ERASMUS UNIVERSITEIT ROTTERDAM

Operations Research and Quantitative Logistics

Inventory control at the Erasmus MC

Author:
Joris Wagenaar

Supervisors:
Dr. W. van den Heuvel and Dr. T. Farenhorst-Yuan

October 12, 2011

Abstract

Inventory management has never been researched to full extend at the Erasmus MC. Their only concern was having enough inventory to be able to supply the demand. As soon as the department "Bestel en Informatiepunt Logistiek" (BIL) became responsible for the inventory this policy changed. The new policy is to minimize the costs, while being able to supply the demand.

The first topic to investigate is which articles do the hospital have to keep in stock in their distribution center? There are about 35,000 different articles used in a year and there is only room to store 2,000 articles. We formulated a model to decide this. This model is based on the optimal order size Q, the average demand per order, the average demand per year, the lead time and the risk-class of an article. The model concludes that 54 articles that are currently kept in stock should no longer be kept in stock. It also concludes that 68 articles that are not stored at the moment, should be stored.

The Erasmus MC expects the DC to maintain a service level of at least 98% for the articles they keep in stock. The service level is the percentage of demand that is directly satisfied by the DC. The service level depends on the order size Q and the reorder point R. We minimize the corresponding inventory costs of the 2,000 stored articles, while maintaining a service level of at least 98%.

The 2,000 articles should also be assigned to a location within the DC. Articles that are ordered a lot should be stored in the front of the warehouse. Articles that ordered often together should be stored close to each other. That is why in chapter five a model is proposed that clusters articles with a high similarity. The cluster with the highest popularity is stored in the front of the DC, while the cluster with the lowest popularity is stored in the back of the DC.

This thesis is dedicated to my father, who passed away on February 25, 2007.

Acknowledgements

This thesis has been written in order to finish the master program Operations Research and Quantitative Logistics at the Erasmus University Rotterdam. It is a result of a six months during internship at the Erasmus MC, where I worked at the department "Bestel en Informatiepunt Logistiek". I would like to thank the persons that helped and supported me with this thesis.

First of all I would like to thank Willeke Stekelenburg who arranged an internship at the Erasmus MC and gave me advise about the insights at the Erasmus MC. I would also like to thank the rest of the colleagues at the Erasmus MC for providing a nice environment to work in.

Second I would like to thank my supervisors dr. Van den Heuvel and dr. Farenhorst-Yuan, their suggestions helped me improving my thesis. They also gave me enough space to set up my own research and supported me during the whole process.

My final thanks goes to Irma Tempelaar and Nathalie van Zeijl for reading my thesis and giving useful comments.

CONTENTS

1. INTRODUCTION

The department "Bestel en Informatiepunt Logistiek" (BIL) is a department of the Erasmus Medical Center (the main hospital of Rotterdam). Last year a new company executing system has been introduced at the Erasmus MC, namely Oracle eBS. With the introduction of this new system the BIL became responsible for the inventory management within the Erasmus MC. There has never been a research on the inventory policies at the Erasmus MC. This is mainly due to the fact that hospitals, also the Erasmus MC, did not see inventory management as there core business. As long as there was enough inventory to fulfill the demand, they were not really interested in it. But now BIL has the responsibility, they want to take a closer look at the inventory policy and to minimize the corresponding costs.

The Erasmus MC keeps a central stock at a distribution center in Barendrecht and a decentral stock at the hospital itself. The decentral stock is kept at circa 120 different departments at the hospital. Their content is different at every department and depends on the articles that are used at that department. For example the department gynaecology needs different articles than the department neurology.

The whole hospital uses around the 35,000 different article types per year, the distribution center (DC) currently keeps around 2,000 of those 35,000 article types in stock. During the remainder of this thesis the definition of an article is an article type. An article exist out of a number of products, for example gloves is an article and there are more than 1,000 gloves (products) stored in the DC.

When a department needs an article, they first try to satisfy that demand from their own inventory. If the article is out of stock at the department, or if the article is not kept in inventory at all at the department, they place an order for that article at BIL. When they do supply from their own inventory, it is possible that the inventory drops to a certain reorder point and then they also place an order at BIL.

When BIL receives an order from a department, they place that order either at the DC or at a supplier. Supplying from the central stock is of course a lot faster than supplying from a supplier, so when possible BIL always places the order at the DC. But it is possible that an article is not kept in stock at Barendrecht, or that the article is out of stock. In that case BIL places the order at a supplier.

The DC delivers articles at most departments ones a day, the only exception are the operating rooms. The articles are delivered twice a day there. So the lead time from the DC to the hospital is assumed to be one day.

The stock on hand per article is checked three times a week at the DC and when the stock is below a certain reorder value the DC places an order at BIL to refill their stock up to its maximum level. BIL places the order at one of the suppliers and the lead time from the suppliers is assumed to be five days, this will be explained in the next chapter.

There are three segments of importance to the Erasmus MC. First of all they want to know which articles they have to keep in stock and for which articles is it better to directly order them at a supplier? The lead time from the DC to the hospital is a lot shorter than the lead time from a supplier to the hospital. There are certain articles for which it is absolutely necessarily that they are delivered within one day, because those articles can make the difference between life or death of a patient. Those articles always have to be kept in stock at the DC. For the other articles it has to be decided whether it is profitable to keep them in stock or not. Both keeping articles in stock as placing orders at a supplier has certain costs. So for every article it has to be decided whether it is cheaper to store the article or to direct order the article. This research will be done in chapter 3.

The second important segment is to decide, for every article that has to be stored, when to reorder and how many products to reorder. When the reorder is made every time the article is out of stock, the holding costs will be low but also the service level will be low. The service level is the percentage of demand that is directly satisfied by the DC. If there is a reorder every time there still is a lot of stock on hand, then both the costs and service level will be high. The balance between these two factors is used in chapter 4 to determine the optimal reorder point.

The last important segment is to allocate all the stored articles to a location within the DC. Articles that are stored in the front of the DC are much easier to pick up fast than articles in the back of the DC. So articles that are often demanded by the hospital need to be allocated at the front of the DC. Also articles which are often demanded together need to be allocated near to each other in order to be able to pick them easily together. This is discussed further in chapter 5.

But first an analysis on the data will be done in the next chapter.

2. DATA ANALYSIS

In this chapter an analysis on the data, given by the Erasmus MC, is done. We will first give a short description of the data, after that we will determine the arrival and order size distribution per article.

2.1 General data description

The Erasmus MC saved every order made in the past 14 months. When a department places an order the following data is known about that order:

- The date and the time of the order
- The department who placed the order
- Either the article number or the article supplier number
- The number of ordered products per article demanded
- The price of the ordered article(s)
- The supplier (if applicable)

The range of the data is from 1 March 2010 up to 1 May 2011. In this time there are more than 500,000 orders placed. As mentioned before, an article can be ordered at the DC or directly at the supplier. Both orders are saved and an article number is recorded when the article is ordered at the DC and an article supplier number is recorded when the article is directly ordered at a supplier. The article (supplier) number is the identifier of the article, without this number it is not known which article is ordered.

The article supplier number is not always recorded. Most of the time that means that a service is demanded instead of an article, for example the restoration of a X-ray machine. But there is a problem with the data for around 5,000 orders. In those orders an article is demanded instead of a service, but it is unknown which article is demanded because there is no identifier. These orders are omitted during the rest of the study.

Orders placed for the same article at the same department within one hour of each other are aggregated, because of possible order mistakes by the department. This is also done by the BIL department when an actual order comes in.

We will first give a short general description of the data:

- There are around 35,000 different articles ordered at the Erasmus MC the past year.
- The article that is ordered most in the last 14 months is article 37106, in total 968 times.
- 123 of the current articles in stock are only demand once.
- 22,281 of the current articles that are not in stock are only demanded once.
- There are 882 deliver addresses.
- Deliver address AB-H-1127 placed the most orders, namely 13,324 times.
- 148 deliver addresses have only placed one order.
- The most expensive article in the DC is article 1047056 and costs 1,238.12 euro.
- The most expensive article in total is article 2010/S 93-139388 and costs 895,000 euro.
- There are 14 articles with price zero, the hospital can get these articles for free.
- Currently 2,086 different articles stored in the DC.
- The order costs are equal to €25.68 per order according to Zimmerman [2010].
- The holding costs are equal to 25% of the price of an article per year.

In figure 2.1 the average demand per year per article currently stored in the DC is shown. On the horizontal axis the articles are allocated and on the vertical axis the average demand per year is allocated. As can be seen the average demand per year differs between 1 and 150,000. So there is a big difference between the demand of the different articles.

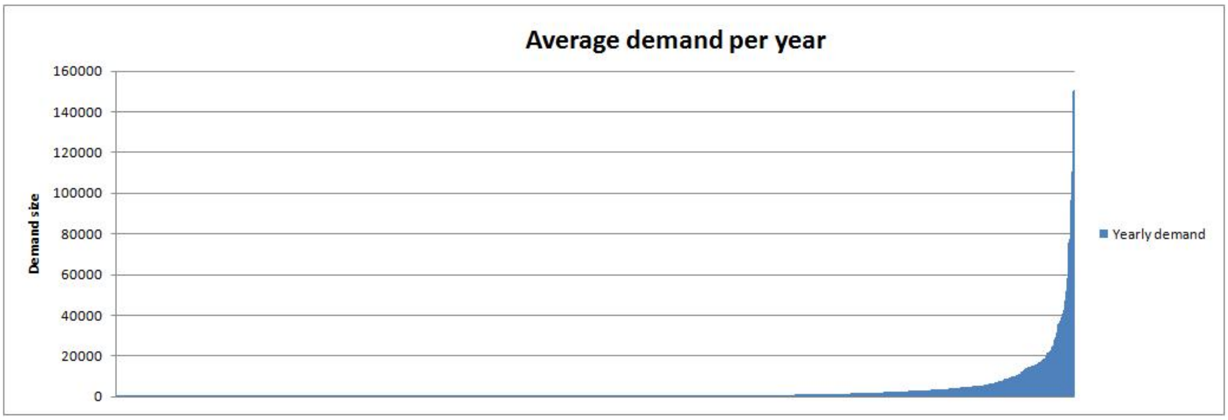

Fig. 2.1: The average yearly demand per article

2.2 Estimating the distribution of the demand

At the moment 2,086 articles are being kept in stock at the DC. Before we are able to improve the inventory processes within the Erasmus MC, we first need to have information about the demand arrival and the order size process per article.

The demand arrival process is the process of the time between order arrivals. It is most likely that the time between order arrivals fluctuates. So the distribution of the time between order arrivals has to be estimated per article.

The order size process is the process of the demand size per order arrival. It is also most likely that the demand size fluctuates per order, for example the hospital needs hundred gloves today, but only sixty gloves tomorrow . So also the distribution of the order size has to be estimated per article.

The processes will be tested on a hypothesized distribution in the following sections with the Anderson Darling test.

2.2.1 Anderson Darling test

The Anderson Darling (AD) test is a statistical test, it tests whether a sample has been drawn from a specified cumulative distribution function or not (see Anderson and Darling [1951]). The test is a modification of the Kolmogorov-Smirnov (K-S) goodness of fit test (see Kolmogorov [1933] and Smirnov [1948]). The AD test makes use of the specific distribution to calculate the critical values, this has the advantages of a more sensitive test. That is why we decided to use the AD test instead of the K-S test. The test statistic is defined as follows:

$$A^2 = -N - S \tag{2.1}$$

where:

$$S = \sum_{i=1}^{N} \frac{2i-1}{N}[\ln F(Y_i) + \ln(1 - F(Y_{N+1-i}))] \tag{2.2}$$

and N is the sample size, Y_i are the data ordered from small to large and F is the cumulative distribution function of the hypothesized distribution. The test statistic is based on the distance between the hypothesized distribution and the distribution of the data and then calculating the probability of obtaining data which have an even larger value than the value observed, assuming the hypothesized distribution is true.

The hypothesis that the data follows the specified distribution is rejected when the test statistic is greater than the given critical value for a chosen significance level α.

2.2.2 Arrival process

The demand arrival process is expected to follow a Poisson process. A Poisson process is a collection of random variables $N(t) : t \geq 0$, where $N(t)$ are the number of arrivals up to time t. The number of arrivals between time a and b, $N(b) - N(a)$, has a Poisson distribution. The Poisson process possesses a few properties:

- $N(0) = 0$, at the start there are no arrivals.
- The arrivals are independent of each other. This means that an arrival now has no influence on when the next arrival will be.
- The arrivals are stationary, this means that the number of arrivals only depend on the length of the interval.
- There are no simultaneous arrivals.

According to the properties of the Poisson process, the inter arrival times have to be Exponentially distributed and independent of another. So we have to test whether the interarrival times between the demand of an article are Exponentially distributed and independent of another in order to test whether the arrival process follow a Poisson process.

There are two types of Poisson arrival processes, a homogeneous and a non-homogeneous Poisson arrival process. With a homogeneous Poisson process, the intensity parameter λ is constant over time. The intensity parameter λ is equal to the expected number of arrivals per unit of time. When the rate λ changes over time the Poisson process is called non-homogeneous.

At the Erasmus MC it could be possible that some articles follow a homogeneous Poisson process and other articles follow a non homogeneous Poisson process. For example catheters likely follow a homogeneous Poisson process, while cooling agents are more likely to be demanded more during the summer than during the winter. So first every article will be tested on whether the article follows a homogeneous Poisson process or not (if not, that does not automatically mean that the article follows a non homogeneous Poisson process). After that, the remaining articles will be tested whether they follow a non-homogeneous Poisson process.

It is only possible to test on articles which are demanded more than seven times in total (see Lewis [1961], Stephens [1970]). There are 275 articles current in stock which are demanded less than eight times, these articles are omitted in the tests and are assumed to be homogeneous Poisson distributed.

There are 1811 articles left which are first tested on a homogeneous Poisson process. That will be done by testing on Exponentiality of the interarrival times with the AD test. As can be seen in Table 2.1 1,370 articles are homogeneous Poisson distributed according to the test.

Number of articles	Number of homogeneous articles	Percentage
1811	1370	75.65

Tab. 2.1: Testing on Exponentiality of the inter-arrival times

The remaining 441 articles are tested on a non-homogeneous Poisson process. That means that the arrival rate could fluctuate in time. We will test whether the arrival rate is different per quarter (so per three months), this is because there is not enough data to be able to test the arrival rate per month. First the data is splitted into five different datasets, the first four datasets contain data per three months and the last dataset contains the data of the last two months. Then per dataset, per article the inter arrival times are tested on the Exponential distribution with the AD test. In Table 2.2 can be seen that 72.74% of the remaining articles follow a non-homogeneous Poisson process according to the test. We assume the other 27.26% articles also to be non-homogeneous Poisson distributed. That are about 120 articles.

Number of tests	Number of successful tests	Percentage
2205	1604	72.74

Tab. 2.2: Testing on Exponentiality of the inter-arrival times

We conclude that the order arrivals do follow a Poisson process, either a homogeneous Poisson process or a non-homogeneous Poisson process.

2.2.3 Demand size

Every time an order for an article arrives, this order has a certain order size. It could be possible that the order has the same size every time, or that the size differs per order. There are three different Poisson processes, with respect to the demand size, and throughout this article the following definitions will be used for these different Poisson processes:

- When the order always has an order size equal to 1 the total process is called a pure Poisson process.
- When the order size is always the same, but larger than 1, the process is called a constant Poisson process.
- When the order size differs per order, then the total process is called a compound Poisson process.

If an article follows a compound Poisson demand process, it is unknown how large the demand will be every time an order arrives, so in order to work with it we need to determine the underlying distribution.

For the pure and constant Poisson articles no distribution has to be determined, because every time an order arrives the demand is of the same size and this order size is known. So the first thing to do is to find all the articles with a pure or constant Poisson demand process. There are also articles that are only ordered once in the past 1.5 year. These articles are assumed to be pure or constant Poisson distributed as well, depending on the size of the one order. This is done because there is no further data available about them, and until now the order size is constant. Of the 2,086 articles, 212 articles follow a pure or constant Poisson process.

For the remaining 1,874 articles an underlying distribution has to be estimated. As explained before, there have to be at least eight demand arrivals before it is possible to test on a certain distribution with the AD test. There are 149 articles ordered less than eight times. When the index of dispersion (see Darwin [1957]) is below one it is allowed to assume that the demand size of the articles follow a Multinomial distribution. The index of dispersion is the ratio between the variance and the average demand size of an article:

$$ratio = \frac{\sigma^2}{\mu} \tag{2.3}$$

Where σ^2 is the variance and μ is the average demand size.

All the 149 articles have an index of dispersion below one as can be seen in figure 2.2. This means that we can assume a Multinomial distribution for the 149 articles which are ordered less than eight times.

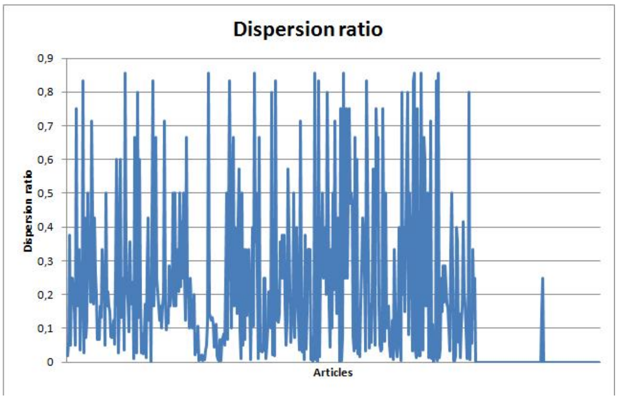

Fig. 2.2: The index of dispersion per article.

For the remaining 1725 articles it is unknown which distribution the order size has. So the AD test will be done on the most common demand size distributions: the Normal, Gamma, Weibull and the Uniform distribution. When an article follows a homogeneous Poisson arrival process, it is not certain that the article also has a homogeneous demand size distribution. So the first thing to test is which articles do have a homogeneous Normal, Gamma, Weibull or Uniform order size distribution. In table 2.3 a summary of the Anderson-Darling tests on a homogeneous distribution are given.

In total, 1,282 articles follow the homogeneous Gamma distribution. We test the remaining 443 articles only on a non-homogeneous Gamma distribution, because a negligible number of articles do follow these distributions.

