

Master of Science Thesis

DISTRIBUTED CONTROL OF REFRIGERATORS FOR THE SMART GRID

JOS VAN DEN HASPEL

Distributed Control of Refrigerators for the Smart Grid

MASTER OF SCIENCE THESIS

For the degree of Master of Science in Mechanical Engineering at Delft
University of Technology

J. van den Haspel

December 6, 2013

Faculty of Mechanical, Maritime and Materials Engineering (3mE) · Delft University of
Technology

Abstract

Weather-dependent energy sources such as wind and solar power lack the flexibility to adjust their power supply in response to the demand. To facilitate the integration of these renewable sources the electricity demand needs to respond to the available electricity supply. Thermostatically controlled loads such as refrigerators and electric water heaters are useful for this due to their autonomous operation, thermal buffer property, and high aggregate power demand. The work in this thesis is aimed at controlling a large-scale population of these thermostatically controlled loads to ensure that the aggregate power demand of these appliances matches the available renewable power supply.

Instead of a large centralized controller a distributed approach is chosen to provide robustness, modularity, and scalability. By using distributed model predictive control (DMPC), local controllers are able to anticipate on future events and to directly integrate local constraints such as user comfort levels. However, the majority of thermostatically controlled loads is constrained to binary on/off inputs, complicating the DMPC optimization due to discrete ‘jumps’ in the negotiations among controllers. Existing DMPC approaches make use of additional heuristics to terminate the optimization when these jumps result in non-converging oscillations.

The contributions of the work in this thesis are twofold. First, the potential of demand response using thermostatically controlled loads is investigated and translated into initial guidelines for a feasible match between the renewable supply and aggregate demand. Secondly, a serial implementation of feasible-cooperation DMPC is proposed to control the appliances. Closed-loop simulations show that unlike parallel DMPC methods from the literature such an approach converges towards the centralized MPC solution. Moreover, communication and computation requirements of the serial approach scale linearly with the amount of subsystems, without requiring a centralized coordinating agent.



The work in this thesis was supported by Almende B.V. Their cooperation is hereby gratefully acknowledged.



Copyright © Delft Center for Systems and Control (DCSC)
All rights reserved.

Contents

1	Smart Grids and Demand Response of Thermostatically Controlled Loads	1
1.1	Electrical grids	1
1.2	Smart grids	3
1.3	Demand response	6
1.4	Thermostatically controlled loads	7
1.5	Literature review	7
1.6	Thesis objective and outline	11
2	Thermostatically Controlled Load Models	13
2.1	Simulation model	13
2.2	Prediction model	16
2.3	Open-loop simulations and validation	17
2.4	Conclusions	18
3	Model Predictive Control for Demand Response	21
3.1	Controller design	21
3.2	Mixed integer linear program	22
3.3	Closed-loop simulations	24
3.4	Reachable power profiles	25
3.5	Model predictive control of multiple appliances	28
3.6	Conclusions	29
4	Distributed Model Predictive Control for Demand Reponse	31
4.1	Global system and objective	31
4.2	Feasible-Cooperation DMPC	32
4.3	Formulating objective functions	33
4.3.1	Aggerating local objectives	33
4.3.2	Decomposing a global objective	33
4.4	Parallel and serial update algorithms	37
4.5	Conclusions	38
5	Simulations and Comparison	41
5.1	Simulated control methods	41
5.2	Performance criteria	42
5.3	Synthetic reference supplies	42
5.3.1	Performance evaluation	43
5.4	Case study with real data	50
5.4.1	Performance evaluation	51
5.5	Conclusions	54
6	Conclusions and Future Research	55
6.1	Conclusions	55
6.2	Discussion and future research	56
	Bibliography	57

Chapter 1

Smart Grids and Demand Response of Thermostatically Controlled Loads

This chapter will briefly describe the context of this thesis. The chapter concludes with the statement of the research objective, based on conclusions drawn from a literature survey.

1.1 Electrical grids

There are few manmade systems around that are so complex and yet so present in everyone's life as the electric grid. People have grown so accustomed to the availability of electricity that it has become a resource most take for granted.

The tree-like structure of the electrical grid has remained relatively unchanged since its creation. In such a structure, large centralized power plants generate electricity that is distributed through high voltage lines. At the end of those high voltage lines, smaller distribution networks are used to deliver power to end-users: households and businesses (Figure 1.1).

One of the key characteristics of the grid is that the demand and the supply of electricity should be balanced at all times. Differences between the two will result in an increase or decrease of the grid frequency, and ultimately black-outs. A thorough coordination between consumption and generation is therefore required to keep the grid balanced and operational.

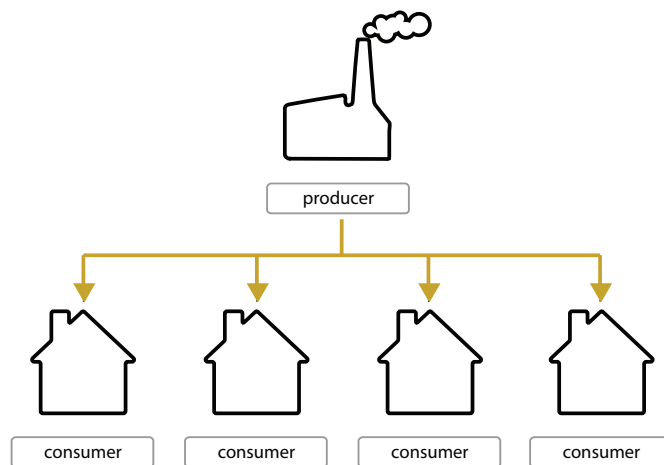


Figure 1.1: Conventional grid structure




Figure 1.2: Demand-driven supply chain

Electricity markets

For a long time, the coordination of generation and demand has been the responsibility of governments. In the past two decades many developed countries have liberalized their grid system, replacing centralized coordination by wholesale electricity markets. Here, large buyers and sellers negotiate market clearing electricity prices to balance supply and demand. The markets operate on various time scales; long-term contracts can be used to assure reliable supplies in the future, whilst real-time imbalances can be solved on spot markets. Within liberalized electricity markets four main actors can be identified:

- *Electricity producers*
Producers generate electricity using mainly chemical combustion, nuclear fission, and to a lesser extent renewables such as hydro-power, solar, and wind. They are generally large in order to benefit from economies of scale.
- *Transmission system operators*
Transmission system operators (TSOs) are natural monopolies that manage (a section of) the grid infrastructure. They are responsible for keeping the grid stable. When imbalances occur TSOs can offer price incentives on a separate balancing market to resolve them.
- *Electricity retailers/utilities*
Electricity retailers buy electricity at the markets and sell it to their contracted customers. They pay fees to use the infrastructure of the TSO and risk fines if they do not stick to their forecasted demand to prevent imbalances.
- *Consumers*
Most end-users such as residents and businesses buy electricity based on long-term contracts with their utility companies. Nowadays consumers usually pay time-of-use prices with a fixed high rate during the day and a lower rate during the night.

Demand driven supply chain

The demand curve for electricity is extremely inelastic. Even in the case of real-time pricing, changes in electricity prices will not significantly change consumer demand. Another way of putting this is that the consumer demand drives the electricity supply chain. This is illustrated by the simplified scenario that has been illustrated in Figure 1.2.

1. An end-user puts on a kettle to boil some water for a cup of tea.
2. The electricity retailer that has contracted the end-user needs to provide the required electricity at the previously agreed rate. If the retailer has not anticipated on this event, additional power needs to be purchased at an electricity market.
3. The increased demand of the utility company will drive up the market clearing price of electricity until it becomes profitable for producers to increase their output.
4. Producers generate the electricity required to balance the additional demand.

Unpredicted peaks or drops of the customer demand require power plants to ramp their production up or down very abruptly. As only a limited number of power plants, most notably gas turbine power plants, are able to do this and costs of such services are high, market clearing prices on spot markets may soar.




Figure 1.3: Projected development of renewable generation in the EU

Trends and challenges

The demand-driven supply chain leads to inefficiencies when forecasts are not accurate or large peaks in the demand occur. On top of this, recent trends in the energy landscape may result in additional stresses on the conventional grid structure, raising the question whether it is ready for the future.

- *Increasing demand and dependency*

The global demand for electricity is expected to double in the coming 25 years [19]. Apart from the increasing quantitative demand, the qualitative demand, i.e. dependency on a reliable supply of the electricity, has increased as well. A century ago, a blackout would require people to light candles to continue their work. Today it will most probably lead to a complete shutdown of operations. The growing quantitative demand may result in larger imbalances. The capacity and responsiveness of balancing mechanisms will need to increase significantly to ensure a qualitative electricity supply.
- *Increasing use penetration of renewables*

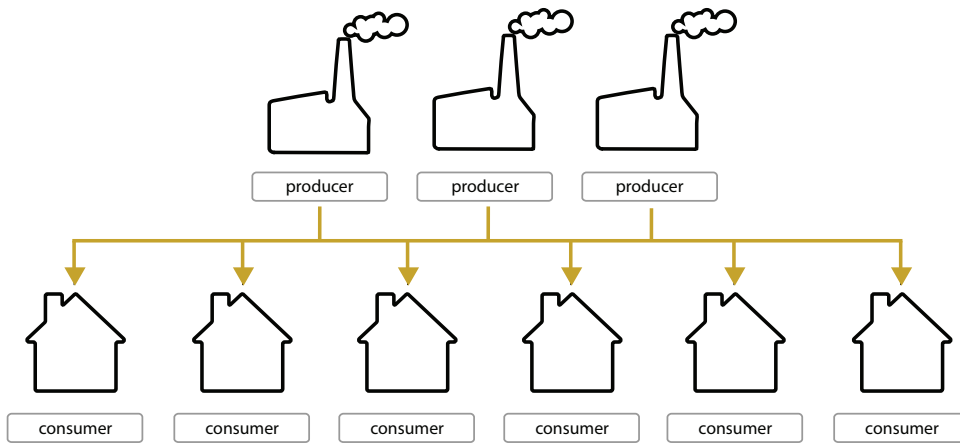
Partially driven by the increasing demand, an increasing fraction of electricity is produced using renewable sources such as wind and solar. Member states of the European Union have agreed to strive towards an average contribution of 20% of their energy mix from renewable energy sources by 2020 [9], [10]. The expected composition of these renewable energy sources in 2020 shows that over 50% of these sources depend on weather factors such as wind and sun intensity (Figure 1.3). Unlike fossil fuel based power plants, wind turbines and photovoltaic installations cannot easily be adjusted to temporarily increase their output. Uncertainty in wind predictions and volatility of wind energy production are therefore among the main concerns of TSOs [12].
- *Increasing distribution of electricity resources*

While traditional fossil-fueled power plants tend to be large due to benefits of scale, these advantages do not necessarily apply to renewable-based generation such as wind turbines and photovoltaic installations. In many western countries consumers have started investing in personal photo-voltaic installations, wind turbines, and micro combined heat and power systems, turning them into so-called 'prosumers'. This dramatically increases the amount and distribution of active components in the power grid, making it much harder to coordinate balancing, pricing, and the prevention of (local) congestions.

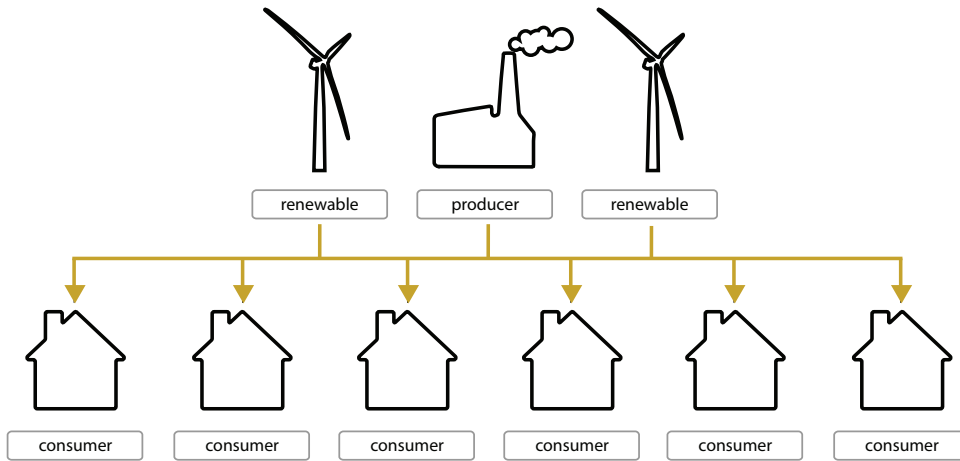
The classical grid may not be able to handle the described developments. Some of them, such as the increased distribution of electricity resources, may also enable new possibilities for a future, smarter grid.

1.2 Smart grids

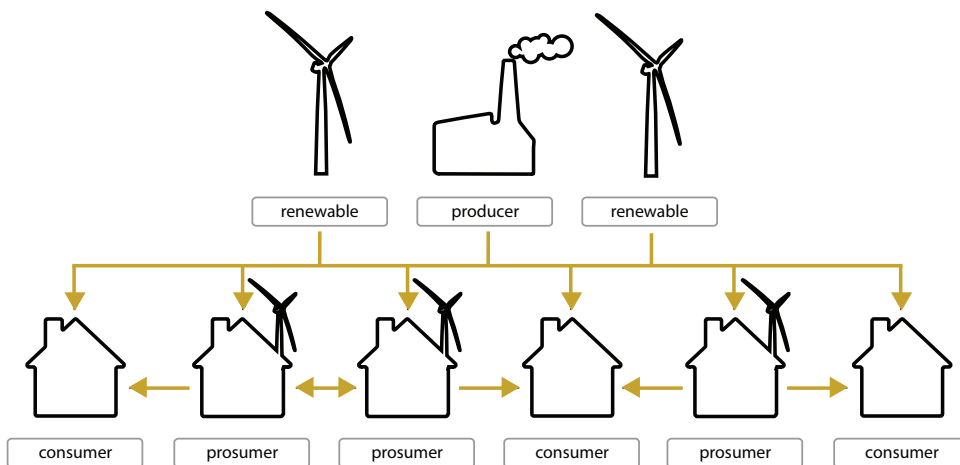
To improve the efficiency of the current grid and to cope with the upcoming challenges a smart grid has been proposed in [1]. This smart grid is a combination of the electric grid infrastructure and a



(a) Increasing demand



(b) Increasing penetration of renewables



(c) Increasing distributed generation

Figure 1.4: Three main trends and challenges to the existing grid

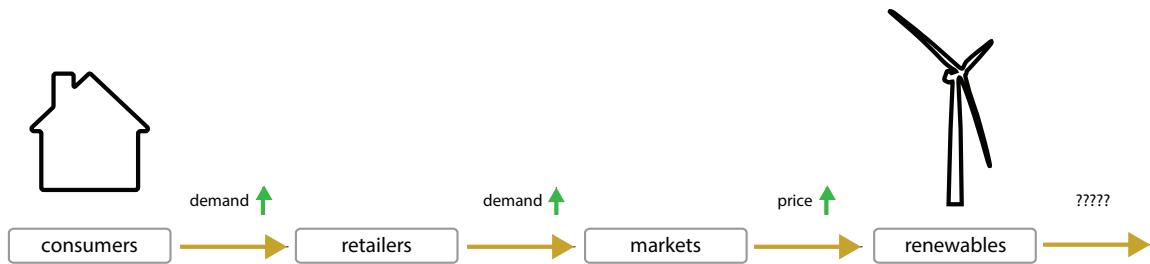


Figure 1.5: Demand driven supply chain in which weather dependent generation is unable to respond to an increasing demand



Figure 1.6: Reversed, generation driven supply chain

digital information and communications network. Individual components such as smart meters, wind turbines, and other nodes in the grid network can be fitted with some processing power and the ability to communicate with their neighboring elements. This opens up an enormous amount of opportunities for monitoring, efficiency, security, and control.

