GETTING A BETTER GRIP ON FORECASTING INTERMITTENT PARTS

ON THE EVALUATION OF DIFFERENT FORECASTING TECHNIQUES AND THE RESULTING INVENTORY CONTROL PARAMETERS: A COMPARATIVE STUDY

Erwin van Wingerden, BSc
March 21, 2012

Erasmus University Rotterdam & Gordian Logistic Experts B.V.

Supervisors:
Dr. Ir. R.J.I. Basten
Prof. Dr. Ir. R. Dekker
W. van Jaarsveld, MSc
Dr. Ir. J.W. Rustenburg

Management Summary

To raise their competition, providing good after sales service plays an important role for many companies. In order to provide this service, companies need to have the right amount of spare parts on hand. When there is too much stock on hand, the holding costs become too high; when there is not enough stock on hand, this will result in unavailability of the capital good, and thus the primary process of its user. To be able to make a good trade-off, it is key to accurately forecast demand. This has always been difficult, but due to improvements in technology failure rates generally decrease, which makes it even more difficult to do so. As a result, selecting the best method to determine the amount of spare parts on hand is very important. In order to find the best forecasting method to determine the inventory level for each part, we evaluate five different traditional forecasting methods and two methods based on an Empirical distribution. Based on the inventory costs and service level that are achieved for each method, we suggest a classification that enables us to improve the forecasting of spare parts.

We use the data of the following three companies in our study:

1. NedTrain, the maintenance organisation performing the maintenance, service, overhaul, and cleansing of rolling stock for the Netherlands Railways

2. Alstom Transport, a manufacturer of various assets that are used within the rail sector

3. The Defence Material Organisation, a service center responsible for all material within the Defence organisation

We analyse the data, by looking at the demand patterns, the supply patterns, and the price distribution. An important observation is that all three companies deal with many intermittent parts, parts that have almost no demand, although at NedTrain Haarlem there is also a relatively large group of fast movers.

We then look at a variety of traditional forecasting methods, namely Moving Average, Single Exponential Smoothing, Double Exponential Smoothing, Croston, and the Syntetos-Boylan Approximation. We fit these forecast results to a Normal Distribution, and use the Mean Squared Error as measure for the variation of demand.

We expect that this way of forecasting will not work well for intermittent parts. As these parts are rarely requested, the Normal Distribution might not give a good fit to the actual demand pattern.

Especially for the intermittent parts, we introduce two methods based on an empirical distribution. The first method is proposed by Porras and Dekker (2008). It derives an empirical distribution by sampling consecutive demands for a fixed lead time value. We also introduce the Empirical plus method, which is able to take the variation in the lead time into account. Instead of always sampling using the same lead time, this method takes the different lead times, that have occurred in the past, into account by sampling over different lengths.

These methods are more likely to perform better for the intermittent parts, as we derive the distribution immediately, and do not require the fitting to a common distribution.

The traditional forecasting methods, and the two methods using an empirical distribution require different steps to be taken in order to obtain the required inventory level. Figure 1 gives an overview of the different steps that need to be taken.

We use simulation in order to calculate the performance in terms of (order line) fill rate, and holding costs. As inventory policy we decide to use an (R, nQ) policy, with Q the minimum order quantity or modular order quantity, if there is one specified, and 1 otherwise.

We then try to make a classification by looking at the influence of the following criteria on the performance of the different methods:

1. Average inter-demand-interval
2. Squared coefficient of variation of demand
3. Average Lead time
4. Average demand size
5. Price

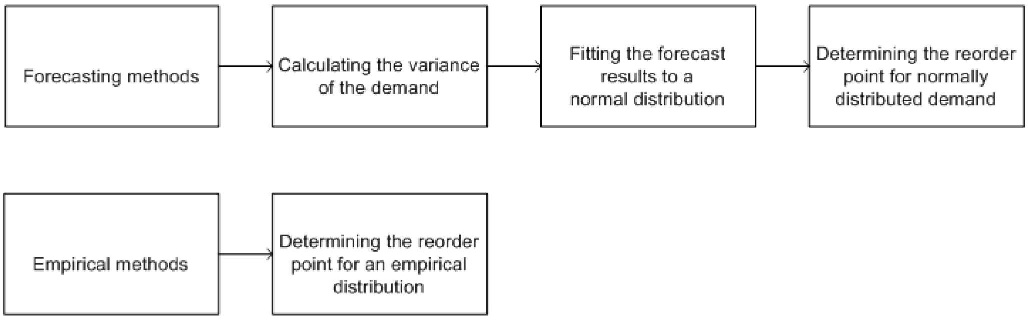

Figure 1: Overview of two different methods to obtain the reorder point.

From these criteria, we find that the average inter-demand-interval, and the squared coefficient of variation of demand have the most influence on the performance of the different methods. We find that the Syntetos-Boylan Approximation gives the best performance in almost all different classes of parts, except for the slow moving parts with a low coefficient of variation in the demand. For this last group, the Empirical methods give the best performance. After having determined the breaking points for these two classification criteria, we suggest the classification given in Figure 2. Using this classification enables us to decrease the holding costs while maintaining the same amount of service, or increase the amount of service while maintaining the same amount of holding costs.

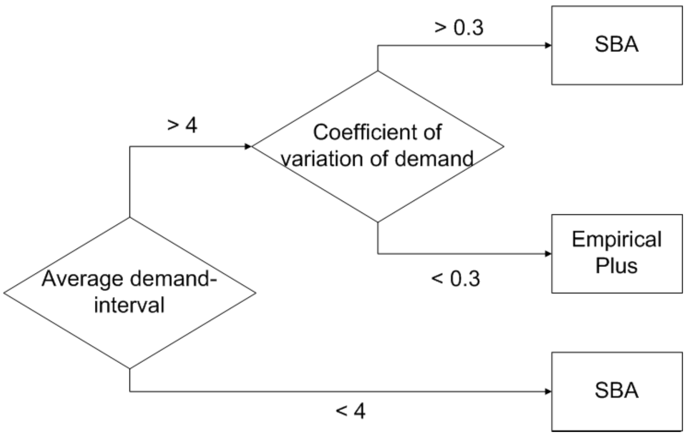

Figure 2: Classification for spare parts based on the average demand-interval and the coefficient of variation of demand

Contents

Chapter 1

1 Introduction

1.1 General introduction

The after sales service plays an important role for companies to stay competitive. However, to be able to provide this service against the lowest costs requires the companies to have a good inventory control. When there is not enough stock on hand to perform the maintenance, this will result in unavailability, downtime, or delayed maintenance. The costs of unavailability, especially in the case of capital goods, are often very high. However, when there is too much stock on hand, the holding costs become too high. There is a trade-off between the costs of unavailability and holding costs. To be able to make a good trade-off, it is key to accurately forecast demand.

Due to improvements in technology, however, there are less failures over time, which makes it more difficult to accurately forecast the demand. Companies are faced with these challenges, and are struggling to tackle these challenges (see, e.g., Cohen, Agrawal, and Agrawal (2006)). The reasons behind these problems are clear as companies, especially the Small Maintenance Organisation (SMOs), lack adequate knowledge and experience in the field of spare parts management. They are having difficulties to maintain spare parts planning experts, and are faced with an ever increasing complexity in terms of supply chain dynamics and technology. They do not have the scale of business to invest in knowledge and training, and are not able to afford the required information technology. Because of these reasons, there is an increasing demand for logistic professionals, whereas the number of logistic professionals available tends to decrease (Supply Managers, 1997).

To challenge these problems, Gordian initiated a project called Planning Services. The goal of this project is to build a Spare Parts Planning Control Tower (SPPCT), which is able to take over (a part of) the inventory control from one central location. This tower will be controlled by expert planners of Gordian. Figure 3 represents this idea.

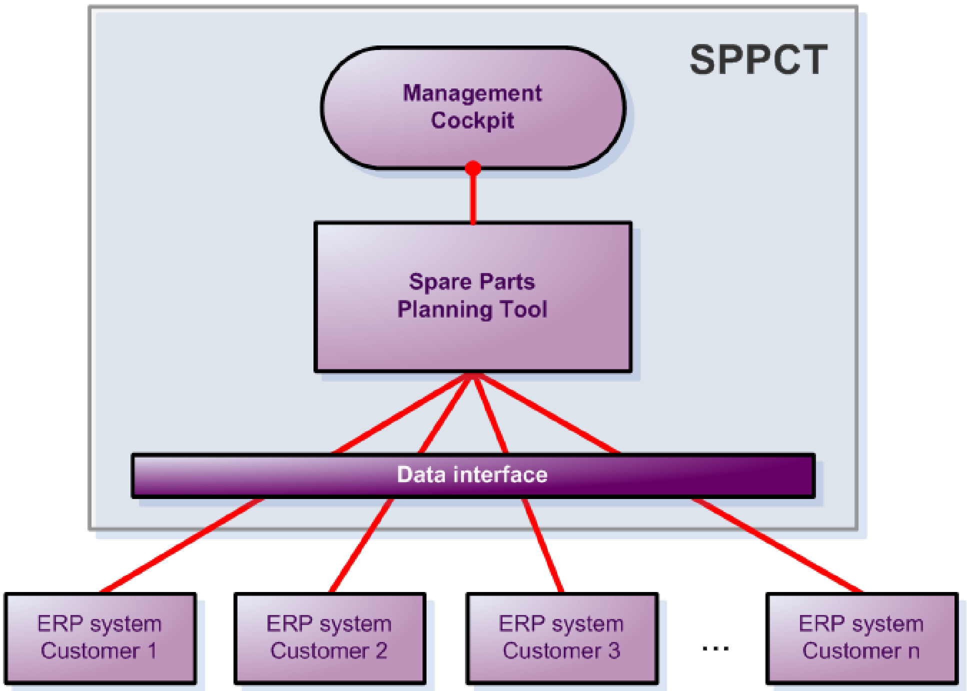

Figure 3: Spare Parts Planning Control Tower

Although the SPPCT is mostly intended for SMOs, it could also be used at Large Scale Maintenance Organisations (LSMOs). In that case the LSMO will have its own SPPCT. The SPPCT should be able to handle all the data from different SMOs, to do the planning of the spare parts, and to actually take over the ordering of the parts. Wouters (2010) has already done some research by investigating the requirements of the SPPCT, and to make clear what is needed to achieve this goal. The handling of the data, and the design of the management cockpit described in Figure 3 will not be a part of this research.

We go more into detail about the Spare Parts Planning Tool (SPPT), which is a part of the SPPCT, in Section 1.2. We will focus mainly on the forecasting and inventory control. In Section 1.3, we make clear what problems we need to deal with, and what we will investigate in order to increase the performance of the SPPT. In Section 1.4, we describe our methodology, and in Section 1.5, we will give a summary of what we have delivered. Finally in Section 1.6, we give an outline of this Master Thesis.

1.2 Spare Parts Planning Tool

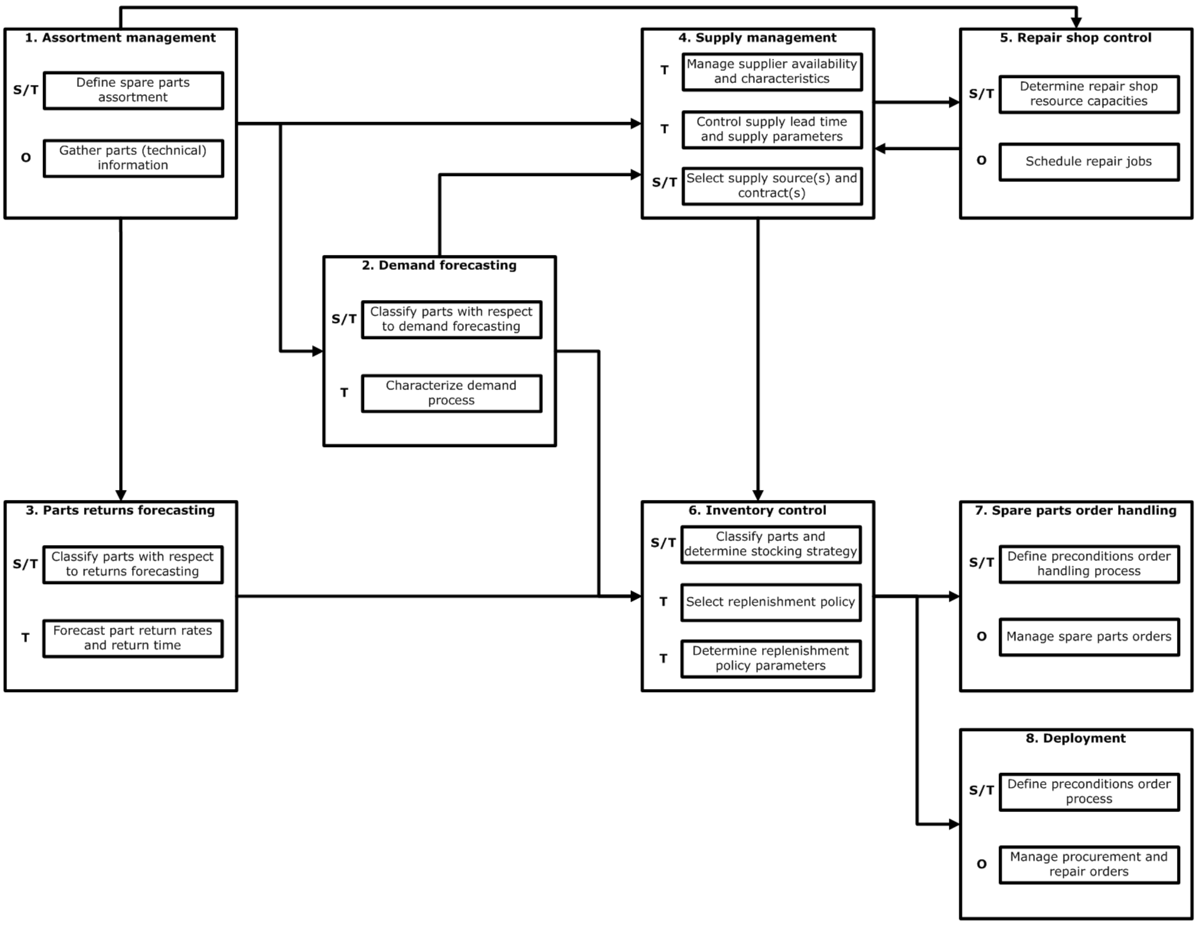

Figure 4: Overview and clustering of decisions in maintenance logistics control.

As mentioned, the SPPT should be able to do the planning of spare parts. As the planning consists of several steps, we first give a schematic overview of the requirements of the SPPT, based on the Framework for maintenance spare parts planning and control by Driessen, Arts, van Houtum, Rustenburg, and Huisman (2010). This overview is given in Figure 4, where S stands for strategic, T stands for Tactical, and O stands for Operational.

There will be three different versions of the SPPT. The first version consists of the actual forecasting, inventory control, spare parts order handling, and deployment. In our research, we focus on the forecasting and inventory control: the two clusters that are required in order for the SPPT to function properly. The other clusters will be implemented in a later version of the SPPT.

The methods and amount of clusters that will be included in each version, depends on the needs and interests of the companies. Companies may not be interested or able to outsource the complete process when the first version of the SPPT is built. They first would like to be sure that the SPPT works correctly, and get familiar with the SPPT before they decide to outsource their entire process.

In order to be able to make a forecast, we need historical demand data for all the parts, and also information regarding the lead time, Minimum Order Quantity (MOQ), Module Quantity (MOD), and the price. Based on this information, we can determine the inventory policy, and the corresponding parameters. There are different ways to obtain these results; we will explain how we obtained our results, and what inventory policy we use, in Chapter 5.

1.3 Problem Description

As becomes clear from the previous section, it is important that all the SPPT clusters work well together. Because the scope would be too big to look at all the clusters we focus on the demand forecasting and the inventory control.

There has already been quite some research regarding the forecasting methods, and how to determine the accuracy based on indicators such as Mean Squared Error (MSE) and Mean Absolute Percentage Error (MAPE) (see, e.g., Willemain, Smart, Shockor, and DeSautels (1994), Ghobbar and Friend (2003), and Eaves and Kingsman (2004)). However, performance measures such as the MSE do not necessarily give good insights into the performance in terms of achieved service levels and holding costs. Teunter and Sani (2009) do not only look at the MSE but also at the achieved service level. Also Boylan, Syntetos, and Karakostas (2008) classify such that the forecasting methods perform better in combination with inventory control. Still they remark that more work has to be carried out.

For parts that have a more regular demand, it is easier to put the right amount of stock on hand, since there is less variation in the demand. However, for parts with intermittent demand, we expect that the standard way of forecasting will not give good results as it is more difficult to accurately fit the demand for these parts to a distribution.

Another weak point in the recent forecasting methods is that it is often assumed that the lead time is deterministic, where in reality the lead times are not always equal. It is also the case that the forecast is made for one period, whereas the lead time is often not equal to one period. In order to obtain a different length of the period, they multiply the forecast and variation. This could lead to a lower performance.

The research question we would like to answer is:
How to forecast demand for intermittent parts in such a way that it enables us to put the right amount of stock on hand?
To be able to answer this question, we try to answer the following questions:

1. How to determine the Lead Time Demand (LTD) distribution as accurately as possible?

2. How to classify parts such that we choose the right forecasting method for each part?

1.4 Methodology

To answer the questions stated in Section 1.3 we use the data of the following three companies:

- NedTrain Haarlem
- Alstom Transport Ridderkerk
- The Naval Maintenance Company

We use the data to perform a data analysis as described in Chapter 2 for all three companies. Based on these results, we can obtain some expectations of what results we will obtain in our simulation study. We first look at different methods to forecast the demand in Chapter 3. In Chapter 4 we then look at an Empirical method that uses an empirical distribution to obtain inventory policy parameters. We also extend this method such that it is able to take realised lead times into account.

In our simulation study, we do not look at performance measures such as the MSE or the MAPE; we look at the service level and holding costs we achieve for each method, given a fixed service level target. Based on the results we make a classification that enables us to choose the best method for each different group of parts.

1.5 Deliverables

This research results in the following deliverables:

- New method to determine the inventory policy parameters using an empirical distribution
- Insights into the performance of different methods
- Classification for spare parts
- Implementation of the simulation and forecasting methods in Visual Basic in Excel
- This Master's thesis

1.6 Outline

In Chapter 2 we give some background information of the three companies of which we obtain the data. We also analyse the data in order to obtain some characteristics of the data. In Chapter 3 we look at a number of traditional forecasting methods, and describe how we use these forecast results to fit the demand according to the normal distribution. We then describe how we obtain the reorder point using this distribution. We also describe two methods that use an empirical distribution in order to obtain the reorder point in Chapter 4. In Chapter 5 we explain the inventory policy which we use, and how our simulation is set up. We then evaluate the results we obtain from the simulation in Chapter 6. In Chapter 7 we give our conclusions and recommendations for further research.

Chapter 2

2 Companies and their data

In this chapter, we will give some insights into the companies that participate in this study. In Section 2.1 we first look at NedTrain, which performs the maintenance, service, overhaul, and cleansing of rolling stock for the Dutch Railways (NS). We then look at Alstom Transport, a manufacturer of various assets that are used within the rail sector, in Section 2.2. Finally we look at the Defence Material Organisation (DMO), a service centre responsible for all material within the Defence organisation, in Section 2.3. For each company, we describe its main activities, and then go more into detail about the department of which we obtain our data. For this part we have used company websites, information from Gordian's consultants, and power point presentations, which have been supplied by the companies. We then finally describe and analyse the data.

2.1 NedTrain

NedTrain is specialised in maintenance, service, cleansing, modernisation, overhauling, and damage recovery for rolling stock for the passenger railway. With 3,100 FTE's distributed over 43 sites, NedTrain maintains their customers' rail road cars and locomotives 24/7. NedTrain has service locations distributed across the country, and a Head Office located in Utrecht. An overview of the locations is given in Figure 5.

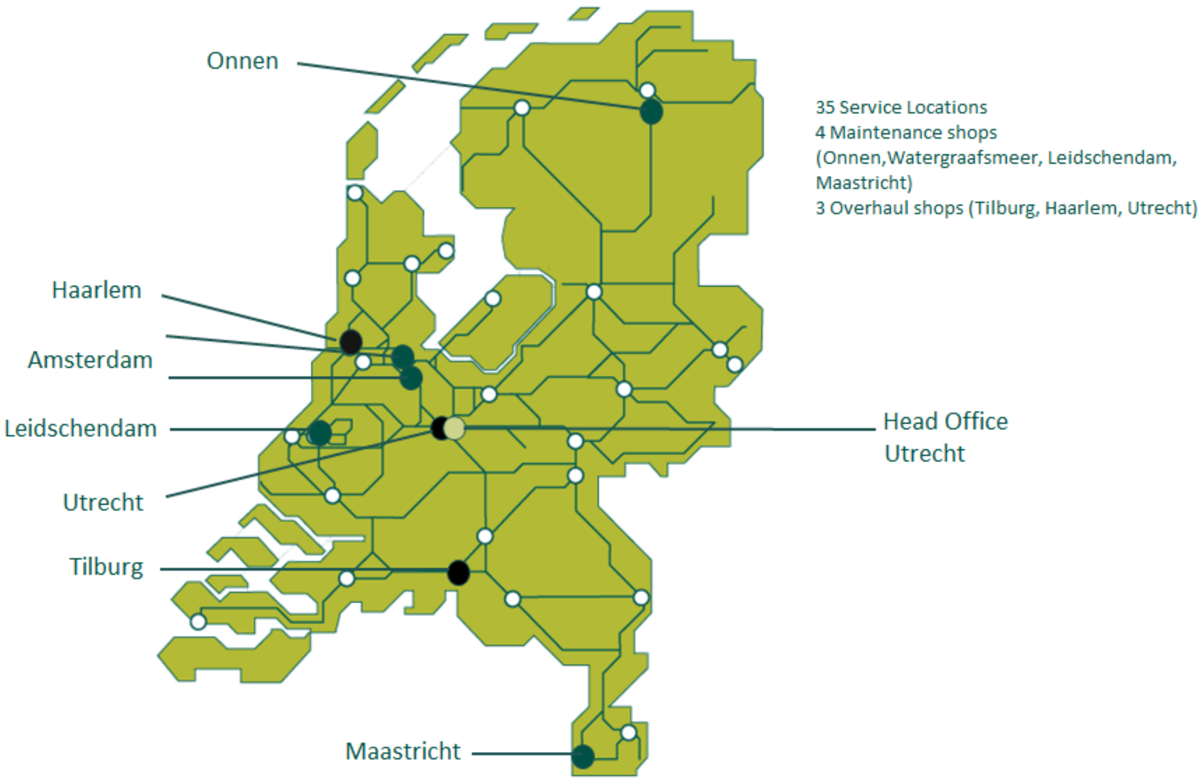

Figure 5: Locations of NedTrain

NedTrain has many customers, of which the Dutch Railways (NS) is the largest customer. The service locations (Dutch: *Servicebedrijven*) are responsible for the daily maintenance of the train. A train needs to be serviced every 13 weeks. When a train needs to be serviced, it will be sent to the Maintenance Shop (Dutch: *Onderhoudsbedrijf*). The maintenance shops have more equipment and parts compared to the service locations. When the maintenance for the wheelsets or bogies is likely to take more than two hours, or when the Maintenance Shop is not available to perform the maintenance, the train will be sent to one of the overhaul shops (Dutch: *revisie bedrijven*). Beside this, the Overhaul shops are also responsible for the overhauling, converting, and upgrading of main parts (Dutch: *Hoofddelen*). The main parts are repaired by replacing or repairing sub parts, which have one of the following characteristics:

1. Repairable (Dutch: *Wisseldelen*), these parts are repaired and reused after a failure
2. Consumable (Dutch: *Slijtdelen*), parts that are not reused after a failure

Each part also has one of the following characteristics:

1. '*Maakdelen*', parts that are manufactured by NedTrain
2. '*Inkoopdelen*', parts that are purchased from an external manufacturer

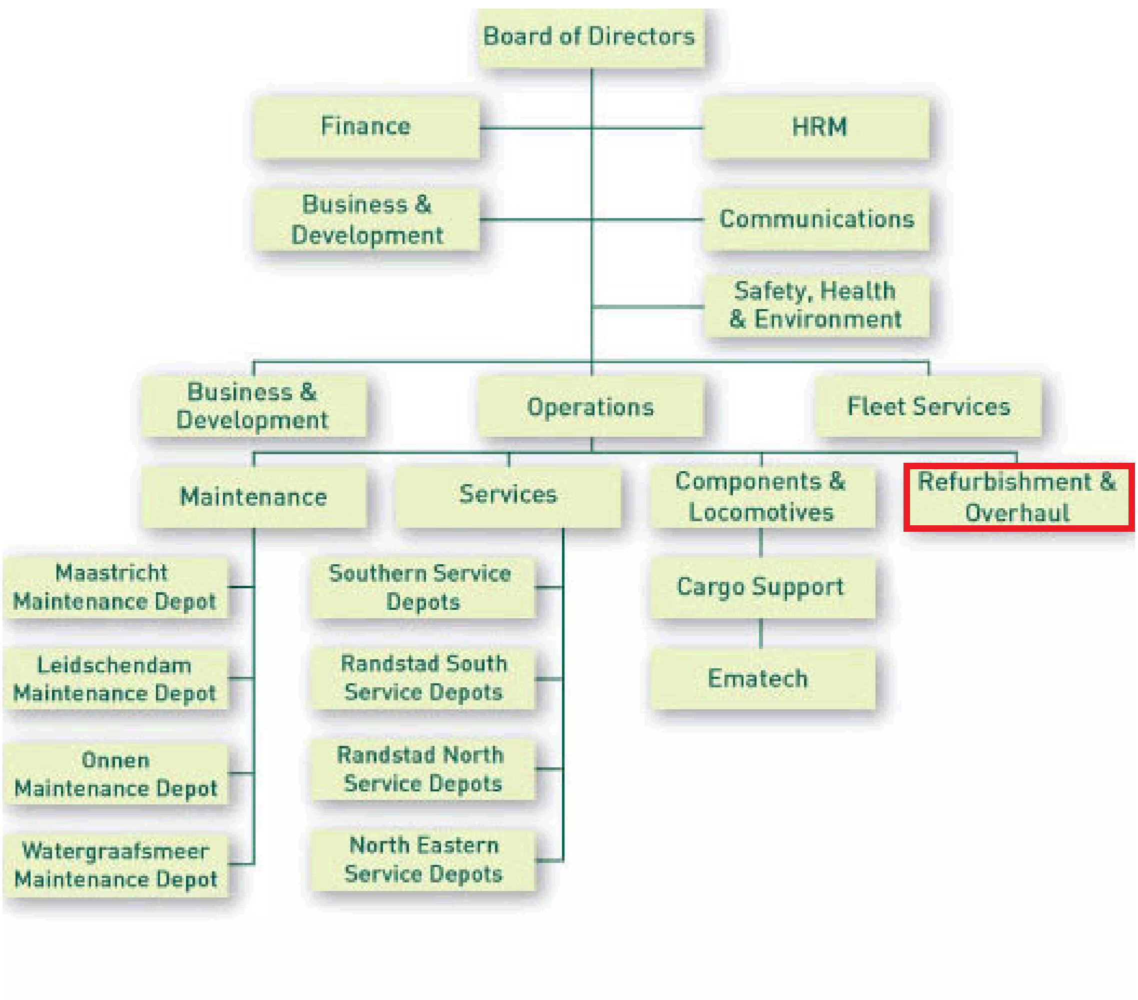

Figure 6: Organigram of NedTrain

2.1.1 NedTrain Haarlem

From now on, we will focus on NedTrain Haarlem (NTH). NedTrain Haarlem belongs to the Refurbishment & Overhaul department, see Figure 6. There are currently working over 700 FTE, making it NedTrains's largest repair shop. Its main activities are:

1. Overhauling, converting, and upgrading of trains/trams/metro cars
2. Overhauling of bogies (Dutch: *Draaistellen*), and wheel sets (Dutch: *Wielstellen*)
3. Repairing of collision, and fire damage
4. Manufacturing of parts, i.e. the *maakdelen*

For example, NTH is responsible for the overhaul of approximately 3,000 wheel sets and 1,100 bogies annually.

2.1.2 Detailed description of the data and the underlying demand processes

NedTrain Haarlem has provided us with data of 4,677 parts with their corresponding demand, and supply orders. The data consists of all demands from 04/07/2005 to 01/03/2011. NedTrain Haarlem currently uses Baan IV as its ERP system to store all the data, and it also uses Gordian tooling to perform calculations that Baan IV is not able to do. More detailed information of the data can be found in Appendix B.1.

Because most of the '*maakdelen*' have a very low price, we will focus only on the '*inkoopdelen*'. The planned lead time at NedTrain Haarlem consists of the contractual lead time, and safety time. The safety time is added to compensate for the uncertainty of the supplier, and for the internal administration. We will use the planned lead time in our research.

