Master Thesis

Industrial Engineering & Management

Optimizing the base stock levels in a multi-echelon multi-item inventory/production system

Kasper van Ekeris

July 2023

[Part of the data from this thesis is left out or altered due to confidentiality. No product names are mentioned, y-axis are left out of the graphs, and random numbers are used in the analysis]

Supervisors
University of Twente
Dr. D.R.J. Prak
Dr. P.B. Rogetzer

Company
N. Spoelder

Management summary

Context

Nedap livestock management is a part of Nedap N.V., a Dutch technology company. One of Nedap livestock management's key products is the SmartTag. The SmartTag is a monitoring device for cows, attached to a collar, which provides valuable insights into their health and behavior. For our research we focus on the inventories of the production process of the SmartTag. Since all configurations of the SmartTag are produced at Nedap in two steps, we are dealing with a multi-echelon multi-item inventory/production system.

Currently, Nedap has little insights into good inventory management for their situation. Their production system provides multiple positions to place inventory. However, currently no safety stock is kept apart from a month of demand at the component level. This decision is solely based on the experience of employees and not on analysis. Therefore, we conduct our research to answer the following research question:

How can the inventory be optimally managed for the production process of the SmartTags?

Method

We have built a model to simulate the production process on a tactical level. Our simulation model evaluates the performance of different base stock levels placed on different echelon levels. Moreover, by adjusting parameters such as lead time, demand (mean and standard deviation), and production capacity, we created opportunities to assess the effect of the base stock levels on various scenarios. Our objective is to minimize the average inventory on stock while maintaining a certain service level.

We have found multiple models from literature which provide a method to determine base stock levels: Two versions of the model suggested by Chopra & Meindl (2007) and the guaranteed service model (GSM) by Graves & Willems (2003). Additionally, we used the current safety stock Nedap uses to determine a set of base stock levels and we evaluate a scenario in which no safety stock was kept at all.

Next to these five models, we also proposed an optimization model of our own. This model uses simulated annealing (SA) in combination with our simulation model to optimize the base stock levels.

Results

With the help of our simulation, we determined the objective values of the different models. The SA algorithm was able to obtain the lowest average inventory while maintaining a service level above 95%. When we compare our optimization model with Nedap's current model in terms of average inventory, we see a decrease of 30.7%. Moreover, not only does the average inventory becomes lower, but the service level also experiences a noticeable increase, reaching 96.3% compared to the previous 88.4%, respectively.

Furthermore, our analysis indicates that the current system is limited by the availability of the components rather than the production capacity. This suggests that the production capacity is sufficient to fulfill the expected demand. From our sensitivity analysis we found that an increase of more than 40% in demand or a decrease of more than 20% in production capacity leads to overloading the current system. Therefore, shifting the limiting factor from component availability to production capacity.

Conclusion

We created an optimization model which is able to outperform Nedap's current model and standard benchmark models obtained from literature. Outperforming these models is possible due to the effective placing of the base stock levels. Unlike these other models, our simulation model incorporates production constraints, which is crucial in determining the optimal base stock levels. Our model uses both echelon levels to account for the production constraints, which results in better objective values. Furthermore, with the support of our simulation model we are able to optimize and evaluate different decisions. This simulation is not only useful for optimizing the current base stock levels, but the simulation can also be used by Nedap to simulate future scenarios. Hence, our simulation model can advise Nedap to take and evaluate tactical decisions.

When conducting our sensitivity analysis, we found that the lead time has a great impact on the performance of the base stock levels. Therefore, Nedap has to make sure this lead time is known and accurate. In addition, the current production utilization showed that the system is currently not limited by the production capacity. From our analysis we were able to determine that a 40% increase in demand, makes the production capacity the limiting factor.

Furthermore, we recommended Nedap to implement the base stock levels into their ERP system. Additionally, we encourage Nedap to increase their KPI monitoring. Especially, the inventories and service levels are of importance, since they represent the performance of our decisions (base stock levels). The supplier lead time should also be monitored closely in order to make sure that the simulation model has accurate lead time values. The more accurate information is put into the simulation the more accurate information comes out. Therefore, it is of the essence that the input is as accurate as possible, which requires adequate monitoring.

Finally, our research provides two main opportunities for future research. The first opportunity is to evaluate other options to optimize the base stock levels. Machine learning algorithms or other (meta)heuristic are alternatives for optimizing the base stock levels. The second opportunity is to include more stochasticity in the model, especially stochastic lead times are interesting to incorporate in the simulation model, since we concluded that the effect of the lead time is significant.

Contents

List of tables

List of figures

List of abbreviations

Abbreviation	**Definition**
ERP	Enterprise resource planning
MRP	Material requirements planning
BOM	Bill of material
MTO	Make-to-order
MTS	Make-to-stock
PCB	Print circuit board
KPI	Key performance indicator
CODP	Customer order decoupling point
OPP	Order penetration point
ATO	Assemble-to-order
ETO	Engineer-to-order
DNN	Deep neural network
RL	Reinforced learning
GA	Generic algorithm
SA	Simulated annealing
MOQ	Minimum order quantity
C&M MTO	Chopra & Meindl Make-to-order model
C&M ATO	Chopra & Meindl Assemble-to-order model
GSM	Guaranteed service model

1. Introduction

This introduction provides an overview of Nedap, a Dutch technology company. In particular the business unit livestock management. One key product from livestock management is the focus of this thesis, the SmartTag. We place our focus on the supply chain of the SmartTag. From this, we state the research problem together with the research question and the sub-questions. Finally, we discuss the scope of the research.

1.1. Background information

The research for this master thesis is conducted at Nedap. Nedap is a technology company which provides hardware and software solutions for people in order to improve their work experience. Nedap consists of multiple business units: livestock management, healthcare, security, and more. This assignment is part of livestock management. Livestock management supports professional livestock farmers, both cow and pig farmers, in order to provide them with tools to run a profitable and sustainable business. This is done by enabling farmers to automate daily tasks and take decisions based on individual animal data. This empowers the livestock farmers to fulfil the increasing demand in protein, again, while maintaining a profitable business and meeting sustainability and animal welfare standards.

One way of achieving these goals is with the use of Cow Control. Cow Control is one of the propositions within livestock management that maximizes performance of livestock while reducing time, labor, and costs. An example of a product that contributes to this proposition is the SmartTag, this device hangs around the cow's neck and collects data from that cow. The data supports the farmer with information about the health, fertility, identification, and location of the cows. This SmartTag is sold on the global market and has a wide range of configurations due to customer-specific labelling and tag functionality.

The SmartTag is produced by Nedap SMART, which is the production facility of Nedap. Although it is part of the Nedap group, Nedap SMART is considered a separate supplier. Livestock management places orders at Nedap SMART for their SmartTags. This process is the same for the other parts of the Cow Control solutions Nedap offers. Livestock management places orders at the suppliers for their products. The one key difference is that Nedap SMART is located at the same place at the headquarters of livestock management, which makes communicating decisions accessible. Also, Nedap SMART is part of the Nedap group, this results in a cooperation in which decisions are made in harmony. These decisions include safety stock levels, production quantities, demand forecasts, and the production planning/process. Because of this cooperation it is possible to have a significant impact on the entire supply chain. A more detailed description can be found in Chapter 2, in which the current system is analyzed.

At the moment, there are a lot of factors influencing the supply chain of the SmartTag. One of these factors is the worldwide microchip shortage, which resulted in an increased backlog. Together with the shortage there is also an increase in lead times, which makes it difficult to create and maintain inventory. Moreover, the microchip shortage is applicable to one key component of the SmartTag. Because this component cannot be delivered, the SmartTag production could not produce at its maximum rate anymore. These factors forced the production to be make-to-order, which is currently performed as follows. First, the expected demand in the form of a forecast is entered into the Enterprise Resource Planning (ERP) system, which is the foundation for the planning. This forecast is then used in a Material Requirement Planning (MRP) to output planning propositions. These propositions are reviewed and executed on a weekly basis. After the operational purchase is completed, a production planning is created based on the current planning parameters. Finally, when the production of the SmartTag is completed, it is shipped to the warehouse, in which orders are picked and shipped to the intended destination. Currently, the lead time for the entire process is set to twenty weeks.

Although the past years have been a challenging time in terms of supplies and lead times, the future is expected to be promising. Suppliers should be able to deliver their supplies on time and in the right quantities. Furthermore, an additional production line is on its way to increase production capacity. Also, a second version of the same SmartTag will be created to reduce the risk of supplies complications. This second version has the same functionality but requires different components, which means that risks are spread. Finally, a search for an additional supplier is in place with the goal to increase the component availability.

1.2. Problem identification

At this moment, a lot is unclear about the inventory, production, and forecast of the SmartTags. First of all, there have been few to no possibilities to build up any stock in the last period. This is primarily due to the unexpected increase in lead time of components. This unexpected increase forced the production strategy to be make-to-order (MTO), rather than make-to-stock (MTS). Moreover, an attempt to accomplish an MTS strategy is exceedingly difficult right now due to the limited production capacity and a high number of back orders. However, it is expected that the future creates possibilities in the form of lower component lead times and an increased production capacity. Furthermore, small production gaps currently occur in the planning, which create possibilities for transitioning to an MTS strategy. Currently, Nedap lacks knowledge about the SmartTag's inventory management and production. No evident inventory levels are set, lead times to customers cannot be promised, and there is no clear plan on how the future changes should be handled. From interviews with the employees, it became clear that everyone would benefit from a structured inventory management. Thus, the goal is to find how the inventory management of the SmartTag can be improved.

1.3. Research problem

In general, Nedap is not able to create a structured planning for producing the SmartTag. Currently, decisions are made based on experience and gut feeling. Being able to provide a more data-driven approach will enable Nedap to create a more intelligent planning. A good planning will help them to provide customers with shorter lead times and a higher reliability. Also, Nedap can directly profit from a better planning in terms of costs. Optimal safety stock levels and order sizes can lead to a significant decrease in inventory cost. Based on these problems, a set of research questions is established, which are explained in Section 1.4.

1.4. Research questions

The main research question is stated as follows:

How can the inventory be optimally managed for the production process of the SmartTags?

In order to be able to answer this question, other research questions with sub-questions should be answered. In Chapter 2, the basis of the process is mapped in order to be able to create a suitable solution. This results in the first research question:

How is the supply chain of the SmartTag currently orchestrated at Nedap?

From this question the following sub questions are derived.

- How does the general process of selling a SmartTag look like?
- How does the production of the SmartTag look like?
- How is the bill of material (BOM) from the SmartTag structured?