At this point the development of such a smart grid is widely considered to be the key to create a future-proof electric grid. Many large enterprises have dedicated specialized divisions to the development of these solutions. In the academical field many of the main universities and research institutions in the United States and the European Union have started smart grid research programs and dedicated conferences and journals have sprung up.

The scale of the smart grid developments is enormous. Besides technical challenges and the vast amount of stakeholders, there are still many political, legal, administrative and business hurdles to be taken. The focus of this research will therefore be on a specific branch of smart grid research that deals with some of the main challenges that have been laid out so far: demand side management.

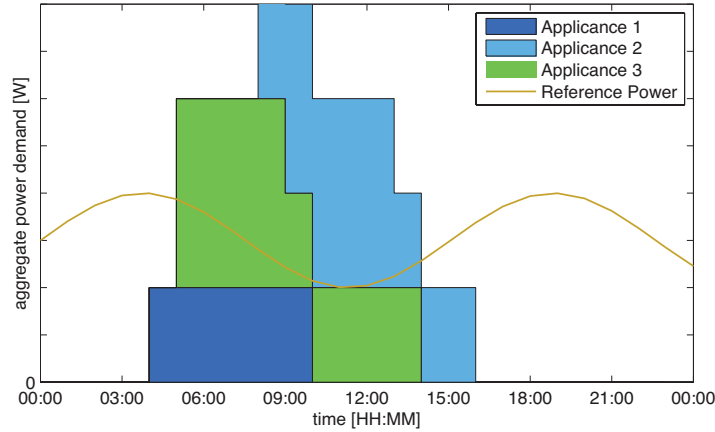
Demand side management

The demand driven supply chain of the electric grid is possible because power plants are able to adjust their output based depending on price levels. Wind and sun, on the other hand, do not respond to these pricing levels, making it harder to integrate renewable-based generation on a large scale. This can be seen in Germany, where a high penetration of PV currently leads to large grid imbalances and the use of expensive gas-turbine power plants to compensate these. The inability of renewable generation to comply with the conventional demand-driven supply chain is illustrated in Figure 1.5.

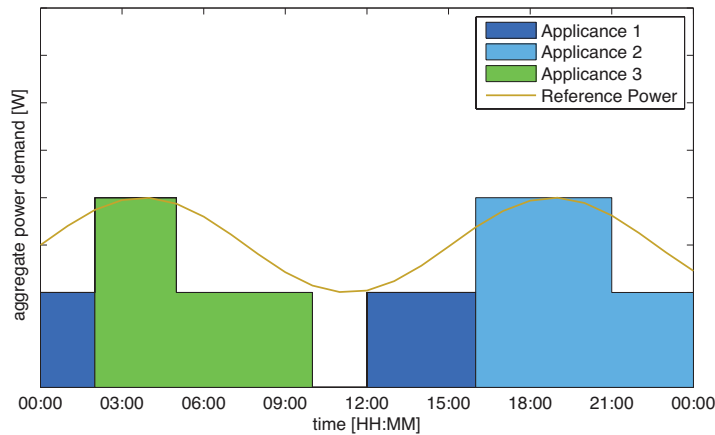
Grid energy storage can be extremely useful to buffer the temporary imbalances that may occur. At this point, however, the total grid energy storage capacity, mainly pumped-storage hydroelectricity, is marginal with respect to the total electricity throughput. Although the development of plug-in electric vehicles is expected to increase this capacity, additional methods will be needed.

A widely suggested alternative is to manage the demand to match a certain generated supply. This requires a paradigm shift and reversal of the conventional supply chain, this is illustrated in Figure 1.6.

It should be noted that in Figure 1.6 the function of the electricity retailer is to convert market price levels into certain incentives for consumers. If consumers (or prosumers) gain direct access to electricity markets the role of the electricity retailer in its current form may become obsolete. Another, more general interpretation of the retailer can therefore be that of an aggregator of consumers that can act on their behalf on wholesale electricity markets.



(a) Power demand of unscheduled appliances



(b) Power demand of scheduled appliances

Figure 1.7: Illustration of demand response principle

Different approaches are possible when it comes to the implementation of demand side management. Based on a centralized perspective, an utility could gain direct load control of customers appliances such as pool pumps. Such a centralized approach may have disadvantages with respect to privacy, acceptance, and adaptation by the end-users of the appliances. The work in this thesis is based on a distributed perspective, in which an utility offers incentives to initiate a locally controlled response of the demand.

1.3 Demand response

The general objective of demand response is that the customer demand for electricity adapts to a certain reference power that is determined by the coordinating party: a TSO, retailer, or other aggregator. This principle is illustrated in Figure 1.7. The customer demand might initially exceed the reference power (Figure 1.7a). Rescheduling the operation of these appliances will result in a better distribution and fit of the reference power (Figure 1.7b).

Demand response relies on user awareness of opportunity costs and their ability to make good decisions. This may be achieved by offering clear overviews of, e.g. the potential savings when a user operates its washing machine later than initially planned. A downside of this approach is that users will need to stay informed and repeatedly decide to respond or not. An alternative is that electric loads respond autonomously, on behalf of their users.

Obviously not all appliances are suitable for this: a television that switches itself off while the user is watching is unlikely to be appreciated.




Figure 1.8: By staying on for a longer time (t_2) this TCL decreases its temperature so it can stay off for longer (t_3 vs. t_1)

1.4 Thermostatically controlled loads

Thermostatically controlled loads (TCLs), such as electric water heaters, refrigerators, and air-conditioning systems, are designed to control the temperature of a conditioned space or mass. They are especially useful for automatic demand response because of some characteristic properties:

- *High impact*
Heating and cooling generally requires a lot of energy. In the United States TCLs comprise roughly 50% of electricity consumption [14]. In the EU-27, residential electricity consumption makes up 29.7% of the total electricity demand and similar to the U.S. TCLs make up about 57% of this domestic electricity demand [3]. Since TCLs comprise such a large fraction of the electricity demand, using them for demand response is likely to have a large impact.
- *Continuous availability*
Most TCLs are connected to the grid around the clock. Unlike a television that operates on demand of its user, the TCLs can turn on/off at any time. This means that the appliance is continuously available for demand response and not restricted by user-defined time constraints.
- *Autonomous operation*
TCLs are controlled autonomously; the thermostat controls the switching of the heat pump without user interference. Autonomous demand response is therefore unlikely to cause any inconvenience to the user, as long as the temperature stays within user-defined comfort bounds.
- *Storage capabilities*
TCLs have the ability to temporarily store energy in the form of temperature gradients with their environment. They can therefore anticipate on the future electricity supply, e.g. a refrigerator is able to temporarily lower its temperature to extend the time it can stay off afterwards. This principle is illustrated in Figure 1.8.

1.5 Literature review

The use of thermostatically controlled loads for demand response is not a new idea, and various control methods have been proposed in recent research. This section will briefly review and classify the work that has been done to date and identify new opportunities.

Demand response of single thermostatically controlled load

In the majority of the available literature the demand response problem is interpreted as a scheduling problem. Main examples of this approach can be found in [23, 26, 32].

These approaches all assume some form of static appliance properties. Thermostatically controlled loads, on the other hand, all have individual temperature dynamics. Based on these dynamics, accurate predictions of future states can be made and optimal inputs can be chosen accordingly. Model predictive control (MPC) is a control approach that determines optimal inputs based on the known system dynamics [30]. It provides the ability to anticipate on future events and to directly include operational or comfort constraints.

In [38] model predictive controller is developed for an electric heating system of a house. The controller objective is to reduce the cost of required electricity in a market with real-time pricing schemes. In [39] this approach is extended to optimize the use of locally generated electricity using a PV-installation.

In [2] a model predictive controller have been developed for a domestic freezer. In [4] a similar approach has been used to control a refrigerator. The objective of these controllers is to respond to signals coming from the utility.

The TCL models that are used to predict the future states all contain integer input variables. This results in (mixed) integer problems that are NP-hard. Moreover, due to the integer variables these problems are non-convex. Depending on their scale, these problems can be solved using solvers like CPLEX or, for linear problems, the GNU Linear Programming Kit. It should be noted that these solvers make use of heuristics such as the branch and bound algorithms, optimality can therefore no longer be guaranteed.

Demand response of multiple thermostatically controlled loads

Proposed model predictive control methods for multiple thermostatically controlled loads can be classified in three main areas, as seen in the key review papers [33] and [8]. The general architecture of these methods is illustrated in Figure 1.9.

- *Decentralized control*

In decentralized control all subsystems optimize their own objective function, without taking the other subsystems into account (Figure 1.9a).

This method has been used in [34]. Each TCL is equipped with a temperature sensor and the capability of measuring the frequency of the grid. When the frequency deviates from its normal value, indicating grid imbalances, the appliance can react by turning on or off. The appliances do so depending on their temperature, i.e. the refrigerators with the lowest temperatures will turn off first. This method has been implemented in Dynamic Demand¹, a pilot project with 1000 refrigerators.

The decentralized control approach can work well but the lack of communication and ability to coordinate actions in decentralized control is likely to result in a globally unstable system [31]. The lack of communication also makes it impossible to change the objective of the controller online, making this a very rigid approach.

- *Centralized control*

In centralized control the dynamics of all subsystems are combined in a single prediction model (Figure 1.9b). Although this is favorable in terms of convergence and stability, the size of the model will quickly result in mixed-integer problems that become intractable to solve.

In [6, 20, 24] state-bin models have been developed to generate a simplified representation of aggregated TCLs. In [25] this approach is extended to minimize the required information and communication among units, enabling the possibilities for centralized control of large-scale systems.

Disadvantages of a centralized approach are not necessarily computational, but of a more fundamental nature [35]. Depending on the application, a centralized approach might have shortcomings with respect to privacy, reliability, and robustness, since everything relies on a single control agent.

¹Dynamic Demand, www.dynamicdemand.co.uk

Moreover, conservative assumptions on the heterogeneity of subsystems might be needed to reduce the complexity of the global models.

- *Distributed control*

In distributed control all subsystems locally optimize an objective function whilst coordinating the optimization with other subsystems [7]. This approach has favorable properties with respect to scalability and robustness whilst maintaining coordination and stability. Although this generally comes at a price of extra communication and possible non-convergence of the negotiations, the distributed model predictive control approach appears promising and will be therefore be reviewed in more detail.

Distributed model predictive control of thermostatically controlled loads

distributed model predictive control (DMPC) lets subsystems determine their inputs locally yet coordinate their actions by communicating with other subsystems (Figure 1.9c). The goal of the communication is to reach a consensus on the values of the interconnecting variables by negotiation over multiple iterations. The type of connection can be through the dynamics, inputs, constraints, and/or the objective functions.

Dual-decomposition DMPC

Dynamic dual-decomposition techniques have been developed for DMPC in [16, 29] and extended in [15]. In these papers it has been shown that negotiations among linear systems converge to a global consensus assuming a strongly convex optimization problem.

This control approach has also been applied to demand response cases. In [5] this method is used to develop a distributed model predictive controller for a complex refrigeration system that includes display cases with inputs that are constrained to be binary variables. In [21] this method is used to develop a controller for a network of micro combined heat and power systems that also have binary input signals.

In both cases problems with using this method arise due to the binary inputs of the thermostatically controlled loads. Changes in the discrete inputs may result in large jumps in the locally optimal value of the interconnecting variables. Since the resulting optimization problem is no longer convex, this makes it more difficult for the negotiations to converge. To overcome this, additional heuristics can be applied that terminate the negotiations after a maximum number of iterations [5] or whenever the inputs start to oscillate [21].

Feasible-Cooperation DMPC

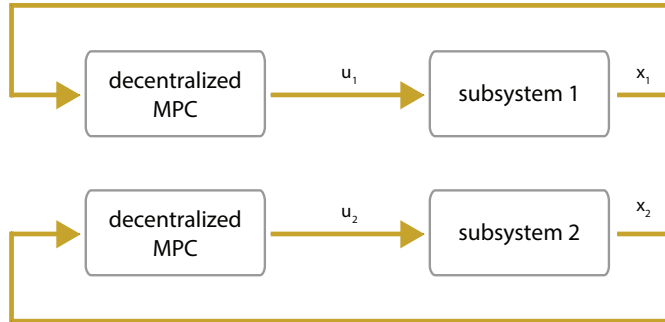
Feasible-Cooperation DMPC (FC-DMPC) is an alternative method proposed in [31, 35, 37] that has regularly been applied to control of electric power networks and complicated systems like hydro-power valleys. In this method the local objectives are replaced by a global objective function. At each iteration, each subsystem optimizes the global objective function with respect to its local input variables. During this optimization it is assumed that the input variables of the other subsystem remain equal to those in the previous iteration. After each optimization step the local optimal control inputs are broadcasted to the other controllers. Under certain conditions this method has been shown to converge to optimal centralized performance. To ensure feasibility and stability terminal penalties and constraints are added to the optimization.

Although feasible cooperation DMPC is able to cope with a larger variation of subsystems, integer inputs can still cause big jumps in the global objective function. However, the generality of this approach provides a good starting point for further research.

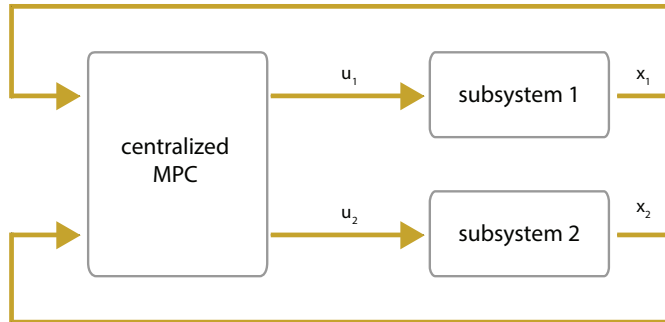
DMPC relies on the convergence of negotiations on interconnecting variables to reach a consensus. If the subsystems have discrete inputs this may result in big jumps in the objective functions, possibly leading to non-convergence of these negotiations.