Our data consists of the activities of the overhauling of bogies and wheel sets. Within this group, there are in general two reasons to cause demand for a part:

1. **Projects**: When a project is taking place, many main parts of the same type come in for planned overhaul within a period of two to three years. The arrival of the main parts is always planned a few months in advance. When a main part comes in for overhaul some parts are always needed, whereas other parts are needed only occasionally. Whether the latter parts are actually needed for a certain main part is only known at the moment the main part has arrived and has been inspected. In order to predict this demand, NedTrain Haarlem has made an estimate of the probability that a part will be needed. This estimate is based upon a few try-out tests, and the engineers opinion. During the overhaul, this estimate may change. As the number of historical observations is limited, this probability is therefore not always estimated very accurately.

2. **Corrective maintenance**: When a main part comes in for corrective maintenance at a maintenance shop and the expected time to repair the part is longer than two hours, or the part cannot be repaired by the maintenance shop, it will be sent towards an overhaul shop. It is only known at the moment the main part arrives, what parts are needed for the maintenance.

When a part is not available, this can result in two different scenarios.

1. The most common scenario is that the employees of NedTrain Haarlem can no longer continue with their process, and the complete production line is stopped, which causes a significant delay in the maintenance, and requires rescheduling of the process.

2. If it is not possible to reschedule the process, the engineers sometimes try to find a substitute for the part. As a result, some demands have not been registered, because the engineers used another part instead. The alternative part is registered, although there was no actual request for that part. As it is not possible to retrieve this information, we will not correct the data for this.

As both scenarios have a negative impact on the efficiency and performance, it is important to have enough stock on hand.

2.1.3 Analysing the data

In this section, we will analyse the data. We will look at:

- The price of each part
- The number of requests each part has had over the last five years
- The size of each demand that has occurred
- The planned lead time of all parts
- The relation between the planned lead time and the price of a part
- The difference between the planned lead time and the realised lead time for each order

In Figure 7 we have made a histogram of the price of each part. The filled bar contains the total number of parts that have a price that is in the range given on the x-axis. As the size of the range varies, we also added the relative size, which is the total number of parts divided by the size of the range. From Figure 7, we see that there are more cheaper parts compared to expensive parts. We also see that there is a large number of parts that have a price of zero. These parts have not been ordered yet or have not been ordered for a long time. Therefore, the price is not recorded in the data. We will leave all parts without a price out of our research.

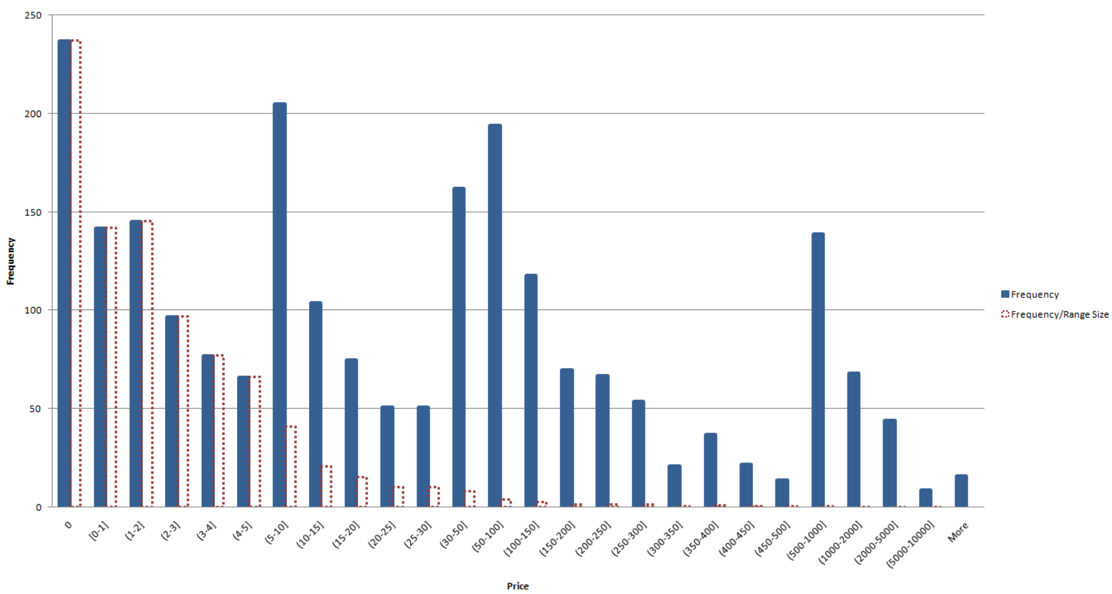

Figure 7: Prices of the parts in Euro at NedTrain Haarlem

We are also interested in the number of times a demand has occurred for each part over the last five years. When many parts are requested often, it will probably be easier to forecast the demand. Figure 8 presents a histogram of the number of times each part has been requested over the last five years. As NedTrain performs the maintenance of many different types of trains which are often used for many years, we expect there would be a regular and/or recurrent pattern in the demand. We do notice that relatively many parts have had over 300 requests over the last 5 years. These parts are probably easy to forecast as they are used very often. However, we see that there are more parts that have had no requests during the last five years, and there are also many parts that have had less than ten requests over the last five years. These parts are installed and used, but rarely needs to be replaced.

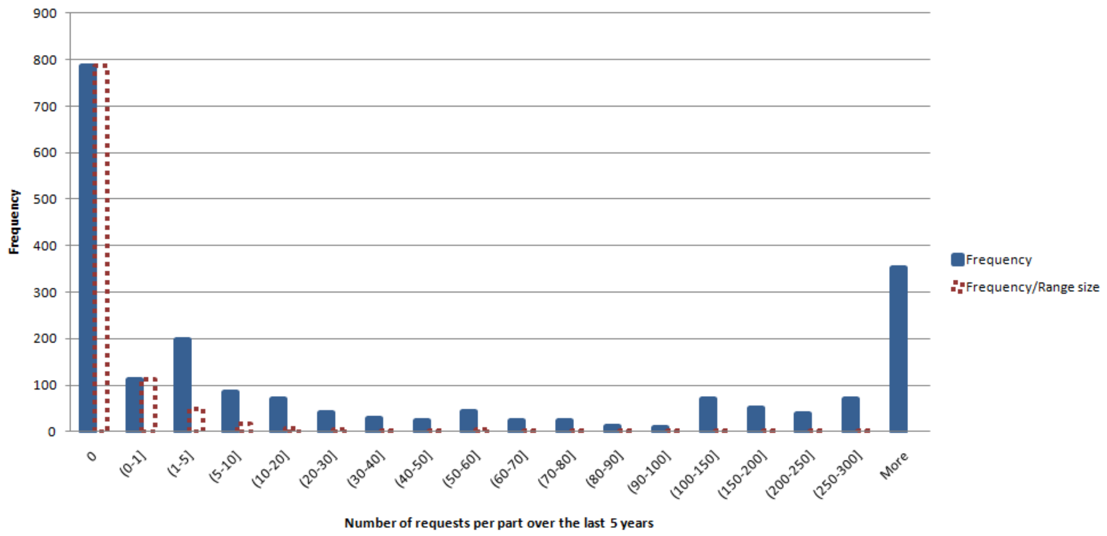

Figure 8: Total number of requests per part over the last five years at NedTrain Haarlem

We are also interested in the size of the demands. It could be that when a part is broken, all parts of the same type will be replaced during maintenance. For instance, NTH always replaces two wheel tires at once. This will avoid the same main part to return within a short time period because the other wheel tire needs maintenance. Considering that two wheel sets have to be placed inside of one bogie, and each wheel set consists of two wheel tires, we expect that demands will often be a multiple of two or four.

In Figure 9 we see a histogram of the size of the demands over the last five years. We see that the demand is most likely to be a multiple of four. When a part always needs to be replaced in groups of two or four, it is not useful to put three or five units on stock. It can be a good idea to take this information into account, when putting these parts on stock. However, even though many times a part is requested in a multiple of four, sometimes also different demands occur for that part. Therefore, we will not take this into consideration. Besides positive demand, there is also negative demand present in the data. Negative demand in the data represents the return of parts. These parts have not been used, although they initially thought it was needed. When we face negative demand during the simulation, we will put the amount on stock, as we consider the part being returned.

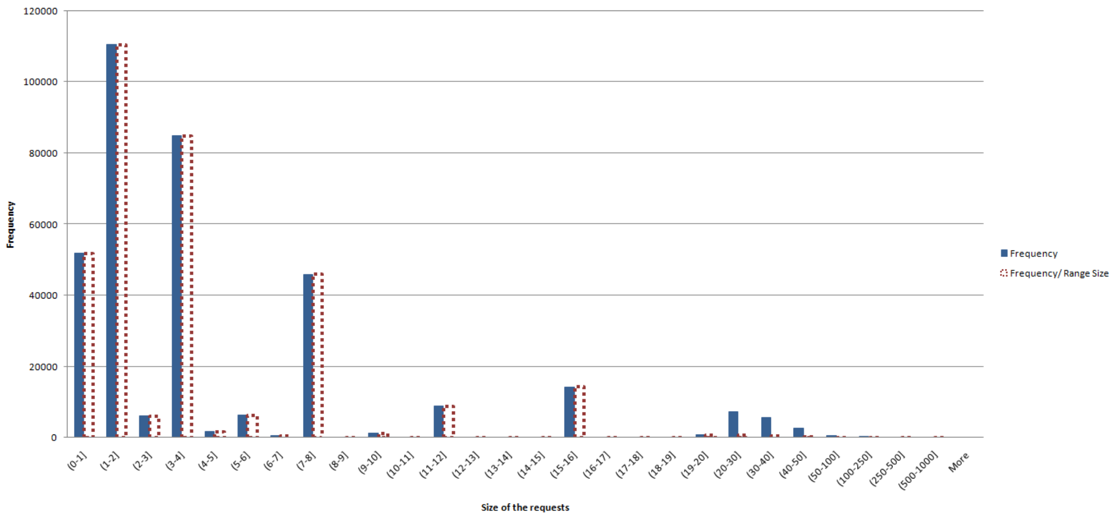

Figure 9: Demand sizes that have occured at NedTrain Haarlem

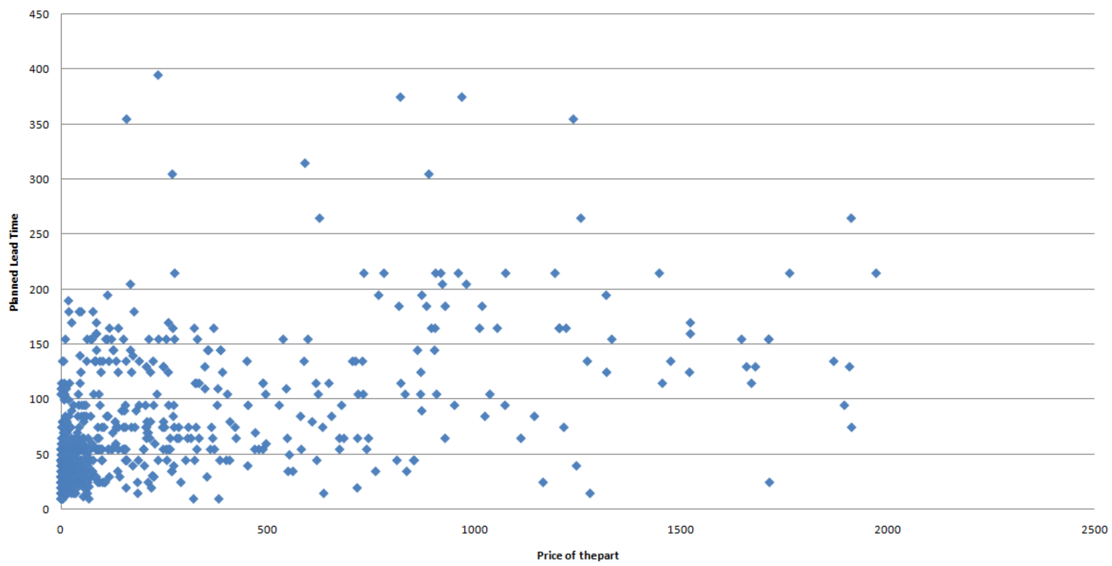

Figure 10: Relation between the price and planned lead time at NedTrain Haarlem

When we look at the planned lead time, we would expect that it is less than a couple of months for many parts, and only some parts, in general the more expensive parts, will have a lead time that is larger. However, when we look at Figure 10, we see that there is no direct relation between the price and the lead time.

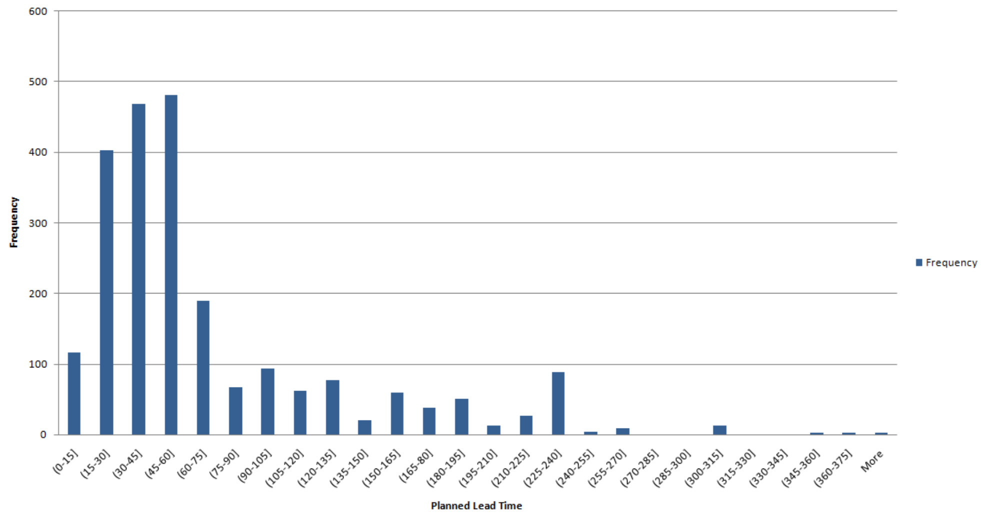

Figure 11: Planned lead time in working days at NedTrain Haarlem

Beside this scatter plot, we have also made an overview of the distribution of the lead time for all parts in Figure 11. We see that a large number of parts have a lead time less than three months. We also see a peak at the range (225-240]. When a part has not been ordered yet, or has not been ordered for a long time period, the contractual lead time is manually set to 222. If we then add safety time, this results in planned lead times between the range (225-240]. As these lead times are not based on the contractual lead time, we will leave all parts with a contractual lead time of 222 out of our research.

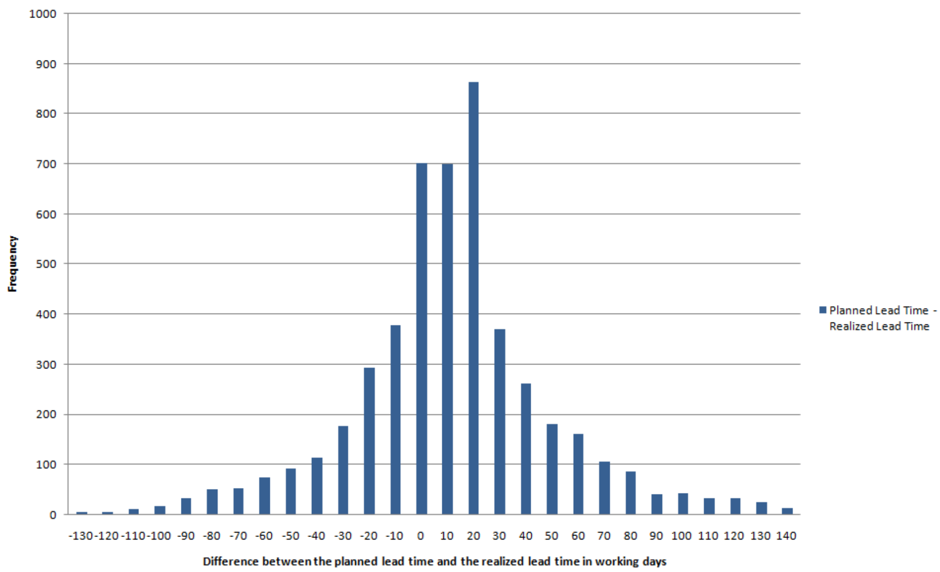

Figure 12: Difference between the planned lead times and the realised lead times for NedTrain Haarlem

Besides the average lead time, we are interested in the difference between the planned lead time and the realised lead time. If the realised lead time is shorter than the planned lead time, this will result in having too much stock on hand. If the realised lead time is longer, this will result in a lower service level. We expect that the planned lead time on average does not differ a lot from the realised lead time, but that there will be some variation. Figure 12 shows a histogram of the difference between the planned lead times and the realised lead times. As the planned lead time is given in working days, and the realised lead time is in calendar days, we adjusted this by dividing the planned lead time by 1.4 in order to convert working days to calendar days. A negative value in this Figure represents an order of which the realised lead time was longer than the planned lead time. We see that the average is close to zero, but that there is some variation.

An overview of the statistics after filtering the data is given in Table 1. We see that at NedTrain Haarlem about one out of five orders have a realised lead time that differs more than a month from the planned lead time, either too early or too late. Note that the planned lead time in Table 1 is given in working days, and the difference between the planned and realised lead time is given in calendar days.

	Price	Demand size	Lead Time in working days	Difference between planned and realised lead time in calendar days
Minimum	0.01	−2.00	10.00	−763.00
Maximum	54,586.24	1,440.00	375.00	482.00
Average	433.01	19.27	64.90	5.42
Standard deviation	3,091.82	50.78	62.10	55.23
10th Percentile	1.63	2.00	25.00	−42.00
90th Percentile	622.70	39.00	135.00	56.00
Number of values	790.00	348,313.00	790.00	5,133.00

Table 1: Overview of general statistics of the data at NTH

2.2 Alstom Transport

In this section, we will first give some general information about Alstom Transport. We then focus on Alstom Ridderkerk, which is one of its plants. Next, we discuss the data that has been provided. Alstom Transport has a diverse range of systems, equipments, and services including:

1. Rolling stock, for instance Metros and High-speed trains
2. Rail infrastructures, i.e. track laying and electrification
3. Signalling, for instance control systems and passenger information
4. Services, i.e. Maintenance, Part supply and Modernisation
5. Turnkey systems, i.e. a combination of the services above

Alstom Transport has over 60 locations worldwide, with a total of approximately 27,000 employees. A distribution of the employees over the world is given in Figure 13. From this Figure we see that most activities take place in Europe, of which France has the largest contribution. Alstom Transport has about 210 customers, together responsible for a yearly revenue of 5.5 billion € in 2009.

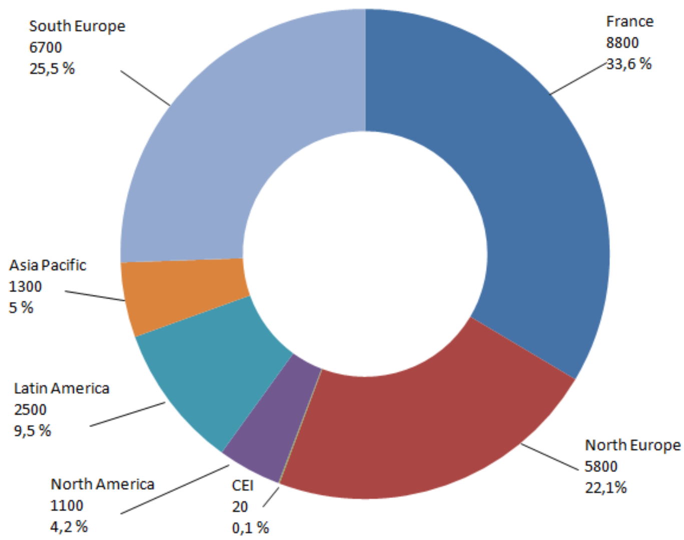

Figure 13: Geographical distribution of the employees

2.2.1 Alstom Ridderkerk

In this section, we will focus on Alstom Ridderkerk. Alstom Ridderkerk is responsible for:

1. The design and manufacturing of electric propulsion and electronic control systems for trains, metros, and trams
2. The modernisation of traction systems
3. Warehousing spare parts & Repairs
4. Commissioning & Service

The total sales volume for Parts & Repairs at Alstom Ridderkerk was about 13 million € in 2009/2010. The equipment of Alstom Ridderkerk has been installed in many countries. For example: Mexico (25 trains) and Switzerland (42 trains). The installed base consist of equipment build in the early sixties until recently.

2.2.2 Detailed description of the data and the underlying demand processes

The data we have obtained from Alstom Ridderkerk consists of 3,957 parts. Alstom Ridderkerk previously used Triton as its ERP system, but recently switched over to SAP. Because of this switch, there was some inconsistency in the data due to the migration, as we will see later on in this section. More detailed information of the data is found in Appendix B.2. The planned lead time consists of the contractual lead time, added with extra time for the quotation, for the ordering process, time to form goods receipt, and for the quality assurance. This planned lead time is used our research. Within the demands, we distinguish two different SAP types, the '261' and the '601' orders. When a part is used in a production process to repair or manufacture a main part it will be booked as a '261' movement. The data consists of all demands from 20/04/2006 to 20/04/2011.

A part booked as a '601' movement has been delivered to a customer based upon a sales order. In this case, Alstom Ridderkerk sends the part(s) to the customer when they have enough stock. A schematic overview of this process is given in Figure 14.

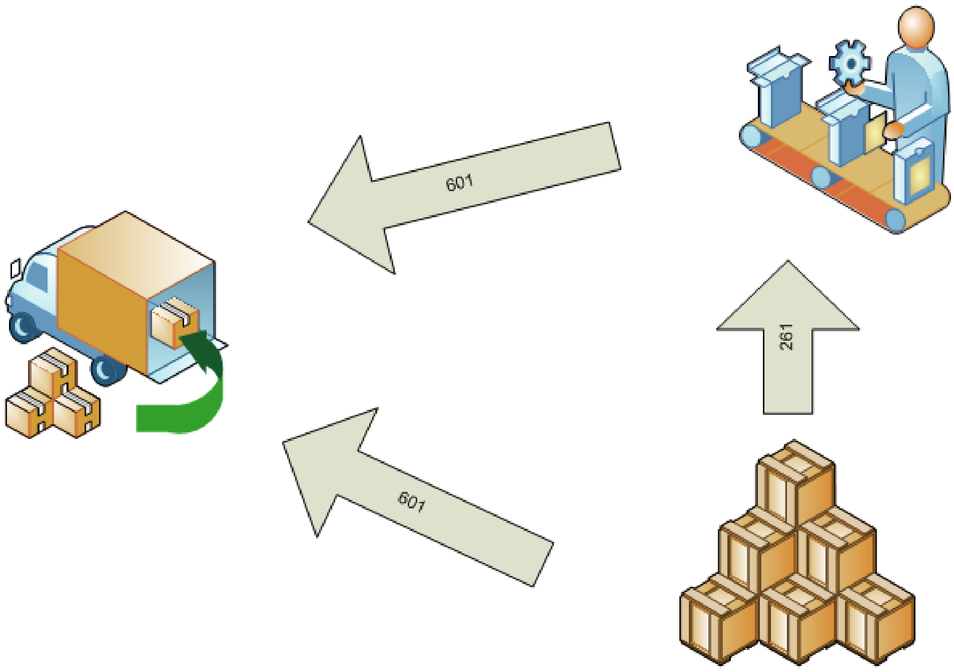

Figure 14: Different demand orders at Alstom

When a main part comes in for repair or maintenance, Alstom Ridderkerk first books an Repair Report (Dutch: *Reparatie Rapport*). When Alstom Ridderkerk is able to do the repair, it will perform the repair themselves. The spare parts that are used will be booked with 261 movement. When Alstom Ridderkerk is not able repair the part, it sends the broken part to the manufacturer. When the part has been repaired by Alstom Ridderkerk, or by the manufacturer, Alstom Ridderkerk will again inspect the part or perform some additional (integration) test. When the part is repaired properly, it will be sent back to the customer. An overview of this process is given in Figure 15. Because it is not known in advance what items are needed, this demand in general is unplanned.

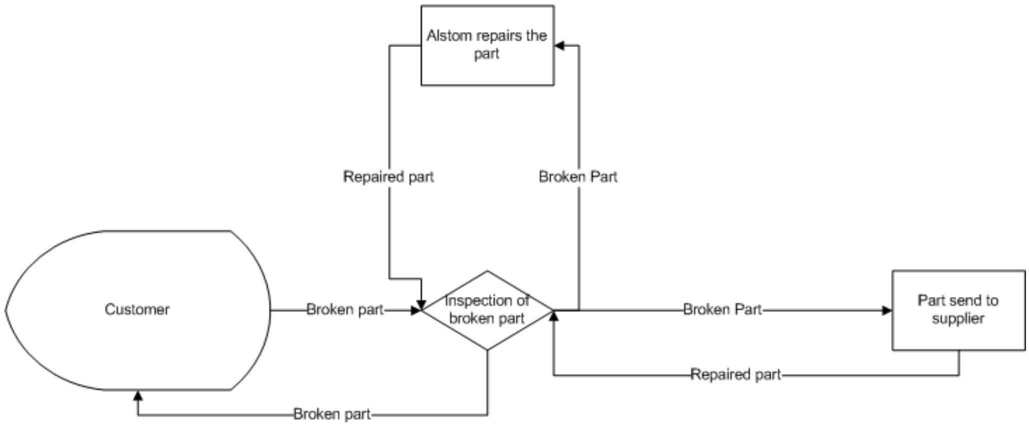

Figure 15: Flowchart for a repairable main part

Alstom Ridderkerk also performs modernisation projects, which are planned long before the actual modernisation takes place. Alstom then knows in advance what items are needed for the modernisation. In this case, we consider that the demand is planned. Based upon the price, leadtime, and contractual requirements, the MRP parameters are set in a way that the project will not be disturbed by quality issues due to the supplier or defects during the production or commissioning process. Sometimes, the modernisation project is scheduled backward, or forward in time. But this will only happen if the parts that are required are available, and will always be done in mutual agreement. As we do not need forecasting methods to predict this demand, we will leave these demands out of our research.

2.2.3 Analysing the data

Now that we have some insight into the company and the data, we will analyse the data in more detail. We will perform the same data analysis as we have done for NedTrain Haarlem. We will begin by looking at the prices of the parts. Figure 16 gives an overview of the prices of the parts at Alstom Ridderkerk.

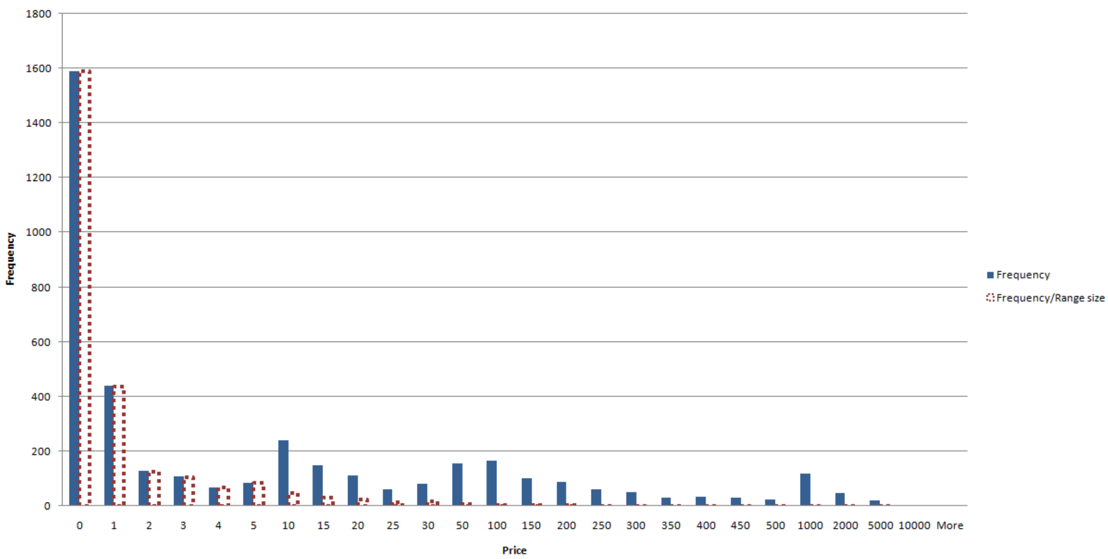

Figure 16: Prices of the parts in Euro at Alstom

The distribution of the prices is very similar to that of NedTrain Haarlem. Also at Alstom Ridderkerk, we notice that there are many parts that have a recorded price of zero in the database. At Alstom Ridderkerk there are two explanations for having prices of zero:

1. The price of the part is not known

2. The price of the part is lost due to the migration of the ERP system

We leave these parts out of our research.

We are further interested in the number of times a demand has occurred for each part, over the last five years. At NedTrain Haarlem, there is a relatively large group of parts with frequent demands. These parts have often been used in a project. We expect that Alstom Ridderkerk has less parts with this pattern, as we do not have demands for projects in our data.