- Which insights can be obtained from the historical data?

After that we consult the literature to conduct research in Chapter 3, with the purpose to obtain knowledge on how to improve the current situation. Therefore, the second research question is:

What literature is available concerning the optimization of inventory management in a production environment?

From this question the following sub questions are derived.

- How does the safety stock placement influence the inventory management?
- How can optimal safety stocks levels be calculated?
- How does the production process influence the safety stock levels?
- Which methods exist for optimizing inventory-production problems?

When the knowledge of the literature is acquired, the next step is to convert this knowledge to a method. This is done in the Chapter 4, which presents the third research question:

Which method is able to improve the inventory management for the current system?

From this question the following sub questions are derived.

- Which input is required by the method?
- How can solutions be generated and evaluated?
- Which output variables should be provided?

Next to that, we want to evaluate how our method performs. Moreover, the impact of the input on the output should be evaluated. In Chapter 5, the output and the effect of the input are examined to answer the fourth research question:

How does our method perform and how does the input relate to the output?

From this question the following sub questions are derived.

- What is the optimal output for the given situation?
- How do the different inputs affect the output?
- How do assumptions impact the output of the model?

Finally, in Chapter 6 we conclude the research and provide recommendations to Nedap. In addition, we discuss the limitations and impact of the research.

1.5. Scope

This research focusses solely on the supply chain of SmartTags. The demand, inventory decisions, production and purchases of the components of the SmartTags will all be included in the scope. Thus, the shipment of the products from and to the warehouse are excluded. In order to focus on the best planning, it is important to focus on a single crucial item for Nedap, the SmartTag. However, orders are highly likely to include more different items than just SmartTags. Also, because the SmartTag is the only product that is produced in-house, it is chosen as the focus of this research. Moreover, the

SmartTags can be seen as the bottleneck due to the components shortages it faces. With this knowledge, the goal is to provide a solution for the SmartTag, with the opportunity to adapt that solution for other products within livestock management. Furthermore, the current situation (extreme scarcity of components) is such a unique and temporary scenario, which is why we do not use these values for our method. It makes more sense to focus on a future scenario, in which values are far more reasonable. Therefore, we assume that this future scenario will have "normal" values and we use this future scenario to base our method on. This means lower lead times, ample possibilities to generate stock, no shortages and no more backlog. The goal is to provide a method which evaluates multiple scenarios based on these expected parameters. Also, since no forecast model is provided, demand is estimated based on the historical data and expectations from the sales team.

2. Current system analysis

In this chapter, the supply chain of the SmartTag is explained in more detail. This answers the research question on how the supply chain is orchestrated. First, in Section 2.1 an overview of the general process is provided. This process includes the steps from a customer order to delivery. Second, in Section 2.2 we explain the production process of the SmartTag. Next to that, the product layout is discussed in Section 2.3. Furthermore, in Section 2.4, data is gathered in order to gain quantitative information about the SmartTags and its production. Finally, Section 2.5 concludes the current system analysis.

2.1. General process

In order to gain an understanding of how the SmartTags are currently processed, an overview is provided in Figure 1. The first step in this process is the arrival of an order from a customer. The customers are business partners from Nedap, who use Nedap's products for their projects. Nedap takes this order in and returns a confirmation with an expected shipping date. Currently, this expected shipping date is more than twenty weeks after the intake of the order. This means that Nedap, currently has a lead time of over twenty weeks for the SmartTags. However, this is largely due to the shortages from the last few years. Before this, the lead time was close to five weeks. After this, the MRP plans the items and generates purchase orders to the supplier. All of this is part of the process when an order has been received. Next to that forecast, Nedap creates its own forecast in order to predict demand for the MRP. The general/forecasted orders together with Nedap's forecast form the input for the MRP. The output of the MRP is used as input at the supplier. In this case the supplier of Nedap livestock management is Nedap SMART and although it is part of Nedap, they are treated the same way as the other suppliers. Furthermore, Nedap SMART has their own MRP system which sends orders to their suppliers and creates a production planning. When the planning is fixed and the products are available, the production process can start. After this, the products are shipped to an external warehouse (grey box in Figure 1). The products stay at the warehouse until the shipping date. Finally, the customer receives the product together with the packing list.

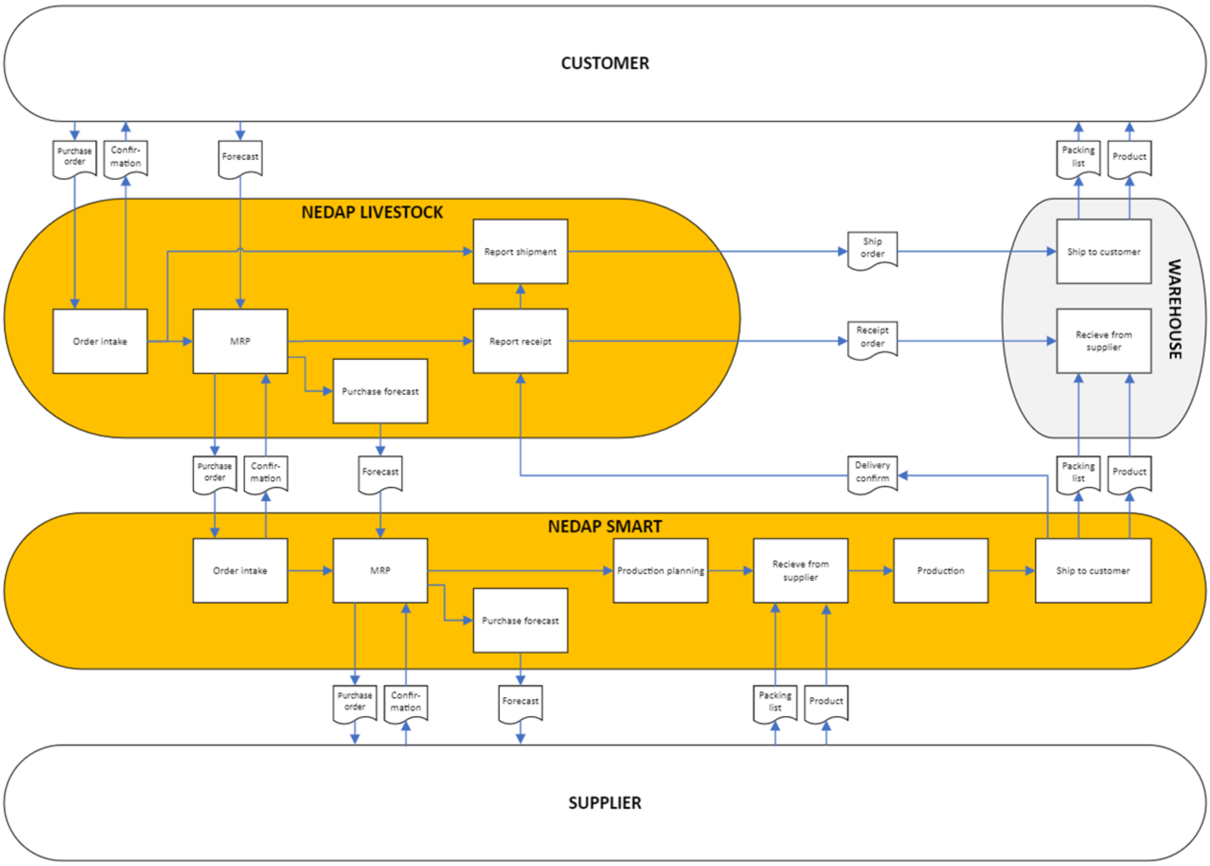

Figure 1. Overview general process at Nedap

Now that we have provided a general overview of the process, we explain the individual processes in more detail. First, we explain which actions are taken when the sales orders come in. Second, we discuss the forecast. After that, the purchase orders and warehouse deliveries are explained in more detail. In addition, the production process itself is explained separately in Section 2.2.

2.1.1. Incoming sales orders

The first step in the process is the incoming sales orders, in other words, the purchase orders from the customers. The key action in this process is deciding on when the order is ready to be produced and shipped. In the ERP system, the earliest possible shipping date is provided, based on the MRP. Nedap aims to ship the SmartTags within five weeks to the customer. Moreover, some customers work on project basis, which means that they place orders far in advance, because they know when SmartTags are needed for a project. Therefore, some planning in advance can be done.

2.1.2. Forecast products

Another input for the MRP is Nedap's own forecast of the SmartTags. A forecast is produced for every individual end product. However, the forecast is inaccurate at times, since the demand on end product level consist of many irregularities. Another reason for an inaccurate forecast is wrong historical data, because of the shortages of the previous years. The forecast is based on the sales orders which are actually shipped. And although the data showed a dip in sales, the demand was still there, but could not be delivered. These irregularities make the forecast unusable, which is the reason that the current forecast program is discontinued. Furthermore, Nedap currently relies on the predictions from the sales team. At this point, the products are forecasted, based on these predicted values. These values are predicted with the assistance of experts and can therefore be assumed as an accurate forecast for the coming period.

2.1.3. Purchase orders

The incoming sales order and forecast are used by the MRP to generate purchase proposals. These go from Nedap livestock management to SMART and from SMART to their suppliers. The suppliers decide together with SMART on product lead times and order quantities. Next to that, SMART decides on their lead times and production capacity with livestock management. The corresponding values are provided in Section 2.4.

2.1.4. Warehouse delivery

Nedap uses an external warehouse to arrange the storage of their end products. Suppliers send their products together with the packing list to this external warehouse. The party that is in charge of this warehouse sends a confirmation to livestock when the products arrive. This information is used to confirm the shipment to the customer. The last step in the entire process is to ship the products from the warehouse to the customer's location, with the packing list.

2.2. Production

The SmartTag is one of the only products within livestock management which is still produced in-house. The entire production can be divided into two parts: the first part is about the housing with the printed circuit board (PCB), the second about labeling and covering the housing to make it a usable tag. The first production step of the first part is the potting and testing of the PCB. After that, the PCB is put into a housing and casted in special material, after which the housing is closed via mirror welding. The last step of the first production part consists of programming the tag and lasering the housing. After this, a small container with a PCB is ready to be used. However, this small container is not ideal to place around the cow's neck. That is where the second production process comes into place. This part has two purposes, one is to frame the tag in such a way that the cow has the least amount of discomfort, another is to cover/label the tags for the right customer. This last purpose is important, since some customers have their own color and logo on their tags. This is because business partners sell the tags under their name. Nevertheless, some smaller customers do not have customized tags and use the general Nedap SmartTags. In addition, not all customers demand all different tag types, some customer order only one type of SmartTag. For the current analysis it is irrelevant which customers order which type of SmartTag.

