university of groningen

faculty of science and engineering

KWR

Reliability Analysis and Optimal Maintenance Scheduling for Heat Assets in Energy System Modelling

by

Bastiaan J. Kok

A Thesis
Submitted in Partial Fulfilment of the Requirements for the Degree
Master of Science

Faculty of Science and Engineering

Supervisors:

First	MSc. Mehran Mohebbi
Second	prof. dr. ir. Jacqueline M.A. Scherpen
Daily	dr. Juan E. Machado Martinez
Company	dr. ir. Daniël J. Bakker
	dr. ir. Dirk Vries

Groningen
June 23, 2023

Abstract

In this work, a maintenance and reliability model is developed for the virtual heat assets in the digital twin of WarmteStad's heat plant. The maintenance schedule of the assets is optimised using a maintenance optimisation algorithm that minimises the cost of maintenance and the cost of unexpected failures of the assets while imposing a lower bound on the reliability of the assets over a given time horizon. A unit commitment algorithm is developed to simulate the activation and deactivation of the heat assets based on the digital twin simulation data. The feature of modelling maintenance and random failures of the assets, and compensating for the production loss caused by downtime of the assets is included in the algorithm. Monte Carlo simulation is performed to quantify the availability of the assets and reliability of the entire heat plant. The performance of the maintenance optimisation model is assessed in terms of *substandard heat supply hours* that can potentially lead to a *heat supply failure* of WarmteStad's heat plant. Results of the experiments showed that substandard heat supply occurs in 85 out of 10,000 simulated years; hence, a heat supply failure is unlikely to occur. The observed high reliability of the heat plant is explained by the fact that WarmteStad has built-in redundancy in the number of production assets in the heat plant and that the asset with the greatest thermal production capacity is highly reliable.

Keywords— Reliability analysis, Preventive maintenance scheduling, Mathematical optimisation, Monte Carlo simulation, Multi-energy systems, Digital twin

Contents

List of Abbreviations

CHP	Combined heat and power
CM	Corrective maintenance
CSV	Comma separated value
DHN	District heating network
GH	Gas heater
HP	Heat pump
KPI	Key performance indicator
kW	Kilowatt
kWh	Kilowatt hour
MCS	Monte Carlo simulation
MT-ATES	Mid-temperature aquifer thermal energy storage
MWh	Megawatt hour
O&M	Operations and maintenance
OEM	Original equipment manufacturer
PM	Preventive maintenance
pp	Percentage point
SLA	Service level agreement
SOS2	Special ordered set of type 2
TB	Thermal buffer

Chapter 1

Introduction

Industrial systems are in general subject to deterioration through usage and exposure to environmental factors over time. This degradation ultimately leads to system failure, with safety issues, equipment damage, and unexpected machine unavailability as a consequence [1]. Reliability is a key performance measure with a profound impact on the economy and safety of industrial systems [2]. System reliability generally depends on age and usually decreases as components deteriorate. To overcome the effects of deterioration on system reliability, maintenance is performed to the system by improving the condition of its components [3],[4].

Maintenance can be classified into two main categories. The first is corrective maintenance, which is usually performed when the machine has failed [5]. The second is preventive maintenance, which corresponds to the scheduled actions performed at regular intervals to avoid machine breakdowns [6]. During the maintenance period, the system is in an unproductive state [7]. Hence, preventive maintenance activities are costly when frequently performed as they increase the system's downtime [8]. Therefore, it is desirable to reduce maintenance costs without significantly reducing the system's reliability [9]. In this research, work is done on developing a reliability assessment model and optimisation of maintenance for the industrial assets of a digital twin of a heat plant in Groningen.

KWR is a not-for-profit research institute that 'bridges science to practice' and works at the interface of science, business and society [10]. The research group Energy and Circular Systems within KWR has developed a digital twin of the heat plant of WarmteStad, a heat provider located in Groningen that operates a heat plant with combined heat and power (CHP) units, heat pumps (HP) and gas heaters (GH). These industrial machines are referred to as assets and are used to generate heat and electricity from natural gas and increase the temperature of residual heat from two nearby data centres. The generated heat is transported to the customers through a district heating network.

The digital twin is a model used to simulate the assets of the heat plant. Currently, there is no restriction to the utilisation of the virtual assets, i.e. the simulated assets in the digital twin. However, this is not an accurate representation of the physical heat plant. In the physical heat plant, assets are subjected to deterioration through usage, hence, require maintenance and experience downtime. KWR wants to improve the accuracy of the digital twin by implementing constraints on the availability of the assets. Therefore, this research is conducted to provide a method for assessing the availability of the assets in the heat plant.

The scope is limited strictly to the heat-producing assets. A maintenance model is developed to optimise the maintenance schedule of the heat assets with respect to costs, while imposing lower bounds on the system reliability. Based on this optimal maintenance schedule, the availability of the assets in the virtual heat plant is assessed. Furthermore, the maintenance costs can be quantified more specifically as the number of maintenance actions is computed by the model. This allows for cost optimisations of the heat plant by relaxing the constraint imposed on the reliability of the heat plant.

The contribution of this research is the novelty to combine maintenance optimisation and Monte Carlo simulation for the reliability analysis of heat assets in energy system modelling. In this research, maintenance optimisation is applied to the assets in a virtual heat plant of a digital twin. Monte Carlo simulation is used to assess the availability of the assets and to quantify the reliability of the entire heat plant.

In the last decades, numerous papers have been published on preventive maintenance modelling and optimisation. In [11], the authors introduce hazard rate and age improvement factors for a sequential preventive maintenance policy. It is assumed that the failure time follows a Weibull distribution, and preventive maintenance reduces the age of the system. [12] generalises an imperfect maintenance optimisation problem to multi-state systems, where reliability is defined as the ability to satisfy a given demand. The effective age reduction concept characterises the imperfect preventive maintenance actions. A genetic algorithm, an evolutionary search algorithm, is used as an optimisation technique to obtain the optimal sequence of maintenance actions. [13] proposes two models to optimise the maintenance interval for multi-component systems connected in series and parallel. One model minimises costs subject to a reliability constraint. The other model maximises reliability subject to a budget constraint. The model consists of elementary analytical equations and is solved using Microsoft Excel solver that uses a generalised reduced gradient algorithm. [8] studies preventive maintenance and renewal scheduling for multi-unit systems. An integrated optimisation method is developed to schedule maintenance and renewal activities by grouping them and finding an optimal balance. The problem is modelled as a pure integer linear program that minimises maintenance and renewal costs and downtime costs over a planning horizon. In [7], a preventive maintenance schedule for high-speed trains is optimised using a binary programming model. The problem is solved using a meta-heuristic algorithm called simulated annealing, a stochastic search algorithm. The performance of simulated annealing is compared to an exact method using Gurobi solver. It is observed that simulated annealing is better suited for large-scale problems in terms of solving time and that the solution gap between the two methods is almost negligible. In [14], a district heating network is designed through a thermofluid dynamics and reliability modelling approach. Unexpected failures are modelled using an exponential distribution and the reliability of the flow network is assessed through Monte Carlo simulation.

The outline of this report is structured as follows: In Chapter 2, a problem analysis is performed and the design goal of the thesis is presented. Chapter 3 describes the materials and tools, such as mathematical model formulations and data required for this research. In Chapter 4, the operationalisation of the maintenance and reliability modelling is explained. Chapter 5 describes the scenarios for the three experiments and provides the results of the experiments, including an analysis. The discussion and future work are provided in Chapter 6 and the final conclusion is presented in Chapter 7.

Chapter 2

Problem Analysis

This chapter presents the problem analysis that forms the basis of the research. First, a research context is provided, where the functioning of WarmteStad's heat plant and the different assets is explained. Then, the system description and scope are presented followed by the goal. A conceptual model is used to describe the relationship of the elements of the artefact developed during this research. The chapter ends with a review of the business case and an explanation of functioning of the thermal buffer.

2.1 Research Context

WarmteStad is a heat provider situated in Groningen. Their goal is to contribute to a more sustainable future for Groningen by providing sustainable heat to the customers through a district heating network (DHN). This DHN is under development and will eventually be expanded to provide heat to 12,500 customers in 2026 [15]. The heat is generated in a heat plant located at Zernike in the north west of Groningen, where residual heat from two nearby data centers is used as one of the heat sources. In the future, solar thermal energy will be used as an additional sustainable heat source in combination with mid-temperature aquifer thermal energy storage (MT-ATES). The combination of solar thermal and MT-ATES is crucial, as solar energy is widely available in the summer when the heat demand is low. However, in the winter, the converse is true; the solar thermal yield is low, and heat demand is high. Therefore, MT-ATES is required as a seasonal buffer.

The heat plant is set up to produce heat as follows: heat pumps extract thermal energy from a low-temperature residual heat supply from two nearby data centres QTS and Bytesnet. The electricity for the heat pumps is generated using gas-fired combined heat and power (CHP) units. These units consist of a large reciprocating gas engine that is connected to a generator. The electricity is used for the heat pumps or can be provided to the electricity grid. The thermal energy produced in the process is used for the DHN. A buffer that can store thermal energy for one day is used for smoothing the heat demand profile such that the heat-producing assets can work at a near-constant and efficient production capacity.

The expansion of the DHN has the consequence that WarmteStad must increase its thermal energy production capacity to guarantee reliable heat supply in periods of high demand. As a result, the heat plant must be expanded by deploying more heat-producing assets and increasing the size of the thermal buffer. Due to the complexity of the energy system with varying heat demand and production capacity, a digital twin can be used to optimise the future design of the heat plant. A digital twin is a digital representation of a physical entity that enables the same functional services as its physical counterpart [16]. Driven by data and modelling, the digital twin can perform simulation and optimisation of the heat plant with respect to the heating strategy, control of the assets, efficiency and costs [17]. KWR is working on the development of a digital twin for WarmteStad [18]. The digital twin is modelled using graph theory, where the nodes in the graph represent the different assets, and the edges that connect the nodes represent the thermal energy flow between those assets. A schematic overview of the heat plant is provided in Figure 2.1.

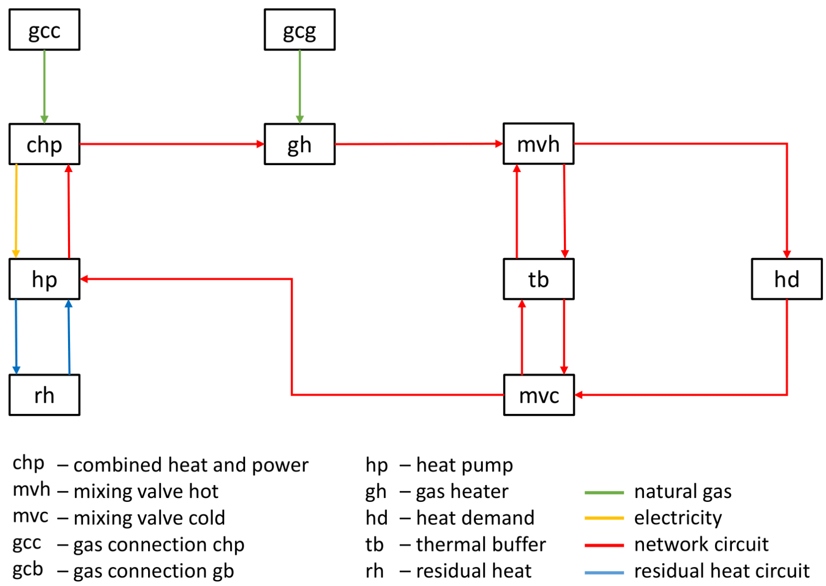

Figure 2.1: A schematic overview of the structure of the assets in WarmteStad's heat plant, modelled in the digital twin. The system consists of two main circuits: one circuit that provides residual heat from the data centres to the heat pumps, and one network circuit that connects the heat assets, thermal buffer and DHN. Energy is provided to the CHP and GH by the gas grid in the form of natural gas.

Currently, only the fundamental physics that describes the working principle of the physical assets is modelled for the assets in the digital twin. This means that these virtual assets function the same as their physical counterpart in terms of production capacity and efficiency. However, what is not included is the downtime associated to maintenance or unexpected failures of the assets, meaning that the asset availability in the digital twin is 100%. The availability of the assets in an energy system is the ratio of the up-time to the up-time plus downtime. The reliability of the system is determined by the ability of the system to provide an adequate supply of energy for a given period [19].

In the physical heat plant, assets are subjected to degradation by wear through usage. Hence, the assets require maintenance on a regular basis. The maintenance protocols for the assets handled by WarmteStad are primarily based on operating time, and some inspections are performed based on calendar time. More detailed information on the current maintenance protocols is provided in Subsection 3.3.3.

Performing maintenance implies that the assets in the physical heat plant experience downtime throughout their operational life. KWR wants to include these features in the digital twin, as the availability of the assets has a profound impact on the reliability of the entire energy system.

2.2 System Description and Scope

The system for this research is the digital twin of WarmteStad's heat plant. However, in this digital model of the heat plant, the assets are not modelled individually. Though, individual assets are required for the maintenance and reliability modelling. Therefore, the schematic overview of the digital twin in Figure 2.1 is modified to the one given in Figure 2.2. Here the individual assets are considered, which are four CHP units, five heat pumps and three gas heaters.

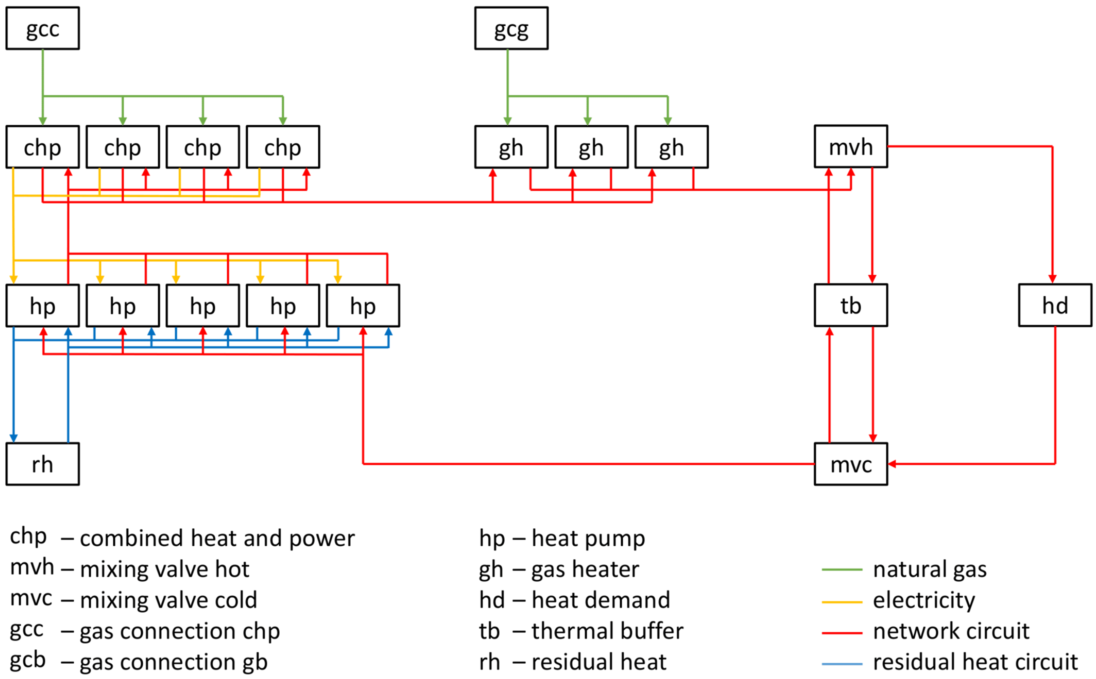

Figure 2.2: The system diagram for this research. The blocks and arrows are elaborated in the legend.