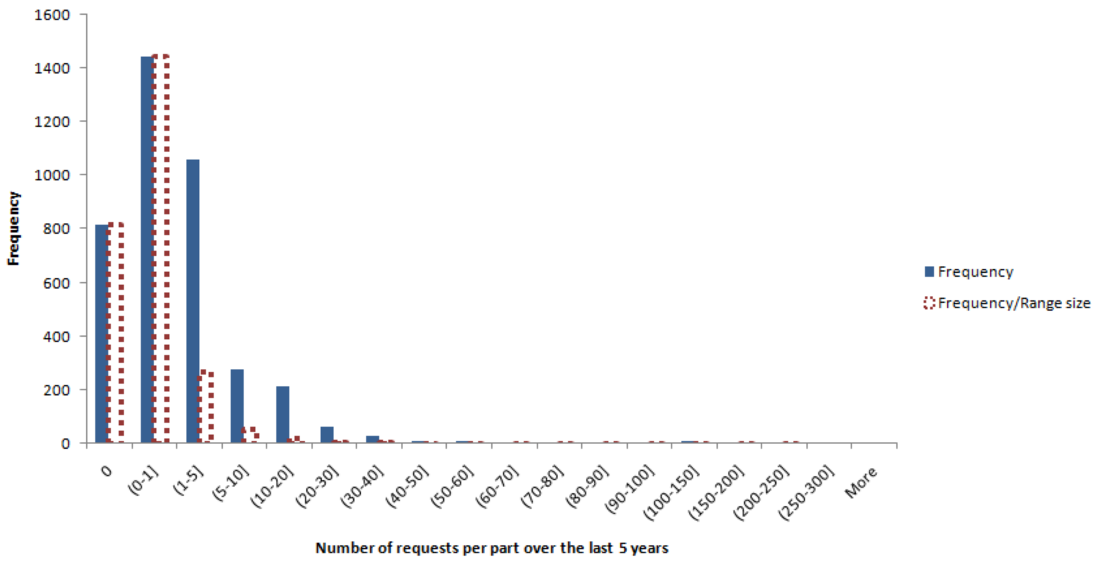

Figure 17: Total number of requests per part over the complete period at Alstom Ridderkerk

Figure 17 shows the number of demands for each part at Alstom Ridderkerk. From this Figure, we see that Alstom Ridderkerk has less parts with many demands compared to NedTrain Haarlem.

When we look at the size of the demands, we expect that most demands will be small as it is not common that parts are requested in large numbers very often. When we look at Figure 18, we see that even demands occurrence more often than uneven demands. There is no clear explanation why this pattern occurs, and it is not known that certain parts are always required in a fixed quantity.

When we look at the planned lead time, we expect that most lead times will be less than a few months. We have made a histogram of the planned lead time for all parts in Figure 19. We see that many parts have a planned lead time less than a few months.

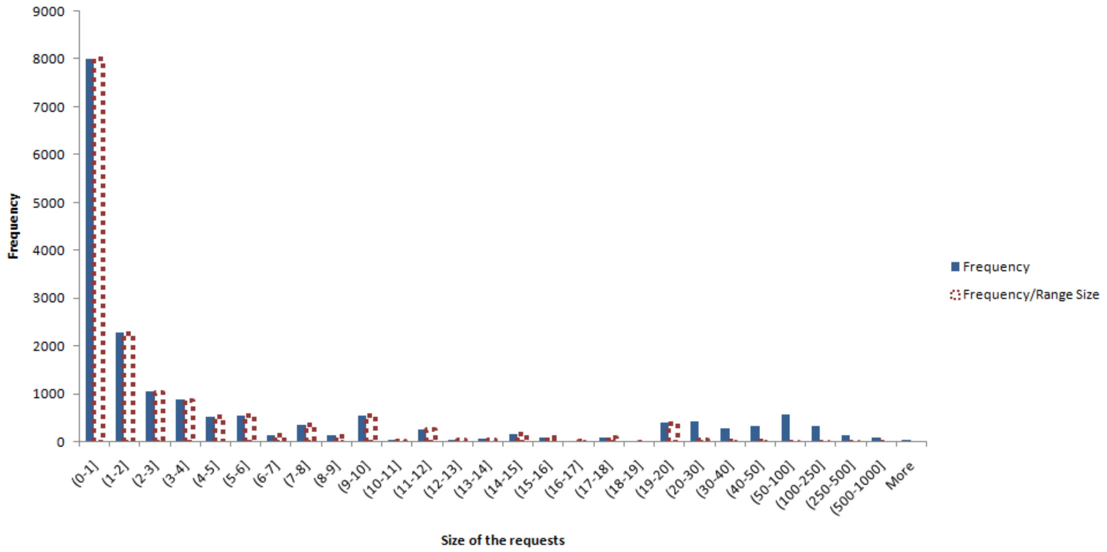

Figure 18: Demand sizes that have occurred at Alstom Ridderkerk

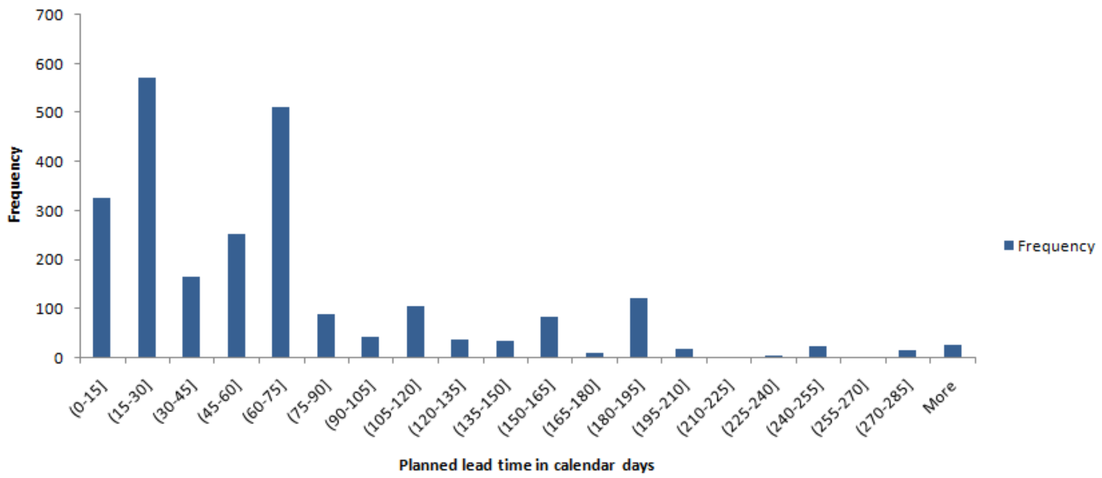

Figure 19: Planned lead times in calendar days for Alstom

We expected that the more expensive parts will have longer lead times compared to the less expensive parts, as these parts often are more specific, and/or takes more time to produce. However, this was not the case at NedTrain Haarlem. When we do the same analysis for Alstom Ridderkerk, we see from Figure 20 that this is also not the case at Alstom Ridderkerk.

We conclude that there is no clear relation between the price of a part, and the planned lead time. What cannot be seen from the data is that many parts are produced on order, specifically for the railbusiness for which in some cases the raw material is scarce.

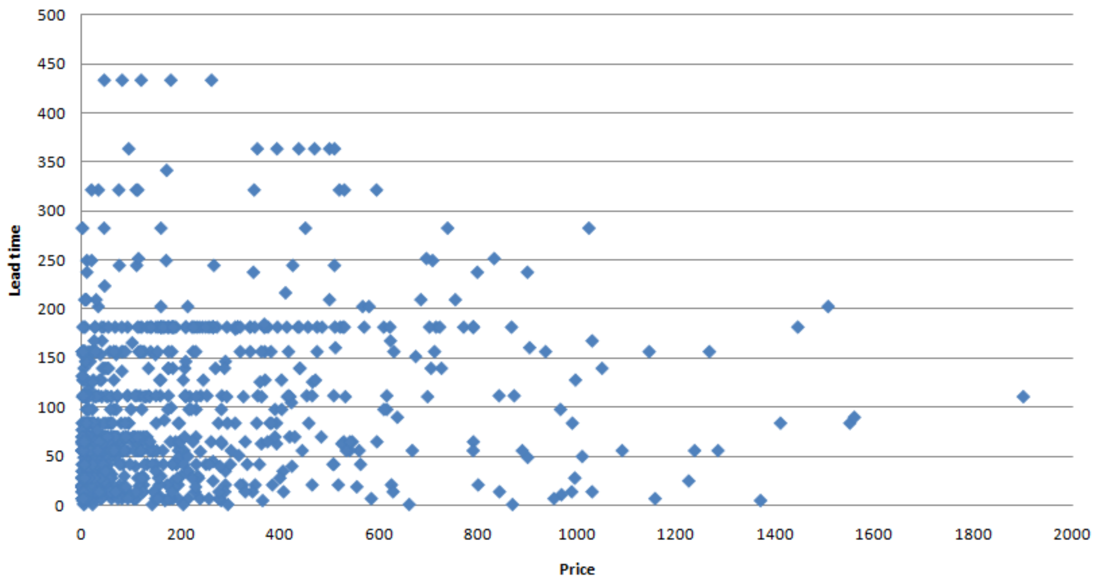

Figure 20: Relation between the price and lead time at Alstom Ridderkerk

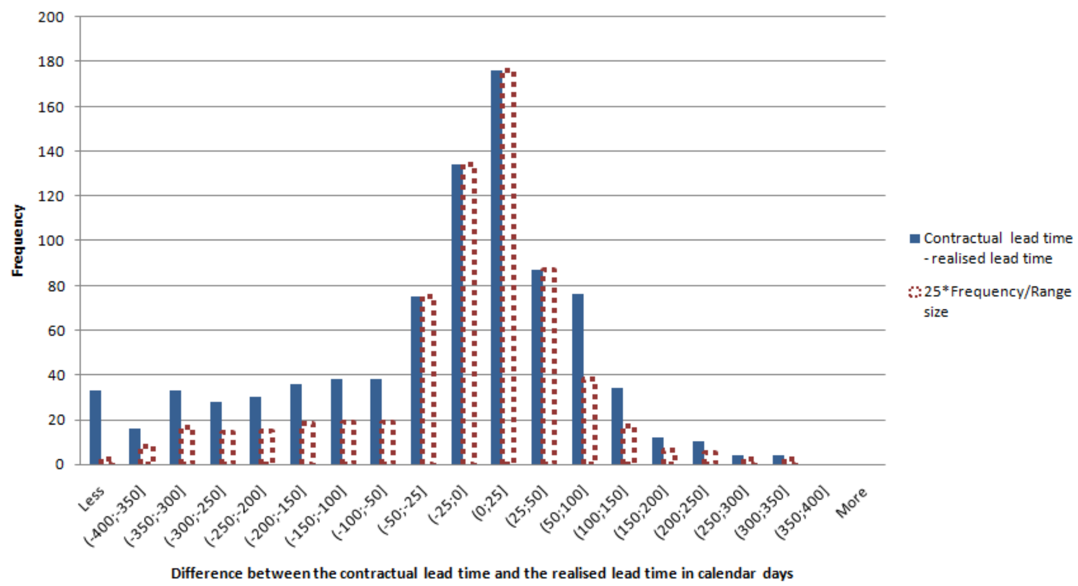

Figure 21: Difference between the planned lead time and the realised lead times in calendar days

We are further interested in the difference between the planned and realised lead time. We expect to obtain more or less the same results as we have found for NedTrain Haarlem. When we look at Figure 21, we see that an arrival of an order, which is on time or up to 25 days early, has the largest probability to occur.

However, unlike NedTrain Haarlem, the tails are much larger. The right tail at Alstom Ridderkerk is smaller compared to the left tail. This means that Alstom Ridderkerk has more orders that arrive too late than too early.

Table 2 gives an overview of the characteristics we have looked at after filtering the data. We see that Alstom Ridderkerk has very large demand sizes in the data. Even a demand of size 5,500 has occurred. Looking at these demand sizes, we would think these demands are outliers. We cannot say for sure that this is the case here. The part for which this demand has occurred had more requests of similar sizes. Most likely this demand is in grams, which is usual for copper braid. Therefore, we have left these demands in our data. We also see that there is much more variation in the lead time compared to NedTrain Haarlem. Because of this we expect that using a method using a stochastic lead time will be more useful at Alstom Ridderkerk than at NedTrain Haarlem.

	Price in €	Demand size	Lead Time in calendar days	Difference between planned and realised lead time in calendar days
Minimum	0.01	0.03	1.00	−536.00
Maximum	1,900.00	5,500.00	434.00	333.00
Average	83.90	35.47	73.82	−57.54
Standard deviation	189.24	186.54	66.23	148.84
10th Percentile	0.43	1.00	14.00	−307.40
90th Percentile	240.13	52.00	162.50	70.00
Number of Values	748.00	7,610.00	748.00	594.00

Table 2: Overview of general statistics of the data at Alstom Ridderkerk

2.3 Defence Material Organisation

In this section, we give some general information of the Defence Material Organisation (DMO). We then go more into detail about the Naval Maintenance Company (NMC) (Dutch: *Marinebedrijf*), a service centre responsible for all naval material within the Defence Organisation. Next, we follow the same procedure as we have done for the other two companies. The Defence Material Organisation is a relatively young division of the Defence Organisation, with its headquarters located in The Hague, and with other logistics units, and maintenance operations dispersed across the country. The Defence Material Organisation is staffed by civilian, and military personnel with backgrounds in the navy, army, air force and military police (Dutch: *Marechaussee*). There are about 6,000 employees working for the Defence Material Organisation, of which approximately 80% is civilian and 20% military personnel.

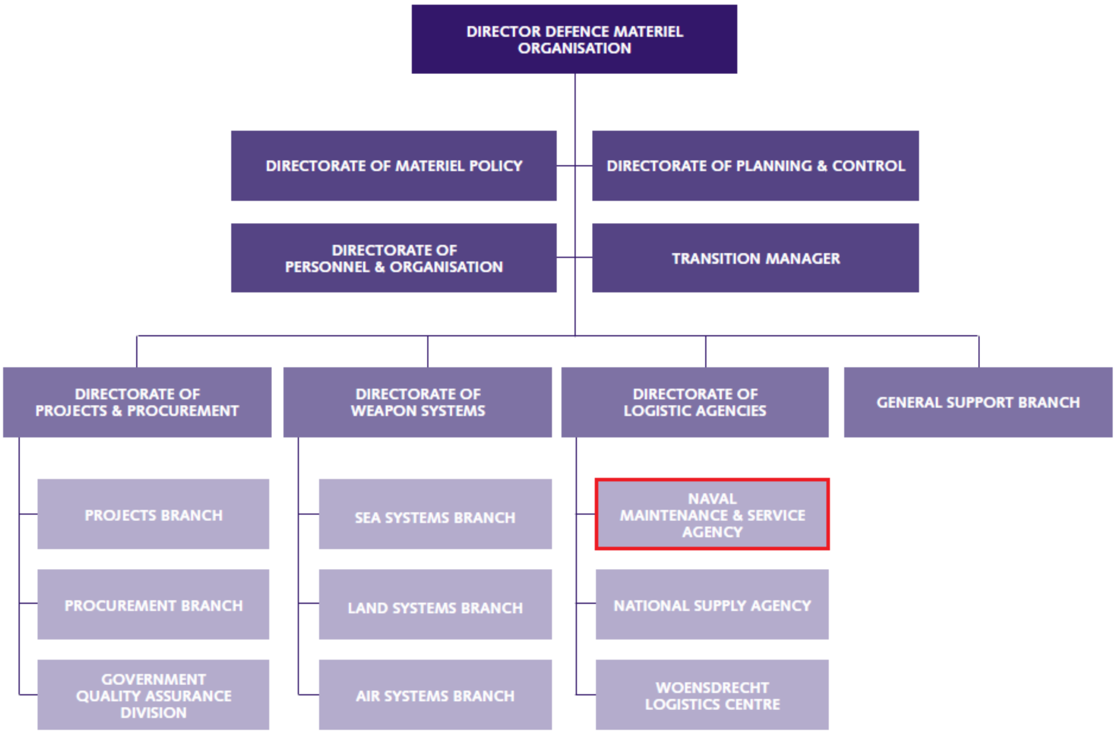

Figure 22: Organigram of the Defence Material Organisation

2.3.1 Naval Maintenance Company

In this section, we focus on the Naval Maintenance Company (NMC), which is part of the Defence Material Organisation. An organigram of the Defence Material Organisation is given in Figure 22. The NMC is the sustainment agency, with the emphasis on maritime material from the Navy Command. It carries out work for the Army Command, and the Air Force Command. The NMC can be divided into three divisions:

1. Fleet
2. Special Products
3. Logistic Services

The Fleet division manages the maintenance and supply of the ship hulls, the propulsion system, and the electronic and weapon systems on board of the ships. It also provides technical advice and support, to ships anywhere in the world. It advises on, and participates in large-scale material projects (new-build). The Special Products division has a broad range of activities in a variety of product groups, for a diverse user group. For example, the Special Products division works on Commando Corps night-vision goggles, it calibrates instruments for the F-16, and maintains guided weapons. The Logistic Services division is responsible for the supply of operational goods and spare parts (procurement, administration, storage, and distribution) for both operational users, and the maintenance process.

In short, the Naval Maintenance Company maintains, repairs, supplies, advises and modifies.

2.3.2 Description of the data

The data we obtained from the NMC consists of 258,018 parts. Of each part we also have the corresponding demands of the last ten years. For the NMC, we look at all the demands that occurred from 01/07/2005 to 29/4/2011 such that we obtain a more or less equal time from of NedTrain Haarlem. More detailed information of the data is found in appendix B.3.

2.3.3 Analysing the data

Now that we have a general picture of the NMC, we use the same approach as we have done before by looking at the following characteristics:

- The prices of the parts
- The number of requests each part has had from 01/07/2005 to 29/04/2011
- The size of each demand that has occurred
- The planned lead time of all parts
- The difference between the planned lead time and the realised lead time for each order

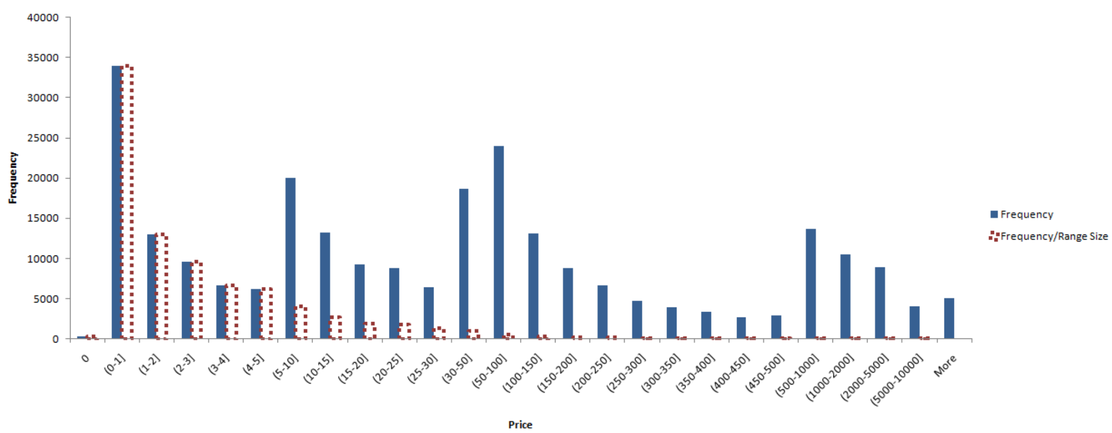

Figure 23: Price distribution of the parts in Euros at the Naval Maintenance Company

When we look at the price of parts, we expect similar results as we have obtained from NedTrain Haarlem, and Alstom Ridderkerk. In Figure 23, we see that this indeed seems to be the case. We see a decline in the number of articles when the price increases, and also see that some parts have a price of €0. Although there are not many parts with a price of €0 we still like to know the cause. After asking the data analyst at the Naval Maintenance Company, we obtained that all parts that costs less than €0.005 are rounded down by the system to a price of €0. We leave these parts out of our research.

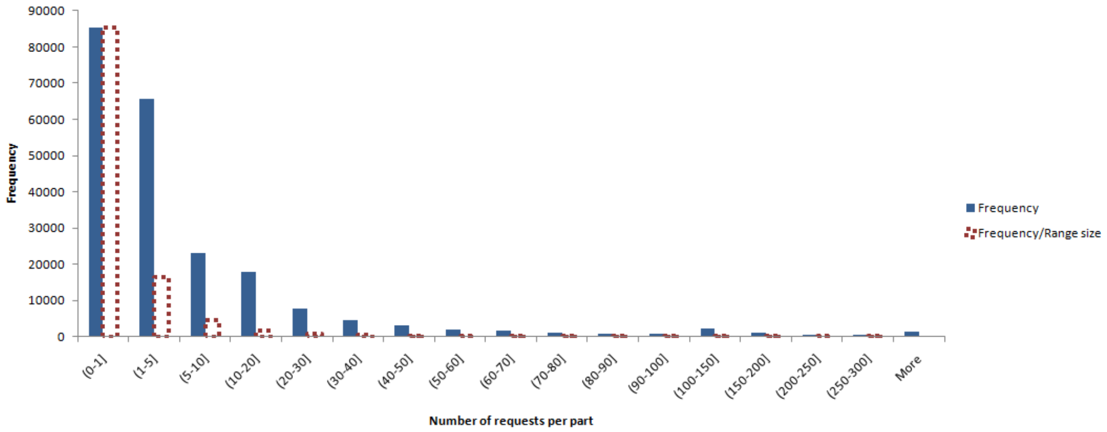

Figure 24: Total number of requests per part at the Naval Maintenance Company

We are also interested in the demand frequency for each part over the last five years. Parts that are requested on a more regular basis are probably easier to forecast. Figure 24 represents the number of times a part has been requested in the last five years. We see that the parts follow the same pattern as we have found for the other two companies.

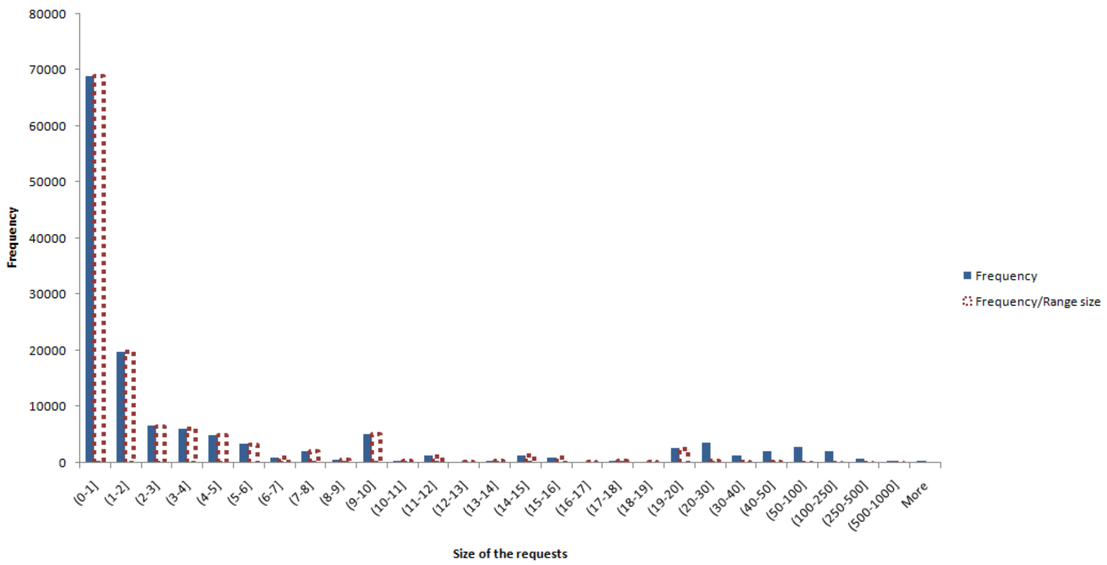

Figure 25: Size of requests at the Naval Maintenance Company

Next, we look at the size of each demand. We expect that most demands are small, as it is not common that parts are frequently used in large numbers. When we look at Figure 25, we see that most requests are of size 10 or smaller. We also see that the number of request at the range (9,10], and the range (19,20] occur more often than expected. It seems that there is a typical demand quantity. However, after asking the data analyst, we can conclude that this is often caused by either a minimum package size, or because the demand is used as a replenishment for the floor stock (Dutch: *grijpvoorraad*). As parts sometimes have minimum package sizes, it is reasonable to assume that these parts should also be ordered with a typical order quantity. However, as the minimum order quantity (MOQ), and the modular order quantity (MOD), are not available, we cannot take this into consideration.

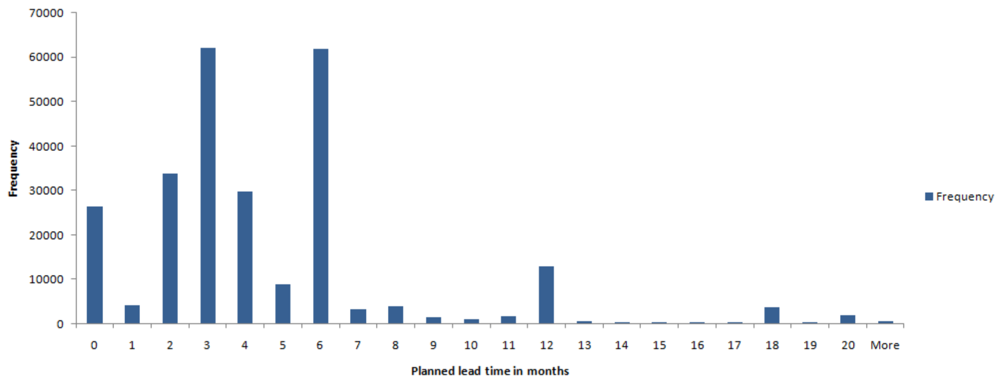

Figure 26: Planned lead time at the Naval Maintenance Company

When we look at the lead times, we expect similar results as we have found for NedTrain Haarlem, and Alstom Ridderkerk. In Figure 26 we see that most parts have a lead time smaller than six months. Compared to NedTrain Haarlem, and Alstom Ridderkerk, the Naval Maintenance Company has relatively long lead times for their parts. We also see that there are peaks at twelve and eighteen months. An explanation could be that manufacturers round up their lead times to a common value, such as a half year. For instance, parts with a lead time of eleven months could be rounded up to one year by the supplier. Therefore, we expect that the realised lead time is often shorter than the planned lead time. We have made an histogram to represent this difference in Figure 27. The difference is noted in calendar days, and when the difference is negative, this represents that an order has arrived later than planned. We see that most orders have a realised lead time that is shorter than the planned lead time, as we expected.

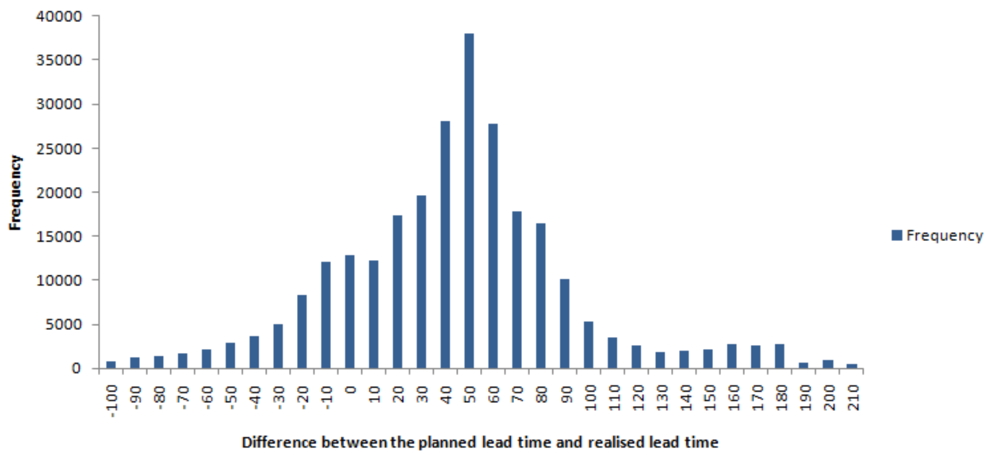

Figure 27: Difference between the realised lead times and planned lead time in calendar days at the Naval Maintenance Company

Table 3 gives an overview of the statistics at the NMC after filtering the data. We see that the most expensive part of the NMC has a value of 2.7 million €. This value is far more expensive than the most expensive part of NedTrain Haarlem. Although this price is very large, we see that only ten percent of the parts exceed the price of € 6,498.00. Compared to the most expensive part, these parts have only a small influence on the total inventory costs. We also notice that a demand of 1,440 has occurred. These demands can represent an order for the floor stock. As these parts have more requests of similar sizes, we have kept these demands in our data. We also see that some parts have lead times longer than three years. These parts are probably very specific, which results in producing on demand, and a relatively long production time. These parts are more difficult to control as the variation of the demand during the lead time increases when the lead time becomes longer. Besides the long lead time, these parts also often have intermittent demand patterns. We also see that on average, parts arrive more than a month too early, and only one out of ten orders arrives more than 8 days too late.

	Price in €	Demand size	Lead Time in months	Difference between planned and realised lead time in calendar days
Minimum	0.01	0.01	0.01	−1265.00
Maximum	2,697,070.10	15,000.00	36.80	1,100.00
Average	4,723.65	19.25	4.34	47.24
Standard deviation	43,379.69	151.20	3.29	85.35
10th Percentile	2.28	1.00	2.00	−8.00
90th Percentile	6,498.18	24.00	8.00	122.00
Number of values	5,191.00	39,408.00	5,191.00	25,890.00

Table 3: Overview of general statistics of the data at the NMC

2.4 Conclusions

We have seen that at all three companies the distribution of the prices is quite similar. There are many relatively inexpensive parts, and only a few expensive parts. However, the range of the prices of the parts differs between the companies, but we expect that this will not influence the performance of the different forecasting methods.