Distributed model predictive control of hybrid systems

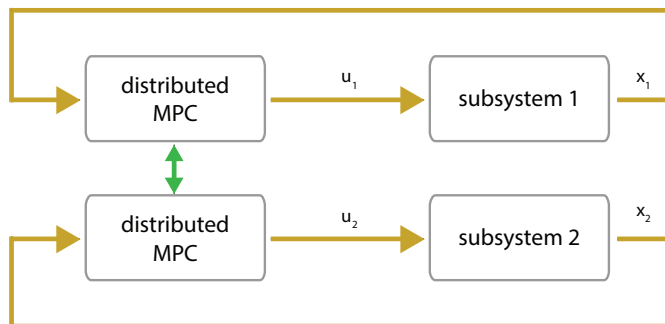
Due to their integer/binary inputs and continuous temperature dynamics, TCLs can be classified as hybrid or switched systems. Unfortunately, research on DMPC of hybrid systems has been very limited [17]. In [27] the possibilities for distributed model predictive controller for a network of hybrid



(a) Decentralized control



(b) Centralized control



(c) Distributed control

Figure 1.9: Methods for control of multiple systems

subsystems are explored. To overcome the non-convergence of the negotiations several heuristic methods are proposed. These methods are aimed at increasing the granularity of the discretization, terminating the negotiations at certain points, or increasing the penalty on deviations in the interconnecting variables during iterations. Whilst these methods have been shown to work in some cases, analytical work on the subject has not been found. The work in this thesis will therefore focus on new (combinations of) DMPC methods aimed at controlling hybrid systems.

Conclusions

A controllable demand is useful to facilitate the integration of renewable electricity generation. Thermostatically controlled loads are extremely useful to respond to a given electricity supply. MPC is useful to anticipate on forecasted energy supplies. To control multiple loads, DMPC offers advantages such as scalability, heterogeneity of the loads, and robustness of the implementation. Integer inputs complicate the convergence of the DMPC negotiations among subsystems. No analytical work on DMPC of hybrid systems has been found.

1.6 Thesis objective and outline

Building on the conclusions of the literature review, the objective of this thesis work is:

Develop a distributed model predictive control approach for a large-scale population of thermostatically controlled loads. The global control objective of the controller is to match a forecasted power supply whilst ensuring convergence of the negotiations among the loads.

The main focus of this thesis work is the study of the control architecture rather than the practical implementation of such a controller. Therefore the following initial assumptions are made:

1. All system states are fully observable. This means that all appliances are equipped with the necessary sensors. Disturbances are not taken into account.
2. All individual loads cooperate towards a global objective. This cooperation can be based on a community of owners, i.e. a neighborhood that wants to maximize the use of locally generated electricity, a single owner that own all appliances, or a utility that gives individual users incentives to cooperate.
3. Each load is able to communicate with all other loads without delays and loss of information.

Thesis outline

The structure of this thesis will be as follows:

- **Chapter 2: Thermostatically Controlled Load Models**
Models are useful for simulation and the prediction of future states. In this chapter the required models are formulated and validated using an experimental setup.
- **Chapter 3: Model Predictive Control of Demand Response**
The prediction model is used to synthesize a model predictive controller for a single TCL. The performance of the controller and its contribution to demand response is evaluated.
- **Chapter 4: Distributed Model Predictive Control of Demand Response**
To control multiple TCLs a distributed controller is developed based on existing DMPC methods. These methods are extended to cope with the characteristics of the TCLs.
- **Chapter 5: Simulation and Comparison**
The developed DMPC method is evaluated with respect to a centralized controller and alternative DMPC methods. To give additional insights in the usability and practical impact of the proposed method it is applied to two case studies.

- **Chapter 6: Conclusions and Future Research**

Based on the results of this thesis appropriate conclusions will be drawn, including a discussion and recommendations for further work.

Chapter 2

Thermostatically Controlled Load Models

In this chapter two different models are formulated that fit their respective purposes. They are validated using measured data from an experimental setup.

1. *Detailed simulation model*

This model is used to simulate the behavior of TCLs and to assess the practical performance of the proposed control solutions.

2. *Simplified prediction model*

This model will be used to synthesize a model predictive controller. A simplified model will be made to predict future states at a relatively fast rate, enabling the implementation of the model in real-time control.

2.1 Simulation model

In general, a thermostatically controlled load can be seen as a conditioned space or mass with a controlled temperature $\theta(t)$ that can be different from the ambient temperature θ^{amb} .

Based on basic thermodynamic principles, the passive temperature dynamics will result in a decrease of the temperature gradient $\theta(t) - \theta^{\text{amb}}$ over time. By activation of a heat pump, energy can be added or withdrawn from the system, respectively increasing or decreasing the temperature of the load. The vast majority of TCLs are based on a heat pump that can only turn on or off, e.g. the compressor of a domestic refrigerator. A schematic representation of the energy flows is presented in Figure 2.1.

The rate of the energy flows is determined by the parameters of the load. A parametrization that is commonly used in the literature [2, 6, 20, 25] is:

- *Absolute thermal resistance*




Figure 2.1: Schematic energy flows of a TCL that decreases the temperature of the load when switched on

The absolute thermal resistance R in [K/W] is a measure of isolation of the load. It determines how fast the passive temperature dynamics evolve.

- *Thermal capacity*

The thermal capacity C in [J/K] specifies the amount of heat required to change the temperature of the TCL. The thermal capacity is a collection of intervening variables such as mass and specific heat capacity.

- *Rated electric power*

The rated electric power $p(t)$ in [W] is the power required by the TCL in its active state. It can be a function of time.

- *Efficiency factor*

The dimensionless efficiency factor η indicates the fraction of the rated electric power $p(t)$ that can effectively be used to add or withdraw thermal energy from the load. The factor can take negative values for cooling appliances such as refrigerators and positive values for heating appliances such as electric water heaters.

Based on these parameters the passive and active energy flows can be expressed as:

$$\begin{aligned}\dot{E}_{\text{passive}}(t) &= -\frac{1}{R}(\theta(t) - \theta^{\text{amb}}(t)) \\ \dot{E}_{\text{active}}(t) &= \eta p(t)u(t)\end{aligned}\tag{2.1}$$

in which $u(t) \in \{0, 1\}$ is a binary control input that denotes whether the heat pump of the TCL is turned on or off. The sum of these two energy flows is the net energy change of the TCL, resulting in a temperature change of the load:

$$\dot{E}(t) = \dot{E}_{\text{passive}}(t) + \dot{E}_{\text{active}}(t) = C\dot{\theta}(t)\tag{2.2}$$

Temperature dynamics

The temperature dynamics of the load can be formulated by substituting (2.1) in (2.2), resulting in:

$$C\dot{\theta}(t) = -\frac{1}{R}(\theta(t) - \theta^{\text{amb}}) + \eta p(t)u(t)\tag{2.3}$$

which can be rewritten to:

$$\dot{\theta}(t) = -\frac{1}{RC}\theta(t) + \frac{1}{RC}\theta^{\text{amb}} + \frac{1}{RC}\eta R p(t)u(t)\tag{2.4}$$

This continuous-time expression can be discretized using a sample time t_s . For the sake of notational clarity we define the aggregated thermodynamic property:

$$\alpha = e^{-t_s/RC}\tag{2.5}$$

The discrete-time temperature dynamics can then be represented by:

$$\theta(k+1) = \alpha\theta(k) + (1-\alpha)\theta^{\text{amb}} + (1-\alpha)\eta R p(k)u(k) \quad (2.6)$$

or, formulated as a discontinuous differential equation:

$$\theta(k+1) = \begin{cases} \alpha\theta(k) + (1-\alpha)\theta^{\text{amb}} & \text{if } u(k) = 0 \\ \alpha\theta(k) + (1-\alpha)\theta^{\text{amb}} + (1-\alpha)\eta R p(k) & \text{if } u(k) = 1 \end{cases} \quad (2.7)$$

Power dynamics

Based on the method presented in [2] it is assumed that the electric power demand $p(k)$ can be modeled using an exponential equation, if we define the variable:

$$\beta = e^{-t_s/t_{c,p}} \quad (2.8)$$

in which $t_{c,p}$ is the time constant of the power dynamics. The overall power can then be expressed using a formulation similar to (2.7):

$$p(k+1) = \begin{cases} 0 & \text{if } u = 0 \\ \beta p(k) + (1-\beta)p_{\text{ss}} & \text{if } u = 1 \end{cases} \quad (2.9)$$

Thermostat control

The switching of the control signal u is based on hysteresis control. The objective is to keep the temperature of the conditioned load at a certain reference temperature θ^{ref} , allowing deviations within a temperature deadband θ_{Δ} . For a cooling appliance with $\theta^{\text{ref}} < \theta^{\text{amb}}$ the control logic can be formulated as:

$$u(k+1) = \begin{cases} 1 & \text{if } \theta(k) > \theta^{\text{ref}} + \frac{1}{2}\theta_{\Delta} \\ 0 & \text{if } \theta(k) < \theta^{\text{ref}} - \frac{1}{2}\theta_{\Delta} \\ u(k) & \text{if } \theta^{\text{ref}} - \frac{1}{2}\theta_{\Delta} \leq \theta(k) \leq \theta^{\text{ref}} + \frac{1}{2}\theta_{\Delta} \end{cases} \quad (2.10)$$

Hybrid automaton

Until now the TCL descriptions contain discrete inputs $u(k) \in \{0,1\}$ and continuous temperature and power dynamics $\theta(k)$ and $p(k)$. This combination is classified as a hybrid model. Various ways to formulate such models have been described in [11]. One of the basic model formulations that can easily include state resets is the hybrid automaton.




Figure 2.2: Hybrid automaton of a thermostatically controlled load

If we define a state vector $\mathbf{x}(k) = [\theta(k), p(k)]^T$ the dynamics in (2.7) and (2.9) can be combined:

$$\mathbf{x}(k+1) = \begin{cases} \begin{bmatrix} \alpha & (1-\alpha)\eta R \\ 0 & \beta \end{bmatrix} \mathbf{x}(k) + \begin{bmatrix} (1-\alpha)\theta^{\text{amb}} \\ (1-\beta)p_{\text{ss}} \end{bmatrix} & \text{if } u(k) = 1 \\ \begin{bmatrix} \alpha & 0 \\ 0 & 0 \end{bmatrix} \mathbf{x}(k) + \begin{bmatrix} (1-\alpha)\theta^{\text{amb}} \\ 0 \end{bmatrix} & \text{if } u(k) = 0 \end{cases} \quad (2.11)$$

Combining the discrete input $u(k)$ and continuous states results in a hybrid automaton (Figure 2.2). The control law as seen in (2.10) is now implemented by using the guards G . The power is reinitialized at each state change by the resets R , which can also be expressed as:

$$p(k) = \begin{cases} p^{\text{max}} & \text{if } u(k) - u(k-1) < 0 \\ 0 & \text{if } u(k) - u(k-1) > 0 \end{cases} \quad (2.12)$$

2.2 Prediction model

To synthesize a model predictive controller a simpler model is made, based on an abstraction of the hybrid automaton of the simulation model. The power demand is assumed to have a constant value $p(k) = p^c$ when the TCL is switched on, as seen in [6, 20, 24]. Now the state-reset policy can be omitted and the state-space model is can be simplified to:

$$\begin{aligned} \theta(k+1) &= \underbrace{\begin{bmatrix} \alpha \end{bmatrix}}_A \theta(k) + \underbrace{\begin{bmatrix} (1-\alpha)\eta R p^c \end{bmatrix}}_B u(k) + \underbrace{(1-\alpha)\theta^{\text{amb}}}_f \\ y(k) &= \underbrace{\begin{bmatrix} 1 \\ 0 \end{bmatrix}}_C \theta(k) + \underbrace{\begin{bmatrix} 0 \\ p^c \end{bmatrix}}_D u(k) \end{aligned} \quad (2.13)$$

in which the output $y(k) = [\theta(k), p(k)]^T$. This is a linear affine state-space model with binary inputs that can be used to predict future states in a very straightforward way.




Figure 2.3: Experimental setup

2.3 Open-loop simulations and validation

Both the simulation model and the prediction model are now validated using measured data from an experimental setup. Since the thermal properties of different TCLs will vary the objective of the validation is not to get an exact model fit but to assess the ability to reproduce the main dynamics and characteristics. The experimental setup consists of an ordinary domestic refrigerator. Both state variables are measured and sent to an online database, hosted and managed by Sense Observation Systems¹.

1. *Temperature*

The temperature inside the refrigerator is measured using a simple LM35 temperature sensor that is connected to an Arduino² board. The Arduino board is equipped with a WiFi connection using an Electric Imp³, this enables it to send all data to the online database.

2. *Power*

The electric power demand of the refrigerator is measured using a plug-in kit manufactured by PlugWise⁴. The collected data is also sent to the online database.

The sensor data is accessible through an online monitor on a custom website⁵, and can be imported directly to MATLAB for analysis. Figure 2.3 gives an impression of the experimental setup.

¹<http://www.sense-os.nl>

²<http://www.arduino.cc>

³<http://www.electricimp.com>

⁴<http://www.plugwise.com>

⁵<http://www.sensepower.herokuapp.com>

Table 2.1: TCL parameters

R	1.7	K/W
C	3200	W/K
θ^{ref}	5.25	°C
θ_{Δ}	3.7	°C
η	-0.5	-
p^{max}	82	W
p^{min}	50	W
$t_{c,p}$	300	s

Table 2.2: Variance accounted for per model and state

model type	state	calibration: VAF [%]	validation: VAF [%]
simulation	temperature θ	97.2	88.2
	power p	52.5	33.2
prediction	temperature θ	87.8	91.7
	power p	39.5	51.6

Model parameters

Initial estimates of the TCL parameters are based on physically feasible values. During a second iteration these variables are then manually adjusted to fit the calibration data from the measurement setup. The resulting values are shown in Table 2.1.

It should be noted that the refrigerator that was used for the experimental setup was almost empty. More content will significantly increase the thermal capacity C . For practical applications the variation of these parameters should arguably be accounted for by online parameter estimation.

Open-loop simulations

Both the simulation model and the prediction model are now used for open-loop simulations. Figure 2.4 shows the results of those simulations together with measured data from the experimental setup.

With respect to the TCL temperature, both the simulation model and the prediction model are able to approximate the measured values with acceptable accuracy. It seems that some small hysteresis effects at the input switchings are not accounted for in the models.


Looking at the simulated power demand, the simplifications that were made to formulate the prediction model are clearly visible. Using this prediction model for model predictive control will not enable the TCL to anticipate the relatively high start-up power. However, it is unclear whether the high initial power demand is present across different types of TCLs or limited to refrigerators and freezers. To analyze the performance of the models quantitatively the variance accounted for (VAF) of each simulated output variable \hat{y}_i with respect to the measurements y_i is computed using:

$$\text{VAF}_i = \left(1 - \frac{\text{var}(y_i - \hat{y}_i)}{\text{var}(y_i)} \right) \cdot 100\% \quad (2.14)$$


Computing these values for all states independently yields the values in Table 2.2. These results show that the temperature dynamics can be described accurately using both models. With respect to the power dynamics the prediction model counterintuitively performs better than the simulation model.

2.4 Conclusions

In this chapter two models have been constructed. A detailed simulation model is used to simulate the dynamics of a TCL. A simpler prediction model, in which the power demand of the TCL is assumed to



(a) Comparing open-loop simulation results and measured calibration data



(b) Comparing open-loop simulation results and measured validation data

Figure 2.4: Comparison of the TCL model simulations with measured data

have a constant value, is made to facilitate the design of a model predictive controller. Both models have been used for open-loop simulations. The simulation results have been compared to measured data from an experimental refrigerator, showing that the models are capable of reproducing the main dynamics and characteristics of a TCL.

