot{m}_2 x_2}{\dot{m}_0 + \dot{m}_2} \tag{3.33}$$

$$\dot{x}_i^h = \dot{m}_i^h c^h (x_{in} - x_i^h) \mp d^h v^h \frac{T_i^h}{x_r^h - x_i^h} P^h \tag{3.34}$$

$$\dot{x}_r^h = -a^h x_r^h \pm b^h v^h \frac{T_r^h}{x_r^h - x_i^h} P^h \tag{3.35}$$

The output of this system was split into two pipes, however the temperature that entered these pipes were the same and the only thing one had to keep track of was that the mass flows were correct which was obtained by simple solution of coupled, linear, continuity equations, which in the electric analog is the same as Kirchoffs current law. That is, the final two states became

$$x_3 = (x_i^h(t-t_0)\mu + (1-\mu)x_{in}(t-t_0))\mathrm{e}^{-\beta_4} \tag{3.36}$$

$$x_4 = (x_i^h(t-t_0)\mu + (1-\mu)x_{in}(t-t_0))\mathrm{e}^{-\beta_1} \tag{3.37}$$

Collecting all the equations (3.29)-(3.37) gave the full state space model of this network which could be written in the following compact form,

$$z = A^t x(t - t_0) + B^t u(t - t_0) \tag{3.38}$$
$$\dot{y} = A' x(t) + B'(y) u(t) \tag{3.39}$$

where A^t, B^t and A' are matrices with scalar entries. $B'(y)$ is a matrix which depends on y and

$$z = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} \tag{3.40}$$

$$y = \begin{bmatrix} x_r^h \\ x_i^h \\ x_r^c \\ x_i^c \end{bmatrix} \tag{3.41}$$

are vectors of the states that only depend on the time-delayed state variables and non-time delayed state variables, respectively.

3.5 Verification of the Model

Now that the entire model had been obtained, one could ask whether the network acts in a desirable way - such that, at least, a part of the energy was transported between the actors in the network. First, the water that was injected from the battery to the network have a temperature that was the same as the surrounding environment, which in this thesis will be assumed to be 288K, and then the input to the heat pump of both buildings was set so that they reach the desired room temperature. This was done by setting the equations (3.31), (3.32), (3.34) and (3.35) to zero, calculating the heat pump inputs as a function of the incoming water temperature, and then using the remaining equations of (3.38) to calculate the water temperature throughout the grid.

In Figure 3.13, one can see that the temperature of the cooled building started of at 287K and the temperature of the heated building during the same time period is 293K, see Figure 3.14. After some time, the desired temperature of the cooled building changed to 277K and the appropriate control signal was applied. It was expected that the cool building would want to deposit more heat to the grid and consequently increase the grid temperature. This could be seen by looking at the internal state temperature in Figure 3.14, where the temperature went up.

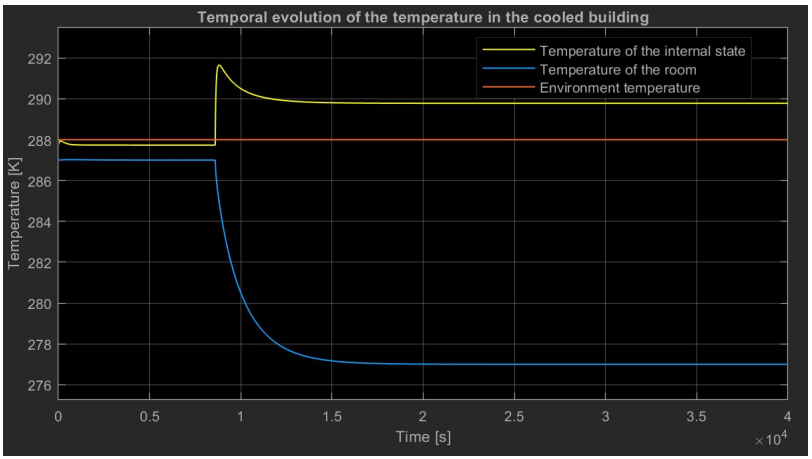

Figure 3.13 The graph shows the temperatures in the cooled building. When the cool building lowers its temperature, the heat pump transfers more energy to the thermal grid. This can be seen through the internal temperature, which has increased its temperature.

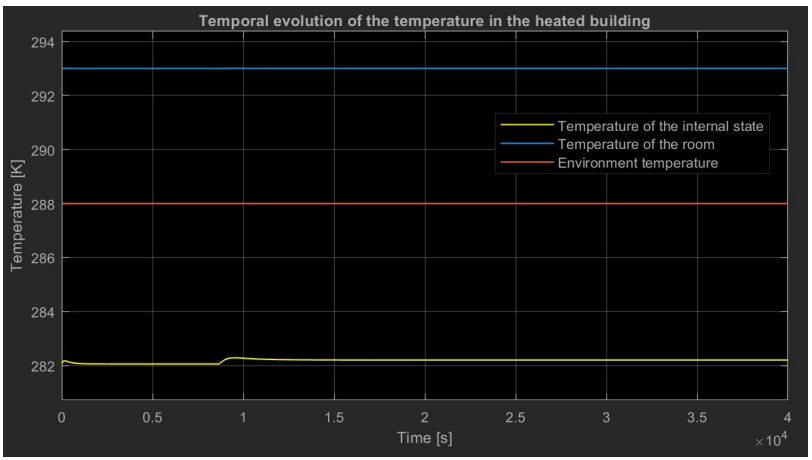

Figure 3.14 The graph shows the temperature in the heated building. It can be seen that after the cooled building lowers its temperature, more heat energy is supplied to the heated building since the internal state increases its temperature.

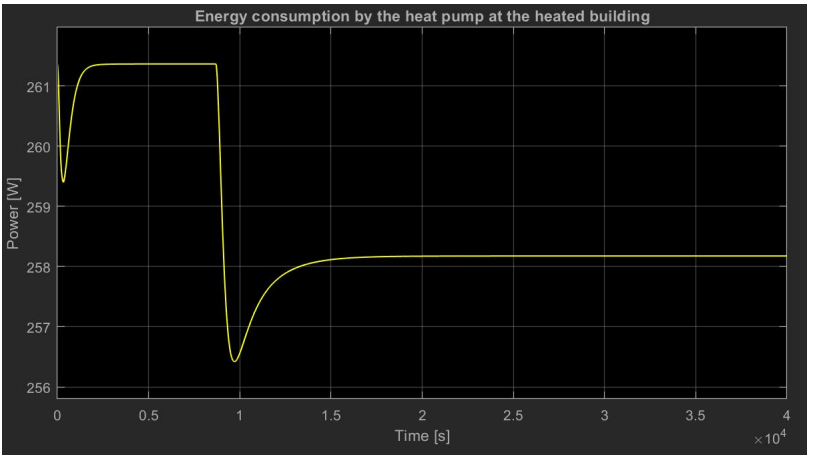

Figure 3.15 The energy consumption that the heat pump in the heated building uses to transfer energy between the internal state and the room temperature.

As a proof that some of the extra heat energy had made its way to the heated building could be seen by looking at its energy usage, see Figure 3.15. The input signal to the heat pump was lowered meaning that less power was required to heat it. This new relative abundance of heat reduces the energy usage of the heated building. One could thus say that the cooled building has in effect sold off some of its heat to the other user in the network.

4

Stationary Operation

In the previous chapter, it was shown that it was possible to send in signals to the heat pumps in order to reach some desired temperature of the building. In principle, given the temperature that the battery sets for the grid, it is always possible to calculate the control signal for the heat pumps. However, a much smarter approach could be made.

Suppose that there is a high energy demand during a time when electricity that drives the heat pumps is expensive. One could then use the battery to make sure that the energy consumption of the heat pumps becomes as low as possible by setting the correct temperature. Another scenario might be that during a time of low price on electricity and the battery wants to restock hot or cold water in its reservoir, it could set the temperature so that it sends in a minimal amount of thermal energy to the grid, while it is being resupplied either by a power plant or by the prosumers in the network. These scenarios require a very different temperature at the inlet of the grid.

One could think of this problem as a first step towards analyzing control of a highly delayed and nonlinear model that is this relatively simple network. By finding the optimal temperatures throughout the grid given some objective function with respect to minimizing some objective function under constraints, the optimal values for the input to the heat pumps can be generated.

4.1 Finding Optimal Stationary Points

During consideration of stationarity, it was assumed that the desired room temperatures had been fulfilled. This meant that the time derivatives of equation (3.38) were set to zero. From this, stationary values of the inputs u and x_i were obtained from the last four equations, which could be used to calculate how much heat they drew from the grid. This lowered the temperature of the pipe that was passing through the building. Time delays in this section were skipped, since they do not have an impact at stationarity.