The scope of this research is limited to the three production assets: the CHP units, heat pumps and gas heaters. These three assets are the most complex assets in the heat plant, in terms of their number of components and moving parts. Hence, they are most prone to failure. Therefore, developing a reliability and maintenance model for these assets will contribute the most to assessing the reliability of the entire heat plant.

2.3 Goal Statement

The goal of this research is to develop a maintenance and reliability model for the virtual assets in the digital twin of WarmteStad's heat plant. Optimising the maintenance intervals of the assets with respect to costs and reliability ensures cost-optimal operation of the heat plant with improved up-time. Furthermore, uncertainty in the reliability of the assets is quantified through the application of Monte Carlo simulation. By analysing the results of ten thousand simulation runs, the likelihood of heat supply failure is assessed through statistical inference.

2.4 Conceptual Model

A conceptual model is developed to display the interrelationship between the different elements that make up the maintenance and reliability model to satisfy the design goal and is presented in Figure 2.3.

The digital twin is used to generate simulation data of the heat plant for a certain scenario. This data includes the hourly average heat production of the assets in kWh. The data is pre-processed to extract the hourly production and utilisation of the individual assets, as it is not an output variable in the simulation data.

Then, the asset utilisation information is used in a unit commitment algorithm to commit the required number of assets to satisfy the heat demand and keep track of their operating time. The reliability of the assets is based on the operating time and a parametric lifetime distribution of the assets.

The maintenance interval is determined by the maintenance optimisation algorithm. Here, an optimal maintenance schedule is computed based on the optimisation of a cost function that captures the costs of corrective and preventive maintenance, while imposing lower bounds on the reliability of the assets. When maintenance is due, preventive maintenance is scheduled, and the asset is taken out of operation.

Lastly, the model also includes unexpected failures of the assets. These failures are stochastic and occur according to the failure distribution corresponding to the assets. When an unexpected failure occurs, the asset is out of operation, and corrective maintenance is applied. Monte Carlo simulation is performed to quantify the risk of a heat supply failure of the entire heat plant, by iterating the maintenance and reliability model 10,000 times.

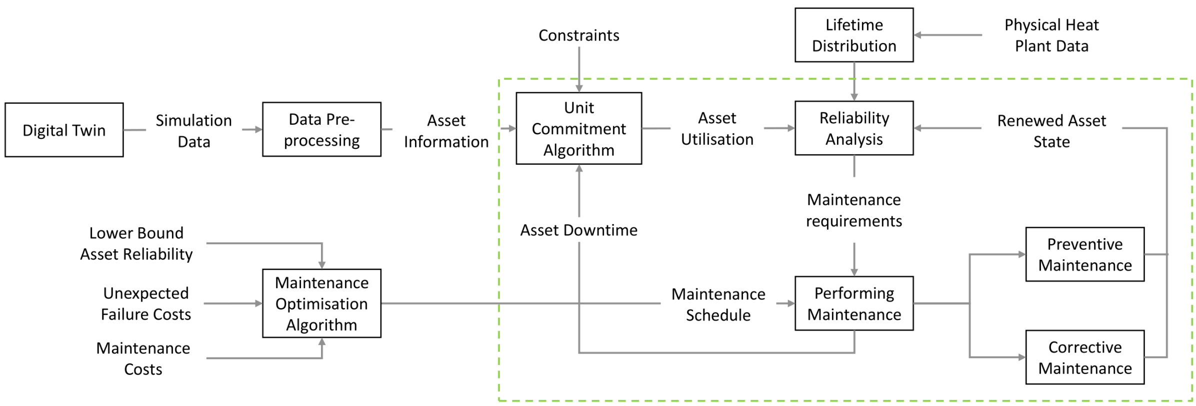

Figure 2.3: A conceptual model that describes the relationship between the different elements of the maintenance and reliability model. The green frame demarcates the part of the maintenance and reliability model that is iterated during MCS.

2.5 Risks of Heat Supply Failure

WarmteStad is primarily focused on supplying heat reliably, as they do not want to leave their customers in the cold. To emphasise the importance of this core task of heat suppliers, the legislator is working on a heat law 2.0 where reliable heat supply is enacted in the law. The legislation is incentivised by attributing a financial penalty to the event of a heat supply failure under special conditions.

When the event of a heat supply failure occurs, WarmteStad is obliged to financially compensate its customers, which can lead to a significant expense. The implications of a heat supply failure are quantified using the number of substandard heat supply hours. This is the number of consecutive hours WarmteStad fails to supply heat to its customers.

According to a concept version of the heat law 2.0, a heat supplier has to compensate its customers if[1]:

- The heat supply failure lasts 8 to 12 hours. The heat supplier has to compensate those affected by it 35 EUR.
- For each subsequent period of 4 hours, the heat supplier has to compensate 20 EUR.
- No compensation is necessary when:
  - The heat supply failure lasts less than 24 hours and is the first failure that year.
  - The heat supply failure is due to an extreme situation that cannot be attributed to the heat supplier.
  - The heat supply is interrupted due to planned work.

Assessing the expense attributed to a heat supply failure is a difficult task as it requires knowledge of the Dutch legal system, which is beyond the scope of this research. Therefore, quantifying the costs attributed to a heat supply failure is not part of this research. Though, the amount of substandard heat supply is quantified in this research. Based on this information, an estimation of the risk of incurring costs attributed to a heat supply failure is possible.

2.6 Thermal Buffer

The thermal buffer at WarmteStad consists of two large tanks of 750 m^3 each. For simplicity, the thermal buffer is considered as one large tank of 1500 m^3. In the top of the buffer, the thermal fluid is stored at a maximum temperature of 93 °Celsius. In the bottom of the buffer, the thermal fluid has a minimum temperature of 60 °Celsius. With this 33 °C temperature difference, the thermal buffer stores roughly 58 MWh of thermal energy when completely saturated. This is roughly 10% of the average energy demand during the top 10% highest energy demand days.

A schematic representation of the buffer is presented in Figure 2.4. A thin layer with a high temperature gradient called a thermocline is in between the “hot” and “cold” top and bottom of the buffer. The level of the thermocline indicates the saturation of the buffer. The energy content of the thermal buffer is a variable computed by the digital twin. Based on this information, the amount of substandard heat supply can be quantified by subtracting the

[1]The criteria are taken from a KWR internal document.

amount of lost production from the energy content at the beginning of the period.

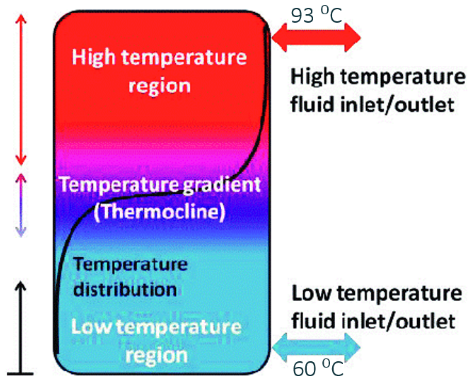

Figure 2.4: Schematic representation of WarmteStad's thermal buffer with a temperature gradient in its thermal fluid of 33 °Celsius. The left right double arrows indicates the inlet/ outlet of the buffer. Figure retrieved from [20] and modified by author.

Chapter 3

Methods and Materials

This chapter provides a theoretical background to the maintenance optimisation problem and how the problem is solved. Then, Monte Carlo simulation (MCS) is introduced to provide a method for quantifying the reliability of the heat plant through statistical methods. The chapter ends with an analysis of the data available during this research. This includes data from the digital twin, asset production data and WarmteStad's maintenance information.

3.1 Maintenance Schedule Optimisation

In a repairable system, maintenance is required to keep the reliability of the system within an acceptable range. When a repairable unit has failed, it is restored to an operating condition by applying corrective maintenance to its components without replacing the entire unit [21]. Due to the costs attributed to maintenance actions and downtime of the system, it is desired to optimise maintenance for a specific time horizon. In this section, the underlying mathematics for maintenance optimisation are presented. The optimisation model is based on the ones described in [22] and [23]. Here, a maintenance schedule is optimised for a system of components for a specific time horizon, where costs attributed to maintenance and unexpected failure of the system are minimised while imposing lower bounds on the system's reliability.

3.1.1 Mathematical Formulation

Notation

A. Parameters

N: number of assets
T: planning horizon
J: number of intervals
β: shape parameter Weibull distribution
η: scale parameter Weibull distribution
α: improvement factor maintenance
F: cost of unexpected failure of an asset
M: cost of maintenance
R_{min}: lower bound asset reliability

B. Decision variables

$X_{i,j}$: effective age of asset i at start of period j
$X'_{i,j}$: effective age of asset i at end of period j

$$m_{i,j} = \begin{cases} 1 & \text{if asset } i \text{ is maintained at period } j \\ 0 & \text{otherwise} \end{cases}$$

Life time modelling

The lifetime distribution of the assets is assumed to follow a Weibull distribution, as it is the most commonly used distribution to model the age of industrial machines and components and provides a good prediction for many types of lifetime distributions [24]. A Weibull distribution can be fitted to both complete and right censored data, i.e. data of assets that have not yet failed [25]. Statistical estimation methods such as the maximum likelihood estimate, least square regression and method of moments can be used to estimate the parameters of the Weibull distribution [26], [27], [28]. The probability density function of a two-parameter Weibull distribution is

$$f(t) = \beta\eta^{-\beta}t^{\beta-1}\mathrm{e}^{-\left(\frac{t}{\eta}\right)^{\beta}}, \tag{3.1}$$

where $\beta > 0$ is the shape parameter, $\eta > 0$ is the scale parameter and $0 \leq t \leq \infty$. The shape parameter of a Weibull distribution determines the shape of the hazard rate function. With $0 < \beta < 1$, the hazard rate is a decreasing function of time. This indicates early life failures, so-called "infant mortality". With $\beta = 1$, the hazard rate is constant over time, indicating the "useful life" period of a component. With $\beta > 0$, the hazard rate is increasing; hence the probability of failure increases with time [29].

The objective is to find a schedule of future maintenance actions for each asset over a specific time horizon $[0, T]$. The interval is divided into J discrete intervals of length T/J. At the end of period j, the asset is either maintained or no action is performed. The notion of virtual age is introduced to model the effect of maintenance actions, which was first introduced in [30]. It is assumed that an asset starts with zero virtual age when put into operation and will experience degradation throughout its operational life. Maintenance actions performed on an asset are assumed to remove these damages using an improvement factor, effectively reducing the system's age. The real age of a system is the time elapsed since the asset was put into operation [30]. Throughout this research, the notions virtual age and effective age are used interchangeably.

For simplicity, it is assumed that maintenance activities are instantaneous, i.e., the time required to perform maintenance is negligible with respect to the length of the interval and is, therefore, zero in the optimisation problem. However, a cost associated to maintenance actions is considered. The initial age for every asset is set to zero. Then, let $X_{i,j}$ denote the effective age of asset i at start of period j, and $X'_{i,j}$ denotes the effective age of component i at the end of period j, thus

$$X'_{i,j} = X_{i,j} + \frac{T}{J} \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T. \tag{3.2}$$

Effect of Maintenance

When an asset i is maintained at the end of period j, the maintenance action effectively reduces the age of that asset in the next period as

$$X_{i,j+1} = \alpha X'_{i,j+1} \quad \text{for } i = 1, \ldots, N; \quad j = 1, \ldots, T; \quad \text{and } 0 \leq \alpha \leq 1. \tag{3.3}$$

Here, α denotes the improvement factor of the maintenance action, effectively reducing the age of the asset. For $\alpha = 0$, the effective age is reduced to zero, implicating that the asset is "as-good-as new". When $\alpha = 1$, the maintenance action has no effect, meaning that the asset is "as-bad-as-old". The maintenance action at the end of period j leads to an instant failure rate reduction. When maintenance is performed to component i in period j, a constant cost M_i is incurred at the end of the period.

Do nothing

If no maintenance action is performed in period j, the effective age of the assets propagates through the periods as

$$X'_{i,j} = X_{i,j} + \frac{T}{J} \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T, \tag{3.4}$$

$$X_{i,j+1} = X'_{i,j} \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T. \tag{3.5}$$

Cost of unexpected failure

When optimising the maintenance schedule for a system, the costs caused by unexpected failure must be incorporated into the objective function. As the age of a system increases, a higher rate of occurrence of failure is inevitable, resulting in higher costs of unexpected failure. Conversely, when the effective age of an asset is sufficiently small, the probability of an asset failure is low, thus yielding a lower cost of failure. To take the costs of an unexpected failure into account, the probability of an asset failure P_f is calculated in each period for each asset in the system as

$$P_f[N_{i,j}] = \int_{X_{i,j}}^{X'_{i,j}} f(t)\,\mathrm{d}t \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T. \tag{3.6}$$

Under a Weibull distribution for the lifetime of an asset, the probability of failure occurring in a given time interval is given by

$$\begin{aligned} P_f[N_{i,j}] = \int_{X_{i,j}}^{X'_{i,j}} \beta\eta^{-\beta}t^{\beta-1}\mathrm{e}^{-\left(\frac{t}{\eta}\right)^{\beta}}\,\mathrm{d}t = \mathrm{e}^{-\left(\frac{X_{i,j}}{\eta}\right)^{\beta}} - \mathrm{e}^{-\left(\frac{X'_{i,j}}{\eta}\right)^{\beta}} \\ \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T. \end{aligned} \tag{3.7}$$

The cost attributed to an unexpected failure of an asset at time i is denoted by F_i in units of cost per failure event. Hence, the costs attributed to the failure of a component i in period j is given by

$$F_{i,j} = F_i \cdot P_f[N_{i,j}] = F_i\left(\mathrm{e}^{-\left(\frac{X_{i,j}}{\eta}\right)^{\beta}} - \mathrm{e}^{-\left(\frac{X'_{i,j}}{\eta}\right)^{\beta}}\right) \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T. \tag{3.8}$$

Component reliability

When the lifetime of a component follows a Weibull distribution, the reliability of that component at any given time is the probability that the component has not yet failed. The reliability function is given by [29]

$$R(t) = \lim_{k\to\infty} \int_t^k \beta\eta^{-\beta}t^{\beta-1}\mathrm{e}^{-\left(\frac{t}{\eta}\right)^{\beta}}\,\mathrm{d}t = \mathrm{e}^{-\left(\frac{t}{\eta}\right)^{\beta}}. \tag{3.9}$$

In the optimisation model, the assets are subjected to a minimum reliability constraint. The reliability of an asset over a given time horizon can be computed by taking the product of the reliability in the intermediate time steps [29]. Using (3.9), the reliability of an asset over time is given by

$$R_i(X'_{i,j}) = \prod_{j=1}^{T} \mathrm{e}^{-\left(\frac{X'_{i,j}}{\eta}\right)^{\beta}} \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T. \tag{3.10}$$

3.1.2 Mathematical Optimisation Model

With the formulations introduced in the previous section, it is possible to set up a mathematical optimisation problem that finds an optimal preventive maintenance schedule for multi-component systems. The objective is to minimise the preventive and corrective maintenance costs, while imposing a lower bound on the asset reliability.

Problem formulation

$$\text{minimize} \quad \sum_{j=1}^{T} \left(F_{i,j} + M_i \cdot m_{i,j}\right) \tag{3.11a}$$

subject to

$$X_{i,1} = 0 \quad \text{for } i = 1, \ldots, N, \tag{3.11b}$$

$$X_{i,j} = (1 - m_{i,j-1})X'_{i,j-1} + m_{i,j-1}\alpha X'_{i,j-1} \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T, \tag{3.11c}$$

$$X'_{i,j} = X_{i,j} + \frac{T}{J} \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T, \tag{3.11d}$$

$$\prod_{j=1}^{T} \mathrm{e}^{-\left(\frac{X'_{i,j}}{\eta}\right)^{\beta}} \geq R_{\min}, \tag{3.11e}$$

$$m_{i,j} = 0, \quad \text{or } m_{i,j} = 1 \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T, \tag{3.11f}$$

$$X_{i,j}, X;_{i,j} \geq 0 \quad \text{for } i = 1, \ldots, N; \quad \text{and } j = 1, \ldots, T. \tag{3.11g}$$

The objective function is a cost function that sums the cost of unexpected failures of the assets and the cost of preventive maintenance over the time horizon. The first constraint sets the initial age of the assets to zero. The second and third constraints compute the effective age of the assets recursively. The fourth constraint is the model's main constraint and ensures that an asset's reliability is greater than or equal to the minimum required reliability. Constraints five and six restrict the decision variables to be binary and positive.

