Load scheduling for decentralized CHP plants

Henrik Aalborg Nielsen, Torben Skov Nielsen, and Henrik Madsen

Informatics and Mathematical Modelling
Technical University of Denmark
Lyngby
June 2001

IMM

This document was prepared with LaTeX and printed by IMM, DTU, Lyngby.

Acknowledgements

First of all the authors would like to thank the Danish Energy Agency for financial support of this project under contract 1323/98-0025, Danish Energy Research Program. Also many thanks to Sønderborg Kraftvarmeværk I/S, Sønderborg and Energi- og Miljødata, Aalborg.

Special thanks to Jens Christian Hansen, Leo Thams, and Tue G. Christensen, from Sønderborg Kraftvarmeværk for supplying the practical example used in the report, to Anders N. Andersen, Energi- og Miljødata who guided us using his wide experience within the field of optimal operation of decentralized power plants, and to Steffen Moe, Sønderborg Fjernvarme for inspiring discussions. We would like to thank the following persons

Anders N. Andersen, Energi- og Miljødata
Bente Andersen, Farum Fjernvarme
Egon Erlandsen, Frederiksberg Varmeværk
Helge Faurby, Danmarks Meteorologiske Institut
Henrik R. Hansen, Vestegnens Kraftvarmeselskab
Jens Chr. Hansen, Sønderborg Kraftvarmeværk
Jens H. Hansen, Rambøll
Klaus W. Hansen, ELSAM
Dorthe R. Jensen, Centralkommunernes Transmissionsselskab
Gert Jensen, Frederiksberg Varmeværk
Oluf Jørgensen, Esbjerg Kommune
Peter Laursen, ABB Energi & Industri
Steffen Moe, Sønderborg Fjernvarme
Hans Ravn, Elkraft
Jan Strømvig, Odense Fjernvarme
Jesper Thiesen, Danmarks Meteorologiske Institut

for their contribution to the project and for active participation in a meeting, held in Spring 2000, concerning the subject of this report. At the meeting, approaches to load scheduling of decentralized combined heat and power plants were discussed. The approach used in this report reflects this discussion.

Summary

This report considers load scheduling for decentralized combined heat and power plants where the revenue from selling power to the transmission company and the fuel cost may be time-varying. These plants produce both heat and power with a fixed ratio between these outputs. A heat storage facility is used to be able to deviate from this restriction.

The load scheduling must be performed with only approximate knowledge about the future. At present in Denmark this uncertainty is only associated with the heat demand, but in the future revenues of produced energy and the fuel costs might also be uncertain and dependent on time. It is suggested to use a combination of background knowledge of the operator and computer tools to solve the scheduling problem. More specificly it is suggested that the plant is equipped with (i) an automatic on-line system for forecasting the heat demand, (ii) an interactive decision support tool by which optimal schedules can be found given the forecasts or user-defined modifications of the forecasts, and (iii) an automatic on-line system for monitoring when conditions have changed so that rescheduling is appropriate. In this report the focus is on methods applicable for items (ii) and (iii). For item (i), see (Madsen & Nielsen 1997, Nielsen & Madsen 2000).

The approach taken in this report is explicitly to describe how the total revenue from running the plant depends on the schedule for the heat and power producing units of the plant. Hereafter optimization theory, in this case dynamic programming, is applied to find the optimal schedule. To take the uncertainties into account it might be considered to use stochastic dynamic programming. However, it is argued that this is unpractical because the forecasting system will need to be integrated into the optimization system, whereby a modular design of the software cannot be obtained. Furthermore, we believe that all relevant forecasting methods are far too complicated to allow for this integration; both uncertainties originating from the dependence of heat load on climate and from meteorological forecasts need to be taken into account.

Instead we suggest that the decision support system allows the operator to investigate the sensitivity of the optimal schedule to variations in the input. Furthermore, we suggest that the system is equipped with the possibility to simulate realistic realizations of the heat demand based on the actual forecast and previous forecast errors. By letting the system find optimal schedules for each of these realizations the operator can gain some insight into the importance of the uncertainties.

It is shown that with modern personal computers (e.g. 1 GHz Pentium III), operating systems (e.g. RedHat Linux 6.0), and compilers (e.g. GNU C 2.91) the calculations can be performed quickly enough to allow use to be applicable in practice. One optimal schedule covering one week can easily be found within 5 to 10 seconds. When considering many possible realizations of the future heat demand some techniques are needed to reduce the amount of CPU time required. The results indicate that it is possible to find optimal schedules for 100 realizations of heat demand using less than 3 minutes of CPU time. Furthermore, the methods allow for massive use of parallel processing.

Resumé

Nærværende rapport beskæftiger sig med driftplanlægning for decentrale kraftvarmeværker, hvor brændselsudgiften og indtægten ved salg af elektricitet til transmissionsselskabet kan være tidsafhængig. For sådanne værker er forholdet mellem el- og varmeproduktion konstant, og et varmelager anvendes for at kunne afvige fra denne begrænsning over kortere perioder.

Planlægningen af driften må foretages ud fra delvis viden om fremtiden. På nuværende tidspunkt er denne usikkerhed kun til stede for såvidt angår varmebehovet, men i fremtiden kan brændselsudgifter og indtægten stammende fra produceret energi også være usikre og tidsafhængige. Det foreslås at løse planlægningsproblemet vha. en kombination af EDB-værktøjer og baggrundsvinde fra værkets operatør. Specifikt foreslås det, at et værk udstyres med (i) et automatisk on-line system til forudsigelse af varmebehovet, (ii) et interaktivt beslutningsstøttesystem vha. hvilket optimale driftforløb kan findes, givet disse forudsigelser eller bruger-bestemte modifikationer af disse forudsigelser, og (iii) et automatisk on-line system til overvågning af hvornår drift-betingelserne er ændret så meget at, driftplanen bør revideres. I denne rapport fokuseres der på metoder til anvendelse i forbindelse med punkterne (ii) og (iii). For punkt (i) henvises til (Madsen & Nielsen 1997, Nielsen & Madsen 2000).

Fremgangsmåden anvendt i denne rapport er eksplicit at beskrive, hvordan værkets totale indtjening afhænger af driftplanen for de el- og varmeproducerende enheder på værket. Herefter anvendes optimeringsteori, i dette tilfælde dynamisk programmering, for at finde en optimal driftplan. For at tage hensyn til usikkerhederne kan det overvejes at anvende stokastisk dynamisk programmering. I rapporten argumenteres der for, at dette er uhensigtsmæssigt, fordi systemet til forudsigelse af varmebehovet da må integreres med optimeringssystemet, hvorved programmet ikke kan opbygges modulært. Desuden mener vi, at alle relevante metoder til forudsigelse af varmebehov er så komplicerede, at denne integration vil være overordentlig vanskelig; både usikkerheder stammenende fra varmebehovets klimaafhængighed og fra meteorologiske forudsigelser bør tages i betragtning. I stedet foreslår vi, at beslutningsstøttesystemet giver operatøren mulighed for at undersøge følsomheden af den optimale løsning mht. variationer i input. Desuden foreslår vi, at systemet udstyres med muligheden for at simulere realiseringer af varmebehovet, baseret på den faktiske forudsigelse og tidligere fejl. Ved at lade systemet finde optimale driftplaner for hver af disse realisationer kan operatøren opnå indsigt i betydningen af usikkerhederne.

Det vises, at med moderne personlige computere (f.eks. 1 GHz Pentium III), operativsystemer (f.eks. RedHat Linux 6.0) og kompilerer (f.eks. GNU C 2.91) kan beregningerne udføres hurtigt nok til anvendelse i praksis. En optimal driftplan for en uge kan let findes indenfor 5 til 10 sekunder. Når mange mulige realisationer af det fremtidige varmebehov skal gennemregnes, er det dog nødvendigt at benytte nogle teknikker til reduktion af den nødvendige beregningstid. Resultaterne i denne rapport indikerer, at det er muligt at finde optimale driftplaner for 100 realisationer ved anvendelse af mindre end 3 minutters CPU tid. Endelig tillader metoderne massiv brug af parallel processering.

Contents

Acknowledgements **iii**

Summary **v**

Resumé **vii**

1 Introduction **1**

2 General considerations **3**

2.1 Optimal operation of decentralized CHP plants 3

2.2 Decision support system 6

3 Dynamic Optimization **9**

3.1 The dynamic optimization problem 9

3.2 Dynamic programming 10

3.3 Discrete formulation 12

4 Model of Sønderborg CHP plant **15**

4.1 Description of the system 16

4.2 Revenues and costs related to the operation 19

4.3 Formulation of the optimization problem 23

4.4 Numerical values 23

5 Sønderborg CHP plant – examples 27

5.1 Results 27

5.2 Computation time 34

6 Model of Sønderborg CHP plant – varying load 37

7 Stochastic effects 41

7.1 Simulation of heat demand 42

7.2 Efficient computations 44

8 Conclusion and discussion 47

Bibliography 51

Chapter 1

Introduction

In this report we consider load scheduling for decentralized combined heat and power (CHP) plants. Such plants are common in district heating systems in Denmark because they allow the heat to be produced near the end consumer and at the same time are able to produce electricity which can easily be transmitted over long distances. In this way the total efficiency of the energy production can be increased as compared to plants producing only electricity.

A common configuration for a particular district heating system is that one CHP plant supplies heat to the district heating (DH) network and electricity is sold to a power transmission company. Furthermore, some peak-load boilers are often placed in the DH network or near the CHP plant.

The CHP plant consists of a number of power production units, often based on engines running on natural gas (NG) or on bio gas (BG). However, also steam turbines supplied with heat from waste combustion or bio-fuels (e.g. wood chips) and gas turbines exist. Peak load boilers can be considered in parallel to the remaining units if they are placed at the CHP plant. The revenue from the electricity produced depend on the time of day and of the day of week. Basicly, for working days, the revenue is low during the night and high (two levels) during the day. For this and other reasons CHP plants are equipped with a heat storage facility which allows the heat and power to be produced during periods in which the revenue from the electricity sale is high. Some CHP plants also have the possibility to dispose of heat by cooling. This may be advantageous if the revenue from sale of the electricity is high enough.

Due to the heat storage facility the optimal operation of the plant depend on the future heat load which is unknown. However, to some extend the heat load may be predicted on-line, see e.g. (Nielsen & Madsen 2000), and these predictions and the associated uncertainty should be used when deciding how to run the CHP plant. Since such predictions are updated at regular intervals it makes sense to reconsider the planned operation of the plant at regular intervals. This can be done in a pure on-line fashion in which an automatic control system uses predictions of heat load, together with cost-information to find an optimal load schedule, which then is automatically transmitted to the SCADA system. However, in this report the aim is somewhat more moderate / realistic. We will aim at describing a system in which an operator decides the load schedule for the plant. However, exactly as for the automatic control system, it makes sense to reconsider the load schedule at regular intervals, maybe one or two times a day or when expectations for the future changes. Examples of changing expectations includes (i) large changes in the expected heat demand, e.g. due to large changes in the weather forecast, (ii) changes in the fuel cost, (iii) time-varying fuel costs, (iv) changes in the prices at which the electricity is sold. In these situations the operator needs a decision support tool by which different cost-structures and restrictions on the operation can be investigated. The tool must also help the operator to get an overview of the uncertainties inherent in the decision problem and help taking these into account. At present mainly the future heat demand is uncertain. In the future the revenue of the electricity produced at a particular time of day and day of week may not be given on beforehand. This is due to the fact that the whole Danish market for power moves towards a spot and future/forward market. In such a situation the planning would have to take even more stochastic effects than the future heat load into account. Similar remarks apply for the cost of natural gas. The suggested procedure is prepared to cover this expected situation as it allows all costs and revenues to be time-varying in an arbitrary manner. However, to investigate the effect of the uncertainties more simulations than considered in this report might be necessary.

Chapter 2

General considerations

2.1 Optimal operation of decentralized CHP plants

Operation of decentralized combined heat and power (CHP) plants is characterized by the fact that the electricity produced are sold to an power distribution company at a price which depend on time of day and week, cf. Table 4.1 on page 20. Except for peak load boilers, the plants are often the sole suppliers of heat to a local district heating network. Other sources of cooling than the district heating network is often not available and thus the ratio between the heat and power production is fixed. Using a heat storage facility the heat and power production are only tied together over long horizons.

The heat storage facility is a tank with hot water on top and cold in the bottom. When the storage is charged hot water is let in at the top and cold water is tapped from the bottom. There is no physical boundary between the hot and cold water and hence the difference in specific mass of the hot and cold water is solely responsible for keeping the temperature levels separate. So-called diffusers are used to ensure that the temperature gradient down trough the tank is preserved as well as possible. Due to the construction of the storage it will always be charged with water at a temperature equal to the temperature of the hot water already in the storage. Thus, it is only possible to change this temperature when the storage is empty, i.e. full of cold water from the return pipeline of the network.

The problem now consists of determining how the future heat load is allocated between the different heating plants in order to minimize the operational costs for an entire district heating system consisting of heating plants and distribution network. The objective is to minimize the expected operation costs within the planning horizon considered given an (uncertain) predicted heat load. The planning horizon will depend on the configuration of heating plants for the district heating system in question, but for system with heat accumulators the necessary planning horizon will be in the magnitude of days.

The discrete nature of a start/stop schedule, the long planning horizon, the number of restrictions imposed on a solution for the entire system and finally the complexity of the models describing the distribution network give rise to a problem of very considerable size. Thus the subject/task of identifying a cost function for the operation of an entire district heating system with multiple heating plants and following that finding a feasible solution to the posed optimization problem will in most cases be close to impossible due to the size of the problem which

In order to make the solution of the optimization problem feasible it is suggested to separate the optimization of the entire system into a scheduling between the different heat (and power) producing units including eventual heat accumulators (long planning horizon) followed by a control problem for the distribution network (considerably shorter control horizon). The potential gains by optimal scheduling between several production units will typically outweigh the potential gains by optimal operation of the distribution network by a considerably margin. Hence it makes sense to let the operation of the distribution network be subordinated the scheduling even at the cost of a (slightly) sup-optimal solution compared to an optimization which encompasses the entire district heating system.

The purpose of the scheduling is to derive a plan for each heating plant stating when the plant should be running as well as the heat production level. Depending on the configuration of production facilities the scheduling horizon will typically be up to five or seven days ahead. The scheduling between the different production units is done on basis of the following input:

- Predictions of heat load covering the horizon considered in the scheduling.
- Predictions covering the scheduling horizon of the necessary minimum supply temperature in order to fulfill the consumer requirements.
- The heat production costs for the different productions units. These may vary with time and production level. Also the start/stop costs has to be considered.

- Limitations in the available heat production capacity. The heating plants may be subject to contractual obligations, which may restrict the minimum or maximum heat production. These restrictions may be time-varying. Also limitations in the allowable rate of change of the heat production (or supply temperature) has to be considered.