The production process of SmartTag creates three places to place stock: at the component level (the PCBs), at the half fabric/assembly level (housing without cover), and at the end product level (SmartTag with customer specific cover). The first production part (casting) is the most restrictive production step in terms of capacity. The second production part (covering) is able to cover roughly 20% more than the casting. Furthermore, the casting is always done in a fixed batch size and the covering is done in variable batches sizes. Finally, Figure 2 shows an example of how these SmartTags are produced.

[CONFIDENTIAL]

Figure 2. Overview production SmartTag (example)

2.3. Product lay-out

As can be seen in Figure 2 there is a wide range of configurations of the SmartTag. In order to get a better understanding on how the configurations differ from each other, an illustrative example of the BOM of a SmartTag is provided in Figure 3. This figure shows the three echelon levels of the production. At the component level four different components exist. Moreover, these components results in six different half-fabrics.

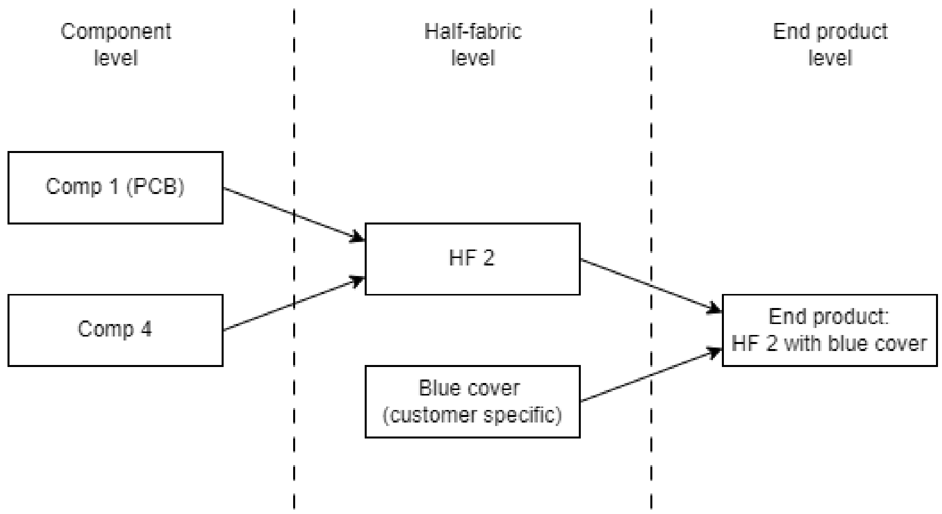

Figure 3. Example BOM SmartTag (HF 3)

2.4. Historical data

Now that we know how the supply chain, production, and BOM look like, we obtain quantitative data. This quantitative data consists of historical demand data and key performance indicators (KPI). The historical demand tells us about the distribution/importance of the products and customers. The KPIs provide insights into the performance of the current situation.

2.4.1. Demand data

The demand of the different SmartTags is not evenly distributed. Some tag types are more popular, and some customers demand more tags than others. In order to get a better understanding of which products are critical, demand information is required. Moreover, a distinction is made between fourteen different customers. In Table 1, an overview is created about the demand shares per product. This data is based on the last two years, since this time span represents the most accurate ratio between the products of the current situation. From this table it becomes clear that certain products are requested more than others.

Table 1. Demand share per product

[CONFIDENTIAL]

Next to that, Nedap currently uses different safety stock values for various components, based on their demand. Furthermore, Nedap has closer contact with some large customers, in order to provide a high service level for those customers with high demand. This contact results in a larger planning horizon/forecast from the customers. Since these customers have a higher demand, they plan their orders ahead of time. This is done to increase the chance that their order is delivered on time. When a sales order is planned early, Nedap has time to plan ahead and order the components early. Otherwise, components have to be purchased on the basis of a forecast, which comes with some uncertainties.

Besides deviation of products and customers, it is also interesting to see how product demand evolved over the years. In Figure 4 the demand from the individual and aggregated demand from the past two years can be seen. No clear trends or seasonality can be seen from Figure 4

[CONFIDENTIAL]

Figure 4. Demand Neck Tags last two years

2.4.2. Relevant parameters and KPIs

In addition to the demand data, we discus relevant parameters and KPIs. These include, lead times, safety stocks, and (minimum) order quantity. These values are provided on three distinct levels. Component, half-fabric and end product level.

Component level

The components are bought from external suppliers and each component has different values for parameters, which can be seen in Table 2. The lead time is supposed to be 6 weeks. Also, Nedap aims for a safety stock of roughly a month of demand. Finally, there exists a minimum order quantity which depends on the supplier and a fixed order size.

Furthermore, the customer specific covers are also components, although they are used on the second production part. Nedap SMART orders the covers from an external supplier, however, since they are cheap, have a short lead time, and are supplied by reliable suppliers, the covers are not evaluated within the process of the safety stock placement.

Table 2. Overview parameters of the components

Component	Lead time	Safety stock	Minimum order quantity
Comp 1	6 weeks	[Confidential]	[Confidential]
Comp 2	6 weeks	[Confidential]	[Confidential]
Comp 3	6 weeks	[Confidential]	[Confidential]
Comp 4	6 weeks	[Confidential]	[Confidential]

Half-fabric and end product level

Next to that, the half-fabrics are all produced by Nedap SMART from the components; hence all half-fabrics have the same lead time and batch size. The current safety stock for each individual half-fabric is zero, since Nedap decided not to stock any half-fabrics apart from the work in progress inventory. The end products are also produced by SMART and have a short production time and variable batch size. Finally, there is also no safety stock for the end products, since there has been no capacity for producing safety stock for the end products.

2.5. Conclusion current system analysis

In conclusion, the SmartTag has a wide range of configurations, which are made from four different components and colored covers. These SmartTags are produced in two steps. The demand of the individual SmartTags is difficult to estimate, however, we did find that the demand on half-fabric level was more predictable. Also, we argued that the covers are not a bottleneck, since the covers are cheap, have a short lead time, and have reliable suppliers. However, knowing that the covers are not the bottleneck is not a reason to eliminate the entire second production part. Therefore, we are dealing with a multi-echelon multi-item inventory system problem in a production environment. Hence, we need to conduct research on how we can approach this problem.

3. Literature review

In this chapter, literature is reviewed with the goal to gain more insights into multi-echelon multi-item inventory problems in a production environment. First, we need knowledge on the different production processes, see Section 3.1. Next to that, Section 3.2 is about the placement of safety stocks, with the focus on multi echelon safety stock placement in production. Moreover, in Section 3.3 the influence of capacitated production plannings on inventory decisions is discussed. Furthermore, Section 3.4 is about the application of simulation and heuristics in inventory optimization. Finally, Section 3.5 concludes our literature review by summarizing our key findings.

3.1. Safety stock placement

We are dealing with a multi-echelon system, which means we have multiple options to place safety stocks. In the case of Nedap there are three echelon levels. We have to determine which echelon level provides the best point to place safety stock. This depends on the decoupling point. From this decoupling point, a manufacturing strategy is derived.

3.1.1. Decoupling point

In order to decide on the most fitting production process, information is required on what distinguishes the different strategies. The key difference between the strategies is the placement of the safety stock. This is known as the customer order decoupling point (CODP), which is the point in the production process where the product is assigned to a customer specific order. Some papers refer to the CODP as the order penetration point (OPP) (Mihiotis, 2014). This determines essentially the production planning type, meaning if the production is MTS, MTO, Assemble-to-Order (ATO) or Engineer-to-Order (ETO). A general rule in this case is that the CODP needs to match the most important stock position (Olhager, 2010).

The CODP represents the point where customer order-driven activities are separated from the forecast-driven, see Figure 5. This influences the part in the supply chain that follows a push-strategy or a pull-strategy (Mason-Jones & Towill, 1999). The placement of the CODP comes down to finding a balance between cost and lead time. On the one hand, moving the CODP upstream results in a longer and less reliable lead time for the customer. On the other hand, moving the CODP downstream means that there will be higher stock costs, due to the increase in variety and longer Make-to-Stock time. Furthermore, van Donk (2001) provided four product and market characteristics to determine the decoupling point. These four points are the required delivery reliability, required delivery time, predictability of demand, and the specificity of demand. First, when the required delivery reliability is high, the decoupling point should be downstream in the supply chain. Placing stocks close to the customers promises high reliability, since there is less chance of disruption. Secondly, if the required delivery time is low, the decoupling point should also be downstream. Maintaining stock downstream the supply chain results in lower lead times, because no time is spent waiting for the production or supplies. Thirdly, if the demand is very predictable, the stock is also better kept downstream the supply chain. Keeping stock downstream the supply chain asks for an accurate and predictable forecast. When this is not the case, keeping stock more upstream results in reducing the risk of an inaccurate forecast by being order-driven rather than forecast-driven. However, keeping stock upstream the supply stream is only possible if the production times are acceptable. Fourthly, a large number of different end products asks for a more upstream decoupling point. When a great variety of products exists, it is better to keep stock upstream the supply chain with the purpose of reducing the risks, when product commonality is applicable (Hillier, 1999).

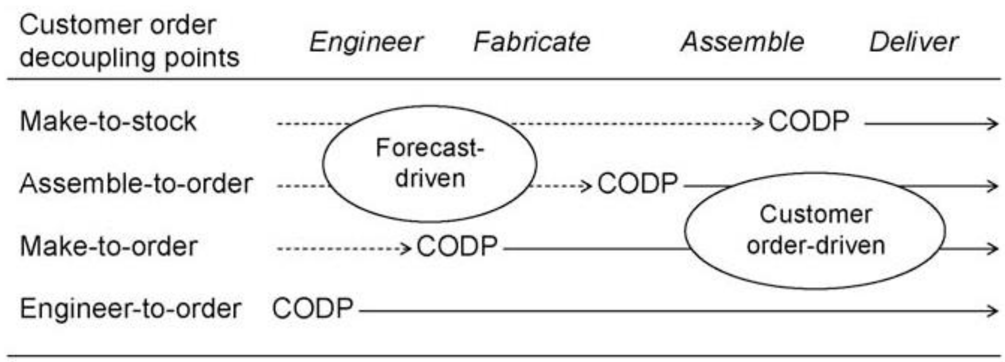

Figure 5. Customer order decoupling point
source: Olhager (2010)

3.1.2. Manufacturing strategies

Figure 5, shows four different manufacturing strategies, i.e., MTS, ATO, MTO, and ETO. Smart Tags are already engineered the same way for all customers, and because ETO means that customers order something that would be engineered to their liking, ETO is not relevant for this case. Therefore, we are only researching the remaining strategies: MTS, ATO, and MTO. Moreover, we also research hybrid options, hybrid options are for example combining MTS and MTO (Olhager & Prajogo, 2012; Rajagopalan, 2002; Zaerpour et al., 2008).