From Tables 1, 2 and 3 we see that the average demand size per request is almost the same at NTH and the NMC, but much higher at Alstom. The demand size at NTH is often a multiple of 4 because these parts are used for bogies or wheelsets. This also means that the variation in demand size is much lower at NTH compared to the other two companies. Alstom Ridderkerk has the largest average demand size and variation in the demand size. Because this makes it more difficult to put the right amount of stock on hand, we expect that the achieved fill rate will be lower at Alstom compared to the other two companies.

At NTH we further noticed that there are relatively many parts with a high number of requests, compared to the other companies. We can thus conclude that Alstom Ridderkerk, and the NMC have to deal with a majority of parts which belong to the group of intermittent parts. As a result, we expect that we obtain better results at NTH. We also expect that the average achieved fill rate at NTH will be higher compared to the other companies, because of the MOQ and MOD.

The lead times at the Naval Maintenance Company are much longer than at the other two companies. As longer lead times make it more difficult to put the right amount of stock on hand, we expect that the achieved fill rate at the other companies will be higher.

We also looked at the difference between the realised lead time and planned lead time. At all three companies, the realised and the planned lead time differ quite often. At Alstom, orders arrive more often too late than too early, whereas this is the other way around at the NMC. At NTH, orders arriving too early and too late occur with approximately equal probability. Because of this, we expect a slightly higher order fill rate at the NMC, and a lower fill-rate at Alstom.

Chapter 3

3 Forecasting methods

In this chapter, we describe some traditional forecasting methods, which we can use to obtain the reorder point. Remark that with the traditional methods, we also include Croston's method and the Syntetos-Boylan approximation (SBA), although, for example, Syntetos, Boylan, and Croston (2005) do not include these methods when talking about traditional forecasting methods. An overview showing the steps that need to be taken in order to obtain the reorder point using the traditional methods is given in Figure 28.

In Section 3.1, we look at different forecasting methods to estimate an average of the demand based on historical demand values. Once we have calculated the average, we determine the variation of the demand in Section 3.2. In Section 3.3 we fit the average and variation to the normal distribution. Based on this distribution, we then are able to determine the reorder point, which is explained in Section 3.4.

There are many different forecasting methods, and there are also many other distributions which we could have chosen instead of the normal distribution. Parts that are used regularly are easier to fit to a certain distribution. However, when a part only moves a few times a year, it is difficult to fit the demand to a common distribution. We expect that because of this, this way of forecasting does not give good results for the latter type of parts. Therefore, we introduce the Empirical method in Chapter 4, which uses an empirical distribution and does not require the fitting to a distribution.

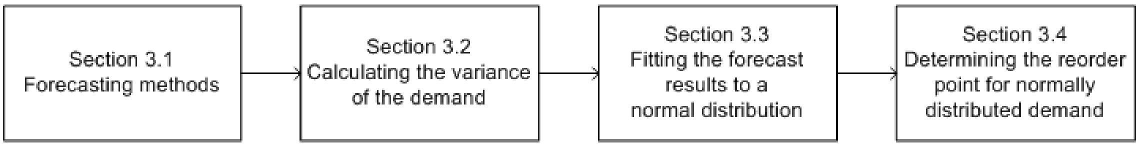

Figure 28: The traditional method to obtain the reorder point.

3.1 Forecasting methods

In order to be able to put the right amount of stock on hand, we should know the demand in the upcoming periods. However, as this is often not known in advance, people have been trying to predict future demand as good as possible (Syntetos, Boylan, & Disney, 2009).

In this section, we look at different forecasting methods that are most used in the literature. We first describe the constant and trend model, which form the underlying demand model(s) of the forecasting methods. We then look at Simple Moving Average (SMA), Single Exponential smoothing (SES), Double Exponential Smoothing (DES), Croston, and the Syntetos Boylan Approximation (SBA). We use these forecasting methods to obtain a forecast of the average demand per month.

3.1.1 Constant model

In the constant model, demands in different periods are represented by independent random deviations from an average that is assumed to be stable over time compared to the random deviations (Axsäter, 2004). This model is most appropriate for products that are in a mature stage of a product life cycle and are used regularly. If we do not expect a trend or seasonal pattern, it is in most

cases reasonable to assume a constant model.
Let us introduce the following notation:

x_t = observed demand in period t, $t \in \mathbb{N}_0 = \mathbb{N} \cup \{0\}$
a = average and expected demand per period
ϵ_t = independent random deviation with expectation zero

When assuming that the demand follows a constant model, the demand in period t, x_t, can be represented as:

$$x_t = a + \epsilon_t \tag{1}$$

3.1.2 Trend model

If the demand tends to increase or decrease systematically, we may not assume that the demand follows a constant model. It is better to extend the constant model by also considering a linear trend in the demand. Let

$a_0 =$ average demand in period 0
$b =$ trend, that is the systematic increase or decrease per period

The demand is then modelled as:

$$x_t = a_0 + bt + \epsilon_t \tag{2}$$

Note that the constant model is a special case of the trend model where $b = 0$

3.1.3 Simple Moving Average

When using the Simple moving average (SMA) we assume that the demand follows the constant demand model which is represented by Equation 1. To estimate the average demand, a, we take the average over the N most recently observed values. Let:

$N =$ the number of observed values of which the average is taken (we take $N = 12$)
$\hat{a}_t =$ estimate of a after observing the demand in period t
$\hat{x}_{t,\tau} =$ forecast for period $\tau > t$ after observing the demand in period t

We obtain:

$$\hat{x}_{t,\tau} = \hat{a}_t = \frac{1}{N} \cdot \sum_{n=t-N+1}^{t} x_n \tag{3}$$

where $t \geq N - 1$.
Note that the value of $\hat{x}_{t,\tau}$ is the same for any value $\tau > t$ because it is assumed that the demand follows the constant model. When we take a longer time period for N, we have a shorter time period for our simulation. Taking a shorter time period is also not beneficially, as we may miss information, especially for parts with seasonality.

3.1.4 Single Exponential Smoothing

Just as in the case of the SMA, it is assumed that the demand follows the constant model, represented by equation 1. We initialise the forecast x_0 by taking the average over the first $N = 12$ periods. To update the forecast in period $\tau > t$, we use a linear combination of the previous forecast and the most recent realised demand x_t. Let:

$\alpha =$ smoothing constant with a value of $(0 \leq \alpha \leq 1)$

$$\hat{x}_{t,\tau} = \hat{a}_t = (1 - \alpha)\hat{a}_{t-1} + \alpha x_t$$

If we choose an α of zero, this results in having a constant forecast, equal to the initial forecast. An α of one results in a forecast equal to the demand in the last period.

3.1.5 Double Exponential Smoothing

When we expect that the demand increases or decreases slowly over time, it is better to assume the trend model, represented by equation (2), instead of a constant model. Let:

$\alpha =$ smoothing constant with $(0 \leq \alpha \leq 1)$
$\beta =$ smoothing constant with $(0 \leq \beta \leq 1)$
$\hat{b}_t =$ estimate of b after observing the demand in period t
There are different methods to update the parameters $\hat{a}_t$ and $\hat{b}_t$ but we consider the method suggested by Holt (2004), which is often used in the literature. The double exponential smoothing (DES), is then as following:

$$\hat{a}_t = (1 - \alpha)(\hat{a}_{t-1} + \hat{b}_{t-1}) + \alpha x_t$$

$$\hat{b}_t = (1 - \beta)\hat{b}_{t-1} + \beta(\hat{a}_t - \hat{a}_{t-1})$$

The forecast for time period $t + k$ is then obtained by

$$\hat{x}_{t,t+k} = \hat{a}_t + k\hat{b}_t$$

3.1.6 Croston's method

When the time between two consecutive demands is large, the forecast using Single Exponentional Smoothing (SES) slowly goes to zero, even when the demand size is often large. Croston (1972) suggested an adjustment to SES to handle such situations. When using Croston, the forecast is only updated in periods with positive demand. When demand occurs, two averages are updated using exponential smoothing: the size of the positive demand, and the time between two periods with positive demand. Let:

k_t = the number of periods since the preceding positive demand
$\hat{k}_t$ = the estimated average of the number of periods between two positive demands at the end of period t
$\hat{d}_t$ = the estimated average of the size of a positive demand at the end of period t
$\hat{a}_t$ = estimated average demand per period at the end of period t.

We update $\hat{k}_t$ and $\hat{d}_t$ according to:

If $x_t = 0$

$$\hat{k}_t = \hat{k}_{t-1}$$
$$\hat{d}_t = \hat{d}_{t-1}$$

Otherwise ($x_t > 0$):

$$\hat{k}_t = (1-\alpha)\hat{k}_{t-1} + \alpha k_t$$
$$\hat{d}_t = (1-\beta)\hat{d}_{t-1} + \beta x_t$$

We then obtain the forecast for the average demand per period as

$$\hat{x}_{t,t+\tau} = \hat{a}_t = \hat{d}_t/\hat{k}_t \tag{4}$$

3.1.7 Syntetos Boylan Approximation

Croston's method described in Section 3.1.6 is shown to be biased by Syntetos and Boylan (2001), especially when the time between two consecutive demands (p) is large or when there is much variation in the demand. In order to deal with this bias, Syntetos and Boylan (2001) propose the Syntetos-Boylan Approximation (SBA). This method has been developed based on Croston's idea, and has been shown to outperform Croston's method on generated data. The new estimator is shown to be approximately unbiased and improves when p becomes larger (Syntetos & Boylan, 2005b). We use the same definitions as described in Section 3.1.6. The Syntetos Boylan Approximation uses the same estimation of the average, as described in Equation 4 but multiplied by a factor $(1-\alpha/2)$ in order to compensate for the bias. The estimation is then as follows:

$$\hat{a}_t = (1-\alpha/2)\hat{d}_t/\hat{k}_t$$

3.2 Obtaining the variation of the demand

So far we only looked at estimating the average demand in the future. To be able to model the demand, we also need to know the variation in the demand. We use the MSE as a measure for the variation, where we only look at the last $N = 12$ periods for the same reasons as described in Section 3.1.3. The MSE is then calculated as follows:

$$MSE_t = \frac{1}{N}\sum_{n=0}^{N-1}\left(x_{t-n} - \hat{x}_{t-n-1,t-n}\right)^2 \tag{5}$$

We obtain the standard deviation by taking the square root:

$$\hat{\sigma}_t = \sqrt{MSE_t}$$

3.3 Fitting the forecast results to the normal distribution

Once we have obtained our average and variation per month, we need to fit the parameters to a common distribution. Porras and Dekker (2008) have found that using a normal distribution performs better than the Poisson distribution. They mention that this is because the normal distribution is less sensitive to changes in the service level, and thus more conservative re-order points are advised when using the normal distribution compared to the Poisson distribution. Therefore, we use the normal distribution. However, we should note that the normal distribution might lead to bad predictions on the actual fill rate as they also mention that the actual achieved fill rate differs more from the desired fill rate compared to the Poisson distribution.

As the average and variation is based on a period of one month, and the lead time is often larger than one month, we first need to adjust the average and variation such that it corresponds to the demand in the lead time. After we have done this, we model the demand according to the normal distribution.

3.3.1 Determining the average demand and variation of the demand during lead time

The forecast results are estimations of average demand and variation for a period of one month. As the planned lead time (see Chapter 2) is not always equal to one, we need to adjust the forecast results, such that we obtain an estimation of the average demand and variation during the lead time. Let
μ' = forecasted average demand during lead time
σ' = forecasted variation of the demand during lead time
L = the planned lead time

We then obtain the μ' and σ' as follows (see Axsäter (2004, p. 33)):

$$\mu' = \hat{x}_{t-1,t} \cdot L \tag{6}$$

$$\sigma' = \hat{\sigma}_t \cdot \sqrt{L} \tag{7}$$

3.3.2 The normal distribution

The standardized normal distribution with mean $\mu' = 0$, and standard deviation $\sigma' = 1$ (see Axsäter (2004, p. 85)) has the following density and distribution function:

$$\varphi(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}, -\infty < x < \infty \tag{8}$$

$$\Phi(x) = \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}} du, \tag{9}$$

When we have different values for the mean and variation we obtain the density function by:

$$(1/\sigma')\varphi((x-\mu)/\sigma), -\infty < x < \infty$$

and the distribution function as

$$\Phi((x-\mu')/\sigma')$$

Once we have modeled the demand according to the normal distribution, we can determine the reorder point.

3.4 Determining the reorder point for normally distributed demand

Let us first introduce the following notation: $R =$ Reorder point
$Q =$ Order quantity

In order to calculate the fill rate, we need to calculate the probability that the Inventory Level (IL) is below 0. In other words, we would like to calculate the following:

$$P(IL \leq x) = F(x) = \frac{1}{Q}\int_{R}^{R+Q}\left[1 - \Phi\left(\frac{R - x - \mu'}{\sigma'}\right)\right] du$$

Let us now first introduce the loss function $G(x)$ as:

$$G(x) = \int_{x}^{\infty}(v - x)\varphi(v)dv = \varphi(x) - x(1 - \Phi(x))$$

Note that

$$G'(x) = \varphi(x) - 1$$

We then rewrite F(x) as:

$$F(x) = \frac{1}{Q}\int_{R}^{R+Q}\left[-G'\left(\frac{R - x - \mu'}{\sigma'}\right)\right] = \frac{\sigma'}{Q}\left[G\left(\frac{R - x - \mu'}{\sigma'}\right) - G\left(\frac{R + Q - x - \mu'}{\sigma'}\right)\right]$$

We then are able to calculate the *fill-rate* for a given Q, R, μ', and σ' by

$$S_2 = 1 - F(0) \tag{10}$$

see (Axsäter, 2004) for a more detailed explanation. Note that companies are often more interested in the order line fill-rate instead of the part fill-rate. However, determining the reorder point using an order line fill-rate target is theoretically more difficult. However, in our simulation we also keep track of the order line fill-rate, which we use to present our results in Chapter 6.

In order to determine R for a given service level target ($S2$), we use the following iterative approach:

1. Start by setting $R = 0$
2. Calculate the part fill-rate according to (10)
3. If the part fill-rate is high enough, stop
4. Otherwise, increase R by 1 and go back to step 2

Chapter 4

4 Empirical methods

Instead of using forecasting methods and a normal distribution, we also look at two methods that use an empirical demand distribution in order to determine the inventory policy parameters. The empirical distribution is obtained by sampling consecutive demands from historical demands. This distribution does not require regular or repeating patterns, and is also able to represent demand which does not fit to a common distribution. This is often the case for intermittent parts. The Empirical method described in Section 4.1 uses a constant lead time, whereas the Empirical plus method described in Section 4.2, which we propose, also takes the realised lead times into account. In Section 4.3, we then describe how we obtain the reorder point based on the empirical distribution.

4.1 Empirical model

The first method using an empirical distribution that we consider has been proposed by Porras and Dekker (2008). This method constructs a histogram of demands over the lead time without sampling. Demands are taken directly from the data set over a fixed number of periods, equal to the lead time. Because this method does not sample individual demands, it automatically captures autocorrelations and fixed demand intervals due to preventive maintenance. Porras and Dekker (2008) have shown that the empirical distribution slightly outperforms the method of Willemain, Smart, and Schwarz (2004) on total inventory costs, while achieving a similar service level.

We now give an example of how this works in practice. Consider that we have a part with demand represented by Figure 29, which has a lead time of 3 periods.

We first determine a starting point, in this case period 1. It then samples the demand over the periods 1 to 3, which leads to a LTD of 1. In the next sample, we start at period 17. We then sample the demand over the periods 17 to 19, which leads to a LTD of 2. Our third sample begins at period 7, which results in a LTD of 2. We repeat this process many times, for instance 500. In this way, we obtain an empirical distribution of the LTD.

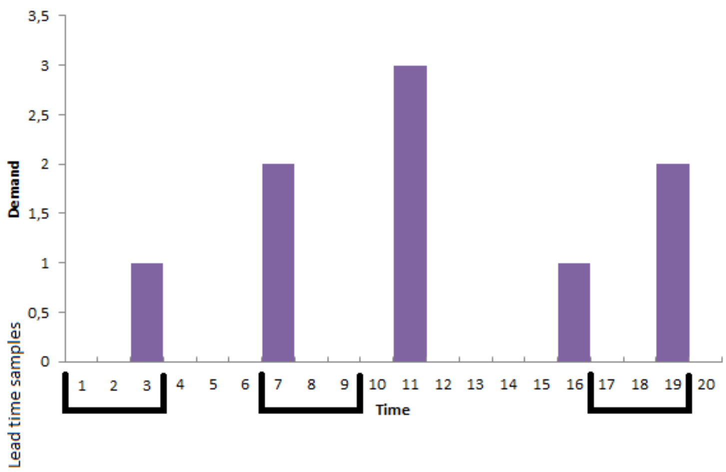

Figure 29: A figure representing three samples of the Empirical method, with a lead time of 3 periods

4.2 Empirical plus model

In reality the lead time can differ a lot each order. We have seen in Chapter 2 that the planned lead time can sometimes be quite different from the realised lead time. As all previous methods do not take this into consideration, we have thought of another method that does take this into consideration. Instead of always taking demands from the data over a fixed number of periods, we vary the number of periods. The number of periods is determined by randomly picking one of the realised lead times for each sample. When there is no realised lead time, we simply take the planned lead time instead. We call this method the Empirical Plus method from now on.

We now give an example of how this works in practice. Consider that we have a part with demand represented by Figure 30, which has three realised lead times of 1,2, and 4 periods. We first determine a starting point, in this case period 1. It then randomly picks one of the realised lead times, for instance 2. It then samples the demand over the periods 1 to 2, which leads to a LTD of 0. In the next sample, we start at period 16, we randomly pick a lead time of 4 periods. We then sample the demand over the periods 16 to 19, which leads to a LTD of 3. Our third sample begins at period 7, and randomly pick a lead time of 1 period, which results in a LTD of 2. We repeat this process many times, for instance 500. In this way, we obtain an empirical distribution of the LTD.

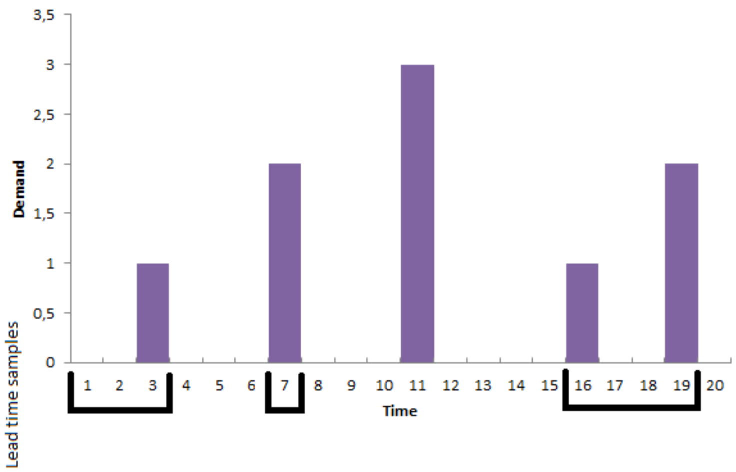

Figure 30: A figure representing three samples of the Empirical plus method, with three different lead time values

4.3 Determining the reorder point for an empirical distribution

Let us first introduce the following notation:
x = a sample of the lead time demand obtained from the empirical method

We then calculate the fill rate (FR) as following:

$$FR = 1 - \frac{\sum_{x|x>R}(x-R)}{\sum_{x}(x)} \tag{11}$$

Note that this method does not consider that the starting stock could be less than the reorder point, which some other methods do consider. In Appendix C.2, we describe this alternative method to determine the reorder point, and discuss why we chose to use equation (11).

After we have obtained our reorder point we use the same iterative approach as described before to determine the reorder point, such that the fill rate target is met, which is as following:

1. Start by setting $R = 0$
2. Calculate the service level
3. If the service level is high enough, stop.
4. Increase R by 1 and go back to step 2

After this iterative procedure we have obtained all parameters that are needed in order to do the simulation.

Chapter 5

5 Simulation set-up

In this chapter, we explain how the simulation study is set up. In Section 5.1, we decide what inventory policy we use, and how we determine the corresponding order quantity. In Section 5.2 we explain how we obtain our results using simulation.

5.1 Inventory policy and order quantity

Let us first introduce some notation:
Inventory level (IL) = Stock on Hand - Back orders
Inventory position (IP) = IL + Outstanding orders

In our simulation, we use an (R, nQ) model with a daily review period (see Axsäter, 2004, p. 88). Whenever the IP becomes equal to or less than R, an order is placed of size nQ, where $n \in \{1, 2, \cdots\}$, such that $R < IP \leq R + Q$. When the companies do not have ordering costs but may have restrictions of the supplier, such as a minimum order quantity (MOQ), or a modular order quantity (MOD), this inventory policy is known to be optimal (see Axsäter, 2004, p. 138). We always use a $Q = 1$, except when there is a MOQ or MOD specified in the data. In this case, we adjust the Q such that the requirements are met. By choosing a small value of Q, the size of the order has less impact on the total performance. If Q becomes too large, there will be (almost) no difference between the different methods, as the performance is mostly determined by the order size. However, we should note that even though we do not have data regarding ordering costs, the companies do have to pay some ordering costs.

5.2 Obtaining results by simulation

In order to obtain the results by simulation we:

- Aggregate the daily demands to periods of one month, as we make our forecasts for the traditional methods based on monthly demand
- Filter all parts that have had no demand or that have had only demand in the first 24 months. These parts have no demand during our simulation, and therefore there is no fill-rate defined for these parts
- Filter all parts that only had demands in a time-frame shorter than 6 months. These parts have probably been used for a short running process. As this is often known in advance, we assume the stock of these parts is not determined using forecasting methods

We calculate the holding costs based on the total stock on hand. An alternative could be to determine the holding costs based on the stock on hand and the total stock on order. As the companies often pay the suppliers the moment the stock arrives, calculating the holding costs based on the stock on hand would give a more realistic representation for the holding costs. We calculate the holding costs using a holding cost rate of 20%.

When we compare the performance of different forecasting methods, there is a trade-off between the service level obtained and the total holding costs. When we thus have only one value of the fill-rate and holding costs for each method, it is not always possible to determine the best method. Therefore, we run the simulation for each method for each of the following service level targets: {50%,55%,60%,65%,70%,75%,80%,85%,90%,95%,96%,97%,98%,99%,99.5%}.

For the initialisation of our forecasting methods we use the first two years of the data. Our first forecast is determined by taking the average over the first 12 months. We then update the forecast each month, for 12 months. At the end of this initialisation, we have a forecast for month 25, using the first 24 months.

When using one of the traditional forecasting methods described in Chapter 3, we determine the smoothing parameter(s) by minimising the MSE. We first make a forecast over all the data for different smoothing parameter(s) between 0.1 and 0.3 with step size of 0.05. We then choose the smoothing parameter(s), such that we obtain the lowest average MSE.

We also evaluate a method that chooses one of the traditional forecasting methods described in Chapter 3, and the corresponding smoothing parameter(s), by minimising the MSE for each part. We call this method the 'MSE' method from now on.

We use a discrete event based simulation with events $t \in \{0, \ldots, T\}$ to obtain our results. We distinguish three types of events:

- 'Order arrival', represents the arrival of an order from the supplier
- 'Demand', represents the demand for a spare part
- 'Update', represents the update of the reorder point, which occurs every first day of the month

Notice that all demands are fixed, as they are based on historical demand data.

In the simulation we use the following sequence of events. First, we handle the 'Order arrival' event. We add the current inventory by the size of the order that has arrived and update the total holding costs.

Then, we handle the 'Demand' event. We update the holding costs and remove the parts from the total stock on hand. If the total stock on hand is below the reorder point, we place a new order. If we order during the simulation, this results in an added event of the type 'Order arrival'.

Finally, if it is the first day of the month, we update the reorder point. For the traditional forecasting methods, we make a new forecast of the average and calculate the MSE over the last 12 months. Based on these two values, we calculate the new reorder point by applying the formulas given in Chapter 3. For the Empirical methods we repeat the process described in Chapter 4 in order to obtain a new reorder point. At each moment, there is a certain period over which the forecast is based, and a period over which we evaluate the result of the forecasts. Note that these two do not overlap, but do change in size over time.

We next introduce the following notation:

- P (Price) = The price of a part
- Q = Lot size
- h (Holding costs rate) = 20%
- LIC (Last inventory change) = The last event when the inventory changed
- IL (Inventory level) = The current number of parts on stock - the number of back orders
- IP (Inventory position) = IL + outstanding orders
- HC (Holding costs) = The total holding costs
- Total number of orders = Number of orders that have taken place so far
- Total demand = Total demand that has taken place so far

- Total on hand = Total demand that has been delivered from stock on hand so far
- Total orders on hand = Total orders that have been completely delivered from stock on hand so far
- Part fill-rate = The part fill-rate calculated after the simulation
- Order-line fill-rate = The order-line fill-rate calculated after the simulation
- $Event_t \in \{'\text{Demand}', '\text{Update}', '\text{Order arrival}'\}$
- $Demand_t$ = Demand for $Event_t \in$ Demand
- $Order\ size_t$ = The size of the order for $Event_t \in$ Order arrival
- $Date_t$ = The date event t takes place
- R_t = Reorder point

The simulation is then as follows:

Step 1 Initialisation

- Calculate R_0 based on one of the forecasting methods described in Chapter 3 or Chapter 4, using the first two years of our data
- $HC := 0$
- $LIC := 0$
- $IP := R_0 + Q$ [1]
- $IL := R_0 + Q$
- Total number of orders $:= 0$
- Total demand $:= 0$
- Total on hand $:= 0$
- Total order on hand $:= 0$

Step 2 $t := t + 1$; If $t >= T$: go to step 3

- If $event_t = '\text{Order arrival}'$:
  (a) $HC := HC + \frac{(Date_t - Date_{LIC})}{365} \cdot P \cdot h \cdot \max\{IL, 0\}$
  (b) $IL := IL + Order\ size_t$
  (c) $LIC := t$
  (d) Go to step 2
- If $event_t = '\text{Demand}'$:
  (a) $Update\ HC$:
  (b) $HC := HC + \frac{(Date_t - Date_{LIC})}{365} \cdot P \cdot h \cdot \max\{IL, 0\}$
  (c) $IL := IL - Demand_t$
  (d) $LIC := t$
  (e) Total Demand $:=$ Total Demand $+ Demand_t$
  (f) Total on hand $:=$ Total on hand $+ \max\{\min\{Demand_t, IL\}, 0\}$

[1] Teunter and Sani (2009) mention that different authors have proposed different ways of determining the starting stock, and that there is no clear idea of how this should be done.

(g) Total number of orders := Total number of orders $+ 1$
(h) If $IL \geq 0$:
   i. Total orders on hand := Total orders on hand $+ 1$
(i) If $IP \leq R$:[1]
   i. Determine n, such that $R < IP + nQ \leq R + Q$
   ii. $IP := IP + nQ$
   iii. Determine the time the order arrives
(j) Go to step 2

- If $event_t =$ 'Update$'$:
  (a) $R := R_t$
  (b) If $IP \leq R$:
     i. Determine n, such that $R < IP + nQ \leq R + Q$
     ii. $IP := IP + nQ$
     iii. Determine the time the order arrives
  (c) Go to step 2

Step 3 Calculate the achieved service levels

- Part fill-rate $= \frac{\text{Total on hand}}{\text{Total Demand}}$
- Order-line fill-rate $= \frac{\text{Total orders on hand}}{\text{Total number of orders}}$

[1] When a demand occurs, and an order has been placed, it is possible that at the same day the reorder point is updated to a lower value, such that $IP \geq R + Q$. If we would have first updated the reorder point we would not have placed an order. In our simulation we still let the order arrive when this occurs.

Chapter 6

6 Results

In this chapter we use the results obtained from the simulation described in Chapter 5, to make a classification for spare parts. First we look at recent literature about the classification of spare parts in order to obtain possible classification criteria, in Section 6.1. In Section 6.2, we evaluate the possible classification criteria, and suggest a classification based on the most important classification criteria. We finally also check whether we should determine the smoothing parameters for the traditional forecasting methods based on the MSE, or whether it is better to use constant smoothing parameters, in Section 6.3.

6.1 Recent literature on the classification of spare parts for inventory control

In this section, we describe a paper and a Master's thesis that are strongly related to the current research. We will use the literature as a starting point for our classification.

Boylan et al. (2008) classify spare parts, such that the performance of the inventory control can be improved when choosing the right forecasting methods. They look at the inter-demand interval in months, and at the squared coefficient of variation of demand, CV^2. When calculating the CV^2, they only look at non-zero demands. The CV^2 is then calculated as following:

$$CV^2 = \left(\frac{\text{Standard deviation of demand}}{\text{Average demand}}\right)^2$$

They compare two different forecasting methods for slow moving items: Croston's method, and the SBA. They also compare two different forecasting methods for fast moving items: SES and SMA. They then determine the breaking points, the point at which another method starts to perform better, based on the errors of the forecasting methods. As a measurement they use the geometric root mean squared error (GRMSE) instead of the MSE. See Boylan et al. (2008) for a detailed description on how to calculate the GRMSE. Once the breaking points have been determined they use the SBA and SMA as the forecasting methods to compare the results. They look at the service level that is achieved, and at the amount of stock that is put on hand when applying service level targets between 93% and 97 %. We should note that the breaking points have been determined based on the results, using a single data set only. When they would have used multiple data sets they could have obtained different breaking points (for each data set). Figure 31 represents the classification of Boylan et al. (2008).They do also mention that other decision variables could be used in the classification to extend the framework.