The stochastic nature of the heat load and supply temperature predictions should be taken into account by the scheduling algorithm. Hence the scheduling could be formulated as a stochastic optimization problem, where the correlation structure of the prediction errors is included in the formulation. Possible methods to solve such a problem include stochastic dynamic programming and Monte Carlo simulations. However as described in Chapter 1 we do not aim at a fully automatic system and therefore the stochastic nature of the problem does not necessarily have to be included directly into the mathematical formulation of the optimization problem, cf. Section 2.2.

The outcome of the scheduling is a plan for the various heating plants covering the scheduling horizon. Only the first part of the plan corresponding to the horizon considered by the distribution network control is used as input to the distribution network controller. For each of the (running) plants the schedule consists of a set of (time-varying) constraints and reference values used as input to the distribution network controller:

- Maximum values for the permissible supply temperature. The maximum restriction corresponds to the minimum supply temperature constraint used in the scheduling.

- Flow rate for all but one of the heating plants. The low level pressure control in the network will normally require, that the flow rate only is allowed to vary freely for one of the heating plants. The latter should normally be the plant with the lowest production costs.

- The desired redistribution of heat load.

This report deals with the problem of load scheduling. Control of supply temperature is considered in e.g. (Madsen, Nielsen & Søgaard 1996), see also (Madsen & Nielsen 1997).

2.2 Decision support system

As explained in Chapter 1 this report does not focus on an automatic control system for controlling the production on the power production units. Instead we believe that the operator should have tools available which supports the decision making regarding load scheduling. Advantages of this approach includes:

- The ability for the operator to use auxiliary and not well defined information.
- The system need not to be able to handle all exceptions.
- The methods need not to be integrated into one system which takes all effects into account. However, the tools supplied should be integrated into one system with a common user interface.

The main disadvantage of the approach is that the operator need to use the decision support system at a regular basis. To facilitate this approach an automatic system for monitoring a plan could be implemented. This system should then alert the operator when conditions have changed so much that it is sensible to reconsider the current plan. Figure 2.1 summarizes this approach to identify and update the load schedule. The load schedule is based on forecasts of the heat load, which is not considered in this report, see e.g. (Nielsen & Madsen 2000). Using a decision support tool the operator identifies a load schedule which is implemented in the SCADA system. This current load schedule is then monitored on-line, and compared with updated forecasts of the heat load. The automatic monitoring system alerts the operator if the current plan seems to be sub-optimal as compared to updated information. In this case the operator is requested to revise the load schedule.

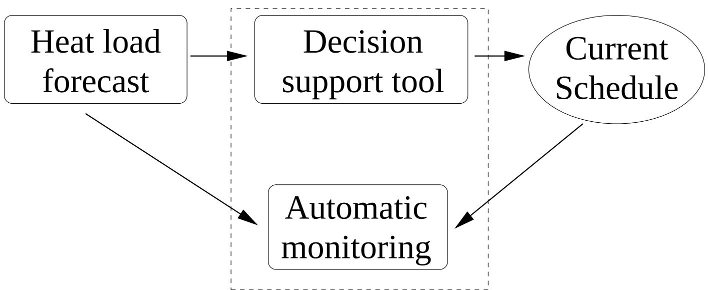

Figure 2.1: Identification and monitoring of a load schedule.

Decision support tool: The decision support tool should allow the operator to:

- Find optimal schedules given heat load forecasts, time-varying revenue from selling power, time-varying cost of fuel, and requirements on the operation of the plant.
- Analyze the sensitivity to deviations from the forecasted heat load.
- Simulate the future heat load and find optimal schedules given these simulated heat loads.

In all cases the optimal schedules can be found using deterministic dynamic programming. This is exemplified for Sønderborg Kraftvarmeværk in Chapters 4 and 5. However, this approach assumes the future heat load, revenue per unit power, and fuel costs to be known. Consequently, the load schedules identified are not guaranteed to be optimal in the actual setting where the future heat load is only known to some extend. The decision support tool circumvents this by allowing the operator to modify the forecast and find the resulting optimal schedule. A more stringent approach is to simulate a number of future paths of the heat load and find optimal schedules for each of these. Based on these the operator can identify a schedule which take the uncertainties into account. Section 7.1 considers how to simulate the heat load independently of the actual system producing the forecasts; this approach allows for a modular design of the software. The issue of uncertainty, or stochastic effects, are discussed in Chapter 7. The approach above corresponds to item 3 on page 41.

Automatic monitoring: The automatic system for monitoring the current load schedule should alert the operator if the schedule becomes sub-optimal. Strictly speaking, this is not possible without a system which can automaticly find optimal schedules. Since the entire approach is based on the assumption that it is unpractical to develop such a system the monitoring system is proposed to be based on secondary indicators and alarm levels. The monitoring system should display:

- The most recent heat load forecast together with the forecast in effect when the current load schedule was decided upon.
- The actual content of the heat storage together with the content forecasted when the current load schedule was decided upon.
- The actual load of the production units as compared to the planned.

- The optimal schedule (in the deterministic setting) given the most recent forecasts of the heat load together with the current schedule.

The type of alarms and the alarm levels should be configurable. Some relevant alarms could be based on:

- Feasibility of the current load schedule as compared to the most recent forecast.
- Difference in the actual and expected content of the heat storage.
- Difference in forecasts.
- Difference in revenues, given the most recent heat load forecast, of the current schedule and the optimal schedule given the most recent heat load forecast.

Chapter 3

Dynamic Optimization

To make the report self-contained the dynamic optimization problem is introduced and its solution obtained by the use of dynamic programming is introduced in this chapter. We have chosen dynamic programming for solving the optimization problem since it can be applied without restrictions on the properties of the criteria function or on the set of restrictions. Compared to many other methods for dynamic optimization the cost of this decision is (i) a relatively slow algorithm, and (ii) the need to formulate a discrete version of the optimization problem. In Section 5.2 these aspects are considered for the particular application. This chapter is based on (Ravn 1996).

3.1 The dynamic optimization problem

The dynamic optimization problem in discrete time, also termed *the discrete time optimal control model*, consists of a criteria function additive in the time steps $i = 0, \ldots, N-1$, a number of control variables $\mathbf{u}_i \subseteq \mathbb{R}^m$, a number of state variables $\mathbf{x}_i \subseteq \mathbb{R}^n$, an equation describing the evolution of a state vector, a set of restrictions on the control and state vectors, and a restriction on the final value of the state. This can be formulated as

$$\max \sum_{i=0}^{N-1} f_i(\mathbf{x}_i, \mathbf{u}_i) + f_N(\mathbf{x}_N) \tag{3.1}$$

w.r.t.

$$\begin{aligned} \mathbf{x}_{i+1} &= \mathbf{d}_i(\mathbf{x}_i, \mathbf{u}_i) \quad i = 0, \ldots, N-1 \\ (\mathbf{x}_i, \mathbf{u}_i) &\in \mathbb{V}_i \qquad\qquad i = 0, \ldots, N-1 \\ \mathbf{x}_N &\in \mathbb{X}_N \end{aligned}$$

where $\mathbf{u}_1, \mathbf{u}_1, \ldots, \mathbf{u}_{N-1}$ should be chosen to maximize the criteria function. $\mathbb{V}_0$ must be chosen to reflect the fact that the initial state $\mathbf{x}_0$ is known.

In the context of this report we can think of the summation in the criteria function as describing the revenue from running the CHP plant over the horizon considered. The last term in the criteria function can be used to assign a value to the final storage content. The state $\mathbf{x}$ could contain the content of the storage and, possibly, the state of the gas turbine with the purpose of including start/stop costs. The multivariate control variable $\mathbf{u}$ will contain the load of the heat and/or power producing units on the plant.

3.2 Dynamic programming

The dynamic structure of the optimal control problem (3.1) is shown in Figure 3.1.

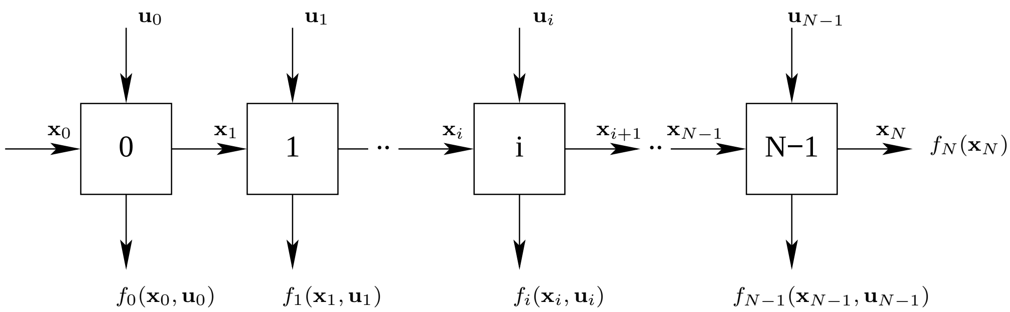

Figure 3.1: Dynamic structure of the optimal control problem.

It is seen that if $\mathbf{x}_{N-1}$ is known then the optimal $\mathbf{u}_{N-1}$ can be found by considering the criteria function

$$f_{N-1}(\mathbf{x}_{N-1}, \mathbf{u}_{N-1}) + f_N(\mathbf{x}_N) = f_{N-1}(\mathbf{x}_{N-1}, \mathbf{u}_{N-1}) + f_N(\mathbf{d}_{N-1}(\mathbf{x}_{N-1}, \mathbf{u}_{N-1})).$$

This optimization problem will have m control variables, as compared to mN for the original problem. Let $\mathbf{u}_i^*$ denote the optimal control at time i. It is then clear that the optimal control at time $N-1$ is a function of the state at time $N-1$ and we may write

$$\mathbf{u}_{N-1}^*(\mathbf{x}_{N-1}).$$

Furthermore, it is clear that the optimal value of the criteria function depends on $\mathbf{x}_{N-1}$ only and we can write the contribution at times $N-1$ and onwards to the criteria function as

$$F_{N-1}(\mathbf{x}_{N-1}) = f_{N-1}(\mathbf{x}_{N-1}, \mathbf{u}_{N-1}^*(\mathbf{x}_{N-1})) + f_N(\mathbf{d}_{N-1}(\mathbf{x}_{N-1}, \mathbf{u}_{N-1}^*(\mathbf{x}_{N-1}))). \tag{3.2}$$

For completeness we define

$$F_N(\mathbf{x}_N) = f_N(\mathbf{x}_N). \tag{3.3}$$

Given $\mathbf{x}_{N-2}$ we can then find $\mathbf{u}_{N-2}^*(\mathbf{x}_{N-2})$ by considering the criteria function

$$f_{N-2}(\mathbf{x}_{N-2}, \mathbf{u}_{N-2}) + F_{N-1}(\mathbf{d}_{N-2}(\mathbf{x}_{N-2}, \mathbf{u}_{N-2})),$$

since

$$\mathbf{x}_{N-1} = \mathbf{d}_{N-2}(\mathbf{x}_{N-2}, \mathbf{u}_{N-2}).$$

Again it is seen that the optimal control at time $N-2$ is a function of the state only and thus we may write

$$\mathbf{u}_{N-2}^*(\mathbf{x}_{N-2})$$

and furthermore

$$F_{N-2}(\mathbf{x}_{N-2}) = f_{N-2}(\mathbf{x}_{N-2}, \mathbf{u}_{N-2}^*(\mathbf{x}_{N-2})) + F_{N-1}(\mathbf{d}_{N-2}(\mathbf{x}_{N-2}, \mathbf{u}_{N-2}^*(\mathbf{x}_{N-2}))). \tag{3.4}$$

It is seen that this process, also called the *backwards recursion*, can be continued until the known initial state $\mathbf{x}_0$ is reached.

When the backward recursion is terminated at the initial (known) state the initial optimal control $\mathbf{u}_0^*$ will be known. From this the state at time 1 can be found as $\mathbf{x}_1^* = \mathbf{d}_0(\mathbf{x}_0, \mathbf{u}_0^*)$ whereby the optimal control at time 1 can be found as $\mathbf{u}_1^*(\mathbf{x}_1^*)$. This *forward recursion* can be continued until all optimal controls are known.

For a more formal introduction to dynamic programming, including *Bellman's principle of optimality*, the reader may consult (Ravn 1996, Section 7.4). See also (Bellman 1957).

3.3 Discrete formulation

In the previous section we considered the functions $\mathbf{u}_i^*(\mathbf{x}_i)$ and $F_i(\mathbf{x}_i)$. In the general case, i.e. without assumptions regarding the parametric form of the criteria function, these can not be expressed in a simple way. In this report we therefore calculate the functions at a number of discrete values of the state $\mathbf{x}_i$ and use multidimensional interpolation (Press, Teukolsky, Vetterling & Flannery 1992) to obtain values for arbitrary arguments. Using this approach the number of optimization problems which must be considered during the backwards recursion will grow exponentially with the dimension of the state vector. However, not all of these need to be solved since some combination of states may not be feasible. The optimization problems at each stage of the backwards recursion is of dimension one in this report and for this we employ golden section search (Press et al. 1992). For the examples considered in Chapter 5 we only need to determine if the gas turbine should be stopped or not.

Due to the large number of calculations which must be performed the core of the optimization algorithm need to be implemented in a relatively low-level computer language such as C or Fortran. When the number of states increases it will be necessary to use parallel programming to be able to obtain a solution within a reasonable time. Fortunately, the calculations which must be performed at each stage of the backward recursion are not dependent. For this reason it is relatively uncomplicated to implement the method so that parallel processing can be applied. However, at the present point in time, we believe that the cost of appropriate hardware will be too high for use in production planning at decentralized CHP plants. Furthermore, the exponential growth will always be prohibitive for some number of state variables. A more promising approach, which do not exclude parallel processing, is suggested by Chen, Ruppert & Shoemaker (1999) who attack the problem by not considering all points in the grid spanned by the discrete values of the state variables.

As an example consider a function of two scalar arguments and assume that 10 levels of each argument are considered. The levels are denoted $x_i;\ i = 1, \dots, 10$ and $y_j;\ j = 1, \dots, 10$ If all combinations is to be investigated the function need to be evaluated 100 times. Now if it is assumed that the function is additive in the arguments, i.e.

$$f(x_i, y_j) = \mu + \alpha_i + \beta_j$$

where $\sum_i \alpha_i = \sum_j \beta_j = 0$, then the values at the levels can be described by use of 19 coefficients. These coefficients can be obtained by evaluating the function at 19 of the 100 combinations of levels, see also Figure 3.2.