Chapter 3

Model Predictive Control for Demand Response

Model predictive control (MPC) is an advanced type of control in which a sequence of future inputs is determined by optimizing a user-defined objective that may be a function of predicted model states, outputs, and/or inputs. The system dynamics are predicted using a dynamic model of the controlled system. Among the main benefits of this type of control is the ability to anticipate on future states and/or references and to optimize current control actions accordingly. This is an extremely useful property for demand response. Moreover, the single optimization problem provides an easy way to directly include additional constraints and objectives.

The amount of time steps into the future that the dynamics are predicted can be limited by a prediction horizon N_p . The amount of input variables that is optimized for can be limited by the control horizon N_c , such that $N_c \leq N_p$; this principle is illustrated in Figure 3.1.

At every time step k the open-loop optimal input sequence \mathbf{u}^* is computed. Of this sequence only the first input $u(k)$ is applied to the system. At the next time step $k + 1$ the optimization problem is solved again. This method is commonly referred to as a receding-horizon approach.

3.1 Controller design

Objective function

The main objective of the TCL controller is to match the power demand of the TCL to a given power supply. It is assumed that a forecast of the available power is provided. We define the predicted power demand $\hat{\mathbf{p}}(k)$ of the TCL and the forecasted power supply $\mathbf{p}^{\text{ref}}(k)$ as two vectors. Since the power demand is directly related to the inputs, these vectors contain predicted or forecasted values for every timestep

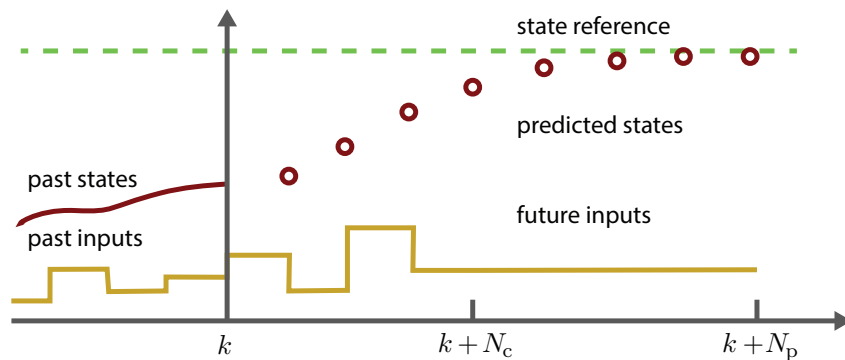


Figure 3.1: Model predictive control

$k + \ell|k$ of the prediction horizon such that $\ell = \{0, 1, 2, \dots, N_p - 1\}$, we have:

$$\hat{\mathbf{p}}(k) = \begin{bmatrix} \hat{p}(k|k) \\ \hat{p}(k+1|k) \\ \hat{p}(k+2|k) \\ \vdots \\ \hat{p}(k+N_p-1|k) \end{bmatrix}, \quad \mathbf{p}^{\text{ref}}(k) = \begin{bmatrix} p^{\text{ref}}(k|k) \\ p^{\text{ref}}(k+1|k) \\ p^{\text{ref}}(k+2|k) \\ \vdots \\ p^{\text{ref}}(k+N_p-1|k) \end{bmatrix} \quad (3.1)$$

The main objective can then be formulated as an L1 norm, equal to the sum of the absolute differences for each step of the prediction horizon:

$$J(k) = \sum_{\ell=0}^{N_p-1} |p^{\text{ref}}(k+\ell|k) - \hat{p}(k+\ell|k)| = \|\mathbf{p}^{\text{ref}}(k) - \hat{\mathbf{p}}(k)\|_1 \quad (3.2)$$

In addition to the main objective, secondary objectives can be added to this objective function, i.e. maintaining a certain reference temperature or minimizing the amount of switches. Such a composite objective function might force the controller to make compromises with respect to the power reference tracking performance. To assess this performance the secondary objectives are therefore included only at a later stage in this thesis.

Constraints

The optimization of the objective function is subject to a main constraint; the temperature of the TCL should never exceed certain temperature limits. These limits are often user-defined, e.g. comfort limits for temperature control of rooms or the temperature of hot water in electric water heaters. We define the predicted temperature trajectory vector as:

$$\hat{\boldsymbol{\theta}}(k) = \begin{bmatrix} \hat{\theta}(k+1|k) \\ \hat{\theta}(k+2|k) \\ \vdots \\ \hat{\theta}(k+N_p|k) \end{bmatrix} \quad (3.3)$$

and the constraint as:

$$\theta^{\min} \leq \hat{\boldsymbol{\theta}}(k) \leq \theta^{\max} \quad (3.4)$$

It should be noted that although the comfort limits in this constraint are assumed to be static over the entire prediction horizon, an extension to time dependent constraints can be implemented in a straightforward and similar way.

3.2 Mixed integer linear program

To optimize the objective function in (3.2) the future states need to be predicted. Moreover, the objective function contains absolute terms that complicate the optimization. In this section these issues will be resolved by rewriting the problem into a mixed integer linear program (MILP).

Predicting future states

Prediction of the future temperature $\hat{\boldsymbol{\theta}}(k)$ and power demand $\hat{\boldsymbol{p}}(k)$ is done using a recursive application of the prediction model in (2.13). Based on the current temperature $\boldsymbol{\theta}(k)$ and a certain input $\boldsymbol{u}(k)$ the next state can be expressed as:

$$\hat{\boldsymbol{\theta}}(k+1|k) = A\boldsymbol{\theta}(k) + B\boldsymbol{u}(k) + f \quad (3.5)$$

The temperature in the time step after that can be predicted based on the preceding state:

$$\hat{\boldsymbol{\theta}}(k+2|k) = A\hat{\boldsymbol{\theta}}(k+1|k) + B\boldsymbol{u}(k+1) + f \quad (3.6)$$

Substituting (3.5) in (3.6) yields:

$$\hat{\boldsymbol{\theta}}(k+2|k) = A(A\boldsymbol{\theta}(k) + B\boldsymbol{u}(k) + f) + B\boldsymbol{u}(k+1) + f \quad (3.7)$$

Generalizing this expression for all $j \in \{1, 2, \dots, N_p\}$ results in:

$$\hat{\boldsymbol{\theta}}(k+j|k) = A^j\boldsymbol{\theta}(k) + \sum_{i=0}^{j-1} A^{j-1-i} (B\boldsymbol{u}(k+i) + f) \quad (3.8)$$

Predicting future outputs

Knowing the state development of the temperature, the future outputs for $\ell \in \{0, 1, 2, \dots, N_p - 1\}$ can be predicted using the output equation of the affine state space prediction model in (2.13):

$$\hat{\boldsymbol{y}}(k+\ell) = C \left(A^\ell\boldsymbol{\theta}(k) + \sum_{i=0}^{\ell-1} A^{\ell-1-i} (B\boldsymbol{u}(k+i) + f) \right) + D\boldsymbol{u}(k+\ell) \quad (3.9)$$

Collecting the terms that result from this summation yields the matrices M_1 , M_2 , and M_3 . These matrices can be used to compute the outputs based on the current temperature $\boldsymbol{\theta}(k)$ and the sequence of future inputs $\boldsymbol{u}(k)$ in an affine way:

$$\hat{\boldsymbol{y}}(k) = M_1\boldsymbol{\theta}(k) + M_2\boldsymbol{u}(k) + M_3 \quad (3.10)$$

These predicted future outputs can be used to extract $\hat{\boldsymbol{\theta}}(k)$ and $\hat{\boldsymbol{p}}(k)$.

Slack variables

The absolute terms of in the objective function can be rewritten as a linear objective function subject to additional constraints by introducing an auxiliary slack variable ρ such that:

$$-\rho(k) \leq p^{\text{ref}}(k) - p(k) \leq \rho(k) \quad (3.11)$$

By introducing a new vector of optimization variables that includes the slack variables, $\mathbf{v}(k) = [\mathbf{u}(k), \boldsymbol{\rho}(k)]^T$, the objective (3.2) can then be rewritten in a linear form:

$$J(k) = [0 \quad \mathbf{1}] \begin{bmatrix} \mathbf{u}(k) \\ \boldsymbol{\rho}_p(k) \end{bmatrix} = \boldsymbol{\lambda} \mathbf{v}(k) \quad (3.12)$$

Synthesizing MILP

The expressions in (3.4), (3.10), and (3.11) can be aggregated in a single set of linear inequalities:

$$A_{\text{ineq}} \mathbf{v}(k) \leq b_{\text{ineq}} \quad (3.13)$$

The resulting optimization problem is a mixed integer linear program (MILP) of the form:

$$\begin{aligned} & \min_{\mathbf{v}(k)} \boldsymbol{\lambda} \mathbf{v}(k) \\ \text{s.t.} & \\ & A_{\text{ineq}} \mathbf{v}(k) \leq b_{\text{ineq}}(k) \\ & \mathbf{u}(k) \in \{0, 1\}^{N_p} \end{aligned} \quad (3.14)$$

This MILP is solved in MATLAB¹ using the MPT 3.0 toolbox [18] and the GLPK solver².

3.3 Closed-loop simulations

The TCL parameters used in the closed-loop simulations are identical to those that have been used in the open-loop simulations (Table 2.1). Since the the TCL inputs are restricted to binary variables, if the control horizon is smaller than the prediction horizon, the TCL inputs at $k+j$ for $j = N_c, \dots, N_p - 1$ will be constrained to either all *off* or all *on*. This will result in a continuous increase or decrease of the TCL temperature, may result in infeasible values or suboptimal solutions. This is illustrated in Figure 3.2. The control horizon is therefore chosen to be equal to the prediction horizon.

The length of both horizons is a trade-off between the speed of computation and the ability to anticipate on future events. Initial values of $N_p = N_c = 8$ are taken because they yield relatively fast results. Moreover, a sample time of $t_s = 60$ s enables the controller to anticipate on upcoming changes in the power supply. The controller is tested with three different power reference signals:

1. A square wave power reference supply
2. A cosine power reference supply
3. A constant power reference supply

¹MATLAB Release 2013a, The MathWorks, Inc., Natick, Massachusetts, United States.

²GNU Linear Programming Kit, <http://www.gnu.org/software/glpk/>




Figure 3.2: A control horizon that is smaller than the prediction horizon is likely to yield infeasible results due to the binary inputs of the TCL

Figure 3.3a and Figure 3.3b show that once the reference power becomes higher than approximately $0.5p^c$ the TCL is likely to switch on. The simulation results in both figures show cases in which the lower temperature constraint restricts the TCLs ability to switch on and to optimally track the reference power supply.

In Figure 3.3c the constant reference power supply is relatively low. Since the binary control action only allows values of either 0 or p^{\max} the TCL tries to stay off. At some point the upper temperature constraint is threatened to be violated and the TCL temporarily switches on to prevent this. To illustrate the impact of the restriction to binary inputs another closed-loop simulation is done in which the binary constraint is relaxed to allow real inputs. Note that the simulation model only has discrete input states $u(k) \in \{0, 1\}$ and therefore these closed-loop simulations are done using the *prediction model*.

The results of these closed-loop simulations (Figure 3.4) show a reference tracking performance that is significantly better. However, in all cases the lower temperature constraint makes it impossible to keep tracking the reference signal at all times.

3.4 Reachable power profiles

The closed-loop simulations have shown that regardless of the hybrid properties of the model, certain power profiles cannot be matched due to the temperature constraints. While this may not necessarily be problematic when this occurs infrequently, a structurally unreachable power profile will make the demand response controller unusable. This issue is largely unmentioned in the literature.

In this section critical characteristics of the power profiles are investigated, resulting in initial guidelines to understand the limitations of certain combinations of supply and TCL demand. These guidelines can be used when assessing the performance of the control methods later in this thesis. The following three characteristics have been found:



1. *Average power*

To compute the minimum and maximum average power demand³ we relax the binary constraint on the inputs and allow real valued inputs, enabling a steady state temperature at which $\theta(k+1) = \theta(k)$. For a cooling TCL (with $\eta < 0$ and $\theta^{\min} < \theta^{\max} < \theta^{\text{amb}}$) at the minimum or maximum steady state temperature we assume respectively a maximum and minimum power demand:



$$p^{\max} = \frac{\theta^{\min} - \theta^{\text{amb}}}{\eta R} \quad p^{\min} = \frac{\theta^{\max} - \theta^{\text{amb}}}{\eta R} \quad (3.15)$$

We assume that these values are equal to the minimum or maximum average power demand $\bar{p}(k)$



³Theoretically any real input between 0 and 1 can be achieved by using binary inputs and pulse width modulation. However, this approach requires rapid switching which is physically unfeasible due to a minimum dwell time.



(a) Square reference supply




(b) Cosine reference supply




(c) Constant reference

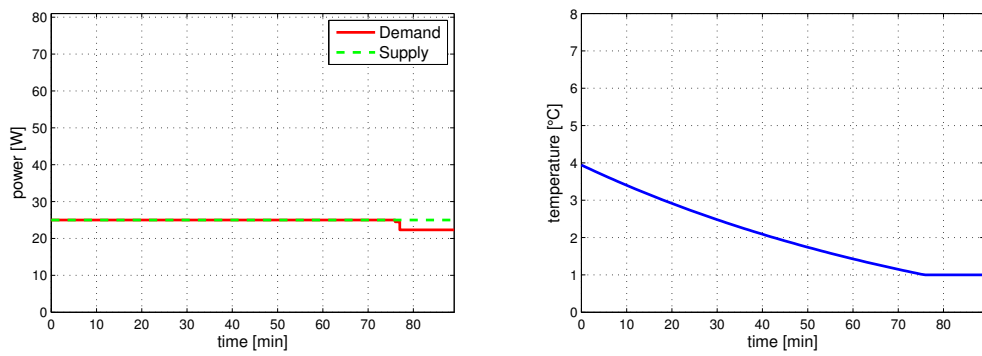
Figure 3.3: Closed-loop simulations using a TCL with binary inputs and various references.



(a) Square reference supply



(b) Cosine reference supply



(c) Constant reference supply

Figure 3.4: Closed-loop simulations using a TCL with real valued inputs and various references

for a TCL with infinitely fast switched inputs at either of the temperature boundaries. One of the criteria for a usable controller is therefore that the average power supply \bar{p}^{ref} is between this minimum and maximum average demand.

2. *Power variations*

Apart from average power values there are limitations to the variation with respect to that average. Naturally the reference power cannot exceed the minimum or maximum power demand of the TCL, so that $p^{\text{min}} \leq p^{\text{ref}}(k) \leq p^{\text{max}}$.

3. *Periodicity of the power profile*

The reference supply profile is also restricted by the speed of the TCL temperature dynamics. Given certain dynamics and constraints, all TCLs have a maximum time that can be taken to move the temperature from its upper to lower temperature and back. Since TCLs cannot stay on or off for longer, any periodical power supply with a period that is longer than the maximum on/off cycle time of a TCL cannot be matched its demand.

3.5 Model predictive control of multiple appliances

Up to this point the work has been focused on the control of the demand response for a single TCL. Two of the main advantages of controlling multiple TCLs are:

1. *Increased impact*

Controlling more TCLs has the advantage of having a larger impact. However, as discussed in the literature review, aggregating many individually controlled appliances in a decentralized way may result in instabilities. Therefore some coordination among the appliances is required.

2. *Reduced relative demand response error*

In the previous closed-loop simulations we have seen that due to the binary inputs of the TCL, the controller is not able to accurately match the reference signals. By having more appliances, the relative error due to the binary inputs can be reduced.