First of all, Meisel & Bierwirth (2014) point out that MTO is a great strategy for dealing with stochastic demand. When order arrivals and sizes are difficult to predict, it is beneficial to wait with production until an order actually arrives. However, producing according to a MTO strategy is only possible if the production time allows it. I.e., the production time should not exceed the lead time promised to the customer. Moreover, the production capacity should also be capable of handling demand without extending this promised lead time.

Second, MTS is considered a proactive strategy, in which items are pushed to the stock. Because it is not precisely known how much should be pushed, it is also called a speculative process (Chopra & Meindl, 2001). Applying MTS requires some decisions to be made about inventory policies, safety stock levels, order sizes, etc. Most of these decisions are based on the forecasts of the products. According to Köber & Heinecke (2012), especially companies that sell products with a high demand volume, standardized products, and low coefficient of variation can profit from an MTS environment. Furthermore, Nel & Badenhorst-Weiss (2010) make a distinction between functional and innovative products. Functional products result in a stable demand with a low uncertainty and the opposite holds for innovative products. Hence functional products suit an MTS environment better, and innovative products fit better to MTO. This is due to the characteristics of the products, as mentioned above.

Next to that, ATO can provide a balance between the MTS and MTO strategy. This is supported by the paper of Ghrayeb et al. (2009), who concluded that the ATO strategy inherits the strengths and conceals the weaknesses of both the MTS and MTO strategy. The ATO strategy has the customer order decoupling point in the middle of the production process. This means that the production follows a push strategy until the decoupling point and is pulled by customer orders after the decoupling point. This is especially beneficial in cases where assembly is not considered an intensive task and customers expect some form of customization (Du et al., 2005).

Finally, hybrid options exist where each product has its own manufacturing strategy. A common hybrid solution is the combination of MTO and MTS. There is a lot to be gained by combining these two strategies (Peeters & van Ooijen, 2020). Especially in situations where different products have different characteristics. Moreover, Sun et al. (2008) conclude that there is a trade-off between the maximum manufacturing efficiency that the MTS operations provides and the minimum inventory investment that comes with MTO operations, while maintaining a high service level.

3.2. *Safety stocks*

When the CODP is chosen, safety stock placement should be in line with this decoupling point. However, safety stock is not the only option to incorporate safety, this can also be done by incorporating safety time. Both safety stock and safety time build a buffer to handle the variation in demand. Safety stock does that by stocking more items than expected to account for the variability. Safety time however, adds more time than expected to account for the demand variability. Both methods are similar, one accounts for safety with stocks, the other with time. According to Iannone et al. (2022), safety time is favored over safety stock when the uncertainty is mostly in the timing of demand or replenishment. However, safety stock is favored over safety time in case the uncertainty lies particularly in the quantity of demand or supply. In the current situation uncertainty in demand is more applicable than the timing. Moreover, Grasso & Taylor (2007) argue that safety time is never preferred over safety stock. Hence, the safety stock is the focus of this research. Therefore, we conduct research on how to determine safety stock levels, after we researched the risk pooling effect.

3.2.1. *Risk pooling*

Grouping suppliers, products, or facilities in order to reduce the standard deviation of the demand is considered risk pooling. Hillier (1999), shows that commonality of products can have a great impact on the total inventory cost, which is considered risk pooling. Although risk pooling knows many more applications, the key concept remains the same: reducing deviation. Sobel (2008) summarizes this idea as follows: "The standard deviation of a sum of interdependent random demands can be lower than the sum of the standard deviations of the component demands.". This is done with the goal to reduce business risks and uncertainties. Furthermore, the paper of Benjaafar et al. (2005) examined inventory pooling in production-inventory systems. This showed the impact of utilization, demand and process variability, control policy, service levels, and the structure of the production on the safety stock level. Most results of Benjaafar et al. (2005) are in line with those of Eppen (1979). Eppen's research shows that a pooled system yields a lower cost when compared to a distributed system. Also, the risk pooling effect (of a pooled system) increases when the variance of demand increases. The opposite holds for the increase in correlation between demands. He uses the following formula, see Eq. (1), to calculate the pooled standard deviation with correlation, assuming normally distributed demand, in which the σ_{pooled} is the standard deviation of the pooled demand. This standard deviation of pooled demand is constructed by taking the square root of the sum of all the standard deviations of the products i that are included in the pooled demand and the correlation between the items. If there is no (significant) correlation between the different products, Eq. (1) can be rewritten as Eq. (2).

$$\sigma_{pooled} = \sqrt{\sum_{i=1}^{n} \sigma_i^2 + 2\sum_{i=1}^{N-1}\sum_{j=i+1}^{N} \sigma_i \sigma_j \rho_{ij}} \tag{1}$$

$$\sigma_{pooled} = \sqrt{\sum_{i=1}^{n} \sigma_i^2} \tag{2}$$

3.2.2. *Safety stock level*

In order to determine safety stock levels placed at the selected safety stock locations (in the supply chain), multiple formulas exist. Two common variations of deterministic safety stock formulas are assessed. The first one is from Chopra & Meindl (2007), see Eq. (3), this is considered as the basic safety stock calculation, which takes into account the following parameters of product j; the safety factor z_j, the standard deviation of the demand σ_j^D, and the lead time L_j.

Chopra & Meindl model:

$$SS_j = z_j \sigma_j^D \sqrt{L_j} \tag{3}$$

Second, Graves & Willems (2003) took another approach to determine the safety stock levels. This approach is called the Guaranteed Service Model (GSM). The GSM tries to find the ideal location and level of safety stock by evaluating the service times. For instance, when a product should be sent to the next stage in week 4, the outbound service time s_j^{out} is 4. When the product (ideally) arrives in week 2, the inbound service time s_j^{in} is 2. In the formulation of the GSM the (production) lead time PL_j is also taken into account, this is the lead time or production time at the current stage. With this information together with the safety factor z_j and the standard deviation of the demand σ_j^D, the safety stock can be determined by Eq. (4).

GSM by Graves & Willems:

$$SS_j = z_j \sigma_j^D \sqrt{s_j^{in} + PL_j - s_j^{out}} \tag{4}$$

From this formula it can be seen that if the inbound service time is the same as the outbound, the lead time is the only factor in the square root, and thus reduces to the C&M model. Moreover, the inbound and outbound service times are both decision variables in a minimization problem. By minimizing the inventory costs, the optimal service times are decided and thus also the safety stock levels. The GSM is a non-linear problem, which means that it is more difficult to solve because not all equations are linear.

3.3. *Capacitated production planning*

Inventory in a capacitated production system is impacted by variables such as production batch sizes, maximum production, and production times. When production time increases, total lead time increases as well, which will have an effect on the safety stock. Also, maximum production can influence lead time and thus the inventory level (Sitompul et al., 2008). For example, if the production capacity is reached by the demand, production is moved to the next period. This leads to increased lead time and thus also increased safety stock. Furthermore, batch sizes can influence the average inventory level. Larger batch sizes result in a higher average inventory and thus the total holding cost increases. Finally, Kumar & Aouam (2018) also show that reduced setup costs result in lower lead times and lower safety stocks.

Another model is provided by Kapuściński & Tayur (1998), who used a simulation-based approach in order to evaluate stock levels with respect to production capacity. The simulation followed a sequence of events. Each period starts with either inventory or backlog. Based on this value, production is executed, within its capacity limitation. After the start of production, the demand arrives, which concludes the period. Three basic insights were found according to the simulation study. First, an increase in demand results in an increase in the order up-to level. Second, a decrease in capacity caused the order up-to level to increase. And finally, an increase in demand variance also increases the order up-to level.

3.4. Inventory-production problem

As discussed in Section 3.3, production constraints have a significant impact on inventory decisions. Therefore, we conduct research about existing methods for the inventory-production problem, which incorporates both inventory and production constraints.

3.4.1. Inventory-production methods

Inventory-production models are considered NP-hard, which means that the optimal solution cannot be obtained in polynomial time (Bylka & Rempala, 2007). Therefore, different simulation models, machine learning algorithms and heuristics exist in order to obtain (near-) optimal values for safety stock values and other decisions variables (Jung et al., 2004).

Taleizadeh et al. (2016) consider an approach to a three-level distribution network problem. Their paper consists of a production-inventory model, which includes backlogging and the possibility of rejected production batches. A heuristic algorithm is used to minimize the total cost of the problem. Next to the heuristic algorithm, the literature also provides us with machine learning algorithms. Pirhooshyaran & Snyder (2020), apply deep neural networks (DNN) to optimize inventory decisions in multi-echelon supply chains. According to their research, their application is able to outperform alternative methods in many complex supply chains. In addition, Zhao & Sun (2010) provide a multi-agent reinforced learning (RL) model to optimize inventory levels, which in most cases is able to obtain better results than genetic algorithms (GA). Another algorithm that is commonly used for inventory-production problem is simulated annealing (SA). Haddock & Mittenthal (1992) investigated the feasibility of using SA in combination with a simulation model. Also, in the paper of La Fata & Passannanti (2017), SA is applied to an inventory optimization problem. This paper shows the application of simulated annealing, in combination with a Monte Carlo simulation, on inventory profiles.

3.4.2. Simulated annealing

Applying the SA algorithm effectively requires a sufficient understanding of the parameters involved, since these parameters influence the performance of the algorithm. Therefore, we require more information about the effect of the SA parameters. For this information we consult the paper of Brusco (2014). The first decision which is discussed in the paper, is the generation of the initial solution. Simulated annealing needs an initial solution in order to be able to create a neighbor solution. This initial solution can be generated completely randomly, or according to a certain heuristic. Secondly, the start temperature $T(1)$ and end temperature $T(Q)$ are discussed. The temperature determines the balance between exploring and exploiting. In this context exploration is about extending the search area in the solution space, whereas exploitation is about focusing on the current solution and improving it. Simulated annealing might accept worse neighbor solutions in order to escape local optima and explore different areas. Whether a worse solution gets accepted is determined by the Metropolis rule. This rule accepts the solution based on a randomly generated number, the difference in solution values, and the current temperature. Choosing the right temperature is important, because temperatures which are too high only explore and never converge to an optimal value. However, temperatures which are too low, most likely get stuck in a local optimum. Next to that is the cooling scheme, this determines how quickly the temperature decreases (with a factor c). The cooling schemes discussed by Brusco (2014) are the linear, exponential, reciprocal, and logarithmic approach. The cooling scheme in combination with the temperature range determines the number of iterations Q. Another key decision is the Markov length m, this determines how many neighbors are examined for a single temperature. Increasing the Markov length results in a longer run time, but a higher chance of finding a better solution, since more (different) solutions are evaluated. Finally, the last key decision is how the neighbors are selected. This decision depends heavily on how solutions are determined in the first place. An overview of the simulated annealing algorithm is given in Figure 6.