The data is splitted for those 443 articles in four datasets per three months and one dataset with two months of data. These datasets will be tested on the Gamma distribution and this results in $443 * 5 = 2,215$ different tests.

Distribution	Number of tests	Number of successful tests	Percentage
Normal	1,725	15	0.94%
Gamma	1,725	1282	74.32%
Weibull	1,725	37	2.31%
Uniform	1,725	40	2.50%

Tab. 2.3: Testing on homogeneous order sizes

As can be seen in table 2.4 83.79% of the remaining 443 articles follow a non-homogeneous Gamma distribution, we assume that the other 16.21% also follow a non-homogeneous Gamma distribution.

Distribution	Number of tests	Number of successful tests	Percentage
Gamma	2,215	1,856	83.79%

Tab. 2.4: Testing on non-homogeneous Gamma distribution

We conclude that the order size distribution is either Multinomial, for the articles that are demanded less than eight times, or Gamma for the other articles.

2.3 Summary

In the following tables a short summary of the AD tests on the arrival and order size distribution are given:

Goal	Distribution	% homogeneous	% non homogeneous	% assumed	% small
Arrival process	Poisson	65.68	15.44	5.70	13.18

Tab. 2.5: Summary of demand arrival testing

Goal	Distribution	% Pure and constant	% homogeneous	% non homogeneous	% assumed	% small
Order size	Gamma	10.16	61.46	17.99	3.44	7.14

Tab. 2.6: Summary of order size testing

% homogeneous means the percentage of articles that follow a homogeneous distribution, % non homogeneous means the percentage of articles that follow a non-homogeneous distribution, % assumed means the percentage of articles that are assumed to follow a distribution and % small means the percentage of articles that have too few orders to be able to test on a distribution.

As can be seen only a small percentage of the total articles are assumed to follow a certain distribution or are ordered less than eight times, the rest of the articles have significant prove that they follow a Poisson or a Gamma distribution. These distributions will be used in chapter 4 to determine the reorder point.

3. APPLYING THE JIT POLICY AT THE ERASMUS MC

In this chapter we will look at the question which articles need to be stored in the DC. We will develop an algorithm that determines which articles should be kept in inventory, based on multiple criteria. We start this chapter with an introduction and literature research, then the algorithm will be proposed with an example and finally the results are discussed.

3.1 Literature research

The Erasmus MC has multiple different suppliers, different articles and different departments who continuously order multiple articles. The BIL department receives those orders and distributes the articles either from the DC in Barendrecht or directly from the suppliers to the different departments. Articles that are directly distributed from the supplier to the departments are called just in time (JIT) articles (see Danas et al. [2002]), while articles who are distributed from the DC to the departments are just called DC articles. These articles are first distributed from the supplier to the DC and stored there, and when there is a demand from a department they are distributed from the DC to the department. In figure A.1 an overview of the different article flows within the hospital is given. In the left figure the DC articles are shown, as can be seen the articles are first distributed from a supplier to the DC and then from the DC to the departments. In the right figure the article flow of the JIT articles is shown, as can be seen the articles are directly distributed from the suppliers to the departments.

The JIT policy is mainly used in the manufacturing industry and is a demand pull system. This means that demand is directly satisfied from the supplier, most of the time small orders will be delivered by the suppliers just before they are needed. This system requires a good relation between the supplier and the customer, the communication between those two is crucial for the system to work. Because the communication with the suppliers is crucial, only few suppliers are needed for the JIT system to work (see Kowalski [1991]; Rossetti [2008]; Kua-Walker [2010]). The more suppliers there are, the harder it is to maintain a good relation with all of them.

One of the reasons to apply a JIT system is that it could be beneficial. The costs associated with inventory will significantly improve (see Kim and Schniederjans [1993]; Kim and Rifai [1992]; Whitson [1997]). The costs of unused inventory will be redundant. Supplies in stock can no longer be damaged, lost or expired, because there are almost no articles in stock.

The number of employees with inventory as their concern are reduced and the costs of moving and managing supplies become neglected (see Kim and Rifai [1992]). The

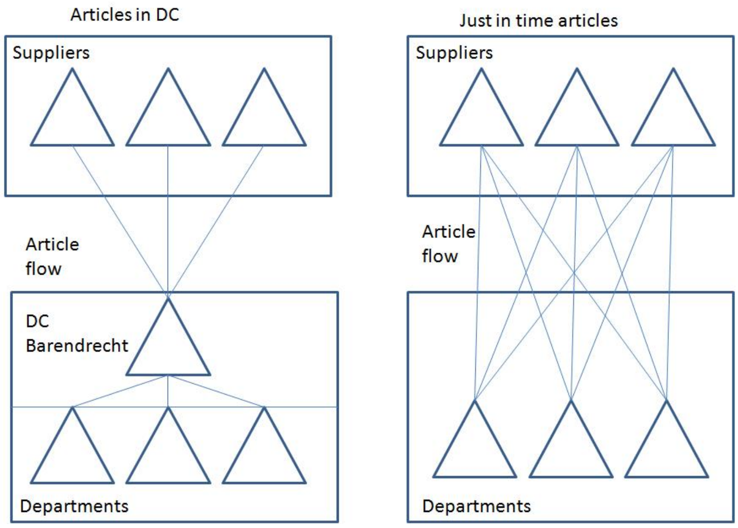

Fig. 3.1: Difference between direct delivery at the departments and first deliver at the DC and then to the departments.

rooms which were used for inventory can be used for other activities, or can be completely eliminated. At the Erasmus MC there is no possibility to completely eliminate the room, because it is a separate DC and certain articles always have to be in stock. We will discuss this later.

A problem with JIT inventory systems within hospitals is that the demand is uncertain and unpredictable (see Whitson [1997]; Kua-Walker [2010]). Certain articles within the hospital can mean the difference between life and death. When these articles are out of stock and the lead time of the supplier is for example more than two days patients could die because the article is delivered too late.

It is possible to keep a safety stock at a warehouse in order to avoid stock out (see Whitson [1997]). This means that hospitals always order directly at the supplier, unless the article is absolutely needed before the supplier could deliver it. In that case the hospital will order from the warehouse. Another possibility is to use JIT only on items which do not make a difference between death and life. Nicholson et al. [2004] investigated the JIT policy within a three stage echelon model at hospitals. He concludes that outsourcing only the non-critical medical supplies results not only in inventory cost savings, but also in higher service levels. At the Erasmus MC this is not the case, because the lead time from the supplier is longer than the leadtime from the DC.

Applying the JIT policy at the Erasmus MC has some complications. The first complication is that the lead times of the suppliers are not fixed, but could differ per order. For example the lead time could be two days one time and the next time 14 days, so the lead time of a supplier is unknown. According to the Erasmus MC the expected lead time of a supplier is five days, this is based on an average of the past year. For applying JIT policies the lead times have to be known and have to be short in order to meet the demand in the right way.

Another complication with applying the JIT policy at the Erasmus MC is that there are minimum amounts to order for some articles. For example if the demand at a unit is equal to 200 and the minimum amount to order is 1000, then 800 articles still have to be stored. These two factors make the JIT policy inapplicable for certain articles at the Erasmus MC, but the JIT policy could be applicable for other articles and could reduce the costs for supplying those articles.

Zimmerman [2010] investigated which criteria the Erasmus MC should use to determine whether the JIT policy could be applied at an article or not. He formulated a tool to decide per article whether it is necessarily to keep that article in stock or not. This tool is based on whether the articles make the difference between life and death and on whether the articles have a high "deliver risk". The deliver risk depends, among others, on whether there are substitutes for the article, how many suppliers deliver the article, does the supplier keeps inventory only for the hospital and is the supplier reliable? His tool can currently not be applied at the Erasmus MC, because there is not enough data available. But in the next section an algorithm based on this tool will be proposed.

3.2 Model

Michael Zimmerman (2010) formulated a tool to determine whether to take an article in stock at the DC in Barendrecht or not. His tool is based on the 'deliver risk' and the risk class of an article. The deliver risk is the risk of delivering the article in time, or finding substitutes in time. The risk class of an article means the risk for a patient when the article is not delivered in time. If the article makes a difference between life and death then the article falls within a high risk class. In Appendix A Zimmermans complete tool is explained.

Not all required data for Zimmermans tool is currently available at the Erasmus MC, so his tool can currently not be implemented. That is why we propose a new model based on the available data:

1. Determine the optimal reorder quantity Q and compare it with the average demand per order D
   - If $Q \leq 2 * D$, go to step 2
   - If $Q \geq 2 * D$, go to step 5
2. Compare the expected time T until a new order arrives with the lead time L
   - If $T \geq L$, go to step 3
   - If $T \leq L$, go to step 5
3. Check whether the article has a high patient risk
   - If that is the case, go to step 5

- If not, go to step 4

4. Apply the JIT policy on this article
5. Store the article in the DC

We will explain the details of every step:

(1) The optimal reorder quantity Q is the optimal quantity that needs to be ordered every time the DC places an reorder. Optimal means that the the total costs (order + holding costs) are minimized. Q is calculated with the EOQ formula, developed by Harris [1913]:

Q	=	Optimal order quantity
D	=	Annual demand of the article
P	=	Purchase cost per unit
F	=	Fixed order costs
h	=	Holding costs as percentage of the purchase costs

$$Q = \sqrt{\frac{2DF}{hP}} \tag{3.1}$$

The EOQ formula determines the most beneficial quantity to order (Q). If this quantity is lower than or equal to the average demand size (D), then it is better to order at the supplier every time a demand arises. If Q is larger than D, then it is better to keep that article in stock. But BIL does not want to place a reorder almost every time there is an order, that is why after consulting BIL we decided that Q has to be at least twice as large as D before taking an article in stock.

(2) The lead time (L) of an article is the time it takes for a supplier to deliver the article. At the Erasmus MC L is expected to be five days for every article. These five days is the average of all lead times of all suppliers. The actual lead time per supplier is not known. So it could be possible that a certain article has a lead time of one day and an other article has a lead time of 9 days. When the expected time until a new order arrives (T) for a certain article is shorter than the L of that article, then it is safer for the hospital to store the article in the DC.

(3) Articles with a high patient risk have to be stored in the DC. These articles can make the difference between life and death for patients, so the risk is too high to not store them. The Erasmus MC handed me a list of articles that have to be kept in stock. It would be more efficient to have this step first, to decrease computational time in the current first step. But it is not exactly clear at the Erasmus MC which articles have a high patient risk and if the hospital decides to check all the articles that are stored only according to the risk class they can find them easily by checking this step. That is why this is the third step.

(4) For the articles in this step it is profitable to apply the JIT policy on them.

(5) For the articles in this step it is profitable or safer to store them in the DC.

3.3 Example

In this section the previous discussed model is demonstrated with a small example. The model will be applied on five existing articles within the Erasmus MC. In order to apply the model, the average yearly demand, the price, the average order size, the expected number of arrivals per week and the patient risk has to be known. This information for all the five articles is given in table 3.1.

Article	**Identifier**	**Demand per year**	**Price**	**EOQ**	**Demand size**	**Interarrival time**	**Patient risk**
1	1121002	17.14	9.6	20	20	250	No
2	467104	0.86	400	1	1	250	No
3	1051316	6058	2.78	670	71.39	3.1	No
4	1025106	0.85	257	1	1	16.13	Yes
5	1177230	105	57.24	20	6.47	15.15	No

Tab. 3.1: Data information for five articles.

With the above information it is possible to apply the model:

1. Step 1 is comparing the optimal reorder size Q with the average demand size. Articles 3 and 5 have a Q more than twice as high as the average order size, so those articles have to be stored in the DC. Articles 1, 2 and 4 do not have to be stored according to the Q and these articles go to the step 2.

2. Step 2 is comparing the lead time of 5 days with the expected interarrival time between demands of an article. None of the remaining articles have to be stored in the DC, because the expected interarrival times are 250, 250 and 16.13 days, this is a lot higher than the lead time of five days. If article 3 would not be stored according to step 1, then he would have been stored according to this step because his expected interarrival time is 3.1 days, and this is lower than the lead time.

3. Step 3 is checking if an article has a high patient risk. Only article 4 has a high patient risk and so this article have to be stored in the DC.

4. Step 4 concludes to apply the JIT policy on article 1 and 2.

5. Step 5 concludes to store articles 3, 4 and 5 in the DC.

3.4 Results

The model is applied on both the articles who are currently kept in stock and on the articles on which currently the JIT policy is applied. In this section the results of the model will be shown and discussed.

3.4.1 Current articles in stock

First the model will be applied on all 2,086 articles that are currently kept in stock. Out of these 2,086 articles, 54 articles do not have to be kept in stock any longer. This is because Q is not much larger than the average demand size as can be seen in figure 3.2.

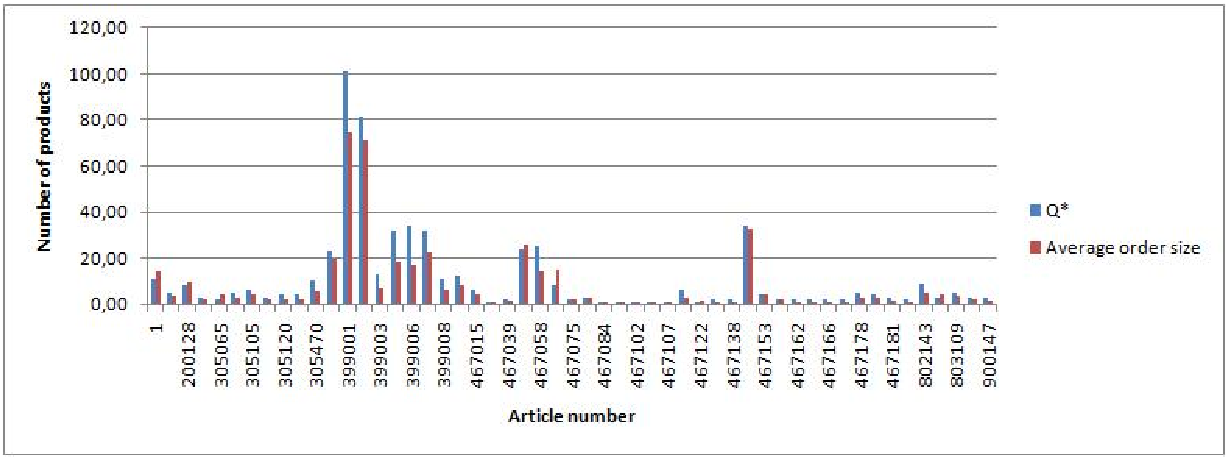

Fig. 3.2: Difference between optimal order size Q and average demand per order

It is not profitable to store them in the DC according to the Q. Also for these articles the expected number of demand arrivals is less than 1 per week and the articles are not critical. So the advice to the hospital is to not store those articles in the DC.

For the other 2,033 articles either the optimal order size is at least two times the average order size, the articles are expected to be ordered more than once a week or the articles have a high patient risk. There are 43 articles stored only because they have a high patient risk, if they did not have a high patient risk they would not be kept in stock.

3.4.2 Current JIT policy articles

There is one additional claim with the current JIT articles when comparing them with the current articles in stock. There are a lot of articles only ordered once in the past fourteen months, and it is unlikely that these articles are ordered again very soon. But it could be possible that according to the EOQ formula these articles have to be kept in stock, so the additional claim is that the articles have to be ordered more than once before keeping them in stock.

Then applying the model on the articles that currently are not kept in stock results in 68 articles for which it is profitable to keep them in stock. As can be seen in figure 3.3 the optimal reorder size Q is a lot higher than the average order size by the hospital.

So 54 of the current articles in stock do not have to be kept in stock any longer and 68 new articles have to be kept in stock. In the next chapter the consequence of this new policy will be expressed in terms of costs.

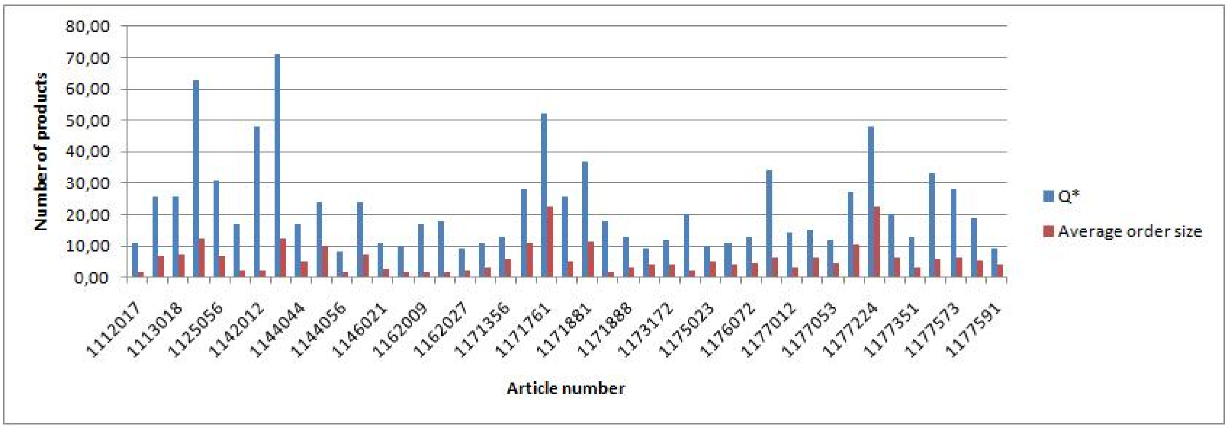

Fig. 3.3: Difference between optimal reorder size Q and average demand per order for a subset of the JIT policy articles

4. OPTIMAL REORDER POINT

In the previous chapter the reorder quantity Q was already determined, which means every time the DC places an order Q units will be ordered. However according to Zheng [1992], using the EOQ formula to determine Q in a stochastic environment instead of the optimal Q, leads to a maximum increase of the costs of $\frac{1}{8}$. The increase even vanishes when the order costs are significant relative to other costs. In the beginning of this chapter the Q determined by the EOQ will be used, but in section 4.4 the EOQ Q will be compared with the Q determined with joint optimization.

In this chapter we will first determine when the Erasmus MC has to place an order. At which inventory level do they have to reorder articles in order to be able to keep satisfying the demand from the departments. The maximum inventory level when the reorder has to be placed is called the reorder point. So when the stock drops to or below the reorder point a reorder will be placed.