Van Duren (2011) also made a classification to increase the performance of inventory control. However, where Boylan et al. (2008) focus mainly on the classification to select appropriate forecasting methods, Van Duren (2011) is mainly interested in a classification to choose the right inventory control policy. Van Duren (2011) also considers the framework of Boylan et al. (2008), and extends this with other criteria. He first looks at all possible criteria on which the classification can be made, and then looks at how they could be used to enable him to choose better inventory control policies. He then makes a classification, based on some of the criteria. For each class he uses a different inventory control policy. We should note that the results of Van Duren (2011) are obtained using a single data set, consisting of 2 years of data from KLM Engineering & Maintenance.

From these two papers we obtain the following classification criteria:

1. Average inter-demand-interval

2. Squared coefficient of variation of demand

3. Average Lead time

4. Average demand size

5. Price

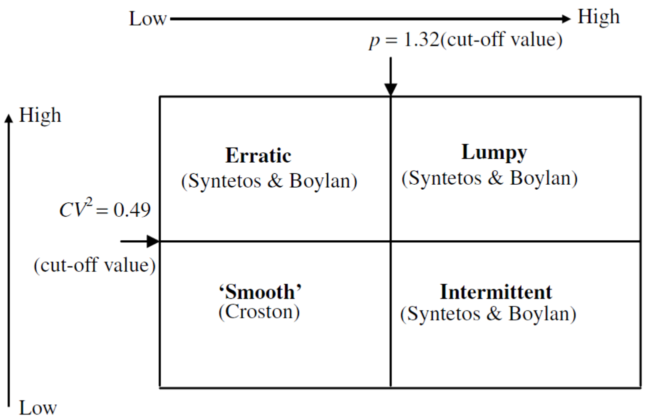

Figure 31: Possible classification of Boylan et al. (2008). The recommended forecasting method is given between the brackets. CV^2 represents the coefficient of variation in demand, excluding demands of 0. p is the average inter-demand interval in months.

6.2 Classification for spare parts

In this section we come up with a classification for spare parts based on the results from the simulation described in Chapter 5, and the classification criteria mentioned in the previous section. We first look at the overall performance of the different methods to get some first insights. We then evaluate the classification criteria and determine the most important classification criteria. After we know the most important classification criteria, we look at how the methods behave for different values of these criteria. Finally, we determine the breaking points, and propose a classification for spare parts.

6.2.1 Overall performance

In Figures 32, 33, and 34 we see the overall performance of all methods at respectively NTH, Alstom Ridderkerk, and at the NMC. On the horizontal axis, we see the achieved average order line fill-rate. On the vertical axis we see the average holding costs. We obtain more or less similar results at all three companies, except for the Empirical methods, which do not give good overall results at NTH, whereas it does give good overall results at the other companies. In Chapter 2, we noticed that NTH has many parts with short average demand-intervals. We further noticed in Chapter 4, that the way we obtain the reorder point for the Empirical methods might not always work well for parts with a low average demand-interval. Because of this, the Empirical methods do not give good overall results at NTH.

We also see that the traditional forecasts all show a similar curve in the figures. This can be explained by the fact that these methods all use the normal distribution to determine the reorder point. However, we see a clear difference in the performance. At all three companies we see that the SBA, performs best among the traditional forecasting method, followed by respectively

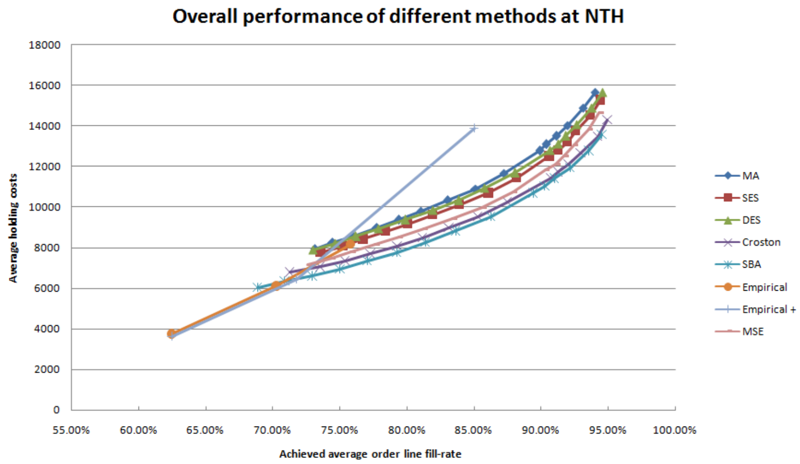

Figure 32: Relation between the achieved order line fill-rate and average holding costs for different methods at NTH

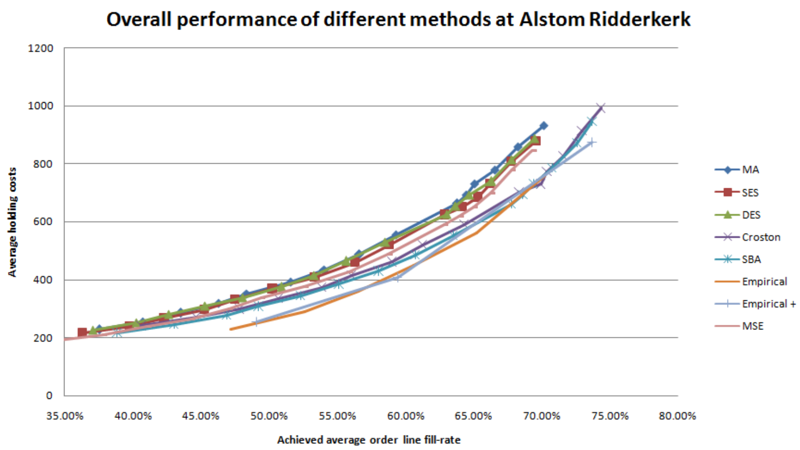

Figure 33: Relation between the achieved order line fill-rate and average holding costs for different methods at Alstom Ridderkerk

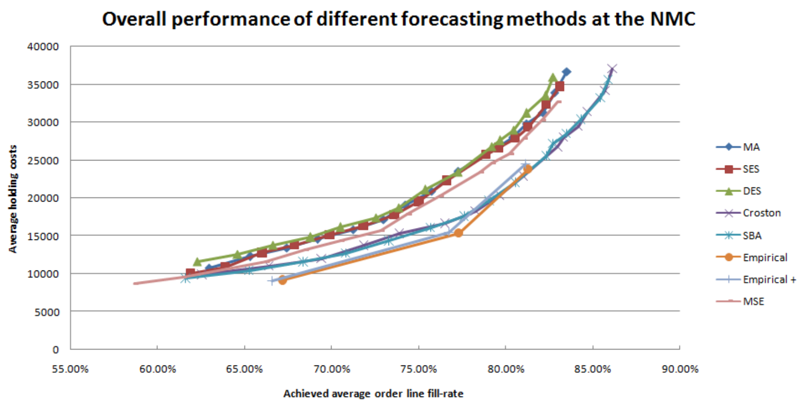

Figure 34: Relation between the achieved order line fill-rate and average holding costs for different methods at the NMC

Croston's method, the MSE method, and SES. MA gives the lowest performance, except at the NMC, where DES gives the lowest performance.

Based on these results we find that, it is better to always choose the SBA as forecasting method, instead of choosing the method and the smoothing parameter(s) for each part based on the MSE. However, in the literature the method and the smoothing parameters are often determined based on the MSE (see, e.g. Ghobbar and Friend (2003), and Eaves and Kingsman (2004)).

6.2.2 Evaluation of the classification criteria

We now have some overall insights into the performance, but we already mentioned that different parts would probably require different methods in order to obtain a better performance. Therefore, we first divide all parts into three equally sized groups based on the classification criteria mentioned in section 6.1. In this way, we obtain 15 different groups of parts for each company. For each classification criteria we look at the two groups with either the lowest values or the highest values, as the results of the middle group is probably a mixture of the other two. Note that because we determine the cut-off values in this way, we obtain different cut-off values for each company. Extensive results can be found in Appendix D, here we give the most important results.

The first key observation is that the average demand-interval and the squared coefficient of variation of demand have the most influence on the performance for the different methods. Remark that these are the same classification criteria Boylan et al. (2008) find. We find that the other classification criteria have less impact on the performance of the different methods.

For parts that have a large average demand-interval, the SBA overall performs better compared to the other traditional forecasting methods. At the NMC, a decrease in the total inventory costs can be achieved of approximately 36% using the SBA instead of MA while maintaining the same service level. We also see that the Empirical methods perform well, which is in line with our expectations. As the empirical methods does not require fitting to a distribution, it is better at capturing the lead time demand distribution for these parts as fitting to a distribution is difficult.

When the average demand-interval is short, the SBA still performs better than the other traditional forecasting methods, although the difference is smaller. At the NMC, a decrease in the total inventory costs can be achieved of approximately 18%, by using the SBA instead of MA, and maintaining the same service level. However, the Empirical methods now perform less good compared to the traditional forecasting methods, as they are not able to obtain higher order line fill-rates.

For large values of the CV^2, the traditional forecasting methods give better results than the empirical methods. This can be explained by the fact that the empirical methods are not able to put more stock on hand than the largest lead time demand sample. However, we should note that Willemain et al. (2004) propose a method called jittering, ie adding some random variation to non zero demands.

Note that the difference in holding cost is an average, and does not necessarily mean that a certain method is always best for that part. Based on one observation, we can only say that a method is better when both the costs are lower and the fill-rate obtained is higher. When we look at each part of NTH individually, approximately 5% of the time, the SBA is better than MA. For approximately 4 % of the time, MA is better than the SBA. For all other parts, we cannot determine which method gives the best performance based on one result.

Let us now look at an example showing the difference between MA and the Empirical method, represented in Figure 35. The first line represents the demand for a part at Alstom Ridderkerk, which has a lead time of 8 months. For the calculation of the reorder point we set a service level target of 90%. The second line represents the reorder point obtained using MA, whereas the last line represents the reorder point obtained using the Empirical method. At the beginning of the simulation the Empirical method starts with a higher reorder point compared to MA. However, after the demands occur, the reorder point of MA increases rapidly, whereas the reorder point calculated by the Empirical method is more stable. As a result the Empirical method obtains both a higher service level, and lower holding costs. However, this is not always the case, as most of the time we cannot determine the best method based on only one observation.

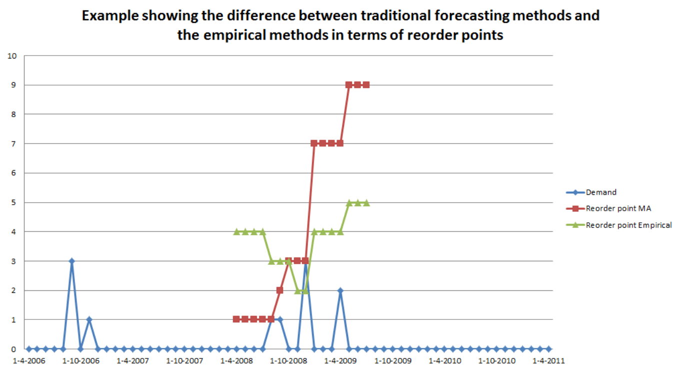

Figure 35: Example showing the differences between MA and the Empirical method in the reorder points obtained

6.2.3 Determining the breaking points

In this section we look at determining the breaking points, such that it enables us to suggest a good classification for spare parts. We find a breaking point for the average demand-interval (p) of 4 and a breaking point for the squared coefficient of variation of demand (CV^2) of 0.3. However, if we would have determined the breaking points based on the data of one company only, the breaking points may have been (slightly) different. Note that these breaking points are different from the breaking points Boylan et al. (2008) recommend in Figure 31. This is probably because they have different data and also consider different methods.

An overview of the distribution of the number of parts at each company and group is given in Table 4. We see that NTH has relatively many parts that have an average demand-interval less than 4, whereas Alstom Ridderkerk and the NMC have more parts with an average demand-interval > 4. We also see that the NMC has relatively few parts with a high CV^2 in the group of slow movers compared to the other companies.

Number of parts with:	NTH €	Alstom Ridderkerk	NMC
$p > 4$ and $CV^2 < 0.3$	52	217	1043
$p > 4$ and $CV^2 > 0.3$	78	346	611
$p < 4$ and $CV^2 < 0.3$	176	29	76
$p < 4$ and $CV^2 > 0.3$	482	146	247

Table 4: Distribution of the number of parts at all three companies

Our second key observation is that for the parts that have a p larger than 4, and a CV^2 less than 0.3, the Empirical Plus method overall performs best. However, at NTH the difference between the Empirical methods, Croston's method and the SBA is not that large. At the higher service level targets these methods even perform better than the Empirical Plus method. However, as the amount of parts at NTH in this group is not that large compared to the other companies, and the Empirical Plus method clearly performs better than the forecasting methods at

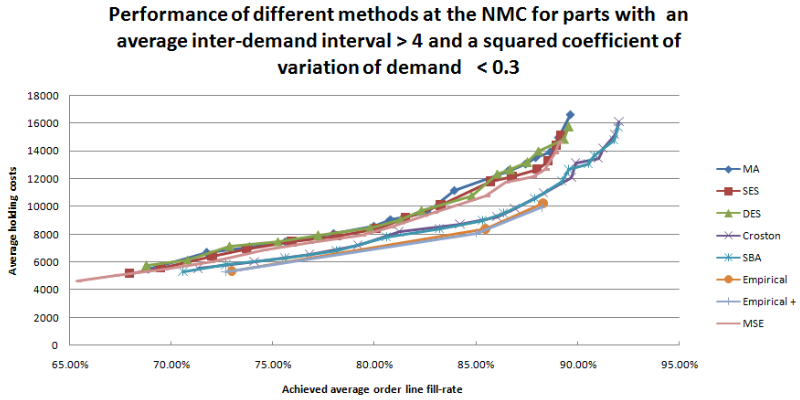

Figure 36: Relation between the achieved order line fill-rate and average holding costs for different methods at the NMC

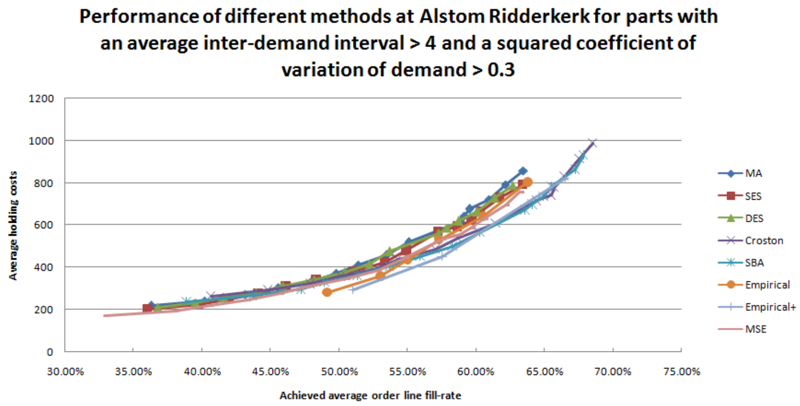

Figure 37: Relation between the achieved order line fill-rate and average holding costs for different methods at Alstom Ridderkerk

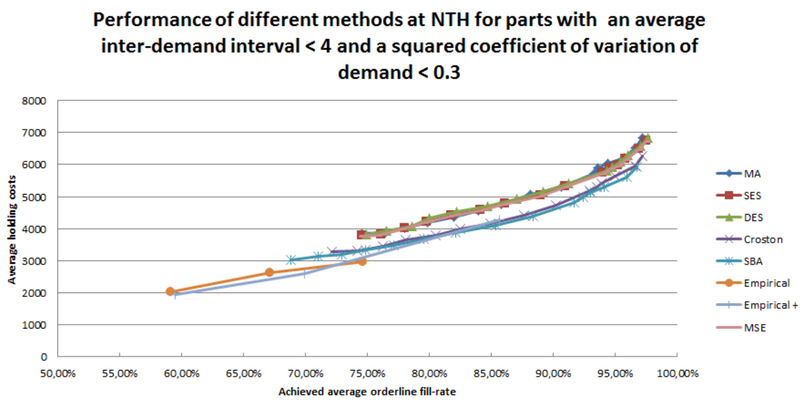

Figure 38: Relation between the achieved order line fill-rate and average holding costs for different methods at the NTH

the other companies, we still recommend this method. Figure 36 shows the results of this group of parts at the NMC. When there are many zero demands, but there is not much variation in the demand, the forecasting methods are more likely to get a higher variation, due to the difference between the forecast and the zero demand periods. As the Empirical methods do not have this problem, and are also good capable of forecasting intermittent demand patterns, this could explain why the Empirical methods perform better especially for these group of parts.

Although the Empirical Plus method gives similar performance compared to the Empirical method at NTH and at the NMC, we see that at Alstom Ridderkerk the Empirical Plus method gives slightly better results than the Empirical method. This can be explained by the fact that at Alstom the realised lead times are often longer compared to the planned lead times. Because of this, the total demand in the lead time using the Empirical method is more likely to be underestimated.

Our third and last key observation is that for all other groups of parts, we should use the SBA. Note that the performance of Croston's method is more or less similar to that of the SBA. For the parts that have an p larger than 4, and a CV^2 larger than 0.3, the Empirical method performs better at the lower service level targets. However, when the service level target becomes slightly higher, the SBA and Croston's method give better results. Figure 37 shows the result of this group of parts at Alstom. For the parts that have an p less than 4, and a CV^2 less than 0.3, we should note that we obtain different results at each company. However, at NTH we see that the SBA and Croston's method give the best results for these parts. When we take into consideration that the amount of parts at Alstom Ridderkerk and at the NMC is limited in this group compared to NTH, and the fact that the difference between the methods is small at the other companies, we suggest to use the SBA for these parts. Figure 38 shows the results of this group of parts at NTH.

Finally, based on these observations, we recommend the classification depicted in Figure 39.

6.3 Evaluating the influence of the smoothing parameters

In the previous section we found that the MSE method, which chooses both one of the traditional forecasting methods and the smoothing parameters based on the MSE, does not give the best results. This is especially the case for parts with a long average demand-interval. As we find that the SBA is for many parts the best method, we will look at what smoothing parameters should be chosen for this method. Although we have two smoothing parameters for this method, we used the same value for α as β. If we would have a business where the demands will (rapidly) increase or decrease over time, we would rather use higher values for the smoothing parameters. For businesses where the demand is more or less equal over time, we will rather use small values for the smoothing parameters. We also looked at using different values for α and β, but obtained similar results.

When we look at the overall results at all three companies, we see that if we determine the smoothing parameter values based on the MSE, we obtain the best results. Figure 40 shows the overall results using different smoothing parameters at Alstom. Similar results are also obtained at the other two companies. Therefore, we recommend to determine the parameters based on the MSE when using the SBA method.
We also look at whether this holds when using SES as forecasting method. However, as we can see in Figure 41, we would be better off using a constant value between 0.05 or 0.15, if we use SES for all parts. We also found that using larger values for the smoothing parameters is not recommended as we can see from the results in Figure 41.

Although we see that the smoothing parameter(s) have some impact on the results, the impact is relatively small compared to choosing the right forecasting method. Teunter and Sani (2009) also find that the impact of different smoothing parameters is not that big.

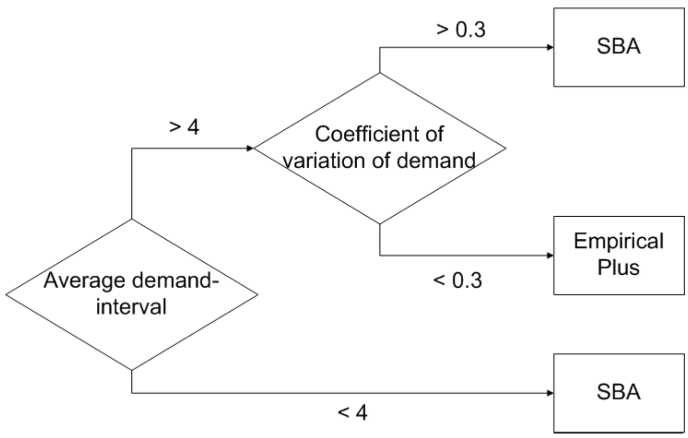

Figure 39: Classification for spare parts based on the average demand-interval and the coefficient of variation of demand

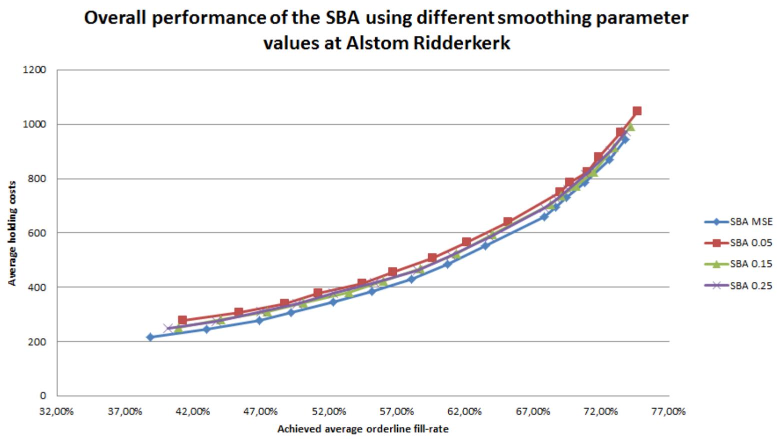

Figure 40: Relation between the achieved order line fill-rate and the different smoothing parameters for the SBA at Alstom Ridderkerk

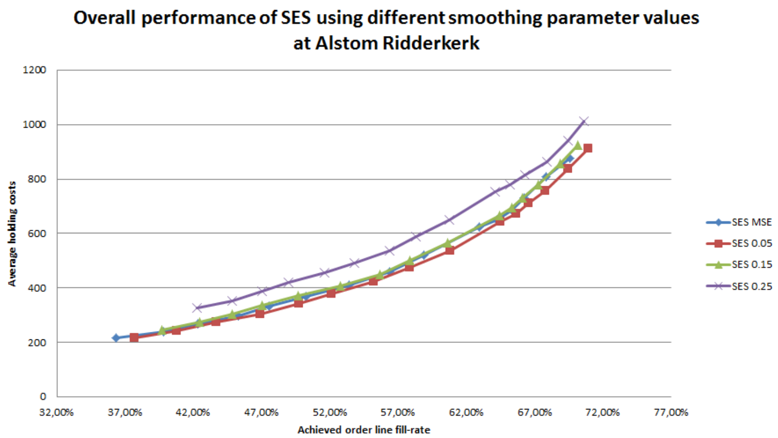

Figure 41: Relation between the achieved order line fill-rate and the different smoothing parameters for SES at Alstom Ridderkerk

Chapter 7

7 Conclusions and recommendations for further research

In Section 7.1 we present our key observations, and our proposed classification based on the simulation. In Section 7.2, we then give our recommendations for future research.

7.1 Conclusions

Our first key observation is that the average inter-demand interval and the squared coefficient of variation of demand (CV^2) have the most impact on the performance of different methods. We looked at different classification criteria, but these two classification criteria have the most influence on the performance on the different methods.

Our second key observation is that for parts that have a large average inter-demand interval, but have a low CV^2, the Empirical Plus method gives the best results. Note that we could also use the Empirical method as the performance of these two methods are more or less similar.

Our third and final key observation is that for all other parts, the SBA gives the best results and that we should determine the smoothing parameter for the SBA based on the MSE. However, we should note that Croston could also be a good alternative, as there is not much difference in the results between these two methods.

Based on these results, we suggest the classification for the spare parts described by Figure 42. This classification enables us to decrease the holding costs while maintaining the same amount of service, or increase the amount of service while maintaining the same amount of holding costs.

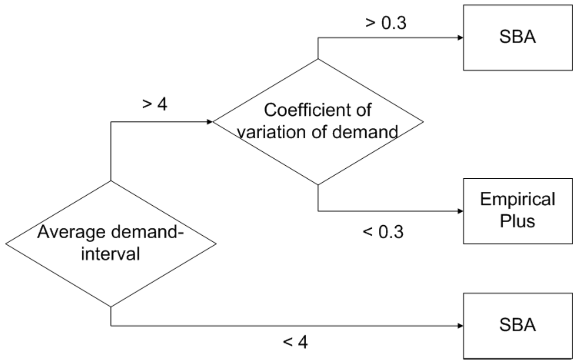

Figure 42: Classification for spare parts based on the average demand-interval and the coefficient of variation of demand

7.2 Recommendations for future research

During our research we found some topics that could be interesting for future research. For instance, in Section 5.1, we decided to use an order quantity of one if we did not have any information about the MOQ or MOD. However, in reality it is not common to always order with an order size of one as we also have ordering costs. This made it difficult to determine the reorder point for the Empirical methods in a good a way. Therefore, it could be interesting to determine whether the same methods perform best if we would use another method to determine the order size, for instance by calculating the EOQ.

In Section 3.1 we mentioned that we had to adjust the initialisation for Croston and the SBA, as we would otherwise underestimate the demand. However, the way we initialise the demand also has an impact on the performance of the different methods. Therefore, it could be interesting for future research, to look at the importance of a proper initialisation, and how we can initialise the demand as good as possible.

In this research we only looked at the normal distribution, and the empirical distribution to model the demand. It might be useful to look at different distributions. For instance the (Compound) Poisson, the log normal distribution or the Gamma distribution.

During this research, Gordian developed another method called the Typical Demand Quantity (TDQ) method. It assumes that the time between requests is exponentially distributed, and that the size of each demand has an empirical distribution, based on the sizes of the demand that have occurred in the past. This method could especially be useful for NTH as we found that at NTH demands are likely to be a multiple of 4. If the average inter-demand-interval is large, this method still puts 4 on stock, whereas the traditional forecasting methods may put less or even more on stock. We expect that this method will give similar results as the Empirical methods. It could be interesting to compare the results we would obtain using this method, to the results we have found in Chapter 6.

In our data analysis of Chapter 2, we found some strange values. For instance, parts with a price of zero, parts with extremely high demand, and parts that have not been requested at all. It could be interesting for Planning Services to do some further research on methods and techniques to smartly identify strange values. When we are able to easily detect and filter these demands, this results in a better reliability of the data. As a result, the performance of the forecasting methods may increase.

In our study, we focus on the spare parts level only. Romeijnders, Teunter, and Van Jaarsveld (2012) develop a method that also takes the component repairs into consideration to improve the forecast. It might be interesting to obtain the necessary data, such that we are able to apply this forecasting method, as they show that this method is better than the SBA in terms of MSE.

8 References

Alfredsson, P., & Verrijdt, J. (1999). Modeling emergency supply flexibility in a two-echelon inventory system. *Management Science*, *45*(10), 1416-1431.

Axsäter, S. (2004). *Inventory control*. Springer New York.

Babai, M. Z., Syntetos, A. A., Dallery, Y., & Nikolopoulos, K. (2009). Dynamic re-order point inventory control with lead-time uncertainty: analysis and empirical investigation. *International Journal of Production Research*, *47*(9), 2461-2483.

Beasens, B., Mues, C., Martens, D., & Vanthienen, J. (2009). 50 years of data mining and or: upcoming trends and challenges. *Journal of the Operational Research Society*, *60*, S16-S23.

Bitran, G. R., & Hax, A. C. (1977). On the design of hierarchical production planning systems. *Decision Sciences*, *8*, 28-55.

Boylan, J. E., & Syntetos, A. A. (2009). Spare parts management: a review of forecasting research and extensions. *Journal of Management Mathematics*, *21*, 227-237.

Boylan, J. E., Syntetos, A. A., & Karakostas, G. C. (2008). Classification for forecasting and stock control: a case study. *Journal of the Operational Research Society*, *59*, 473-481.

Braglia, M., Grassi, A., & Montanari, R. (2004). Multi-attribute classification method for spare parts inventory management. *Journal of Quality in Maintenance Engineering*, *10*(1), 55-65.

Cavalieri, S., Garetti, M., Macchi, M., & Pinto, R. (2008). A decision-making framework for managing maintenance spare parts. *Production Planning and Control*, *19*(4), 379-396.

Cohen, M. A., Agrawal, N., & Agrawal, V. (2006). Winning in the aftermarket. *Harvard Business Review*, 129-238.

Cohen, M. A., Zheng, Y.-S., & Agrawal, V. (1996). Service parts logistics: a benchmark analysis. *IIE Transactions*, *29*(8), 627-639.

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly*, *23*(3), 289-303.

Deloitte Research. (2006). *The service revolution in global manufacturing industries.*

Driessen, M. A., Arts, J. J., van Houtum, G. J., Rustenburg, W. D., & Huisman, B. (2010). Maintenance spare parts planning and control: A framework for control and agenda for future research. *BETA working paper v. 20.*

Duchessi, P. (1998). A conceptual approach for managing of spare parts. *International journal of physical distribution and materials management*, *18*(5), 8-15.

Eaves, A. H. C., & Kingsman, B. G. (2004). Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operational Research Society*, *55*, 431-437.