The above example is a simplification of the approach used by Chen et al. (1999)

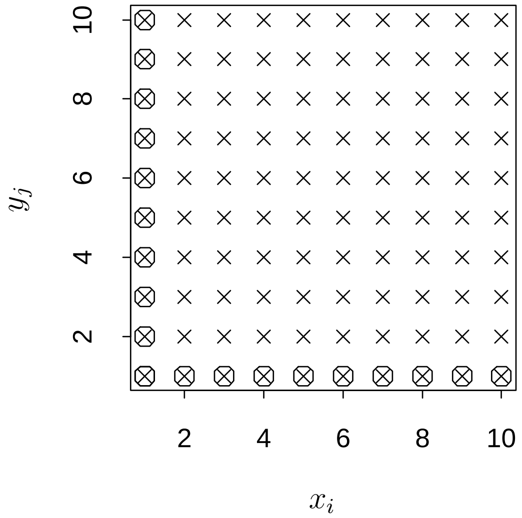

Figure 3.2: Points needed for an additive model (○) and a non-additive model (×).

to reduce the number of combinations which must be investigated at each stage of the backward recursion. However, the function is not assumed to be fully additive in the arguments. Instead, it is assumed that some low-order interactions between state variables might be present and the grid-points are selected by use of an experimental design tuned to detect such interactions. Furthermore, regression splines are used to generate a continuous representation of the functions. Since some low-order interactions are allowed for the approach outlined above is only beneficial when the number of state variables is high.

We have not used the approach in this report since for the example considered we have two state variables only, however, see Chapter 7. Furthermore, it will be quite time-consuming to implement the methods.

Chapter 4

Model of Sønderborg CHP plant

The combined heat and power plant in Sønderborg, Denmark consists of a waste combustion unit, a gas turbine, a steam turbine, and a heat storage facility. The steam turbine receives steam from the waste combustion unit and, via an exhaust gas boiler, from the gas turbine. Furthermore, some cooling towers are installed. Figure 4.1 shows an outline of plant without the cooling towers. The power produced is sold to the power transmission company Eltra at a price depending on the time of day/week and the natural gas are purchased from the distribution company Naturgas Syd I/S at a price which in the future will be varying to some extend.

In this chapter we will formulate a model of Sønderborg CHP plant which allows us to determine the optimal operation of the gas turbine given the heat demand, the amount of steam coming from the waste combustion unit, and the amount of cooling. Cooling will not be explicitly mentioned since it can be handled as an extra heat demand. For other plants cooling might introduce a loss of revenue from selling heat.

The aim of this chapter is to demonstrate how optimization can be applied to find good solutions for arbitrary time-variations of the power- and gas-price. Consequently, the focus has not been on modelling the revenue/cost-structure of Sønderborg CHP plant in all details. As will be clear from Section 4.1 the important aspect is that the total revenue can be expressed as a function of the amount of natural gas used in each time step.

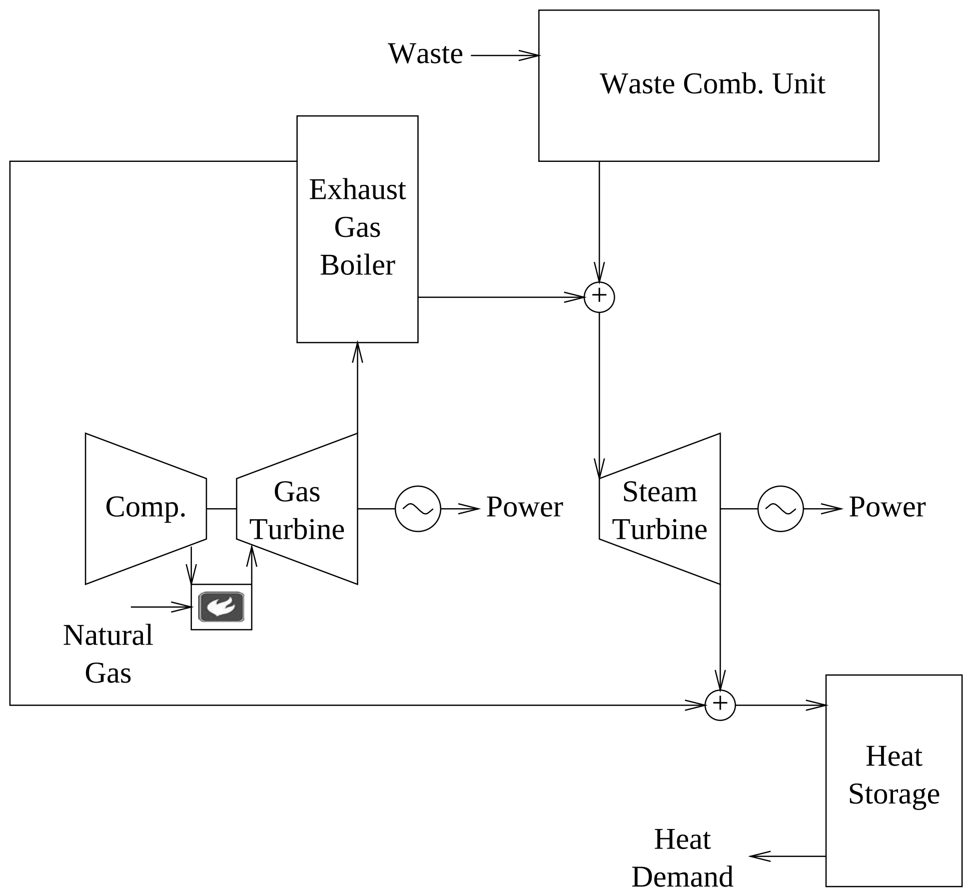

Figure 4.1: Outline of Sønderborg CHP plant. Cooling towers are not shown.

4.1 Description of the system

A graphical description of the mathematical model is depicted in Figure 4.2. The model has the following inputs:

- F_g; the amount of natural gas used by the gas turbine.
- H_w; the amount of steam coming from the waste combustion unit.

The model has the following outputs:

- $\delta_w H_w$; the amount of energy needed for pre-heating in the waste combustion unit.

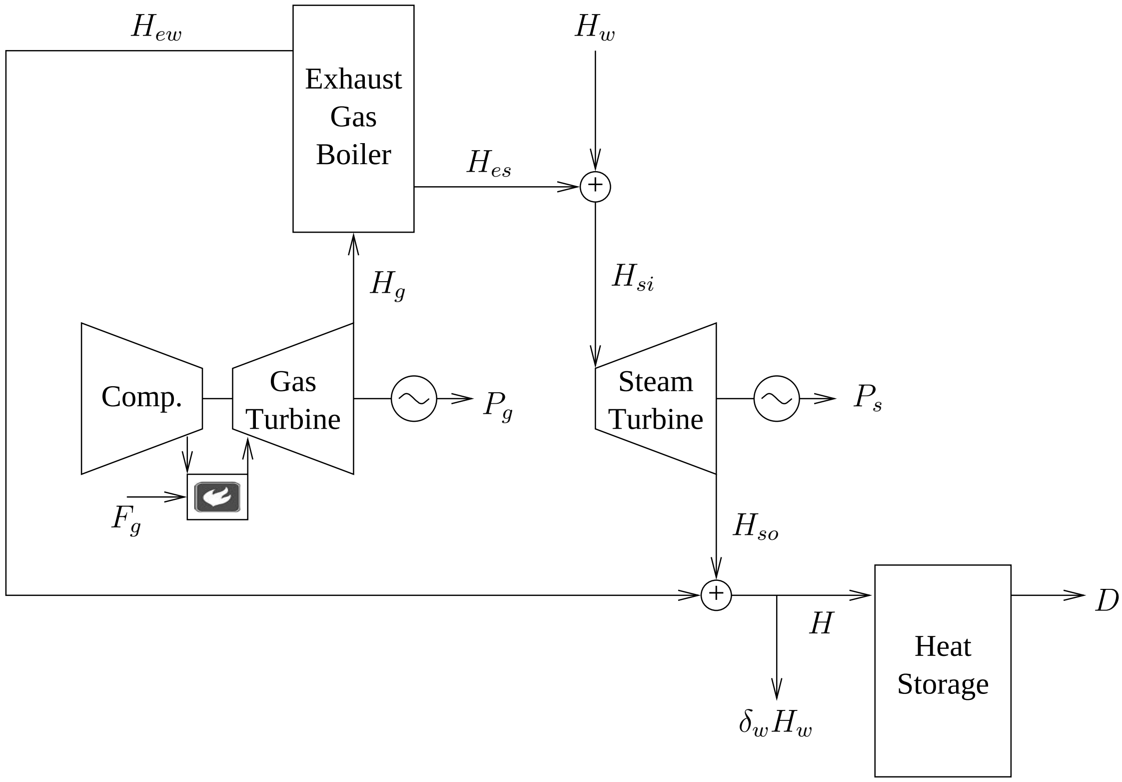

Figure 4.2: Graphical description of the model for Sønderborg CHP plant.

- D; the heat demand of the district heating network connected to the plant.
- P_g; the power produced by the gas turbine.
- P_s; the power produced by the steam turbine.

Finally the model has the following internal variables:

- H_g; the amount of heat in the exhaust from the steam turbine.
- H_{es}; the amount of steam produced by the exhaust gas boiler.
- H_{ew}; the amount of hot water produced by the exhaust gas boiler.
- H_{si}; the amount of steam going to the steam turbine.
- H_{so}; the amount of heat coming from the steam turbine.

- H; the amount of heat going to the heat storage facility.
- A; the amount of heat stored in the heat storage facility.

For the variable A the unit is MWh and for the rest the unit is MW.

The amount of power which is sold to the power transmission company is

$$P = (1 - \delta_p)(P_s + P_g), \tag{4.1}$$

where δ_p denotes the fraction of the power used by the plant itself.

The gas turbine is either not running or it is running at a constant load. Therefore it can be characterized by two constants η_{gp} and η_{gh} indicating the fraction of F_g which is turned into power (P_g) and heat (H_g), respectively:

$$\begin{array}{rcl} P_g & = & \eta_{gp} F_g \\ H_g & = & \eta_{gh} F_g \end{array} . \tag{4.2}$$

Hence the exhaust gas boiler can also be described by two constants η_{es} and η_{ew} indicating the fraction of H_g which is turned into steam (H_{es}) and hot water (H_{ew}), respectively:

$$\begin{array}{rcl} H_{es} & = & \eta_{es} H_g \\ H_{ew} & = & \eta_{ew} H_g \end{array} . \tag{4.3}$$

For the steam turbine it is necessary to know how much of the steam H_{si} is turned into power (P_s) and heat (H_{so}) at different loads. It is assumed that this can be handled by letting these fractions depend on the load in terms of H_{si}, i.e.

$$\begin{array}{rcl} P_s & = & \eta_{sp}(H_{si}) H_{si} \\ H_{so} & = & \eta_{sh}(H_{si}) H_{si} \end{array} \tag{4.4}$$

The terms $\eta_{sp}(H_{si})$ and $\eta_{sh}(H_{si})$ are called *efficiency curves*. The quantity $1-\eta_{sp}(H_{si})-\eta_{sh}(H_{si})$ is the relative loss from the steam turbine and the generator connected to it. The input to the steam turbine H_{si} is given by

$$H_{si} = H_{es} + H_w. \tag{4.5}$$

The total amount of heat produced H is given by

$$H = H_{so} + H_{ew} - \delta_w H_w, \tag{4.6}$$

which can be regarded as a function of the amount of natural gas used F_g and given a number of constants (mainly efficiencies), including the amount of steam coming from the waste combustion unit H_w.

Finally, the heat accumulator must be described by a dynamic equation. Here a discrete-time formulation will be used:

$$A_t = (1 - \eta_a \Delta t) A_{t-\Delta t} + \Delta t \left(H(F_{g,t}) - D_t \right), \tag{4.7}$$

in which the variables are indexed by time and Δt is the length of the time interval. Hence A_t is the amount of heat accumulated at time t, η_a is the relative heat loss during Δt, $\Delta t D_t$ is the heat demand over the period from $t - \Delta t$ until t and $\Delta t H(F_{g,t})$ is the amount of heat produced during the same period.

The CHP plant will impose some restrictions on the operation. As already mentioned the gas turbine will either be stopped or it will run at a certain level, i.e.

$$F_g \in \{0, f_g\}, \tag{4.8}$$

where f_g is the amount of gas used when the turbine is running. The heat accumulator has a given capacity in terms of a volume, which is converted in to MWh by use of the following relation

$$A_{max} = \frac{\rho V C_p \Delta T}{3.6 \times 10^9 J/MWh}, \tag{4.9}$$

where ρ is the specific mass of water, C_p is the specific heat capacity of water, V is the volume of the tank ($11500 m^3$), and ΔT is the difference between the supply and return temperatures of the water in the district heating system. This results in the following restriction on A_t for all points in time:

$$A_t \in [0, A_{max}]. \tag{4.10}$$

If not stated otherwise we use $\Delta T = 40.58\ °C$ and the values of ρ and C_p at $62\ °C$ (Incropera & DeWitt 1985, Table A.6), i.e. $\rho = 982.3\ kg/m^3$ and $C_p = 4186\ J/(kg \cdot K)$. Hence $A_{max} = 533\ MWh$.

4.2 Revenues and costs related to the operation

The revenues from running the plant consist of a revenue from selling power to the transmission company and of a revenue from the combustion of waste. Since we will consider the heat supplied from the waste combustion unit as given it is not necessary to take this into account when optimizing the operation of the gas turbine. Likewise, the heat demand can not be influenced by the operation of the plant, and for this reason it is not necessary to consider the revenue from selling heat to the district heating network.

Power: The revenue from selling power depends on the time of day on which the power is produced; an example is depicted in Figure 4.3 and Table 4.1 list the actual prices and periods. However, for the CHP plant considered here the periods may deviate from what is shown in Table 4.1. This is decided by the power transmission company and the information is available to the CHP plant one week in advance. Here this problem is handled by letting the revenue per MWh power sold depend on the running time in a totally unrestricted fashion, i.e. when the period from time t until $t + K\Delta t$ is considered the revenue from selling power is

$$\sum_{k=1}^{K} c^{(p)}_{t+k\Delta t} \Delta t P_{t+k\Delta t}, \tag{4.11}$$

where the time index $t + k\Delta t$ indicates the period from $t + (k-1)\Delta t$ until $t + k\Delta t$, $c^{(p)}_{t+k\Delta t}$ is the revenue per MWh in the period just mentioned, and $P_{t+k\Delta t}$ is the average power in MW in the same period of time, Δt is measured in hours. Note that no restrictions on the development of $c^{(p)}$ over time is assumed.

	Load		
Mon. – Fri.	Low (153 DKr/MWh)	High (359 DKr/MWh)	Peak (483 DKr/MWh)
Oct. – Feb.	21:00 – 06:30	06:30 – 07:30 12:00 – 17:00 18:30 – 21:00	07:30 – 12:00 17:00 – 18:30
Oct. – Feb.	21:00 – 06:30	06:30 – 07:30 12:00 – 21:00	07:30 – 12:00
The following days are always defined as low load: Weekends, Maundy Thursday, Good Friday, Easter Monday, Friday four weeks after Easter (St. Bededag), Ascension Day, With Monday, and December 25-26.			

Table 4.1: Revenue from selling power.