Solving the optimisation problem

The optimisation problem is a nonlinear mixed-integer programming model due to the binary decision variables $m_{i,j}$ and the nonlinear term in the objective function and reliability constraint. The solution space for a system of N components and T intervals grows exponentially as $2^T N$. Enumerating all possibilities is no viable option for solving the problem, as it becomes computationally intractable quickly as T increases. Therefore, more efficient techniques for solving such problems are required. Solvers for mixed integer programs do not explicitly examine every possible solution but instead examine a subset of possible solutions and use optimisation theory to prove that the optimal solution is found [31]. While providing an extensive review of the literature on solving mixed-integer programs is beyond the scope of this research, it is essential to emphasise that these types of problems are inherently difficult to solve. The fact that this problem is nonlinear and mixed-integer further reduces the number of solvers that can handle this type of problem.

3.2 Monte Carlo Simulation

Various reliability analysis methods exist for energy systems [32], such as analytical models, Markov models or Monte Carlo modelling. From these methods, MCS is a powerful tool for modelling the reliability of engineering systems in order to quantify the uncertainty and risk associated to the performance and reliability of the assets [33].

Assessing the system reliability through analytical models is often too restrictive, as it relies on simplified assumptions, and interdependencies are easily overlooked. An alternative approach could be based on Markov models, which takes into account a wide range of dependencies [32]. Compared to analytical methods and Markov processes, MCS can handle more reliability evaluation conditions, making it more suitable for large-scale systems [34].

The principle of this method is the generation of certain random events in the model to create a realistic lifetime scenario of the system. Reliability analysis aims to evaluate the probability of failure in an engineering system. Using MCS, the uncertainty about the probability of failure can be formally quantified through Bayesian statistics [35]. Furthermore, MCS can be used to perform a sensitivity analysis on the input parameters to measure how uncertainty in the model parameters affects the system's reliability [36].

The stochastic failure behaviour of the assets in WarmteStad's heat plant is modelled using a Weibull distribution. The random failure times of the assets are computed by sampling from its probability density function, using the *inverse transform method.* This gives one outcome for the failure behaviour of the assets and does not provide comprehensive insight into the reliability of the entire heat plant. Therefore, MCS is performed by iterating this process numerous times, generating a wide range of possible outcomes. From this information, it is possible to assess the reliability metrics of the heat plant through statistical inference.

3.2.1 Inverse transform method

A random variate is the outcome of a random variable following a given distribution. Random number generators essentially provide random variates following a uniform distribution in the interval between $[0,1]$. The technique used for generating random variates that follow a Weibull distribution is the inverse transform method [37]. Its procedure is given in (3.12) to (3.15).
Consider the Weibull distribution from (3.1). Its cumulative probability function is given by

$$F(x) = 1 - \mathrm{e}^{-(\frac{x}{\eta})^{\beta}}. \tag{3.12}$$

By the inverse transform method, it follows that

$$U = F(x) = 1 - \mathrm{e}^{-(\frac{x}{\eta})^{\beta}}, \tag{3.13}$$

such that

$$X = F^{-1}(U) = \eta(-\ln(1-U))^{\frac{1}{\beta}}. \tag{3.14}$$

Since $(1-U)$ distributes uniformly in the same way as U in the interval between $[0,1]$, it follows that

$$X = \eta(-\ln(U))^{\frac{1}{\beta}}, \tag{3.15}$$

where U is a uniformly distributed random variable, and X follows a Weibull distribution.

3.2.2 Convergence

The MCS has converged when the measured index attains stable values. In this work, the measured index is the number of unexpected failures of the assets. The stabilisation of the value of an index is measured by the standard error [38] and is given by

$$\xi = \frac{\sigma}{\sqrt{N_{mc}}}, \tag{3.16}$$

where σ is the standard deviation of the index and N_{mc} is the number of MCS's.

Convergence has occurred when the standard error drops below a predetermined value ϵ. Thus, the stopping criterion for the MCS is met when

$$\xi \leq \epsilon. \tag{3.17}$$

3.3 Input Data

3.3.1 Digital twin model data

The input data for the reliability and maintenance modelling comes from a simulation run of the digital twin. In this simulation run, a scenario was taken with four combined heat and power (CHP) units, five heat pumps and three gas heaters. The output data of this simulation run is stored in comma-separated value (CSV) files. Two data files are relevant to this research. The first relevant CSV file contains information on the energy flows between the assets, residual heat from the data centre to the heat pumps and the heat demand. The second relevant CSV file contains the energy balance of the assets, the gas connection to the CHP unit and gas heater. All data points are the average hourly production values of the assets.

3.3.2 Asset production data

The CHP unit produces at its maximum capacity 700 kW net electrical power. The thermal energy production that can be recovered and used productively is 942 kW. The gas heater has a maximum capacity of 6,000 kW of thermal energy at 80°C. It then consumes 711 Sm^3 of natural gas[1] at an efficiency of 96%. There are two types of heat pumps used in the heat plant. One type is used for the residual heat provided by QTS, and the other type is used for Bytesnet. At maximum capacity, the heat pump of QTS requires 461 kW electrical power and has a coefficient

[1] Where Sm^3 is a standard cubic meter of natural gas, at a standard pressure of 101.325 kPa and 288.15 K. https://www.iso.org/standard/65049.html

of performance of 3.4[2]. The heat pump used for the residual heat from Bytesnet requires 459 kW electrical power at maximum capacity and has a coefficient of performance of 3.7[3]. The performances of the assets are sourced from their datasheet, which is available at KWR.

3.3.3 WarmteStad Maintenance Information

Maintenance information of the physical heat assets is required to parameterise the Weibull distribution that is used for the maintenance and reliability modelling. After consultation with WarmteStad, they shared information on the maintenance prescription of the assets provided by their respective original equipment manufacturer (OEM). The maintenance prescription and associated costs are captured in a service level agreement (SLA) with the OEM of the CHP unit, heat pump and gas heater.

WarmteStad stated that the costs of maintenance from the SLA are confidential. Therefore, the costs are omitted from the document. Furthermore, a shortened version of the maintenance prescription is provided to avoid sharing any sensitive contractual information. The impact of this decision on the results of the research is negligible.

Combined heat and power units

The maintenance actions performed on the CHP unit have a time between maintenance of 1,500 hours or a multiple thereof. These are the more frequent maintenance actions where small maintenance such as oil change, lubrication, valve adjustment, or cleansing of components is performed.

For a longer operating period of the CHP unit, performing only the regular maintenance actions is not sufficient due to grime deposition and wear on the components. Therefore, more thorough maintenance is required where components are overhauled c.q. replaced. Overhaul of the equipment ensures a high reliability and a longer lifespan of the assets. During maintenance, the downtime can take around four to seven days depending on the maintenance activities.

Heat pumps

The heat pumps mainly require maintenance to the ammonia detection system and to the piston compressors. The ammonia detection system is maintained twice per year and the piston compressors roughly every 3,000 hours. A heat pump contains two piston compressors which are overhauled after a few regular maintenance cycles to extend their lifetime. The maintenance actions take two to four days depending on the number of maintenance actions performed.

Gas heater

The gas heater is a relatively simple asset with respect to the CHP units and heat pumps, as it does not contain many moving parts. For this reason, the gas heater has low maintenance requirements and is maintained once per year. The maintenance actions take one day.

[2] The COP is provided by the manufacturer with ΔT is 51°C. True COP varies under operating condition.
[3] The COP is provided by the manufacturer with ΔT is 53 °C. True COP varies under operating condition.

Chapter 4

Operationalisation

This chapter provides the parameterisation of the model input parameters and describes the operationalisation of the maintenance optimisation model and unit commitment algorithm. This includes formulas and algorithms that define the developed models.

4.1 Basic Assumptions

This section provides an overview of the assumptions made throughout the research for the maintenance and reliability modelling. A description of how these assumptions are derived is provided the sections below.

1. The shape parameter of the Weibull distribution that describes the lifetime of all three assets is $\beta = 2.5$. The scale parameter is $\eta = [40, 30, 100]$ weeks for the CHP unit, heat pump and gas heater, respectively.
2. The improvement factor $\alpha = 0$.
3. The downtime attributed to the maintenance actions is fixed for all assets.
4. All CHP units have the same lifetime distribution.
5. All heat pumps have the same lifetime distribution.
6. All gas heaters have the same lifetime distribution.
7. The time between maintenance for the heat pumps is limited to one interval.
8. Production loss of the heat pumps and CHP units is compensated by the gas heaters.

4.2 Parameterisation of the Model Input Parameters

The maintenance and reliability model has several input parameters that must be predetermined for the simulation. This section is used to assign values to those parameters based on the data presented in Section 3.3, parameters found in the literature and information provided by WarmteStad's maintenance engineer.

4.2.1 Parameterisation of the Weibull distribution

In the previous chapter, it is assumed that the arrival rate of unexpected failures of the assets follows a Weibull distribution for its characteristic property, which is the ability to model different phases of the lifetime of a component for a given shape parameter. Failure data of the assets in the physical heat plant is required to parameterise the Weibull distribution that is used for the modelling in this research. Consultation with WarmteStad contributed to useful information on the prescribed maintenance for the assets. However, unfortunately, WarmteStad was not able to provide extensive information on the failure rates of the assets, as the heat plant is operating in its initial phase. As a result, especially the heat pumps (HP) are not yet operating under optimal conditions and experience more downtime than expected. The maintenance engineer said that the combined heat and power (CHP) units had been operating stable in the past half year and experienced no failures. The two heat pumps currently in operation experienced six failures in total. When these failures occur, the safety protection system lets the heat pump run out slowly. Then, the ammonia refrigerant must evaporate for one day, and the heat pump can be put into operation again if no additional maintenance is required. The gas heaters (GH) were operating at a reduced capacity and have not experienced any failures.

Based on this information provided by WarmteStad, it was not possible to fit a Weibull distribution to the failure data. The amount of failure data was inadequate and most of the information was provided on the prescribed maintenance by the manufacturer. A literature search for the failure distribution of similar assets results in the parameters that are provided in Table 4.1.

In these papers, the failure rate is given using an exponential distribution, which is a special case of the Weibull distribution, where $\beta = 1$ [29]. The values found in the literature are on the low side with respect to the prescribed maintenance intervals provided by WarmteStad and the information provided by the maintenance engineer on the reliability of the assets. Therefore, based on a combination of the literature, the maintenance data and the information provided by the maintenance engineer, the Weibull parameters are estimated to be $\beta = 2.5$ for the shape parameter of all three assets, and $\eta = [40, 30, 100]$ weeks for the scale parameter of the CHP, HP and GH respectively.

Table 4.1: Failure rate parameters of assets similar to the assets in WarmteStad's heat plant found in the literature.

Asset	Failure rate	Reference
Gas heater	5.00E-3 $\frac{1}{\text{hour}}$	[14]
CHP	1.65E-3 $\frac{1}{\text{hour}}$	[14]
CHP	5.16E-4 $\frac{1}{\text{hour}}$	[14]
Heat pump	0.5 $\frac{1}{\text{year}}$	[39]
Heat pump	0.0013 $\frac{1}{\text{hour}}$	[40]
Gas engine	0.008037 $\frac{1}{\text{day}}$	[41]
Auxiliary boiler	0.000903 $\frac{1}{\text{day}}$	[41]

4.2.2 Parameterisation for the maintenance optimisation model

The maintenance optimisation is performed for a time span of one year, which is equal to the simulation time of the digital twin. Hence, the planning horizon is 52 weeks. Due to the complexity of the optimisation problem, it is not possible to optimise the maintenance schedule on an hourly or even daily basis. It would simply take too long to find a solution. Therefore, the number of intervals is set to 52, resulting in a resolution of the maintenance optimisation algorithm of one week.

From the prescribed maintenance information of the assets provided by WarmteStad, it is observed that regular maintenance includes the replacement of components, which implies an improvement of the state of that component to as good as new. However, the replacement of some components does not restore the entire asset to "as-good-as-new". This is observed by the fact that overhaul of the assets is performed after tens of thousands of hours for the CHP unit.

These long intervals imply that overhaul is only performed after a few years of operation. This renders it difficult to assess the value of the improvement factor that brings the state of the asset to somewhere in between "as-bad-as-old" and "as-good-as-new". The fact that the assets must be overhauled after some years of operation means that the assets are durable, and maintenance improves the asset significantly. Therefore, the improvement factor is set to zero, which means the asset state is restored to as good as new after maintenance. The same holds for the heat pump and gas heater; thus, the improvement factor for those assets is set to zero as well. To avoid confusion, an improvement factor of zero reduces the age of the asset to zero; hence, the asset is improved to "as-good-as-new".

Based on the maintenance information provided by WarmteStad, the maintenance intervals are set to 1,500 hours for the CHP units, 3,000 for the heat pumps and 5,000 hours for the gas heaters. The maintenance costs provided by WarmteStad are confidential; hence, it was not permitted to use that information for this research. Therefore, fictitious values are taken for the maintenance costs, though still in a realistic ratio with respect to the true maintenance costs. The maintenance costs are set to 15,000 EUR for the CHP units, 10,000 EUR for the heat pumps, and 2,000 EUR for the gas heater.

The maintenance costs are captured in a service-level agreement. As a result, a fixed price is paid for the maintenance actions applied to the assets. Therefore, the cost of unexpected failures is equal to the cost of maintenance, which is confirmed by the maintenance engineer.

Table 4.2 presents the results from a literature search for the operation and maintenance (O&M) costs for the assets in WarmteStad's heat plant. The table contains the non-fuel O&M costs of the assets and expected cost per asset per year based on the O&M costs. This information is used as a reference value to compare to the maintenance costs from the experiments, as a method of model validation.

Table 4.2: Expected O&M costs of for WarmteStad's assets based on values found in the literature. The average thermal capacity of the heat pumps is estimated to be 1,500 kW, at a COP of 3.3.

	CHP	HP	GH
Capital expenditures (CAPEX)	n/a	400 EUR/kW$_{th}$	n/a
Non-fuel O&M costs	0.008-0.024	5	3.5
Unit	EUR/kWh$_e$	% CAPEX/year	EUR/kW$_{th}$/year
Source	[42]	[43], [44]	[45]
Expected costs per asset per year	20,462-61,385	30,000	21,000

4.2.3 Parameterisation for the unit commitment algorithm

Downtime of the assets attributed to maintenance is considered in the unit commitment algorithm. Downtime is usually modelled using a log-normal distribution [33],[46]. However, WarmteStad has no sufficient data on the duration of downtime of the assets. Therefore, the downtime is considered constant with a value of 100 hours for the CHP unit, 75 hours for the heat pumps and 25 hours for the gas heaters. These values are based on the information in Section 3.3 and were provided by WarmteStad's maintenance engineer.

A common standard error of $\epsilon = 0.05$ is used as a stopping criterion to guarantee small variance in the results of the Monte Carlo simulation (MCS) [47]. Convergence of the MCS must be assessed in hindsight, as the standard deviation of the number of failures is assessed based on the results of the simulation. Therefore, the initial number of MCS's is set to 10,000. This means that one year of simulation time is iterated 10,000 times. This many iterations are necessary to increase the variation in the number of unexpected failures and the time at which they occur.