Consider the following problem,

$$\underset{T}{\text{minimize}} \quad T^T QT + T^T W \tag{4.1}$$

subject to

$$g(T,P,T_e) \leq 0$$
$$h(T,P,T_e) = 0$$

with an appropriate, symmetric, positive, semidefinite Q-matrix, and under linear constraints in T, $h(T,P,T_e) = A(P,T_e)T - B(P,T_e)$. This problem could be formulated as, for example, minimize the heat loss due to interaction with the environment. By following ordinary stationary optimization as shown in [Böiers, 2010], and under the assumption that there are no inequality constraints (or assuming that all inequalities will be fulfilled), the solution of the minimization can be found as the solution to the following matrix relation

$$\begin{bmatrix} 2Q & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} T \\ \lambda \end{bmatrix} = \begin{bmatrix} -W^T \\ B \end{bmatrix} \tag{4.2}$$

where λ is an array of Lagrange multipliers.

When inequality constraints are present in the optimization, it is not difficult to expand the problem posed in equation (4.1). For each of those inequality constraints, one has to add them to the Lagrangian multiplied by a Lagrange multiplier, similarly to the equality constraints. However, additional constraints are imposed on the problem, namely $\mu_i \geq 0$ and $\mu_i g_i(T,P,T_e) = 0$. The solution is thus evaluated on every edge of the solution space, if one chooses the right combination of non-negative and zeros on μ_i.

In order to find the optimal stationary operating point, one must take into consideration different operating modes. This means that since each optimization is made for one particular makeup of heat exchangers, heat pumps and cooling machines, one has to try out all other viable combination of modes in order to find the minimal cost. In practice, this means that both the objective functions, i.e., Q and W matrices and the constraints are changed in order to find the best operating points.

At stationarity, all the time derivatives in equation (3.38) become zero, giving enough constraints to make the problem formulation of (4.1) only have one degree of freedom, i.e., one variable to optimize against, namely the input temperature at the beginning of the network. This control strategy requires the central power plant to obtain information about the desired reference temperature in each node. This particular control scheme requires a lot of overhead communication between the nodes and the central plant as well as the room specifications, such as room size or dissipation coefficients pipe lengths needs to be known. This implies that - should a major shift in the coefficients occur, such as installing better insulating windows, it would have to be reported to the production plant.

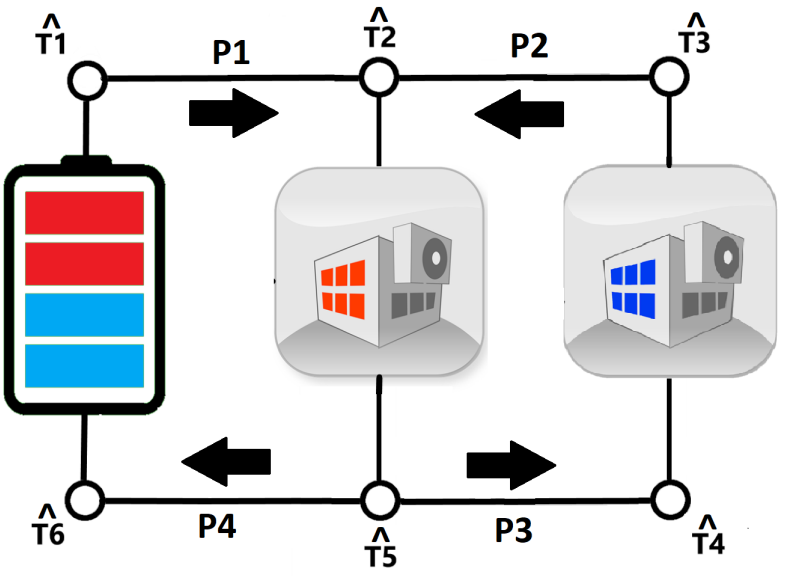

Figure 4.1 The figure shows the location of the states that are defined in (4.3).

4.2 Example: Optimal Stationary Inputs

Referring to Figure 4.1 for a definition of the state variables, let

$$\hat{x} = \begin{bmatrix} \hat{x}_1 \\ \hat{x}_2 \\ \hat{x}_3 \\ \hat{x}_4 \\ \hat{x}_5 \\ \hat{x}_6 \\ x_i^h \\ x_i^c \end{bmatrix}, \quad \hat{x}_k = \hat{T}_k - T_e \tag{4.3}$$

Suppose we want to minimize the following expression

$$\underset{\hat{x}}{\text{minimize}} \quad \gamma \hat{x}_1^2 + \gamma \hat{x}_6^2 + \varepsilon (x_i^h - x_r^h)^2 + \varepsilon (x_i^c - x_r^c)^2 \tag{4.4}$$

subject to

$$(\dot{m}_0+\dot{m}_2)\hat{x}_2 = m_0 e^{-P_1}\hat{x}_1 + \dot{m}_2 e^{-P_2}\hat{x}_3 \tag{4.5}$$
$$\hat{x}_6 = e^{-P_4}\hat{x}_5 \tag{4.6}$$
$$\hat{x}_4 = e^{-P_3}\hat{x}_5 \tag{4.7}$$
$$\hat{x}_5 = C_1\hat{x}_2 - (1-C_1)T_e \tag{4.8}$$
$$\hat{x}_3 = C_2\hat{x}_4 - (1-C_2)T_e \tag{4.9}$$
$$\kappa_1\hat{x}_4 = x_i^c - (\kappa_1-1)T_e \tag{4.10}$$
$$\kappa_2\hat{x}_2 = x_i^h - (\kappa_2-1)T_e \tag{4.11}$$
$$\tag{4.12}$$

where C_k, P_k and κ_k are a collection of physical constants, which were obtained by setting equations (3.29)-(3.37) to zero.

Expression (4.4) could be equivalently written as

$$\underset{\hat{x}}{\text{minimize}} \quad \hat{x}^T Q\hat{x} + \hat{x}^T W \tag{4.13}$$

Where the cost matrices are

$$Q = \begin{bmatrix} \gamma & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \gamma & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \varepsilon & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \varepsilon \end{bmatrix}$$

$$W = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ -2\varepsilon x_r^h \\ -2\varepsilon x_r^c \end{bmatrix}$$

where ε is the cost proportional to the input signal at the heat pumps and where $\gamma = \dot{m}_0^2 (c_p)^2 d$ is the cost for the plant to heat up the water. d is some discount factor that is related to how much cheaper it is to heat up water at the plant instead of using electricity to heat the water.

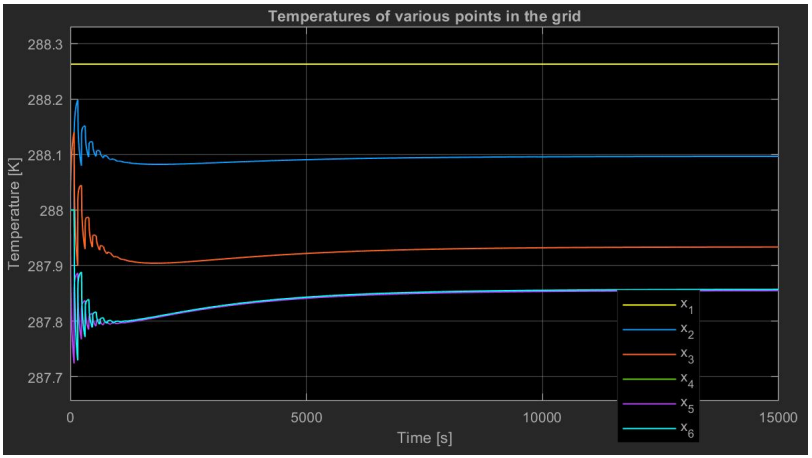

Figure 4.2 The graph shows the evolution and the stationary values of the temperatures in the thermal network. The states are shown in Figure 4.1

.

By writing the relations (4.5)-(4.12) as $Ax = B$, it was possible to use equation (4.2) to obtain the following temperatures throughout the grid,

$$\hat{T} = \begin{bmatrix} 288.2628 \\ 288.0966 \\ 287.9336 \\ 287.8571 \\ 287.8549 \\ 287.8571 \\ 288.7804 \\ 282.3632 \end{bmatrix} \tag{4.14}$$

for some estimated values of the physical constants used in (4.5)-(4.12).

A simulation of the grid using $\hat{T}_1$ as the input set by the battery and P^h and P^c, which were calculated from the respective x_i and x_r, is shown in Figure 4.2. Here, all nodes started at the environment temperature. at $t = 0$, they obtained a new reference value, and applied the inputs to the grid and heat pumps. One can see that after about three hours, the buildings had finally reached their reference values in Figure 4.3 and Figure 4.4.