Flores, B., & Whybark, D. (1989). Implementing multiple criteria abc analysis. *Journal of operations management*, *7*(1), 79-85.

Gajpal, P. P., Ganesh, L. S., & Rajendran, C. (1994). Criticality analysis of spare parts using the analytic hierarchy process. *International journal of production economics*, *35*, 293-297.

Ghobbar, A. A., & Friend, C. H. (2003). Evaluation of forecasting methods for intermittent parts demand in the field of aviation: a predictive model. *Computers and Operations Research*, *30*, 2097-2114.

Holt, C. C. (2004). Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, *20*(1), 5-10.

Hua, Z. S., Zhang, B., Yang, J., & Tan, D. S. (2007). A new approach of forecasting intermittent demand for spare parts inventories in the process industries. *Journal of the Operational Research Society*, *58*, 52-61.

Huiskonen, J. (2001). Maintenance spare parts logistics: Special characteristics and strategic choices. *International Journal of Production Economics*, *71*, 125-133.

Kennedy, W. J., Wayne Patterson, J., & Fredendall, L. D. (2002). An overview of recent literature

on spare parts inventories. *International Journal of Production Economics*, *76*, 201-215.
Kranenburg, A. A., & Houtum, G. J. van. (2007). Effect of commonality on spare parts provisioning costs for capital goods. *International Journal of Production Economics*, *108*, 221-227.
Kwakernaak, R. (2008). *Inventory management of spare parts on a dredging ship.*
Petrovic, D., & Petrovic, R. (1992). Sparta ii: Further development in an expert system for advising on stocks of spare parts. *International journal of production economics*, *24*, 291-300.
Pinçe, C., & Dekker, R. (2010). An inventory model for slow moving items subject to obsolescence. *European Journal of Operational Research*, *213*(1), 83-95.
Porras, E., & Dekker, R. (2008). An inventory control system for spare parts at a refinery: An empirical comparison of different re-order point methods. *European Journal of Operational Research*, *184*, 101-132.
Ramathan, R. (2006). Abc inventory classification with multiple-criteria using weighted linear optimization. *Computers & operations research*, *33*, 695-700.
Romeijnders, W., Teunter, R., & Van Jaarsveld, W. (2012). *A two-step method for forecasting spare parts demand using information on component repairs.*
Sani, B., & Kingsman, B. G. (1997). Selecting the best periodic inventory control and demand forecasting methods for low demand items. *Journal of the Operational Research Society*, *48*, 700-713.
Schultz, C. R. (1987). Forecasting and inventory control for sporadic demand under periodic review. *Journal of the Operational Research Society*, *38*(5), 453-458.
Silver, E. A., Pyke, D. F., & Peterson, R. (1998). *Inventory management and production planning and scheduling*.
Supply Managers. (1997). Logistics suffers from skilled staff famine. *Supply Management*, *2*(13), 13.
Syntetos, A. A., Babai, M. Z., & Altay, N. (2010). Modelling spare parts demand: an empirical investigation. *8th International Conference of Modeling and Simulation*.
Syntetos, A. A., & Boylan, J. E. (2001). On the bias of intermittent demand estimates. *International Journal of Production Economics*, *71*(1-3), 457-466.
Syntetos, A. A., & Boylan, J. E. (2005a). The accuracy of intermittent demand estimates. *International Journal of Forecasting*, *21*, 303-314.
Syntetos, A. A., & Boylan, J. E. (2005b). The accuracy of intermittent demand estimates. *International journal of forecasting*, *21*, 303-314.
Syntetos, A. A., Boylan, J. E., & Croston, J. D. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, *56*, 495-503.
Syntetos, A. A., Boylan, J. E., & Disney, S. M. (2009). Forecasting for inventory planning: a 50-year review. *Journal of the Operational Research Society*, *60*, S149-S160.
Syntetos, A. A., Nikolopoulos, K., & Boylan, J. E. (1996). Judging the judges through accuracy-implication metrics: The case of inventory forecasting. *International Journal of Forecasting*, *26*, 134-143.
Teunter, R. H., & Babai, M. Z. (2010). Determining order-up-to levels under periodic review for compound binomial (intermittent) demand. *European Journal of Operational Research*, *203*, 619-624.
Teunter, R. H., & Duncan, L. (2008). Forecasting intermittent demand: a comparative study. *Journal of the Operational Research Society*, *60*, 321-329.
Teunter, R. H., & Sani, B. (2009). Calculating order-up-to levels for products with intermittent demand. *International Journal of Production Economics*, *118*, 82-86.
Van Duren, J. (2011). *Differentiated spare parts management: an application in the aircraft industry.*

Watson, R. B. (1987). The effects of demand-forecast fluctuations on customer service and inventory cost when demand is lumpy. *Journal of the Operational Research Society*, *38*, 75-82.

Willemain, T. R., Smart, C. N., & Schwarz, H. F. (2004). A new approach to forecasting intermittent demand for service parts inventories. *International Journal of Forecasting*, *20*, 375-387.

Willemain, T. R., Smart, C. N., Shockor, J. H., & DeSautels, P. A. (1994). Forecasting intermittent demand in manufacturing: comparative evaluation of croston's method. *International Journal of Forecasting*, *10*, 529-538.

Wouters, S. (2010). Functional specifications of a collaborative spare parts planning system. *Msc thesis University Maastricht*.

A Appendix
A.1 List of Figures

A.2 List of Tables

A.3 Symbols & Acronyms

GRMSE geometric root mean squared error

LSMO Large Scale Maintenance Organisation

LTD Lead Time Demand

MAPE Mean Absolute Percentage Error

MOD Module Quantity

MOQ Minimum Order Quantity

MSE Mean Squared Error

SBA Syntetos-Boylan Approximation

SES Single Exponentional Smoothing

SMA Simple moving average

SMO Small maintenance organisation

SPPCT Spare Parts Planning Control Tower

B Data obtained from the companies

In this Appendix, we will describe the describe the data we have obtained from the companies. In Appendix B.1 we will describe the data we have obtained of NedTrain Haarlem. In Appendix B.2 we describe the data we have obtained of Alstom. In Appendix B.3 we will also do this for the Naval Maintenance Company.

B.1 Available data NedTrain

- Information regarding the parts
  - Code of the part
  - Description of the part
  - Group of the part (Main parts, consumables or repairables)
  - Variety of the part (‘*inkoopdeel*’ or an ‘*maakdeel*’)
  - Price (€)
  - Unit of measure
  - Minimum order quantity (The minimum amount that has to be ordered when an order needs to be placed)
  - Modular order quantity (The size of the order must be a multiple of this value)
  - Lead time (The lead time in working days calculated from the contractual lead time added with some safety time)
- Information about demand orders
  - Order number
  - Code of the part
  - Quantity
  - Date of the demand
  - Planned demand (This code represents whether the demand was for a project. In general all parts installed for a project are planned demand. However NedTrain Haarlem did not add the project code consistently until 2009. Therefore, not all planned demand can be filtered.)
- Information about supply orders
  - Order number
  - Code of the part
  - Date the order has been placed
  - Expected date of arrival
  - Date of arrival
  - Quantity

B.2 Available data Alstom

- Information regarding the parts
  - Code of the part
  - Price (€)
  - Lead time (The lead time given in calendar days)
  - Unit of measure
  - Minimum order quantity (The minimum amount that has to be ordered when an order needs to be placed)
  - Modular order quantity (The size of the order must be a multiple of this value)
- Information about demand orders
  - Order number
  - Code of the part
  - Order type (Used for production or delivered to a customer directly)
  - Quantity
  - Date of the demand
  - Planned demand (This code represents whether the demand was known in advance)
- Information about supply orders
  - Order number
  - Code of the part
  - Creation date of the order
  - Expected date of arrival
  - Date of arrival
  - Quantity

B.3 Available data Naval Maintenance Company

- Information regarding the parts
  - Nato Stock Number (NSN)
  - Purchase price (€)
  - Lead time (The planned lead time given in months)
- Information about demand orders
  - Order number
  - Code of the part
  - Order type
  - Quantity
  - Date
  - Planned demand (Demand can be regular or incidental)
- Information about supply orders
  - Order number
  - Code of the part
  - Creation date of the order
  - Expected date of arrival
  - Date of arrival
  - Quantity

C Validating the simulation study

In order to validate the simulation study described in Chapter 5, we divide the validation into three steps. In Appendix C.1, we looked at whether given the demands, we obtain the proper forecasting results. After we have obtained the forecasting results, we look at whether we obtain the proper reorder point based on these results in Appendix C.2. Finally, we validate whether we obtain the results we wish to obtain given a certain reorder point and service level target in Appendix C.3.

C.1 Validating the forecasting methods

In order to validate the forecasting models we randomly select some parts and compare the forecasting results to the results we obtain by doing the calculations by the hand.In this way we validate the MA, SES, DES, Croston, and the SBA method. During this step we found that the initialisation of Croston and SBA was not always very good. We adjusted this initialisation as described in Chapter 3.
For the empirical methods we can not manually calculate the forecast result, as it consists of 1000 sampled lead time demand values. In order to validate these models, we use the debugger in Visual Basic, and validate whether the demand is sampled correctly over the lead time. In order to validate this, we used some simple dummy data. During this debugging we found some small bugs, as this sampling process in first instance was not always taking samples of exactly the lead time.

C.2 Validating the reorder point calculations

In order to validate the reorder point calculations, we used Appendix 2 of Axsäter (2004). We randomly picked some values for the average and variation of demand, and then manually calculated the reorder point. We compared this reorder point to the reorder point we calculate when determining the reorder point as described in Chapter 3. As the empirical distribution determines the reorder point in a different way, we also need to validate this process. We initially used the following method to calculate the reorder point for the Empirical distribution: First we calculated the expected shortage (ES):

$$ES(R) = \sum_{x|x>R} (x - R)f(x)$$

We then calculate the fill-rate (FR) as:

$$FR = 1 - \frac{ES(R)}{Q}$$

We used the debugger in Visual Basic and looked at whether the expected shortage was calculated correctly. We found that it was working correctly, however the results we found for this method were worse than the results of all the other methods. After discussing this problem, we found that the problem lies in the fact that we assume that $Q = 1$ for parts that do not have an MOQ. Due to this, the empirical method often puts too much stock on hand. Let us give an example: When we have a sample of 5 lead time demands of {0,0,1,1,10}. If we have that $R = 5$ and $Q = 1$, we obtain the following: $ES(5) = 0 * 0.4 + 0 * 0.4 + 5 * 0.2 = 1$ When you then calculate the fill-rate we obtain: $FR = 1 - \frac{ES(5)}{Q} = 1 - \frac{1}{Q} = 0$ Note that this fill-rate is lower than the actual fill-rate we would expect in this case. In order to deal with this problem we adjusted the method as following:

$$ES(R)_+ = \sum_{x|x>R} (x-R) f(x) ES(R)_- = \sum_{x|x>R} (x-R+Q) f(x) ES(R) = ES(R)_- - ES(R)_+$$

We then calculate the fill-rate as:

$$FR = 1 - \frac{ES(R)}{Q}$$

After running the simulation for the empirical distribution, we again found that the results were worse than we have found for the other methods. Therefore, we introduced the method described in Chapter 4.

C.3 Validating the simulation process

We also validated the simulation process, where we looked at whether the simulation follows the steps as described in Section 5.2. First we used the debugger in Visual Basic and followed the simulation for a few parts over the entire process. During this debugging we found a few bugs, which we adjusted. After the debugging we added some code to export information to a worksheet during the simulation. We check whether we see anything strange, and plotted a figure representing the inventory level over a certain time period. At this step we could conclude that the simulation follows the steps as described in Section 5.2.

D Extensive results

In this Appendix, we will present all results we have obtained from our simulation. We divide this appendix into three different sections. In Appendix D.1 we will look at the results we have obtained for NedTrain Haarlem. We look at different groups based on the classification criteria mentioned in Appendix 6.1. In Appendices D.2, and D.3 we will also do this for Alstom Ridderkerk, and for the Naval Maintenance Company.

D.1 NedTrain Haarlem

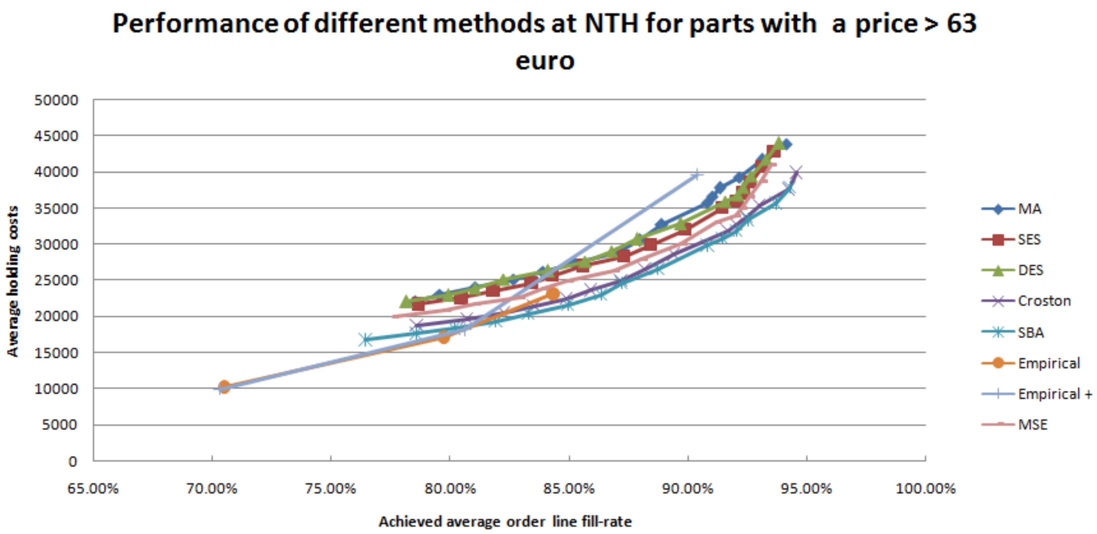

Figure 43: Relation between the achieved order line fill-rate and average holding costs for expensive parts at NTH

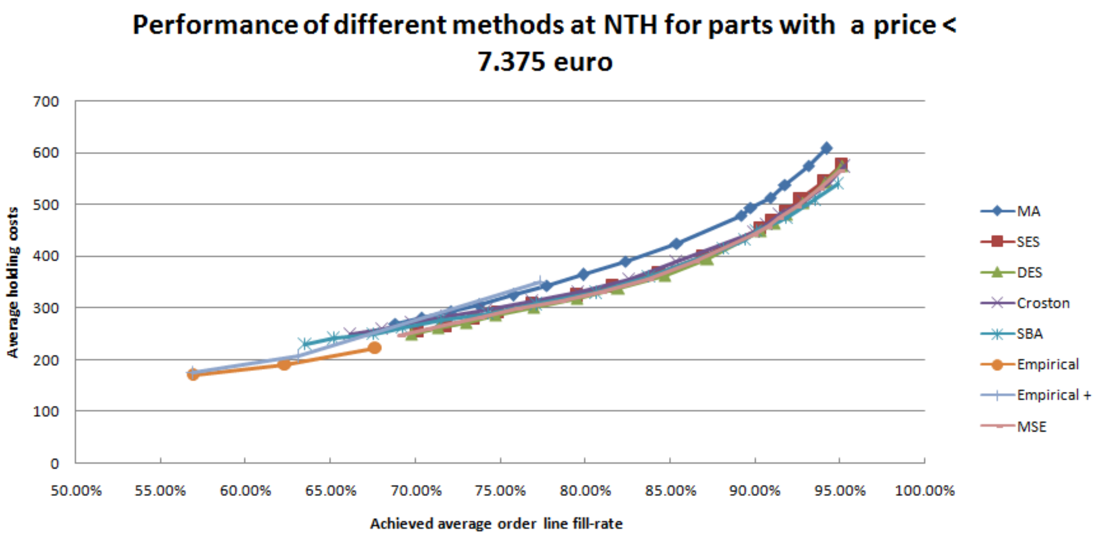

Figure 44: Relation between the achieved order line fill-rate and average holding costs for cheap parts at NTH

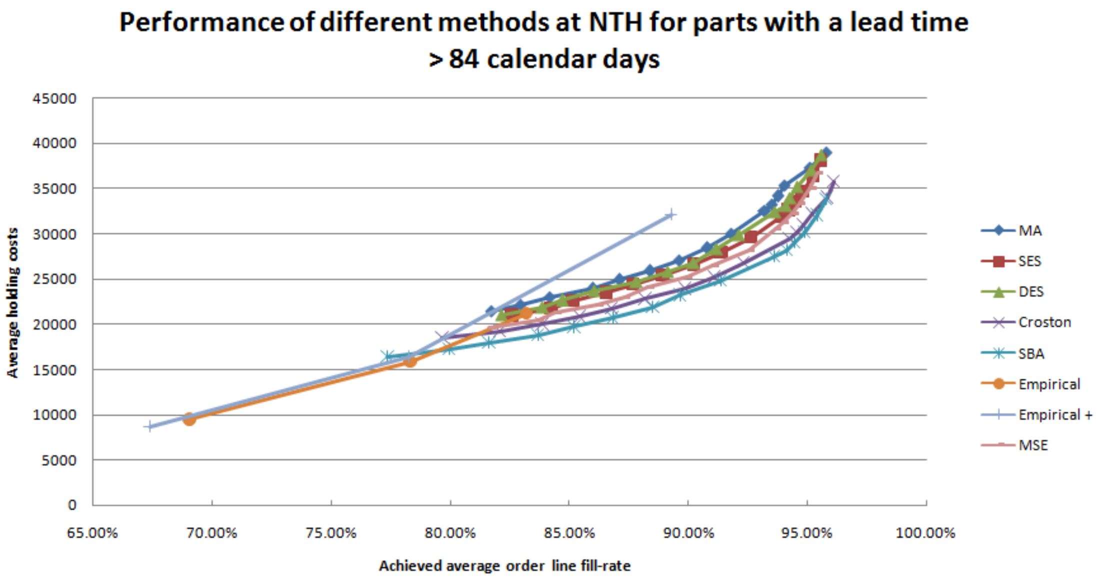

Figure 45: Relation between the achieved order line fill-rate and average holding costs for parts with a long planned lead time at NTH

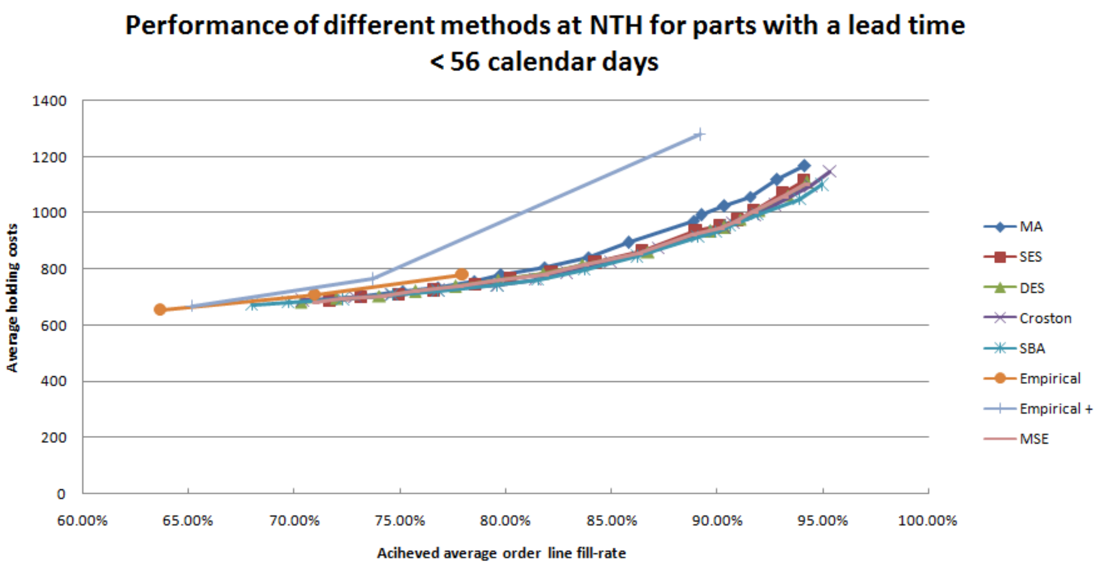

Figure 46: Relation between the achieved order line fill-rate and average holding costs for parts with a short planned lead time at NTH

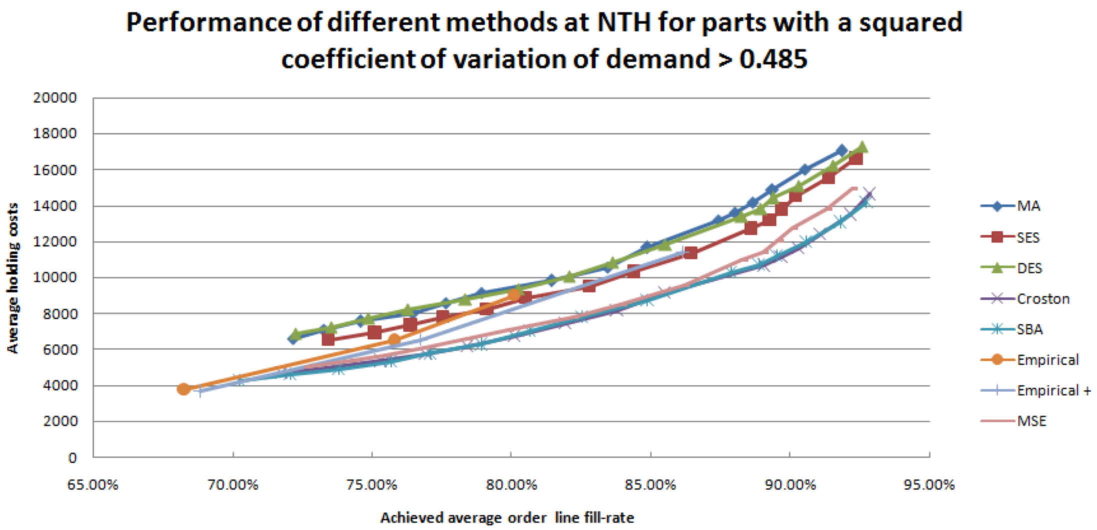

Figure 47: Relation between the achieved order line fill-rate and average holding costs for parts with a large squared coefficient of variation of demand at NTH

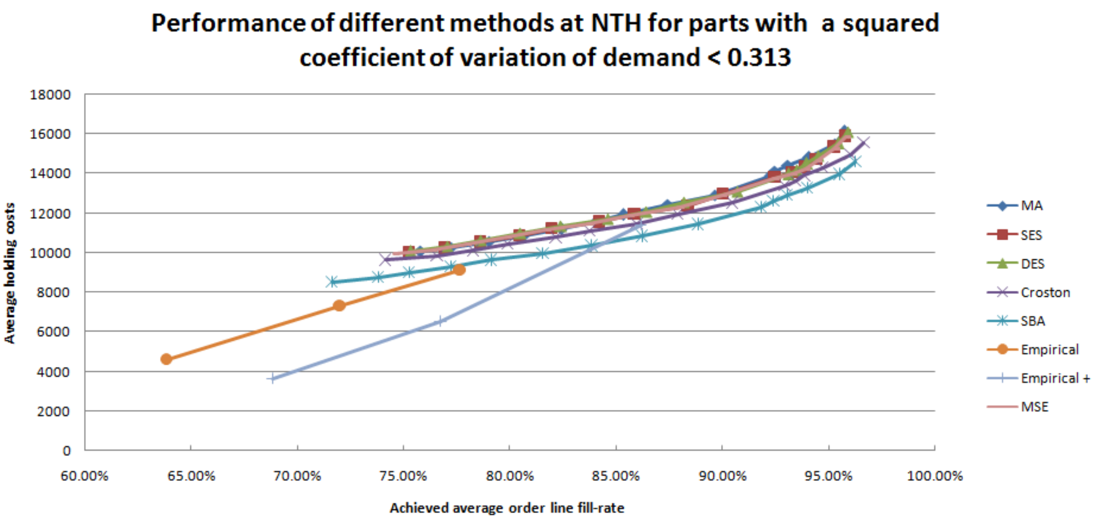

Figure 48: Relation between the achieved order line fill-rate and average holding costs for parts with a small squared coefficient of variation of demand at NTH

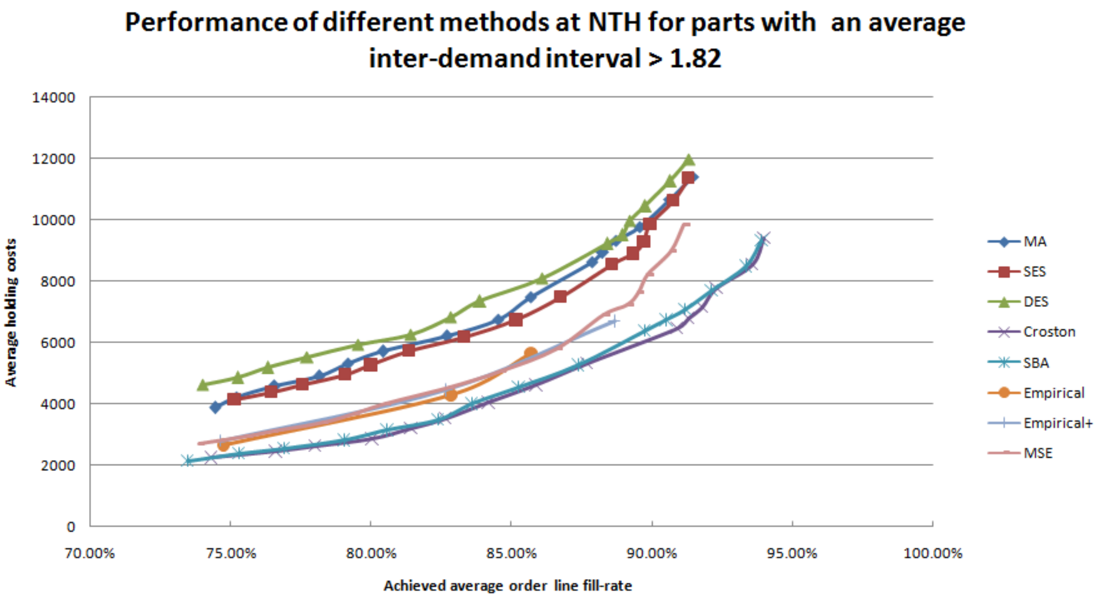

Figure 49: Relation between the achieved order line fill-rate and average holding costs for parts with a large average demand-interval at NTH

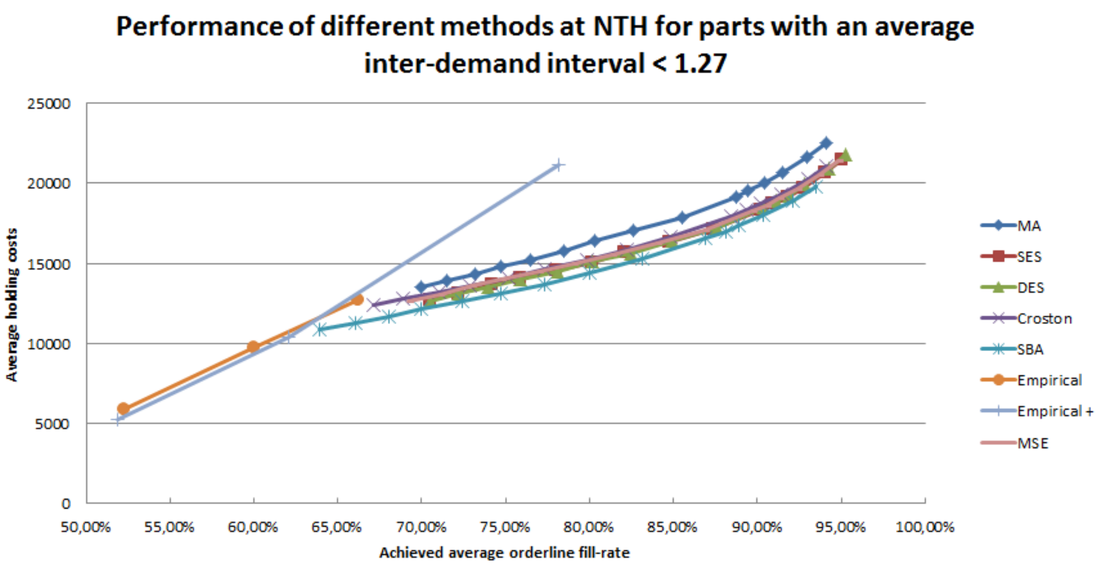

Figure 50: Relation between the achieved order line fill-rate and average holding costs for parts with a small average demand-interval at NTH