4.3 Maintenance Optimisation Algorithm

The optimisation problem in Subsection 3.1.2 must be solved to find the optimal maintenance schedule. In this research, the optimisation problem was solved using Matlab and the toolbox for optimisation YALMIP[48]. An additional solver was used, which is more efficient than the built-in branch and bound algorithm. Furthermore, this solver can handle nonlinear nonconvex optimisation problems by using a piecewise affine (i.e. piecewise linear) approximation of the function. This third-party-solver is Gurobi[1] and was provided under an academic licence.

The reliability constraint in (3.10) is nonlinear and nonconvex, which cannot be handled by Gurobi directly. Therefore, the nonlinear function must be approximated using a piecewise linear function. First, the constraint is rewritten as

$$R_i(X'_{i,j}) = \prod_{j=1}^{T} \mathrm{e}^{-\left(\frac{X'_{i,j}}{\eta}\right)^{\beta}} \iff \ln(R_i(X_{i,j})) = \sum_{j=1}^{T} -\left(\frac{X'_{i,j}}{\eta}\right)^{\beta}. \tag{4.1}$$

This modification is necessary for the solver as Gurobi cannot handle the product of a function as a constraint. The left-hand side in (4.1) is mathematically equal to the right-hand side of the equation, though, with the product becoming a summation, Gurobi can now handle this type of constraint.

Then, a piecewise linear approximation of the nonlinear function is defined as a special ordered set of type 2 (SOS2) [49]. The right-hand side of (4.1) is computed on a sufficiently fine grid, and the data points are used to define a SOS2 constraint. Gurobi can handle SOS2 constraints and interpolates linearly between the data points in the set. The same procedure is taken to make a piecewise linear approximation of the nonlinear part of the objective function in (3.8).

4.4 Asset Utilisation

For the maintenance and reliability model, it is required to have information on how many assets are in use at a given time. This information is not directly available from the digital twin simulation data. However, based on the simulation data, it is possible to determine the number of assets that were used in the simulation.

[1]https://www.gurobi.com/

In order to compute the number of CHP units required, the total electrical energy produced in one hour is computed first. The electric energy production is not part of the output data of the digital twin, thus must be computed using the energy flows and balances. The electrical energy generated by the CHP unit is computed as

$$E_{\text{CHP,e}} = E_{\text{gcc,g}} + E_{\text{hp\_to\_chp,th}} - (-E_{\text{EB,CHP,th}}) - E_{\text{chp\_to\_gh,th}}, \tag{4.2}$$

where $E_{\text{CHP,e}}$ is the electrical energy generated by the CHP unit, and $E_{\text{EB,CHP,th}}$ is the imbalance of the energy balance over the CHP assets. The fact that it is negative can be interpreted as energy that is lost to the surroundings, for example, by the emission of hot flue gasses. $E_{\text{gcg,g}}$ denotes the energy equivalent of the natural gas provided by the gas connection to the CHP units. $E_{\text{hp\_to\_chp,th}}$ and $E_{\text{chp\_to\_gh,th}}$ denote the thermal energy flow from the heat pumps to the CHP units and the thermal energy flow from the CHP units to the gas heaters, respectively. All energy terms are expressed in kWh.

In the digital twin, the heat pumps and the CHP units are coupled electrically. Thus, the CHP units produce as much electricity as is required for the heat pumps. Therefore, now that the electricity generation of the CHP units is available by (4.2), it is possible to compute the number of heat pumps and CHP units active at all time steps in the digital twin simulation data.

Combined heat and power units

The CHP units generate at maximum capacity 700 kW of electricity see Subsection 3.3.2. The CHP units have the highest efficiency at maximum capacity; therefore, the hourly number of active CHP units in the digital twin simulation data is computed by

$$\mathrm{N}_{\text{CHP,required}} = \left\lceil \frac{E_{\text{CHP,e}}}{700} \right\rceil, \tag{4.3}$$

where $\lceil x \rceil$ denotes the ceiling function, that maps the variable x to the nearest integer greater than, or equal to x.

Heat pumps

There are two types of heat pumps that have nearly the same electricity consumption, 459 kW and 461 kW see Subsection 3.3.2. Therefore, the heat pumps are averaged to use 460 kW of electrical power each. The number of heat pumps active at each hour in the digital twin simulation data is computed by

$$\mathrm{N}_{\text{HP,required}} = \left\lceil \frac{E_{\text{CHP,e}}}{460} \right\rceil. \tag{4.4}$$

Gas heaters

The hourly number of active gas heaters is based on the gas consumption of the gas heaters in the digital twin simulation data. Its computation is given by

$$\mathrm{N}_{\mathrm{GH,required}} = \left\lceil \frac{E_{\mathrm{gcg,g}}}{711 \cdot 8.792} \right\rceil, \tag{4.5}$$

where the value 711 is the natural gas consumption of one gas heater at maximum capacity in Sm^3, and the value 8.792 is the lower heating value of one Sm^3 natural gas in kWh, see Subsection 3.3.2.

4.5 Unit Commitment Algorithm

Based on the digital twin simulation data, the total number of assets required to fulfil the heat demand was determined for each hour of simulated time. For the reliability modelling of the heat plant, it is necessary to determine the individual utilisation of the assets, i.e. which assets are in operation, which assets are on standby, which assets are unavailable due to a breakdown or scheduled maintenance, and what is the operating time of the assets. This type of information cannot be derived from the digital twin simulation data. Therefore, it is necessary to develop a unit commitment algorithm for the twelve assets based on the digital twin simulation data. The unit commitment algorithm is an algorithm that commits the number of required CHP units, heat pumps and gas heaters based on the digital twin simulation data. The algorithm is based on a for loop that cycles through every hour of simulation time. This loop contains different components such as activation and deactivation of the assets, execution of maintenance for the assets, modelling unexpected failures of the assets and compensating for lost production when a heat pump or CHP unit fails, and no other unit is available.

4.5.1 Activation and deactivation of assets

The activation and deactivation of assets is classified in three categories: CHP units, heat pumps and gas heaters. The process of activating and deactivating assets is explained for assets in general and can easily be applied to any category of assets. The pseudo-code of the algorithm for activating heat pumps is given by

Line 1: The algorithm reads the number of required heat pumps in time step j. Line 2: If the number of required heat pumps exceeds the number of active heat pumps in time step $j-1$, a heat pump must be activated. There is no algorithm or logic behind determining the best heat pump to commit; therefore, the first available heat pump will be activated in numerical order. This is actually also the sequence in which the assets in the physical heat plant are activated, according to its control system description. Line 4: this loop ensures that the assets are activated in numerical order. Line 5 ensures that only assets are activated when necessary. Lines 6-8 check if the asset is available, commit the asset and reduce the number of assets to be activated with one. If the number of required assets is less than the number of active assets,

Algorithm 1 Pseudo-code for activation the assets

```
read number of required assets
if number of active assets < number of required assets then
    activate = number of required assets - number of active assets
    for i = 1:total number of assets do
        if activate > 0 then
            if asset i is available then
                activate asset i
                activate += -1
            end if
        end if
    end for
end if
```

an asset must be deactivated. This is done in reversed numerical order, meaning that the last active asset in the sequence is deactivated first.

4.5.2 Compensating production loss

Consider the case where all heat pumps are in operation. When one of the heat pumps fails, the production loss cannot be compensated by another heat pump. Therefore, this production loss must be compensated by another asset. The CHP unit and heat pump are coupled electrically; however, if one or the other fails, the assets can be connected to the electricity grid. Therefore, in theory, a CHP unit can compensate the thermal energy production loss of a heat pump and vice versa. However, in the model, it is assumed that only the gas heater is allowed to compensate for this production loss. This assumption is made since the supply side of low-temperature residual heat of the heat pump cannot be increased on demand. And for the CHP unit, its prime purpose is to generate electricity for the heat pumps, not for the electricity grid.

Now consider the case where all gas heaters are active and running at full capacity. If production is lost due to an unexpected failure, another gas heater cannot be activated. As a result, the heat plant fails to match the heat production demand which is called *substandard heat production*. It is good to note that this does not necessarily imply that no heat is delivered to the customer, as there is still energy stored in the thermal buffer.

The algorithm for compensating heat production loss is given in Algorithm 2. The first step is to calculate the amount of production loss (line 1), then compute the residual capacity of the gas heaters if those are activated (line 2). If this residual capacity is sufficient to compensate for the production loss of a heat pump or CHP unit, the production capacity of the active gas heaters is increased. If the residual capacity is not sufficient to compensate for the loss of production (line 3), another gas heater is activated (line 4). If no other gas heater is available (line 5), the amount of production loss is stored as substandard heat production (line 6).

Algorithm 2 Pseudo-code for compensating production loss

```
do compute production loss
do compute residual capacity gas heater
if residual capacity gas heater < production loss then
    do activate additional gas heater
    if gas heater is unavailable then
        store heat production loss
    end if
end if
```

4.5.3 Performing Maintenance

The maintenance schedule is optimised using the maintenance optimisation algorithm. This algorithm generates an optimal maintenance schedule for the given input parameters over a given time horizon. Based on this schedule, the optimal effective age for maintenance is determined and used in the unit commitment algorithm. Here, maintenance of the assets is performed at the optimal effective age. If there is an asset already in maintenance, the maintenance is postponed. Otherwise, the asset is taken out of operation into maintenance. A downtime for the respective asset is attributed to the maintenance action, and the asset cannot be committed during this time. When the downtime has passed, the asset is set on standby and can be committed when it is required.

Algorithm 3 Pseudo-code for performing maintenance

```
if maintenance is due then
    if any asset is already under maintenance then
        do not perform maintenance
    else
        do perform maintenance
        do assign downtime to maintenance action
        do update commitment matrix
    end if
end if
```

4.5.4 Postponing Maintenance

There are two situations where performing maintenance can negatively impact the ability of the heat plant to fulfil the heat demand. The first situation is when an asset is already under maintenance, it is sensible to wait until completion of the maintenance, as maintaining multiple assets at the same time can greatly impact the ability of the heat plant to fulfil the heat demand.

The second situation is when there is a period of high heat demand upcoming. In this situation, it is better to postpone the maintenance of the assets until the period of high heat demand has passed. It is expected that the benefits of postponing maintenance by a few days outweigh the negative effects. The negative effects are an increased risk of failure of the asset during this period while performing maintenance certainly results in unavailability of the asset.

When one of both situations occurs, the unit commitment algorithm waits to perform maintenance until the conditions for maintenance are satisfied. That is either when no asset is in maintenance or when a period of lower heat demand is expected. The following subsection explains how a period of lower heat demand is detected based on past heat demand data.

4.5.5 Detecting period of lower heat demand

Detection of a period of lower heat demand is performed using a *rule-based* algorithm. This algorithm is used to impose a maintenance constraint on the assets that restricts performing maintenance during a period of high heat demand. A moving average of the heat demand is computed for the past three days for smoothing the heat demand. The pseudo-code for the rule-based algorithm is given in Algorithm 4.

Consider the case where maintenance of an asset is due; then, when the 72-hour-average heat demand (72AHD) is below 10 MW, maintenance is always allowed. If this is not the case, and the heat demand is above 20 MW, maintenance is never allowed. If the 72AHD is between 10 MW and 20 MW, it is checked if the heat demand trend of the past 72 hours is *decreasing*. If these two conditions are met, maintenance on the asset is allowed.

Algorithm 4 Pseudo-code for detecting a period of lower heat demand

$\text{72AHD}_i = \frac{1}{72}\sum_{j=i-72}^{i} \text{hd}(t_j)$ ▷ compute the 72-hour-average heat demand
if maintenance is due **then**
    if $\text{72AHD}_i <= 10$ MW **then**
        do perform maintenance
    else if $\text{72AHD}_i <=$ 20MW **and** trend $[\text{72AHD}_{i-72} \ldots \text{72AHD}_i]$ is decreasing **then**
        do perform maintenance
    else if $\text{72AHD}_i > 20$ MW **then**
        do not perform maintenance
    end if
end if

4.5.6 Unexpected failure of assets

At this point, the algorithm can commit assets and schedule maintenance when it is due; however, at any given time, there is a probability that an asset will fail randomly. To capture the effect of these random failures in the model, stochastic failure behaviour is added to the assets. The propensity of an asset failure occurring is based on the probability density function of the lifetime of that asset. Therefore, it is possible to model random failures by sampling from this probability density function.

The inverse transform method is used to draw a sample from the Weibull distribution, see Subsection 3.2.1 for further details. For the initial model run, it is assumed that the effective age of the assets is zero. Therefore, a uniformly distributed sample is drawn in the interval between $[0, 1]$. This probability corresponds to a specific lifetime at which the asset is expected to fail and is computed using (3.1). If the effective age at which the asset is expected to fail is below the age at which maintenance is applied according to the optimal maintenance schedule, the asset fails unexpectedly. If this age is above the optimal maintenance age, no unexpected failure occurs, and the asset is maintained when maintenance is due.

When maintenance has been applied to an asset, a new sample is drawn from its lifetime probability density function to simulate its new expected failure point. However, the effective age of the asset can be nonzero now, depending on the improvement factor. Therefore, the lower bound of the uniform distribution is shifted to the probability that corresponds to the effective age of the asset.

Chapter 5

Results and Analysis

This chapter starts with a description of the scenario's for the simulation experiments conducted in this research. Then, the results of the experiments are provided including an analysis. The interpretation of the results is provided in the discussion.

5.1 Scenario's for Experiments

This section describes the three scenarios for the simulation experiments performed in this research. First a base scenario using WarmteStad's maintenance schedule is performed as a model benchmark, then a simulation is performed using the maintenance optimisation algorithm, and lastly, a simulation is run where a constraint is imposed that restricts maintenance during high heat demand.

Naturally, there are endless variations to the experiments that can be conducted using the maintenance and reliability model developed during this research. However, the objective of the experiments is to show the model's key features for a realistic scenario of WarmteStad's heat plant.

5.1.1 Experiment 1: Base Scenario

In this experiment, a base scenario is simulated using the prescribed maintenance intervals for the assets provided by their manufacturers. This is a base scenario, as WarmteStad is currently using this information for its maintenance planning. Therefore, no maintenance optimisation is performed in this experiment, and the results are used as a benchmark for the model. The maintenance interval for the combined heat and power (CHP) units is 1,500 hours, the heat pump 3,000 hours, and the gas heater 5,000 hours, based on the information provided by WarmteStad. Furthermore, the constraint that maintenance is not permitted during a period of high heat demand is *inactive*. This experiment is referred to as the *base scenario* experiment.

5.1.2 Experiment 2: Maintenance Optimisation

This experiment is conducted to assess the performance of the maintenance optimisation algorithm. As in the base scenario, the constraint that maintenance is not permitted during a period of high heat demand is *inactive*. This experiment is referred to as the *maintenance optimisation* experiment.

The input parameters used for this experiment are given in Table 5.1. This table provides an overview of the parameters defined in Section 4.2. The lower bound for the reliability of the assets is set to 80% for the CHP, 80% for the HP and 80% for the GH. These values are chosen arbitrarily for demonstration purposes and can take any value between 0% and 100%.

Table 5.1: Model input parameters used for the maintenance optimisation experiment. The parameters between brackets correspond to the CHP, HP and GH respectively.

Input parameters					
$\eta =$	[40, 30, 100]	weeks	$\beta =$	2.5	
F =	[15000, 10000, 2000]	EUR	$\alpha =$	0	
M =	[15000, 10000, 2000]	EUR	T =	52	weeks
R =	[80, 80, 80]	%	MCS =	10000	iterations

5.1.3 Experiment 3: Maintenance Constraint

This experiment explores the scenario where maintenance optimisation is performed and in addition a maintenance constraint is imposed that restricts maintenance during periods of high heat demand. The maintenance constraint is imposed by a *rule-based* algorithm defined in Subsection 4.5.5. The model input parameters are equal to the parameters used in the maintenance optimisation experiment and are presented in Table 5.1. The experiment is performed to assess the effect of this constraint on the substandard heat supply, and is referred to as the *maintenance constraint* experiment.