The cheapest solution is to set the reorder point equal to $-Q$, so every time the virtual stock drops to or below $-Q$ there is a reorder of Q units. This means that the maximum level of stock is equal to zero and there never is stock available at the DC. No stock means no holding costs, but a reorder point of $-Q$ leads to a low service level. The service level is the percentage of demand that can be met directly from the stock. A low service level can lead to dangerous situations within the Erasmus MC, for example if the demand can only be satisfied in ten percent of the cases, then in 90 percent of the cases the hospital is waiting for its products and patients could be in danger. Another possibility is to set the reorder point really high, which of course leads to a high service level. But this policy tends to unnecessary high costs.

There has to be a balance between the costs and the service level. The Erasmus MC wants to maintain a service level of 98% for the articles they keep in stock. In this chapter the optimal reorder point while maintaining a service level of at least 98% will be determined. We start with an introduction and literature research, then we introduce formulas to determine the reorder point, after that we show our results and compare these results with the current situation and finally we take a look at the joint optimization.

4.1 Introduction and Literature research

The main decisions to make concerning inventory management are the decisions when to reorder and how much products to reorder. The decision of how much to reorder is discussed in the previous chapter and leads to an optimal reorder quantity Q. The decision of when to reorder is mostly based on the stock on hand, but also on the outstanding orders and the backorders. Outstanding orders are orders that are already placed, but not yet delivered. Backorders are demand from the customers that is currently not satisfied, because there were not enough products in stock. The so called inventory position (IP), see formula 4.1, is used to determine when to reorder. But the corresponding holding costs will depend on the inventory level (IL), see formula 4.2.

$$Inventory\ Position = Stock\ on\ hand + Outstanding\ orders - backorders \tag{4.1}$$

$$Inventory\ Level = Stock\ on\ hand - backorders \tag{4.2}$$

When the IP drops to or below the reorder point an order of size Q is placed. The order will not be delivered immediately, it depends on the lead time of the supplier. The lead time is the time it takes for a supplier to deliver.

An example to explain the IP and IL:

In table 4.1 the data for the reorder point, Q, L, the current IP and the current IL are given for a certain article. Now a demand for 11 products arrive from the hospital, that leads to the fact that both the IL and the IP drop to 9. Because the IP is below 10 a reorder of 10 products is placed at the supplier and the IP increases to 19. The next day a demand of 11 products arrives again from the hospital. The IL will drop to -2, this means that there is a backorder of 2 products, and the IP will drop to 8. Because the IP is again beneath 10, a reorder of 10 products is made at the supplier and the IP increases to 18, etc. The reorder arrives after five days and then the backorder is fulfilled and the IL increases to 8, one day later the next reorder arrives and the IL increases to 18.

Reorder point	Q	L	IP	IL
10	10	5	20	20

Tab. 4.1: Data for example

There are three different inventory costs: holding costs, order costs and backorder costs. Holding costs are the costs of holding one article in inventory per unit of time, order costs are the costs of placing an order and backorder costs are the costs of having no stock on hand while there is a demand. An example of backorder costs is giving the customer a discount on the price, because he has to wait. At the Erasmus MC there are no backorder costs.

This means that there are only holding and order costs. The holding costs depend on the reorder point, the larger the reorder point is the larger the holding costs will be. The reorder point has no influence on the order costs, only the order size Q has influence on the order costs. This means that the total costs will be larger when the reorder point increases. But a low reorder point leads to a low service level, and the Erasmus MC wants to have a service level of at least 98%, so the reorder point can not be too low. The goal is to minimize the reorder point while maintaining a service level of at least 98%.

There are two different inventory ordering policies, namely the (s, S) policy and the (R, Q) policy. Both policies have a different way to decide when to order and how much to order. But for both policies there exist algorithms to determine the reorder point.

4.1.1 (s, S) Policy

The (s, S) policy involves a reorder point s and an order up to level S. Every time the IP drops to or below s there will be a reorder. The order size is chosen such that the IP increases exactly up to S. A lot of heuristics are proposed in the literature to determine the parameters s and S. For example Roberts [1962], Wagner et al. [1965], Schneider [1978] and Tijms and Groenevelt [1984] all proposed different heuristics, a heuristic means that the policy is not optimal but close to optimal. There also exist several algorithms to determine the s and S optimal, such as the algorithm proposed by Federgruen and Zheng [1991]. This algorithm is based upon Markov decision problems, where the special structure of (s, S) policies is exploited.

In this thesis the (s, S) policy will not be applied, because some articles have fixed order quantities and that is why it could be impossible to always order up to exactly S.

4.1.2 (R, Q) Policy

The second inventory policy is the (R, Q) policy, every time the IP drops to or below R, Q units will be ordered. When the IP drops every time to exactly R the policy is equal to the (s, S) policy with $s = R$ and $S = R + Q$. But when the IP drops below R there is a difference between the (s, S) and the (R, Q) policy, with the (R, Q) policy there will still be an order of Q units while with the (s, S) policy the order will be an order up to $R + Q$ (S). In this paper we use an (R, n*Q) policy, this means that every time the IP drops below R there will be a reorder equal to $n * Q$, where n is an integer chosen in such a way that the IP is above R again.

Galliher et al. [1959] and Whitin and Hadley [1964] were the first to investigate the (R, Q) policy. They do not give methods to determine the optimal R and Q. Until 1991, as mentioned in Zipkin and Browne [1991], there was no reliable method to compute an optimal (R, Q) policy. There were several heuristics developed, as can be seen in for example Hopp et al. [1996]. Recently Federgruen and Zheng [1992], Matheus and Gelders [2000] and Axsater [2006] developed algorithms to determine optimal values for R and Q when the demand arrives according to a (compound) Poisson process. All the algorithms depend on a given service level, the percentage of demand that at least needs to be satisfied. The best combination of R and Q, while satisfying the service level, is then determined. This combination is the most cost saving combination. In section 4.2.1 the algorithm of Axsater will be discussed in more detail and applied.

4.2 Calculate R for a given Q

4.2.1 Optimal algorithm

In the previous chapter Q was determined according to the EOQ formula. In this section for every article the reorder point R will be determined for a given service level. The service level (S2) is the ratio between the expected satisfied demand and the expected total demand:

$$S2 = \frac{\sum_{k=1}^{\infty} \sum_{j=1}^{\infty} min(j,k) \cdot f_k \cdot P(IL = j)}{\sum_{k=1}^{\infty} k \cdot f_k} \tag{4.3}$$

where $min(j,k)$ is the delivered quantity to the customer. Which is equal to the demand size k when the inventory level j is larger than k and equal to j if k is larger than j. f_k is the probability that the demand is of size k and $P(IL = j)$ is the probability that the inventory level is equal to j.

The algorithm is to increase R until the service level is larger or equal to the given service level, and the corresponding R will be the optimal R.

In the case of the Erasmus MC the formula can be simplified to:

$$S2 = \frac{\sum_{k=1}^{\infty} \sum_{j=1}^{R+Q} min(j,k) \cdot f_k \cdot P(IL = j)}{\sum_{k=1}^{\infty} k \cdot f_k} \tag{4.4}$$

This is because the inventory level can maximum be equal to $R + Q$. Furthermore the definition of f_k must be known. In the previous chapter it is proven that the demand size is Gamma distributed. The Gamma distribution is a continuous distribution, this means that according to the Gamma distribution the demand size does not have to be an integer, but the demand at the Erasmus MC is always an integer. For this we assume demand between $k - 1 + \epsilon$ and k to be a demand of size k. With ϵ a very small number. That is why f_k can not be calculated with the probability density function, but with the difference between the cumulative distribution function of k and of $k - 1 + \epsilon$.

$$f_k = \frac{\gamma(\alpha, k/\theta)}{\Gamma(\alpha)} - \frac{\gamma(\alpha, (k - 1 + \epsilon)/\theta)}{\Gamma(\alpha)} \tag{4.5}$$

Where $\Gamma(\alpha)$ is the complete Gamma function and $\gamma(\alpha, k/\theta)$ is the incomplete Gamma function (see Abramowitz and Stegun [1962]). This only holds for the articles with a known continuous distribution. As explained before, the other articles are assumed to be Multinomial distributed, this will be discussed later.

The formula for the inventory level is equal to:

$$P(IL = j) = \frac{1}{Q} \sum_{k=max(R+1,j)}^{R+Q} P(D(L) = k - j) \tag{4.6}$$

So the inventory level depends on Q and on the distribution of demand during lead time. Adelson [1996] gives a recursion scheme to determine the distribution of the demand size during lead time:

$$P(D(L) = z) = \frac{\lambda \cdot L}{z} \sum_{j=0}^{z-1} (z - j) \cdot f_{z-j} \cdot P(D(L) = j) \tag{4.7}$$

and

$$P(D(L) = 0) = exp(-\lambda \cdot L) \tag{4.8}$$

Now for a given service level, Q, demand and arrival distribution, the R can be calculated. Start with $R = 1$ and increase R until the service level is at least equal to the given service level. Usually the policy is to start with $R = -Q$, but in the previous chapter we determined for which articles there always has to be stock on hand. So the minimum R is equal to 1 and that is why we start the algorithm with $R = 1$.

4.2.2 Heuristic

The formula of the previous section has a minor difficulty when applying at the Erasmus MC. As can be seen in chapter 3 there are some articles with an optimal reorder size Q larger than 1,000, for these articles it takes too much time to apply the formulas in section 4.2.1. It is necessarily to develop a heuristic in order to determine the reorder point for these articles.

First a simulation of the situation at the DC is developed:

1. Generate a large number of inter arrival times, depending on the arrival rate of the article.

2. Generate an order size for every arrival, depending on the Gamma parameters of the article.

3. Start at $Time = 0$.

4. As long as $Time \leq$ Number of years to simulate, do for every demand or reorder arrival:

   - $Time = Time +$ interval time.
   - Determine the holding costs during the time interval.
   - Check whether there is a reorder or a demand arrival.
   - If it is a reorder do:
     - Increase the inventory level with the reorder.
     - Check whether there is a backorder that has to be fulfilled.
     - If so, fulfill the backorder from the inventory level.
   - If it is a demand arrival do:
     - Supply the demand, as far as possible, from the stock.
     - Check whether the stock is below zero.
     - If so place a backorder and set the stock equal to 0.
   - Check whether the inventory position drops below R.
   - If so reorder $n * Q$ products, and add them by the inventory position, also determine the order costs.

5. Determine the service level of the article.

The heuristic depends on the above simulation. If the simulation is runned for a large number of years, the service level will approach the expected service level. So instead of using the formula of section 4.2.1 the simulation can be used. Simulating only once can lead to an outlier, so instead of one simulation there will be 1,000 simulations and the average service level of those 1,000 simulations will be used. The simulation starts with a stock of $R + Q$ products. In practice the stock will almost never be equal to $R + Q$, so starting with $R + Q$ is not very truthfully. The first 100 years of the simulation will be used to set the stock to its stationary level, from then on the service level will be recorded. So the simulated time will be 1100 years, with only the last 1000 years used to determine the service level.

The simulation usually takes longer than the formula of section 4.2.1, only for articles with a large Q the heuristic is preferable. Besides the speed, the heuristic does not give an optimal solution. So when possible always apply the formula of section 4.2.1, which is optimal.

The heuristic:

1. Start with $R = 0$.
2. Simulate 1000 times 1100 years and determine the service level (SL).
   - If $SL \geq 0.98 \rightarrow$ Step 4
   - If $SL < 0.98 \rightarrow$ Step 3
3. $R = R + X$ and go to step 2.
4. $R = R - 1$ and go to step 5.
5. Simulate 100000 years and determine the SL.
   - If $SL \geq 0.98 \rightarrow$ Step 4
   - If $SL < 0.98 \rightarrow$ Step 6
6. $R = R + 1$ and stop.

where:

$$X = max(0.10 * Q, 1) \tag{4.9}$$

In step 2 the average service level of 1000 simulations is determined and this average service level is compared with the required service level (0.98), if the average service level is lower than 0.98 R will be increased with a number X. X depends on the order size of the article. If Q is equal to 100,000 then increasing R with one has no big influence. But if Q is equal to 5 then increasing R with 100 has a way too large influence. That is why R increases with a percentage of Q, namely with $0.10 * Q$.

If the average service level is larger than the 0.98 then R is decreased with one. This is done until R is again below 0.98 and then increasing R with one leads to the minimum R with a service level above 0.98.

4.3 Results

In this section the results of applying the formulas and the heuristic of the previous section will be proposed. Also these results will be compared with the current policy applied at the Erasmus MC.

4.3.1 New policy

The Erasmus MC stated that the service level per article has to be at least 98%. The formulas of section 4.2.1 are used to determine an optimal value for the reorder point R, given the required service level. The value of the optimal order size Q is determined with the EOQ formula (see formula 3.1). As discussed in the previous section there are some articles with a Q too large to be able to apply the formulas of section 4.2.1 on this articles heuristic 4.2.2 is applied.

As can be seen in section 2.2.2 there are 441 articles that follow, or assumed to, a non-homogeneous Poisson arrival process. If the non-homogeneous arrival rate is used then the reorder point R will differ every three months. The Erasmus MC does not want to work with different reorder points every three months. That is why the largest arrival rate is chosen as 'fixed' arrival rate to determine the reorder point. In that case the service level will be around 98% in the three months with the highest arrival rate and larger in the other months.

This leads to the following results.

In figure 4.1 the R and the Q are presented for the articles with a R below 100. This are 1816 different articles, the different articles are shown on the x-as. As can be seen the Q differs between the one and the 7,876. The average Q is equal to 860.

Figure 4.2 presents the R and the Q for which the R is between 100 and 1,000. This are 243 different articles, the different are shown on the x-as. The Q for these articles differs between the 38 and 11284. The average Q is equal to 1758.

Figure 4.3 shows the R and the Q with a R larger than 1000. There are only 23 articles with a R above 1000, the different are shown on the x-as. The Q differs for these articles between the 2027 and the 29148. The average Q is equal to 10427.

So when the R increases, the Q also increases most of the time. That makes sense, because fast moving articles need a large R and Q to be able to meet the required service level of 98%. Q is most of the time much larger than R. If the Q and R are almost equal, either the reorder point is very large or the order quantity is low. A large reorder point leads to large holding costs and a low order quantity leads to large order costs. So it makes sense that the Q and the R differ from each other.

In figure 4.4(a) the service levels for all the articles, when determining the service level with the above presented R and Q, are presented. As can be seen the service level is always above 0.98 (98%). The average service level is equal to 0.983. In figure 4.4(b) the variance of the service level, determined every five years, is shown. This variance is never larger than 0.01, that means that the service level does not differ much every five years.

For some articles the service level is almost equal to 0.98, but for other articles the

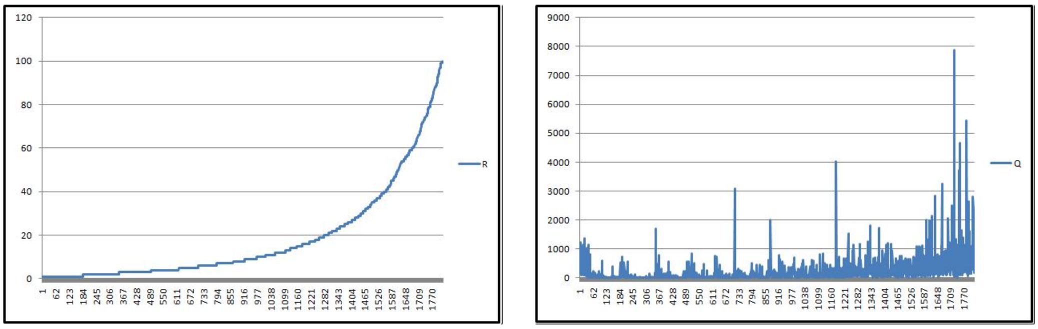

(a) R for articles with a R below 100 (b) Q for articles with a R below 100

Fig. 4.1: The Q and R for articles with a R below 100

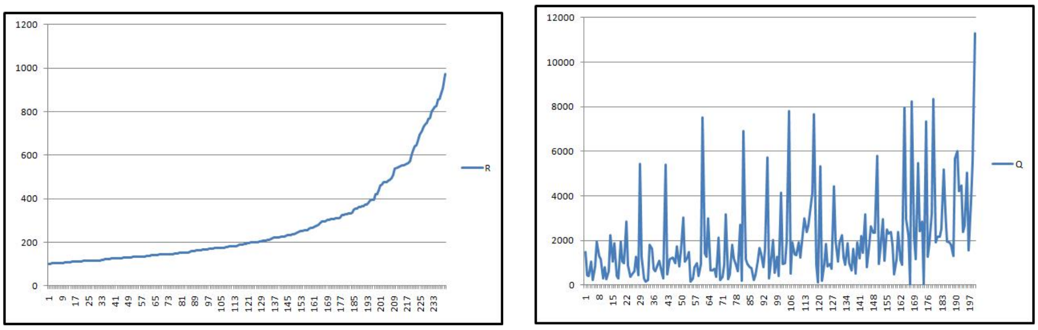

(a) R for articles with a R between 100 and 1000 (b) Q for articles with a R between 100 and 1000

Fig. 4.2: The Q and R for articles with a R between 100 and 1000

service level is close to 1. Articles with a large Q and a large demand have a service level almost equal to 0.98, because increasing or decreasing R with one does not have a large effect on the service level. The service level, determined in formula 4.4, will slowly increase towards 0.98 and can be stopped just above 0.98.

Articles with a small Q and relative few demands most of the time have a service level larger than 0.98, because increasing R with one has a large effect. That leads to the fact that the service level, determined again with formula 4.4, will exceed 0.98.

There is one exception: articles with a price equal or almost equal to zero. These articles have a very large Q because that is profitable according to the EOQ formula (see formula 3.1). Because the Q is very large and the arrival rate is not that large, the service level will be far above 0.98.

Next to the service level it is important for the Erasmus MC to know what the costs of the new policy are. As explained before there are two different types of costs when looking at the inventory at the Erasmus MC, namely holding and order costs. The holding costs depend on the number of products in stock, so on the inventory level. The order costs depend on how many times an order is placed, so on Q and the expected demand per year.