This mode of open loop control was very time consuming. If it took three hours to heat a room, people would try to hold on to as much of that heat as possible. A smart system, where the room is heated when you arrive there, will have to start preheating hours before you arrive. The consequence will be that a lot of heat energy has to be stored inside buildings all the time, similar to how it is today. It was in our interest to try and find some better controller that could both quickly heat up a room

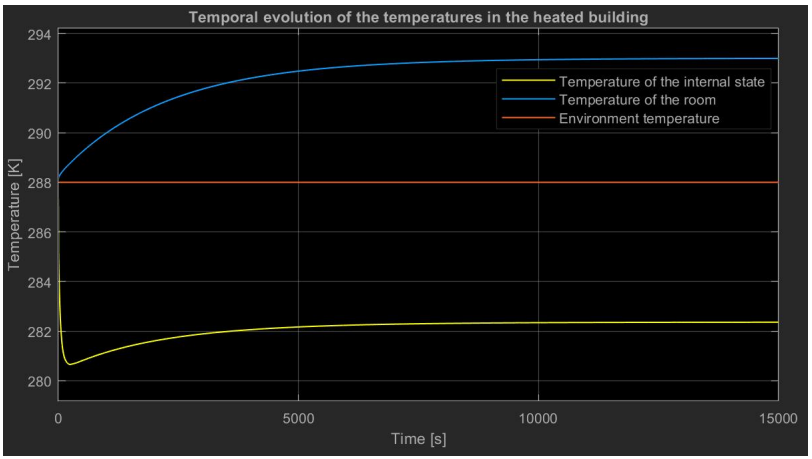

Figure 4.3 Temporal evolution of the temperature in heated building after a constant input temperature has been set.

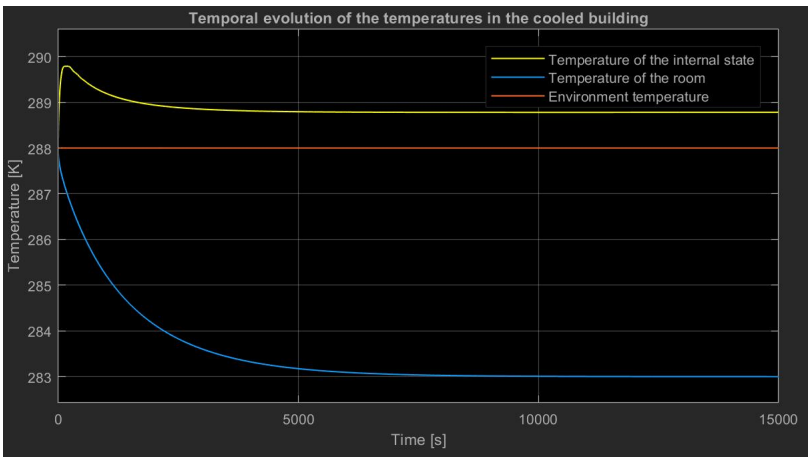

Figure 4.4 Temporal evolution of the temperature in cooled building after a constant input temperature has been set.

and then quickly cool it down when the heat was needed elsewhere. This type of controller will be discussed in in the subsequent chapters.

Another issue with this method was that the control signal needs to be calculated by using all other values of the grid. This implies that information about every state in the network had to be sent to all the controllers, in order to be able to calculate such an input. While a solution might be to use cloud services to connect everyone to each other, it would also introduce some security concerns. What would happen if connectivity to the cloud was lost due to server maintenance or power outage? There would still be a need for being able to provide heating and cooling.

Also, when it comes to these types of open control loops, any errors or changes in the models would produce wrong results, making it virtually impossible to guarantee that the temperatures would reach the correct values. It would also be impossible to deal with disturbances, which is with dealt in the next chapter.

5

Disturbance Suppression

The main advantage of having a heat sharing network could also be seen as a major flaw. Everyone has a heating or cooling cost that is inherently dependent on everybody else, which could be exploited by an attacker to drive up the cost. For example, an attacker could draw out as much energy as possible from the grid during times of high demand, making it very inefficient for the other users to extract energy. On top of that, it might also be difficult to deduce which node in the network is an attacker, since from an energy perspective, everybody increases their usage. A way to suppress this effect could be to synthesize controllers that make better use of the thermal grid when applying a control signal.

5.1 Disturbance Rejection with P-controllers

An important feature of large networks is that they need to be able to suppress disturbances in any way they may enter into the system. A simple example of a disturbance might be that your neighbor opens their window in the middle of winter, all while his heating pumps are trying to keep the room temperature way above the environment temperature. This implies that the grid temperature drops, and your own controller quickly needs to change its signal to address that error, which implies that it will need to increase its power consumption since there is less energy in the grid to use for heating of your room. A P-controller with a large gain is a highly reactive controller that is able to quickly address the temperature drop in your own room. Thus, by letting the change in the control signal from the stationary input be

$$\Delta P = K(r_r - x_r) \tag{5.1}$$

where K is the gain of the P-controller. A simulation of the heated building getting some load error is shown in Figure 5.1. One can see that a small stationary error appeared after the load disturbance entered. This error could be removed over time by adding a integrator. Here however, a controller that quickly removed as much

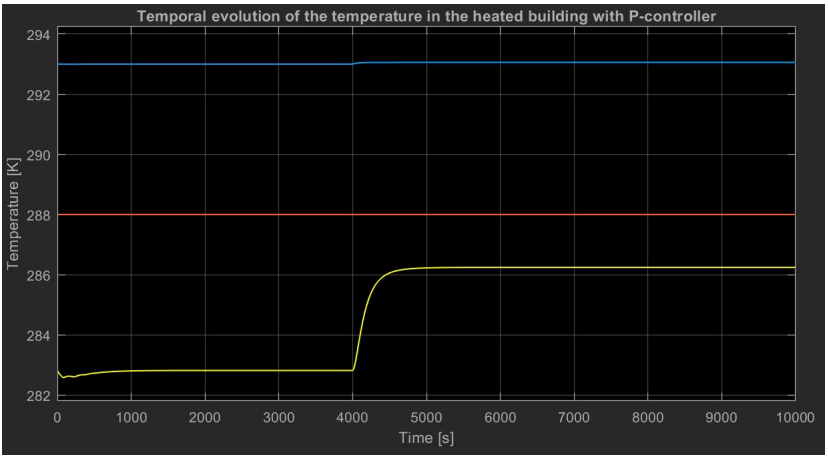

Figure 5.1 The figure shows how the heated building is affected by some internal load error. Since there is a P-controller maintaining the Temperature in the room, only the internal state changes its temperature.

stationary error as possible without integration was of interest. This was done in order to avoid oscillations, which will be shown to be a consequence of the integrator in Chapter 6.

Figure 5.2 shows the impact of the neighboring state. It can be seen that the P-controller here also removed the impact on the room temperature fairly quickly at the expense of increasing the power input, Figure 5.3. This high dependence on all the other prosumers in the network posed some security concerns. In this scenario it was possible for an attacker, who had control of one of the nodes, to fiddle with the energy usage in order to affect your controller without you knowing about it.

Conversely, there are privacy issues that enter on the other side of the argument. In this system, an attacker could simply observe the input signal to the heat pump to figure out if the heating or cooling machine is on in the other building, which could indicate whether or not someone is home. It is clear that a better controller was needed!

5.2 Model Simplification

As mentioned at the end of Section 3.3, difficulties arose when working with time-delayed systems. Before a good disturbance rejection controller was obtained, a few simplifications to the model was needed. The proposal was to approximate the transfer function of time delays by some Taylor expansion of the exponential function. This created the following approximate state space model of equation (3.38),

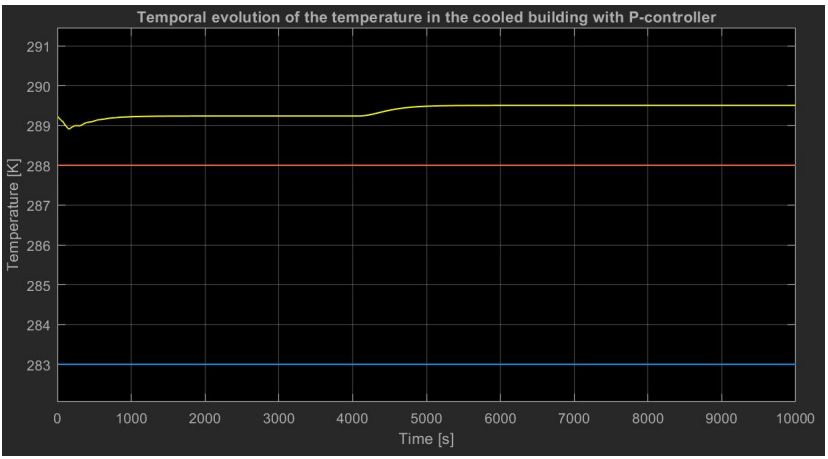

Figure 5.2 The graph shows how the load error propagates to a neighboring building. It can be seen that the internal state temperature rises as a response to the load disturbance.