5.2 Base Scenario vs Maintenance Optimisation

The first step of analysing the simulation result is to assess if the Monte Carlo simulation (MCS) has converged. The stopping criterion for the MCS is based on the standard error, which is computed using (3.16). The standard errors of experiments 1 and 2 are provided in Table 5.2. This table shows that the MCS has converged for all three assets in both experiments.

Table 5.2: This table presents the values of the standard error ξ for experiments 1 and 2 and states if the MCS has converged.

	Base scenario			Maintenance optimisation		
Asset	CHP	HP	GH	CHP	HP	GH
ξ	0.0040	0.0136	0.0022	0.0040	0.0053	0.0017
$\xi \leq \epsilon$	true	true	true	true	true	true

Histograms of the number of failures of the assets are presented in Figure 5.1. It is observed that the failure rate of the CHP units in the maintenance optimisation experiment has increased slightly with respect to the base scenario. This implies that the number of maintenance actions on the CHP units is slightly lower than in the base scenario. The failure rate of the heat pumps has decreased significantly in the maintenance optimisation experiment with respect to the base scenario. This is due to a significant increase in maintenance actions applied to the heat pumps. The gas heaters are highly reliable overall due to the low failures in both experiments. Though, the gas heater failed slightly less often in the base scenario.

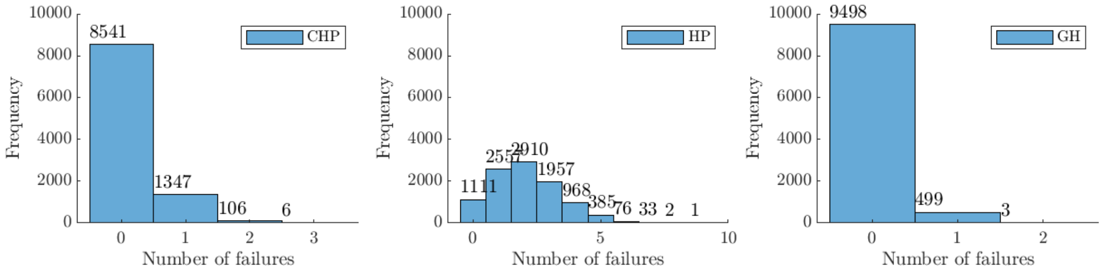

(a)

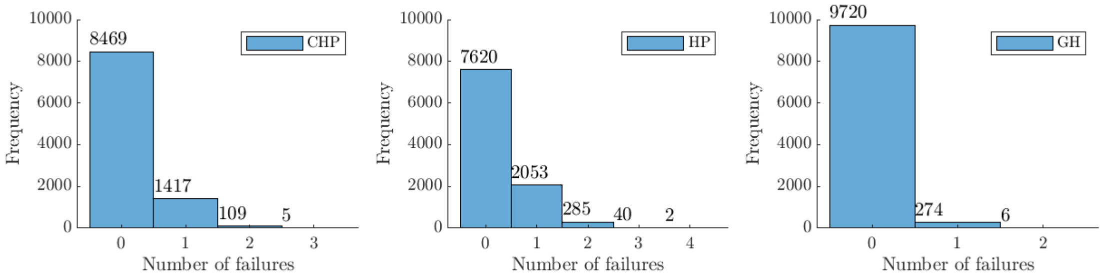

(b)

Figure 5.1: Histograms of the number of unexpected CHP, HP and GH failures during MCS. The histograms of the base scenario are presented in Figure 5.1a, and the histograms of the maintenance optimisation experiment are presented in Figure 5.1b.

An important KPI to assess the reliability of the entire heat plant is the number of substandard heat production hours. This occurs when the heat plant cannot fulfil the required heat production demand. Histograms of the number of substandard heat production hours during both experiments are given in Figure 5.2. This figure shows that in the base scenario, more substandard heat production hours occur than in the maintenance optimisation experiment. Furthermore, the duration of the substandard heat production is shorter. Further analysis of the lost heat production is necessary to determine if substandard heat production resulted in a heat supply failure in any of the simulations.

The maximum heat production loss in one consecutive period is computed to quantify a potential heat supply failure. The histograms of the maximum heat production loss during MCS are presented in Figure 5.3. These histograms show that heat production loss occurs significantly more often in the base scenario than in the maintenance optimisation experiment. A significant proportion of the heat production loss is 0-20 MWh in the base scenario, while in the maintenance optimisation experiment, there is only a small peak in 0-5 MWh of production loss.

To assess if the heat production loss resulted in a *heat supply failure*, the heat production loss is subtracted from the energy content of the thermal buffer at the beginning of a heat production loss period. If the heat production loss is greater than the energy content of the thermal buffer, the production loss results in *substandard heat supply*, where the heat demand of the DHN cannot be fulfilled. This is true in 83.1% and 29.9% of the cases in the base scenario and the maintenance optimisation experiment, respectively. The histograms of substandard heat supply are presented in Figure 5.4. The histograms show that in the maintenance optimisation

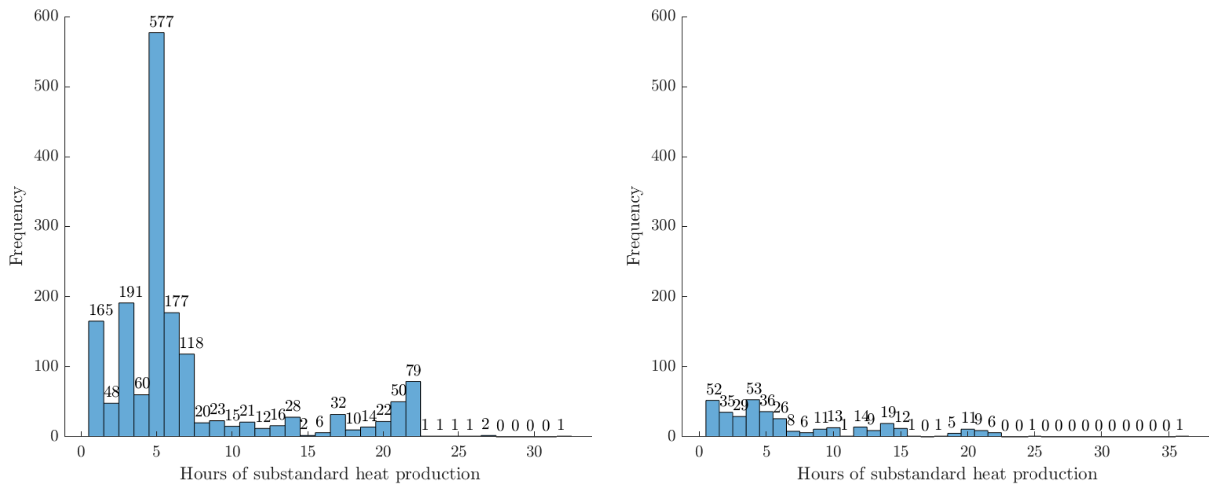

(a) (b)

Figure 5.2: Histograms of the number of substandard heat production hours. The histogram of the base scenario is presented in Figure 5.2a, the histogram of the maintenance optimisation experiment is presented in Figure 5.2b.

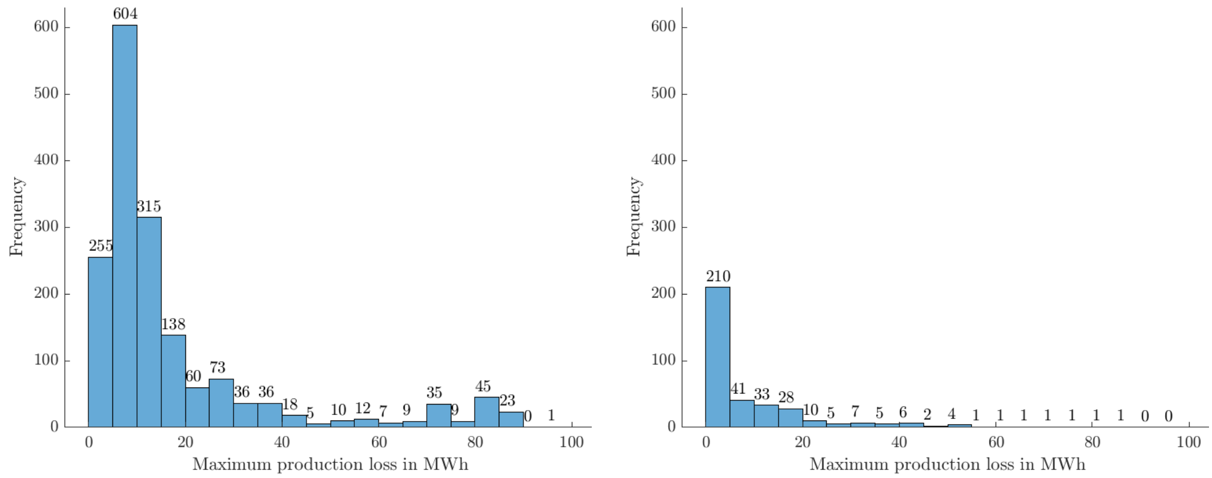

(a) (b)

Figure 5.3: Histograms of the maximum heat production loss in one consecutive period in MWh. The histogram of the base scenario is presented in Figure 5.3a, and the histogram of the maintenance optimisation experiment is presented in Figure 5.3b.

experiment, significantly fewer heat supply failures occur than in the base scenario. In total, there is a reduction of 85.9% in the occurrences of heat supply failure.

Furthermore, the quantity of substandard heat supply during a heat supply failure is higher in the base scenario than in the maintenance optimisation experiment. In 15.4% and 84.4% of the cases in the base scenario than in the maintenance optimisation experiment, respectively, a heat supply failure was due to unexpected failures of assets. In the remainder of the cases, the heat supply failure was due to scheduled maintenance.

Lastly, the preventive and corrective maintenance costs are presented in Figure 5.5. From these bar plots, it is observed that the costs for maintenance and unexpected failures of the CHP

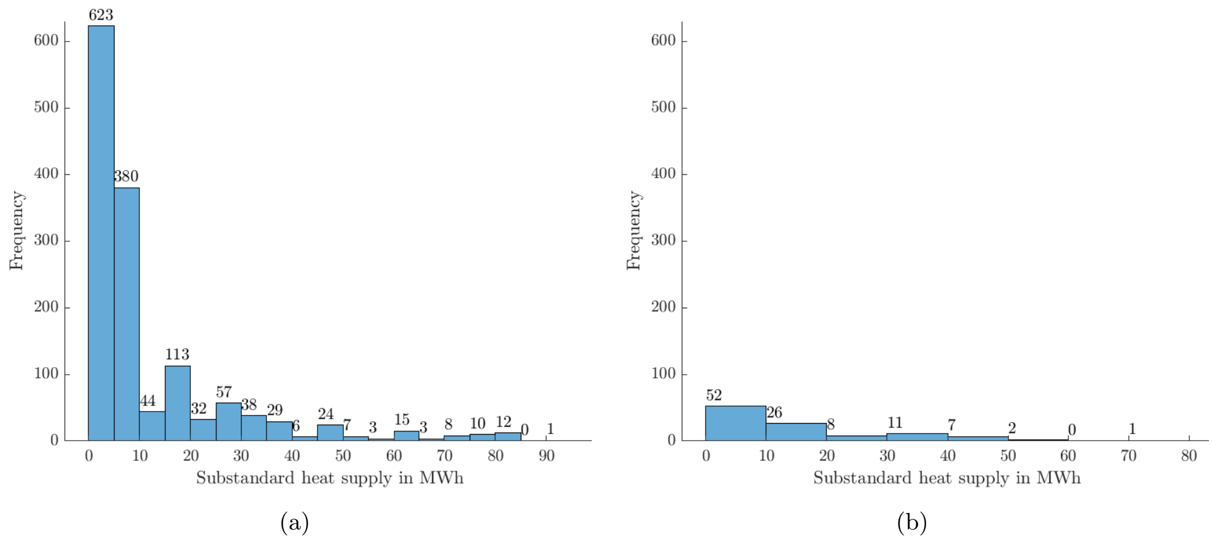

Figure 5.4: Histograms of the substandard heat supply in MWh. The histogram of the base scenario is presented in Figure 5.4a, and the histogram of the maintenance optimisation experiment is presented in Figure 5.4b.

units are nearly equal, as the number of maintenance actions for the CHP units was nearly equal in both experiments. The maintenance costs for the heat pumps have significantly increased in the maintenance optimisation experiment with respect to the base scenario. As a result, the cost of unexpected failures has decreased significantly. Overall, the combined costs for preventive and corrective maintenance are significantly higher in the maintenance optimisation experiment than the base scenario. The maintenance costs for the gas heater is slightly, though negligible higher in the maintenance optimisation experiment than in the base scenario.

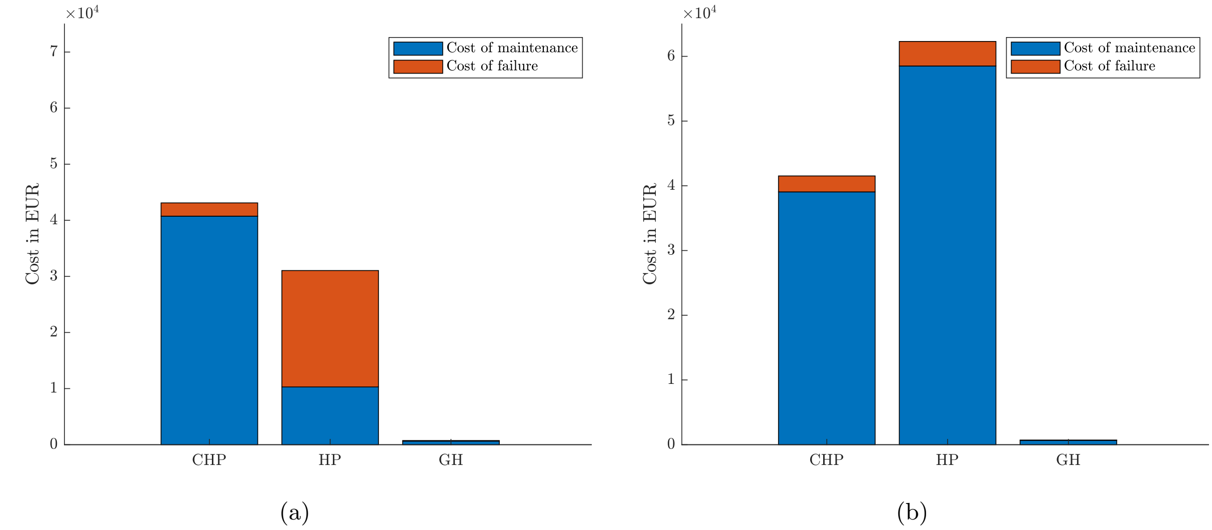

Figure 5.5: Stacked bar plots of the average costs of maintenance and unexpected failures of assets. The costs are the average costs per individual asset over 10,000 MCS runs. The bar plot of the base scenario is presented in Figure 5.5a, and the bar plot of the maintenance optimisation experiment is presented in Figure 5.5b.

5.3 Maintenance Optimisation vs Maintenance Constraint

In the maintenance constraint experiment, maintenance is not permitted during a period of high heat demand. The effect of this constraint on the periods where maintenance is permitted is plotted in Figure 5.6. This graph shows the three-day moving average of the heat demand of one year of simulation time. The moving average is used for smoothing the heat demand. No forecast of the heat demand is made, implying that only information on past heat demand is available during the simulation.