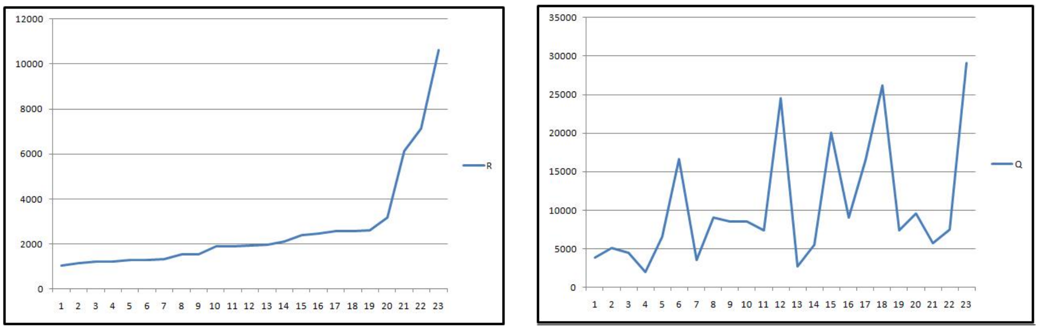

(a) R for articles with a R above 1000 (b) Q for articles with a R above 1000

Fig. 4.3: The Q and R for articles with a R above 1000

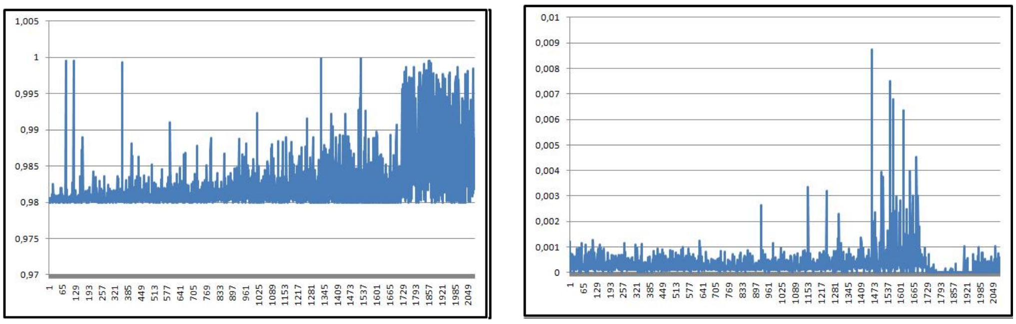

(a) Service level of all the articles (b) Variance per year of the service level

Fig. 4.4: The service level and variance of the new policy

The total costs are equal to the summation of both:

$$Order\ costs = \frac{K * \lambda * 253 * D}{Q} \tag{4.10}$$

$$Holding\ costs = h * p * \sum_{j=0}^{\infty} j * P(IL = j) \tag{4.11}$$

$$Total\ costs = Holding\ costs + Order\ costs \tag{4.12}$$

The upper part of formula 4.10 is the expected demand per year. K is equal to the fixed order costs, λ is the arrival rate, 253 are the number of working days in a year and D is the average demand per order. In formula 4.11 h is equal to the holding costs per year as a percentage of the price, p is the price of the article and $P(IL = j)$ is determined as in formula 4.6.

Formula 4.12 then determines the theoretic costs with the demand assumed to be constant. In practice the demand is not constant, but varying. That is why next to the theoretic costs also the simulation costs are determined. The simulation costs are the average costs per year after simulating 100,000 years. In figure 4.5 both the simulated and theoretic costs of per article are shown.

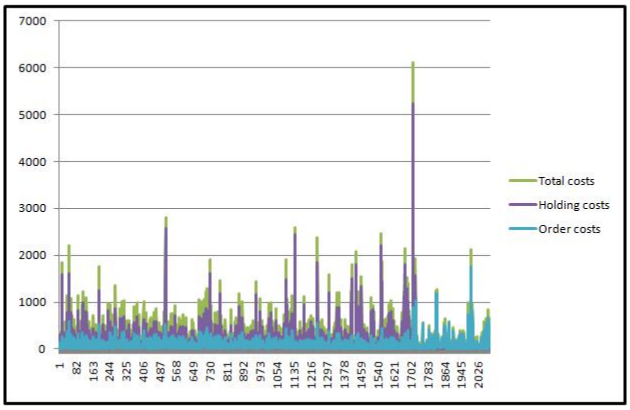

(a) Total, holding and order costs determined by simulation

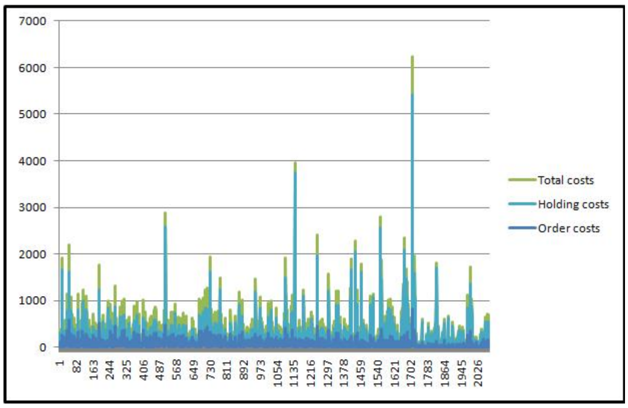

(b) Total, holding and order costs determined by theory

Fig. 4.5: The theoretic and simulated costs of the new policy

According to the theoretic formula the total costs are equal to €555,671.41 and according to the simulation equal to €550,514.26. The simulated and theoretic costs are almost equal to each other.

4.3.2 Comparing the new policy with the current policy

For the Erasmus MC it is convenient to compare the results of the new policy with the current policy. We will compare the two policies on two factors: the service level and the costs. As explained before the service level has to be at least 98%, no matter what the associated costs are.

The results of the previous section were for 2086 different articles. There are current 1903 different articles stored in the DC, that means that in the past 1.5 year 183 articles are converted to JIT articles. In this section we will compare the results of the previous section only on the 1903 articles that are current kept in stock.

The best solution is the solution with a service level above 98% and low costs. As can be seen in the results of the previous section, the service level of the new policy is always above 98%.

In Figure 4.6 and 4.7 the current used R and Q are presented. In the first figure there are 1793 articles on the x-as, all these articles have a Q below 1,000. In the second figure there are 110 articles on the x-as, all these articles have a Q above 1,000. The most remarkable fact is that the Q and the R do not differ that much from each other compared with the Q and the R in figures 4.1, 4.2 and 4.3.

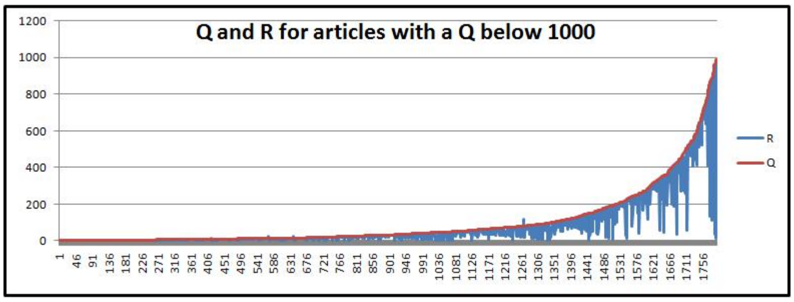

Fig. 4.6: Small Q and R current policy

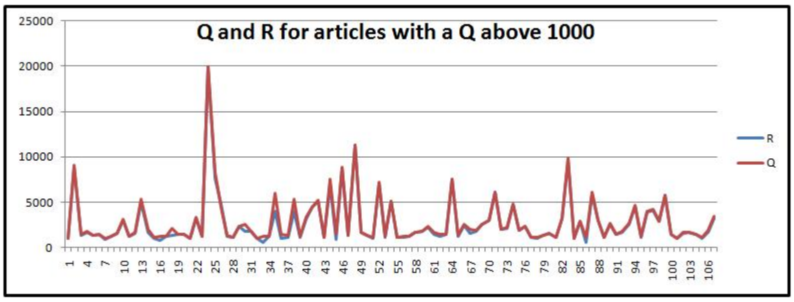

Fig. 4.7: Large Q and R current policy

In figure 4.8 and 4.9 the percentage difference between the Q and R of the new and current policy is shown. The percentage difference is calculated in the following way,

$$\frac{Q(R)_{new} - Q(R)_{current}}{Q(R)_{current}} \tag{4.13}$$

On the y-axis this percentage difference is stated, while on the x-axis the 1903 different articles are stated. The Q is in the new policy almost always higher than the Q in the current policy, while the R in the new policy is most of the time smaller than the R in the current policy.

The difference in Q and R is not very important for the Erasmus MC, only the service level and costs matter. The service level of the current policy can be determined with the

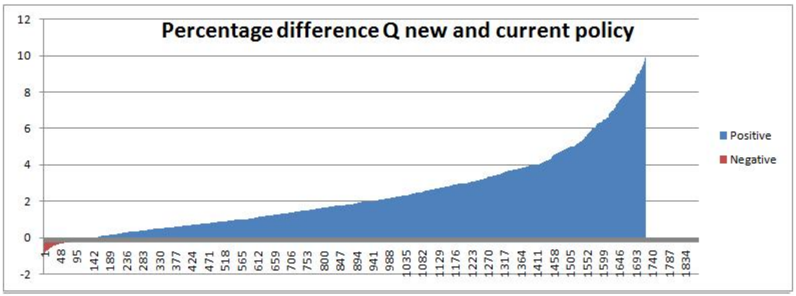

Fig. 4.8: Difference in percentages between the Q of the new and current policy

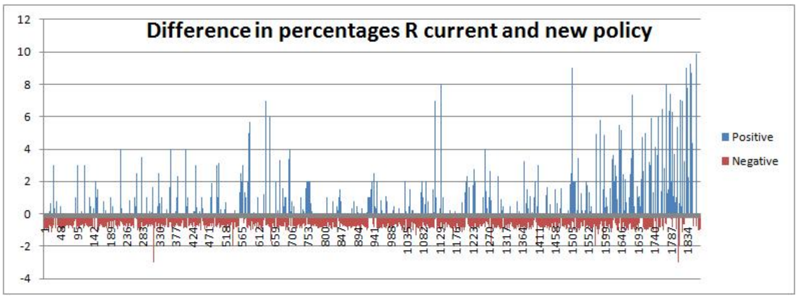

Fig. 4.9: Difference in percentages between the R of the new and current policy

formulas as stated in section 4.2.1. The R and the Q are given, just as the arrival rate and the Gamma parameters (see chapter 2). Applying the formulas leads to the service levels as stated in figure 4.10. The minimum service level is equal to 0,02%.

There are a lot of articles with a service level below the 98% (714 to be precise). The average service level of the current policy is 90.7%, so the current values for R and Q are not sufficient to meet the required service level.

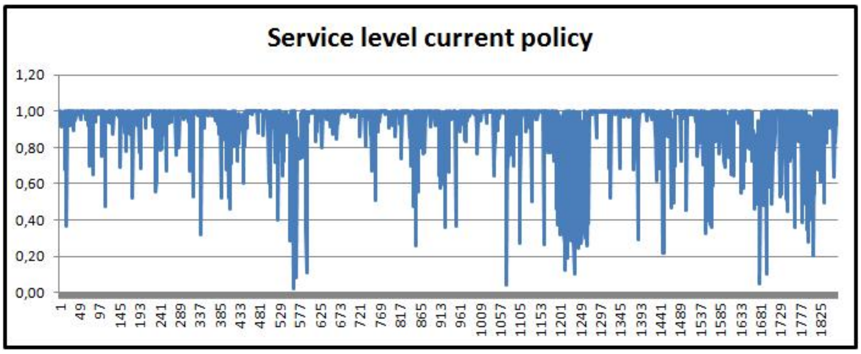

Fig. 4.10: Service level for the current policy

The second important aspect to compare the current policy with the new policy are the costs. The costs are determined with the simulation, with input values the current R, Q, the known arrival rate, Gamma parameters and the price per article. Simulating 100,000 years per article and determining the average costs per year leads to the costs as shown in figure 4.11.

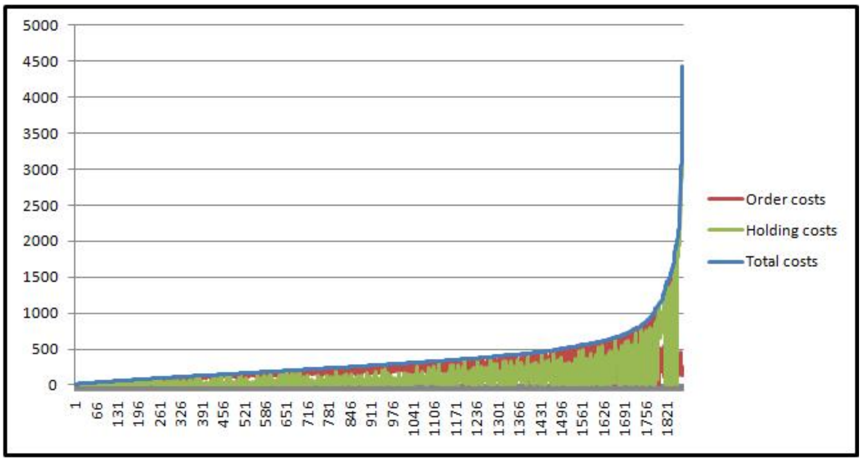

Fig. 4.11: Costs of the current policy per year

The total costs are equal to €689,908.70 per year. This is €139,394 per year larger than the total costs of the new policy. But that is the difference with the new policy based on 2086 articles, instead of the 1903 current stored articles. The costs of keeping the 1903 articles in stock in the new policy is equal to €489,205.39 per year. The difference between the current and new policy is than equal to €200,703.30 per year.

4.3.3 Summary and conclusion

Table 4.2 gives a short summary of the presented results of the current and new policy. The costs and service level are only determined for the 1903 articles that are current kept in stock. This is because there are at the moment 1903 articles stored in the DC.

	Costs per year	Service level
Current policy	€689,908.70	90.7%
New policy	€489,205.39	98.3%
Difference	€200,703.30	7.6%

Tab. 4.2: Summary of the results

4.4 Joint optimization

In section 4.3 the results of determining the R with the Q from the EOQ formula are presented. As stated in the beginning of this chapter according to Zheng [1992], using the EOQ formula to determine Q in a stochastic environment instead of the optimal Q, leads to a maximum increase of the costs of $\frac{1}{8}$. The increase even vanishes when the order costs are significant relative to other costs. In this section we will optimize Q for a given R in order to see how much the costs will decrease. The following algorithm is used to apply to joint optimization:

1. Fix R and decrease Q with one until the service level is beneath 98% or the total costs increases.

2. Fix Q and decrease R with one until the service level is beneath 98% or the total costs increases.

3. Go back to step 1. if step 2. did have effect.

In figure 4.12 the difference between the EOQ Q and the joint optimization Q can be seen. The maximum difference between the two is equal to 20. That does not seem to be a big difference. The R does not differ much from the proposed R, so it is not worth to show the differences between the different Rs.

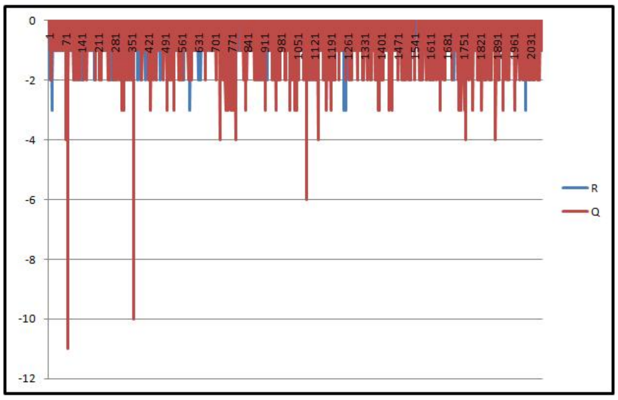

Fig. 4.12: The difference between the Q and R of the normal and the joint optimization policy

But the convenient way to compare the two policies is by looking at the costs. In figure 4.13 the difference in costs between the joint optimization and the new policy is shown.

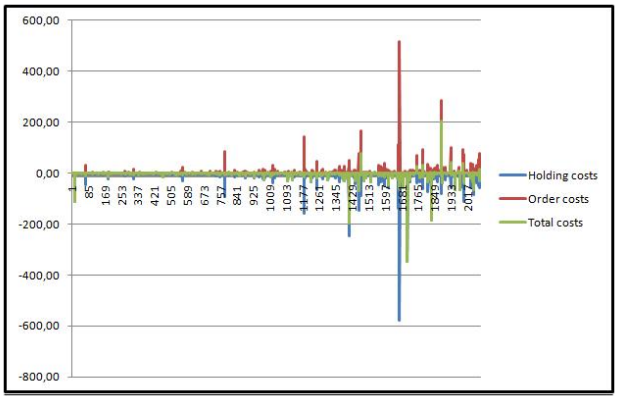

Fig. 4.13: The percentage difference between the costs.

The difference in costs is equal to 3.7% per year between the joint optimization and the earlier proposed policy. The joint optimization takes much more time, so it is not worth to apply joint optimization instead of the proposed policy at the Erasmus MC.

4.5 Benefit of storing/ ordering directly

In chapter 3 we proposed a model to decide which articles should be kept in inventory and which articles should be directly ordered at a supplier. As can be seen in section 3.4.2 there are 68 articles that should be kept in inventory which are now always ordered at a supplier. Besides that there are 54 articles that should always be ordered, but they are currently stored at the DC.

An adjustment to the simulation has been made in order to determine the effect of directly ordering an article at a supplier. The total number of demand arrivals during the simulation are recorded. At the end of the simulation this number is multiplied with the order costs and divided by the number of simulated years. In that way we estimate the average costs per year when the article is always direct ordered at a supplier.

Figure 4.15 shows the difference in costs between storing and directly ordering for the articles that are currently stored in the DC. Ordering these articles direct at a supplier leads to a cost reduction of €17,041 per year. There are no holding costs anymore and because the article is not ordered very often, the order costs are not that large as well.

Figure 4.14 presents the costs for the articles that are currently directly ordered at a supplier. It shows the difference between the costs when applying the current policy and when the article is stored at the DC (the Q and R are determined in the same way as the previous articles). It leads to a total cost reduction of €84,101 per year. This is because the articles are ordered quite often and they are kept in stock now, so the order costs are reduced.