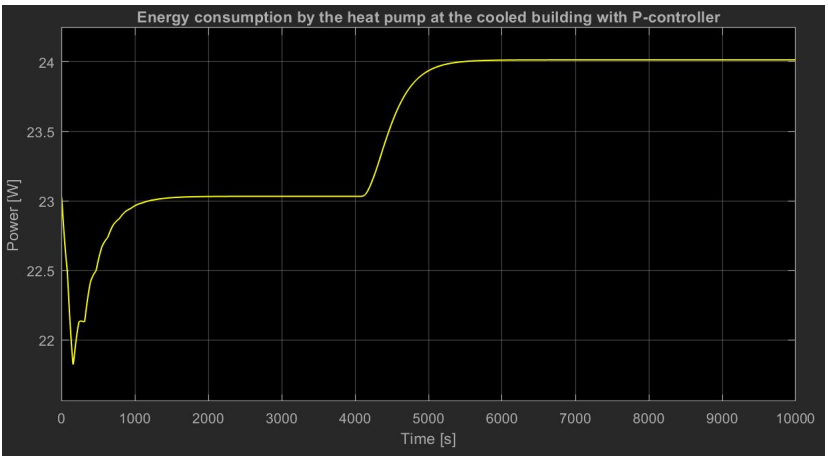

Figure 5.3 The graph shows how the load disturbance of the heated building changes the power consumption of the heat pump in the neighboring building to compensate for the energy loss.

$$z = A^t x(t-t_0) + B^t u(t-t_0) \Rightarrow$$
$$Z(s) = \mathrm{e}^{-st_0}\left(A^t X(s) + B^t U(s)\right) \approx \frac{1}{1+st_0}\left(A^t X(s) + B^t U(s)\right)$$
$$\dot{z} = \frac{1}{t_0}\left((A^t - I)x(t) + B^t u(t)\right)$$

which could be combined with equation 3.39 to obtain

$$\dot{x} = Ax + B(x)u \tag{5.2}$$

A simple way to control a non-linear system in the dynamical domain is to linearize the mathematical model around some operating point and create a controller that can makes the system behave in desirable ways. One may use this linearized model to either move with small increments between operating points, or to reject disturbances and errors as long as the controllers are able to respond to them quickly enough, so that the states do not move far away from the linearized operating point. By linearizing the model, one opened up to the possibility of many different control schemes that generally were difficult to synthesize for the full non-linear model. In the future, the linearized model dynamics will be denoted as,

$$\Delta\dot{x} = A^0\Delta x + B^0\Delta u \tag{5.3}$$

5.3 Stability

There were some natural questions that arose regarding this system, such as stability. One could determine the stability algebraically, by looking at the poles of the A^0 matrix and checking that they all had negative real parts. The problem was that this would only be valid in some unknown local surrounding of the operating point that the system was being linearized around. In this case, given known bounds of the entries in the matrices, Kharitonov's Theorem was applied to determine stability for the range of parameter values. The theorem states the following,

THEOREM 5.3.1—KHARITONOV'S THEOREM
An interval polynomial is the family of all polynomials

$$p(s) = a_0 + a_1 s + a_2 s^2 + \ldots + a_n s^n \tag{5.4}$$

where each coefficient $a_i \in R$ can take any value in the specified intervals $l_i \leq a_i \leq u_i$.

An interval polynomial is stable (i.e. all members of the family are stable) if and only if the four so-called Kharitonov polynomials

$$k_1(s) = l_0 + l_1 s^1 + u_2 s^2 + u_3 s^3 + l_4 s^4 + l_5 s^5 + \cdots$$
$$k_2(s) = u_0 + u_1 s^1 + l_2 s^2 + l_3 s^3 + u_4 s^4 + u_5 s^5 + \cdots$$
$$k_3(s) = l_0 + u_1 s^1 + u_2 s^2 + l_3 s^3 + l_4 s^4 + u_5 s^5 + \cdots$$
$$k_4(s) = u_0 + l_1 s^1 + l_2 s^2 + u_3 s^3 + u_4 s^4 + l_5 s^5 + \cdots$$

□

are stable.

A proof of this theorem may be found in [Yeung and Wang, 1987]. Application of this theorem to the system linearization in equation (5.3), was done by finding the bounds on every matrix element in the A^0 matrix. Then the characteristic polynomials was calculated from the matrices whose elements are all possible combinations of the low and high bounds on every element in the matrix. The highest and lowest value of each coefficient to the characteristic polynomial were stored and then used as the bounds on the coefficients as stated by Kharitonov's Theorem. The calculations were done with the assumption that the desired room temperatures ranged between 270K and 310K, which gave that the network was stable.

5.4 LQR Synthesis

By using the linearized model, derived in Section 5.2, an optimal controller that minimized some cost function was derived. For every type of operating point, a separate controller was needed, since the linearization was different at every operating point. Considering the following cost function,

$$\underset{\Delta u}{\text{minimize}} \int_{t_0}^{t_f} \Delta x(t)^T Q \Delta x(t) + \Delta u(t)^T P \Delta u(t) \mathrm{d}t + \Delta x(t_f)^T W \Delta x(t_f) \tag{5.5}$$

so that

$$\dot{\Delta x} = A^0 \Delta x + B^0 \Delta u$$

where the objective was to minimize the energy input, weighted with some factor of how large error is tolerable. This is the standard optimal control formulation. Finding an optimal control feedback law that minimizes the expression in (5.5) is the equivalent of finding a Linear-quadratic-regulator (LQR) controller when the dynamics are governed by linear differential equations, as it was this case. The solution to equation (5.5) was found by differentiating the Lagrangian of the minimization problem with respect to the input and then by setting it to zero,

$$L = \Delta x^T Q \Delta x + \Delta u^T P \Delta u + \lambda^T (A\Delta x + B\Delta u) \tag{5.6}$$

$$\Rightarrow \frac{\mathrm{d}L}{\mathrm{d}\Delta u} = 2\Delta u^T P + \lambda^T B = 0 \Rightarrow \Delta u = -\frac{1}{2} P^{-T} B^T \lambda \tag{5.7}$$

One can see that the optimal feedback was a function of the Lagrange multipliers. By inserting this expression of the input into the Lagrangian again, one got that the new expression became,

$$L = \Delta x^T Q \Delta x - \frac{1}{4} \lambda^T B P^{-T} B^T \lambda + \lambda^T A \Delta x \tag{5.8}$$

In order to find λ, the following differential equation needed to be solved,

$$\dot{\lambda} = -\frac{\mathrm{d}L}{\mathrm{d}x} = -2Q^T \Delta x - (A^0)^T \lambda \tag{5.9}$$

with the constraint

$$\lambda(t_f) = 2W\Delta x(t_f) \tag{5.10}$$

However, here an infinite time horizon was used for the control, thus $t_f = \infty$. This implies that the continuous algebraic Riccati equation was needed to be solved in order to get R, where $\lambda = 2R\Delta x$ for symmetric R. The Riccati equation is given by

$$A^T R + RA + Q - RBP^{-1} B^T R = 0 \tag{5.11}$$

The feedback law that minimized the cost, as defined by the expression in (5.5), was given by

$$\Delta u = -P^{-T} B^T R \Delta x \tag{5.12}$$

Distributed LQR

A major goal of this thesis was to generate a distributed control scheme with as little information as possible being shared with the battery and the other nodes in the network, making it easy to add and remove additional nodes into the system without the need of completely reshaping the feedback loop of all the other nodes in the network. It was not desirable to have the battery broadcast a feedback signal to all the nodes since, in addition to the concerns raised in Section 4.1, there might be limited bandwidth for systems with lots of nodes which could create time delayed input signals. Ideally, a local controller that was able to only use local information to create a feedback law was desired. However, due to the fact that a node not only directly affected its immediate neighbors, but it also affected their neighbors, although with a much smaller impact, and so on, an ideal controller would need the information of all the states to be able to act optimally, with respect to some cost function. The controller derived in (5.12), did precisely that using global information.

One could think of several ways to approximate an optimal controller with only local variables. A simple way to approximate an optimal controller with local information for feedback was to let the central plant collect all the reference levels of every nodes desired room temperature, determine the A^0-matrix, calculate the optimal feedback law and then broadcast the feedback gains to each controller. The local controller then set the gains of the signals that were not locally measurable to zero, which was approximately valid, since the local measurements had a feedback gain two orders of magnitude larger than the other measurements. This method did not only dictate the feedback law of all the heat pumps in the network, but it also was able to create a feedback law for the battery output. Here, this approach was used and a comparison with the full LQR was made, see Section 5.5.