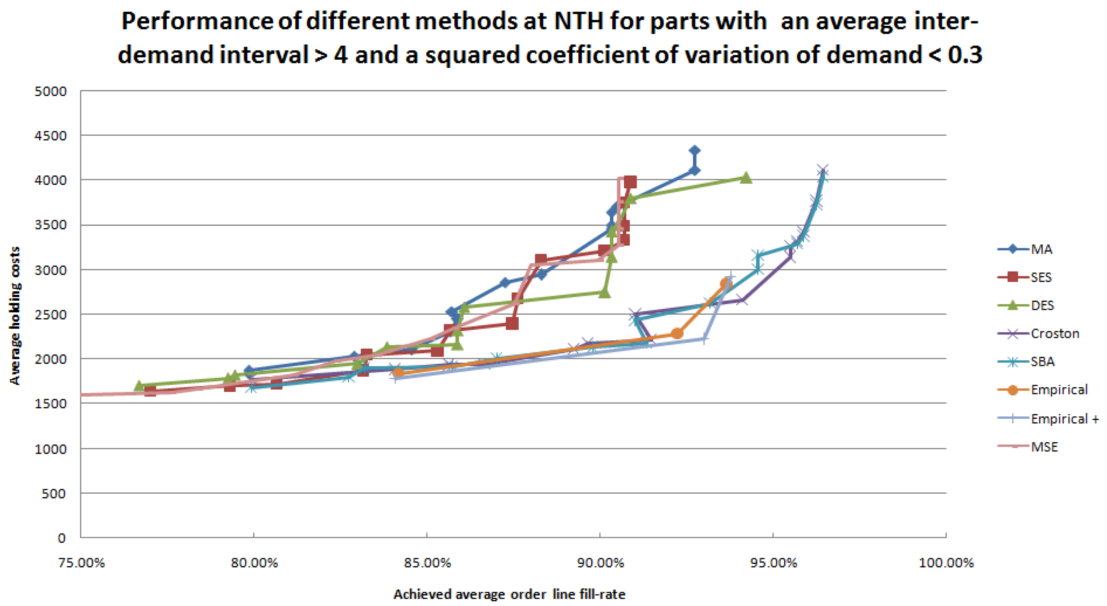

Figure 51: Relation between the achieved order line fill-rate and average holding costs for parts with a $p > 4$ and a $CV^2 < 0.3$ at NTH

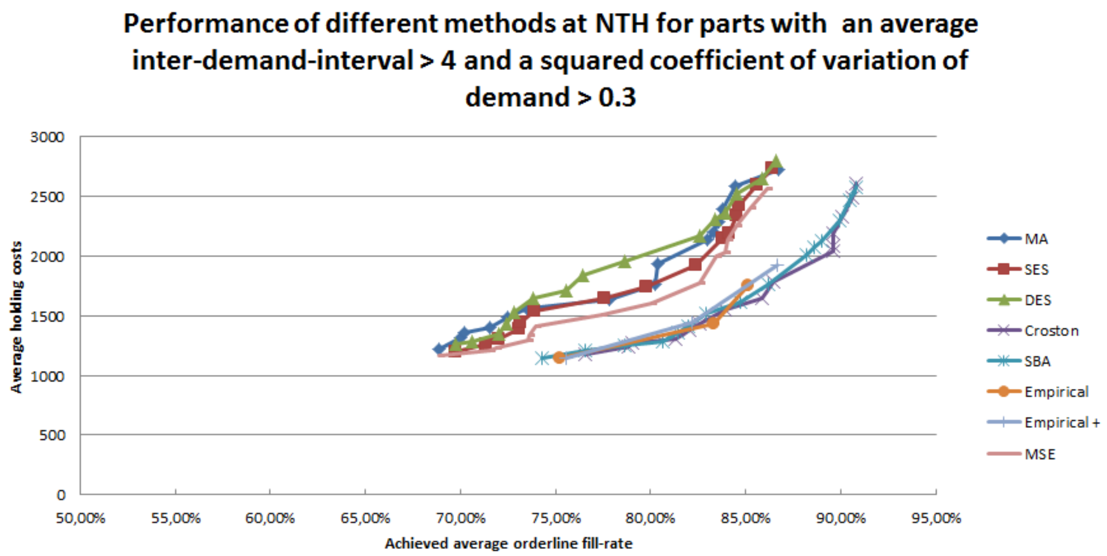

Figure 52: Relation between the achieved order line fill-rate and average holding costs for parts with a $p > 4$ and a $CV^2 > 0.3$ at NTH

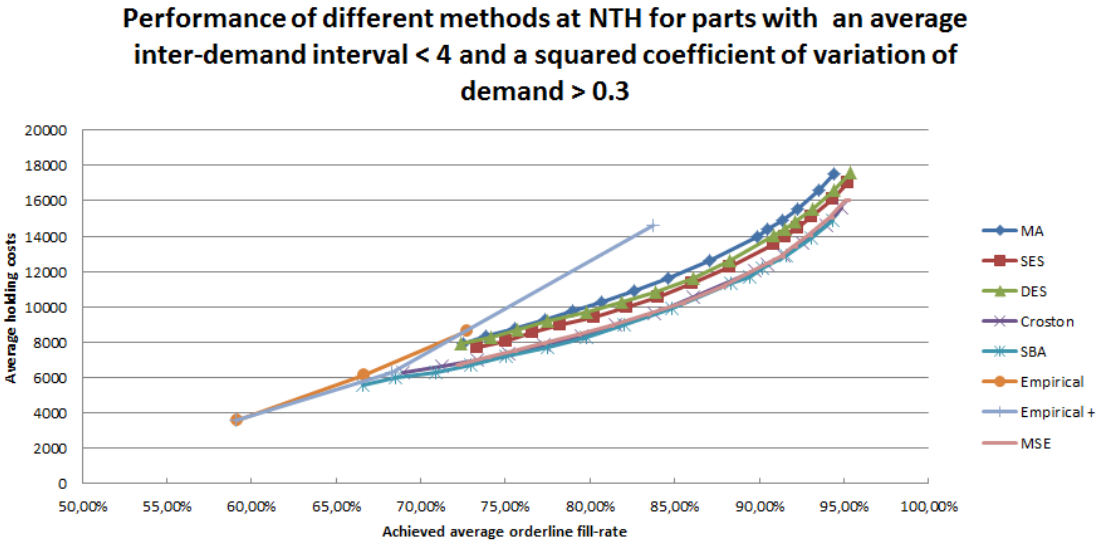

Figure 53: Relation between the achieved order line fill-rate and average holding costs for parts with a $p < 4$ and a $CV^2 > 0.3$ at NTH

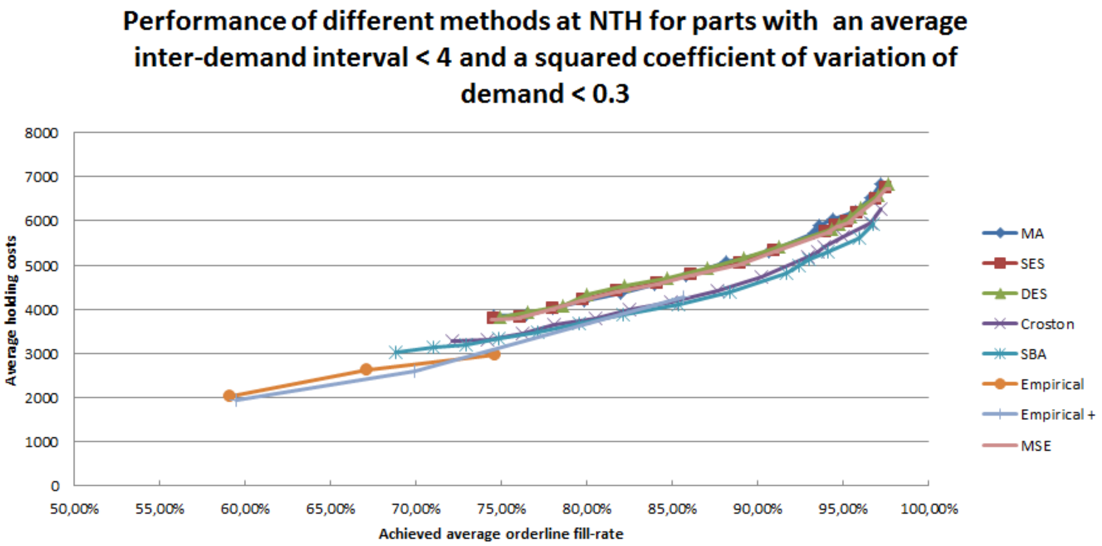

Figure 54: Relation between the achieved order line fill-rate and average holding costs for parts with a $p < 4$ and a $CV^2 > 0.3$ at NTH

D.2 Alstom Ridderkerk

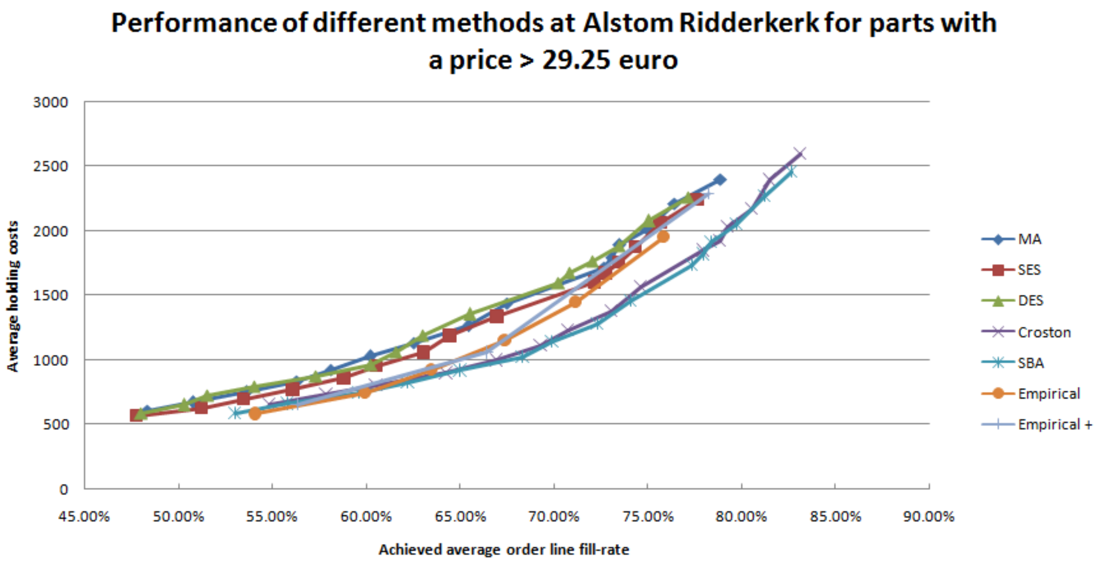

Figure 55: Relation between the achieved order line fill-rate and average holding costs for expensive parts at Alstom Ridderkerk

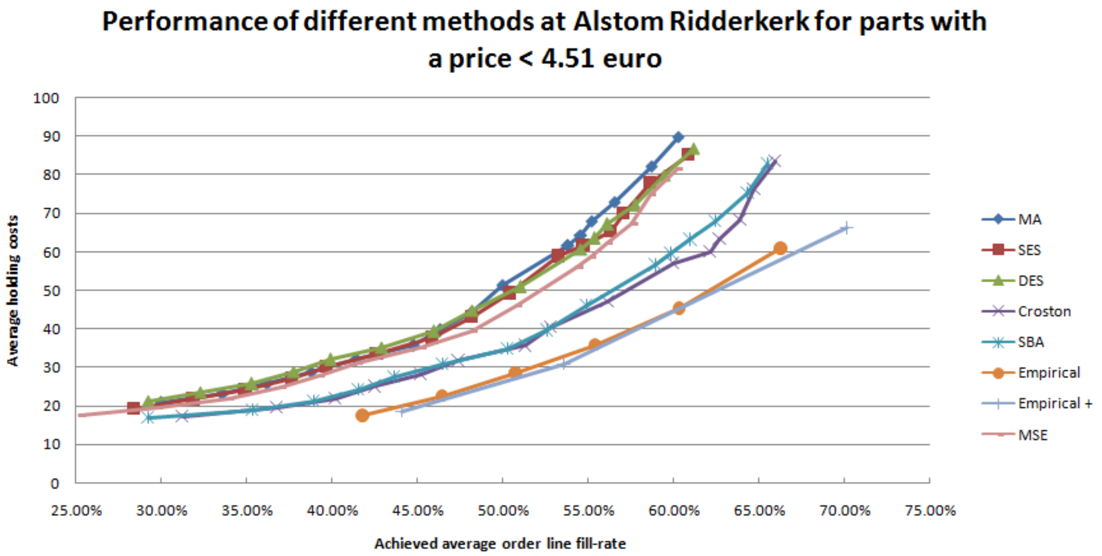

Figure 56: Relation between the achieved order line fill-rate and average holding costs for cheap parts at Alstom Ridderkerk

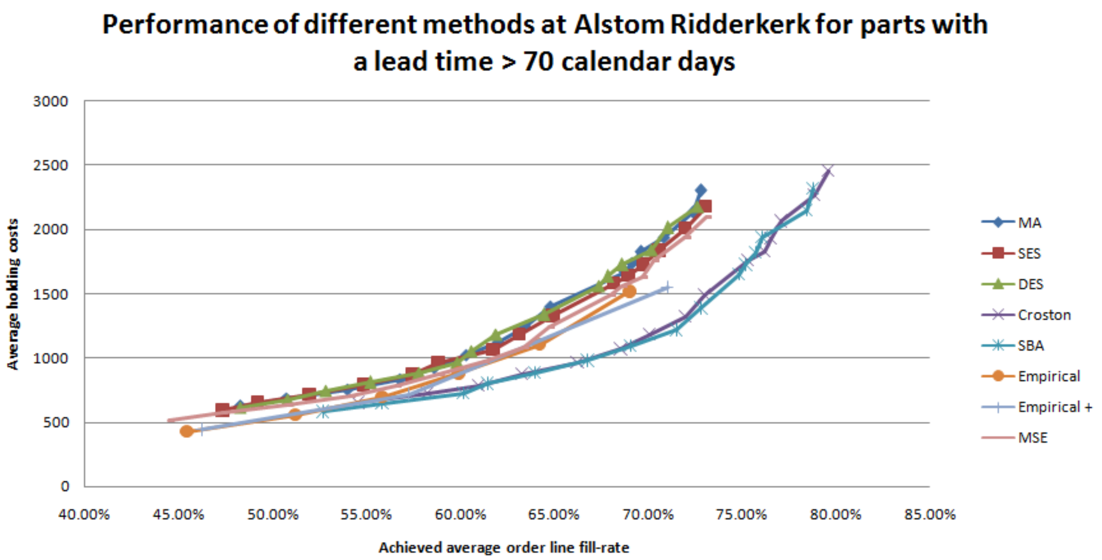

Figure 57: Relation between the achieved order line fill-rate and average holding costs for parts with a long planned lead time at Alstom Ridderkerk

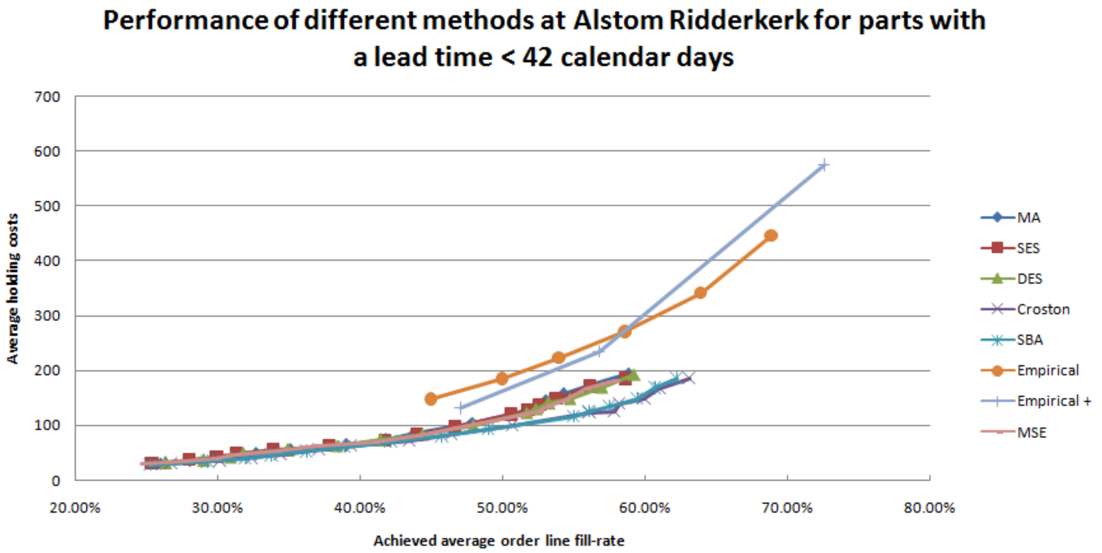

Figure 58: Relation between the achieved order line fill-rate and average holding costs for parts with a short planned lead time at Alstom Ridderkerk

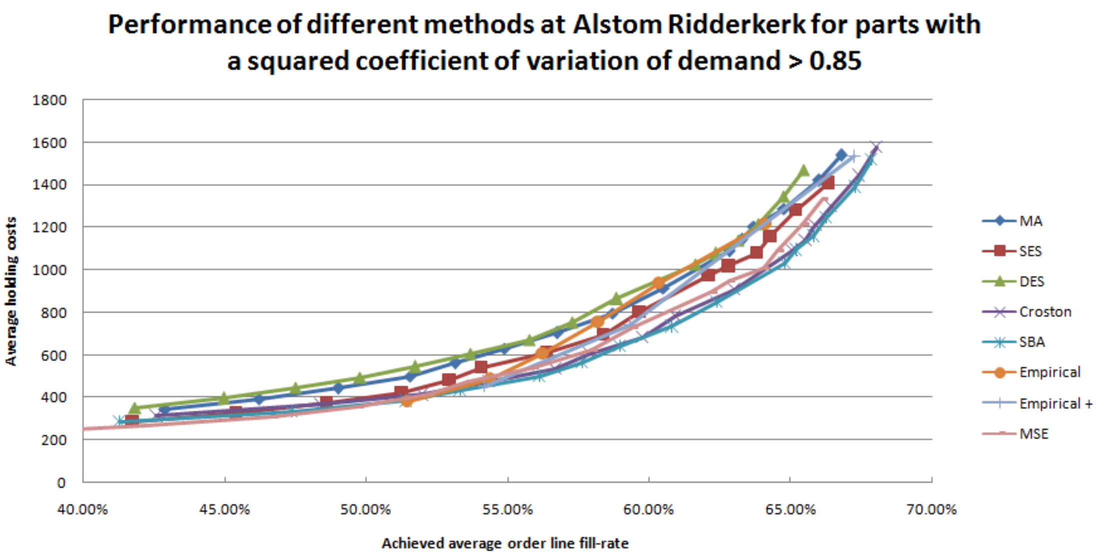

Figure 59: Relation between the achieved order line fill-rate and average holding costs for parts with a large squared coefficient of variation of demand at Alstom Ridderkerk

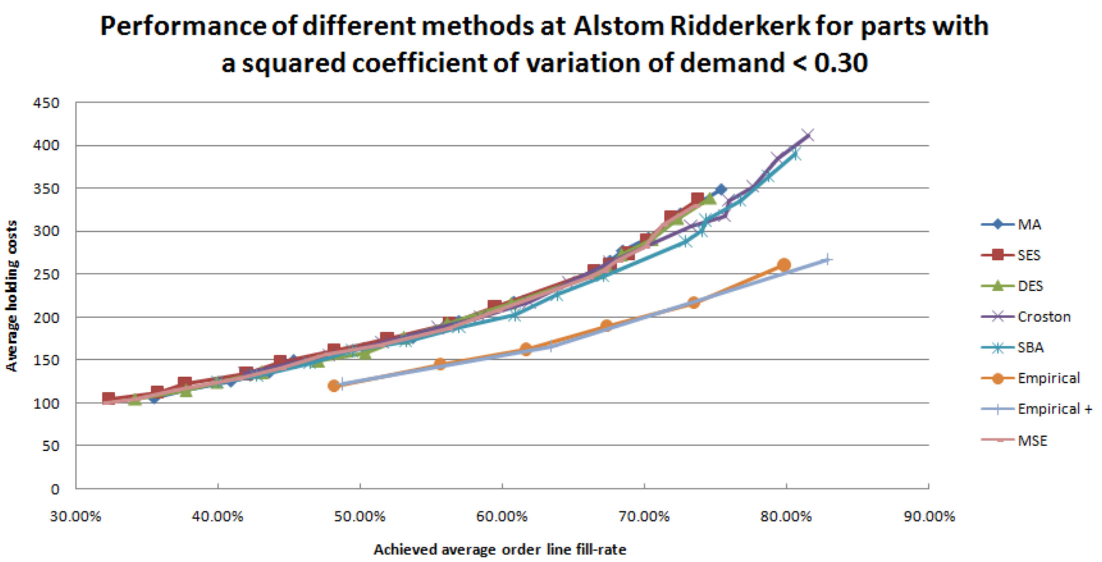

Figure 60: Relation between the achieved order line fill-rate and average holding costs for parts with a small squared coefficient of variation of demand at Alstom Ridderkerk

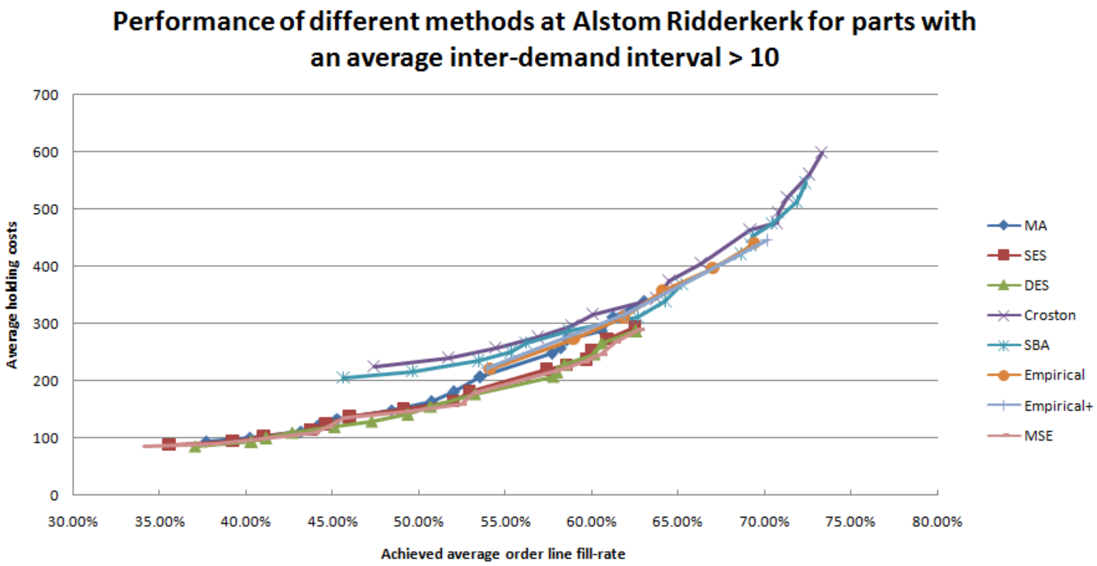

Figure 61: Relation between the achieved order line fill-rate and average holding costs for parts with a large average inter-demand interval at Alstom Ridderkerk

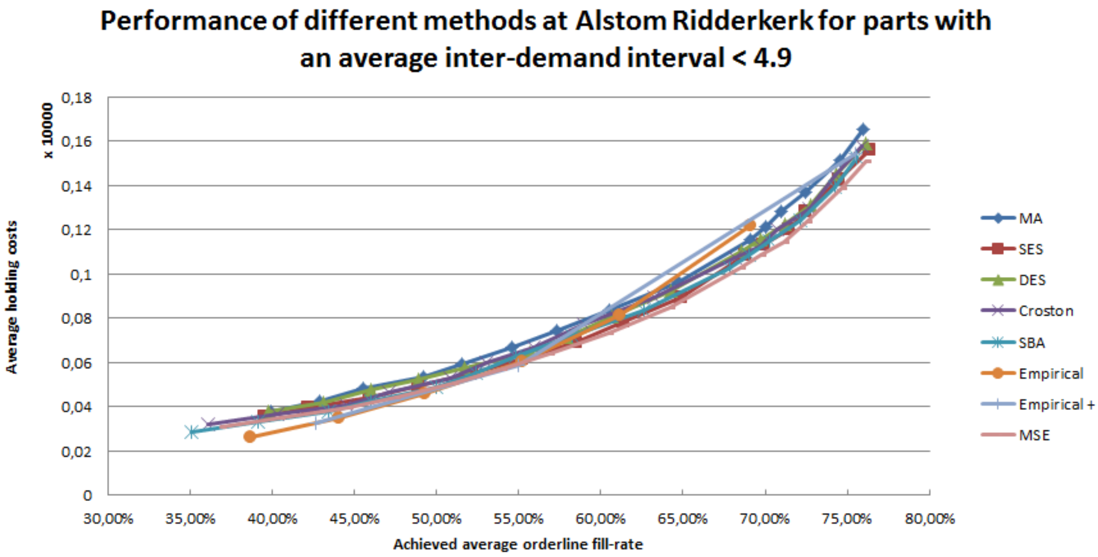

Figure 62: Relation between the achieved order line fill-rate and average holding costs for parts with a small average inter-demand interval at Alstom Ridderkerk

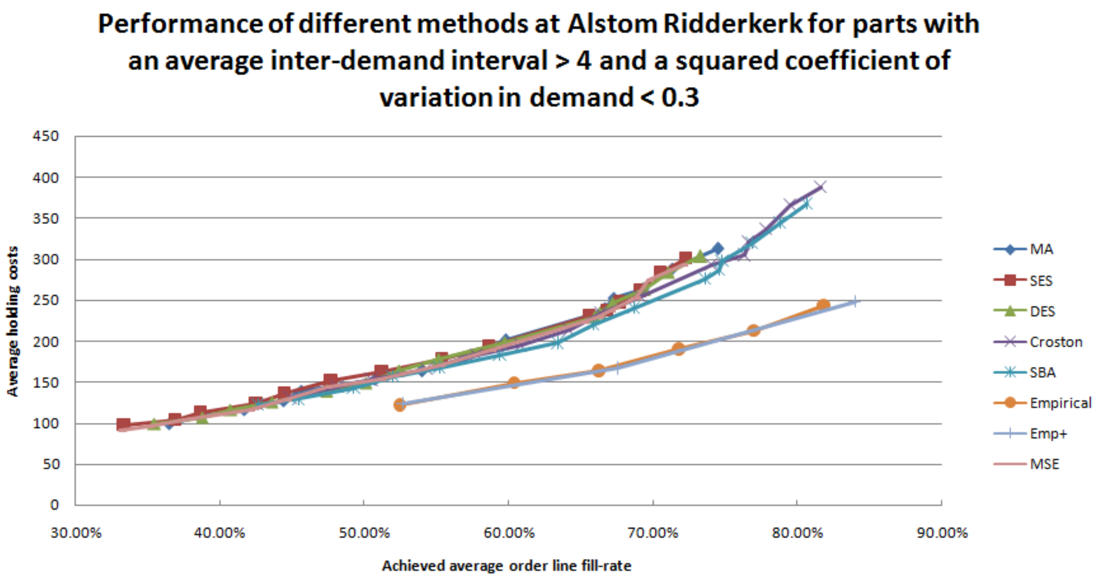

Figure 63: Relation between the achieved order line fill-rate and average holding costs for parts with a $p > 4$ and a $CV^2 < 0.3$ at Alstom Ridderkerk

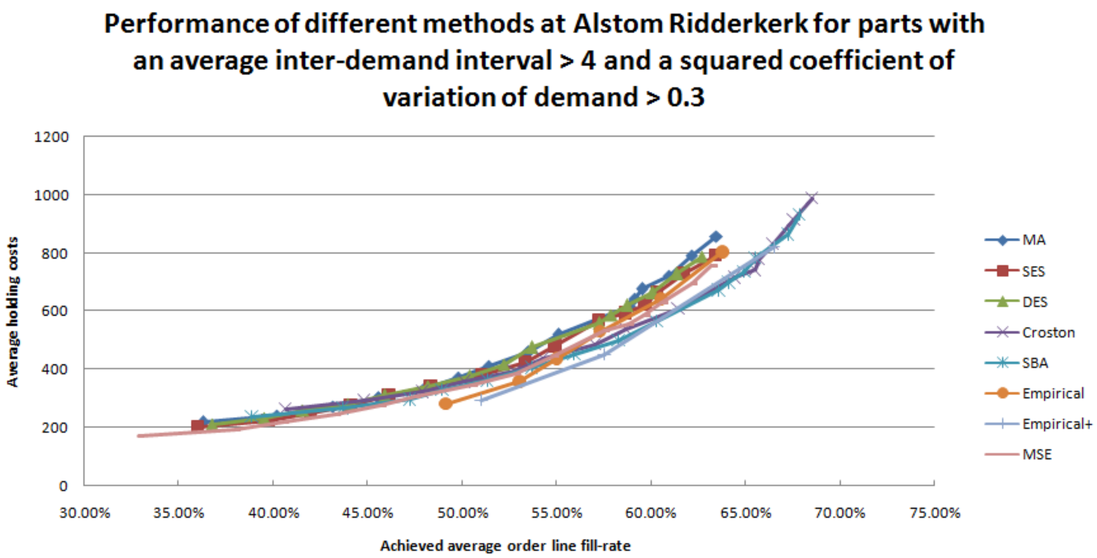

Figure 64: Relation between the achieved order line fill-rate and average holding costs for parts with a $p > 4$ and a $CV^2 > 0.3$ at Alstom Ridderkerk