The total cost reduction of applying the model in chapter 3 is equal to €101,142 per year.

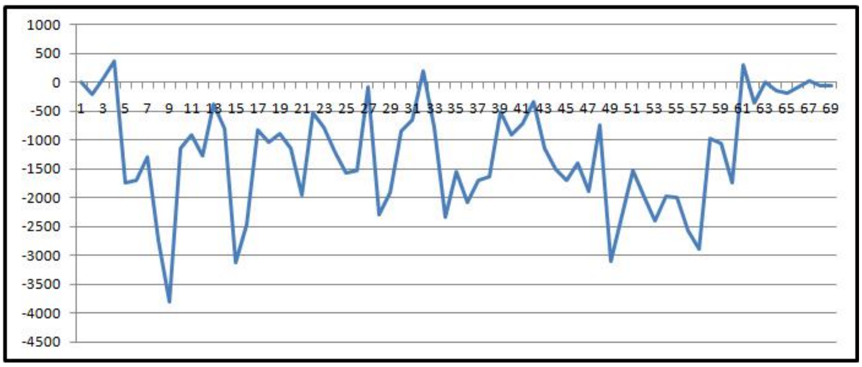

Fig. 4.14: The difference in costs when storing current JIT articles

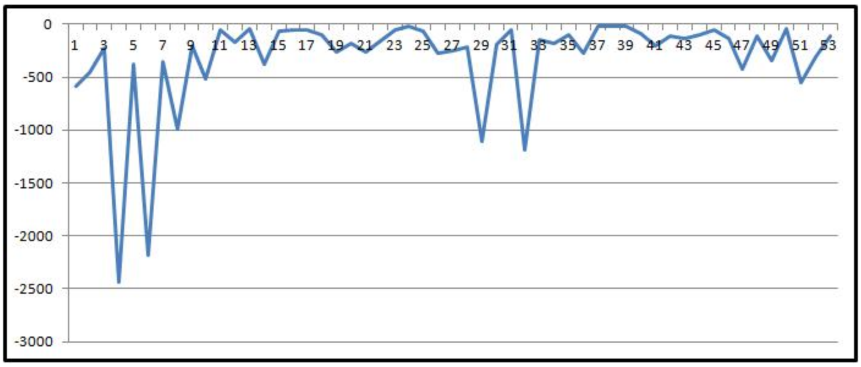

Fig. 4.15: The difference in costs when applying JIT on current stored articles

5. STORAGE LOCATION ASSIGNMENT

In the previous chapter the reorder point R and the optimal reorder size Q per article was determined. This leads to a certain number of products per article that have to be stored in the DC. The maximum number of products, per article, that could be in stock is equal to $R+Q$. So the DC in Barendrecht should have enough room to store $R+Q$ products of every article they want to store.

An important question for the DC is: where to store every article? The articles can be stored in different storage cabinets in the DC. Storage cabinets in the front of the DC can be picked faster than cabinets in the back of the DC, so articles that are ordered a lot should be stored in the front of the DC. It would also be convenient to store articles that are often ordered together close towards each other, in that way they can be picked together.

We will develop heuristics that allocate articles to a storage location in the DC. First an introduction and literature research will be given about this topic. After that the floor plan of the DC will be presented, in order to know how many storage locations there are available. Consecutive a method to determine the similarity between articles will be proposed, in the successive section this similarity will be used to develop a heuristic to allocate articles to a storage location. This heuristic will be compared with an optimal formulation that assigns articles to there optimal location with respect to travel distance. In this way we can check how good the heuristic performs. And finally the results and conclusion will be discussed.

5.1 Introduction and literature research

The management of the DC is controlling two different processes: the storage process and the retrieval process. The storage process contains the following steps: when a reorder for an article is delivered at the DC the products are stored on a temporary location. On this location the promised quantity is checked and if that quantity is correct a sticker is put onto the article. This sticker gives information to the employees about where the article should be stored in the DC. From the temporary location the article is stored by the employees at the instructed location.

The retrieval process is the process where orders from the hospital are collected from the DC. Those orders are placed on a platform until they are transferred to the hospital.

So both for the storage process as for the retrieval process, articles in the front of the DC can be picked or stored a lot faster than articles in the back of the DC. During the rest of the section we only look at the retrieval process. We assume that the best location to store articles for the retrieval process is also the best location to store articles for the storage process.

Order picking, the collecting part of the retrieval process, has been identified as the most labor-intensive and costly activity for a warehouse (see de Koster et al. [2006]). There are two different order picking policies, complete order picking and zoning. Complete order picking is when an employee picks an order entirely and when this is finished he continuous with the next order. Zoning means that every order picker covers a certain area (zone) in the DC. He only collects articles that are stored in his zone. An advantage is that every order picker needs to traverse a smaller area and order pickers become familiar with the articles in their area. A disadvantage of this policy is that orders are split and they must be consolidated again before they are transferred to the hospital. Speaker [1975] and Petersen [2002] did research to zoning in a warehouse and proved that the zone shape has a significant influence on the average travel time within a warehouse. At the Erasmus MC complete order picking is applied.

In this part of the research the main focus will be on the storage assignment of articles. It is important to efficiently store articles in the DC, because the more efficient the articles are stored the less time is needed to pick orders. de Koster et al. [2006] describe five storage methods: random storage, closest open location storage, dedicated storage, full turnover storage and class based storage.

Random storage (see also Petersen [1997]) means that every incoming reorder is placed on a random free spot in the DC. Random storage can only be applied if the storage is done with machines, if the storage is done with manpower random storage will usually turn out to be closest open location storage (see Hausman et al. [1976]). Every incoming reorder is placed on the closest free spot available. The third storage method is dedicated storage (see de Koster and Neuteboom [2001]), this means that for every article a fixed location is reserved, so when a reorder arrives the article is stored on its fixed location. The disadvantage of dedicated storage is that there has to be room for the maximum inventory level. The fourth method is full turnover storage (see for example Heskett [1963]), this method comes down to the fact that articles with the highest turnover are assigned to the location closest to the depot.

The last storage method is class based storage (see Hausman et al. [1976]), this method is a combination of the above policies. For example a combination between dedicated storage, full turnover storage and random storage leads to the ABC policy. Articles are divided into different classes, the fastest moving class contains about 15% of the products stored but contributes to about 85% of the turnover. These are called the A-articles, the next fastest category contains the B-articles and the slowest articles are the C-articles. Each class is assigned to a dedicated area of the warehouse, within that area the articles are randomly stored.

All the above storage policies are article based, that means that the storage policy is based on every article separately. The policies do not take any correlation between articles into account. It is reasonably possible that there are articles which are ordered most of the time together (Frazelle and Sharpe [1989]), for example in a supermarket coffee and coffee milk are often bought together. These articles are stored close to each other in the supermarket. At the hospital there are also articles which are often ordered together and in order to reduce the time needed to pick orders these articles should be stored together. The name of this policy is family grouping, similar products are stored in the same area. There are two types of family grouping methods known (see Bindi et al. [2007]):

1. The complementary based method, this method contains two phases. In the first phase items are clustered together based on their joint demand. In the second phase articles within the same cluster are stored close together.

2. The contact based method, this method is almost the same. Only in the first phase the items are clustered together if an item is often picked just after the other item, or vice versa.

In this research the focus will be on the first method. This is mainly because orders are picked one after another, so the contact frequency does not provide any advantage.

Frazelle [1989] and Tempelaar [2010] investigated similarity indices, which is the core of storage assignment with family grouping. Similarity indices is also the core of the storage assignment in this thesis.

We will propose two algorithms to allocate articles to a storage location within the Erasmus MC. The first algorithm takes the similarity between articles into account in order to cluster them. Then it allocates the cluster that is ordered most to the front of the DC and the cluster that is ordered the least to the back of the DC. The second algorithm also takes similarity of articles and popularity of the clusters into account, but the volume of the articles and the capacity of a shelf are added to this algorithm.

We will start by giving an overview of the floor plan at the DC in the next section.

5.2 *Current situation at the DC*

In this section an overview of the current situation at the DC of the Erasmus MC will be given. The only available data of the current situation is the number of available storage locations at the DC and the place of the storage locations. A storage location is either a location to store pallets or a storage cabinet. It is unknown how large and how long a storage cabinet exactly is or how large a pallet is. Only the number of shelves per cabinet are known and the number of different article types that approximately can be stored on a shelf is known.

The DC exist out of three different warehouses, a pick warehouse, a bulk warehouse and a sterile warehouse. The bulk and sterile warehouses are located on the ground floor and the pick warehouse is located on the first floor. The pick warehouse has the same size as the sterile warehouse and resides exactly above the sterile warehouse.

In figure 5.1 a floor plan of the first floor of the DC is shown. The pick warehouse is the warehouse where non-sterile articles are stored and picked. As can be seen there is one big road in the middle of the warehouse and one road at the left flank of the warehouse.

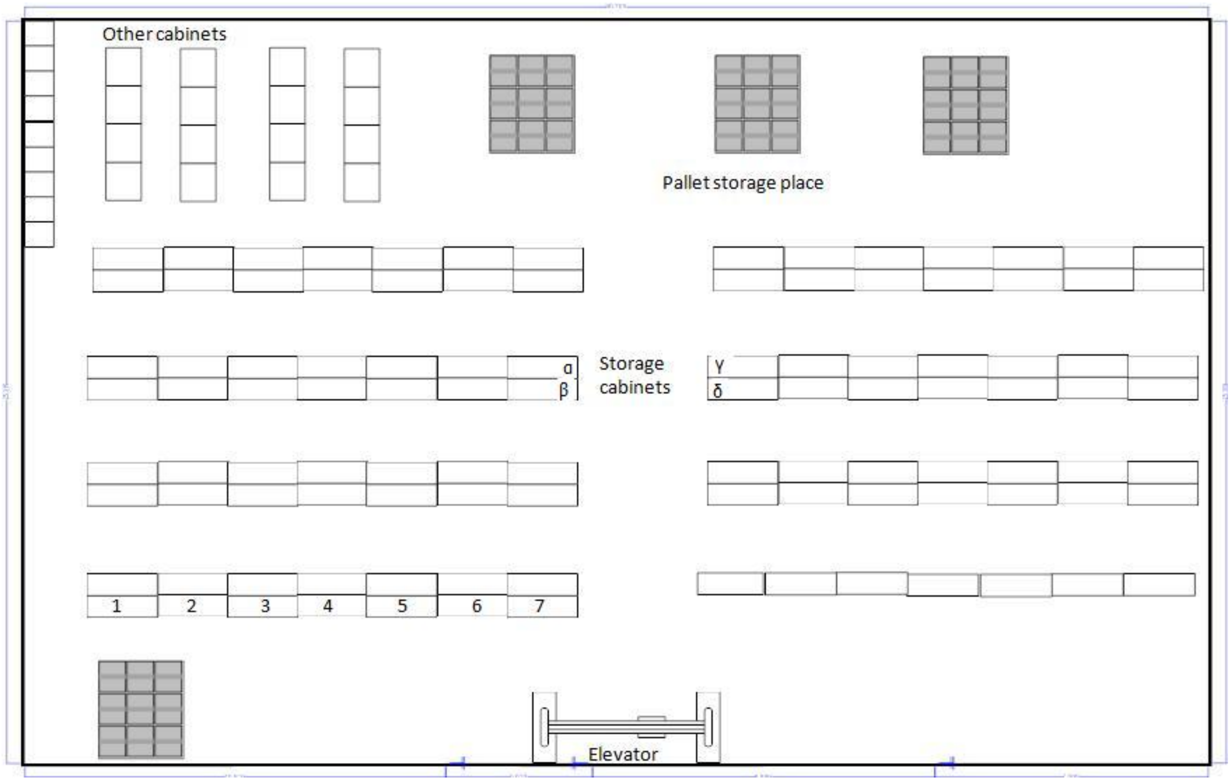

Fig. 5.1: A floor plan of the first floor of the DC

The DC has storage cabinets of different sizes. In the remainder of this section a storage cabinet of size 7 by 3 has seven shelves in the width and three shelves above each other. So such a storage cabinet has 21 shelves in total to store articles.

At the right side of the warehouse there are seven ‘normal’ storage cabinets and at the left side there are eight ’normal’ storage cabinets, they are called normal because these cabinets are used most in the pick warehouse. For example the storage cabinet with the numbers 1 to 7 is the first normal storage cabinet on the left. In total there are fifteen normal storage cabinets, these storage cabinets are of size 7 by 5. That means that all the normal cabinets have 35 shelves in total to store articles. On every shelve there is room to store exactly one article type. That means that there is room to store 525 different article types in the fifteen normal storage cabinets.

Besides these fifteen normal cabinets there are four storage cabinets at the upper left side (called ”Other cabinets” in the figure) of size 4 by 4. There is room to store three different article types per shelf. In total 192 different article types can be stored in these storage cabinets.

There is also one cabinet in the upper left corner of size 9 by 5, with room to store one article per shelf. In total 45 different article can be stored here. At last there are four pallet storage places, on each pallet there is room to store three different articles.

In total there is room for 774 different article types in the pick warehouse.

Articles that are stored close to the elevator can be picked fastest. Also articles stored close to the middle road can be picked faster than articles stored more to the right or left.

The optimal cluster size in the pick warehouse is equal to five, because most storage cabinets have height five. In that case five articles with a high similarity can be stored close to each other. But not all clusters will be of size 5, there will also be clusters with less articles than five in it. Those clusters will lead to left over space in the cabinets with height five. For example if the most popular cluster has four articles in it, this cluster will be placed at the front and one shelf is left over for other articles. In section 5.5 we will explain how this will be intercepted in the allocation model.

We will explain the preferences of certain shelves regarding to other shelves with a small hypothetical situation. In table 5.5 there are four hypothetical storage cabinets. These four cabinets correspond to the storage cabinets with the letters $\alpha, \beta, \gamma, \delta$ in figure 5.1. The shelves are ranked from A to O, where shelf A is the best location to store an article, shelf B the second best, and so on. Storage cabinet β and δ have the best shelves to store articles, because they are closest to the entrance.

At the Erasmus MC the preferable hierarchy of storage locations are unknown. The hierarchy of storage locations are assumed to be comparable with the example, only with more cabinets. The closer to the elevator, the more preferable the location is to store articles.

Tab. 5.1: Storage cabinet α

O	N	L	J	H	F	D
O	N	L	J	H	F	D
O	N	L	J	H	F	D
O	N	L	J	H	F	D
O	N	L	J	H	F	D

Tab. 5.2: Storage cabinet γ

D	F	H	J	L	N	O
D	F	H	J	L	N	O
D	F	H	J	L	N	O
D	F	H	J	L	N	O
D	F	H	J	L	N	O

Tab. 5.3: Storage cabinet β

K	I	G	E	C	B	A
K	I	G	E	C	B	A
K	I	G	E	C	B	A
K	I	G	E	C	B	A
K	I	G	E	C	B	A

Tab. 5.4: Storage cabinet δ

A	B	C	E	G	I	K
A	B	C	E	G	I	K
A	B	C	E	G	I	K
A	B	C	E	G	I	K
A	B	C	E	G	I	K

Tab. 5.5: Hierarchy example of storage locations

In figure 5.2 a floor plan of the ground floor is given. As can be seen there is a sterile warehouse, a bulk warehouse, an office and leftover space.

In the sterile warehouse sterile articles are stored and picked. This warehouse has one big road in the middle, but the entrance of the warehouse is not on the middle road. There are 57 cabinets, one in the back, 28 on the above the middle road and 28 on below the

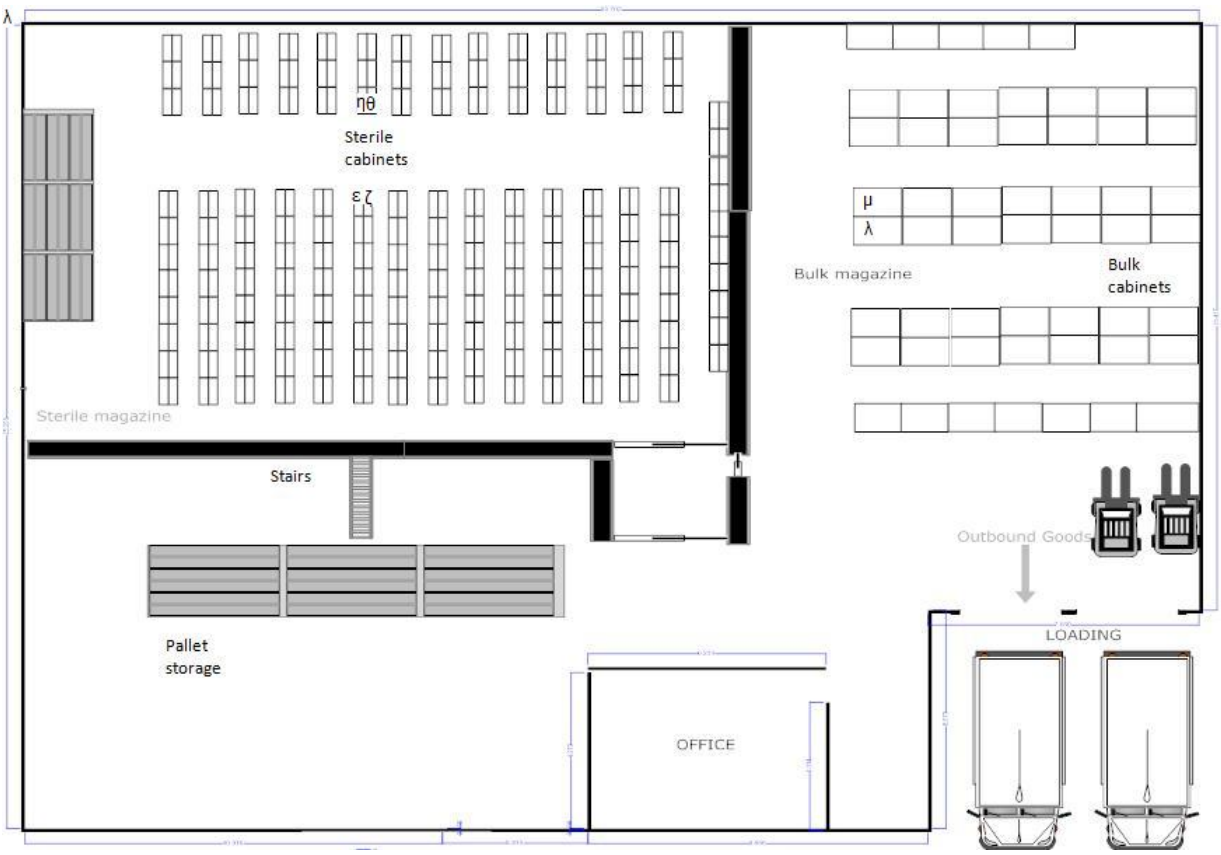

Fig. 5.2: A floor plan of the ground floor of the DC

middle road. The entrance is below the 28 cabinets below the middle road.