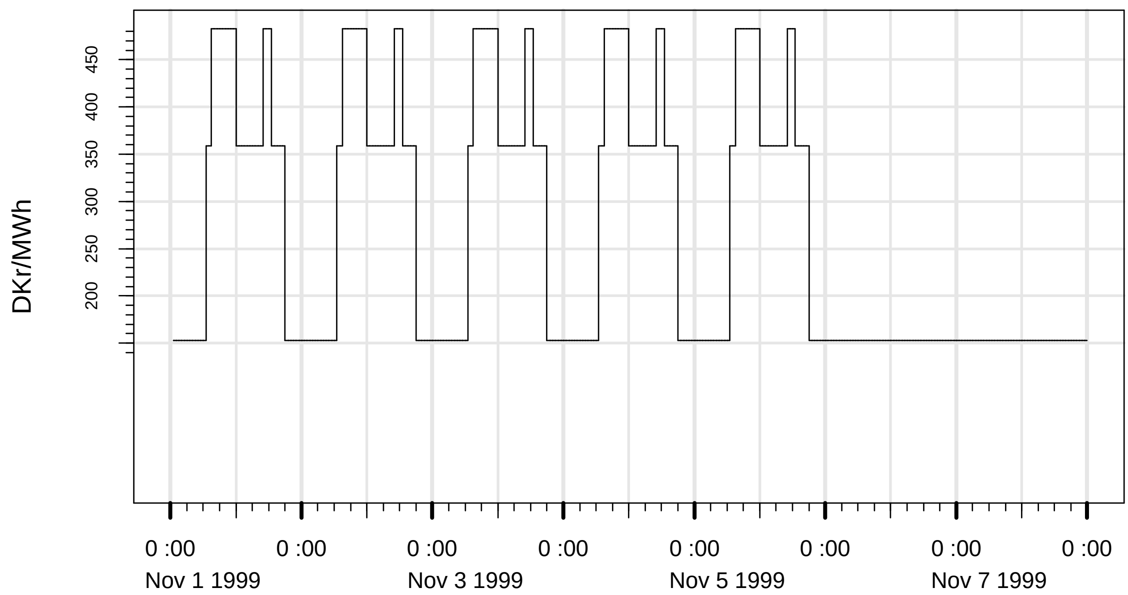

Figure 4.3: Normal revenue from selling power for one week starting at midnight Monday, 1 Oct. 1999.

Natural Gas: The natural gas used is purchased from a distribution company. Since the price of the gas in the future may depend on when it is used we will use a formulation similar to (4.11). However, due to taxation, the price of gas depends on whether the gas is used for producing heat or for producing electricity. Although, this is not physically well defined some reasonable approximations has been agreed on. For the CHP plant in Sønderborg these approximations are described below.

With the notation used above and indicated in Figure 4.2 the heat production originating from natural gas is calculated as

$$H_{NG} = H_{ew} + (H - H_{ew})\frac{H_{es}}{H_w + H_{es}}. \tag{4.12}$$

Considering the time period from t until $t + K\Delta t$ the cost of natural gas can then be expressed by weighting the cost of natural gas used for heat production $c^{(ng,h)}$ and the cost of natural gas used for power production $c^{(ng,p)}$ relative to the amount of natural gas used for heat production, i.e.

$$\sum_{k=1}^{K} \left(c_{t+k\Delta t}^{(ng,h)} \frac{H_{NG,t+k\Delta t}}{H_{t+k\Delta t}} + c_{t+k\Delta t}^{(ng,p)} \left(1 - \frac{H_{NG,t+k\Delta t}}{H_{t+k\Delta t}} \right) \right) \Delta t F_{g,t+k\Delta t}. \tag{4.13}$$

Start/stop: There is a cost associated with the number of times the gas turbine is started. These start/stop costs will be included in the simplest possible way. It is assumed that the cost of starting the unit will not depend on how long time it has been stopped. This cost $c^{(s)}$ should include extra fuel and power needed together with wear and tear from starting and stopping. No explicit cost is assigned when the unit is stopped, but since a start must eventually be followed by a stop the wear and tear originating from stopping should be assigned when starting the unit. Under these assumptions start/stop costs for the time period from t until $t + K\Delta t$ can be expressed as

$$\sum_{k=1}^{K} c^{(s)} I(F_{g,t+k\Delta t} > 0 \wedge F_{g,t+(k-1)\Delta t} = 0), \tag{4.14}$$

where $I(u)$ is one if the logical expression u is true and zero otherwise.

Maintenance: Besides the start/stop cost, there is a maintenance cost associated with the amount of power produced by the gas turbine. This cost is modelled as

$$\sum_{k=1}^{K} c^{(m)} \Delta t P_{g,t+k\Delta t}. \tag{4.15}$$

Criteria function: The revenue from running the plant over the time period from t until $t + K\Delta t$ is the revenue from selling power (4.11) with the cost of natural gas (4.13), start/stop (4.14), and maintenance (4.15) subtracted. As noted on page 18 the total heat produced can be regarded as a function of the amount of natural gas used F_g. This is also true for the power production and for the amount of natural gas used for producing heat. Thus, the total revenue can be expressed as

$$\sum_{k=1}^{K} r_{t+k\Delta t}(F_{g,t+k\Delta t}, F_{g,t+(k-1)\Delta t}), \tag{4.16}$$

where

$$\begin{aligned} & r_{t+k\Delta t}(F_{g,t+k\Delta t}, F_{g,t+(k-1)\Delta t}) = \\ & \quad c^{(p)}_{t+k\Delta t} \Delta t P(F_{g,t+k\Delta t}) \\ & \quad - \left(c^{(ng,h)}_{t+k\Delta t} \frac{H_{NG}(F_{g,t+k\Delta t})}{H(F_{g,t+k\Delta t})} + c^{(ng,p)}_{t+k\Delta t} \left(1 - \frac{H_{NG}(F_{g,t+k\Delta t})}{H(F_{g,t+k\Delta t})} \right) \right) \Delta t F_{g,t+k\Delta t} \\ & \quad - c^{(s)} I(F_{g,t+k\Delta t} > 0 \wedge F_{g,t+(k-1)\Delta t} = 0) \\ & \quad - c^{(m)} \Delta t P_g(F_{g,t+k\Delta t}), \end{aligned} \tag{4.17}$$

is the revenue in the time period from $t + (k-1)\Delta t$ until $t + k\Delta t$.

4.3 Formulation of the optimization problem

Using (4.17), (4.6), and (4.7) the operation of the gas turbine during the time period from t until $t + K\Delta t$ can be decided using the following mathematical formulation which in turn form the basis of a computer implementation

$$\max \sum_{k=1}^{K} r_{t+k\Delta t}(F_{g,t+k\Delta t}, F_{g,t+(k-1)\Delta t}) + c^{(a)}_{t+K\Delta t} A_{t+K\Delta t} \tag{4.18}$$

w.r.t.

$$\begin{aligned} A_{t+k\Delta t} &= (1 - \eta_a \Delta t) A_{t+(k-1)\Delta t} + \Delta t \left(H(F_{g,t+k\Delta t}) - D_{t+k\Delta t} \right) \\ A_{t+k\Delta t} &\in [0, A_{max}] \\ F_{g,t+k\Delta t} &\in \{0, f_g\} \end{aligned}$$

which is solved to find optimal values for $F_{g,t+k\Delta t}$; $k = 1, \ldots, K$. The initial content of the heat storage A_t and the state of the gas turbine $F_{g,t}$ is assumed to be known. In the criteria the terminal content of the heat storage $A_{t+K\Delta t}$ is assigned a value (positive or negative) per energy unit $c^{(a)}_{t+K\Delta t}$. The bounds on the storage content are constant over the time period in the formulation used above. However, it does not complicate the solution of the problem to use varying bounds and it is often relevant to bound the terminal content of the heat storage. Especially, when it is not assigned a value ($c^{(a)}_{t+K\Delta t} = 0$).

Note that when considering period k, i.e. the time period from $t + (k - 1)\Delta t$ until $t + k\Delta t$, the load of the gas turbine in the previous period $F_{g,t+(k-1)\Delta t}$ acts as a state variable. This is clearly seen when considering inputs and outputs from time step k as depicted in Figure 4.4.

Note also that in this section we use a different notation for the time steps than in Chapter 3. Here we generally index a variable by the right endpoint of the time step up to that endpoint. The exception of this is the storage content which in principle is valid for the actual time index only.

4.4 Numerical values

This section lists the numerical values used for Sønderborg CHP plant. These values are identified in cooperation with Sønderborg Kraftvarmeværk.

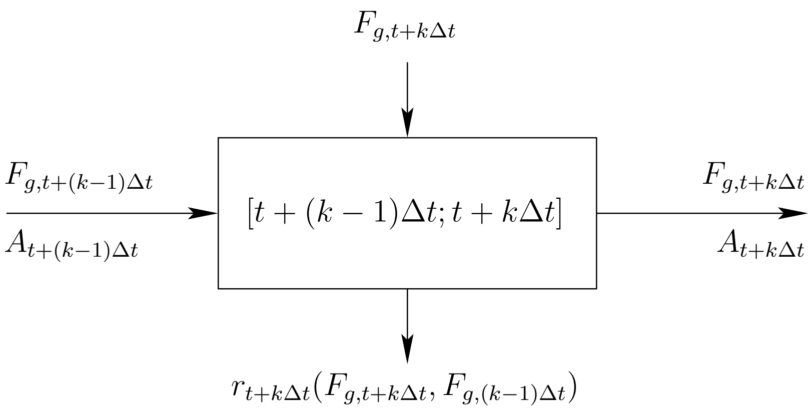

Figure 4.4: Inputs and outputs at time step k, i.e. the time period from $t+(k-1)\Delta t$ until $t+k\Delta t$.

Physical description of the system:

$$\delta_w = 0.1914$$

$$\delta_p = 0.04$$

$$\eta_{gp} = 0.3951$$

$$\eta_{gh} = 0.5749$$

$$\eta_{es} = 0.7351$$

$$\eta_{ew} = 0.2030$$

$$\eta_{sp}(H_{si}) = 0.1674 + 0.0007908\ MW^{-1} \times H_{si}$$

$$\eta_{sh}(H_{si}) = 0.8001 - 0.0007892\ MW^{-1} \times H_{si}$$

$$\eta_a = 0$$

Bounds on operation:

$$f_g = 100\ MW$$

$$A_{max} = 533\ MWh$$

Revenues and Costs:

$c^{(p)}$ as described in Table 4.1.

$c^{ng,h} = 76.8640\ DKr/MWh$

$c^{ng,p} = 25.6213\ DKr/MWh$

$c^{(s)} = 3000\ DKr$

$c^{(m)} = 40\ DKr/MWh$

Chapter 5

Sønderborg CHP plant – examples

Examples of optimal operation of the gas turbine at Sønderborg CHP plant is included in this chapter. In the examples it is assumed that the load of the waste combustion unit is given and that the cooling towers are not used. The underlying model used for maximization of the revenue is described in Chapter 4. The optimization examples are performed using data from the SCADA system of the CHP plant. The heat rate coming from the waste combustion unit is taken as the average of this quantity over the individual half-hour periods for which the optimizations are carried out.

5.1 Results

Figure 5.1 show the hourly heat demand for the last quarter of 1999. From this series a number of periods are selected in the this section. For each of these periods the optimal schedule for the gas turbine is found. As compared to Section 2.2 this does not conform to the actual use of dynamic programming since forecasted or simulated heat demands is supposed to be used. However, because we do not have access to meteorological forecasts we are not able to use actual forecasts of the heat load, cf. (Nielsen & Madsen 2000).

From the data shown in Figure 5.1 four periods of length one week were initially selected. All periods start on Monday at 06:30 (the time at which the high-load period starts) and is of the length of one week. The time interval (Δt) used is 30 minutes; interpolation is used to generate heat demand data corresponding to the interval length,

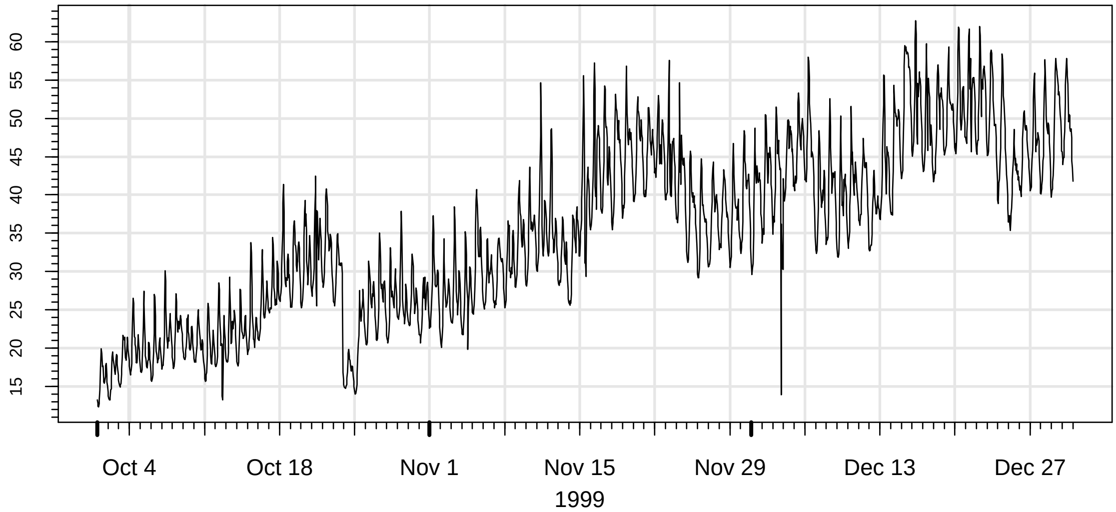

Figure 5.1: Hourly heat demand in MW for the last quarter of 1999.

and it is assumed that the heat storage is empty and the gas turbine stopped at the onset of the periods. The four periods start at:

Period 1: Monday, October 4, 1999, 06:30.

Period 2: Monday, October 11, 1999, 06:30.

Period 3: Monday, November 1, 1999, 06:30.

Period 4: Monday, December 13, 1999, 06:30.

The optimal operation of the gas turbine found using dynamic programming and discrete levels of the heat storage corresponding to steps of 1 MWh are shown in Figure 5.2 (period 1), 5.3 (period 2), 5.4 (period 3), and 5.5 (period 4). No restriction on the terminal content of the heat storage are imposed, instead the terminal storage content is assigned the cost 185 DKr/MWh ($c^{(a)} = -185\ DKr/MWh$).[1]

[1]When the gas turbine runs 1 MWh of heat corresponds to approximately 0.9 MWh of power. If this is produced at 153 DKr/MWh instead of 359 DKr/MWh, then $0.9 \times (359 - 153) = 185$ DKr is lost for every MWh in the heat storage Monday morning.