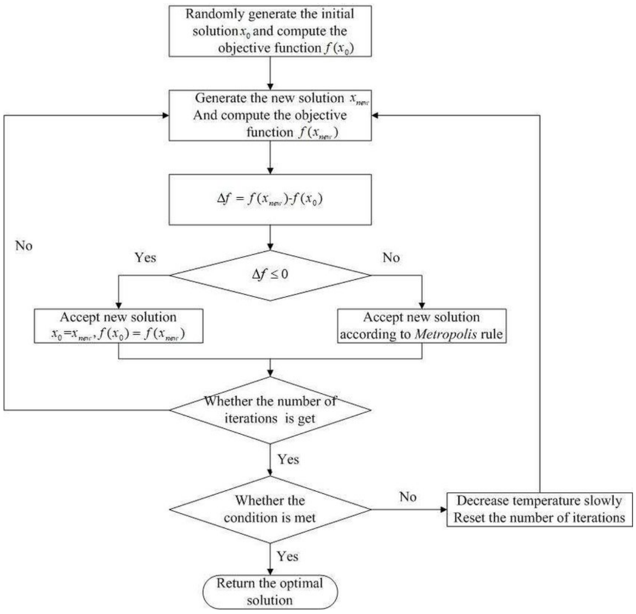

Figure 6. A flowchart of the simulated annealing algorithm
source: Zhou et al. (2018)

Brusco (2014) has conducted two computational studies in order to test different settings. From these studies it can be derived how the settings performed, which form a basis for our settings. First of all, according to the computational studies no initial subset was superior to the other, therefore no best settings can be selected. Secondly, the best initial temperature is given by the average difference between the current and neighbor solution's objective value. Next to that, the best end temperature is given by the start temperature divided by the log value of the total iterations. The study is designed to produce 500,000 iterations. Furthermore, there was no cooling scheme which always outperformed another. However, exponential cooling schemes are considered one of the most popular cooling schemes, hence we choose this one. In Table 3 a summary of the selected settings can be found.

Table 3. Simulated annealing settings

Setting	**Value**
$Iterations: Q$	500,000
$Start\ temperature: T(1)$	avg. $\Delta f(solution)$
$End\ temperature: T(Q)$	$\frac{T(1)}{\log Q}$
$Decrease\ factor: c$	$exp(\frac{log(T(Q)) - log(T(1))}{Q-1})$
$Initial\ subset$	$t.b.d.$

3.5. Conclusion literature

The literature knows many models to determine where safety stock should be placed and how high these stocks should be. We are evaluating the performance of models suggested in Chopra & Meindl (2007) and Graves & Willems (2003). Both models can be used without making any assumptions. However, these models do not include all of the constraints relevant for Nedap's situation. The most important factor that is left out is the production limitation. Sitompul et al. (2008) provide a model which does account for production capacity; however, this is only limited to serial systems. Since the supply chain at Nedap is a network, this does not apply. Finally, Kapuściński & Tayur (1998) propose a simulation-based approach to optimize the safety stock placement in a multi-echelon multi-item system. Therefore, a simulation is used to optimize the safety stock placement and thus the base stock levels. Moreover, simulated annealing promised to be a suited metaheuristic which works well in combination with simulations. In order to get a better understanding of the effect of the different hyperparameters of the SA algorithm, we used the computational studies from Brusco (2014) to determine our settings.

4. Method

The goal of this thesis is to provide a method which can optimize the decisions regarding inventory management in a multi-echelon supply chain. In order to achieve optimal inventory parameters, a simulation is created to evaluate the performance of different approaches. Moreover, multiple assumptions and settings regarding the simulation model are discussed. After the framework of the simulation is created, different models derived from literature are discussed in more detail. Next to that, a simulation optimization framework is provided in order to optimize the inventory parameters for the SmartTags at Nedap. With the simulation model, the existing models are evaluated together with the proposed simulation optimization.

4.1. Simulation model

The first step is to create a simulation model which represents the current process on a tactical level. By simulating the different events that occur and the corresponding decisions that follow, it becomes possible to evaluate certain decisions (input). Since the scope of the simulation is on a tactical level, certain assumptions are made, and operational decisions are left out. In this section the framework of the simulation, together with the input/output and the assumptions is discussed. Also, an overview of some terms within the context of this simulation for Nedap are explained in Table 4. This additional explanation is required because the meaning of these terms is not straightforward in this context.

Table 4. Explanation of specific terms

Term	Explanation
Order	Generated demand for a single type of end product k
Inventory order	An order to produce inventory on half-fabric level
Evaluated order	An order that has been already evaluated in the current week
Unfulfilled order list	A list with all unfulfilled orders (demand)
Cover capacity	The capacity for the covering of the half-fabrics (second step in the production process)
Half-fabric requirements	The required number of half-fabrics as a result from the orders minus the on-hand inventory

4.1.1. Simulation description

In this section the decisions regarding the simulation model (analyses and evaluation) are explained. This simulation can be divided into four parts. The first part is generation of the demand and updating of the orders and variables such as inventory. The second part is evaluation of the demand that can be met from the half-fabrics directly (production step two). The third part is evaluation of the number of batches that should be created from components to half-fabrics (production step one). The last part evaluates the possibilities to create additional batches with the goal to create inventory at half-fabric level. Every week these four parts are evaluated sequentially. When the last part is evaluated, the first part of the next week is evaluated. This process is repeated until the run length of the simulation is reached. A more detailed overview of the entire process can be found in Appendix B.

4.1.2. Input/start

When the simulation model is initiated, all relevant variables and parameters are initialized. In short, this step determines how many weeks the simulation runs, which values the parameters take, and what the base stock levels are. The output at the end is generated based on these inputs, which is why careful assessment of which inputs are taken into account is necessary. In Table 5, Table 6, Table 7, and Table 8 an overview is provided of all relevant sets, parameters, and (decision) variables, respectively.

Table 5. Overview sets

Sets	Explanation
$i \in \{1,2,\ldots, MaxNumberComp\}$	Component i
$j \in \{1,2,\ldots, MaxNumberHF\}$	Half-fabric j
$k \in \{1,2,\ldots, MaxNumberEP\}$	End product (Covered half-fabric) k
$z \in I \cup J$	Element z of set of components i and half-fabrics j
$w \in \{1,2,\ldots, RunLength\}$	Week number w
$m \in \{1,2\}$	Machine m
$p \in \{1,2\}$	Production step p (1: component i to half-fabric j, 2: half-fabric j to end product k)

Table 6. Overview parameters

Parameters	
$MaxNumberComp \in \mathbb{N}$	Number of different components
$MaxNumberHF \in \mathbb{N}$	Number of different half fabrics
$MaxNumberEP \in \mathbb{N}$	Number of different end products
$RunLength \in \mathbb{N}$	Number of weeks the simulation runs.
$WarmUpPeriod \in \mathbb{N}$	Number of weeks the simulation warms up (not including output)
$ProdCapacity_{m,p} \in \mathbb{N}$	The production capacity of machine m for production step p
$BatchSize \in \mathbb{N}$	The batch size
$A_{ij} \in \{0,1\}$	Indicator variable, which equals 1 if component i is required for half-fabric j and 0 otherwise. (Values derived from BOM)
$B_{jk} \in \{0,1\}$	Indicator variable, which equals 1 if half-fabric j is required for end product k and 0 otherwise.
$D_{k,w} \in \mathbb{N}$	Demand follows from a normal distribution (rounded to integer) for week w for end product k with μ_k and σ_k
$\mu_k \in \mathbb{R}$	Mu (mean) for end product k
$\sigma_k \in \mathbb{R}$	Sigma (standard deviation) for end product k
$\mu_z \in \mathbb{R}$	Mu (mean) for component / half-fabric z
$\sigma_z \in \mathbb{R}$	Sigma (standard deviation) for component / half-fabric z
$CSL_z \in [0,1]$	Target cycle service level for component / half-fabric z
$SS_z \in \mathbb{N}$	Safety stock for component / half-fabric z
$L_i \in \mathbb{N}$	Lead time for component i
$MOQ_z \in \mathbb{N}$	Minimum order quantity for component / half-fabric z
$PromisedLeadTime_k \in \mathbb{N}$	The lead time for end product k set by Nedap that is promised to the customer
$ServiceLevel_k \in [0,1]$	Achieved service level for end product k

Table 7. Overview variables

Variables	
$EvalWeek \in W$	The week number that is evaluated
$UnfulfilledOrder_{k,w} \in \mathbb{N}$	The number of end products k, which are still unfulfilled in week w
$BatchesPlanned_{j,w} \in \mathbb{N}$	Batches planned for half-fabric j in week w
$CoveredPlanned_{k,w} \in \mathbb{N}$	Number of end products k covered in week w
$CurrentInventory_z \in \mathbb{N}$	Current inventory of component / half-fabric z
$CurrentInventoryPosition_z \in \mathbb{Z}$	Current inventory position of component / half-fabric z
$Inventory_{z,w} \in \mathbb{N}$	Inventory of component / half-fabric z at the start of week w
$InventoryOrders_j \in \mathbb{N}$	The number of half-fabrics j to produce for increasing the inventory
$RequiredDemand_{j,w} \in \mathbb{N}$	The number of half-fabrics j required to be produced in week w
$OrderArrival_{i,w} \in \mathbb{N}$	Number of components i that arrive at the start of week w
$BackOrder_{k,w} \in \mathbb{N}$	The number of end products k in backorder in week w (unfulfilled demand)
$Delay_{k,w} \in \mathbb{N}$	The total delay for end product k from the order placed in week w
$DemandMetFromInventory_w \in \mathbb{N}$	Auxiliary variable to determine the number of end product covered directly from half-fabric inventory in week w
$InventoryShort_{j,w} \in \mathbb{N}$	Auxiliary variable to determine the number of half-fabrics j short in week w
$DelayOrder_{k,w} \in \mathbb{N}$	The number of weeks the order from week w is delayed for end product k
$OrderInTime_{k,w}$	Auxiliary variable to determine the number of 'late' orders

Table 8. Overview decision variables

Decision Variables	
$S_z \in \mathbb{N}$	Base stock level for component / half-fabric z

4.1.3. Part one (Yellow)

After the initiation of the input variables and the current week are selected, three steps need to be considered in order to start the week. The first one is updating the inventory parameters. For the first week, the inventory(position) is set to the reorder point. The second step is where the demand is generated for the entire week. It is assumed that all the demand arrives at the start of the week and that it is normally distributed (Section 4.2). Moreover, in order to be able to compare different simulations, demand is generated with pseudo random numbers. The third step of this part is to keep track of a list with all demand that is not completed yet. The values in this list represent the demand that is (still) unmet given an end product k at a given week. In Figure 7 a flowchart of these steps is provided.