A different approach might be to only treat the prosumers with a local optimal controller using only information nearby to calculate both the operating points around which it is linearized. The sparsity of A^0 gave only small errors if one treated each prosumer as a independent system. The optimal feedback law still only depended on local measurements and on local parameters. This control scheme was incomplete however, as there was no control law governing how the battery should have behaved. A controller for the battery in a very similar case is derived in [Scholten et al., 2017], although it is not optimal nor distributed. A way to go around this problem could be to use the feedback law derived in (5.12). However, instead of actively measuring all the states it would be possible to use a state estimator to estimate those states. By measuring states x_1, x_2 and x_3 here, one got that the entire system became observable. The downside was that this method could only be used in this particular topology.

5.5 Example: Disturbance Rejection with LQR

One may use the optimal controller derived in Section 5.4 to suppress errors with minimal energy with the help of the linearization of the model made in Section 5.2, by considering the very same situation as in Section 6.1, with the same cost function. However, instead of minimizing the input via the state variables, one can now minimize the actual dynamical input.

Let,

$$x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_r^h \\ x_i^h \\ x_r^c \\ x_i^c \end{bmatrix} \tag{5.13}$$

Using the new cost function (5.5), with the following cost matrices,

$$Q = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \gamma & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \varepsilon & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \varepsilon & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \tag{5.14}$$

$$P = \begin{bmatrix} \gamma' & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{5.15}$$

and following the derivation in Section 5.4 , a control law that not only minimized the error of the room temperature was obtained, but it also minimized the error of the grid temperatures. This cost matrix put a cost associated with the temperature in the pipe so that it minimized the change of its temperature. This makes the other nodes in the network feel less impact of some rogue node causing disturbances. It is worth to mention that the approximate control law at all three nodes in the network depended on the heavily on the x_4 state. The problem was that only the battery was able to directly measure it, which did not make this a distributed control law. However, one could extract additional information from the mathematical models. Given that the controllers at the buildings know the pipe dimensions of pipes P_3 and P_4 and the rate of mass flow, they simply measured the temperature of their outlet (heated building) or inlet (cooled building) flow and simply calculated the temperature at x_4. Those controllers could thus make use of that state in their feedback law as well.

Figure 5.4 shows a simulation of step disturbance that occured on the heated building. It is clear that both its internal state temperature and the input signal were changed. Observation of the cooled building, Figure 5.5, showed that only a small change had taken place on its internal state. One could also see that even though the heated building changed its energy usage just as much as for the case with a P-controller, Section 5.1, the control signal of the cooled building was reduced by about 30%, see Figure 5.6. It seemed as though the heated building extracted more energy from the grid while the cooled building extracted did not change its energy usage as much as before. The reason why this is possible was because of the battery. It took the hit of the energy fluctuations and countered them by raising the temperature and lowering it so that the effect on states in the pipes were minimized.

Section 5.4 discussed some privacy and safety issues that were present when using a P-controller for this type of suppression. With this controller, the attackers were sealed off and were no longer capable of doing any damage. An additional

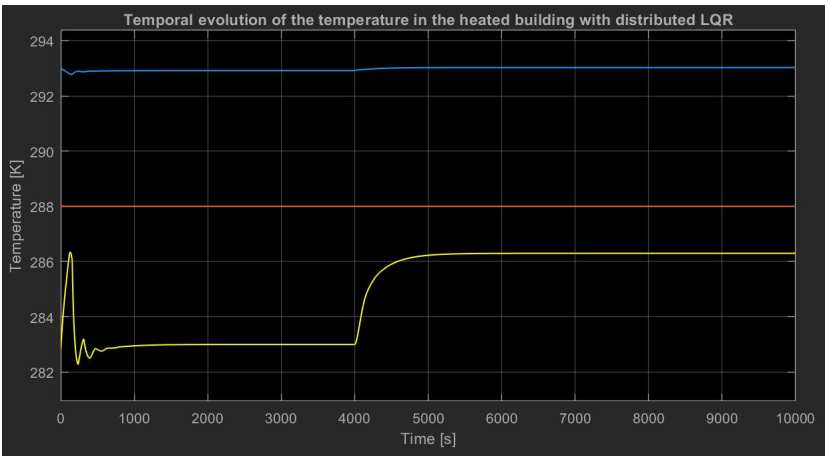

Figure 5.4 The figure shows a step disturbance on the heated building where an LQR is used to reduce the effect on the room temperature. One can clearly see the effect of the load disturbance on the internal state.

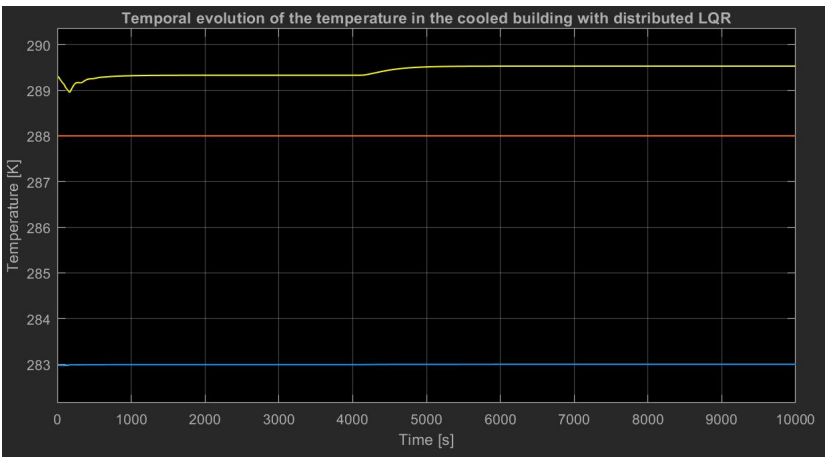

Figure 5.5 The figure shows a step disturbance on the cooled building where an LQR is used to reduce the effect on the room temperature. It can be seen that the keeps the room temperature at the desired setpoint while there is only a small bump in the internal state temperature.

Figure 5.6 The graph shows the energy usage of the cooled building after a load disturbance has affected the heated building.

Figure 5.7 The graph shows the full LQR feedback law in effect in order to compare the difference with the distributed LQR controller shown in Figure 5.4.

feature was that if the history of the control input to each building is saved, one could figure out which node in the network was the malicious one, a posteriori.

For reference, a simulation with the full LQR feedback, without removing any measurements, is shown in Figure 5.7 and Figure 5.8. It may be seen that the difference between these two figures and the distributed LQR were minimal, since the input signals differed at the magnitude of 1% of the full input signal.

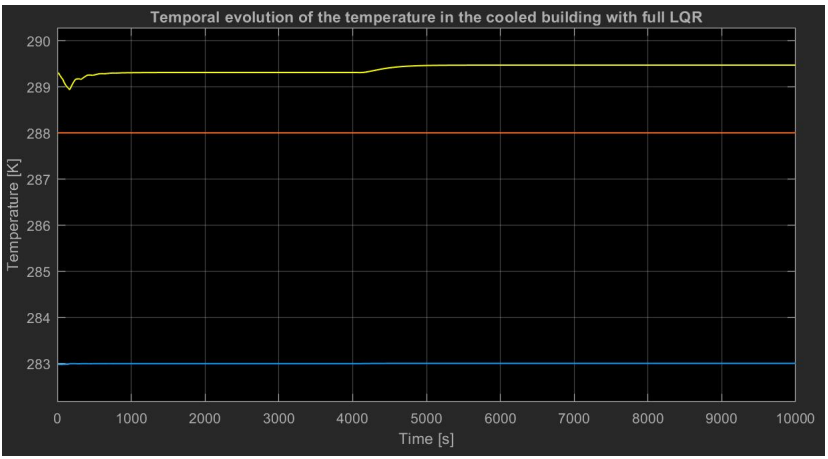

Figure 5.8 The graph shows the full LQR feedback law in effect in order to compare the difference with the distributed LQR controller shown in Figure 5.5.

6

Distributed Control of Changing Operating Points

The previous chapter showed an improvement of the regulation at stationarity. However due to the linearization of the model it was only valid in a small range around the operating point. It is of interest to be able to move between operating points as efficiently as possible with respect to the cost of power and the time it takes to do so. Section 4.2 demonstrated a very slow way to move the stationary points using global information. However, it was not only possible to speed this up with some more aggressive controller, which will be shown in Section 6.1, it was also possible to do it in a distributed way - without the need for any communication between the battery and the nodes, shown in Section 6.2 for an optimal control law.