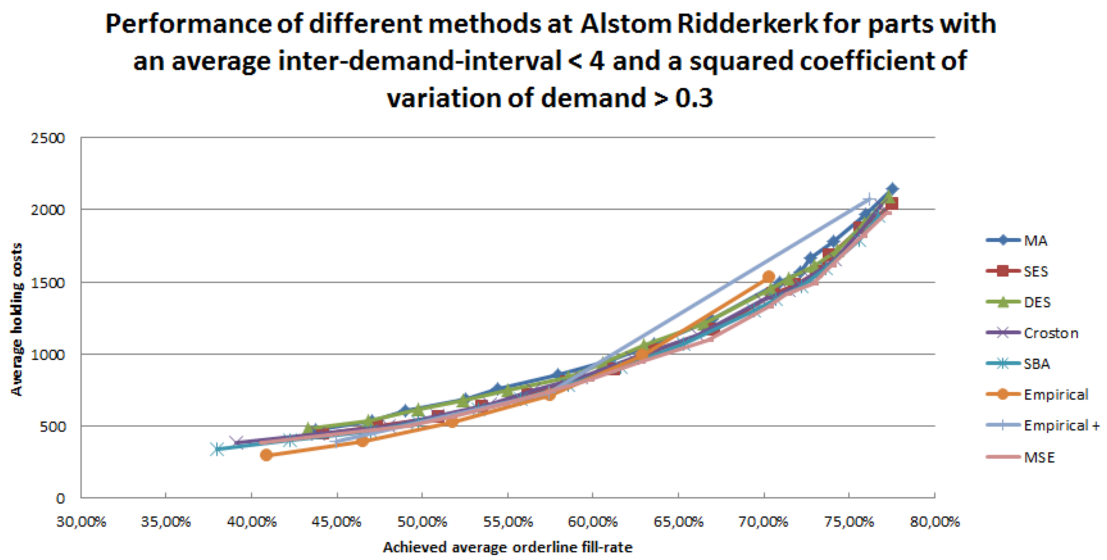

Figure 65: Relation between the achieved order line fill-rate and average holding costs for for parts with a $p < 4$ and a $CV^2 > 0.3$ at Alstom Ridderkerk

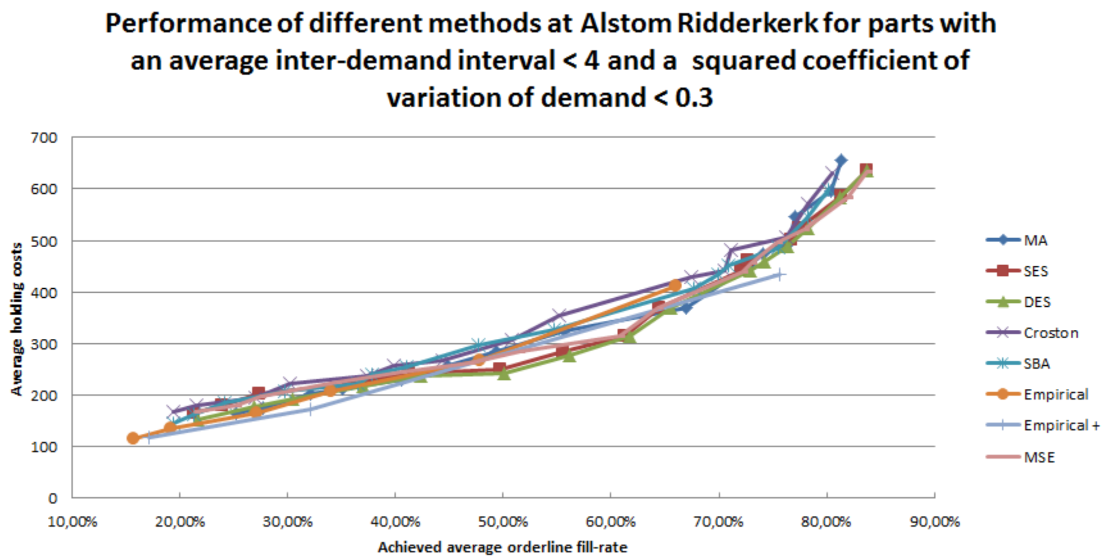

Figure 66: Relation between the achieved order line fill-rate and average holding costs for for parts with a $p < 4$ and a $CV^2 < 0.3$ at Alstom Ridderkerk

D.3 The Naval Maintenance Company

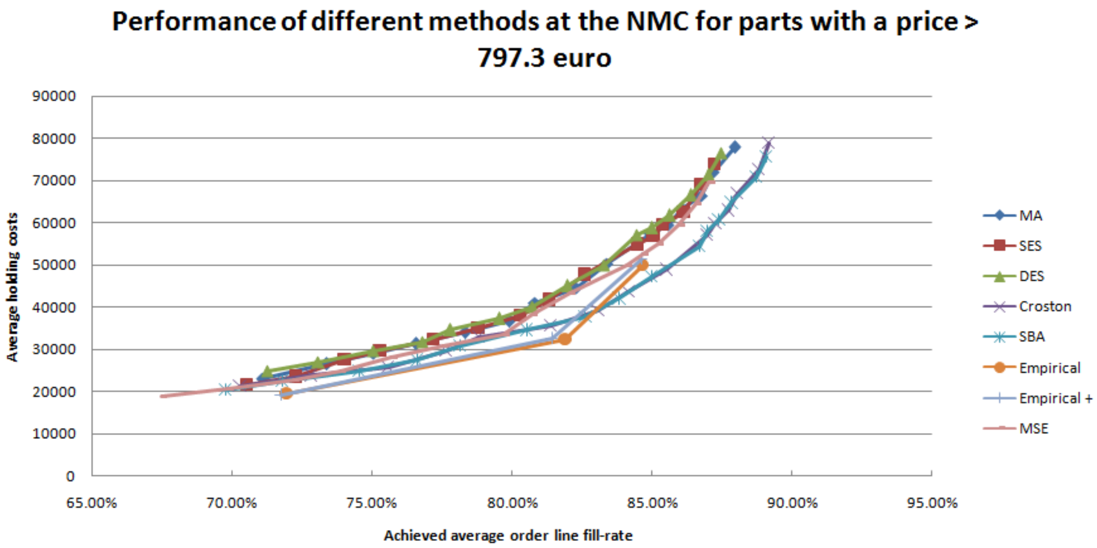

Figure 67: Relation between the achieved order line fill-rate and average holding costs for expensive parts at the NMC

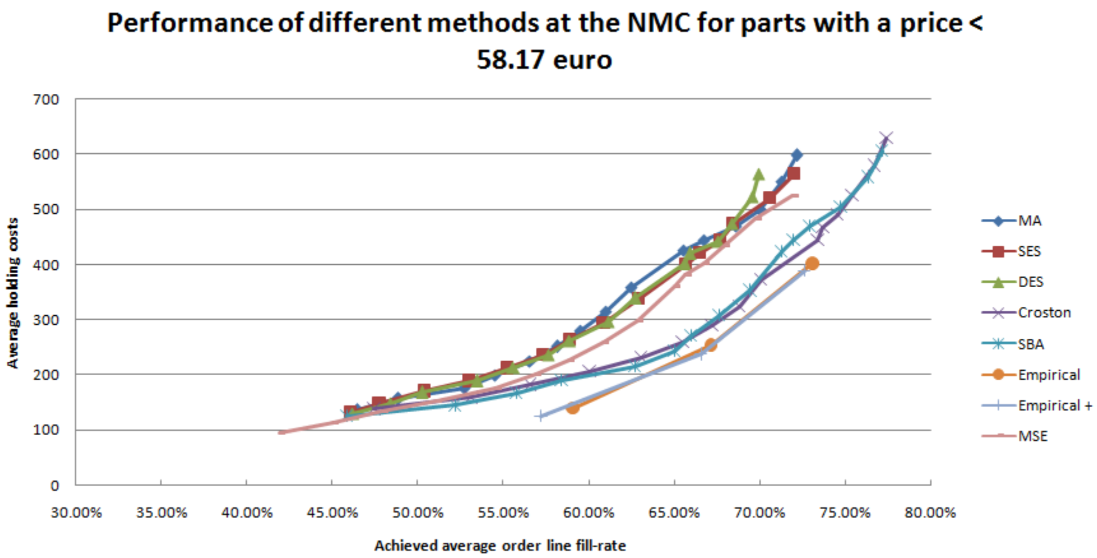

Figure 68: Relation between the achieved order line fill-rate and average holding costs for cheap parts at the NMC

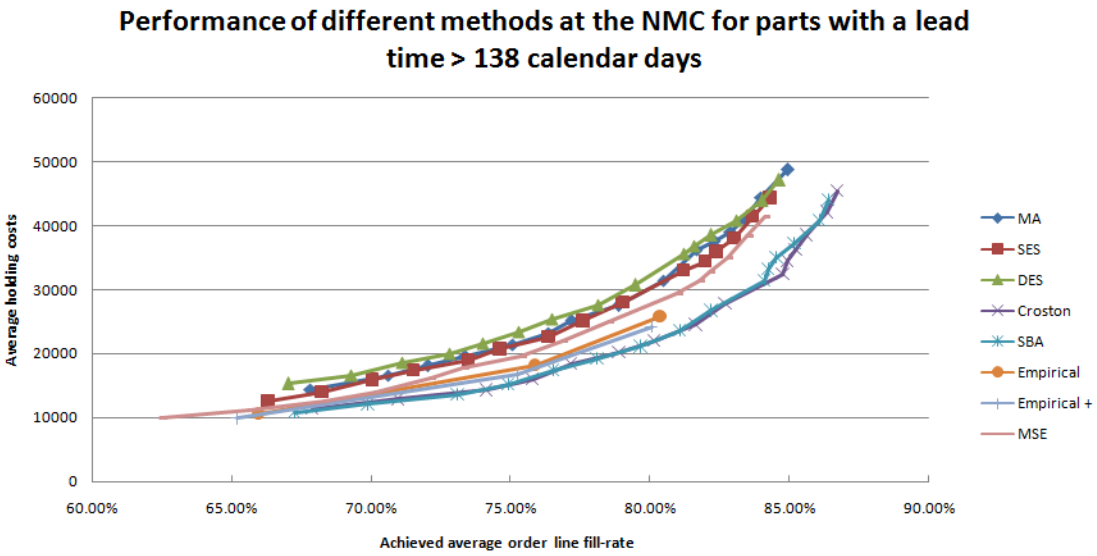

Figure 69: Relation between the achieved order line fill-rate and average holding costs for parts with a long planned lead time at the NMC

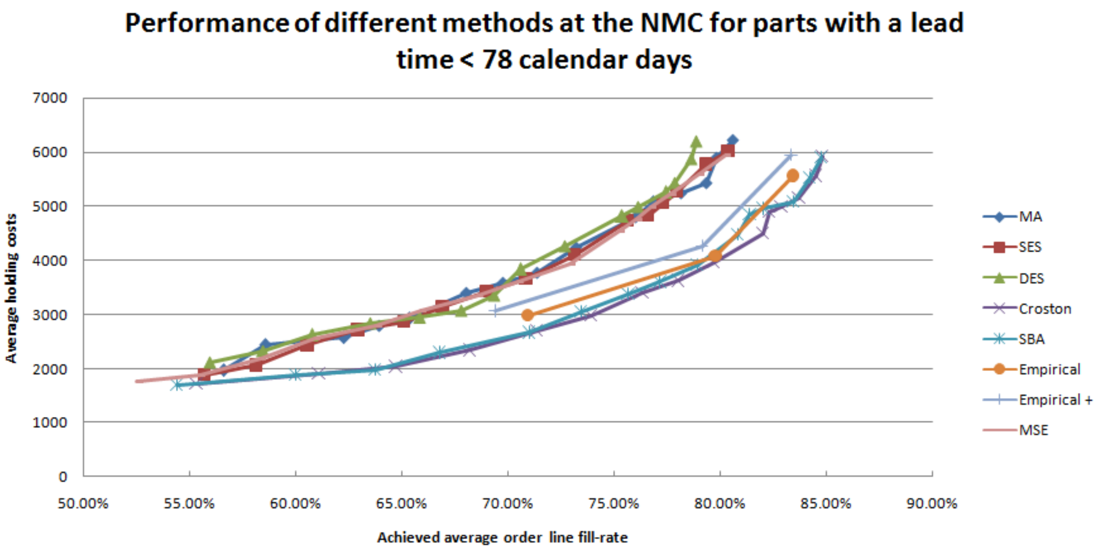

Figure 70: Relation between the achieved order line fill-rate and average holding costs for parts with a short planned lead time at the NMC

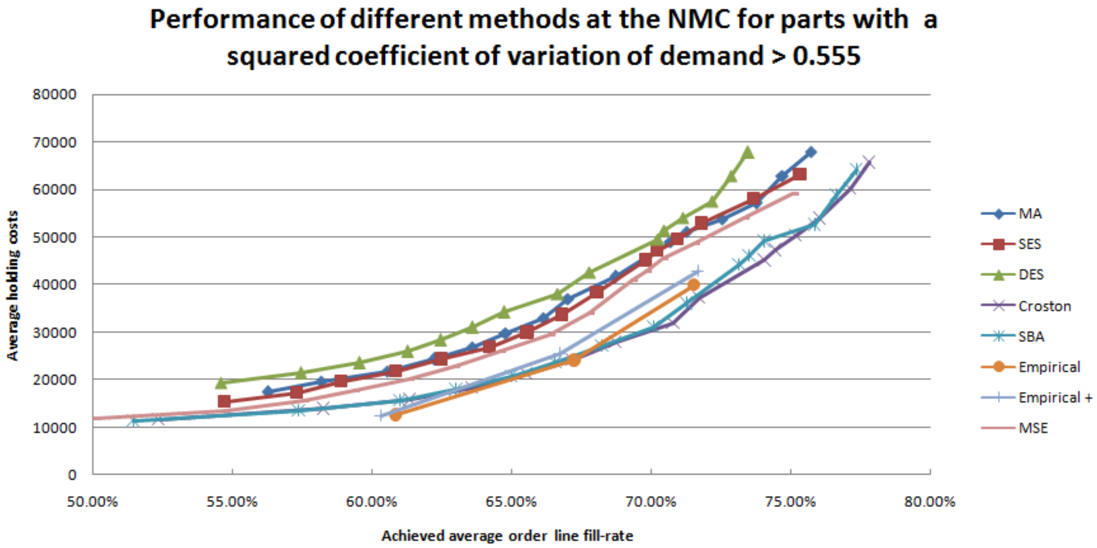

Figure 71: Relation between the achieved order line fill-rate and average holding costs for parts with a large squared coefficient of variation of demand at the NMC

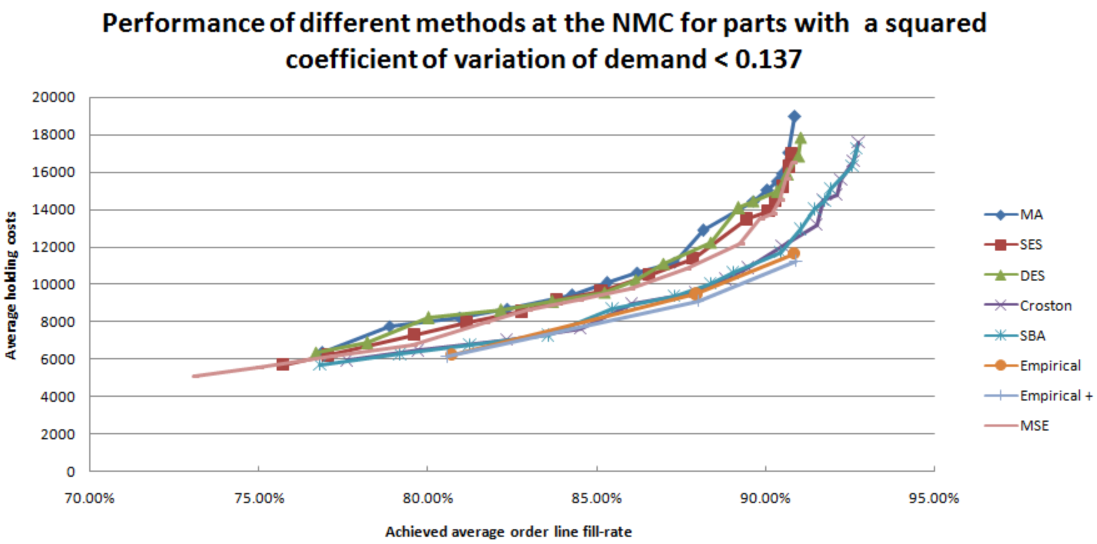

Figure 72: Relation between the achieved order line fill-rate and average holding costs for parts with a low squared coefficient of variation of demand at the NMC

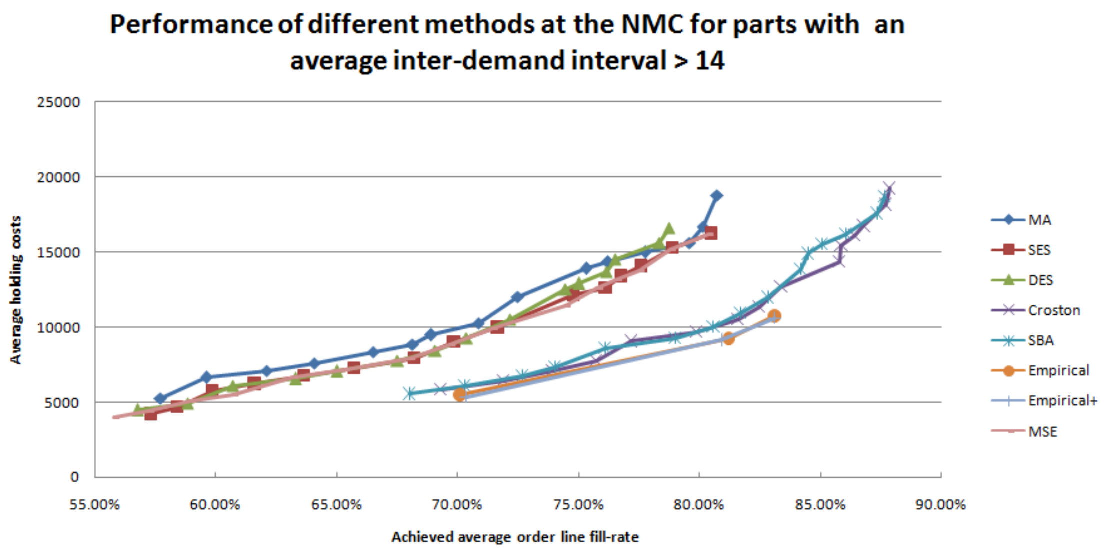

Figure 73: Relation between the achieved order line fill-rate and average holding costs for parts with a long average demand-interval at the NMC

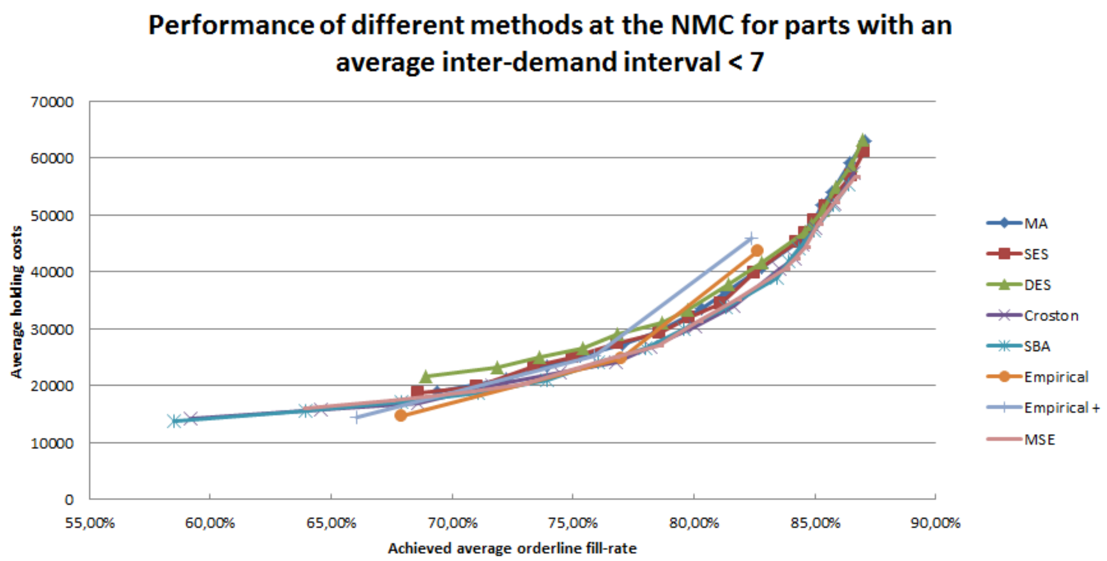

Figure 74: Relation between the achieved order line fill-rate and average holding costs for parts with a short average demand-interval at the NMC

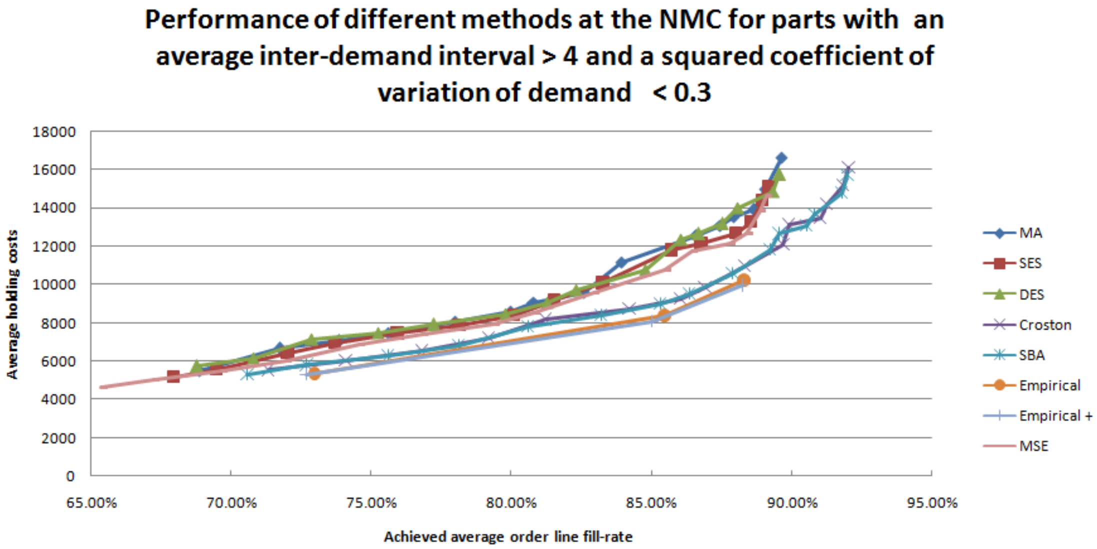

Figure 75: Relation between the achieved order line fill-rate and average holding costs for parts with a $p > 4$ and a $CV^2 < 0.3$ at the NMC

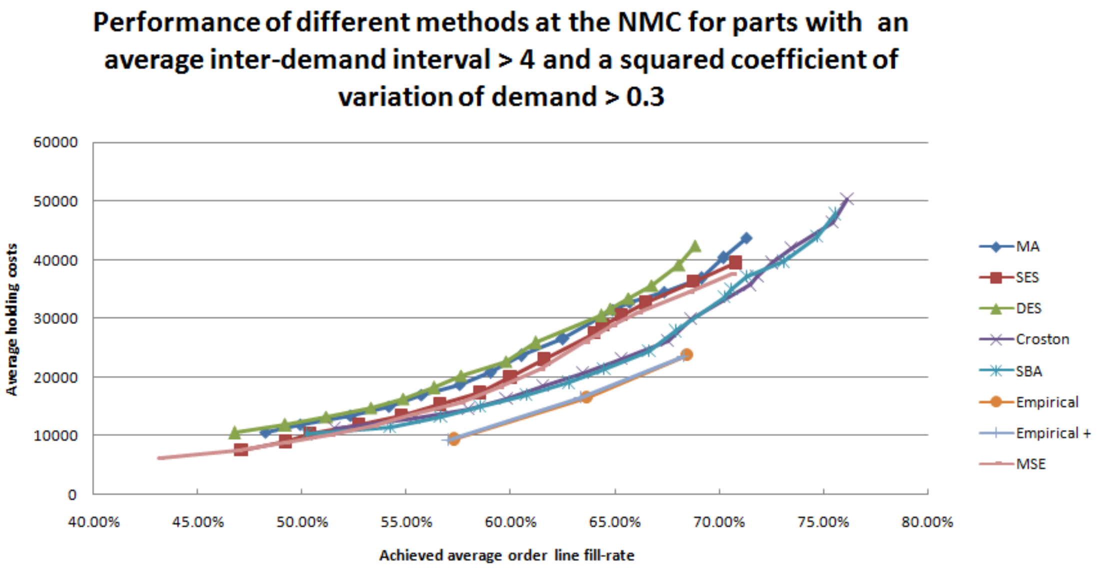

Figure 76: Relation between the achieved order line fill-rate and average holding costs for parts with a $p > 4$ and a $CV^2 > 0.3$ at the NMC

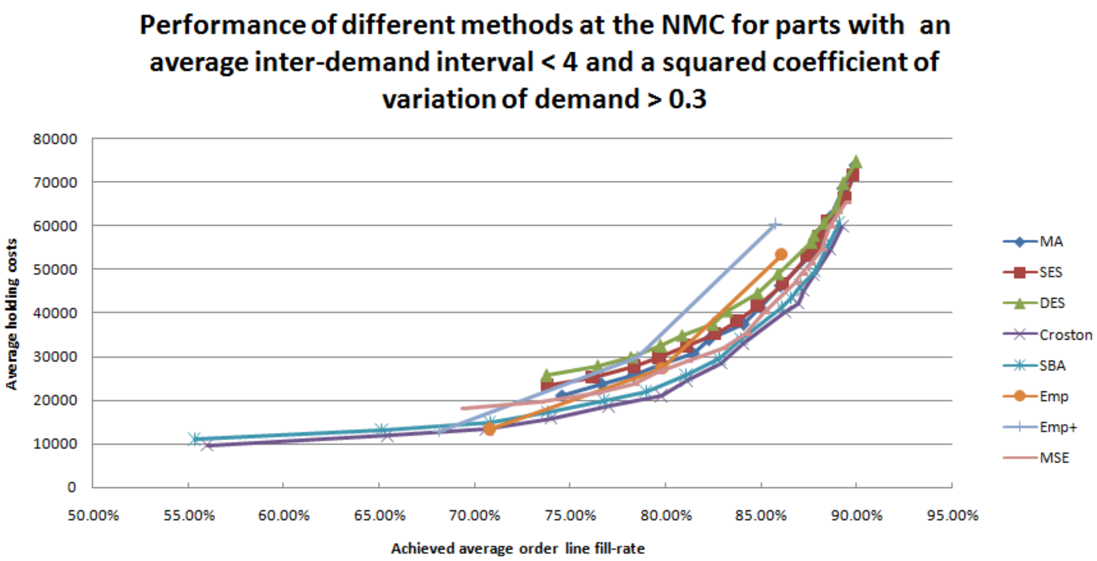

Figure 77: Relation between the achieved order line fill-rate and average holding costs for parts with a $p < 4$ and a $CV^2 > 0.3$ at the NMC

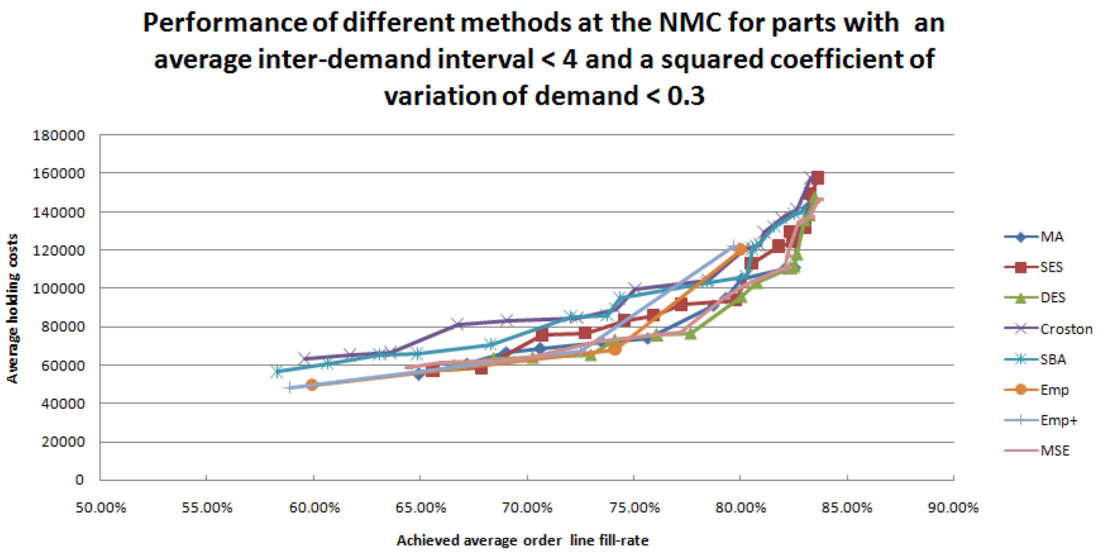

Figure 78: Relation between the achieved order line fill-rate and average holding costs for parts with a $p < 4$ and a $CV^2 < 0.3$ at the NMC