One of the ways to control multiple TCLs is to use centralized MPC. In centralized MPC the dynamics of all TCLs are combined into a single global model. A straightforward approach to formulate such a model is to concatenate the local system states in a single global state vector. Since there is no dynamic coupling or coupling in the inputs this results in block diagonal state space matrices for the global system.

We consider a set of N appliances. By taking the individual appliances $n \in \{1, 2, \dots, N\}$, we can define a global state and global input as:

$$\mathbf{x}_g = \begin{bmatrix} x_1(k) \\ x_2(k) \\ \vdots \\ x_N(k) \end{bmatrix} \quad \mathbf{u}_g = \begin{bmatrix} u_1(k) \\ u_2(k) \\ \vdots \\ u_N(k) \end{bmatrix} \quad (3.16)$$

The evolution of this state and outputs can then be described by:

$$\begin{aligned} \mathbf{x}_g(k+1) &= \begin{bmatrix} A_1 & & & \\ & A_2 & & \\ & & \ddots & \\ & & & A_N \end{bmatrix} \mathbf{x}_g(k) + \begin{bmatrix} B_1 & & & \\ & B_2 & & \\ & & \ddots & \\ & & & B_N \end{bmatrix} \mathbf{u}_g(k) \\ \mathbf{y}_g(k+1) &= \begin{bmatrix} C_1 & & & \\ & C_2 & & \\ & & \ddots & \\ & & & C_N \end{bmatrix} \mathbf{x}_g(k) + \begin{bmatrix} D_1 & & & \\ & D_2 & & \\ & & \ddots & \\ & & & D_N \end{bmatrix} \mathbf{u}_g(k) \end{aligned} \quad (3.17)$$

This state-space formulation is used to synthesize the centralized controller by using the method that has been described in Section 3.1 and Section 3.2.

Since all global dynamics are included in a single MILP, centralized MPC is assumed to yield an optimal solution given that the solver has sufficient time to converge. However, as previously discussed in Chapter 1, the centralized approach has disadvantages with respect to computational speed, robustness, and adaptivity. The next chapter will therefore focus on distributed control methods. To evaluate the performance of these distributed methods, centralized MPC will serve as a reference approach to assess the relative optimality of the distributed approaches.

3.6 Conclusions

Based on the closed-loop simulations the following conclusions are drawn:

1. The hybrid nature of the TCL model with binary inputs severely restricts the ability of a single TCL to track a certain reference supply. By using multiple appliances the relative error may be decreased.
2. The temperature of the TCL with current parameters and binary inputs does not converge to a static value. Instead a certain switching policy or pulse-width modulation approach is required to maintain feasible temperatures.
3. A TCL may not be able to completely match certain power supplies due to its temperature constraints. To avoid this and fully benefit from the demand response potential, both the power supply profile and TCL demand should have similar average power levels and variability.
4. To control multiple TCLs, all individual TCL models can be aggregated in a single global model. This centralized MPC approach is useful to benchmark alternative DMPC methods and to assess their relative (sub)optimality.

Chapter 4

Distributed Model Predictive Control for Demand Reponse

Distributed model predictive control (DMPC) is an alternative for centralized MPC as seen in Chapter 3. The aim of this chapter is to synthesize a robust DMPC method that can handle the hybrid properties of the prediction models that were used in the previous chapters. First, a general global system description is given, together with a general formulation of the global objective of the demand response. After this, a DMPC method from the literature will be applied to this global system and objective. Various local formulations of the global objective and both parallel and serial implementations of the DMPC algorithm will be investigated and the various properties and benefits of these methods will be evaluated. The chapter will conclude with the proposal of a DMPC method to be tested in Chapter 5.

4.1 Global system and objective

As previously seen in Chapter 3, we consider the set \mathcal{N} that contains N appliances. Instead of a single centralized controller, each appliance $n \in \mathcal{N}$ has its own local controller. These controllers are able to communicate with all other controllers $m \in \mathcal{N} \setminus \{n\}$. We assume that all systems $n \in \mathcal{N}$ are dynamically uncoupled and that the local power demand $p_n(k)$ is a function of the local input $u_n(k)$ only:

$$p_n(k) = f_n(u_n(k)) \quad \forall n \in \mathcal{N} \quad (4.1)$$

The entire set of appliances \mathcal{N} is subject to a *global* power reference signal $p_g^{\text{ref}}(k)$ indicating the available power supply. The *global* power demand of the appliances is expressed as the sum of the local power demands:

$$p_g(k) = \sum_{n \in \mathcal{N}} p_n(k) \quad (4.2)$$

The main objective of the demand response is to minimize the mismatch between the global supply and global demand by controlling that demand accordingly:

$$\min_{p_g(k)} |p_g^{\text{ref}}(k) - p_g(k)| \quad (4.3)$$

A schematic representation of the DMPC structure with $N = 3$ appliances is presented in Figure 4.1.




Figure 4.1: Schematic overview of a DMPC structure with $N = 3$ appliances

Based on the literature review in Chapter 1 it was concluded that FC-DMPC as proposed in [37] would be a suitable starting point to develop distributed control methods for demand response.

4.2 Feasible-Cooperation DMPC

Feasible-Cooperation DMPC is a distributed control method proposed in [37]. The core principle of this method is that every controller n locally optimizes the same global objective function J_g for its local optimization variables S_n . This optimization can be done over multiple iterations i within each time step k .

During each iteration i , each controller n assumes that the interconnecting variables S_m of the other controllers m remain equal to those from the previous iteration, S_m^{i-1} . For brevity of the notation in the remainder of this chapter, the dependency on the time variable k is dropped from the notation. The local optimization problem can then be formulated by:

$$S_n^i = \arg \min_{S_n} J_g(S_n, S_m^{i-1}) \quad (4.4)$$

For discrete linear systems this method has been shown to converge to an optimal solution, the performance of which is equal to the centralized MPC solution. Since the local controllers do not compete but cooperate towards a common global goal, this solution is Pareto optimal from a game-theoretic point of view [31]. To reach this optimal solution a large number of iterations may be required. In some cases the sample time may not be long enough to allow this, forcing the iterations to terminate at a suboptimal solution. By introducing a terminal penalty in the optimization problem this sub-optimal solution is assured to be feasible and stable. The combination of these properties make FC-DMPC a good starting point to look at control methods for demand response of hybrid systems.

4.3 Formulating objective functions

Coupling among local systems can be through the objective function, the dynamics, the constraints, or a combination of those. FC-DMPC requires a global objective function J_g . All TCLs are assumed to be dynamically uncoupled as seen in (4.1) and all constraints i.e. on the temperature of the appliance, are purely local. This means that the global objective function should be a function of at least one local variable from each TCL to include it in the FC-DMPC framework.

The way in which the global objective is formulated and translated to locally optimizable parts is an important part of the control system design. From an abstract point of view there are two main approaches to do this:

1. *Bottom-up*
Here, existing local objective functions of the individual subsystems are aggregated into a single global objective function.
2. *Top-down*
Here, a main global objective function is used as a starting point to formulate locally optimizable parts.

In this section both approaches will be reviewed in more detail with respect to the main objective as formulate in (4.3), highlighting possible shortcomings and opportunities.

4.3.1 Aggerating local objectives

In [37] it is stated that the simplest approach to formulate the global objective function is to take a weighted sum of the existing local objectives:

$$J_g = \sum_{n \in \mathcal{N}} w_n J_n(p_n), \quad w_n > 0, \quad \sum_{n \in \mathcal{N}} w_n = 1 \quad (4.5)$$

If we substitute the local objective function (3.2) that has been used in Chapter 3 in this expression this results in:

$$J_g(k) = \sum_{n \in \mathcal{N}} w_n \|\mathbf{p}_n^{\text{ref}}(k) - \hat{\mathbf{p}}_n(k)\|_1 \quad (4.6)$$

The resulting global objective function contains local power reference terms $p_n^{\text{ref}}(k)$. However, the considered system as introduced at the beginning of this chapter is only offered a *global* power reference signal $\mathbf{p}_g^{\text{ref}}(k)$. Therefore, to optimize this global objective function (4.6) locally, the global power reference must be decomposed into N separate references. This means that even when using a bottom-up approach a decomposition is required.

4.3.2 Decomposing a global objective

An alternative to aggregating the local objective functions is to start from the main objective function in (4.3) and to decompose it into locally optimizable parts. The decomposition of a global reference power in an L1 norm objective function has been addressed in [13]. In that paper, the aim is to reduce the required communication among the controllers. The following three methods are proposed:

1. Static local power references with no information
2. Dynamic local power references with local information.
3. Global references with global information.

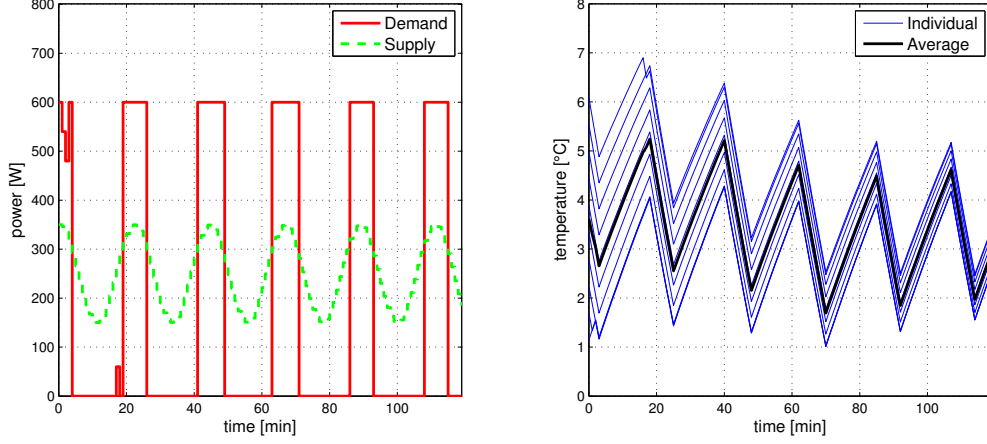


Figure 4.2: Synchronized dynamics of $N = 10$ appliances due to decentralized control resulting from a static decomposition of the global power reference

In this section these three decomposition approaches will be assessed with respect to the demand response problem of TCLs with binary inputs. It should be noted that this approach was originally proposed for systems with real inputs.

Static local power references

A first approach consists of dividing the global power reference into static local parts that are proportional to the average power demand \bar{p}_n of the individual loads:

$$p_n^{\text{ref}}(k) = \frac{\bar{p}_n}{\sum_{n \in \mathcal{N}} \bar{p}_n} \cdot \sum_{n \in \mathcal{N}} p_n^{\text{ref}}(k) \quad (4.7)$$

The result of the static decomposition of the global power reference is that the entire global objective function will become decoupled. The reason for this is that the variables $\hat{\mathbf{p}}_m(k)$ of the other controllers are assumed to be static over each iteration. This results in a constant term in the local optimization problems that can be omitted from the optimization:

$$\begin{aligned} p_n^*(k) &= \arg \min_{\hat{\mathbf{p}}_n(k)} \left(w_n \|\mathbf{p}_n^{\text{ref}}(k) - \hat{\mathbf{p}}_n(k)\|_1 + \sum_{m \in \mathcal{N} \setminus \{n\}} w_m \|\mathbf{p}_m^{\text{ref}}(k) - \hat{\mathbf{p}}_m(k)\|_1 \right) \\ &= \arg \min_{\hat{\mathbf{p}}_n(k)} \|\mathbf{p}_n^{\text{ref}}(k) - \hat{\mathbf{p}}_n(k)\|_1 \end{aligned} \quad (4.8)$$

In the considered case the static decomposition of the power reference will therefore defeat the object of FC-DMPC and essentially transforms it into a decentralized control scheme. In Chapter 3 it was shown that binary inputs pose significant limitations on the ability of a single TCL to track certain reference signals. Having multiple (identical) appliances with a decentralized control scheme will not improve these results. The appliances may even synchronize, resulting in the unstable oscillating behavior as seen in Figure 4.2.




Figure 4.3: Limited communication where a controller exchanges unidirectional information with its neighbors only

Dynamic local power references

An alternative approach is to let the appliances negotiate about the decomposition of the power reference. This is done by expanding the local objective function in (4.8) with terms that allow controller n to trade power references with its neighboring controllers in $\mathcal{N}_n \subset \mathcal{N}$, where we assume that all connections of n with its neighbors in \mathcal{N}_n are managed either by n or by one of its neighbors. The set of connections managed by n is denoted by Δ_n . This principle is illustrated in Figure 4.3 and formalized by:

$$\min_{p_n(k), \delta_{nm}(k)} \left\| \mathbf{p}_n^{\text{ref}}(k) + \sum_{m \in \Delta_n} \delta_{nm}(k) - \sum_{m \in \mathcal{N}_n \setminus \Delta_n} \delta_{mn}(k) - \hat{\mathbf{p}}_n(k) \right\|_1 \quad (4.9)$$

The authors of [13] argue that when the connectivity matrix is sparse, meaning that the number of neighbors in \mathcal{N}_n is smaller than the total number of systems in $\mathcal{N} \setminus \{n\}$, this method may lead to reduced communication requirements. However, with respect to the case of demand response of hybrid systems, two potential caveats can be identified:

1. *Undefined neighborhood sets*

In the case of demand response of a large group of appliances that are all connected to the internet, the set of neighbors \mathcal{N}_n of each appliance is not clearly defined. Indeed, all appliances are assumed to be able to connect to all other appliances regardless of spatial or structural factors. This is not necessarily a problem, as having the ability to connect to all other controllers does not require actually doing so. Instead a subset of neighbors can be chosen based on arbitrary selection criteria and performance may depend heavily on the choice of the neighboring set.

2. *Negotiations among controllers*

Apart from exchanging local inputs controllers communicate the optimal power reference terms δ_{nm} of the connections Δ_n that they manage. In the literature review it was found that existing




Figure 4.4: Global communication where every controller is able to exchange information with all other controllers

DMPC methods for hybrid systems that make use of negotiations on interconnecting variables may not converge due to big jumps in the negotiations whenever local controllers switch their inputs [5, 21]. The approach in [13] is different since every connection among controllers is managed by a unique controller. However, oscillations of the negotiations can still take place due to loops in the communication graph, i.e. three controllers that are connected in a circle can keep passing a certain undesired value around in that circle restricting the convergence.

Global reference

Assuming that all appliances can communicate to each other as seen in Figure 4.4 can give controllers access to global information. In our case this means that while local controllers optimize their own inputs they are aware of the global system state. Instead of decomposing the global objective into separate local objectives, global information also enables the use of the original main objective in (4.3) directly. Since the global power demand is equal to the sum of all local power demands the global objective function can be written as:

$$J_g(k) = \left\| \mathbf{p}_g^{\text{ref}}(k) - \sum_{n \in \mathcal{N}} \hat{\mathbf{p}}_n(k) \right\|_1 \quad (4.10)$$

Locally optimizing this function based FC-DMPC means solving:

$$\hat{\mathbf{p}}_n^i = \arg \min_{\hat{\mathbf{p}}_n(k)} \left\| \mathbf{p}^{\text{ref}}(k) - \hat{\mathbf{p}}_n^{i-1}(k) - \sum_{m \in \mathcal{N} \setminus \{n\}} \hat{\mathbf{p}}_m^{i-1}(k) \right\|_1 \quad (4.11)$$

While global communication within fully connected graphs might be the most straightforward approach to gain global information, alternative approaches are possible that can reduce the amount of required communication.