The 28 cabinets below the middle road are of size 8 by 3. There is room to store one article per shelf. That means that in total 672 different article types can be stored in those cabinets.

The 28 cabinets above the middle road are of size 3 by 3. There is room to store one article per shelf. That means that in total 252 different article types can be stored in those cabinets.

The cabinet in the back of the sterile warehouse is of size 10 by 4. So 40 different article types can be stored in the back. There is also room for three different pallets on the other side of the warehouse, there is room for nine different articles.

In total there is room for 973 different articles in the sterile warehouse.

We will explain the preferences of cabinets in the sterile warehouse again by a small hypothetical situation. In table 5.10 there are four cabinets shown, these four cabinets correspond to the cabinets with the letter $\epsilon, \theta, \zeta, \eta$ in figure 5.2. The shelves in the cabinets are ranked from A to Y, where shelf A is the best shelf and shelf Y the worst shelf to store articles. As can be seen shelves that are located closer to the entrance are ranked higher than shelves further away from the entrance.

The hierarchy at the Erasmus MC is assumed to be comparable, only with a lot more cabinets than in the example.

Tab. 5.6: Storage cabinet ϵ

C	E	K	S	W	U	O	I	G
C	E	K	S	W	U	O	I	G
C	E	K	S	W	U	O	I	G

Tab. 5.7: Storage cabinet η

M	Q	Y
M	Q	Y
M	Q	Y

Tab. 5.8: Storage cabinet ζ

A	B	H	P	T	R	L	F	D
A	B	H	P	T	R	L	F	D
A	B	H	P	T	R	L	F	D

Tab. 5.9: Storage cabinet θ

J	N	V
J	N	V
J	N	V

Tab. 5.10: Hierarchy example of storage locations

The third warehouse is the bulk warehouse. This warehouse has the function to store articles which are already stored in the pick or sterile warehouse, but do not completely fit in those warehouses. For example if there is room to store 500 catheters in a shelf at the pick warehouse, but the DC has to store 2,000 catheters then 1500 catheters are stored in the bulk warehouse. So this warehouse actually keeps stock for the other warehouses.

The bulk warehouse has one big road on the left of the cabinets. There are seven cabinets of size 7 by 4 and one cabinet of size 5 by 4. The difference with the last two warehouses is that there is room to store three different article types per shelf instead of one. In total there is room to store 648 different article types in the bulk warehouse.

The cabinets in the bulk warehouse are a lot higher than the cabinets in the pick or sterile warehouse, so compartments close to the ground are much better to reach than higher compartments. Shelves close to the road can be picked faster than shelves in the back of a side passage. So the best location to store articles is in a shelf at the front of the DC, close to the ground and close to the middle road.

The shelf preferences are again shown by a small hypothetical situation. Tables 5.11 and 5.12 represent two different storage cabinets, which correspond to the storage cabinets with the letters μ, λ in figure 5.2. The letters A to AQ stand for the hierarchy of the shelves, where A is the best shelf to store articles and AQ is the worst shelf to store articles.

There is one problem with setting the bulk warehouse up. There has to be information about the volume, the weight and the capacity of the cabinets in the other warehouses. It is unknown which articles take too much space to be able to store them completely in the pick or sterile warehouse, so it is impossible to allocate articles to the bulk magazine with the current available information.

According to the information the DC gave us there is room to store 1750 different articles in the pick and sterile warehouse. There are 2086 articles that have to be stored. It is also possible to store some article types only in the bulk magazine, very large articles for example. And there is room before the office to store articles, only it is not known

Tab. 5.11: Storage cabinet μ

AE	AG	AI	AK	AM	AO	AQ
Q	S	U	W	Y	AA	AC
C	E	G	I	K	M	O

Tab. 5.12: Storage cabinet λ

AB	AD	AF	AH	AJ	AL	AN
N	P	R	T	V	X	Z
A	B	D	F	H	J	L

how much room there exactly is available. So we assume there is enough room available to store all the 2086 articles.

In the next sections heuristics will be proposed to allocate articles to a shelf in the DC.

5.3 *Similarity indices*

5.3.1 *Theory*

The storage of articles will be based on the relation between two articles. Articles with a high correlation will have to be clustered, so that they can be stored close together. A similarity index gives a numerical value to the similarity between products. Articles with a high similarity index are more similar to each other than articles with a low similarity index. This research will follow most of the steps of the approach proposed in Frazelle and Sharpe [1989].

First the available data has to be transformed into an incidence matrix. This is a matrix that shows the relationship between two classes of objects. The rows indicate all the orders, while the columns indicate all the different articles. The composition of the matrix is binary, where a one indicates that a product is in the order and a zero when it is not in the order. So if (i, j) is one that means that article j is in order i.

With the incidence matrix we can form the so called Belonging Frequency Information (BFI):

$$A : a_{ij} = \sum_{k=1}^{n} a_{ijk} \qquad B : b_{ij} = \sum_{k=1}^{n} b_{ijk}$$

$$C : c_{ij} = \sum_{k=1}^{n} c_{ijk} \qquad D : d_{ij} = \sum_{k=1}^{n} d_{ijk}$$

with:

$$a_{ijk} = \begin{cases} 1 & \text{if articles i and j belong to order k} \\ 0 & \text{otherwise} \end{cases}$$

$$b_{ijk} = \begin{cases} 1 & \text{if article i belongs to order k and article j not} \\ 0 & \text{otherwise} \end{cases}$$

$$c_{ijk} = \begin{cases} 1 & \text{if article j belongs to order k and article i not} \\ 0 & \text{otherwise} \end{cases}$$

$$d_{ijk} = \begin{cases} 1 & \text{if articles i and j do not belong to order k} \\ 0 & \text{otherwise} \end{cases}$$

a_{ij} stands for the number of times article i and j are ordered together, b_{ij} for the number of times article i is ordered without article j, c_{ij} vice versa and d_{ij} stands for the number of times neither article i or j is ordered.

The BFI can be used to make three different similarity indices, in order to determine the similarity between articles:

1. First the most obvious similarity index. That is the fraction of common orders. This is the number of times article combination i and j are ordered together divided by the total number of orders. The mathematical formulation is stated as follows:

$$s_{ij} = \frac{a_{ij}}{a_{ij} + b_{ij} + c_{ij} + d_{ij}} \tag{5.1}$$

2. The second similarity index is the Jaccard index, see Bindi et al. [2007]. The Jaccard index is a common used index to determine the similarity between products. This index is comparable with the fraction of common orders, but now the number of common orders is divided by the number of orders in which product i and/or product j are ordered. The mathematical formulation looks as follows:

$$s_{ij} = \frac{a_{ij}}{a_{ij} + b_{ij} + c_{ij}} \tag{5.2}$$

   The advantage of this similarity index is that the similarity only depends on the orders in which at least one of the articles is ordered and not on all the orders.

3. The last similarity index is the cube per order (CPO) index, see also Bindi et al. [2007]. This similarity index does not only take the popularity (number of times ordered) of the articles into account, but also their external similarities, for example volume or weight. It is possible to use any external similarity, if there is information available about them. For example if the similarity is based on popularity and on

volume, then the CPO is determined as stated in (5.3) and the similarity index as stated in (5.4).

$$CPO_k = \frac{v_k}{\sum_l (a_{kl} + b_{kl})} \tag{5.3}$$

$$s_{ij} = \frac{a_{ij}}{a_{ij} + h(b_{ij} + c_{ij})} \cdot \frac{min(CPO_i; CPO_j)}{max(CPO_i; CPO_j)} \tag{5.4}$$

The intuition behind the CPO index is that articles with almost the same volume, weight or any other criteria will have the same similarity as before, but articles with completely different volumes will have a lower similarity than before. This results into the fact that articles with almost the same volume will be stored closer together.

Another advantage is that articles with a large volume/ weight can easily be stored in the front of the warehouse, to be able to let them form the base of a pallet. This is because articles with the same volume or weight will have a larger similarity, so they are clustered together. According to Bindi et al. [2007] h reduces the impact of coefficient $b_{i,j}$ and $c_{i,j}$ and makes sure that the CPO has a larger impact on the similarity between articles. They use h equal to $\frac{1}{4}$.

In order to determine the similarity between articles at the Erasmus MC a slight adjustment to the indices must be made. A hospital uses sterile and unsterile articles. Sterile articles can not be located near unsterile articles, because that will lead to unsafe situations for the patients. So before determining the similarity indices between the article types, first the sterile articles will have to be separated from the non-sterile articles. So the similarity indices will be determined separately for sterile and for non-sterile articles.

There is one possibility to get s_{ij} equal to zero, only when the articles are never ordered together. Then there is no similarity between the articles and those articles should not be stored close to each other.

In the next subsection a small example will be made to understand the similarity indices better.

5.3.2 Example

In this example there are 5 orders made, each order can contain five different articles. With the content of these orders it is possible to form the incidence matrix. That incidence matrix can be found in table 5.13. The volume and the weight are added to the incidence matrix. All five articles are of the same type, they are either sterile or non-sterile. It does not matter which one of the two, because the cluster procedure is for both the same.

With the help of the incidence matrix the BFI can be formed. In table 5.14 the number of times articles i and j are ordered together is given. This matrix is of course symmetric. Table 5.15 shows the number of times article i is ordered without article j. This matrix is not symmetric, so the whole matrix is presented. Table 5.16 gives the number of times article j is ordered without article. This matrix is of course the transpose of matrix B, but

Orders/ Articles	**1**	**2**	**3**	**4**	**5**
1	0	1	0	0	1
2	1	0	1	1	0
3	1	1	1	0	1
4	0	0	0	0	1
5	0	1	1	0	1
Total times ordered	2	3	3	1	4
Volume	10	2	3	16	21
Weight	25	10	2	15	11

Tab. 5.13: Incidence matrix

Article	1	2	3	4	5
1	-	1	2	1	1
2	1	-	2	0	3
3	2	2	-	1	2
4	1	0	1	-	0
5	1	3	2	0	-

Tab. 5.14: The A matrix

Article	1	2	3	4	5
1	-	1	0	1	1
2	2	-	1	3	0
3	1	1	-	2	1
4	0	1	0	-	1
5	3	1	2	4	-

Tab. 5.15: The B matrix

for the intelligibility the complete matrix is shown. And finally in table 5.17 the number of times neither article i or j are ordered is shown. This matrix is again symmetric.

With the BFI and the volume the three similarity indices can be formed. For example the CPO index between articles 1 and 2 can be calculated as follows:

$$\frac{1}{1+\frac{1}{4}*(2+1)}*\frac{2}{10}=\frac{4}{35} \tag{5.5}$$

These indices can be found in tables 5.18, 5.19 and 5.20. As can be seen according to the common order index and the Jaccard index articles 2 and 5 have the highest similarity. While according to the CPO index articles 2 and 3 have the highest similarity.

Article	1	2	3	4	5
1	-	2	1	0	3
2	1	-	1	1	1
3	0	1	-	0	2
4	1	3	2	-	4
5	1	0	1	1	-

Tab. 5.16: The C matrix

Article	1	2	3	4	5
1	-	1	3	3	0
2	1	-	1	1	0
3	3	1	-	2	0
4	3	1	2	-	0
5	0	1	0	0	-

Tab. 5.17: The D matrix

Article	1	2	3	4	5
1	-	-	-	-	-
2	$\frac{1}{5}$	-	-	-	-
3	$\frac{1}{3}$	$\frac{2}{5}$	-	-	-
4	$\frac{1}{5}$	0	$\frac{1}{5}$	-	-
5	$\frac{1}{5}$	$\frac{3}{5}$	$\frac{2}{5}$	0	-

Tab. 5.18: The common order index

Article	1	2	3	4	5
1	-	-	-	-	-
2	$\frac{1}{4}$	-	-	-	-
3	$\frac{2}{3}$	$\frac{1}{2}$	-	-	-
4	$\frac{1}{2}$	0	$\frac{1}{3}$	-	-
5	$\frac{1}{5}$	$\frac{3}{4}$	$\frac{2}{5}$	0	-

Tab. 5.19: The Jaccard index

Article	1	2	3	4	5
1	-	-	-	-	-
2	$\frac{4}{35}$	-	-	-	-
3	$\frac{4}{15}$	$\frac{8}{15}$	-	-	-
4	$\frac{1}{2}$	0	$\frac{1}{8}$	-	-
5	$\frac{5}{21}$	$\frac{8}{21}$	$\frac{8}{77}$	0	-

Tab. 5.20: The CPO index

In the next section the cluster procedure will be discussed, this procedure depends on the similarity indices.

5.4 *Clustering the articles*

5.4.1 *Model*

In the previous section the Common order, Jaccard and CPO indices were introduced to measure the similarity between articles. In this section we will propose two heuristics to cluster the articles based on Frazelle [1989]. The first heuristic, heuristic 1, is based on a fixed maximum cluster size and the second, heuristic 2, is based on the volumes of articles and the capacity per shelf.

We start with heuristic 1:

1. Create the Incidence matrix.
2. Create the Belonging frequency information (BFI) matrix for every article.
3. Determine either the Common order, Jaccard or CPO index.
4. Cluster procedure:
   (a) Start a new cluster with the article that is ordered most, this is the starting article.
   (b) Cluster this article with another article which has the highest similarity with the article, the similarity has to be larger than zero else go back to step (a).

(c) Cluster the articles with another article that has the highest average similarity with the current clustered articles, this average similarity has to be larger than zero else go to step (a).

(d) If the number of articles in the cluster is lower than the maximum cluster size go to step (c) else go to step (e).

(e) Remove the clustered articles from the list, go to step (a) if there are articles left, else go to step 5.

5. Rank the clusters according to their average popularity.

Step one to three are explained in the previous section, so these steps will not be explained again. Step four is the cluster procedure. Start each cluster with the remaining article that is most "popular", the article that is ordered most. Check which article has the highest similarity with the popular article and cluster that article with it, only if the similarity is larger than zero.

Now the similarity index has to be adjusted for the cluster, this is done by taking the average similarity between the remaining articles and the current cluster. Add the article with the highest similarity with the cluster to the cluster.

Check whether the cluster size is equal to the maximum cluster size, if not repeat the previous step. If so remove all the articles in the cluster from the article list and start over until all articles are clustered. For non-sterile articles the maximum cluster size is equal to five and for sterile articles the maximum cluster size is equal to three, see section 5.2 for the reasons.

In step 5 all the clusters will be put in the sequence from most popular cluster up to least popular cluster. The sequence is determined by taking the average popularity of the articles in a cluster.

As soon as the capacity of the shelves and the volumes of the articles are known the following heuristic, heuristic 2, is more common to use:

1. Create the Incidence matrix.

2. Create the Belonging frequency information (BFI) matrix for every article.

3. Determine either the Common order, Jaccard or CPO index.

4. Cluster procedure:

   (a) Start a new cluster with the article that is ordered most, this is the starting article.

   (b) Find the article with the highest similarity with the current cluster, the similarity has to be larger than zero else go to step (a).

   (c) Check whether adding this article to the cluster satisfies the capacity restriction. If so add the article to the cluster and go back to step (b). If not go to step (d).

Article	1	2,5	3	4
1	-	-	-	-
2,5	$\frac{9}{40}$	-	-	-
3	$\frac{2}{3}$	$\frac{9}{20}$	-	-
4	$\frac{1}{2}$	0	$\frac{1}{3}$	-

Tab. 5.21: Jaccard index with the first cluster

(d) Find the article with the next highest similarity with the current cluster, again the similarity has to be larger than zero or else go back to step (a).

(e) Check whether this article can be added to the cluster, according to the capacity. If so go to step (b), if not go to step (f).

(f) Check whether all articles are attempted to be added to the cluster, if so go to step (g), if not go to step (d).

(g) Check whether all articles are clustered, if so stop the procedure. If not go to step (a).

Every cluster is now as large as the capacity of a shelf tolerates, or a cluster does not have any similarity with other articles. In the next section the previous example will be continued to explain the cluster procedures.

5.4.2 Continued example

In section 5.3.2 an example of how to determine the Common order, Jaccard and CPO indices was given. We will continue this example to explain the cluster procedures. The maximum cluster size for heuristic 1 is equal to three. And for heuristic 2 the capacity of a shelf is equal to 40, that means that the maximum volume that fits in the shelf is equal to 40.

The article that is ordered most is article 5, so we start the cluster with article 5 for both heuristic 1 and 2. According to all the similarity indices article 5 has the best similarity with article 2. The total volume of the possible cluster is equal to 23, that satisfies the capacity restriction so article 2 is added to the cluster according to both heuristics.

Now we can form the average similarity between the cluster and the other articles. As an example the updated Jaccard index can be seen in table 5.21, the other similarity indices are adjusted in the same way.

The cluster has the highest similarity with article 3, according to all similarity indices. The total volume of the new cluster will be equal to 26. So according to both heuristics article 3 is added to the cluster.

The average similarity between the cluster and the other articles have to be determined again. The updated Jaccard index can be seen in table 5.22. The other similarity indices are adjusted in the same way.

Article	1	2, 3, 5	4
1	-	-	-
2,3,5	$\frac{67}{180}$	-	-
4	$\frac{1}{2}$	$\frac{1}{9}$	-

Tab. 5.22: Updated Jaccard index with the first cluster

Cluster	Average popularity
2, 3, 5	$\frac{3+3+4}{3 \cdot 5} = \frac{10}{15}$
1, 4	$\frac{2+1}{2 \cdot 5} = \frac{3}{10}$

Tab. 5.23: Cluster results of heuristic 1

The cluster has the largest similarity with article 1. According to heuristic 1 the cluster is full, so heuristic 1 will not add article 1 to the cluster. The total volume when adding article 1 to the cluster is 36, so according to heuristic 2 article 1 will be added to the cluster.