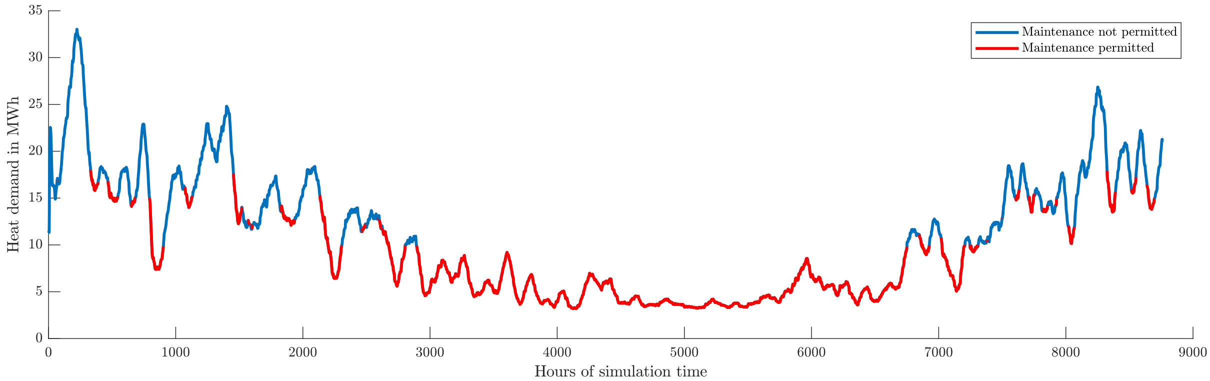

Figure 5.6: Three-day moving average of the heat demand of one year simulation time. The blue line indicates periods of high heat demand where maintenance is not permitted. The red line indicates the periods where maintenance is permitted.

From this figure, it is observed that during peaks of high heat demand, maintenance of the assets is not permitted. This constraint is expected to improve the reliability of the heat plant by reducing the number of maintenance actions during high heat demand; therefore, reducing the impact of scheduled maintenance actions on the heat production. To assess the effect of the maintenance constraint, analysis of the MCS proceeds.

Again, the first step is to assess if the MCS has converged. The standard error in the maintenance constraint experiment is presented in Table 5.3. From this table, it is observed that the MCS has converged for all three assets.

Table 5.3: This table presents the values of the standard error ξ of the maintenance constraint experiment, and states if the MCS has converged.

Asset	CHP	HP	GH
ξ	0.0041	0.0065	0.0017
$\xi \leq \epsilon$	true	true	true

Histograms of the number of failures during MCS are presented in Figure 5.7. This figure shows that slightly more failures have occurred for the CHP unit in the maintenance constraint experiment than the maintenance optimisation experiment. It is expected that this happened as maintenance is in some cases postponed in the maintenance constraint experiment due to the maintenance constraint. This leads to an increased time between maintenance in some cases, which increases the probability of failure of the assets.

The same argument holds for the increased number of failures of the heat pumps. The histogram shows that 1107 more failures occurred during MCS in the maintenance constraint experiment, which is significantly more than the additional CHP unit failures, and is explained by the higher failure rate of the heat pumps. If maintenance of a heat pump is postponed, the increase in the probability of failure is greater than for the CHP or GH.

For the gas heaters, only two more failures occurred during 10,000 MCS runs. This is explained by the fact that the maintenance requirements of the gas heater are already relatively low due to the longer maintenance intervals and relatively short operating time per year. Therefore, the effect of postponing maintenance for the gas heater during periods of high heat demand is negligible.

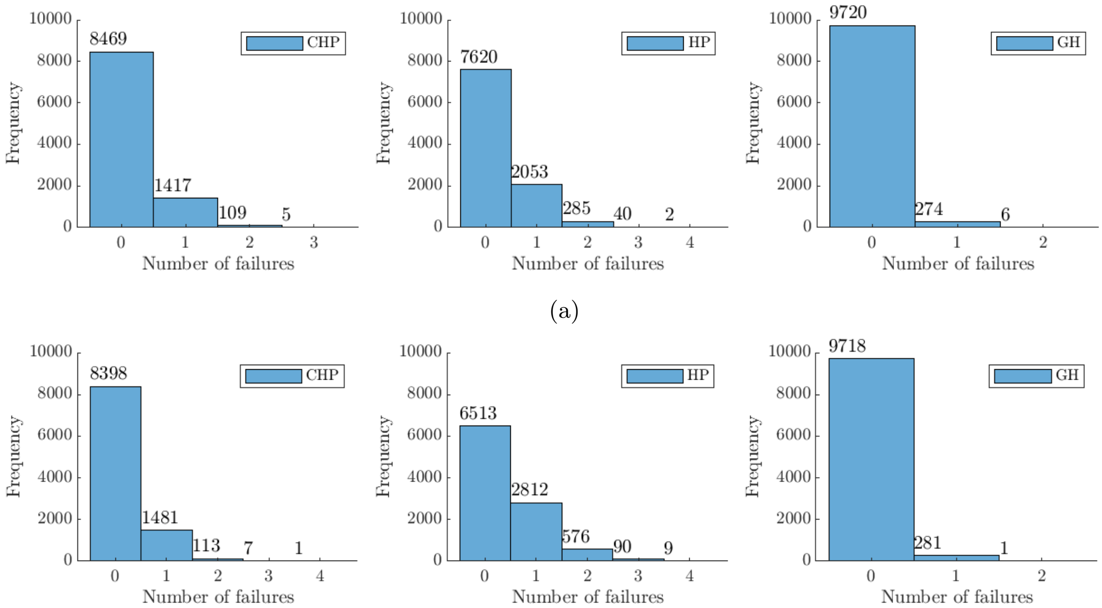

(a)

(b)

Figure 5.7: Histograms of the number of unexpected failures of the CHP, HP and GH during MCS. The histograms of the maintenance optimisation experiment are presented in Figure 5.7a, the histograms of the maintenance constraint experiment are presented in Figure 5.7b.

Histograms of the number of substandard heat production hours are presented in Figure 5.8. These histograms show that overall, the number of substandard heat production hours has further decreased in the maintenance constraint experiment.

Histograms of the heat production loss during experiments 2 and 3 are presented in Figure 5.9. This figure shows that in general, the quantity of substandard heat production measured in MWh has decreased in the maintenance constraint experiment with respect to the maintenance optimisation experiment. However, it is remarkable to observe that a substandard heat production of 80-90 MWh occurred four times more often in the maintenance constraint experiment than in the maintenance optimisation experiment. This is the largest quantity of substandard heat production, thus imposing a higher risk on WarmteStad to incur costs attributed to a heat supply failure.

Another MCS is ran with 20,000 iterations, to investigate if the peak at 80-90 MWh is due

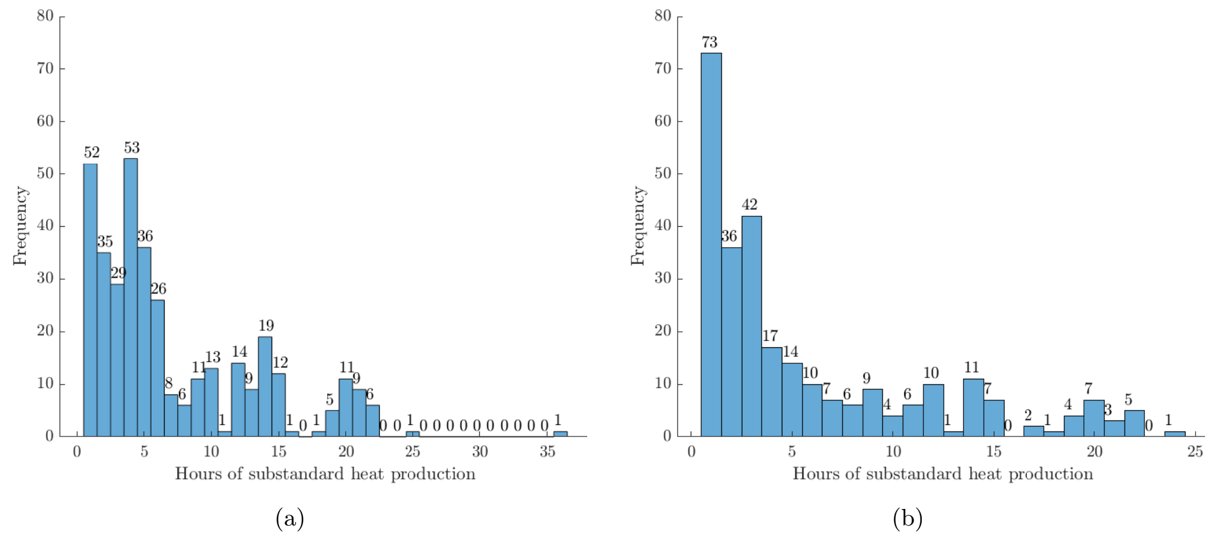

Figure 5.8: Histograms of the number of substandard heat production hours. The histogram of the maintenance optimisation experiment is presented in Figure 5.8a, the histogram of the maintenance constraint experiment is presented in Figure 5.8b.

to the maintenance constraint or due to "bad luck" with the random numbers generated in the simulation. From the results it is observed that 8 events of 80-90 MWh of production loss occur; hence, the peak in the original experiment is assumed to be attributed to "bad luck" in the random number generation during the experiment.

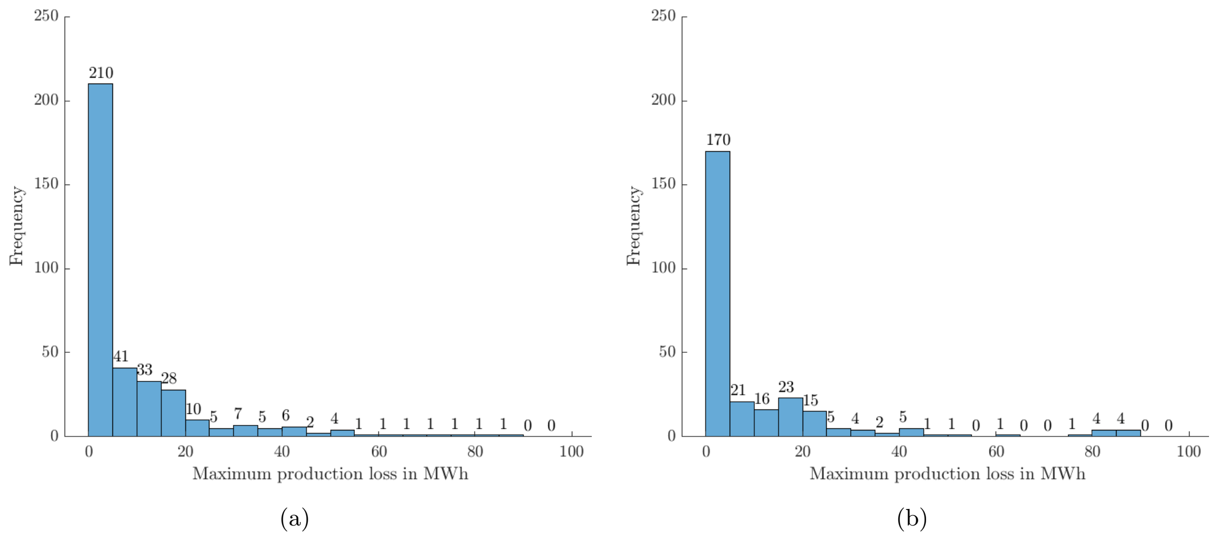

Figure 5.9: Histograms of the maximum heat production loss in one consecutive period in MWh. The histogram of the maintenance optimisation experiment is presented in Figure 5.9a, and the histogram of the maintenance constraint experiment is presented in Figure 5.9b.

In 31.1% of the cases in the maintenance constraint experiment, the heat production loss resulted in a heat supply failure where the energy content of the thermal buffer was inadequate to compensate for the production loss. This is a 1.2 percentage points (pp) increase with respect to the maintenance optimisation experiment, though in total, production loss occurs less often in the maintenance constraint experiment than in the maintenance optimisation experiment.

Furthermore, in 91.2% of the cases in the maintenance constraint experiment, heat production loss is a result of random failures of the assets rather than preventive maintenance actions. This is a 6.8 pp increase in the number of heat supply failures caused by random failures in the maintenance constraint experiment, which implies that the maintenance constraint reduced the number of heat supply failures caused by preventive maintenance actions.

The histograms of the heat supply failures are presented in Figure 5.10. A further reduction of 11.5% in the occurrence of substandard heat supply in the maintenance constraint experiment with respect to the maintenance optimisation experiment. Though again, a substandard heat supply of 70-90 MWh occurs four times in the maintenance constraint experiment and only once in the maintenance optimisation experiment. In the experiment with double the MC iterations, 70-80 MWh of substandard heat supply occurs 4 times as well, without higher values of substandard heat supply occurring, which confirms that the peak in the original experiment is likely due to “bad luck” in the random number generation.

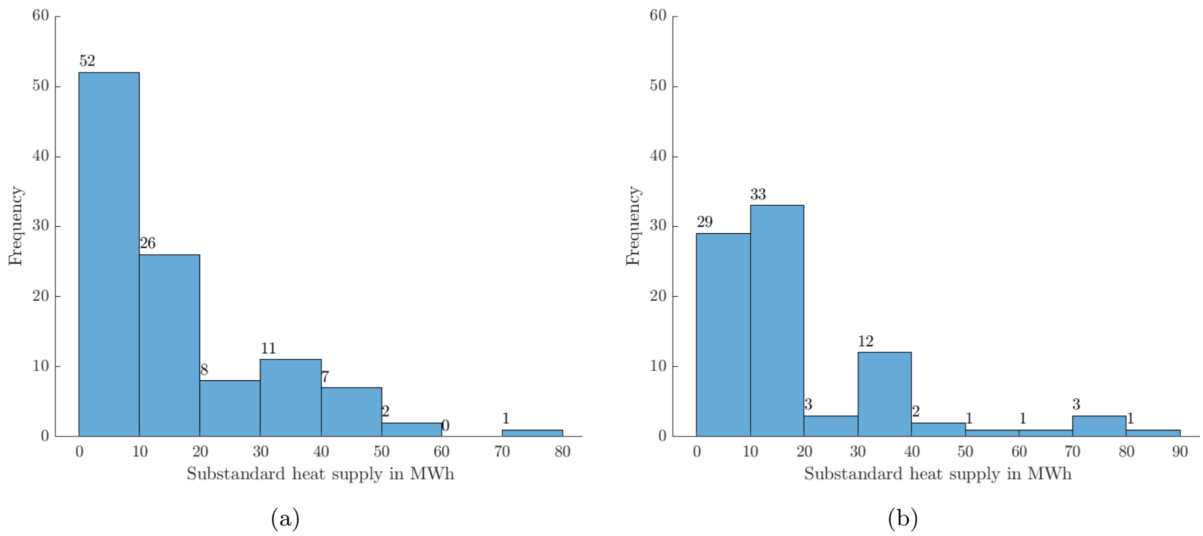

Figure 5.10: Histograms of substandard heat supply in MWh. The histogram of the maintenance optimisation experiment is presented in Figure 5.10a, and the histogram of the maintenance constraint experiment is presented in Figure 5.10b.

Lastly, the costs of maintenance and unexpected failures during the maintenance constraint experiment is compared to the costs in the maintenance optimisation experiment. The bar plots are presented in Figure 5.11. It is observed that the cost of maintenance of the CHP units in the maintenance constraint experiment has decreased slightly with respect to the maintenance optimisation experiment. Postponing the maintenance during periods of high heat demand has reduced the overall maintenance actions applied to the CHP units.

The same holds for the heat pumps. Due to the postponement of maintenance, fewer maintenance actions were applied to the heat pumps, reducing maintenance costs. It is difficult to observe from the bar plot; however, the cost of unexpected failures has increased by 1,500 EUR on average. Again, this is attributed to the fact that maintenance was postponed; thus, the heat pump’s failure probability was slightly higher.

The cost of maintenance and unexpected failures for the gas heater is equal in both experiments. No difference of statistical significance is observed.