The results are quite plausible:

- For period 1 with the lowest heat demand (approximately 3450 MWh) the gas turbine almost only runs during peak-load periods.
- For period 2 with a heat demand of approximately 3950 MWh the gas turbine runs mainly in peak-load periods, but also in some high-load periods, so that the heat storage is full on Friday, October 15 at 21:00. Hereafter, the storage is emptied during the weekend and it is necessary to run the gas turbine during a low-load period and stop it so that the storage is empty Monday morning.
- For period 3 with a heat demand of approximately 4800 MWh the gas turbine runs most high-load and all peak-load periods and it is necessary to start it during the weekend also.
- For period 4 with a heat demand of approximately 8340 MWh the gas turbine runs for almost the entire period. Note however, that opposed to the periods 1–3 it is not optimal to have the heat storage full Friday at 21:00. This is due to the capacity limit of the heat storage. If the storage capacity were approximately 590 MWh, corresponding to a temperature difference of approximately 45 $°C$, then it would be possible to keep the gas turbine running until Friday at 21:00. However, unless the capacity were larger than 800 MWh (corresponding to a temperature difference of 61 $°C$), it would not be possible to avoid the start during the period.

For period 2 the gas turbine is started at 06:00 in the end of the period (Monday, October 18). This seems suboptimal in that the power produced must be sold at the low rate, cf. Table 4.1. Using discrete levels of the heat storage corresponding to steps of 0.1 MWh confirms this because the gas turbine is not started in this case. However, used in a decision support tool we do not regard the suboptimal start described above to be a serious problem in that the operator will be able to take appropriate actions.

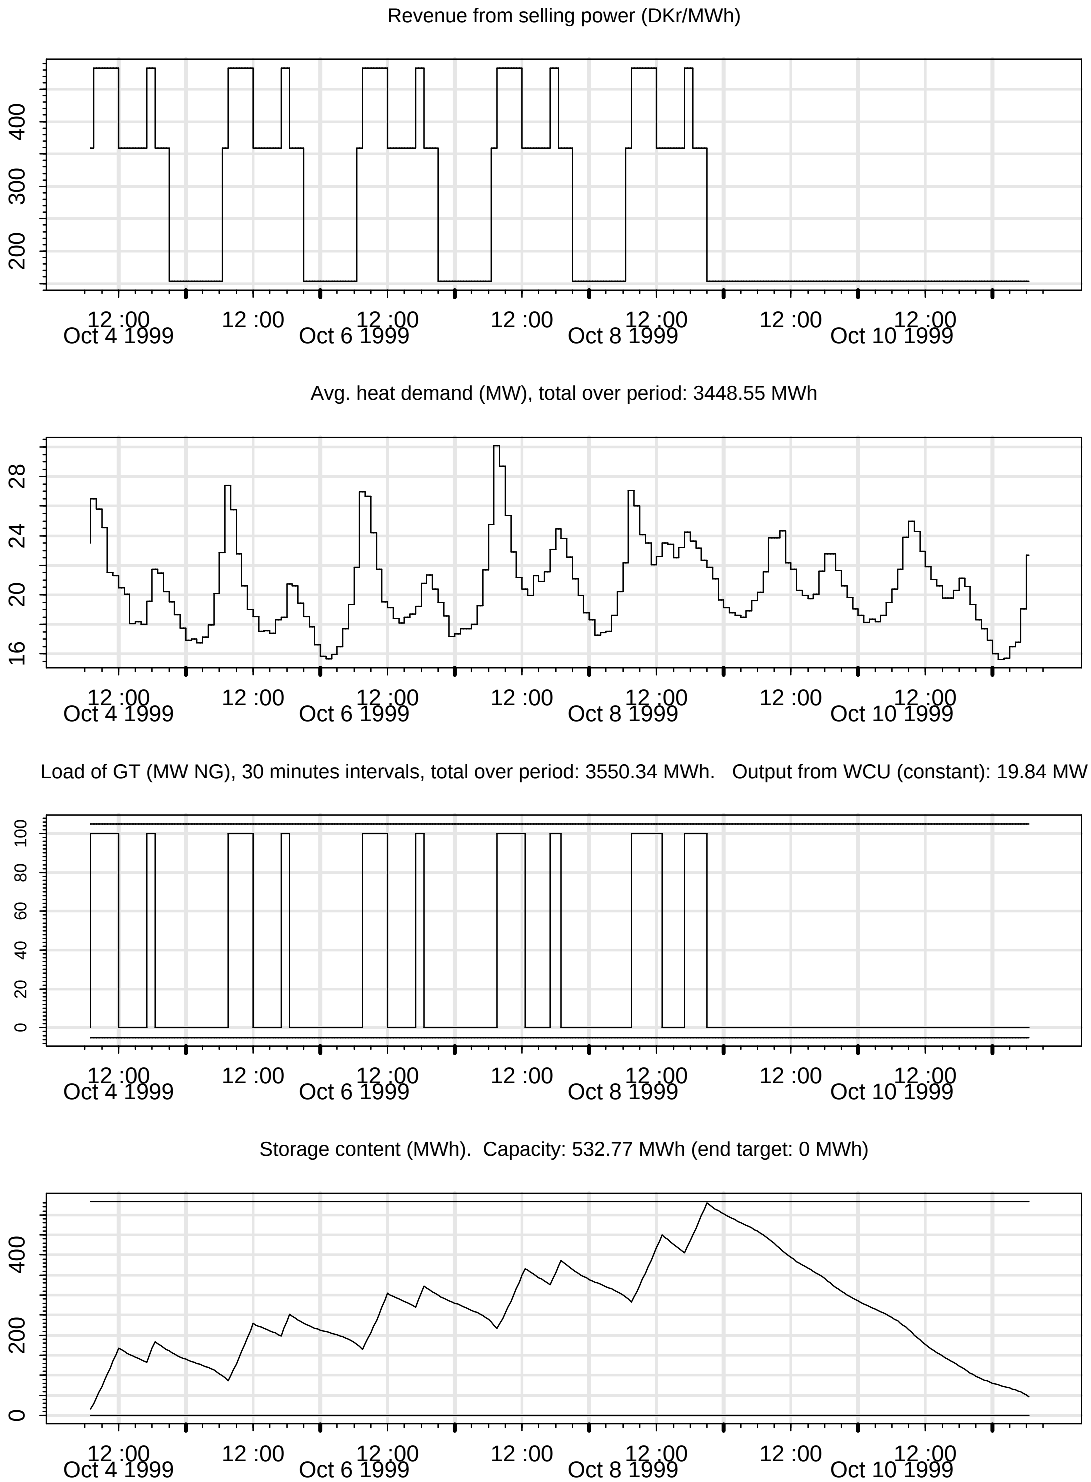

Figure 5.2: Optimal operation of gas turbine for one week starting on 1999.10.04 (criteria function: 612526 DKr).

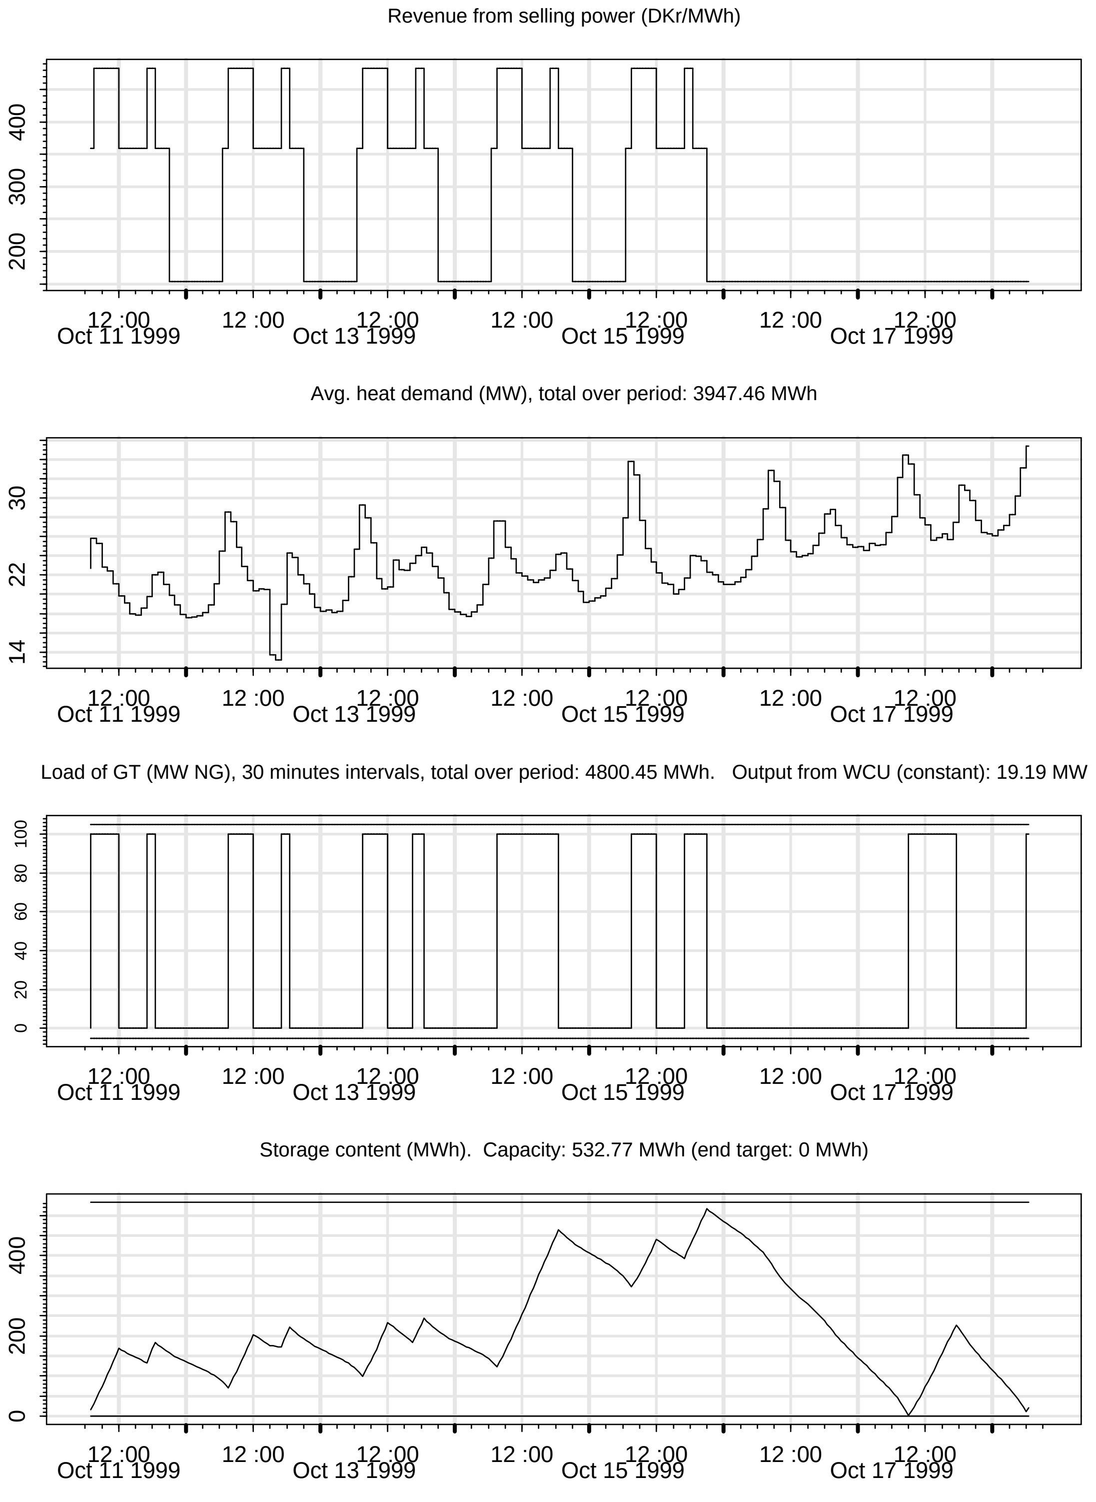

Figure 5.3: Optimal operation of gas turbine for one week starting on 1999.10.11 (criteria function: 631390 DKr).

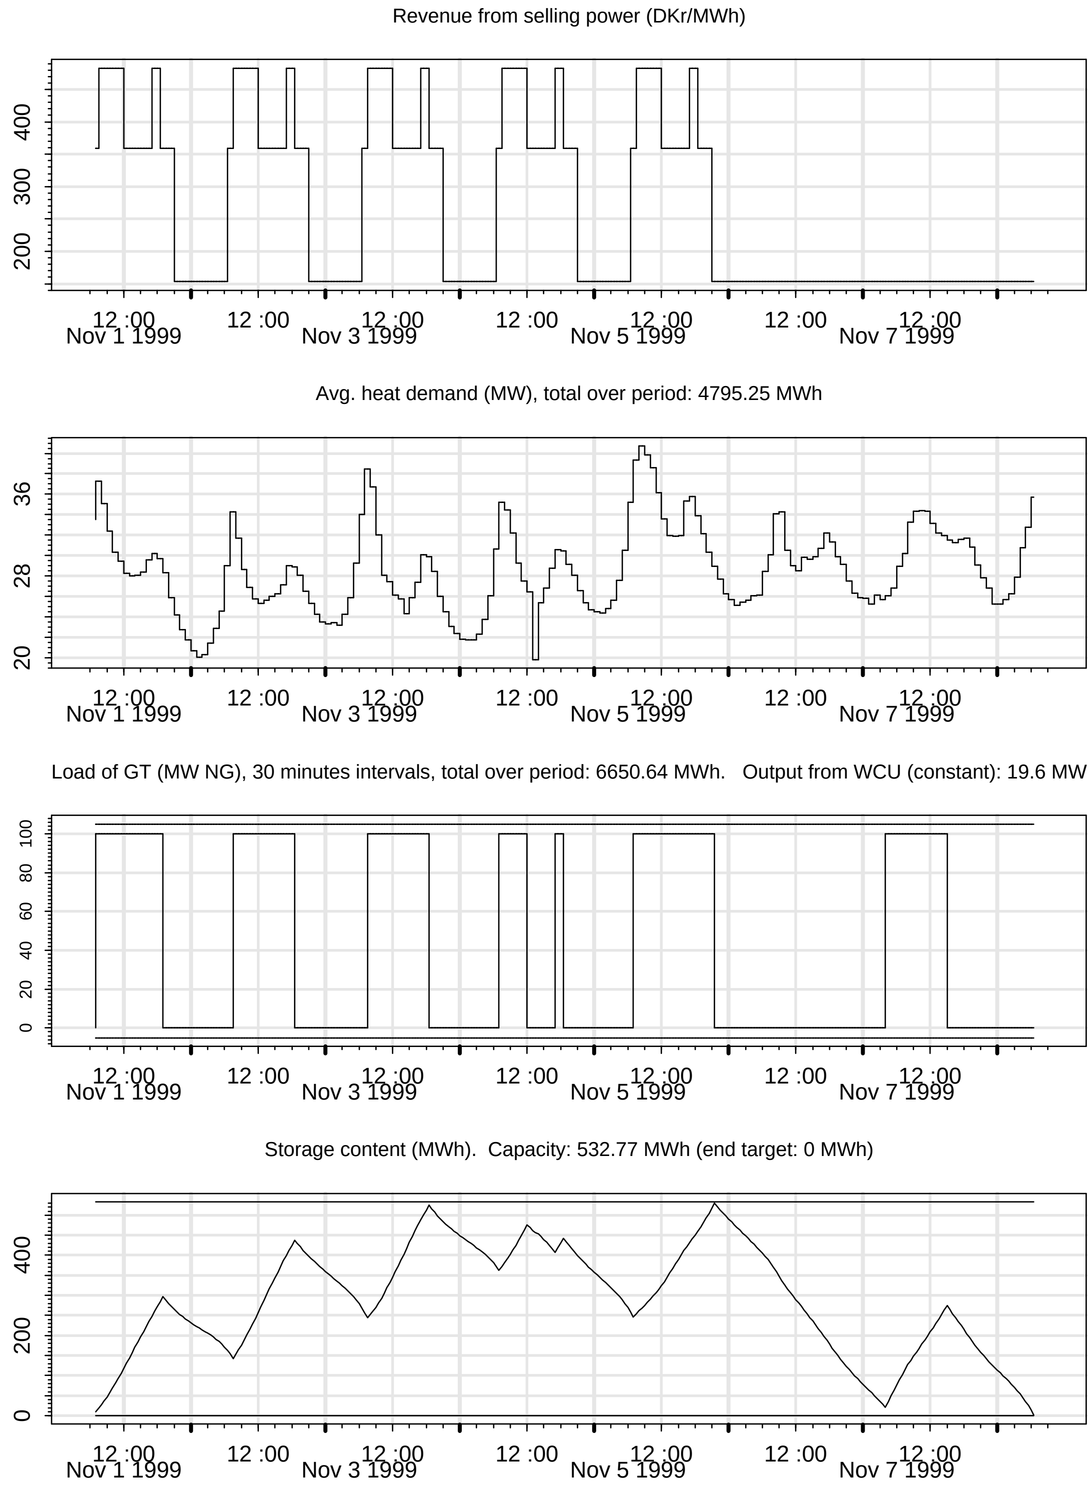

Figure 5.4: Optimal operation of gas turbine for one week starting on 1999.11.01 (criteria function: 793807 DKr).