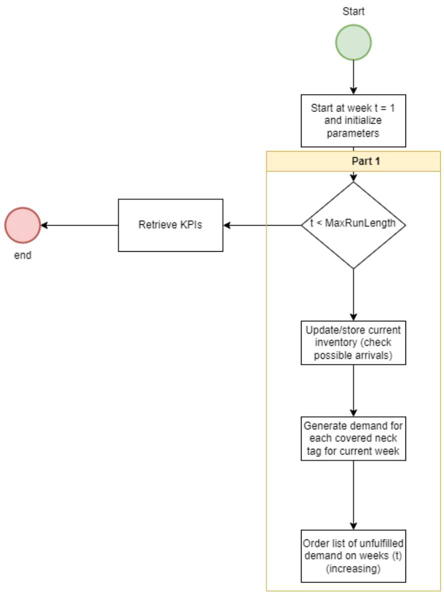

Figure 7. Flowchart part 1 (yellow)

Step 1 (for all items in z):

Inventory (position) for the first week:
$Inventory_{z,1} = CurrentInventory_z = CurrentInventoryPosition_z = S_z$

Inventory for the other weeks:
$CurrentInventory_i = CurrentInventory_i + OrderArrival_{i,w}$
$Inventory_{z,w} = CurrentInventory_z$

Step 2/3 (for all end products k):

Generate demand normally for end product k with μ_k and σ_k. Then add (positive) demand to list of unfulfilled orders. Note that the generated demand $D_{k,w}$ is rounded to the nearest integer. $UnfulfilledOrder_{k,w} = \max\{0, D_{k,w}\}, where\ D_{k,w} \sim N(\mu_k, \sigma_k)$

4.1.4. Part two (Green)

When the list of unfulfilled orders is updated, we start by finding the first week for which unfulfilled orders still exist. For every week in the unfulfilled order list the steps in this part should be taken. When the list is empty or all items in the list are already evaluated, part three is initiated immediately. If there are still unevaluated orders on the list, the following steps are taken.

First, we take the first unevaluated week from the list. For the selected week, we go over the half-fabrics based on expected demand. In practice this means that we first select the half-fabric with the highest expected demand. Next, we determine how many end products can be directly produced from half-fabrics which are on stock. This number depends on the inventory of the half-fabric and production capacity (of the second production part). From this, it follows how many half-fabrics are required to be produced. With this the inventory levels can be updated, and the next covered half-fabric can be selected. Finally, with the updated inventory position we can determine if we have to place an inventory order. When all six end products have been evaluated, we go to part three. This second part of the simulation is shown in Figure 8.

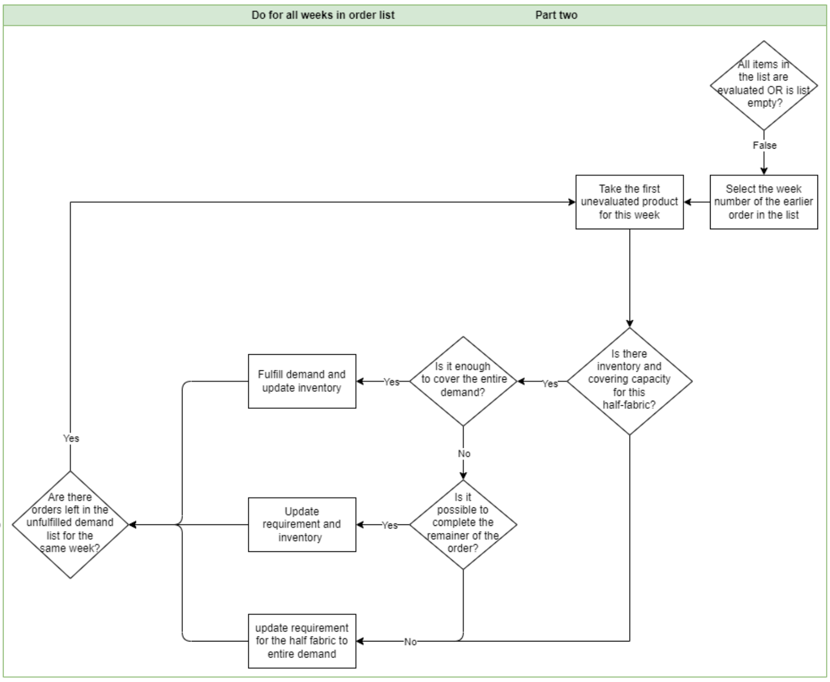

Figure 8. Flowchart part 2 (green)

Step 1:

Check if there are unfulfilled orders on the list, if not go to part three.

Step 2 (do for all half-fabrics *j* (high-low)):

Select the current week to evaluate:
$EvalWeek = 'first\ unevaluated\ week'$

Determine how many end products can be covered from stock:

$$DemandMetFromInventory_w = \begin{cases} \sum_k \left(UnfulfilledOrder_{k,EvalWeek} * B_{jk}\right), & \text{if } \sum_k \left(UnfulfilledOrder_{k,EvalWeek} * B_{jk}\right) < CurrentInventory_j \\ CurrentInventory_j & \text{, else} \end{cases}$$

Check if the production capacity of the covering is not exceeded:

$$DemandMetFromInventory_w = \begin{cases} DemandMetFromInventory_w \\ \quad \text{, if } DemandMetFromInventory_w < \left(CoveredPlanned_{k,w} - \sum_m ProdCapacity_{m,2}\right) \\ CoveredPlanned_{k,w} - \sum_m ProdCapacity_{m,2} \\ \quad \text{, else} \end{cases}$$

Update inventory (position):

$$CurrentInventory_j = \begin{cases} CurrentInventory_j - DemandMetFromInventory_w & \text{, if } DemandMetFromInventory_w < CurrentInventory_j \\ 0 & \text{, else} \end{cases}$$

$$CurrentInventoryPosition_j = CurrentInventoryPosition_j - DemandMetFromInventory_w$$

Update production capacity:

$$CoveredPlanned_{k,w} = CoveredPlanned_{k,w} + DemandMetFromInventory_w$$

Determine how many half-fabrics should be produced:

$$RequiredDemand_{j,w} = \begin{cases} \sum_k \left(UnfulfilledOrder_{k,EvalWeek} * B_{jk} \right) - DemandMetFromInventory_w & , \text{if } DemandMetFromInventory_w < \sum_k \left(UnfulfilledOrder_{k,EvalWeek} * B_{jk} \right) \\ 0 & , \text{else} \end{cases}$$

Step 3 (do for all half-fabrics *j*):

Place an inventory order if the inventory position is below the reorder point:

$$InventoryShort_{j,w} = \max\{0, S_j - CurrentInventoryPostion_j\}$$

$$Inventoryorders_j = \begin{cases} Inventoryorders_j + \max\{MOQ_j, InventoryShort_{j,w}\}, & \text{if } InventoryShort_{j,w} > 0 \\ Inventoryorders_j & , \text{else} \end{cases}$$

4.1.5. Part three (Blue)

When we arrive at this part, we know how many half-fabrics should be produced (see $RequiredDemand_{j,w}$). Based on these values we decide how many batches we want to produce. This, however, is constraint by the production capacity and the inventory of the relevant components. Similar to part two, we go over the half-fabrics (from highest to lowest mean) and determine per half-fabric how many batches are produced. This is done in two steps. First, we try to plan "full" batches, meaning that the entire batch is used to fulfill (part of) the $RequiredDemand$. Next when all the "full" batches have been planned, we plan one more "left-over" batch if there is still $RequiredDemand$ for this half-fabric. We either plan one or none in this case, since we already evaluated if "full" batches should be created. After this, the inventory of the components is evaluated, and we go back to part two. An overview of these steps is provided in Figure 9.

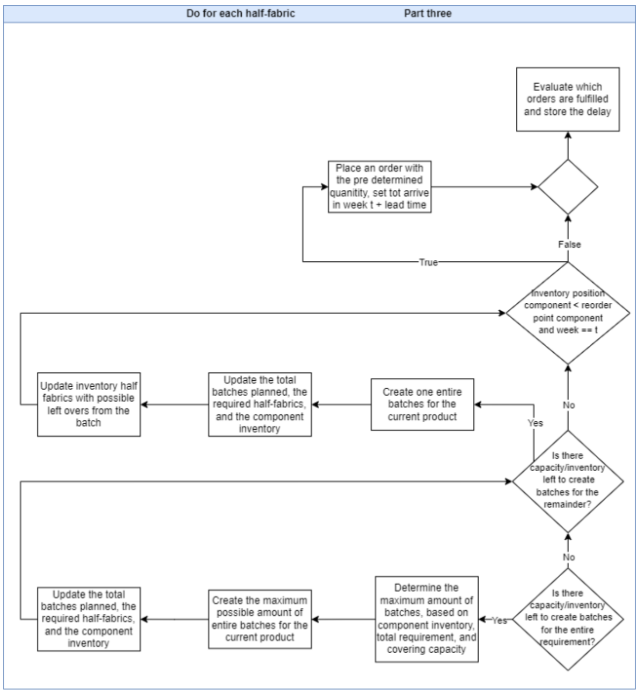

Figure 9. Flowchart part 3 (blue)

Step 1 (do for all half-fabrics j (high-low)):

Calculate the maximum number of batches that can be planned:
$MaxPossibleBatches = \min\left\{\sum_m ProdCapacity_{m,1} - \sum_j BatchesPlanned_{j,w} * BatchSize, \sum_i A_{i,j} * CurrentInventory_i \; div \; BatchSize, \; RequiredDemand_{j,w} \; div \; BatchSize\right\}$

Add the number of "full" batches to the planned batches:
$BatchesPlanned_{j,w} = BatchesPlanned_{j,w} + MaxPossibleBatches$