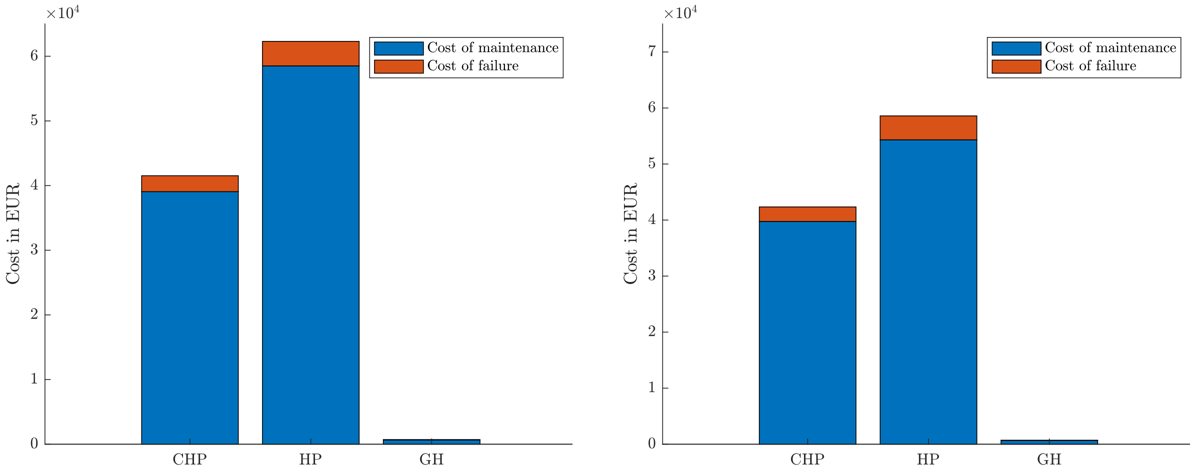

(a) (b)

Figure 5.11: Stacked bar plots of the average costs of maintenance and unexpected failures of assets. The costs are the average costs per individual asset over 10,000 MCS runs. The bar plot of the maintenance optimisation experiment is presented in Figure 5.11a, and the bar plot of the maintenance constraint experiment is presented in Figure 5.11b.

5.4 Boxplot of Substandard Heat Supply

A boxplot of the substandard heat supply during the Base Scenario, Maintenance Optimisation and Maintenance Constraint experiment are presented in Figure 5.12. The statistics attributed to the boxplots are presented in Table 5.4.

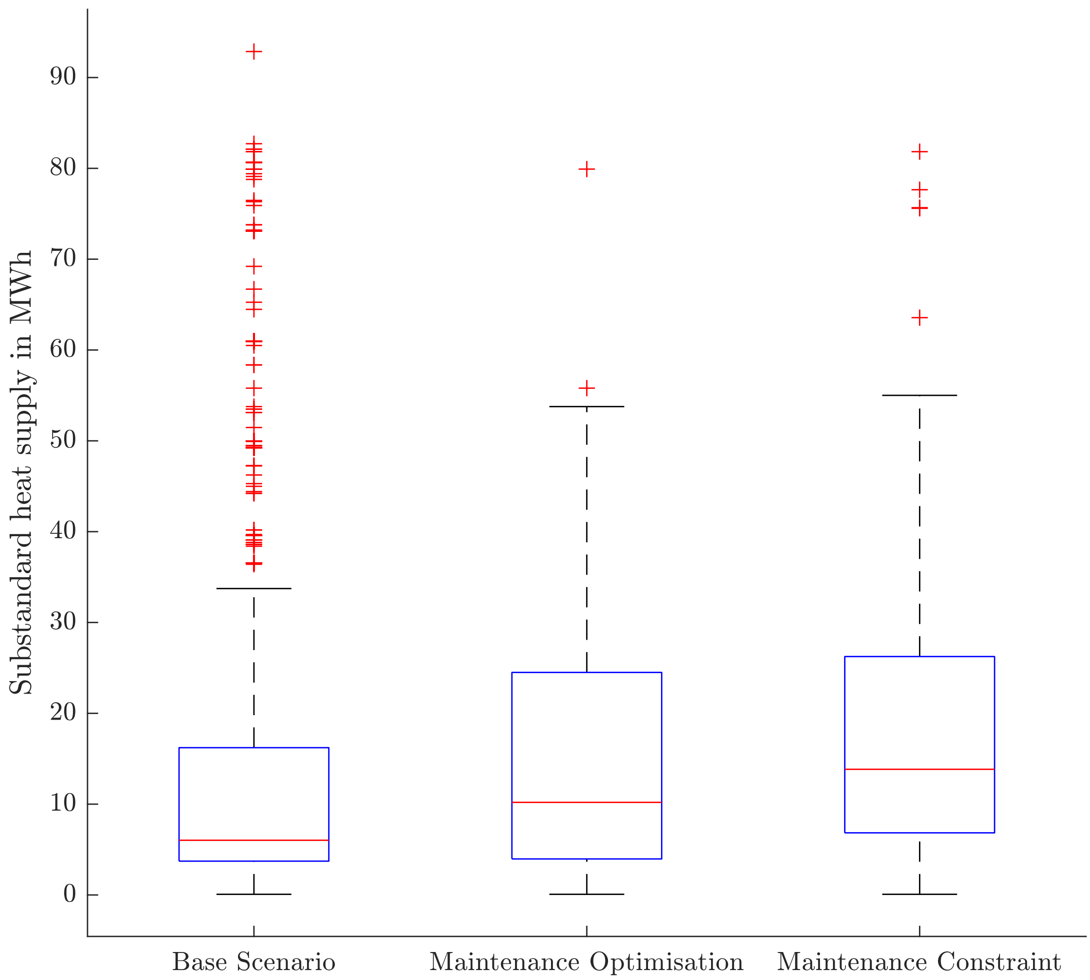

Figure 5.12: Boxplot of the substandard heat supply during the Base Scenario, Maintenance Optimisation and Maintenance Constraint experiment.

Table 5.4: Statistics of the substandard heat supply data that define the boxplots in Figure 5.12.

	Base Scenario	Maintenance Optimisation	Maintenance Constraint
Lower Adjacent	0.073	0.073	0.073
Upper Adjacent	33.735	53.768	55.007
Median	6.024	10.195	13.837
25^{th} Percentile	3.725	3.967	6.841
75^{th} Percentile	16.217	24.502	26.250
Maximum	92.863	79.918	81.837
Num Outliers	118	2	5
Num Point	1405	107	85

From the boxplots it is observed that the dispersion of the substandard heat supply has increased in the maintenance optimisation and maintenance constraint experiment with respect to the base scenario. This increase in dispersion is attributed to the fact that the number of occurrences of substandard heat supply is significantly lower than in the base scenario. Therefore, while the dispersion is greater, the number of occurrences of heat supply is significantly lower.

Furthermore, the number of outliers is in the base scenario significantly higher and of a higher value than in the maintenance optimisation and maintenance constraint experiment. Ultimately, the number of occurrences and amount of substandard heat supply is what imposes a financial risk on Warmtestad. Therefore, this risk is lower in the last two experiments.

5.5 Asset Statistics

The statistics of the assets from all three experiments are presented in Table 5.5. This table includes statistics of the individual assets during MCS, such as the average number of maintenance actions, average downtime of the assets in hours, average operating time in hours and the average availability and reliability. The reliability of the assets is defined as the proportion of assets that have not failed during MCS.

This table shows that the CHP units have nearly the same statistics throughout the three experiments, without any statistically significant difference. Idem ditto for the gas heaters. The heat pumps did undergo significantly more maintenance in the maintenance optimisation experiment than in the base scenario. This resulted in longer downtime and reduced availability.

In the maintenance constraint experiment, less maintenance was performed on the heat pumps than in the maintenance optimisation experiment. This is explained by the maintenance constraint that prohibited asset maintenance during periods of high head demand, increasing the time between maintenance slightly. As a result, the average downtime is slightly less, and the availability is slightly higher.

Table 5.5: Results of the three experiments conducted in this research. All values are mean values of 10,000 MCS runs. The average number of maintenance actions applied to each asset, average downtime and operating time in hours and the average availability and reliability.

Experiment		CHP1	CHP2	CHP3	CHP4	HP1	HP2	HP3	HP4	HP5	GH1	GH2	GH3
1	Maintenance	4.94	3.95	1.98	0	1.77	1.72	0.89	0.76	0	0.95	0	0
	Downtime [h]	501.0	400.1	200.8	0.2	184.4	178.2	100.6	77.9	0.5	25.0	0.1	0.0
	Operating time [h]	8123	6255	3740	756	8537	8310	5618	3861	695	5629	1788	789
	Availability	0.9419	0.9399	0.9490	0.9997	0.9789	0.9790	0.9824	0.9802	0.9992	0.9956	0.9999	1.0000
	Reliability	0.9338	0.9443	0.9710	0.9979	0.4774	0.5001	0.6169	0.7516	0.9929	0.9541	0.9960	0.9995
2	Maintenance	4.94	3.94	1.98	0	10.93	10.89	7.86	4.95	1.05	1.00	0	0
	Downtime [h]	500.3	400.0	200.1	0.3	825.7	822.3	594.7	375.0	80.1	25.5	0.1	0.0
	Operating time [h]	7975	6428	3684	781	7811	7600	5950	4066	1466	5678	1795	787
	Availability	0.9410	0.9414	0.9484	0.9996	0.9044	0.9024	0.9091	0.9156	0.9482	0.9955	0.9999	1.0000
	Reliability	0.9313	0.9396	0.9704	0.9972	0.9276	0.9320	0.9356	0.9555	0.9832	0.9765	0.9956	0.9997
3	Maintenance	4.93	3.71	1.96	0	9.50	8.99	6.07	2.60	0	0.99	0	0
	Downtime [h]	500.3	377.0	199.7	0.2	719.7	682.1	462.3	203.0	2.0	25.4	0.1	0.0
	Operating time [h]	8018	6442	3677	735	7918	7801	5918	4137	1182	5649	1797	787
	Availability	0.9413	0.9447	0.9485	0.9997	0.9167	0.9196	0.9275	0.9532	0.9983	0.9955	0.9999	1.0000
	Reliability	0.9297	0.9402	0.9633	0.9978	0.9066	0.9007	0.9129	0.8974	0.9736	0.9757	0.9964	0.9997

Chapter 6

Discussion

This chapter provides an interpretation of the results obtained from the experiments. The performance of the maintenance and reliability model and the maintenance optimisation algorithm developed during this research is assessed based on the implications of the results. Furthermore, the maintenance and reliability model's limitations are discussed, and suggestions for future work are provided.

Maintenance optimisation

The lower reliability bound for the maintenance optimisation algorithm was set to 80% for one year of operation. With this parameter, it is expected that the probability of the assets surviving for one year is 80%. For the CHP units, in roughly 84 out of 100 simulations during MCS, no failures have occurred for all five combined heat and power (CHP) units. This implies that the reliability is higher than the predetermined 80%, even though it is expected that the reliability constraint was dominant during the maintenance optimisation.

This can be explained by the fact that the 80% reliability is for one year of continuous operation. Naturally, this amount of operating time is not achieved by any CHP unit as the assets require maintenance, hence experiencing downtime a few times per year. However, more importantly, most of the CHP units are not required to operate for all hours of the year, as is observed from Table 5.5. Therefore, the CHP units show a higher reliability than initially expected by the 80% lower bound asset reliability. The same argumentation holds for the heat pumps and gas heaters.

In the maintenance optimisation experiment, the assets are more reliable as the number of random failures is lower than in the base scenario. This is achieved by performing maintenance more frequently. As a result, the availability of the assets reduces as the downtime increases. This is affirmed by the results in Table 5.5. However, it is interesting to observe that the reduced availability of the assets still resulted in a more reliable heat supply, as was observed in Figure 5.4.

The most evident explanation for this observation is the fact that the increased reliability in the maintenance optimisation experiment resulted in fewer random failures with respect to the base scenario, as is observed in Figure 5.1. The reduced number of random failures resulted in a more reliable heat production even though additional maintenance actions replaced the random failures. Conversely, in the base scenario, more random failures occurred due to the less reliable assets. In this case, the random failures were detrimental to the reliability of the heat supply, as they presumably occurred during high heat demand.

The costs attributed to maintenance and unexpected failures are equal due to the SLA of WarmteStad with a third party. As a result, the cost comparison in Figure 5.5 is somewhat trivial. In the maintenance optimisation experiment, the number of maintenance actions has increased, and as a result, the number of random failures is less than in the base scenario. Hence, the cost of maintenance has increased, and the cost of random failures has decreased. The cost of additional maintenance grows faster than the reduction in the cost of unexpected failures.

Maintenance constraint
In the maintenance constraint experiment it was expected that by imposing a constraint that restricted maintenance in periods of high demand, the reliability of the heat plant is improved. The results showed that this constraint lowered the number of heat supply failures. Furthermore, it was observed that in only 8.8% of the cases, a heat supply failure occurred due to preventive maintenance, while the maintenance optimisation experiment, this occurred in 15.6% of the cases. Therefore, an overall reduction in the number of heat supply failures caused by preventive maintenance was perceived.

Although, that does not provide a comprehensive view of the side effects, as the constraint resulted in postponed maintenance, which led to more random failures. For this reason, the number of heat supply failures did not reduce significantly. Furthermore, a few more heat supply failures at the higher end of the spectrum at 80-90 MWh of substandard heat supply were observed, which imposes a higher risk on WarmteStad to incur costs attributed to a heat supply failure. However, this peak was of equal size in a MCS run with double the number of iterations; therefore, the peak may be attributed to “bad luck” in the MCS in the original experiment.

Financial risk WarmteStad
Whether WarmteStad incurs a fine regarding heat supply failure is difficult to say. The quantity of substandard heat supply is computed by the model; however, further research is required to relate substandard heat supply to potential compensation that must be paid for a heat supply failure, as the duration of heat supply failure is also a factor of importance. At maximum capacity, the heat plant produces roughly 25 MW of thermal energy. Therefore, the worst scenarios observed during Monte Carlo simulation (MCS), where 50+ MWh of heat production is lost equates to two hours of maximum production of the heat plant. This occurred 59 times in the base scenario, 3 times in the maintenance optimisation experiment and 6 times in maintenance constraint experiment, out of 10,000 MCS runs. As a result, a heat supply failure is unlikely to occur.

The low probability of a heat supply failure has two main reasons. First, multiple assets per category create redundancy in the system. If one of the assets fails or undergoes maintenance, two or more assets might be available to compensate for the production loss. Second, the gas heaters have significantly more thermal production capacity than the HP and CHP and are the most reliable asset and have the shortest maintenance time. As a result, the heat plant is overall, highly reliable in terms of its ability to fulfil the heat demand.

Model assessment
The maintenance optimisation model worked according to its design requirements. However, not

all features were demonstrated due to the input parameters attributed to WarmteStad's case study. For example, the costs of PM and CM were equal, resulting in the probability of failure not being penalised with respect to the cost of PM in the optimisation of the cost function. As a result, the constraint where lower bounds are imposed on the reliability of the assets was leading. Though, sensitivity analysis of the input parameters can be part of future work.

Furthermore, the time frame for the maintenance optimisation was longer than the operating time of nearly all assets, resulting in more reliable operation than was determined by the "minimum reliability" constraint imposing lower bounds on the reliability of the assets for a specific time horizon.

The unit commitment algorithm functioned appropriately. Many interesting variables are stored, allowing for in-depth analysis of the experimental results. Additional features are easily added to research their effect on the reliability of heat supply of the heat plant.