Heuristic 1 starts a new cluster with the most popular remaining article. This is article 1 (see table 5.13), which is ordered two times in total. Article 4 is the only article remaining and this article has similarity with article 1, so article 4 is added to the cluster according to heuristic 1.

According to heuristic 2 article 4 can not be added to the cluster, because the volume is too large. So article 4 forms a cluster on its own.

Now all the articles are clustered and we have to rank the clusters according to their average popularity for heuristic 1. Heuristic 2 already assigned each cluster to a shelf, so we do not have to rank the clusters according to their average popularity for heuristic 2.

In table 5.23 the clusters ranked to their average popularity can be found.

The cluster with articles 2, 3 and 5 is the most popular cluster, and this cluster should be placed at the front of the DC. The cluster with articles 1 and 4 should be placed on the other available location according to heuristic 1.

According to heuristic 2 the cluster with the articles 1, 2, 3 and 5 should be placed at the front of the DC and article 4 should be placed on the other available location.

5.5 *Assigning clusters to a cabinet*

In the previous section we presented two heuristics to form clusters. Heuristic 2 assigns each cluster also to a shelf or cabinet in the warehouse, but heuristic 1 only forms clusters. In this section we will give a small example to show how the clusters have to be assigned to a cabinet in the warehouse.

We have the same hypothetical situation as in section 5.2 for the pick warehouse, but now with only two cabinets. The hierarchy of the shelves are shown in table 5.26. As can be seen there is room for fourteen clusters with five articles in it.

When there are fourteen clusters of five articles it is easy to allocate them to a shelf. The clusters are ranked according to their popularity and the most popular cluster will be assigned to the shelves with hierarchy A, the second most popular cluster will be assigned to the B shelves, etc.

It becomes more difficult when there are 17 clusters of four articles for example. The first fourteen clusters are assigned in the same way as before. That leads to the fact that one A shelf, one B shelf, one C shelf, etc do not have any article assigned to it. There are also three clusters of four articles left, the most popular cluster of these three will be assigned to shelf A, B, C and D, the next most popular to shelf E, F, G and H and the least popular cluster to shelf I, J, K and L. So the left over space is filled with the least popular clusters.

The common rule to assign clusters to shelves is starting with the most popular cluster and assign it to the shelves highest in hierarchy, continue this principle until there is no hierarchy group of shelves left. Continue with the most popular cluster and assign it to the most popular left over shelves, continue this until all clusters are allocated.

Tab. 5.24: Storage cabinet β

M	K	I	G	E	C	A
M	K	I	G	E	C	A
M	K	I	G	E	C	A
M	K	I	G	E	C	A
M	K	I	G	E	C	A

Tab. 5.25: Storage cabinet δ

B	D	F	H	J	L	N
B	D	F	H	J	L	N
B	D	F	H	J	L	N
B	D	F	H	J	L	N
B	D	F	H	J	L	N

Tab. 5.26: Hierarchy example of storage locations

5.6 IP formulation

In order to be able to test the performance of the heuristics an integer formulation (IP) of the problem is used. This formulation is able to allocate articles to shelves optimal, when the problem instance is small, for given orders. The larger the problem instance, the larger the computational time becomes. So the the problem instance can not become too large. For example if there were ten orders made in the past days and these ten orders contain fifteen different articles, the IP formulation is able to allocate the articles in such a way that the travel distance for picking these orders is minimized.

At the Erasmus MC there are about 2,000 articles and about 500,000 orders, this is way too much data to be able to use the IP formulation. But it is possible to compare the heuristic of the previous section with the IP formulation for smaller problem instances. In that way it is possible to determine how far the heuristic is away from the optimal solution.

There are two assumptions made with the following formulation. The travel distance to pick articles within the same shelf is zero, so it does not matter for the distance whether one, two or three articles are picked from the same shelf. The second assumption made is that the orders are picked independent of each other.

Sets:

$M = \{1, ..., U\}$ is a set of shelves indexed by m and n

$I = \{1, ..., T\}$ is a set of article types indexed by i

$J = \{1, ..., S\}$ is a set of orders indexed by j

Parameters:

$$A_{i,j} = \begin{cases} 1 \text{ if article } i \text{ is in order } j \\ 0 \text{ otherwise} \end{cases}$$

$B_m =$ Distance from start to shelf m

$D_{m,n} =$ Distance between shelf m and shelf n

$V_i =$ The total volume of article type i

$C_m =$ The capacity of shelf m

Variables:

$$X_{i,m} = \begin{cases} 1 \text{ if article } i \text{ is stored in shelf } m \\ 0 \text{ otherwise} \end{cases} \tag{5.6}$$

$$Y_{j,m} = \begin{cases} 1 \text{ if an article from order } j \text{ is stored in shelf } m \\ 0 \text{ otherwise} \end{cases} \tag{5.7}$$

$$Z_{j,m} = \begin{cases} 1 \text{ if articles from shelf } m \text{ are picked first during order } j \\ 0 \text{ otherwise} \end{cases} \tag{5.8}$$

$$L_{j,m} = \begin{cases} 1 \text{ if articles from shelf } m \text{ are picked last during order } j \\ 0 \text{ otherwise} \end{cases} \tag{5.9}$$

$$G_{j,m,n} = \begin{cases} 1 \text{ if the order picker picks shelf } n \text{ right after shelf } m \text{ in order } j \\ 0 \text{ otherwise} \end{cases} \tag{5.10}$$

Objective function:

$$Minimize \ \sum_{j=1}^{S}(\sum_{m=1}^{U}(Z_{j,m} \cdot B_m + \sum_{n=1}^{U} G_{j,m,n} \cdot D_{m,n})) \tag{5.11}$$

subject to

$$\sum_{i=1}^{T} X_{i,m} \cdot V_i \leq C_m, \qquad m = 1, ..., U \tag{5.12}$$

$$\sum_{m=1}^{U} X_{i,m} = 1, \qquad i = 1, ..., T \tag{5.13}$$

$$X_{i,m} \cdot A_{i,j} \leq Y_{j,m}, \qquad m = 1, ..., U, \ j = 1, ..., S \tag{5.14}$$

$$Z_{j,m} + \sum_{m=1, m \neq n}^{U} G_{j,m,n} = Y_{j,m}, \qquad m = 1, ..., U, \ j = 1, ..., S \tag{5.15}$$

$$L_{j,m} + \sum_{n=1, n \neq m}^{U} G_{j,m,n} = Y_{j,m}, \qquad m = 1, ..., U, \ j = 1, ..., S \tag{5.16}$$

$$\sum_{m=1}^{U} Z_{j,m} = 1, \qquad j = 1, ..., S \tag{5.17}$$

$$\sum_{m=1}^{U} L_{j,m} = 1, \qquad j = 1, ..., S \tag{5.18}$$

$$\sum_{m \in K} \sum_{n \in K} G_{j,m,n} \leq \sum_{m \in K} Y_{j,m} - 1, \quad K \subset \{1, ..., U\}, \ j = 1, ..., S \tag{5.19}$$

$$X_{i,m}, Y_{j,m}, Z_{j,m}, G_{j,m,n} \in \{0, 1\} \tag{5.20}$$

Explanation of the variables and constraints:

- 5.7. This variable is necessary because articles stored in the same shelf and ordered together are picked at the same time.
- 5.8. This variable keeps track which article is picked first during an order.
- 5.9. This variable keeps track which article is picked last during an order.
- 5.10. This variable keeps track whether shelf n is picked right after shelf m during an order.
- 5.11. The objective function is to minimize the total distance traveled to pick all the orders.
- 5.12. Every shelf has a certain volume capacity and this constraint makes sure that that capacity is not exceeded.

- 5.13. Every article has to be stored in exactly one shelf.
- 5.14. This constraint makes sure that $Y_{j,m}$ is one when at least one article from shelf m is in order j.
- 5.15. Every shelf that is needed during an order has to be visited once. The shelf can either be visited from another shelf, or from the start.
- 5.16. Every shelf that is needed during an order has to be departed once, the order picker can either go to another shelf or back to the start.
- 5.17. During an order the order picker has to travel from the start to a shelf once.
- 5.18. During an order the start has to be visited once from a shelf (at the end of the order).
- 5.19. This constraint makes sure that an order picker does not walk in circles. No sub tour is allowed while picking articles from the cabinets, see Dantzig et al. [1954] for an extensive explanation of the original TSP constraint. This constraint is slightly adjusted, because not every cabinet is visited during an order. That is why only cabinets that are visited during an order are taken into account in this constraint.
- 5.20 The variables $X_{i,m}$, $Y_{j,m}$, $Z_{j,m}$ and $G_{j,m,n}$ are binary variables.

There is also a small disadvantage of the IP formulation. It only minimizes the order pick distance, but it does not take the composition of a pallet into account. It could be possible that according to the formulation small and light articles are picked first, and so form the base of the pallet. That means that large and heavy articles are picked last and are stored upon the small articles. It is possible to give a penalty to the distance traveled to pick an order according to the weight or volume of an article. The objective function can be adjusted to:

$$Minimize \sum_{j=1}^{S}(\sum_{m=1}^{U} Z_{j,m} \cdot B_m + \sum_{m=1}^{U}\sum_{n=1}^{U} G_{j,m,n} \cdot D_{m,n}) + \frac{1}{a}\sum_{i=1}^{T}\sum_{m=1}^{U} X_{i,m} * V_i * B_m \quad (5.21)$$

So there is a penalty equal to the volume given per distance unit traveled to pick the order. The importance of storing articles with a high volume in the front of the DC can be given with the factor a, the larger a is the less important it is to take volume into account. An important assumption made for this adjustment is that articles in the front are picked immediately during an order, not at last. According to the IP the order picker can visit a shelf at the front of the DC last, but we assume that an order picker will pick that shelf immediately.

5.6.1 Testing the heuristic

The heuristics, the IP formulation and a random policy will be applied to 20 different problem instances. In the first 10 problem instances there are 10 orders made containing 10 different articles. These articles are all sterile or non-sterile, so they are all of the same type. In the second 10 problem instances there are 15 orders made containing 15 different articles. In every problem instance there is room for three articles per shelf.

Figure 5.3 presents the distances between cabinets and the distances between cabinet 1 and the entrance. An assumption of these distances is that when an article that is stored in cabinet five needs to be picked, all the other cabinets can be picked without traveling any further.

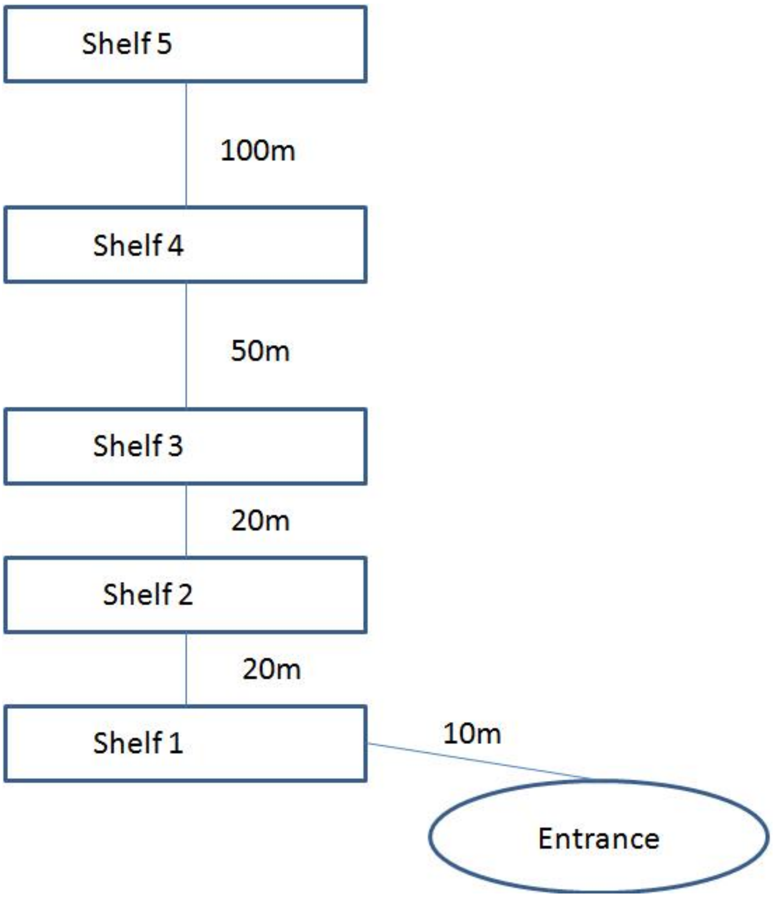

Fig. 5.3: The distances between cabinets

The articles will be assigned to a cabinet using heuristic 1(with both the Jaccard and common order index), the IP formulation and a random policy. We compare the four policies on the total distance traveled to pick all the orders in a situation. In figures 5.4(a) and 5.4(b) the total distance traveled to pick the orders per policy is shown. As can be seen the results obtained from the IP formulation is of course the best, but heuristic 1 (both the Jaccard as the common order index) does not perform much worse. It is clear that heuristic 1 performs a lot better than the random policy.

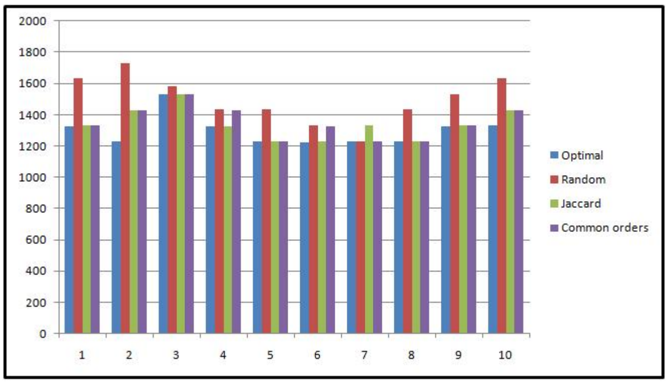

(a) Total distance traveled in situation 1

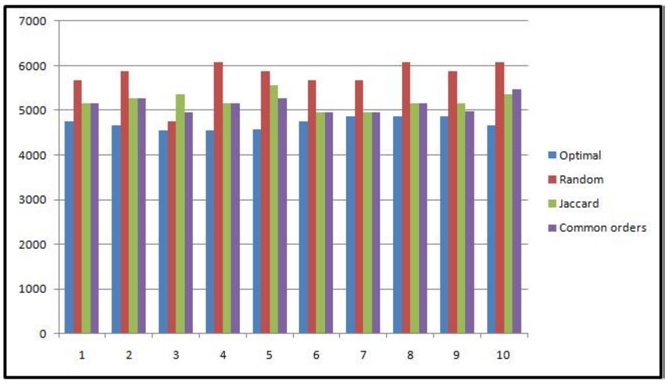

(b) Total distance traveled in situation 2

Fig. 5.4: The total distance traveled per situation per policy

The same conclusion can be drawn by looking at table 5.27. The average distance traveled per policy is compared with the optimal distance in percentages. Again heuristic 1 outperforms the random policy and is not that much worse than the results obtained from the IP formulation. For heuristic 1 it does not really matter whether the Jaccard index or the common order index is used, the results do not differ much from each other.

The IP formulation solves the first problem instance, 10 articles and 10 orders, in less than 5 seconds. The second problem instance, 15 articles and 15 orders, already takes more than 1.5 minute to solve. And if there are 20 articles and 20 orders it takes the IP formulation more than an hour to solve the problem. The formulation can not solve larger problem instances within a reasonable computational time.

That is why testing heuristic 2 is not possible. It is only interesting to test the heuristic with large differences in the volumes and a relative large number of articles, or else there is not much difference with using heuristic 1. The computational time for such a representative situation is too large in order to be able to produce results.

Case	15 or 10 articles	Percentage difference with optimal
Random	10	15.15%
Jaccard	10	3.20%
Common order	10	4.00%
Random	15	22.80%
Jaccard	15	10.70%
Common order	15	9.10%

Tab. 5.27: Average percentage difference between the policies

Articles/Orders	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	0	0	1	1	1	0	0	0	1	1	1	1	1	1	0
2	0	0	1	0	0	0	0	0	1	1	1	1	1	1	0
3	1	1	0	1	0	1	0	0	1	0	1	0	0	0	0
4	0	0	1	1	0	1	1	0	0	1	1	1	0	0	0
5	1	1	0	0	0	0	1	0	1	1	1	1	1	1	0
6	0	1	1	1	0	0	0	0	0	1	1	1	0	1	1
7	1	0	1	1	0	0	1	1	1	1	1	1	0	0	1
8	1	0	1	1	1	0	0	1	1	1	0	0	0	1	1
9	0	1	1	0	1	1	0	0	1	1	1	1	0	0	0
10	0	1	1	1	0	0	0	0	1	0	0	1	1	0	1
11	1	0	0	0	0	1	1	0	0	1	1	1	1	1	1
12	1	1	1	0	0	1	0	1	0	0	0	0	0	1	1
13	1	1	1	0	0	1	1	1	1	1	1	0	0	1	0
14	0	0	0	1	1	0	0	0	0	0	0	0	0	1	1
15	0	0	0	1	1	1	1	1	0	1	1	0	1	0	1

Tab. 5.28: Incidence matrix of the first situation

Based on the test results of heuristic 1 we conclude that the heuristics performs better than assigning the articles to a random spot in the DC.

5.6.2 Testing in detail

In this subsection one of the problem instances of the previous section will be explained in detail. In this problem instance there are 15 orders made containing 15 different articles. All these articles are of the same type, either sterile or non-sterile. The first step is to create the incidence matrix, which can be seen in table 5.28.

With the incidence matrix the BFI can be formed and with the BFI the Jaccard and Common order index. Applying heuristic 1 leads to the results as stated in table 5.29. In this problem the clusters contain the same articles for both the Jaccard as the Common order index. From the table it can be seen that for example articles 7, 8 and 13 are stored cabinet 1, this cabinet is closest to the entrance (as can be seen in figure 5.3).