6.1 Setpoint Change with PI-controllers

While the stationary behavior showed some very interesting things about a fourth generation grid as a mode of transporting energy, it was not very interesting from an automatic control viewpoint. Instead, a better thing study would be the transient behavior with some aggressive controller. The easiest way to obtain an aggressive controller at the heat pumps was to apply a PI-controller. While it was a very naive strategy, it can reach the setpoints quicker without any stationary error. A major advantage was that it only needs a difference between the desired and the current temperature level, which made it a distributed controller. A drawback with the PI-controller was that it was not clear how to obtain a control law for the battery. Here it was assumed that the battery still operated as it did in the stationary case, by collecting the desired set points of each building and then calculating an optimal inlet temperature to the grid.

Figure 6.1 and Figure 6.2 show the temperature evolution of the two rooms in this simple model. In the beginning of the simulation, all the temperatures in the grid were set to the environment temperature, 288K. At $t = 0$, both rooms received

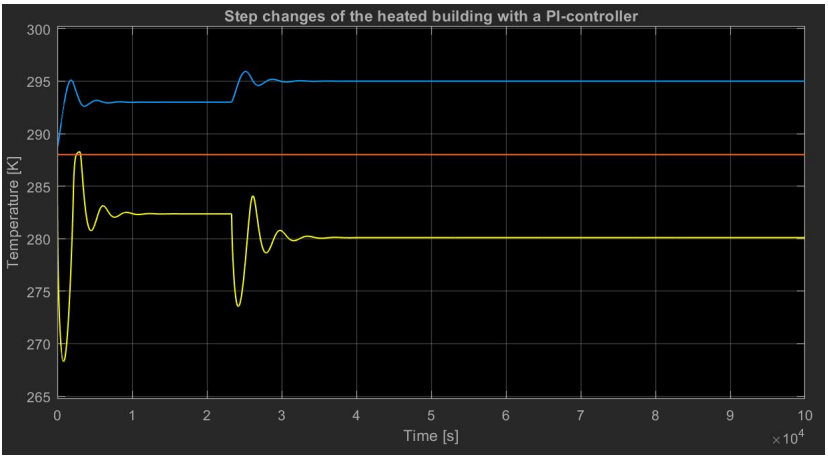

Figure 6.1 The graph shows two step changes for the heated building with a PI-controller. The first step is a step change that simultaneously changes the setpoint of the heated and cold building. The second setpoint change is only applied to the heated building.

a new reference point, 293K and 283K, for the heated and cooled building, respectively. As they came close their reference values, one can see that both of them got an overshoot. In order to cool off, they needed to shut of the heat pumps, as the heat pumps were not able to drive heat in the same direction as the temperature gradient, from the "hotter" side to the "colder". Therefore, there was no interaction here between the grid and the rooms, instead the extra heat in this scenario was dissipated out to the environment. This non-linearity of the control signal was the reason why both buildings started to oscillate heavily afterwards. At $3.8 * 10^4$ seconds, in Figures 6.1 and 6.2, one could also see that the oscillations were still there when only the heated building made a step change for its setpoint temperature level.

As could be seen in Figure 6.1, the oscillations were quite large, and a LQR would have a hard time suppressing those if there were several prosumers in the nodes that were changing their operating setpoints. While it is possible to tune the PI-controller to not have such oscillatory behavior, the trade-off often comes with a less aggressive controller, making the transition times much longer. These oscillations were present in the smallest step change as can be seen in Figure 6.3. Here incremental step changes of 0.1K, 0.5K, 1K and 1.4K were made - to show that the oscillations were present no matter how small of a change is made. This was done in order to show that it was not the limitations of the input signal that caused these oscillations, but rather the combination of the time-delays and the interaction with other prosumers in the system.

After some time, when the PI-controllers had converged to a signal that brought

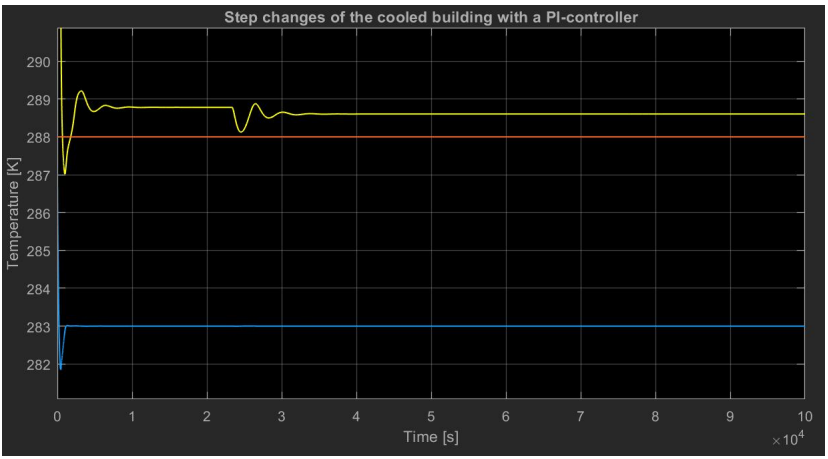

Figure 6.2 The graph shows a step change for the cold building that happens simultaneously as the heated building.

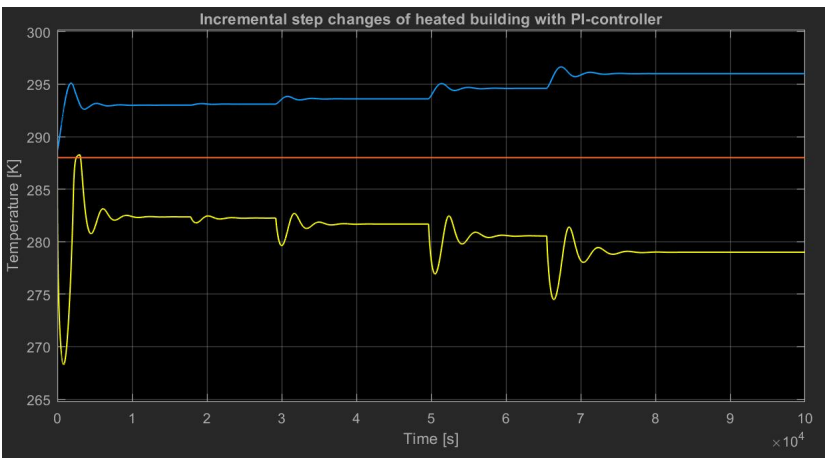

Figure 6.3 The graph shows incremental step changes in the desired room temperature of the heated building. One may see that no matter how small changes were made, there was always an oscillation that followed.

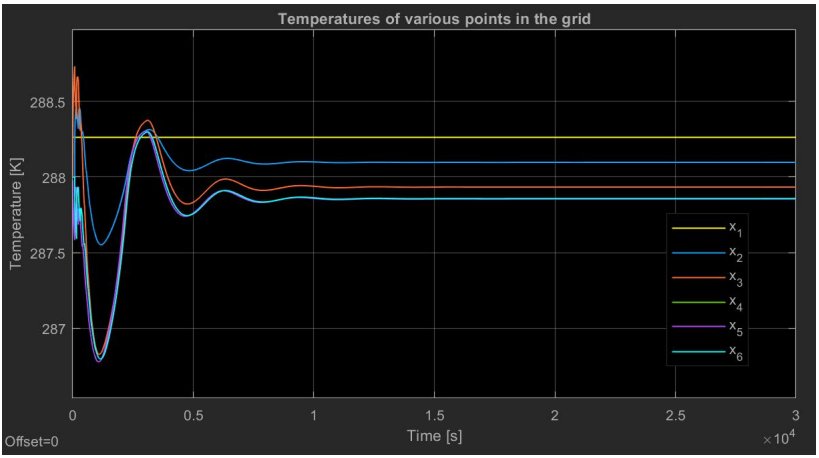

Figure 6.4 The graph shows the temperatures throughout the thermal grid.

the temperature of the buildings to their desired levels, their inputs were in fact the optimal ones calculated by the optimization in Section 4.1, given that the battery input temperature was the optimal one. In Figure 6.4, one can see how the grid temperature changed through the transition. Its final value was the same as predicted by (4.14).