One of these is the introduction of a central blackboard coordinator that is used to aggregate the local demands. Each controller is then able to look up the current global state at this centralized agent instead of requesting information from all other controllers.

4.4 Parallel and serial update algorithms

In the majority of DMPC methods in the literature that have been applied to demand response, all controllers optimize their inputs in parallel [5, 21, 37]. This means that at every iteration i for each time step k all local controllers n update their input schedules synchronously. In this section the consequences of such using this method for the demand response of hybrid systems will be evaluated and alternatives will be evaluated.

Parallel updates

In the parallel approach all local controllers synchronously update their optimal inputs at every iteration i by solving a local MILP. After all controllers have updated their schedules they exchange the required information with other controllers. The parallel update algorithm is presented in Algorithm 1.

Algorithm 1 Parallel update scheme

```

1: initialization:  $p_g(k) = 0$ 
2: for each time step  $k$  do
3:   for each iteration  $i$  do
4:     for each subsystem  $n$  do
5:       solve MILP (4.11)
6:       broadcast required power  $\hat{p}_i^n(k)$ 
7:     end for
8:     update global power  $p_g(k)$  by aggregating all optimized local demands  $p_n(k)$ 
9:   end for
10:  for each TCL subsystem  $n$  do
11:    simulate one time step using optimal input
12:  end for
13: end for

```

Using the parallel implementation of Feasible-Cooperation DMPC with the global objective function in (4.11) results in some complications. At each iteration all controllers will assess the possibility to switch *on* or *off* based on the evaluation of:

$$\left\| \mathbf{p}_g^{\text{ref}}(k) - \sum_{n \in \mathcal{N}} \hat{\mathbf{p}}_n(k) \right\|_1 \quad (4.12)$$

When the difference between the two terms is too large, all controllers will, if their temperature constraints allow it, switch *on* or *off* at the same time. They do this because at every iteration they have the same information. At the next iteration this may lead to an overshoot in the objective function. This overshoot may force controllers to synchronously reconsider their switch, leading them back to their original input. This switching ultimately leads to non-convergence of the iterations and extremely poor performance of the DMPC implementation. This effect is illustrated in Figure 4.5a, here the open loop predictions resulting from all local optimizations oscillate over the iterations without converging to the reference power.

An alternative approach would be to make use of a serial implementation, in which local controllers update their inputs *one after the other*. In [28] the authors have shown that in some cases such an approach is beneficial with respect to a parallel one.

Serial updates

In a serial implementation the local controllers update their inputs in sequence. The global information on the current mismatch between supply and demand is updated and communicated after each local optimization. This scheme has been used in [22, 28].

The serial implementation of FC-DMPC is shown in Algorithm 2. The difference with Algorithm 1 can be found on line 7: using a serial approach the global power demand is updated after each controller updates its control inputs instead of only after *all* controllers have updated their inputs. This principle is also illustrated in Figure 4.5b, here the additional grey lines represent the aggregate planning after device n has updated its schedule. At the first iteration every controller makes its initial planning, the following iterations are used to make adjustments to the initial schedule. In this example the solution converged at the second iteration.

Algorithm 2 Serial update scheme

```

1: initialization:  $p_g(k) = 0$ 
2: for each time step  $k$  do
3:   for each iteration  $i$  do
4:     for each subsystem  $n$  do
5:       solve MILP (3.13)
6:       broadcast required power  $p_n^i(k)$ 
7:       update global power  $p_g(k)$  by aggregating all optimized local demands  $p_n(k)$ 
8:     end for
9:   end for
10:  for each TCL subsystem  $n$  do
11:    simulate one time step using optimal input
12:  end for
13: end for

```

Using a serial implementation of FC-DMPC to control the demand response of a population of TCLs described by a hybrid model is expected to yield the following advantages:

1. *Convergence of the global solution*

The serial scheme is likely to prevent the synchronization of controllers since at the start of each local optimization all controllers have different information, depending on their location in the update sequence. This may facilitate the convergence of the global solution.

2. *Reduced communication whilst maintaining the distributed nature*

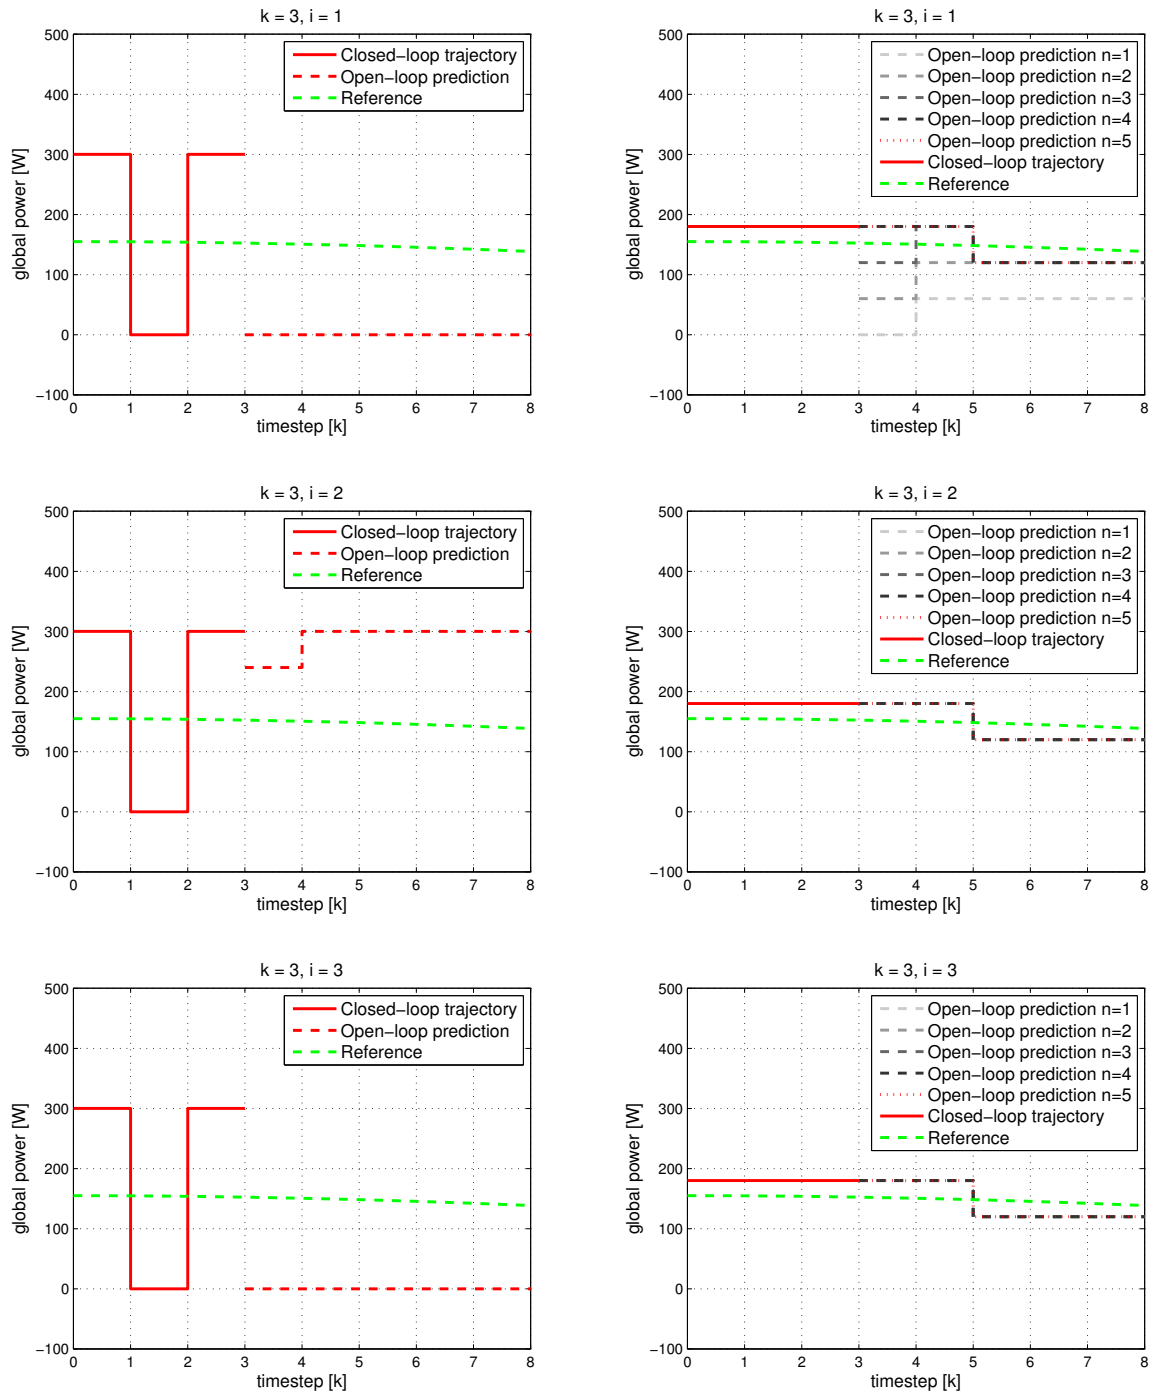
The order in which local controllers update their optimal inputs can be used as the basis to define their neighbors. Every controller can then pass its updated global information on to the next controller, omitting the need for global communication or a centralized blackboard agent to maintain the distributed properties of the controller.

In the next chapter both the serial and parallel implementation of FC-DMPC will be used for closed-loop simulations and compared to alternative methods.

4.5 Conclusions

Based on the findings in this chapter the following conclusions are drawn:

1. The main principle of FC-DMPC is that every controller locally optimizes a global objective function. In this chapter different approaches to formulate such a global objective function have been investigated. In our case, the approaches aimed at decomposing the global power reference into



(a) Parallel updates within each iteration

(b) Sequential updates within each iteration

Figure 4.5: Multiple iterations $i = 1, 2, 3$ within a single timestep ($k = 3$), the sequential scheme has converged at the second iteration iteration while the solution of the parrallel scheme oscillates. The grey open-loop predictions represent the aggregate demand once device n has optimized its schedule

local parts either result in decentralized control or are unlikely to converge due to the hybrid properties of the TCL models.

2. An alternative approach is to maintain a single global reference power and to support local optimization by providing global information locally. Two methods to achieve this are global communication or a centralized blackboard agent.
3. The parallel implementation of FC-DMPC with a coupled objective function to control the demand response of TCLs is likely to lead to synchronization of their actions and non-converging oscillations of the objective function value. The reason for this is that all controllers have access to the same global information indicating whether there is a global shortage or surplus of power. All controllers will therefore have the same incentive, resulting in an undesirable synchronization of the loads.
4. A serial implementation of this method will force controllers to update their optimal schedule *on at the time*. This is expected to yield two benefits: first, it facilitates the convergence of the objective function value and second it is able to relax the need for global communication whilst maintaining the distributed nature of the controller.

Chapter 5

Simulations and Comparison

Both the parallel as well as the serial FC-DMPC approach that have been developed in the previous chapter will be used for closed-loop simulations. The result of these simulations will be compared to a commonly used alternative method from the literature as well as two benchmark methods, based on selected performance criteria.

Two types of closed-loop simulations will be presented. They can be distinguished by the choice of the power reference supply signal:

- *Synthetic reference supply*
To compare the different controllers they will initially be tested using a synthetic reference signal that will yield feasible results with all methods.
- *Reference supply based on real data*
To assess the relevance and potential of the methods they will also be applied to a case study based on real data.

5.1 Simulated control methods

The control methods that will be used for closed-loop simulations can be divided into two main categories: benchmark methods and various alternative DMPC methods. The used benchmark methods are:

1. *Hysteresis control*
This is the original control law of a TCL based on a thermostat as described in Chapter 2. This decentralized method provides a performance reference that can be used to assess whether or not improvements have been made at all.
2. *Centralized MPC*
This centralized control method as described in Chapter 3 is based on a single large model that includes all system dynamics and all constraints. This approach should therefore lead to optimal performance with respect to the optimization of the global objective function. This provides a performance benchmark to assess the relative optimality of the DMPC methods.

The tested DMPC methods are:

1. *Dual-decomposition DMPC*
This commonly used DMPC method from the literature has already been to demand response problems concerning hybrid systems [5, 21]; in this case the implementation as found in [5] is used.
2. *Parallel FC-DMPC*
This is the parallel control approach with a global power reference and global information as presented in Chapter 4.
3. *Serial FC-DMPC*
Here we apply the serial control approach with a global power reference and global information as presented in Chapter 4.

5.2 Performance criteria

The control methods that are used to run the closed-loop simulations will be compared based on the following performance criteria:

1. *Global objective function values*
This means evaluating

$$J_g^{\text{cl}} = \sum_{k=1}^K \left| p^{\text{ref}}(k) - \sum_{n \in \mathcal{N}} p_n(k) \right| \quad (5.1)$$

in which K is the total number of time steps used in the closed-loop simulations.

2. *Computational speed*

The time required to compute the control inputs at each time step k using MATLAB¹ with help of the MPT 3.0 toolbox [18] and the GLPK solver² on a single machine³. A solution is reached when one of the following two stopping criteria is satisfied:

- (a) The solution has converged in the sense that the value of the global objective function does not change in the next iterations:

$$J_g^i = J_g^{i-1} \quad (5.2)$$

- (b) A maximum number of iterations is reached:

$$i = i^{\text{max}} \quad (5.3)$$

Distributed approaches in which multiple optimizations are performed in parallel may be limited by the fact that all computations are done using a single machine. To correct for this, the total computation time will be the sum of the maximum local optimization times per iteration.

3. *Communication requirements*

The amount of communication links that need to be initiated and the size of the transmitted data that is needed per time step k .

5.3 Synthetic reference supplies

Formulation of the reference power profile

Following up on the initial analysis for reachable reference powers in Chapter 3, both synthetic reference signals are constructed using an average power \bar{p}_g that is within the average power range of the TCL population and a certain deviation from that average. Here, the deviation is described by a cosine with period T and an amplitude a that scales with the size of the TCL population N :

$$p_g^{\text{ref}}(k) = \bar{p}_g + a \cdot N \cdot \cos\left(\frac{1}{T} 2\pi\right) \quad (5.4)$$

¹MATLAB Release 2013a, The MathWorks, Inc., Natick, Massachusetts, United States.

²GNU Linear Programming Kit, <http://www.gnu.org/software/glpk/>

³2 GHz Intel Core 2 Duo, 4 GB 1067 MHz DDR3

Apart from using a single cosine reference, an additional reference signal consisting of a sum of cosines is used, described by:

$$p_g^{\text{ref}}(k) = \bar{p}_g + \sum_{i=1}^5 a_i \cdot N \cdot \cos\left(\frac{1}{T_i} 2\pi\right) \quad (5.5)$$

Simulation parameters

For the single cosine power reference signal the prediction horizon and control horizon will be equal to the values that have been used in the previous chapters i.e. $N_p = N_c = 5$. The reason for this is that the number of integer variables in the MILP for centralized MPC increases with the amount of appliances. For populations larger than $N = 5$ the problem becomes too large to solve on the available machine and therefore this population size is used for the initial comparison. The second power reference signal consisting of the sum of cosines is used with a larger population size of $N = 10$ appliances. The reason for this is that complex reference signals need larger TCL populations to decrease the relative error due to their binary inputs, as seen in Chapter 3.