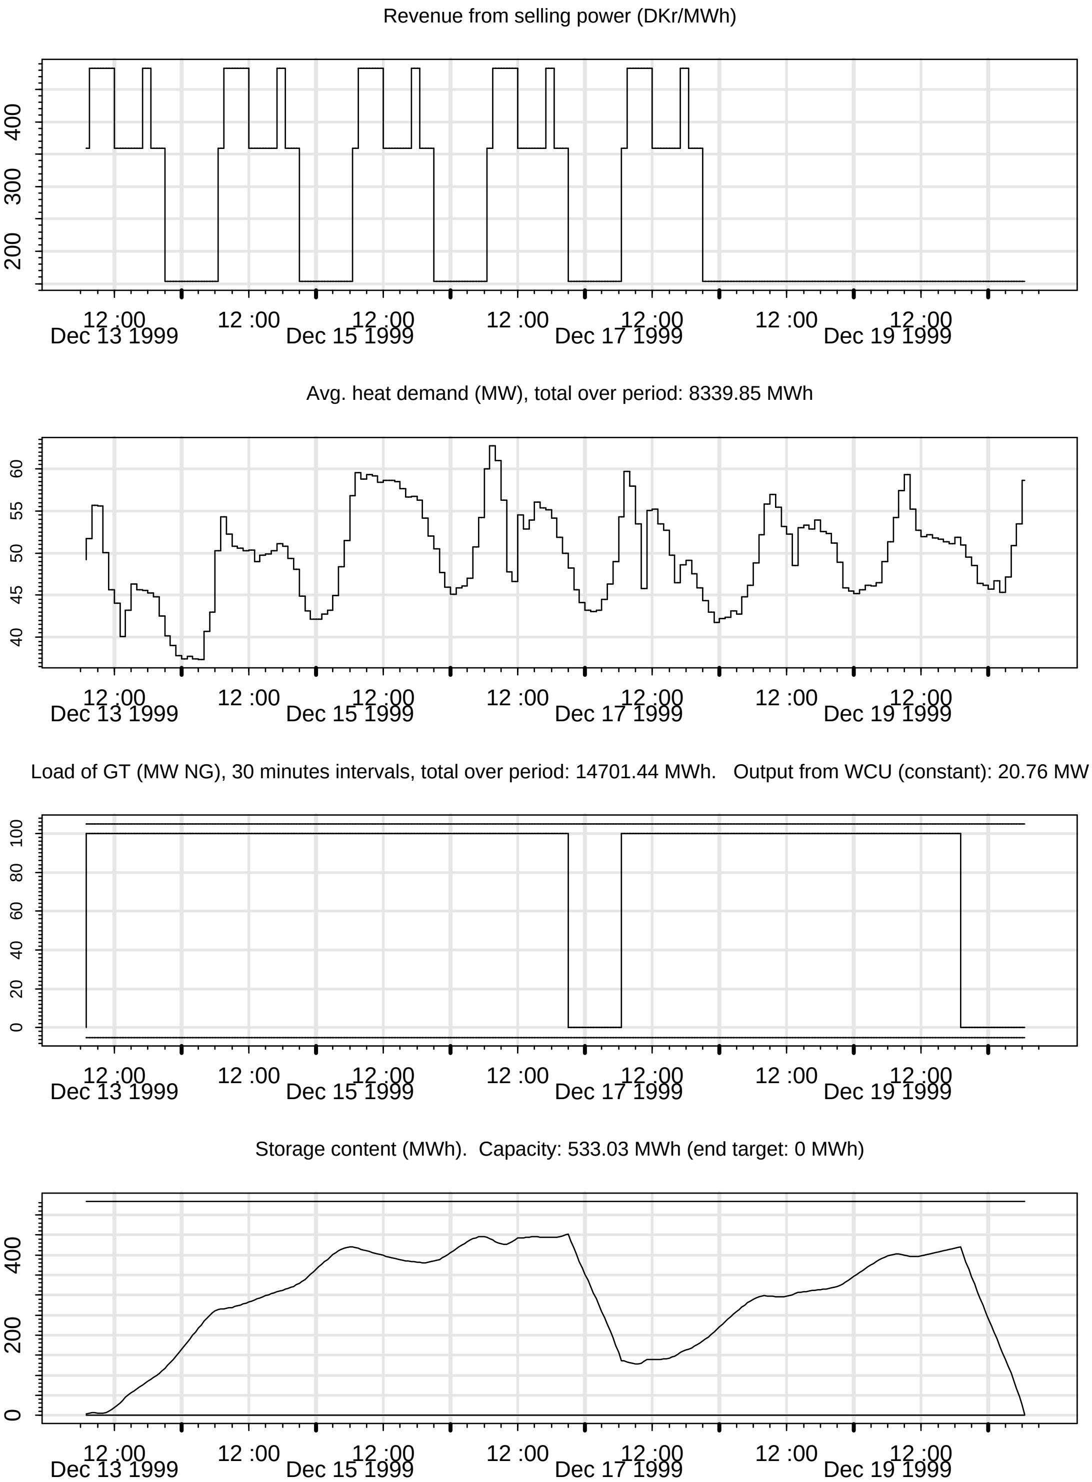

Figure 5.5: Optimal operation of gas turbine for one week starting on 1999.12.13 (criteria function: 940135 DKr).

5.2 Computation time

The examples for which results are presented in Section 5.1 uses approximately two minutes of CPU on a PC with a 450 MH_Z Intel Pentium III processor, running Linux (RedHat 6.0) and using the GNU C compiler (version 2.91) with level 3 optimization of the executable code. However, for the optimization problem considered the implementation is not optimal. This is because the implementation solves a much more complicated optimization problem in which many more restrictions to the operation of the gas turbine can be imposed. Furthermore, the implementation can optimize the operation of a gas turbine capable of varying the load. Hence, it is expected that the solution of the problem can be obtained quicker by using a more dedicated implementation. However, in the following we will consider the actual implementation only.

When solving the optimization problem it is first transformed into a discrete version by only considering a fixed number of levels in the heat storage. This operation is characterized by the interval length in MWh. For small intervals we obtain a solution close to one of the true optimal solutions. For the optimization problem starting in the morning of Monday, October 4, 1999 (cf. Figure 5.2) Figure 5.6 shows the revenue obtained (i.e. the optimal value of (4.16) on page 22) plotted against the CPU time used to find the optimal solution for different values of the discretization used for the heat storage. The discretization is chosen so that it spans the entire interval ranging from 0 MWh to the capacity of the heat storage (533 MWh) and consequently not all step sizes are feasible.

For the fine discretization of 0.1 MWh (24 minutes of CPU time) the revenue obtained is 612500 DKr, this is also true for discretizations up to 5 MWh (24 seconds) and for the discretization corresponding to 14.8 MWh (8 seconds). For 9.9 MWh (12 seconds) and 19.7 MWh (6 seconds) the revenue is marginally lower (550 DKr). Finally, for discretizations corresponding to 29.6 MWh (4 seconds), 38.1 MWh (3 seconds), and 48.5 MWh (2.4 seconds) the revenue is approximately 3000 DKr lower than for 0.1 MWh. For the period starting in the morning of Monday, November 1, 1999 (cf. Figure 5.4) there is a loss of 6000 DKr (0.8%) when going from a discretization corresponding to 9.9 MWh to one corresponding to 15 MWh. Consequently, although there is some variation between periods, a very precise solution should be obtainable using a discretization corresponding to 5 or 10 MWh. This solution can be found using 12 to 24 seconds of CPU time.

However, processors with a clock frequency of 1 GHz will be able to increase the speed with a factor of more than two. Finally, we mention that dynamic programming can easily be solved using parallel programming and the speed will be increased with a factor almost equal to the number of processors, as long as the number of processors is less than two times the number of discrete levels of the heat storage[2] and assuming that the hardware and operating system can efficiently handle multiple processors.

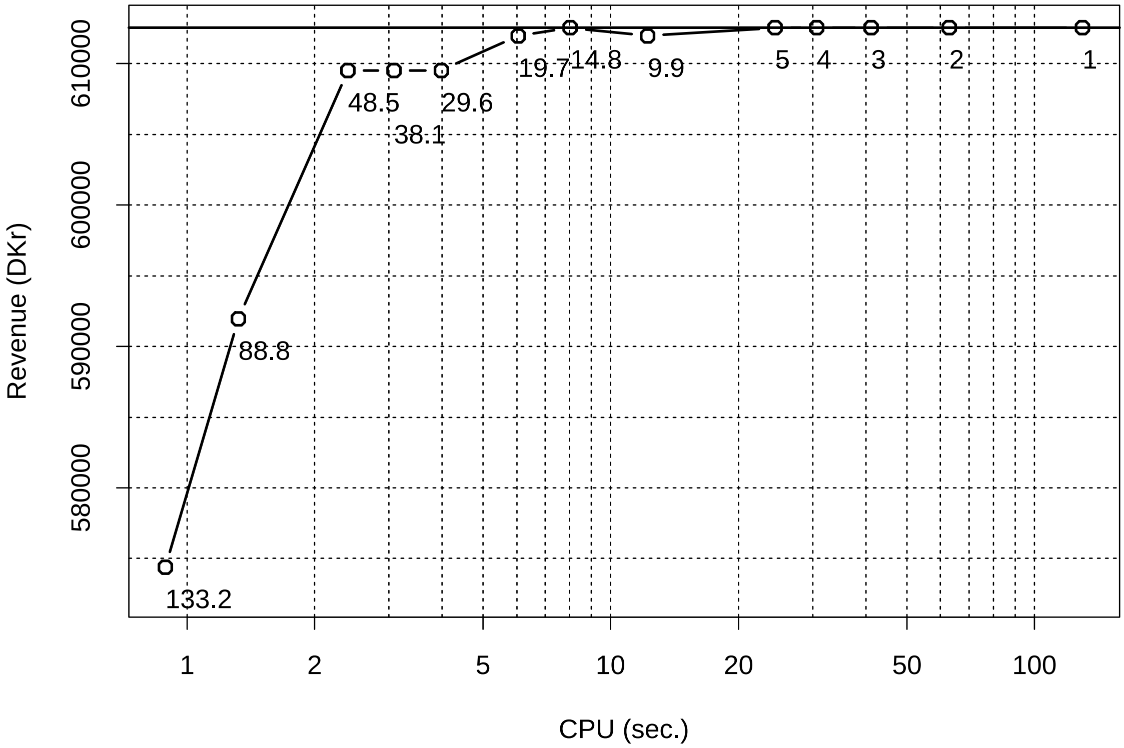

Figure 5.6: Revenue versus seconds of CPU time used on a PC with a 450 MH_Z Pentium III processor running Linux for varying degrees of discretization as indicated below each point (MWh). The horizontal line indicates the revenue obtained using steps of 0.1 MWh (approx. 24 min.). The optimization problem is the one considered in Figure 5.2. Note the log-scale on the horizontal axis.

[2]At each time step, both the cases of running and stopped gas turbine must be considered.

Chapter 6

Model of Sønderborg CHP plant – varying load

It might be considered to run the gas turbine like the one at Sønderborg Kraftvarmeværk at varying load. This will however require knowledge about the efficiencies η_{gp} and η_{gh} in (4.2) on page 18 for a range of values of F_g. Furthermore, it will require information about the implications for the maintenance costs. In this chapter, although this information is not available to us, we will illustrate how such information can be used when finding optimal schedules.

The most simple approach would be to select the load of the gas turbine each time it is started, but not changing the load while the gas turbine is running. The task of scheduling the load would then be to decide when the gas turbine should be started and determine the load for each period in which it is running.

The mathematical formulation of this optimization problem is like (4.18) on page 23, except that instead of the restriction on $F_{g,t+k\Delta t}$ in (4.18) on page 23 we use the two restrictions:

$$F_{g,t+k\Delta t} \in \{0, [f_{g,min}, f_{g,max}]\} \tag{6.1}$$

and

$$(F_{g,t+k\Delta t} - F_{g,t+(k-1)\Delta t}) \times I(F_{g,t+k\Delta t} > 0 \wedge F_{g,t+(k-1)\Delta t} > 0) = 0, \tag{6.2}$$

where $I(\cdot)$ is 1 if its argument is true and 0 otherwise. That is when the gas turbine is running it must run with a minimum load of $f_{g,min}$ and a maximum load of $f_{g,max}$, and the load must not be changed while the gas turbine is running.

In the following we assume that $f_{g,min} = 50MW$ and $f_{g,max} = 100MW$. For the efficiencies we assume that the total efficiency of the gas turbine ($\eta_{gp} + \eta_{gh}$) is 97% (as in Section 4.4). For the power-efficiency we assume that it is constant and equal to the one used in Section 4.4 for $F_g \geq 80MW$; for F_g below 80 MW we assume that it decreases linearly to 30% at 50 MW; see Figure 6.1.