6.2 Nonlinear Feedback

The optimal controller derived in Section 5.4 gave a possibility to suppress disturbances locally. It was possible to solve a similar optimization problem that was not limited to local dynamics of the system. The difference was that this controller depended on the states through some non-linear function. It was derived in a similar way as Section 5.4, using Pontryagin's Maximum principle. Consider the same optimization problem as in Section 5.4, where r is the reference vector for the desired values of the states

$$\underset{u}{\text{minimize}} \int_{t_0}^{t_f} (x(t)-r)^T Q(x(t)-r)+u(t)^T Pu(t)\mathrm{d}t+(x(t_f)-r)^T W(x(t_f)-r) \quad (6.1)$$

so that

$$\dot{x} = Ax+B(x)u$$

Just as in that Section 5.4, an infinite time horizon controller was considered. Because of that, set $t_f = \infty$ and as such, $x(t_f) = r$. First, setting up the Lagrangian gave,

$$L = (x-r)^T Q(x-r) + u^T Pu + \lambda^T (Ax + B(x)u) \tag{6.2}$$

Similarly as in problem (5.5), a differentiation of the Lagrangian (6.2) with respect to the input variables and setting the expression to zero gave an optimal input as a function of the Lagrange parameters, λ.

$$2u^t P + \lambda^T B(x) \Rightarrow u = -\frac{1}{2} P^{-1} B(x)^T \lambda$$

The expression for these Lagrange parameters were obtained through the reinsertion of the optimal input into the Lagrangian (6.2). By differentiating the expression with respect to the state variables and setting the system of equations to zero, it was obtained that the Lagrange parameters needed to solve the following expression,

$$\begin{gathered} L = (x-r)^T Q(x-r) + \lambda^T Ax - \frac{1}{4}\lambda^T B(x) P^{-1} B(x)^T \lambda \\ \Rightarrow \frac{\partial L}{\partial x} = 0 \Rightarrow \\ 2Q^T (x-r) + A^T \lambda + M = 0 \end{gathered} \tag{6.3}$$

where

$$M = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \begin{bmatrix}\lambda_5 & \lambda_6\end{bmatrix} (x_r^h + T_e) H_1(x_r^h, x_i^h) \begin{bmatrix}\lambda_5 & \lambda_6\end{bmatrix}^T \\ \begin{bmatrix}\lambda_5 & \lambda_6\end{bmatrix} (x_i^h + T_e) H_1(x_r^h, x_i^h) \begin{bmatrix}\lambda_5 & \lambda_6\end{bmatrix}^T \\ \begin{bmatrix}\lambda_7 & \lambda_8\end{bmatrix} (x_r^c + T_e) H_2(x_r^c, x_i^c) \begin{bmatrix}\lambda_7 & \lambda_8\end{bmatrix}^T \\ \begin{bmatrix}\lambda_7 & \lambda_8\end{bmatrix} (x_i^c + T_e) H_2(x_r^c, x_i^c) \begin{bmatrix}\lambda_7 & \lambda_8\end{bmatrix}^T \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ M_1 \\ M_2 \\ M_3 \\ M_4 \end{bmatrix} \tag{6.4}$$

$H_1(x_r^h, x_i^h)$ and $H_2(x_r^c, x_i^c)$ were symmetric matrices, whose elements depended on some of the state variables. Although it is difficult to see it here, it turned out that row 5 of $A^T\lambda$ only depended on the states x_r^h and x_i^h. Similarly, row 7 of $A^T\lambda$ only depended on the states x_r^c and x_i^c. The motivation behind setting equation (6.3), or the time derivative of the costate to zero, came from the comparison of the dynamics between $\dot{x}$ and $\dot{\lambda}$. The dynamics of $\dot{\lambda}$ were approximately ten times faster than the ones for $\dot{x}$ and so, the argument is that the transients for the costate have died out, putting the costate in a local equilibrium.

The difference between the linear approach and the non-linear was that the system of equations that needed to be solved were no longer linear, but some were quadratic. It was possible to cancel out the quadratic equations through a linear combination of row 5 and 6, as well as a linear combination of row 7 and 8, which converted the expression into six equations, linearly dependent on λ and two quadratic

equations, where one depended on λ_5 and λ_6, while the other one was composed of λ_7 and λ_8. This implied that there were four possible solutions to this system of equations, which were relatively easy to obtain through some elementary algebra.

Distributed Control

The optimal control law that was obtained from Section 6.2 will not be distributed, instead, every controller at every building would require full knowledge of the entire grid in order to optimally control the system. However, for certain cost functions, one could come close to some local optimal controller with some approximations. For example - given a general cost matrix of the states Q, the sparsity of A allowed the calculation of λ_1, λ_2, λ_3, λ_4, λ_6 and λ_8 as a function of the cost, Qx and the final two parameters λ_5 and λ_7.

From the way the B matrix was structured, the controller to the heated building would only depend on λ_5 and λ_6, while the controller to the cooled building would only depend on λ_7 and λ_8. Similarly, the optimal controller for the battery would only depend on λ_1. Numerical evaluation of the inverse of the system matrix A showed that if the elements of a diagonal cost matrix, for example the one shown in (5.14), were roughly of the same size, then the optimal controller to both buildings depended only on the cost of the room temperature, the room temperature itself and the internal state temperature. In the example that was being used here for the simulations, equation 6.3, it could be shown numerically that the almost optimal controller for the buildings was in fact almost an ordinary P-controller. The sparsity of the A matrix also implied that any other parameters, such as the controllers dependence on pipe length and mass flows, only depended on local components. The control law, could be formulated as,

$$P^h_{opt} \propto \frac{T^h_r}{T^h_i}(T^h_r - r^h_r) \tag{6.5}$$

for the heated building and,

$$P^c_{opt} \propto \frac{T^c_i}{T^c_r}(T^c_r - r^c_r) \tag{6.6}$$

for the cooled building.

Compared to an ordinary P-controller, this one increased its power when the ratio between the state with a higher temperature and the state with a lower temperature increased. If one looks at the previous simulations, for example in Figure 6.1, one could see that the difference between the internal state and the room temperature was at most 27K, where 270K was the lowest temperature. This implied that the optimal controller deviated from the P-controller by a maximum of approximately

$$\frac{297K}{270K} = 10\%$$

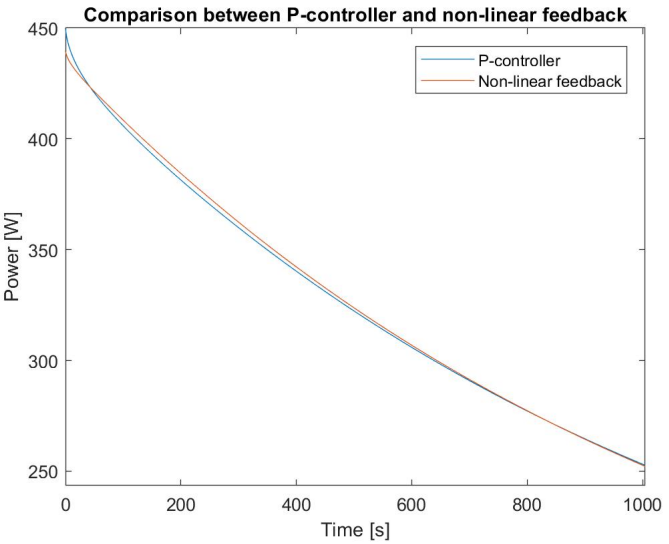

Figure 6.5 The figures show the input to the heated building for the two controllers being compared in this section. It can be seen that there is only a small difference in the beginning, afterwards, the two signals converge.

and so, there should not have been big differences when compared to a P-controller, at least in the stationary case. It was thus expected that this controller also produced a stationary error.

For the controller at the plant, the same calculations showed that it was highly dependent on the cost of all nodes throughout the network. As such, it was not possible to approximate an optimal local controller using this method.

6.3 Benchmarking the Non-linear Feedback

The PI-controller in Section 6.1 showed that oscillations are induced when a step change occurred. The optimal controller derived in the previous Section will be shown to gave better dampened these oscillations. It was be tuned so that without an integrator, it produced a stationary error that was the same as for the proportional part in the PI-controller obtained in Section 6.1. In the simulation of a step change for the heated building using both a ordinary P-controller and the non-linear feedback, one can see that there was no major difference between the two controllers, Figure 6.5, as predicted in the previous section. The small difference between the controllers, went away as the internal state and the room temperature went to their stationary values, creating essentially a P-controller with stationary error.

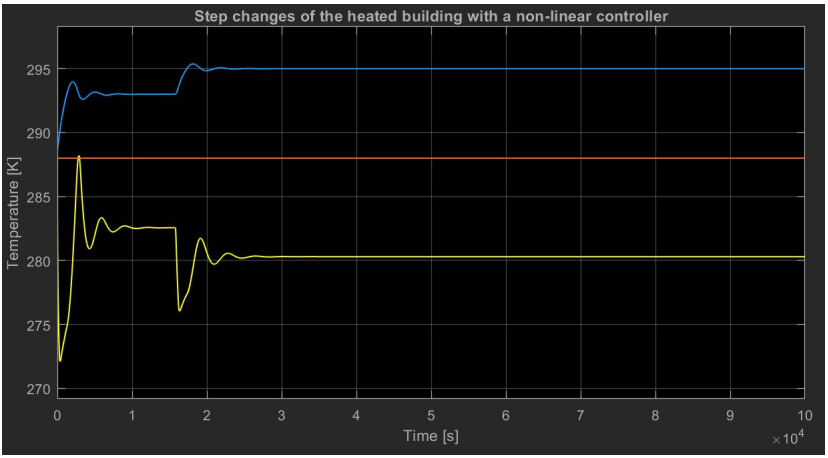

Figure 6.6 The figure shows a step change for the heated building using a nonlinear feedback law, where the gain has been swapped out for the optimal controller derived in the previous section.