The simulation time step size is $t_s = 60$ s. The total number of simulation time steps is $K = 180$. This is sufficiently large to encompass multiple oscillations of the power reference signal and to reduce the effect of the initialization.

The initial conditions of the appliances are equal for all simulations. The initial temperatures of the appliances are evenly spread on the interval $[\theta^{\min}, \theta^{\max}]$.

5.3.1 Performance evaluation

Value function

The closed-loop power dynamics that are presented in Figure 5.2 show that some of the DMPC methods result in a highly variable power demand. The results from dual-decomposition DMPC (Figure 5.1c) and parallel FC-DMPC (Figure 5.1d) show large deviations from the optimal power trajectory defined by the centralized MPC solution presented in Figure 5.1b. Moreover, their performance appears to be even worse than the reference performance of hysteresis control (Figure 5.1a).

Among the simulated DMPC methods the only method that appears to converge towards the centralized performance is serial FC-DMPC. As suggested in Chapter 4, one of the main drivers for this convergence is the fact that the local controllers have access to updated global information whenever they optimize their inputs. Since no controller optimizes at the same time as the others the oscillations as seen in Figure 5.1c and Figure 5.1d can be prevented.


These findings are supported by the simulation results in Figure 5.3. Although there is no centralized MPC solution available⁴ to assess its relative optimality, the serial FC-DMPC method appears to be capable of tracking the complicated reference supply very well.

The overall performance indices as defined in (5.1) are given in Table 5.1. In this table it can also be observed that using the cosine reference the serial FC-DMPC actually performs as well as centralized MPC optimal benchmark.


Table 5.1: Performance indices per method and per simulation

	cosine reference	sum of cos. reference
Hysteresis control	7211	17367
Centralized MPC	2770	-
Dual-decomposition DMPC	21851	34400
Parallel FC-DMPC	22260	23104
Serial FC-DMPC	2770	6761


⁴due to the large time required to compute such a solution




(a) Hysteresis Control




(b) Centralized MPC



(c) Dual-decomposition DMPC




(d) Parallel FC-DMPC




(e) Serial FC-DMPC


Figure 5.2: Temperature dynamics of a population of $N = 5$ TCLs when subject to a cosine reference supply




(a) Hysteresis control




(b) Centralized DMPC



(c) Dual-decomposition DMPC




(d) Parallel FC-DMPC




(e) Serial FC-DMPC


Figure 5.3: Power demand of a population of $N = 10$ TCLs when subject to a reference that is a sum of cosines




(a) Hysteresis Control




(b) Centralized MPC



(c) Dual-decomposition DMPC



(d) Parallel FC-DMPC



(e) Serial FC-DMPC

Figure 5.4: Temperature dynamics of a population of $N = 10$ TCLs when subject to a reference that is a sum of cosines




Figure 5.5: Relative optimality of various DMPC solutions with respect to the centralized MPC solution over multiple iterations

Table 5.2: Computation times to reach a solution per time step using the cosine reference with $N = 5$ appliances

	average time [s]	maximum time [s]	standard deviation [s]
Hysteresis control	0.0007	0.0082	0.0007
Centralized MPC	0.9992	7.6239	1.1449
Dual-decomposition DMPC	0.1513	0.1433	0.0082
Parallel DMPC	0.0238	0.0490	0.0040
Serial DMPC	0.0298	0.0524	0.0062

The relatively poor performance of dual-decomposition DMPC and parallel FC-DMPC is due to the non-convergence of the global objective function value over the iterations. In Figure 5.5 the relative optimality of various DMPC solutions with respect to the centralized MPC solution is presented over multiple iterations. The relative optimality of the DMPC solution J_g^{dmPC} with respect to the centralized MPC solution J_g^{mPC} for multiple iterations i is computed using:

$$\frac{J_g^{\text{dmPC}}(i) - J_g^{\text{mPC}}}{J_g^{\text{mPC}}} \cdot 100\% \quad (5.6)$$

The values have been averaged using data from all time steps. Whilst serial FC-DMPC converges to this optimal solution within a maximum of 3 iterations, both dual-decomposition DMPC and parallel FC-DMPC do not.

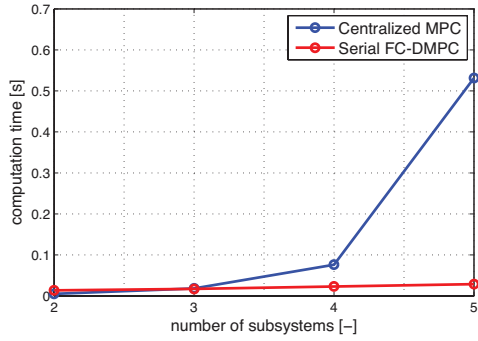
Computational speed

Here we compare the times in which the various methods reach a stopping criterion, i.e. convergence of the global objective function or a maximum number of iterations. Since centralized MPC only obtained a solution using a small population size of $N = 5$ appliances, only the results using that small population reference will be used. The mean value, the maximum value, and standard deviation of the computation times at each time step k are evaluated. The results are presented in Table 5.2 and Table 5.3.

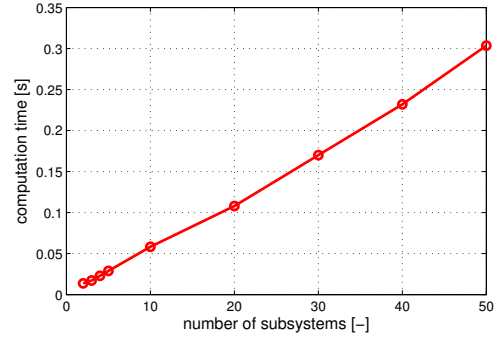
It can be observed that the required computation time for all methods stays well below the sample time of $t_s = 60$ s that has been used in the simulations. The relatively quick computation times result from the relatively small population size that is required to get the centralized MPC solution. However, Table ?? shows that at $N = 5$ the centralized MPC approach already takes significantly more time to reach a solution.

Table 5.3: Computation times to reach a solution per time step using the sum of cosines reference with $N = 10$ appliances

	average time [s]	maximum time [s]	standard deviation [s]
Hysteresis control	0.0010	0.0037	0.0004
Centralized MPC	-	-	-
Dual-decomposition DMPC	0.1921	0.5487	0.0303
Parallel DMPC	0.0224	0.0676	0.0042
Serial DMPC	0.0579	0.1149	0.0128



(a) Centralized MPC vs. serial FC-DMPC



(b) Serial FC-DMPC for larger populations

Figure 5.6: Average time of computation as a function of the number of subsystems

The required computation time depends partly on the size of the MILP that is solved by each controller. A larger amount of integer variables generally results in an MILP that takes longer to solve. For any non-centralized method the amount of integer variables depends solely on the control horizon N_c . For centralized MPC however, this amount is multiplied by the amount of TCLs that are controlled.

For parallel DMPC methods all local MILPs are solved at the same time. The total required computation time is therefore a function of the amount of required iterations only. Since, as seen in Figure 5.5, the value of the global objective function does not converge using these methods, they require the maximum amount of iterations i^{\max} .

For the serial DMPC method the total required computation time also depends on the amount of TCLs, since all local computations are done in sequence.

Since only the solutions of centralized MPC and serial FC-DMPC converge, it is interesting to evaluate how their performance depends on the TCL population size. As mentioned earlier, for population sizes larger than $N = 5$ the centralized MPC optimization problem becomes intractable to solve on the machine used. Figure 5.6a shows the sharp increase in the required computation time for smaller populations. For populations larger than $N = 3$ the serial FC-DMPC method is faster than the centralized MPC method. Moreover, the method is capable of handling much larger populations as the computation time increases linearly, as seen in Figure 5.6b.

Communication

All of the DMPC methods that have been simulated are based on global information that is locally available. The different ways to achieve this have consequences for the required amount of communication. The general results are shown in Table 5.4

The dual-decomposition method as proposed in [5] is based on a centralized pricing agent that determines the price of the interconnecting variables. Each local controller therefore needs to send its optimal inputs and to retrieve the updated prices at every iteration.

Completely distributed parallel FC-DMPC with global communication requires $2 \cdot N \cdot (N - 1)$ communication links to send and retrieve information to and from all other controllers. A centralized coordinating blackboard agent can reduce this to a scheme similar as dual-decomposition DMPC.

As mentioned earlier in Chapter 4, in serial FC-DMPC each controller can pass on the updated

Table 5.4: Communication requirements per time step k

	connections [-]
Hysteresis control	0
Centralized MPC	$2 \cdot N$
Dual-decomposition DMPC	$2 \cdot N \cdot I$
Parallel FC-DMPC	$2 \cdot N \cdot I$
Serial FC-DMPC	$N \cdot I$

global information to the next controller in the updating sequence, requiring N connections within each iteration.

Although the centralized MPC method does not require communication among the controllers to coordinate their actions, all TCL temperatures need to be sent to the central controller and all control inputs have to be sent back to the appliances, this still requires a total of $2 \cdot N$ communication links.

The type of information can be seen as a vector with a length N_p in which each entry is the forecasted local/global power demand per time step, represented by a float (generally 4 bytes). In the case that local temperatures need to be communicated a single float is sufficient.

5.4 Case study with real data

In the previous section it has been shown that serial DMPC is the only investigated distributed MPC method that is able to converge to globally optimal performance whilst handling larger population sizes. In this section the serial FC-DMPC approach is applied to a practical case study using real data.

The aim of this case study is to smoothen the net available supply of the generated wind electricity in Belgium by controlling the aggregate demand of a population of TCLs. The net available supply of electricity is defined as the available wind energy minus the power demand of the TCLs. This is favorable for two reasons:

1. The TCLs, which constitute a major part of the total energy demand, will make optimal use of the available sustainable energy.
2. A smoother net available power supply from renewable generation will facilitate its integration in the energy mix. Currently sharp increases or decreases in the sustainable energy supply require fossil fueled power stations to adjust abruptly to compensate for these fluctuations, causing huge expenses, inefficiencies, and sometimes even blackouts. By smoothening the sustainable energy supply, conventional fossil fueled power plants do not have to ramp up or down so abruptly, improving overall efficiency and reducing costs.

Data

The available dataset⁵ contains the measured electricity generation using wind energy in Belgium in time intervals of 15 minutes. A higher resolution of 180 seconds is obtained by interpolating the available data. For the closed-loop simulations a data set of 120 hours is used based on measurements between the 5th and the 10th of October.

In 2013, Belgium had approximately 11 million inhabitants living in 4.5 million households. Assuming that every household has 1 refrigerator and 1 freezer on average, and approximately every 10 people share a refrigerator at work, this sums up to a total of around 10 million appliances. This means that they have an average power demand of approximately:

$$N \cdot \bar{p}_n = 10.000.000 \cdot 20 \text{ W} = 200 \text{ MW} \quad (5.7)$$

⁵GRID DATA, Elia System Operator NV, <http://www.elia.be/nl/grid-data/productie/windproductie>

Formalizing the case structure

Initially, the net available generated wind energy is equal to the measured value minus the power demand of the TCLs

$$\mathbf{p}^{\text{net}} = \mathbf{p}^{\text{wind}} - \mathbf{p}_{\text{g}} \quad (5.8)$$

The measured wind energy is augmented by using a moving average with a span of 24 hours, yielding a desired smooth signal. This is illustrated in Figure 5.7. The remaining deviation between the smoothed signal and the original measured values is expressed as an error, (5.8) can then be rewritten into:

$$\mathbf{p}^{\text{net}} = \mathbf{p}^{\text{smoothed}} + \mathbf{p}^{\text{error}} - \mathbf{p}_{\text{g}} \quad (5.9)$$

The error should be matched with the aggregate demand of the TCLs. However, as seen in Chapter 3 a reachable power supply should have a long term average equal to the average power demand. Therefore the average power demand is added to the error signal to obtain a reachable power reference supply, this principle is also illustrated in Figure 5.8:

$$\mathbf{p}^{\text{ref}} = \mathbf{p}^{\text{error}} + \bar{\mathbf{p}}_{\text{g}} \quad (5.10)$$

The objective of the controller is to match this reference power supply profile by controlling the global demand \mathbf{p}_{g} . The resulting mismatch $\delta\mathbf{p}$ is:

$$\delta\mathbf{p} = \min_{\mathbf{p}_{\text{g}}}(\mathbf{p}^{\text{ref}} - \mathbf{p}_{\text{g}}) \quad (5.11)$$

If the controller is able to match the reference power perfectly, combining (5.8), (5.10), and (5.11) results in an alternative expression for the net available wind energy, in which the real energy has been smoothed by the controlled power demand of the population of TCLs:

$$\mathbf{p}^{\text{wind}} = \mathbf{p}^{\text{smoothed}} - \bar{\mathbf{p}}_{\text{g}} - \delta\mathbf{p} \quad (5.12)$$

5.4.1 Performance evaluation

The closed-loop simulation results are shown in Figure 5.9. Figure 5.9a shows that the aggregated TCL demand is able to match the available reference power supply. Figure 5.9b shows that in some cases the average temperature moves towards on of the temperature constraints. As previously seen in Chapter 3 this will eventually limit the ability to match a certain power supply. Figure 5.9c shows the error resulting from the serial FC-DMPC method and the benchmark hysteresis approach. The practical implementation of the controller for demand response will always be a trade-off between the following factors:

1. The smoothness of the net power supply. Reducing the span of the moving average will ‘flatten’ the power reference for the population of TCLs, making it easier to match. The downside to this approach is that a smaller smoothing span will result in a rougher net power supply and therefore




Figure 5.7: Measured data and the desired smoothed net power supply

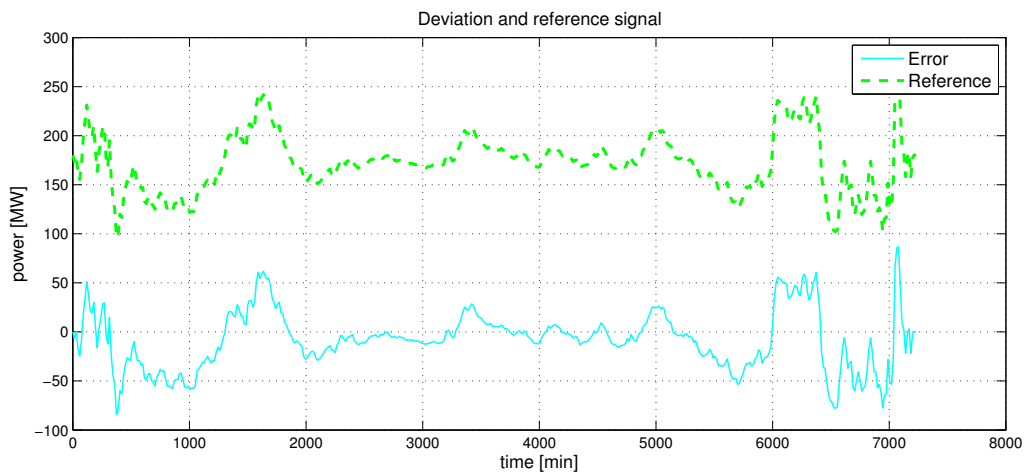





Figure 5.8: Creating a reachable power reference by adding the average power demand to the error



(a) Simulated closed-loop power dynamics for the data-based case study with 100 TCLs



(b) Simulated temperature dynamics using the serial FC-DMPC approach



(c) Error defined by the mismatch between supply and demand

Figure 5.9: Closed-loop simulation results for the Belgian wind case study

contribute less to the overall case objective of smoothing the sustainable power supply in the first place.