Update Required demand:
$RequiredDemand_{j,w} = RequiredDemand_{j,w} - MaxPossibleBatches * BatchSize$

Step 2 (do for all half-fabrics j (high-low)):

Determine of one "left-over" batch can be planned:
$MaxPossibleBatches = \min\left\{\sum_m ProdCapacity_{m,1} - \sum_j BatchesPlanned_{j,w} * BatchSize, \sum_i A_{i,j} * CurrentInventory_i \; div \; BatchSize, \; 1\right\}$

Add the "left-over" batch to the planned batches:
$BatchesPlanned_{j,w} = BatchesPlanned_{j,w} + MaxPossibleBatches$

Update the inventory for the half-fabrics as a result of the "left-over" batch:
$$CurrentInventory_j = \begin{cases} CurrentInventory_j + MaxPossibleBatches * BatchSize - RequiredDemand_{j,w} \\ \qquad \text{, if } MaxPossibleBatches = 1 \\ CurrentInventory_j \\ \qquad \text{, else} \end{cases}$$

Update Required demand:

$$RequiredDemand_{j,w} = \begin{cases} 0 & \text{, if } MaxPossibleBatches = 1 \\ RequiredDemand_{j,w} & \text{, else} \end{cases}$$

Step 3 (do for all components i):

Update the Order arrival and Inventory position:
$$OrderArrival_{i,w+L_i} = \begin{cases} \max\{MOQ_i, S_i - CurrentInventoryPosition_i\} & \text{, if } CurrentInventoryPosition_i < S_i \\ OrderArrival_{i,w+L_i} & \text{, else} \end{cases}$$

$$CurrentInventoryPosition_i = \begin{cases} CurrentInventoryPosition_i + \max\{MOQ_i, S_i - CurrentInventoryPosition_i\} \\ \qquad \text{, if } CurrentInventoryPosition_i < S_i \\ CurrentInventoryPosition_i \\ \qquad \text{, else} \end{cases}$$

Step 4 (do for all end products k and half-fabric j, in which B_{jk}=1):

Update order list, determine delay if order is complete.
$delayOrder_{k,EvalWeek} = w - EvalWeek, \text{ if } RequiredDemand_{j,w} = 0$

$$UnfulfilledOrder_{k,EvalWeek} = \begin{cases} RequiredDemand_{j,w} & \text{, if } RequiredDemand_{j,w} > 0 \\ 0 & \text{, else} \end{cases}$$

4.1.6. Part four (Purple)

At this part the end of the simulated week is reached, after no more unevaluated items are on the order list. This step checks if it is possible to produce an extra batch at the end of the week. This can only be done if all of the following three conditions are met: (1) there is demand for more inventory in the form of an inventory order, (2) there are enough components to make at least one batch of the inventory order, (3) there are still unplanned batches left in the current week). If all these conditions are met, the batch(es) is (are) created and the relevant components are updated. When these steps are taken, it is determined whether the set run length is reached. If the run length is not reached, the next week is simulated. If the run length is reached, the KPIs are calculated. This final part of the flowchart is presented in Figure 10.

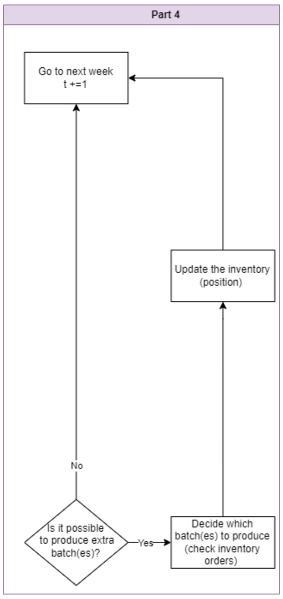

Figure 10. Flowchart part 4 (purple)

Step 1 (for each half-fabric *j* (high-low)):

Check if inventory orders exist and can be produced:

$$MaxPossibleBatches = \min\{(\sum_m ProdCapacity_{m,1} - \sum_j BatchesPlanned_{j,w} * BatchSize)\ div\ BatchSize, \sum_i A_{i,j} * CurrentInventory_i\ div\ BatchSize\ , InventoryOrders_j\ div\ BatchSize\}$$

Update batches planned:

$$BatchesPlanned_{j,w} = BatchesPlanned_{j,w} + MaxPossibleBatches$$

Update inventory orders planned:

$$InventoryOrders_j = InventoryOrders_j - MaxPossibleBatches * BatchSize$$

Update half-fabric inventory planned:

$$CurrentInventory_j = CurrentInventory_j + MaxPossibleBatches * BatchSize$$

Update component inventory planned (for each component i):

$$CurrentInventory_i = CurrentInventory_i - A_{i,j} * (MaxPossibleBatches * BatchSize)$$

4.1.7. Output/end

Finally, if all weeks are evaluated and run length is reached, important data can be stored. This is all the data relevant to determine the KPIs. From the equations below it can be seen which KPIs are used and how they are calculated. The service level and average inventory are used to determine the objective value of the solution. The service level is calculated based on the lead time Nedap ideally promises its customers, which in this case is five weeks. Next to that, in order to make sure that the output satisfies a certain accuracy, the run length, replication, and warmup period are determined. An explanation of the methods used to determine these values can be found in Section 4.3.

Outputs:

$$OrderInTime_{k,w} = \begin{cases} 1 & \text{, if } delayOrder_{k,w} < PromisedLeadTime_k \\ 0 & \text{, else} \end{cases}$$

$$ServiceLevel_k = \frac{\sum_k \sum_w OrderInTime_{k,w}}{NonZeroOrder}$$

, in which $NonZeroOrder$ is the number of orders for which demand was not 0

$$avgServiceLevel = \frac{\sum_k ServiceLevel_k}{MaxNumberEP}$$

$$avgInventory_z = \frac{\sum_w Inventory_{z,w}}{RunLength}$$

$$avgInventory = \frac{\sum_z avgInventory_z}{(MaxNumberHF + MaxNumberComp)}$$

$$avgUtilization = \frac{\sum_w \left(\frac{\sum_j BatchesPlanned_{j,w}}{\sum_m ProdCapacity_{m,1}} \right)}{RunLength}$$

$$avgOrderInterval_i = \left(\frac{NumArrivals_i}{RunLength} \right)^{-1}$$

$$avgOrderSize_i = \frac{\sum_w OrderArrival_{i,w}}{NumArrivals_i}$$

, in which $NumArrivals_i$ is the total number of orders that arrived for component i

4.2. Assumptions

Throughout the simulation, many assumptions have been made in order to model the current situation close to reality on a tactical level. In this section, these assumptions are discussed per category, in order to get a better understanding of the key decisions made in the model.

4.2.1. End products (covered half-fabrics)

First, we made one assumption concerning the end product, which is formulated:

- *There are always ample customer-specific covers on stock, therefore no distinction is made between end products from the same tag type.*

As explained in Figure 2, an end product consists of a half-fabric (tag type) and a customer specific cover. The customer specific cover is always available, due to its low cost and low lead time. Hence, the end products are presented as cover half-fabrics rather than the customer specific end product. In addition, the half-fabrics do need to be covered and this cover production capacity is not infinite. For this reason, we do make a distinction between half-fabrics and end products (covered half-fabrics).

4.2.2. Demand

Second, we have one assumption regarding the demand, which is:

- *Demand of end products is normally distributed*

We decided that the end products are end products rather than customer specific SmartTags. Because of this assumption, we generate demand as end products. As a consequence, demand can be estimated more accurately, due to grouping on half-fabric level (risk pooling). These grouped values are assumed to be normally distributed. Assuming normality also has the benefits of integrating well with the models from the literature. Since these models also assume normal distributed demand, we are able to use the same demand parameters in our model and the literature models. Furthermore, the demand generated from the normal distribution provided values close to the historical data. In Appendix D: Q-Q plot demand, the fit of the normally distributed demand on the historical data is provided.

4.2.3. Orders

Furthermore, we made multiple assumptions regarding the creation and content of an order. These are presented below.

- *Orders consist of only one type of item (an end product)*

In the simulation, the terms order and order list are used to define the generated demand. However, these orders are not simulated from customers, they are generated by the product specific demand. Since the interest of this simulation lies on the products on a tactical level, the demand is generated accordingly.

- *All orders arrive at the beginning of the week*

The demand, and thus the new orders, are generated at the start of the week. Because we simulate on a tactical level, we are interested in the availability of the products rather than the exact fulfillment of customer orders. Therefore, it makes sense to generate the end product demand at the start of the week, since this allows us to plan the generated demand in the current week.

- *Inventory orders are fulfilled at the end of the week*

Inventory orders for the half-fabrics are only produced after the regular demand is handled and are therefore only fulfilled at the end of the week. The produced half-fabrics are added to the inventory of next week, since no more demand follows after the inventory orders in that week.

- *Orders are completed when the entire order is fulfilled.*

Orders cannot always be fulfilled in the same week. When 99% of an order is fulfilled in the first week, but the other percent in the next week, the delay of the entire order is assumed to be one week. In practice, this is much different, since this one large order consists of multiple customer specific orders, meaning that many customers receive their order in the same week. However, the focus is on the specific product, not on the customer. Therefore, it is assumed that the entire order should be fulfilled in order to decide the total delay.

- *The products get evaluated by the simulation in a fixed order*

As can be seen in Section 4.1, the items are evaluated in a fixed order per week. Orders cannot be evaluated at the same time, since planning a batch for one product could mean there is no more capacity for another. For this reason, we assume that the products with the highest expected demand are always evaluated first.

4.2.4. Production

In addition, the assumptions affecting the production are listed below.

- *Production time (component to half-fabric and half-fabric to end product) is zero, if it does not exceed the capacity*

The production time of producing one batch (at production part one) is half a day. The production time for part two depends on the size of the batch, however a production capacity is incorporated. Because these production times are less than one week, they are assumed to be zero. This means that if there is production capacity and the components are available, production can be done within the same week. Next to that, there are currently two different production locations. However, since they have the same type of production line and batch size, it is assumed that the production is done on one big machine. This is a reasonable assumption, since there is no significant difference on a tactical level.

- *Only full batches are created*

Currently, it is decided that only full batches are produced. This also translates to the size in which half-fabrics and components are ordered. Ordering, for instance, two and a half batches worth of components would not make sense, since only entire batches are produced, which results in leaving that half batch as unused inventory.

- *98% of the total production capacity is used, 2% failure rate*

In practice the production can be on hold due to failure and because of planned days off. On average the production is down 2% of the time (on a yearly basis, a year equals 50 weeks). Because of this we assume that every week we have 98% of the total capacity.