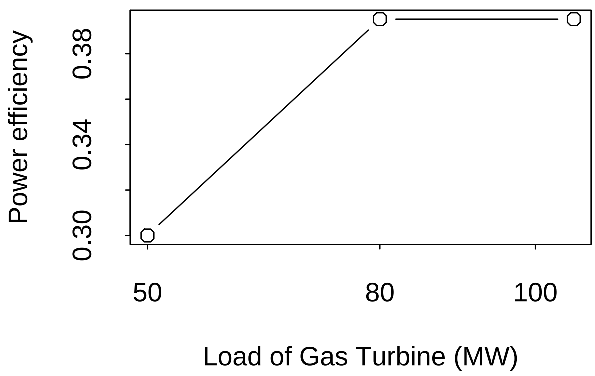

Figure 6.1: Assumed efficiency of producing power (η_{gp}) on the gas turbine against the load of the gas turbine (F_g).

The results are shown for the one-week period starting with an empty storage and a stopped gas turbine on 06:30, Monday, Oct. 1, 1999 the (period 3 in Section 5.1). A step size of 5 MWh for the storage content and 1 MW for the load of the gas turbine was used when solving the optimization problem. The resulting optimal schedule is displayed in Figure 6.2. It is seen that the schedule is almost identical to the schedule found using the original setup, cf. Figure 5.4. The most pronounced difference is that the algorithm selects a load of 80 MW natural gas for the period during the weekend. However, since the efficiency is constant for loads between 80 and 100 MW both solutions are optimal for the setting considered here. The difference in the terminal storage content accounts for the difference in the total amount of natural gas used over the period.

If the power efficiency for a load of 80 MW is increased by approximately 1.5% to 41%[1] the optimal schedule displayed in Figure 6.3 is obtained. It is seen that the solution tries to balance producing power when the price is high with the better efficiency obtained when the load is not so high. Although, the assumed dependence of the efficiency on the load may not be realistic, this example shows that small variations in the efficiencies around the maximum load may alter the optimal schedule. In fact, except for the weekend, the minimum load is 90 MW for which the power efficiency gets 40.4% or only 0.9% higher than for a load of 100 MW.

[1]The assumption is not very plausible but serves to illustrate the method.

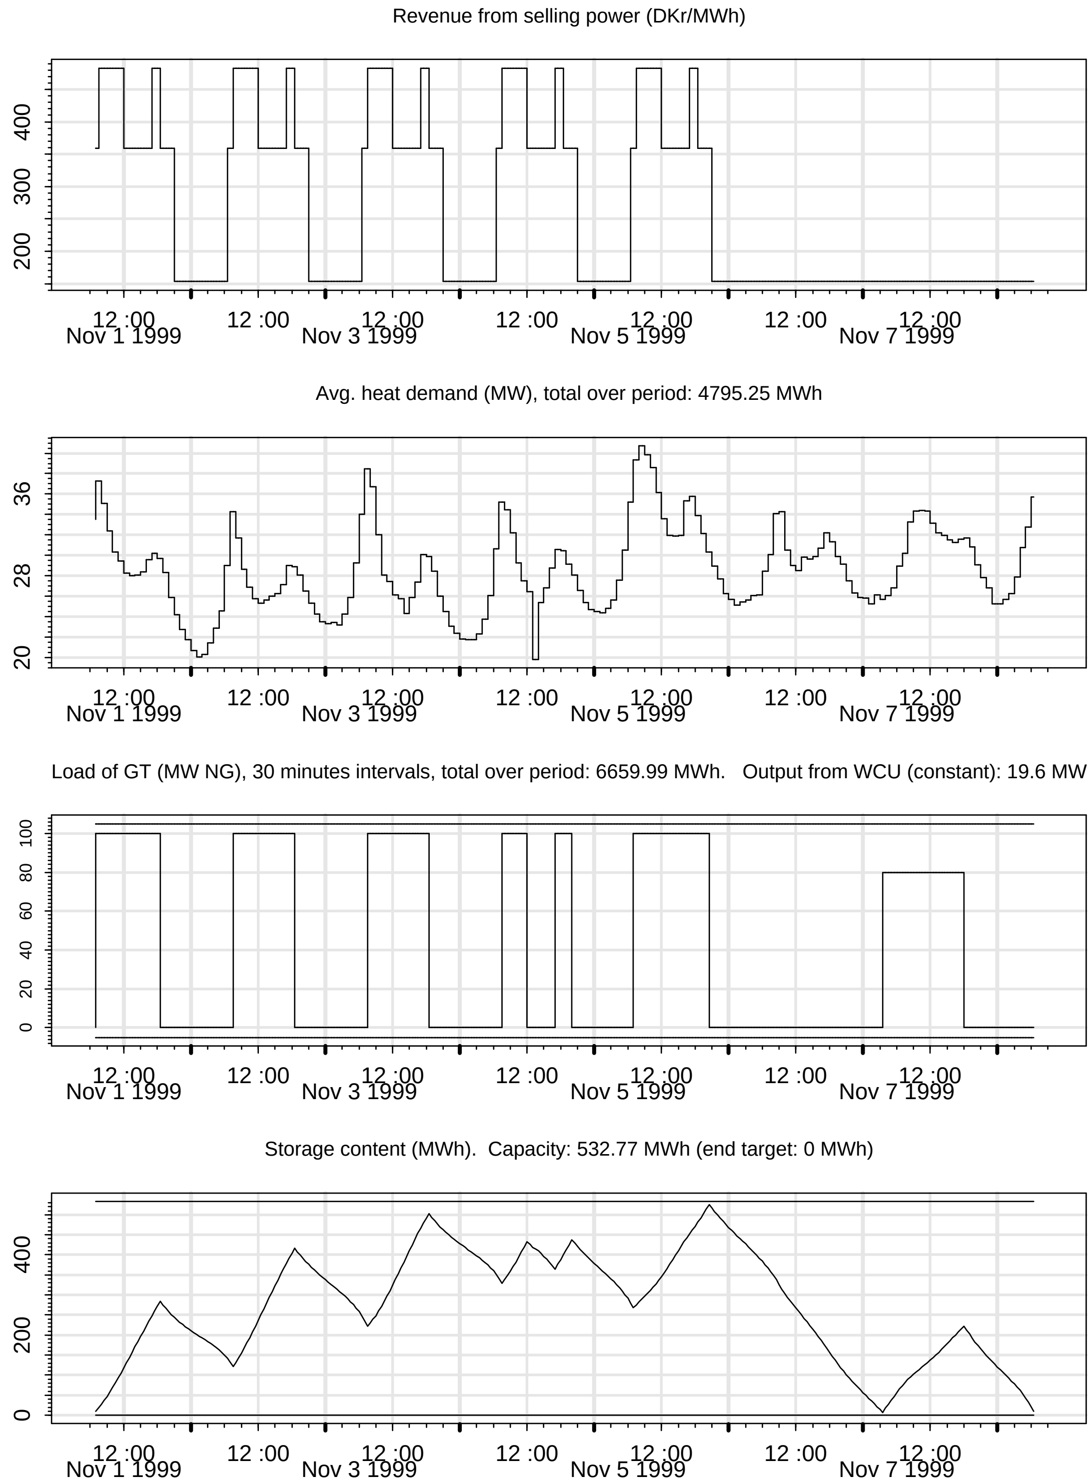

Figure 6.2: Optimal operation of gas turbine for one week starting on 1999.11.01.

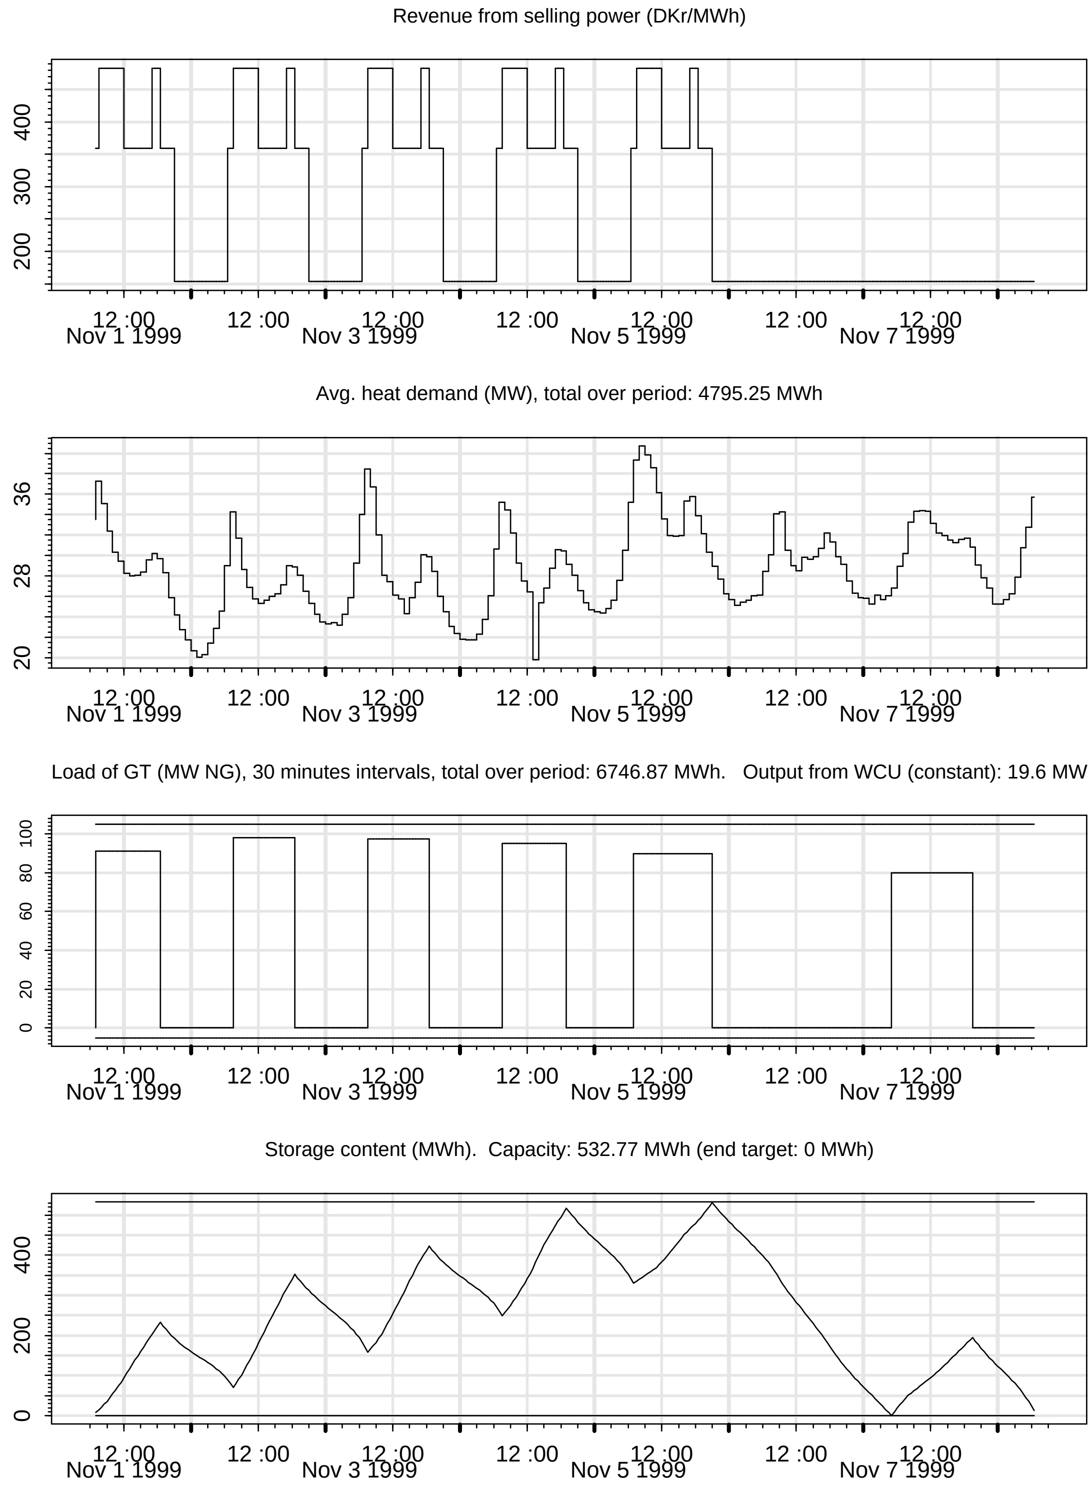

Figure 6.3: Optimal operation of gas turbine for one week starting on 1999.11.01.

Chapter 7

Stochastic effects

In practice the future heat demand, as compared to time t, $D_{t+k\Delta t};\ k = 1, \ldots, K$ is not known, but systems exist which are able to predict the heat demand with some uncertainty (Madsen & Nielsen 1997, Nielsen & Madsen 2000). To account for the stochastic effects in the decision problem a number of approaches exists.

1. Formulate the decision problem as a maximization of the expected revenue from running the plant and use stochastic dynamic programming (SDP) to solve the problem. See e.g. (Ross 1995).

2. Impose superficial bounds on the content of the heat accumulator thereby making the recovery from deviation of the actual values from the predictions possible.

3. Use information from the variance / covariance structure of the predictions to simulate, e.g. 100, possible realizations of the future heat load. Find the optimal solution for each realization and compose these solutions into a decision about how to run the plant.

Re. 1:
Maximization of the expected revenue may be regarded as the best solution to the handling of stochastic effects. However, it is very computational demanding (Chen et al. 1999) and it requires the model used in the prediction method to be build into the solution method. Comparing with (Nielsen & Madsen 2000) it is seen that this will make the state vector very large, whereby complicated methods as described in (Chen et al. 1999) is required. Furthermore, from a more practical point of view, it will link

the solution method tightly with the prediction method, making modular design and implementation difficult or impossible.

Re. 2:
This method is used in (Nielsen 1991). The approach is relatively simple but requires knowledge about the variance / covariance structure for the bounds to be calculated. A more serious reservation against the method is that it makes it difficult to investigate when the heat accumulator should be empty or full.

Re. 3:
The method requires a unique relation between the future heat demand and the optimal solution of the individual deterministic decision problems. The solution to decision problems such as (4.18) may not be unique, but if the summands in the criteria function are multiplied with a factor decreasing with increasing k it is plausible that the solution is unique (although this still needs to be investigated further). Given $0 < \alpha < 1$, and in practice close to 1, the criteria in (4.18) is then replaced with

$$\sum_{k=1}^{K} \alpha^k r_{t+k\Delta t}(F_{g,t+k\Delta t}, X_{t+k\Delta t}) + c^{(a)}_{t+K\Delta t} A_{t+K\Delta t} \tag{7.1}$$

The method does not result in a unique solution for the stochastic decision problem, but it presents a number of good solutions depending on the future heat demand. The approach leaves room for a qualified person to take other information into account and use this to decide on a plan about how to run the plant in the near future. Combined with tools for monitoring such a plan, this approach seems to provide a reasonable combination of automation and input from qualified personnel. A more detailed description of the general features of such a decision support system can be found in Section 2.2.