2. Softening the temperature constraints to increase the ability to match a certain power supply. This can be done by reformulating these hard constraints via a penalty function and adding them to the objective function that is solved in the MILP. The downside to this approach is that it may violate comfort preferences of the end users.
3. Changing the composition of the population of appliances. If the changes in the power reference signal are relatively slow, appliances with a larger heat capacity might be able to provide better results since a large heat capacity results in relatively slow temperature dynamics. Examples of such appliances are room heating systems, air conditioning systems and electric water heaters.

5.5 Conclusions

In this chapter the FC-DMPC implementations that have been developed in the previous chapter have been used for closed-loop simulations. The simulation results are compared to those using conventional hysteresis control, centralized MPC, and dual-decomposition DMPC, a commonly used alternative from the literature. Performance criteria are selected to the relative optimality with respect to the centralized solution, the required computation time, and the communication requirements. Based on the findings in this chapter the following conclusions are drawn:

1. The serial FC-DMPC implementation appears to be capable of converging to the centralized performance within a small number of iterations. In the study, dual-decomposition methods as proposed in [5] as well as parallel FC-DMPC do not converge. They result in an extremely variable power demand and poor performance with respect to the centralized optimal solution.
2. With respect to the computation speed the standard hysteresis controller clearly outperforms all other methods, but it lacks any form of coordination. The DMPC methods that do not converge stop when they reach the stopping criterion of a maximum number of iterations. For a population size larger than 3, serial DMPC outperforms centralized MPC due to the increasing amount of integer input variables in the latter method. For population size larger than 5 the amount of integer input variables in the centralized MPC optimization even results in a problem that is intractable to solve on the machine used, while the required computation time of serial FC-DMPC scales linearly with the size of the population.
3. With respect to the required communication the decentralized and centralized approaches outperform the DMPC methods. While in our case dual-decomposition DMPC and parallel FC-DMPC require a centralized agent to obtain global information, serial FC-DMPC can reduce the communication requirements whilst maintaining the distributed nature.
4. Serial FC-DMPC has also been used in a case study based on real data. The objective of the case study was to smoothen the available power supply of electricity from wind generation. Closed-loop simulation have shown that a population of TCLs can effectively be used to smoothen net power supply of wind generation.

Chapter 6

Conclusions and Future Research

6.1 Conclusions

To keep the conventional grid stable the electricity generation is continuously adjusted to match the demand. Despite the advantages of renewable generation, weather factors like wind and sun cannot be adjusted. To facilitate the integration of these renewable sources a paradigm shift is required: the electricity demand should respond to the available supply. Thermostatically controlled loads (TCLs) like refrigerators and electric water heaters are well suited for this, due to their thermal buffer ability. The work in this thesis has aimed at the development of a distributed model predictive control (DMPC) approach for a large-scale population of these loads. The objective of the controller is to control the aggregate power demand of the appliances to match the available renewable power supply.

A hybrid model of a TCL was developed to synthesize a model predictive controller for a single appliance. Closed-loop simulations using this controller show that the binary input of the TCL severely restricts the ability to track certain reference signals. Moreover, if the power supply profile is structurally different than the power demand, the controller is unusable due to local temperature constraints. Three power profile characteristics that the supply and demand should have in common have been identified: the average power, the power variation, and the periodicity of the power profile.

By controlling multiple TCLs the relative error due to the binary inputs can be reduced. However, the hybrid character of the prediction models may result in big jumps in the negotiations among controllers possibly leading to non-converging oscillations in the optimization. To overcome this a serial feasible-cooperation DMPC (FC-DMPC) approach has been proposed.

The serial FC-DMPC implementation has been compared with reference approaches as well as alternative DMPC methods from the literature. Performance criteria were the relative optimality, the computational burden, and the required amount of communication.

Results from closed-loop simulations show that unlike parallel DMPC methods, serial FC-DMPC of the TCLs can converge towards the centralized MPC solution without making use of additional heuristic stopping criteria that stop oscillations. The reason for this is that at every iteration each controller has access to different global information.

The required computation time for serial FC-DMPC increases linearly with the amount of controlled appliances. Although this is disadvantageous with respect to parallel implementations, it easily outperforms centralized MPC for larger population as the centralized optimization problem quickly becomes intractable to solve.

Serial FC-DMPC also discards the need for global communication or a centralized coordinating agent to obtain global information, since global up-to-date information can be passed around from one controller to the other in sequence.

6.2 Discussion and future research

The presented closed-loop simulations using the serial FC-DMPC implementation show favorable results with respect to convergence, optimality, and the computational burden. In this section the limitations of the approach will be discussed and promising directions for future research will be identified.

One of the findings in this work was that not all given power supply profiles can be matched by a any set of appliances, e.g. the daily fluctuations in the solar energy supply cannot be matched by small refrigerators that have to turn on approximately every hour to avoid a violation of the temperature constraints. Basic requirements for a matchable supply and demand have been investigated but further research towards a characterization of the demand response potential of certain TCL populations is recommended.

In this work relatively short prediction/control horizons have been used to enable fast simulation times and fast computation times. Larger prediction horizons will allow TCLs to anticipate more on future power surpluses and/or shortages. There are various possibilities to extend these horizons. One of them is time-instant optimization MPC as proposed in [36], which converts the MILP in a continuous nonlinear optimization problem that can be solved using multi-start pattern search algorithms. The aim of this approach is to optimize for the real-valued time instants at which the input switches *on/off* instead of optimizing the binary input at each fixed time step up to the prediction horizon.

One of the assumptions that has been made in this thesis is that all systems are completely deterministic and that their states are fully observable. However, there might be additional disturbances or stochastic factors that influence the system dynamics, i.e. when a user opens a refrigerator to put a lot of beer inside or when a user takes an exceptionally long shower. Including these stochastic processes in the control approach, or taking a top-down stochastic control approach might enable new possibilities.

In the serial DMPC approach that has been used in this thesis all appliances update their schedules in a fixed order. It would be interesting to investigate whether changing the order of updates influences the performance. The scalability of a serial approach for extremely large populations (thousands to millions) is also open to discussion since all appliances have to wait for their turn to update their inputs. It would be interesting to see whether a hierarchical approach involving a combination of parallel and serial updates is able to converge for large-scale populations whilst reducing the required computation time.

From a long-term perspective it would be interesting to do more fundamental work towards finding DMPC methods for hybrid systems that can be guaranteed to converge to optimal performance and can be generalized to work in a wider range of applications. Examples of such alternative applications can be process installations with a large amount of valves with binary inputs, large networks of LEDs for lighting purposes, a system of waterways with distributed locks/slucices, or traffic networks with ramp metering.

Bibliography

- [1] S.M. Amin and B.F. Wollenberg. Toward a smart grid: power delivery for the 21st century. *IEEE Power and Energy Magazine*, 3(5):34–41, 2005.
- [2] N. Baghina, I. Lampropoulos, B. Asare-Bediako, W.L. Kling, and P.F. Ribeiro. Predictive control of a domestic freezer for real-time demand response applications. In *Proceedings of the 3rd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies (ISGT Europe '12)*, pages 1–8, 2012.
- [3] P. Bertoldi, B. Hirl, and N. Labanca. Energy efficiency status report. Technical report, EU-27 JRC Institute for Energy and Transport, 2012.
- [4] T. Bigler, G. Gaderer, P. Loschmidt, and T. Sauter. Smartfridge: Demand side management for the device level. In *Proceedings of the 16th IEEE Conference on Emerging Technologies Factory Automation (ETFA '11)*, pages 1–8, Sept. 2011.
- [5] R. Bourdais and H. Guéguen. Distributed predictive control for complex hybrid system, the refrigeration example. In *Proceedings of the 12th IFAC Symposium on Large Scale Complex Systems: Theory and Applications*, pages 388–394, 2010.
- [6] D.S. Callaway and I.A. Hiskens. Achieving controllability of electric loads. *Proceedings of the IEEE*, 99(1):184–199, 2011.
- [7] E. Camponogara, D. Jia, B.H. Krogh, and S. Talukdar. Distributed model predictive control. *IEEE Control Systems Magazine*, 22(1):44–52, 2002.
- [8] P.D. Christofides, R. Scattolini, D. Muñoz de la Peña, and J. Liu. Distributed model predictive control: A tutorial review and future research directions. *Computers & Chemical Engineering*, 51:21–41, 2013.
- [9] Commission of the European Communities. Renewable energy road map. Technical report, 2007.
- [10] Commission of the European Communities. Renewable energy: Progressing towards the 2020 target. Technical report, 2010.
- [11] B. De Schutter, W.P.M.H. Heemels, J. Lunze, and C. Prieur. Survey of modeling, analysis, and control of hybrid systems. In J. Lunze, F. Lamnabhi-Lagarrigue, and eds., editors, *Handbook of Hybrid Systems Control Theory, Tools, Applications*, pages 31–55. Cambridge, UK: Cambridge University Press, 2009.
- [12] K. Dietrich, J.M. Latorre, L. Olmos, and A. Ramos. Demand response in an isolated system with high wind integration. *IEEE Transactions on Power Systems*, 27(1):20–29, 2012.
- [13] M.D. Doan, P. Giselsson, T. Keviczky, B. De Schutter, and A. Rantzer. A distributed accelerated gradient algorithm for distributed model predictive control of a hydro power valley. *Control Engineering Practice*, 21(11):1594 – 1605, 2013.
- [14] Energy Information Administration. Residential energy consumption survey. Technical report, U.S. Dept. Energy, Washington, DC, 2001.

- [15] P. Giselsson, M.D. Doan, T. Keviczky, B. De Schutter, and A. Rantzer. Accelerated gradient methods and dual decomposition in distributed model predictive control. *Automatica*, 49(3):829–833, 2013.
- [16] P. Giselsson and A. Rantzer. Distributed model predictive control with suboptimality and stability guarantees. In *Proceedings of the 49th IEEE Conference on Decision and Control (CDC '10)*, pages 7272–7277, 2010.
- [17] M. Heidarinejad, J. Liu, and P.D. Christofides. Distributed model predictive control of switched nonlinear systems with scheduled mode transitions. *AIChE Journal*, 59(3):860–871, 2013.
- [18] M. Herceg, M. Kvasnica, C.N. Jones, and M. Morari. Multi-Parametric Toolbox 3.0. In *Proc. of the European Control Conference*, pages 502–510, Zürich, Switzerland, July 17–19 2013. <http://control.ee.ethz.ch/~mpt>.
- [19] International Energy Agency. World energy outlook 2011. Technical report, 2011.
- [20] S. Koch, J. L Mathieu, and Callaway D. S. Modeling and control of heterogeneous thermostatically controlled loads for ancillary services. In *Proceedings of the 17th Power Systems Computation Conference (PSCC '11)*, pages 961–968, 2011.
- [21] G.K.H. Larsen, N.D. van Foreest, and J.M.A. Scherpen. Distributed control of the power supply-demand balance. *IEEE Transactions on Smart Grid*, 4(2):828–836, 2013.
- [22] J. Liu, X. Chen, D. Muñoz de la Peña, and P.D. Christofides. Sequential and iterative architectures for distributed model predictive control of nonlinear process systems. *AIChE Journal*, 56(8):2137–2149, 2010.
- [23] T. Logenthiran, D. Srinivasan, and Tan Zong Shun. Demand side management in smart grid using heuristic optimization. *IEEE Transactions on Smart Grid*, 3(3):1244–1252, Sept. 2012.
- [24] J.L. Mathieu and D.S. Callaway. State estimation and control of heterogeneous thermostatically controlled loads for load following. In *Proceedings of the 45th Hawaii International Conference on System Science (HICSS '12)*, pages 2002–2011, 2012.
- [25] J.L. Mathieu, S. Koch, and D.S. Callaway. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Transactions on Power Systems*, 28(1):430–440, 2013.
- [26] A. Mohsenian-Rad, V.W.S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Transactions on Smart Grid*, 1(3):320–331, dec. 2010.
- [27] R.R. Negenborn. *Multi-Agent Model Predictive Control with Applications to Power Networks*. PhD thesis, Delft University of Technology, 2007.
- [28] R.R. Negenborn, B. De Schutter, and H. Hellendoorn. Multi-agent model predictive control for transportation networks: Serial versus parallel schemes. *Engineering Applications of Artificial Intelligence*, 21(3):353–366, April 2008.
- [29] A. Rantzer. Dynamic dual decomposition for distributed control. In *Proceedings of the American Control Conference (ACC '09)*, pages 884–888, 2009.
- [30] J.B. Rawlings. Tutorial overview of model predictive control. *IEEE Control Systems Magazine*, 20(3):38–52, 2000.
- [31] J.B. Rawlings and B.T. Stewart. Coordinating multiple optimization-based controllers: New opportunities and challenges. *Journal of Process Control*, 18(9):839 – 845, 2008.
- [32] S. Salinas, M. Li, and P. Li. Multi-objective optimal energy consumption scheduling in smart grids. *IEEE Transactions on Smart Grid*, 4(1):341–348, March 2013.
- [33] R. Scattolini. Architectures for distributed and hierarchical model predictive control – a review. *Journal of Process Control*, 19(5):723–731, 2009.

- [34] J.A. Short, D.G. Infield, and L.L. Freris. Stabilization of grid frequency through dynamic demand control. *IEEE Transactions on Power Systems*, 22(3):1284–1293, 2007.
- [35] B.T. Stewart, A.N. Venkat, J.B. Rawlings, S.J. Wright, and G. Pannocchia. Cooperative distributed model predictive control. *Systems & Control Letters*, 59(8):460–469, 2010.
- [36] H. van Ekeren, R.R. Negenborn, P.J. van Overloop, and B. De Schutter. Time-instant optimization for hybrid model predictive control of the Rhine-Meuse delta. *Journal of Hydroinformatics*, 15(2):271–292, 2013.
- [37] A.N. Venkat, I.A. Hiskens, J.B. Rawlings, and S.J. Wright. Distributed mpc strategies with application to power system automatic generation control. *IEEE Transactions on Control Systems Technology*, 16(6):1192–1206, 2008.
- [38] Y. Zong, D. Kullmann, A. Thavlov, O. Gehrke, and H.W. Bindner. Application of model predictive control for active load management in a distributed power system with high wind penetration. *IEEE Transactions on Smart Grid*, 3(2):1055–1062, 2012.
- [39] Y. Zong, L. Mihet-Popa, D. Kullmann, A. Thavlov, O. Gehrke, and H.W. Bindner. Model predictive controller for active demand side management with pv self-consumption in an intelligent building. In *Proceedings of the 3rd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies (ISGT Europe '12)*, pages 1–8, 2012.