For a successful transition between two operating points one needs to introduce an integrator that eliminates the stationary error. Looking at Figure 6.1 and recalling the PI-controller from Section 6.1, one can see that the introduction of the integrator induces oscillations in the network, due to the large overshoots and the inability to produce a negative input signal. Similarly, as for the test that were tried on the linearization case, this could potentially be used to harm the other nodes in the network, by making their controller try to follow this oscillating behavior. The difference here was that it was not a necessarily bad actor with the goal of damaging the network, but rather it was the users themselves that could cause the oscillations, simply by increasing the temperature a bit.

Figure 6.6 and Figure 6.7 showed the non-linear feedback together with integral action. The tuning of the controller had been set so that the integrating part was the same as the one as in the PI controllers, in other words, it had been tuned so that the two controllers reach the setpoints within the same time. The feedback gain had been set so that it had the same stationary error as the gain of the P-controller used for the control in Figure 6.5.

One can see that the optimal controller minimized the amplitude of both the overshoot of the room temperature and the amplitude of the oscillations of the internal state and subsequently the grid by about 30-50%.

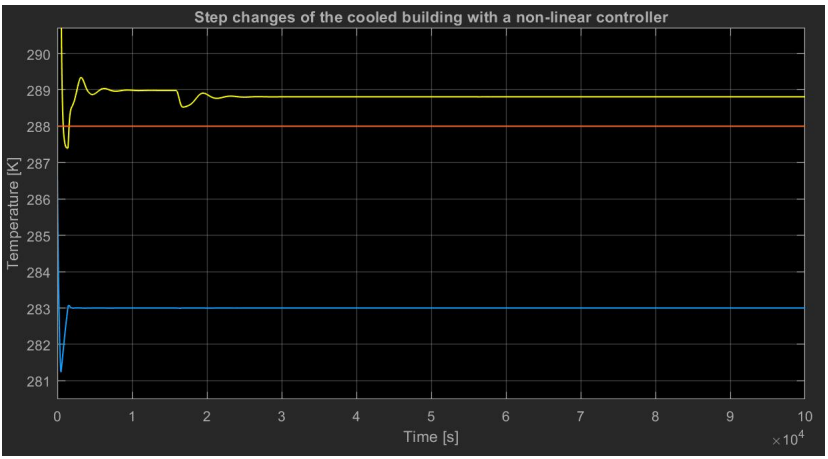

Figure 6.7 The figure shows a step change for the cooled building using a nonlinear feedback law, where the gain has been swapped out for the optimal controller derived in the previous section.

7

Conclusions and Future Work

As there was no clear literature regarding mathematical models of a fourth generation district heating system capable of depositing thermal energy to its neighbors, this thesis managed create such a model. By combining mathematical models of pipes, heat pumps and heat exchangers into a large network, where these components interacted with each other, a resulting network was obtained. This network was a highly nonlinear system, with several temporal delays existing throughout nodes of the system. These effects made it very hard for the network to be controlled in a predictable manner, especially with limited measurements. The goal to find a distributed controller that behaved better than P- and PI-controller was achieved after some simplifications were made. In fact - these controllers turned out to be able to partly help with some other issues regarding security and privacy, partly due to their distributed nature.

Through an ordinary optimization formulation, it was possible to find optimal stationary setpoints for the thermal grid. This made sure that during the vast majority of the time, the system operated by minimizing some cost function, such as energy usage or battery output minimization.

A simplification of the model was made in order to get close to optimal controllers in the dynamic case, when the network needed to transition to its optimal states. The resulting controllers behaved a bit differently from ordinary PI-controller, mainly by being able to more closely follow reference points and dampen oscillations which the PI-controller managed to introduce because of its overshoot. One could thus minimize this effect by using optimal feedback laws which were able to converge to the setpoint within the same time period as the PI-controller, only in a more robust manner.

Finally, disturbance rejection with minimal energy usage around a stationary setpoint was also compared to a P-controller. The optimal controller was able to better reject disturbances when compared to the P-controller by allowing the battery to take care of it instead. A consequence of this optimal controller was that it also

managed to address security and privacy issues, so that the grid could not be used to increase your energy bill nor could it be used to spy on you.

Future expansion of this work could include a deeper analysis of some other controllers minimizing some other cost functions, or using other measurements. One could also try to find better approximate solutions to the equations in Chapters 5 and 6, which were required for the optimal feedback. A major difficulty that arose in this project was the feedback signal for the battery/production plant, since it removed some of the symmetry in the problem. Finding a good controller for that plant would be helpful in removing much of the disturbance errors and minimize the effect nodes have on each other when changing reference points of their room temperatures.

Finally, additional work could be done on the modeling part. A major simplification was the use of approximating time delays with first order systems. Future works could analyze the system using a much better approximation for said delay either through refining the pipes by splitting them up into several parts or by finding some other approximations to the mathematical formulas. A different research direction could try to investigate what possibilities open up when the mass flows in the network are treated as input signals.

Bibliography

Böiers, L.-C. (2010). *Mathematical Methods of Optimization*. Studentlitteratur.

Clement-Nyns, K., E. Haesen, and J. Driesen (2010). "The impact of charging plug-in hybrid electric vehicles on a residential distribution grid". *IEEE Transactions on Power Systems* **25**:1, pp. 371–380. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2009.2036481.

E.ON (2018). *E.ON is building the world's first ectogrid™ at Medicon Village*. http://ectogrid.com/use-cases/medicon-village/. Online, accessed 2018-04-16.

Lund, H., S. Werner, R. Wiltshire, S. Svendsen, J. E. Thorsen, F. Hvelplund, and B. V. Mathiesen (2014). "4th generation district heating (4gdh): integrating smart thermal grids into future sustainable energy systems". *Energy* **68**, pp. 1 –11. ISSN: 0360-5442. DOI: https://doi.org/10.1016/j.energy.2014.02.089. URL: http://www.sciencedirect.com/science/article/pii/S0360544214002369.

Scholten, T., C. D. Persis, and P. Tesi (2017). "Modeling and control of heat networks with storage: the single-producer multiple-consumer case". *IEEE Transactions on Control Systems Technology* **25**:2, pp. 414–428. ISSN: 1063-6536. DOI: 10.1109/TCST.2016.2565386.

Spakovszky, Z. S. (2018). *Heat Exchangers*. http://web.mit.edu/16.unified/www/FALL/thermodynamics/notes/node131.html. Online, accessed 2018-04-13.

Verhelst, C., F. Logist, J. V. Impe, and L. Helsen (2012). "Study of the optimal control problem formulation for modulating air-to-water heat pumps connected to a residential floor heating system". *Energy and Buildings* **45**, pp. 43 –53. ISSN: 0378-7788. DOI: https://doi.org/10.1016/j.enbuild.2011.10.015.

Yeung, K. and S. Wang (1987). "A simple proof of Kharitonov's theorem". *IEEE Transactions on Automatic Control* **32**:9, pp. 822–823. ISSN: 0018-9286. DOI: 10.1109/TAC.1987.1104714.

Lund University
Department of Automatic Control
Box 118
SE-221 00 Lund Sweden

Document name
MASTER'S THESIS

Date of issue
June 2018

Document Number
TFRT-6063

Author(s)
Rijad Alisic

Supervisor
Richard Pates, Dept. of Automatic Control, Lund University, Sweden
Anders Rantzer, Dept. of Automatic Control, Lund University, Sweden (examiner)

Title and subtitle
Distributed Heating Networks

Abstract

There is a shortage of models and analysis methods of fourth generation district heating networks, which are capable of both extracting and depositing heat energy to some thermal network grid. This thesis fills that gap by combining mathematical models of components into a network that is capable of sending heat energy between its nodes. Questions regarding good heating strategies for controlling the nodes were posed, and based on these, some simplifications were made to produce simpler systems to work with. Near-optimal distributed control strategies were produced and tested on simulations of the full mathematical models. For comparison, tuned P- and PI-controllers were also simulated on the full system. The results showed that the optimal controllers induce less oscillations and had less stationary error, however this caused larger control signals and more grid-wise interaction, causing neighboring users to feel more of the impact when a single node changed its operating point. This effect can be suppressed if a heating battery is connected to the system.

Keywords

Classification system and/or index terms (if any)

Supplementary bibliographical information

ISSN and key title
0280-5316

ISBN

Language
English

Number of pages
1-65

Recipient's notes

Security classification