7.1 Simulation of heat demand

It is assumed that predictions of the heat demand are received on-line from an other system such as PRESS (Madsen & Nielsen 1997). Ideally, the model(s) used in the prediction system should be used to generate simulations of the future heat demand. However, to allow for a modular structure of software implementing these methods this path will not be followed. Instead it will be assumed that the optimization system receives the predictions and actual values of the heat load. Based on this it internally generates the simulations. Using this approach it is possible to plug-in a new prediction system without modifying the optimization system. If this changes the covariance structure of the prediction errors the optimization system should automatically learn this new structure, but the simulations can not be trusted until the structure is learned. The simulations should be generated under the assumption that the mean of the prediction errors is zero. If this is not true it should be accounted for by updating the prediction system.

In the following we assume that the time scale is rescaled so that the time step $\Delta t = 1$. Let $e_{t|s}$ denote the prediction error when predicting the heat load at time t given information up to time s. The prediction errors may be collected into a matrix in which the columns $k = 1, \ldots, K$ indicate the prediction horizon and the rows indicate the time at which the predictions are generated, i.e.

$$\begin{bmatrix} \vdots & \vdots & & \vdots \\ e_{t|t-1} & e_{t+1|t-1} & \cdots & e_{t-1+K|t-1} \\ e_{t+1|t} & e_{t+2|t} & \cdots & e_{t+K|t} \\ \vdots & \vdots & & \vdots \end{bmatrix}, \tag{7.2}$$

this constitutes a K-dimensional stochastic process in which simulations can be performed if it is modelled. However, for the purpose of generating simulations corresponding to a particular optimization problem, i.e. a particular t e.g. in (7.1), only row t in (7.2) needs to be simulated. To simulate the row the multivariate cumulative distribution function $F(e_{t+1|t}, e_{t+2|t}, \ldots, e_{t+K|t})$ must be known. Let $\mathbf{e}_t = [e_{t+1|t}, \ldots, e_{t+K|t}]^T$. If $\mathbf{e}_t$ is $MVN(\mathbf{0}, \mathbf{\Sigma})$ then an estimate of $\mathbf{\Sigma}$ is $\mathbf{S} = (1/N)\sum_{t=1}^{N} \mathbf{e}_t\mathbf{e}_t^T$, where N is the number of rows in (7.2). $\mathbf{S}$ is also called the sample covariance matrix function in lag zero of the multivariate process (7.2) (Box, Jenkins & Reinsel 1994).

Normally, prediction methods will be adaptive and therefore the estimate $\mathbf{S}$ should also be allowed to change over time. Assume that row t is the last row in (7.2). Using weights decreasing exponentially as rows become older the adaptive estimate is defined as

$$\mathbf{S}_\lambda(t) = \frac{\sum_{s=1}^{t} \lambda^{t-s}\mathbf{e}_s\mathbf{e}_s^T}{\sum_{s=1}^{t} \lambda^{t-s}}, \tag{7.3}$$

where $0 < \lambda < 1$. Since $\lim_{t\to\infty} \sum_{s=1}^{t} \lambda^{t-s} = 1/(1-\lambda)$, when t is sufficiently large ($\lambda^{t-1} \approx 0$) it holds that

$$\mathbf{S}_\lambda(t) = (1-\lambda)\left(\mathbf{e}_t\mathbf{e}_t^T + \lambda\mathbf{S}_\lambda(t-1)\right). \tag{7.4}$$

This recursion can be used to adaptively estimate $\mathbf{\Sigma}$. At time t the latest estimate of $\mathbf{\Sigma}$ is $\mathbf{S}_\lambda(t-K)$ and hence $\mathbf{e}_t$ can be simulated as a zero mean multivariate normal variable with covariance matrix $\mathbf{S}_\lambda(t-K)$. It is noted that the matrix is symmetric and hence contains $K(K+1)/2$ covariance parameters. It is possible to update some elements of the matrix before others, in practice this possibility should be considered. However, there might be some problems with the non-negative definiteness of $\mathbf{S}$.

7.2 Efficient computations

As indicated in Section 5.2 the deterministic scheduling problem can be solved using 5-10 CPU seconds on a 1 GHz PC. When solving a limited number of optimization problems, corresponding to different paths of the future heat demand, this does not pose a problem. However, for solving the optimization problem for a number of simulated heat demands the amount of CPU time used might pose a problem since it will be hard to justify conclusions regarding 5% and 95% quantiles on less that 100 simulations; this will take more than eight minutes, assuming only one CPU is available. Since the optimization problems can be solved independently it makes sense to use a multiprocessor computer. However, a number of techniques can be used to reduce the total amount of CPU time necessary.

For the discrete dynamic programming problem the amount of CPU time used is dominated by the number of combinations which must be investigated. Hence the CPU time is approximately proportional to

$$K n_g n_a, \tag{7.5}$$

where K is the number of time periods as in (4.18), n_g is the number of state levels investigated for the gas turbine ($n_g = 2$ when considering only on/off schedules), and n_a is the number of levels originating from the discretization of the amount of heat in the heat storage (for steps of 9.9 MWh; $n_a = 55$). For the optimization problems considered in Chapter 5 $K = 336$. In the following we consider how to reduce K and n_a.

The number of time periods K can be reduced by reformulating (4.18) so that the size of the time step Δt varies over the total time period considered in the optimization problem. Consider a time period starting at time t and consisting of time steps $\Delta t_1, \Delta t_2, \ldots, \Delta t_K$, so that the optimization period ends at time $t + \sum_{k=1}^{K} \Delta t_k$. Let

$$\Delta T_k = \sum_{l=1}^{k} \Delta t_l, \tag{7.6}$$

if the time intervals are selected so that the functions describing the revenue is constant during the intervals, then (4.18) can be reformulated as:

$$\begin{aligned}
&\max \sum_{k=1}^{K} r_{t+\Delta T_k}(F_{g,t+\Delta T_k}, F_{g,t+\Delta T_{k-1}}) + c^{(a)}_{t+\Delta T_K} A_{t+\Delta T_K} \qquad (7.7)\\
&\text{w.r.t.}\\
&\quad A_{t+\Delta T_k} = (1-\eta_a \Delta t_k) A_{t+\Delta T_{k-1}} + \Delta t_k \left(H(F_{g,t+\Delta T_k}) - D_{t+\Delta T_k} \right)\\
&\quad A_{t+\Delta T_k} \in [0, A_{max}]\\
&\quad F_{g,t+\Delta T_k} \in \{0, f_g\}
\end{aligned}$$

Except for the content of the heat storage $A_{t+\Delta T_k}$, variables and functions indexed by $t+\Delta T_k$ are constant during the time interval from $t+\Delta T_{k-1}$ to $t+\Delta T_k$. A similar formulation can be used for the optimization problem considered in Chapter 6.

It makes sense to base the selection of $\Delta t_1, \Delta t_2, \ldots, \Delta t_K$ on the time-varying revenue from selling power, cf. Table 4.1. For Sønderborg Kraftvarmeværk, it does not seem reasonable to allow the gas turbine to change its state during peak load periods. If the remaining periods of different revenue are split into time steps of maximum one hour then a full week is split in 148 time steps, corresponding to a 56% reduction in K and approximately the same in the amount of CPU time required.

The heat demand $D_{t+\Delta T_k}$ is the average for the time interval. If the actual heat demand is constant during the time interval then the content of the heat storage will change linearly during the time interval. However, when the actual heat demand is not constant, the content of the heat storage between $t+\Delta T_{k-1}$ and $t+\Delta T_k$ is not controlled. Therefore for large time intervals and strongly varying heat demand it is possible the a schedule found by solving (7.7) could in fact be infeasible in practice. This can however be compensated for by introducing small superficial bounds on the content of the heat storage.

The number of levels originating from the discretization of the amount of heat in the heat storage n_a can be reduced by not using the same step size for all levels of the heat stored. For simulation purposes it is suggested that one solution, corresponding to the forecasted heat demand, is found using the same step size for all levels of the heat storage. Hereafter, for the simulated heat demands the original discretization is used around the optimal content of the heat storage and a larger discretization is used further away from optimum. Using the original discretization for a band corresponding to half the heat capacity and two times the original discretization for the remaining levels results in a reduction in n_a of at least 25%.

If both ways of reducing the CPU time required is implemented, using the values used above, then the total number of combinations which must be investigated is reduced to:

$$(1 - 0.56)Kn_g(1 - 0.25)n_a = 0.33Kn_gn_a,$$

corresponding to a reduction of 67%. Consequently, using the techniques outlined above, it is expected that 100 simulations can be performed using less than 3 minutes on a 1 GHz PC, and less on a multiprocessor machine. Furthermore, as stated in Section 5.2, we also believe that it is possible to make the basic implementation more efficient.

Chapter 8

Conclusion and discussion

In this report load scheduling for decentralized combined heat and power plants are considered for the case where the fuel cost and the revenue from selling power to the transmission company may be time-varying. These plants use a district heating network as their main source of cooling and hence the ratio between the heat and power outputs are fixed by the operating conditions of the plant. To be able to deviate from this restriction a heat storage facility is often used. The load scheduling is performed in order to decide the load of each unit, whereby the charging or discharging of the heat storage is given by the heat demand.

When deciding on a load schedule only approximate knowledge about the future is available. At the present in Denmark, this uncertainty is only associated with the heat demand, but in the future revenues and costs might also be uncertain. The approach taken in this report is to explicitly describe how the total revenue from running the plant depends on the schedule for the heat and power producing units of the plant and on external information such as heat demand, costs, and revenues. Optimization theory, in this case dynamic programming, is then used to find an optimal schedule. However, in order to do so either (a) the external information must be assumed known or (b) a mathematical description of the uncertainty must be build into the optimization problem, which will then aim at maximizing the expected revenue. For (a) the optimal schedule is conditional on the external information which may not be known exactly as it is the case for the heat demand. In principle this will be solved using method (b), but other complications arise in this case. These complications includes, complexity of the mathematical formulation, non-modular software, and the computational time required. For this reason we base the approach on (a), but (b) is further discussed near the end of this chapter.

To circumvent the problems inherently associated with (a) it is suggested that the plant is equipped with

(i) an automatic on-line system for forecasting the heat demand,

(ii) an interactive decision support tool by which optimal schedules can be found given the forecasts or user-defined modifications of the forecasts, and

(iii) an automatic on-line system for monitoring when conditions have changed so much that the schedule needs revision.

In this report the focus is on methods applicable for items (ii) and (iii), whereas methods for predicting the heat demand is considered in (Nielsen & Madsen 2000). To take the uncertainties into account we suggest that the decision support system allows the operator to investigate the sensitivity of the optimal schedule to variations in the input. Furthermore, we suggest that the system is equipped with the possibility to simulate realistic realizations of the heat demand based on the actual forecast and previous forecast errors, cf. Section 7.1. By letting the system find optimal schedules for each of these realizations the operator can gain some insight into the importance of the uncertainties. Of cause not all of (i) – (iii) need to be present for the operator to gain some insight into which schedules are adequate.

In the report it is shown that for a particular decentralized combined heat and power plant, Sønderborg Kraftvarmeværk in Denmark, the calculations can be performed within 5 to 10 seconds, when considering one-week periods. And it is argued that by using a more efficient formulation schedules for 100 possible realizations of the future heat demand one week ahead can be found in less than 3 minutes. Furthermore, we believe that the calculations can be done faster by using a more dedicated implementation than the one used in this report. Finally, the methods considered allows for massive use of parallel processing.

When including a description of the uncertainties into the mathematical formulation, the complexity of the formulation originates from the fact that appropriate forecasting methods will

- use meteorological forecasts of climate variables,
- calibrate models individually for each forecast-horizon, and
- use self-tuning methods.

See also (Nielsen & Madsen 2000). Especially, the two first items complicate the inclusion of the uncertainties into the mathematical formulation, since none of the models used actually describe the stochastic properties of the prediction errors of the heat demand. Assuming the above could be solved, a tight link between the system for predicting the heat demand and the system for load scheduling would be introduced. This may be undesirable from a practical point of view since the software for load scheduling and forecasting cannot be separated into modules which can be maintained separately. Finally, we believe that the complexity of the stochastic process will require a relatively high dimension of the state vector whereby it will be very time consuming to perform the computations and methods like the ones considered by Chen et al. (1999) will be required. However, since the solution to the stochastic problem takes all future realizations of the stochastic process in to account, the computational time is maybe not so critical. This is because the problem needs only to be solved when the properties of the stochastic process changes or when the actual time gets near to the end of the optimization horizon. In practice this will probably require the stochastic scheduling problem to be solved, say, twice every week.

Further research is needed to clarify if stochastic optimization can be applied. One possible path of this research is to use a multivariate normal distribution of the prediction errors as described in Section 7.1. Especially, if the covariance matrix of the prediction errors can be described using only a few parameters this might be a feasible path of future research.

Bibliography

Bellman, R. (1957), *Dynamic programming*, Princeton University Press, Princeton.

Box, G. E. P., Jenkins, G. M. & Reinsel, G. C. (1994), *Time Series Analysis Forecasting and Control*, third edn, Prentice-Hall, Englewood Cliffs, NJ.

Chen, V. C. P., Ruppert, D. & Shoemaker, C. A. (1999), 'Applying experimantal design and regression splines to high-dimensional continuos-state stochastic dynamic programming', *Operations Research* **47**(1), 38–53.

Incropera, F. P. & DeWitt, D. P. (1985), *Fundamentals of Heat and Mass Transfer*, second edn, John Wiley & Sons.

Madsen, H. & Nielsen, T. S. (1997), 'PRESS – PRognose- og EnergiStyrings System', Department of Mathematical Modelling, Technical University of Denmark, DK-2800 Lyngby.

Madsen, H., Nielsen, T. S. & Søgaard, H. T. (1996), *Control of Supply Temperature*, Department of Mathematical Modelling, Technical University of Denmark, yngby, Denmark.

Nielsen, H. A. (1991), Driftoptimering for kraftvarmeværk med akkumulatortank, Master's thesis, Department of Mathematical Modelling, Technical University of Denmark, DK-2800 Lyngby, Denmark. Number 33/91.

Nielsen, H. A. & Madsen, H. (2000), 'Predicting the heat consumption in district heating systems using meteorological forecasts', Department of Mathematical Modelling, Technical University of Denmark, DK-2800 Lyngby, Denmark. Danish Energy Research Program 1323/98-0025.

Press, W., Teukolsky, S., Vetterling, W. & Flannery, B. (1992), *Numerical Recipes in C: The Art of Scientific Computing*, second edn, Cambridge University Press, Cambridge, UK.

Ravn, H. (1996), 'Notes for static and dynamic optimization. (In Danish: Noter til statisk og dynamisk optimering)', Department of Mathematical Modelling, Technical University of Denmark, DK-2800 Lyngby, Denmark. Lecture notes, 378 pages.

Ross, S. M. (1995), *An Introduction to Stochastic Dynamic Programming*, Academic Press, New York/London.