Cluster	Articles	Cluster popularity	Cabinet
1	7, 8, 13	0.64	1
2	1, 2, 5	0.56	2
3	11, 15, 4	0.56	3
4	6, 10, 9	0.51	4
5	12, 3, 14	0.38	5

Tab. 5.29: Heuristic 1 clusters

Cluster	Articles
1	4, 7, 15
2	3, 5, 11
3	8, 12, 13
4	1, 9, 14
5	2, 6, 10

Tab. 5.30: Optimal clusters

Order	Heuristic	Optimal
1	400	100
2	400	400
3	400	400
4	400	400
5	400	200
6	400	200
7	100	100
8	400	100
9	400	400
10	200	400
11	400	400
12	200	400
13	200	400
14	400	400
15	400	400
Total distance	5100	4700

Tab. 5.31: Distance traveled per order

In table 5.30 the optimal allocation to the cabinets according to the IP formulation is given.

In figure 5.3 the distances between the cabinets and the start is shown. This distances will be used to determine the distance an order picker travels per order. It is known which cabinets do need to be visited during an order. In figure 5.5 an overview of how the order picker walks is shown per order. The numbers 1 to 5 stand for the five cabinets. As can be seen the cabinets are not on a straight line as in 5.3, but that is only to make clear which cabinets are visited during each order.

With this information it is possible to determine the distance traveled per order. The results can be found in table 5.31. As can be seen the total distance traveled is lower in the optimal situation, which is logical. But the difference is not very large.

5.7 Summary and conclusion

In this section two different heuristics were proposed to allocate articles to a location in the DC. Both heuristics depend on the similarity between articles and on the popularity of an article. The second heuristic also takes the volumes of articles into account. The first heuristic is compared with an optimal policy and with a random policy in terms of total travel distance to pick all the orders. The results show that the heuristic does not perform much worse than the optimal policy and a lot better than the random policy. It was not possible to compare the second heuristic with an optimal solution, because the computational time in a representative situation is too large.

Unfortunately the distance between cabinets and the weight and volume of articles is not known. So it currently is not possible to apply one of the heuristics at the Erasmus MC. As soon as this information is known the heuristic is able to perform better than the current layout.

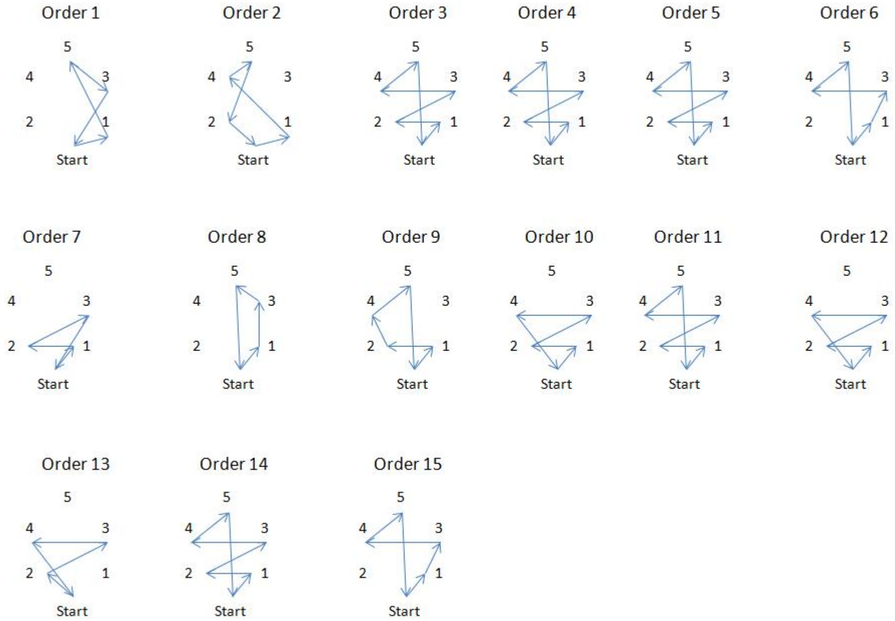

(a) Heuristic total distance traveled

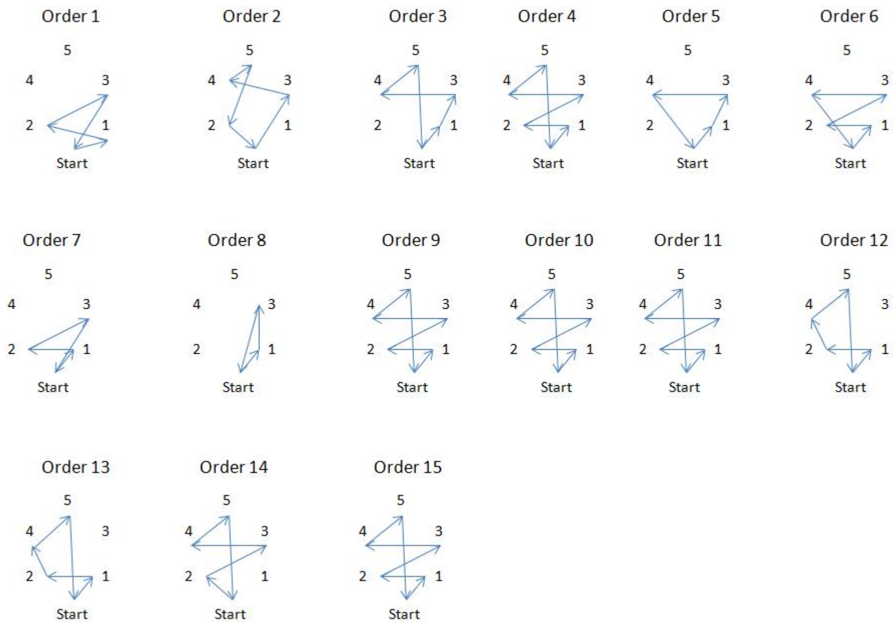

(b) Optimal total distance traveled

Fig. 5.5: The total distance traveled per order

6. CONCLUSION AND RECOMMENDATIONS

6.1 Conclusion

In this thesis we took a closer look on the inventory policies at the Erasmus MC. Three main questions concerning the distribution center of the hospital are answered.

The first question was which articles have to be stored in the DC and which articles should always be directly ordered at a supplier when a demand arrives? Chapter 3 introduces a model to determine this. The model is based on the optimal order size and on the patient risk of storing an article. The model concludes that 54 currently stored articles should always be ordered at a supplier and 68 articles should be stored while they are always ordered at the moment. A simulation shows that this modification leads to a cost reduction of €101,142 euro per year.

What should the reorder point of every stored article be? The reorder point depends on the service level, on the costs and on the reorder size. Chapter 4 determines the optimal reorder point R for a given Q and a service level of at least 98%. The formulas to determine R are based on the formulas of Axsater [2006]. In the current situation the Q and R are almost equal to each other, while in the new situation the Q and the R differ a lot more. A similation shows that this new policy saves €200,703 per year. Besides these costs saving the simulation also shows that the new model leads to an improvement in the service level. The current average service level per article is equal to 90.07% and the average service level of the new policy is equal to 98.30%.

In chapter 5 two heuristics are proposed to allocate the articles to a location within the DC. These heuristics are based on the similarity between articles, because articles with a large similarity have to be stored near to each other. The heuristics are tested on small problem instances and they perform slightly worse than an optimal allocation. They do perform better than a random location assignment. The heuristics could not be applied at the Erasmus MC, because the volume, the distances between shelves and the capacity of the shelves is unknown. But as soon as this information is known the heuristics will most likely lead to an improvement.

6.2 Recommendations for further research

The range of the available data is one year. This was not enough to determine possible seasonality effects. In further research it could be interesting to investigate whether certain articles are demanded more often during the summer or winter for example. Then it is possible to apply a dynamic inventory policy. That means different reorder points during summer and winter.

It was not possible to apply the heuristics of chapter 5 at this moment, because there is no information available about the volumes and/ or weights of articles. There is also no information available about the distances between cabinets, so it is hard to determine which cabinets are preferred over other cabinets. When this information is known it is possible to apply the heuristics.

It is also interesting for further research to develop larger similarity indices. Not only similarity between two articles, but between three or more articles. This could lead to different clustering results. An advantage is that the clustering procedure will be more accurate, it would likely result in better results. A disadvantage is that the computational time will increase, because the similarity indices will become much larger.

The heuristics now depend on either a fixed maximum cluster size or on a fixed shelf capacity. The heuristics can be improved by adding the possibility of using a variable number of shelves per cluster, depending on the similarities.

There currently are separate inventory policies for the departments within the Erasmus MC and the DC. For future research it is interesting to investigate a complete inventory policy containing both the departments as the DC. That will lead to lower costs and a transparent policy.

Appendix A

ZIMMERMANS TOOL

Michael Zimmerman (2010) formulated a tool to determine whether to take an article in stock at the DC in Barendrecht or not. As mentioned in the previous paragraph there are some complications with applying JIT at the Erasmus MC, some of the articles are 'critical' articles. These articles always have to be stored in Barendrecht. The first step of Zimmermans tool is to decide whether an article is critical or not. The easier an article is to get for the hospital the less reason there is to store the article in inventory. There are 4 criteria to decide whether an article is easy to get: 1. Are there substitute possibilities? 2. Is there more than one supplier for the article? 3. How high is the deliver reliability of the supplier(s)? 4. Does the supplier keep an inventory only for the hospital? All 4 criteria get different weights in order of importance:

Substitution possibilities	Weight
Multiple comparable items	-2
Multiple customers	-2
Alternatives	-2
Reliability customers	
98% or more	0
95 to 98%	2
Less than 95%	3
New supplier	2
Inventory for Erasmus MC	-4
Number of suppliers	
1	3
2	2
3 or more	0

These weights are given by an advisor medical supplies and an account manager acquisition. The deliver risk is not the only important factor. Even more important is the risk of the article on a patient. For example if a patient dies when the article is not delivered within a day, then the risk is really high and it makes a lot of sense to store that article in inventory. There are 4 possible risk classes for an article:

- Class 1 → Low risk, for example beds, plaster cast, plasters, etc

- Class 2a → Average risk, for example contact lens, needles, hearing aid, etc
- Class 2b → Average risk, for example anaesthesia machines, blood bags, bone cement, etc
- Class 3 → High risk, for example heart catheter, soluble adhesion, IUD

Based on the delivery risk and the effect on a patient the following decision matrix can be formulated.

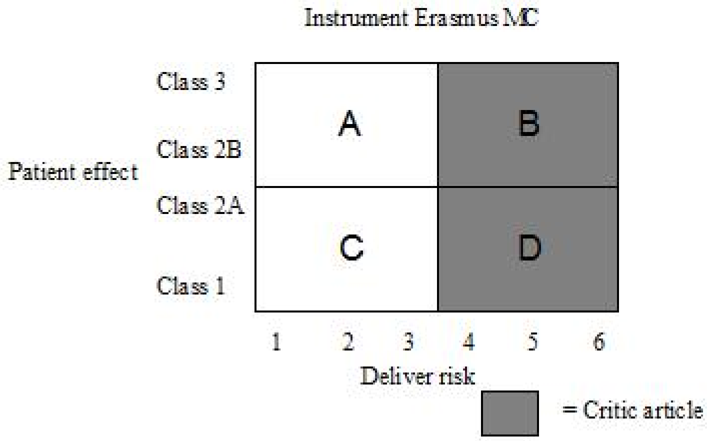

Fig. A.1: Decision matrix to determine whether or not to store an article.

On the y axis the risk class of a patient is shown and on the x axis the sum of the weights of the deliver risks. An article can fall into four different segments, based on their risk.

- A-articles → Big effect on patient, but low risk on late deliver → exceptions define whether or not to store the article in the DC-B
- B-articles → Big effect on patient, but high risk on late deliver → store in DC-B
- C-articles → Small effect on patients and low risk on late deliver → Direct deliver by supplier
- D-article → Small effect on patients, but high risk on late deliver → store in DC-B

So if an article falls within segment B or D then the article absolutely needs to be stored in the dc. If an article falls within segment A or C then the article does not need to be stored in the dc. But it could still be more profitable to keep the article in stock.
The main question for these articles becomes:
Is it more "profitable" to directly deliver from a supplier or to keep the articles in stock and is the lead time lower than the average time for a new demand?

With the EOQ formula the optimal order size Q per article could be determined for the distribution center. This order size means that every time the dc places an order at the supplier how large that order should be. The more times an order is placed the higher the order costs will be, but if there is more often an order the inventory will be lower and so the inventory costs will be lower. So the EOQ formula depends on the order costs, the holding costs and on the demand per year.
If almost every time an order is placed at the DC the dc has to order at the supplier at the same time to complete their inventory, according to EOQ formula. Then it is more beneficial to not store the article at all. So if the optimal order size is equal or lower than the average demand for the article then the article does not need to be stored in the dc.

BIBLIOGRAPHY

M. Abramowitz and I. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications, 1962.

R.M. Adelson. Compound poisson distribution. *Operations Research Quarterly.*, 17, 1996.

T.W. Anderson and D.A. Darling. Asymptotic theory of certain goodness of fit criteria based on stochastic processes. *The Annals of Mathematical Statistics*, 2, 1951.

S. Axsater. *Inventory control*. Operations Research Management Science, 2006.

F. Bindi, R. Manzini, and A. Regatierri. Similarity coefficients and clustering techniques for the correlated assignment problem in warehousing systems. *International conference on Production Research*, 2007.

K. Danas, P. Ketikidis, and A. Roudsar. A virtual hospital pharmacy inventory: An approach to support unexpected demand. *Journal of Medical Marketing: Device, Diagnostic and Pharmacetial Marketing*, 2002.

G. Dantzig, R. Fulkerson, and S. Johnson. Solution of large-scale travelling salesman problem. *Journal of the Operations Research society of America*, 2, 1954.

J.H. Darwin. The power of the poisson index of dispersion. *Biometrika*, 44, 1957.

R. de Koster and A.J. Neuteboom. The logistics of supermarket chains. *Elsevier*, 2001.

R. de Koster, R. Le Duc, and K.J. Roodberg. Design and control of warehouse picking: A literature review. *European Journal of Operational Research*, 182, 2006.

A. Federgruen and Y. Zheng. Finding optimal (s, s) policies is about as simple as evaluating a single policy. *Operations Research*, 39, 1991.

A. Federgruen and Y. Zheng. An efficient algorithm for computing an optimal (r, q) policy in contininuous review stochastic inventory systems. *Operations Research*, 40, 1992.

E.A. Frazelle. Stock location assignment and order picking productivity. *PhD dissertation at the Georgia Institute of technology*, 1989.

E.A. Frazelle and G.P. Sharpe. Correlated assignment strategy can improve order-picking operation. *Industrial Engineering*, 4, 1989.

H.P. Galliher, P.M. Morse, and M. Simond. Dynamics of two classes of coninuous-review inventory systems. *Operations Research*, 7, 1959.

F.W. Harris. How many parts to make at once. *Factory, The Magazine of Management*, 10, 1913.

W.H. Hausman, L.B. Schwarz, and S.C. Graves. Optimal storage assignment in automatic warehousing systems. *Managment Science*, 22, 1976.

J.L. Heskett. Cube-per-order index a key to warehouse stock location. *Transport and Distribution Management*, 3, 1963.

W.J. Hopp, M.L. Spearman, and R.Q. Zhang. Easily implementable inventory control policies. *Operations Research.*, 45, 1996.

G.C. Kim and A. Rifai. Efficient approach to health care industry materiel resource management: An empirical research. *Hospital materiel management quartely*, 13, 1992.

G.C. Kim and M. Schniederjans. Empirical comparison of just in time and stockless materiel management systems in the health care industry. *Hospital materiel management quartely*, 14, 1993.

A. Kolmogorov. Sulla determinazione empirica di una legge di distributione. *iornale dell'Istituto Italiano degli Attuari*, 4, 1933.

J. Kowalski. Inventory to go: Can stockless deliver efficiency? *Healthcare financial management*, 45, 1991.

Y.A. Kua-Walker. Can a just-in-time inventory sytem help reduce costs and increase productivity in hospitals? *Thesis at university of Florida*, 2010.

P.A.W. Lewis. Distribution of the anderson-darling statistic. *The Annals of Mathematical Statistics*, 4, 1961.

P. Matheus and L. Gelders. The (r, q) inventory policy subject to a compound poisson demand pattern. *Internatinal journal of production economics*, 68, 2000.

L.A. Nicholson, J. Vakharia, and S. Erenguc. Outsourcing inventory management decisions in healthcare: Models and application. *European Journal of Operational Research*, 154, 2004.

C.G. Petersen. An evaluation of order picking routing policies. *International journal of Operations & Production Managment*, 17, 1997.

C.G. Petersen. Considerations in order picking zone configuration. *International Journal of Operations & Production Management*, 27, 2002.

D. Roberts. Approximation to optimal policies in a dynamic inventory model. *Studies in applied probability and management science*, 1962.

M.D. Rossetti. Inventory management issues in health care supply chains. *Thesis at university of Arkansas.*, 2008.

H. Schneider. Methods for determining the re-order point of an (s, s) ordering policy when a service level is specified. *Operations Research.*, 29, 1978.

N. Smirnov. Table for estimating the goodness of fit of empirical distributions. *Annals of Mathematical Statistics*, 19, 1948.

R.L. Speaker. Bulk order picking. *Industrial Engineering*, 7, 1975.

M.A. Stephens. Use of komogorov-smirnov, cramer- von mises and related statistics without extensive tables. *Journal of the Royal Statistical Society*, 32, 1970.

I.R. Tempelaar. Achieving efficiency through similarity. *Thesis at Erasmus University Rotterdam.*, 2010.

J. Tijms and H. Groenevelt. Approximations for (s, s) inventory systems with stochastic leadtimes and a service level constraint. *European journal of operations research.*, 17, 1984.

H. Wagner, M. O'Hagan, and B. Lundh. An empirical study of exact and approximately optimal inventory policies. *Management Science.*, 11, 1965.

T. Whitin and G. Hadley. Analysis of inventory systems. *Prentice-Hall*, 1964.

D. Whitson. Applying just-in-time sytems in health care. *IIE solution*, 28, 1997.

Y. Zheng. On properties of stochastic inventory systems. *Management science*, 38, 1992.

M. Zimmerman. Naar een nieuwe norm(en) van voorraadbeheer. *Thesis at Erasmus university in Rotterdam*, 2010.

P. Zipkin and S. Browne. Inventory models with continuous stochastic demands. *The annals of applied probability.*, 1, 1991.