The maintenance and reliability model allows for more indepth analysis of the O&M cost, as the number of corrective and preventive maintenance actions is computed. When comparing the costs of maintenance and unexpected failures in the base scenario to the fixed O&M costs found in the literature, see Table 4.2, it is observed that the costs for the CHP with 43,106 EUR falls in the middle of the interval 20,462-61,385 EUR. The costs of the heat pump with 31,048 EUR is slightly above the 30,000 EUR found in the literature. The costs of the gas heater are with 703 EUR far below the 21,000 EUR found in the literature. Though, for the gas heater it is known from the SLA that the maintenance costs for WarmteStad are are closer to the observed 703 EUR than the 21,000 found in the literature.

The costs for the CHP unit and gas heater in the maintenance optimisation and maintenance constraint experiment are close to the costs in the base scenario. The costs for the heat pump in the maintenance optimisation and maintenance constraint experiment is nearly double the costs in the base scenario. Therefore, it is concluded that the maintenance costs for the CHP unit are accurate based on the costs found in the literature. For the heat pump the costs are accurate for the base scenario, though, for the other two experiments the costs are too high. The costs for the gas heater are too low compared to the values found in the literature, though, it is expected that the literature values are high with respect to the true costs for WarmteStad.

Limitations

A limitation of the maintenance optimisation model is that it is highly abstract. The failure rate of the assets is based on a statistical distribution, and it is assumed that maintenance actions can extend the component's lifetime indefinitely while retaining the same failure distribution.

Furthermore, the solver used to solve the maintenance optimisation algorithm could only handle a one year time frame in one-week intervals. It is unclear to what extent the reduced resolution affected the results, though it is expected to be minimal.

A limitation of the unit commitment algorithm is that it is based on simulation data of the digital twin; as a result, the production of the assets is decoupled from the production in the digital twin. A major drawback of this approach is that the model cannot anticipate on scheduled maintenance of an asset by increasing the thermal production shortly before maintenance. If this anticipation is possible, the buffer can be filled before maintenance is performed; as a result, the amount of lost heat production is expected to be reduced further.

Implications for the digital twin

The models developed during this research allow the user to explore the reliability of WarmteStad's heat plant for different scenarios based on the digital twin simulation data. Though, as mentioned as a limitation, the thermal production of the assets is decoupled from the thermal production in the digital twin. For more accurate results, it is recommended to implement constraints on the availability of the assets in the digital twin. In that case, it becomes possible to anticipate on maintenance by increasing the energy content of the thermal buffer beforehand, which is a significant advantage to determining the probability of a heat supply failure. Furthermore, the feature of random asset failures can be modelled using the Weibull distribution, as is performed in this work.

The disadvantage of implementing constraints to the assets of the digital twin is that it is significantly computationally more expensive than running the "post-processing" MCS maintenance and reliability model developed during this research. This allows for faster scenario exploration with varying asset reliabilities.

Future Work

In the future:

- To increase the performance of the maintenance optimisation algorithm, a heuristic solver such as *genetic algorithm* or *simulated annealing* could be used to solve for a longer time horizon with shorter intervals, increasing the resolution of the optimisation problem. In addition, downtime can be added to the maintenance optimisation algorithm to improve its accuracy.
- The parameterisation of the model input parameters can be further improved upon for the specific case of WarmteStad. For example, currently a fixed downtime is attributed to the maintenance actions. This can be extended to handle varying downtime for different maintenance actions.
- Research is required to study the effect of substandard heat supply on the power transfer to the customers connected to the district heating network. Determine what quantity of substandard heat supply leads to a heat supply failure.
- In addition, a quantitative risk-analysis must be performed to relate the risk of a potential heat supply failure to a financial risk, where the cost of financial compensation of the customers is related to a certain quantity of substandard heat supply.
- The assets are activated numerically per category, as a result, the assets that are activated last, only operate a fraction of the time with respect to the assets activated first. Future work could investigate whether a different commitment sequence that balances the operating time of the assets better, can improve the reliability of heat supply.
- The model can be further validated by performing sensitivity analysis on the model input parameters. For example, on the constraint imposed on the lower bound of the assets reliability, or by using different cost ratios for preventive and corrective maintenance.

Chapter 7

Conclusion

In this work, a maintenance and reliability model for the virtual heat assets in the digital twin of WarmteStad's heat plant was developed. Using a maintenance optimisation model, the maintenance schedule of the assets was optimised by minimising the cost of maintenance and unexpected failures of the assets. The most important condition for the optimisation model was the lower bound imposed on the reliability of the assets over a given time horizon. A unit commitment algorithm was developed to simulate the activation and deactivation of the assets. In this algorithm, features such as random failures of the assets, compensating for production loss and a maintenance constraint were implemented. Monte Carlo simulation was used to quantify the reliability of the entire heat plant by iterating the unit commitment algorithm 10,000 times.

The performance of the maintenance optimisation model was tested by conducting an experiment where an 80% reliability constraint was imposed on the lower bound of the reliability of the assets over a one year time horizon. The results showed that a reduction of 85.9% in the occurrences of substandard heat supply was achieved. This resulted from increased asset reliability, that reduced the number of random failures during critical periods.

Furthermore, a scenario was tested where a constraint was imposed that restricted maintenance of the assets during periods of high heat demand. This further reduced the number of occurrences of true substandard heat supply by 11.5%, with respect to the experiment where the maintenance optimisation was tested. However, besides a reduction in the occurrence of true substandard heat supply, an increase in the high production loss was observed at 70-90 MWh, although, this increase is presumably caused by "bad luck" in the Monte Carlo simulation.

In conclusion, the maintenance optimisation algorithm and unit commitment algorithm functioned according to their design requirements. Though, future work is necessary to refine the parameters for WarmteStad's case study and relate the quantity of substandard heat supply to the financial risk imposed on WarmteStad using a quantified risk-analysis.

The artefact of this research is the maintenance and reliability model that consists of the maintenance optimisation algorithm and the unit commitment algorithm combined with Monte Carlo simulation. The developed model allows for indepth reliability analysis of the virtual assets in the digital twin of WarmteStad's heat plant, based on the digital twin simulation data.

Bibliography

[1] B. De Jonge and P. A. Scarf, "A review on maintenance optimization," *European journal of operational research*, vol. 285, no. 3, pp. 805–824, 2020.

[2] M. Doostparast, F. Kolahan, and M. Doostparast, "A reliability-based approach to optimize preventive maintenance scheduling for coherent systems," *Reliability Engineering & System Safety*, vol. 126, pp. 98–106, 2014.

[3] N. K. Srivastava and S. Mondal, "Development of predictive maintenance model for n-component repairable system using nhpp models and system availability concept," *Global Business Review*, vol. 17, no. 1, pp. 105–115, 2016.

[4] Y. Gao, Y. Feng, Z. Zhang, and J. Tan, "An optimal dynamic interval preventive maintenance scheduling for series systems," *Reliability Engineering & System Safety*, vol. 142, pp. 19–30, 2015.

[5] M. Bevilacqua and M. Braglia, "The analytic hierarchy process applied to maintenance strategy selection," *Reliability Engineering & System Safety*, vol. 70, no. 1, pp. 71–83, 2000.

[6] R. H. Yeh, K.-C. Kao, and W. L. Chang, "Optimal preventive maintenance policy for leased equipment using failure rate reduction," *Computers & Industrial Engineering*, vol. 57, no. 1, pp. 304–309, 2009.

[7] B. Lin, J. Wu, R. Lin, J. Wang, H. Wang, and X. Zhang, "Optimization of high-level preventive maintenance scheduling for high-speed trains," *Reliability Engineering & System Safety*, vol. 183, pp. 261–275, 2019.

[8] F. Pargar, O. Kauppila, and J. Kujala, "Integrated scheduling of preventive maintenance and renewal projects for multi-unit systems with grouping and balancing," *Computers & Industrial Engineering*, vol. 110, pp. 43–58, 2017.

[9] G. Budai, D. Huisman, and R. Dekker, "Scheduling preventive railway maintenance activities," *Journal of the Operational Research Society*, vol. 57, pp. 1035–1044, 2006.

[10] "About KWR," https://www.kwrwater.nl/en/about-kwr/.

[11] T. Nakagawa, "Sequential imperfect preventive maintenance policies," *IEEE Transactions on Reliability*, vol. 37, no. 3, pp. 295–298, 1988.

[12] G. Levitin and A. Lisnianski, "Optimization of imperfect preventive maintenance for multi-state systems," *Reliability Engineering & System Safety*, vol. 67, no. 2, pp. 193–203, 2000.

[13] A. Tam, W. M. Chan, and J. W. H. Price, "Optimal maintenance intervals for a multi-component system," *Production Planning and Control*, vol. 17, no. 8, pp. 769–779, 2006.

[14] M. Badami, A. Fonti, A. Carpignano, and D. Grosso, "Design of district heating networks through an integrated thermo-fluid dynamics and reliability modelling approach," *Energy*, vol. 144, pp. 826–838, 2018.

[15] "Wij zijn warmtestad," https://warmtestad.nl/warmtestad/, 2022.

[16] F. Tao, B. Xiao, Q. Qi, J. Cheng, and P. Ji, "Digital twin modeling," *Journal of Manufacturing Systems*, vol. 64, pp. 372–389, 2022.

[17] B. Xu, J. Wang, X. Wang, Z. Liang, L. Cui, X. Liu, and A. Y. Ku, "A case study of digital-twin-modelling analysis on power-plant-performance optimizations," *Clean Energy*, vol. 3, no. 3, pp. 227–234, 2019.

[18] "KWR en warmtestad werken aan verduurzaming warmtecentrale met een digital twin," https://www.kwrwater.nl/actueel/kwr-en-warmtestad-ontwikkelen-digital-twin-voor-duurzame-warmtecentrale/?highlight=warmtestad, 2021.

[19] A. Lisnianski, G. Levitin, H. Ben-Haim, and D. Elmakis, "Power system structure optimization subject to reliability constraints," *Electric Power Systems Research*, vol. 39, no. 2, pp. 145–152, 1996.

[20] G. Angelini, A. Lucchini, and G. Manzolini, "Comparison of thermocline molten salt storage performances to commercial two-tank configuration," *Energy Procedia*, vol. 49, pp. 694–704, 2014.

[21] B. M. Mun, P. H. Kvam, and S. J. Bae, "Mixed-effects nonhomogeneous poisson process model for multiple repairable systems," *IEEE Access*, vol. 9, pp. 71 900–71 908, 2021.

[22] K. S. Moghaddam and J. S. Usher, "Sensitivity analysis and comparison of algorithms in preventive maintenance and replacement scheduling optimization models," *Computers & Industrial Engineering*, vol. 61, no. 1, pp. 64–75, 2011.

[23] K. S. Moghaddam, *Preventive maintenance and replacement scheduling: models and algorithms.* University of Louisville, 2010.

[24] B. de Jonge, A. S. Dijkstra, and W. Romeijnders, "Cost benefits of postponing time-based maintenance under lifetime distribution uncertainty," *Reliability Engineering & System Safety*, vol. 140, pp. 15–21, 2015.

[25] J. Van Noortwijk and H. Klatter, "The use of lifetime distributions in bridge maintenance and replacement modelling," *Computers & Structures*, vol. 82, no. 13-14, pp. 1091–1099, 2004.

[26] V. V. Krivtsov, "Practical extensions to nhpp application in repairable system reliability analysis," *Reliability Engineering & System Safety*, vol. 92, no. 5, pp. 560–562, 2007.

[27] L. Wang, Y. M. Tripathi, S. Dey, C. Zhang, and K. Wu, "Analysis of dependent left-truncated and right-censored competing risks data with partially observed failure causes," *Mathematics and Computers in Simulation*, vol. 194, pp. 285–307, 2022.

[28] M. A. Al-Fawzan, "Methods for estimating the parameters of the weibull distribution," *King Abdulaziz City for Science and Technology, Saudi Arabia*, 2000.

[29] K. C. Kapur and M. Pecht, *Reliability engineering.* John Wiley & Sons, 2014, vol. 86.

[30] M. Kijima, H. Morimura, and Y. Suzuki, "Periodical replacement problem without assuming minimal repair," *European Journal of Operational Research*, vol. 37, no. 2, pp. 194–203, 1988.

[31] J. C. Smith and Z. C. Taskin, "A tutorial guide to mixed-integer programming models and solution techniques," *Optimization in Medicine and Biology*, pp. 521–548, 2008.

[32] A. C. Marquez, A. S. Heguedas, and B. Iung, "Monte carlo-based assessment of system availability. a case study for cogeneration plants," *Reliability Engineering & System Safety*, vol. 88, no. 3, pp. 273–289, 2005.

[33] D. C. Alexander *et al.*, "Application of monte carlo simulations to system reliability analysis," in *Proceedings of the 20th International Pump Users Symposium.* Texas A&M University. Turbomachinery Laboratories, 2003.

[34] H. Ge and S. Asgarpoor, "Parallel monte carlo simulation for reliability and cost evaluation of equipment and systems," *Electric power systems research*, vol. 81, no. 2, pp. 347–356, 2011.

[35] W. Betz, I. Papaioannou, and D. Straub, "Bayesian post-processing of monte carlo simulation in reliability analysis," *Reliability Engineering & System Safety*, vol. 227, p. 108731, 2022.

[36] S. Xiao and W. Nowak, "Reliability sensitivity analysis based on a two-stage markov chain monte carlo simulation," *Aerospace Science and Technology*, vol. 130, p. 107938, 2022.

[37] W. Li *et al.*, *Reliability assessment of electric power systems using Monte Carlo methods.* Springer Science & Business Media, 2013.

[38] C. Singh and J. Mitra, "Monte carlo simulation for reliability analysis of emergency and standby power systems," in *IAS'95. Conference Record of the 1995 IEEE Industry Applications Conference Thirtieth IAS Annual Meeting*, vol. 3. IEEE, 1995, pp. 2290–2295.

[39] S. Zhang, M. Wen, H. Cheng, X. Hu, and G. Xu, "Reliability evaluation of electricity-heat integrated energy system with heat pump," *CSEE Journal of Power and Energy Systems*, vol. 4, no. 4, pp. 425–433, 2018.

[40] I. Postnikov, "Methods for the reliability optimization of district-distributed heating systems with prosumers," *Energy Reports*, vol. 9, pp. 584–593, 2023.

[41] J. Jiang, X. Wei, W. Gao, S. Kuroki, and Z. Liu, "Reliability and maintenance prioritization analysis of combined cooling, heating and power systems," *Energies*, vol. 11, no. 6, p. 1519, 2018.

[42] K. Darrow, R. Tidball, J. Wang, and A. Hampson, "Catalog of chp technologies; 2015," *US Environmental Protection Agency and the US Department of Energy: Washington, DC, USA*, p. 23, 2017.

[43] S. Meyers, B. Schmitt, and K. Vajen, "The future of low carbon industrial process heat: A comparison between solar thermal and heat pumps," *Solar Energy*, vol. 173, pp. 893–904, 2018.

[44] J. Cox, S. Belding, and T. Lowder, "Application of a novel heat pump model for estimating economic viability and barriers of heat pumps in dairy applications in the united states," *Applied Energy*, vol. 310, p. 118499, 2022.

[45] G. Davies and P. Woods, "The potential and costs of district heating networks," *Pöyry Energy*, p. 101, 2009.

[46] J. Savolainen and M. Urbani, "Maintenance optimization for a multi-unit system with digital twin simulation: Example from the mining industry," *Journal of Intelligent Manufacturing*, vol. 32, no. 7, pp. 1953–1973, 2021.

[47] F. S. Almeida, F. H. Silveira, and S. Visacro, "Stopping criterion for monte carlo method-based simulations of the lightning performance of transmission lines," *Electric Power Systems Research*, vol. 214, p. 108797, 2023.

[48] J. Löfberg, "Yalmip : A toolbox for modeling and optimization in matlab," in *In Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.

[49] C. D'Ambrosio, A. Lodi, and S. Martello, "Piecewise linear approximation of functions of two variables in milp models," *Operations Research Letters*, vol. 38, no. 1, pp. 39–46, 2010.