4.2.5. Components

Next, we made two assumptions about the components, which are:

- *Component lead time is deterministic*

The simulation models only accounts for deterministic lead times, because of three reasons. The first reason is that the models we derived from literature also use deterministic lead times. In order to evaluate the performance, the variables should be the same. Secondly, including stochastic lead times would increase the number of different scenarios significantly. Since we already have an enormous number of scenarios to cover, it would only decrease the accuracy of the model to incorporate even more stochasticity. The last reason is simplicity and lack of accurate information. Because the lack of information about the actual lead time from the suppliers, assuming deterministic times makes for easy calculation. If the simulation model had to incorporate stochastic lead times, it would need accurate numbers for it to add value. However, there is currently no information at Nedap about how the lead time is distributed in practice. Nevertheless, we include the lead time in our sensitivity analysis in Section 5.2, which provides insights into the effect of the lead time on the objective.

- *Components are always delivered in the same amount as ordered, which is always a multiple of the batch size*

Besides assuming that the components always arrive on time, it is also assumed that orders arrive in the same amount as ordered. Practice shows that orders sometimes get partially fulfilled in time. However, the simulation assumes that when an order is placed, the entire order is received on the scheduled time. Moreover, orders are always placed as a multiple of a batch size.

4.2.6. Inventory policy

The final assumptions are about the inventory policy we selected.

- *The holding cost is equal on all levels and for all items*

Currently, Nedap does not care for a specific holding cost. They do not prioritize one inventory placement over another, meaning that it does not matter for the holding cost, if there is one component or one half-fabric on stock.

- *No significant fixed ordering cost for orders with a minimum size are incorporated.*

Next to the holding cost, there are also no significant fixed ordering costs since orders are always placed with a minimum order size at the supplier. Because the minimum order size is relatively large, no fixed ordering cost are considered.

- *A base stock policy is used with a* minimum order quantity *(MOQ)*

The policy used to order from suppliers and create inventory orders at half-fabric level, is a base stock policy. A base stock policy is a great policy, because we do not have a fixed ordering cost which means that we can order as many times as we want. However, there does exist a MOQ, which means that we can only place an order if it is above a certain amount. Therefore, the base stock policy cannot be optimally used. As can be seen in the formula below (from step 3 in part three), the ordered quantity is the maximum of the items short and the MOQ.

$$Ordered\ quantity = \max\{MOQ_i, S_i - CurrentInventoryPosition_i\}$$

- *The base stock levels are always a multiple of the batch size*

Since only full batches are created, there is no need for base stock levels that are not a multiple of the batch size. For example, if there are only enough components left to produce half a batch, these components will not be used. Only when extra components are ordered to create an entire batch, these components will be used. However, having this inventory in the first place only results in higher inventory level, without improving the service level. The decision to choose base stock levels as multiples of batch sizes, results in a significantly smaller solution space.

4.3. Validation/settings

In order to make sure that the simulation models exactly what we want it to model, we validated the model by presenting the simulation model to a small group of experts from the company. These experts were able to validate the decisions and assumptions made in the simulation. Furthermore, the accuracy of the output of the simulation is determined by a couple of parameters. These parameters are warm-up period, run length, and number of replications. These values have been calculated with the sequential procedure (Law, 2014, p. 505). The execution of this approach can be found in Appendix A.

Warm up period	15 weeks
Run length (excl. warm up)	500 weeks (10 ‘work’ years)
Number of replications	400 replications
Expected run time	146.05 seconds

4.4. Literature models description

In Chapter 3 we found different models to determine the safety stock placement and thus base stock levels. However, not every model includes all the relevant details relating to the production process at Nedap. In this section it is explained how different models from the literature can be used to determine the base stock levels.

4.4.1. Basic safety stock calculation

The first way to determine the safety stock placement is by the safety stock calculation provided by Chopra & Meindl (2007). We want to optimize the base stock levels for the entire supply chain of Nedap. In order to do this, we can use the safety stocks to determine the base stock levels. Note that safety stock on end product level is not taken into account, due to the assumptions. This means that we either place safety stock levels at component level or half-fabric levels, which is similar to a MTO or ATO strategy. Both strategies are evaluated separately.

The inputs for the model are as follows. Although it is assumed in the simulation that the production time is zero weeks, in practice not all the resources are always available, which is why we use a production time of one week for this model. We expect this one week to account for delay due to the unavailability of production capacity. However, it is difficult to say if this one week is valid, which is the reason for making a simulation in the first place. Also, the standard deviation of the demand will decrease as the decoupling points moves away from the customer, due to risk pooling, as discussed in Section 3.2.1. The last decision that should be made is about the cycle service level, which is set to the fraction of cycles in which no stock outs occur. Currently, we use a cycle service level of 95%. These values provide enough information to determine the safety stocks and thus the base stock levels. It is shown in the equations below, how these values are determined.

Make-to-Order model (component level) (Hereafter referred to as C&M MTO model)

Base stock level S_i:
$S_i = SS_i + \mu_i * L_i$, in which SS_i is the safety stock and $\mu_i * L_i$ is the demand during lead time. SS_i is calculated as follows:

Safety stock SS_i:
$SS_i = z_i * \sigma_i^D * \sqrt{L_i}$, in which:

$z_i = F_s^{-1}(CSL_i)$, in which F_s^{-1} is the inverse standard normal of the cycle service level CSL_i.

$\sigma_i^D = \sqrt{\sum_j A_{ij} * \sigma_j^2}$, in which $\sigma_j = \sum_k B_{jk} * \sigma_k$

(All correlation coefficients ρ_{ij} are < 0.05, therefore we assume no correlation)

Demand during lead time $\mu_i * L_i$:

$$\mu_i = \sum_k \sum_j A_{ij} * B_{jk} * \mu_k$$

Assemble-to-Order model (half-fabric level) (Hereafter referred to as C&M ATO model)

Base stock level S_j:
$S_j = SS_j + \mu_j * L_j$, in which SS_j is the safety stock and $\mu_j * L_j$ is the demand during lead time. SS_j is calculated as follows:

Safety stock SS_j:
$SS_j = z_j * \sigma_j^D * \sqrt{L_j}$, in which:

$z_j = F_s^{-1}(CSL_j)$, in which F_s^{-1} is the inverse standard normal of the cycle service level CSL_j.

$$\sigma_j^D = \sigma_j = \sum_k B_{jk} * \sigma_k$$

Demand during lead time $\mu_j * L_j$:

$$\mu_j = \sum_k B_{jk} * \mu_k$$

$L_j = L_i + ExpectedProductionTime$, in which $ExpectedProductionTime$ is one week.

4.4.2. Guaranteed service model

Another model we use to determine the safety stock levels is the GSM by Graves & Willems (2003). This model incorporates multi-echelon safety stock placements, by investigating the service times at various stages. This is contrary to the model from Chopra & Meindl (2007), in which only a single level could be evaluated at the same time. However, the GSM does need additional information to calculate the safety stock levels. In addition to the demand mean and standard deviation, lead time, and service level, it also requires the promised outgoing service time from the end products to the customers. For Nedap, the outgoing service time of the end products would be five weeks, since that is currently their ideal lead time. Moreover, since we do not include the end products in the model, we set the outgoing service time of the half-fabrics to five weeks. One downside to this model is that the problem is formulated as a non-linear problem, because of the square root in the objective function of this minimization problem. However, this problem is solved by using the Gurobi solver in Python, which uses piece-wise linear approximation to solve non-linear problems. After the model from Graves & Willems (2003) is solved by the Gurobi solver, the service times for each product at each stage are provided. This service time is used in Eq. (4) to determine the safety stock levels.

Graves & Willems model (Hereafter referred to as GSM)

Base stock level S_z:

$S_z = SS_z + \mu_z * L_z$, in which SS_z is the safety stock and $\mu_z * L_z$ is the demand during lead time. We only calculate the base stock level if the GSM decides to put safety stocks on a certain place. Otherwise, we assume that no base stock levels are considered. SS_z is calculated as follows:

Safety stock SS_z:

$SS_z = z_z \sigma_z^D \sqrt{s_z^{in} + PL_z - s_z^{out}}$, in which z_z and σ_z^D are calculated in the same way as the models above. And where:

$PL_i = L_i$, in which the (production) lead time at component level, is the component lead time.

$PL_j = ExpectedProductionTime$, in which the (production) lead time at half-fabric level is only the $ExpectedProductionTime$, so without the component lead time.

$s_j^{out} = 5$, the (covered) half-fabrics should be available after five weeks (given by Nedap).

s_i^{out} and s_z^{in}, result from solving the GSM.

Demand during lead time $\mu_z * L_z$:

μ_z, is calculated in the same way as the models above.

4.4.3. Zero safety stock / current safety stocks

Besides the two models from literature, there are two other base stock levels which we also want to analyze. One option is putting all safety stock to zero, to analyze the impact of safety stock. In this case we only determine the base stock levels on component level, we do not incorporate half-fabric base stock levels. Another setting that is evaluated, is how Nedap's current approach performs. As stated in Section 2.4.2, the current approach is keeping roughly one month of demand as safety stock per component, and no base stock levels for the half-fabric products. Nedap currently does not work with a base stock policy, but since we know how much safety stock they keep, we can translate it to a base stock policy. Furthermore, since we assume that all the levels are dividable by the batch size, we round the values to the nearest number dividable by the batch size.

Zero safety stock (Hereafter referred to as 'zero' model)

Base stock level S_z:
$S_i = 0 + \mu_i * L_i$, in which $\mu_i * L_i$ is the demand during lead time and is calculated in the same manner as the C&M and GSM.

$S_j = 0$, no base stock levels are considered at half-fabric level.

Nedap's current stock levels (Hereafter referred to as 'Nedap' model)

Base stock level S_z:
$S_i = SS_i + \mu_i * L_i$, in which SS_i is Nedap's safety stock, and $\mu_i * L_i$ is the demand during lead time and is calculated in the same manner as the C&M and GSM.

$S_j = 0$, no base stock levels are considered at half-fabric level.

Nedap's safety stock SS_i:

$SS_i = 4 * \mu_i$, at component level Nedap currently maintains roughly four weeks of expected demand, in which:

$$\mu_i = \sum_k \sum_j A_{ij} * B_{jk} * \mu_k$$

4.4.4. Initial models for simulation

The models found in literature provide a possible solution to the current problem. However, the simulation determines which model provides the best result. In this case, a good result means low inventory levels while maintaining a certain service level. Furthermore, the models are not only evaluated on their performance as individual models. They are also assessed based on their performance as initial model for the simulation optimization explained in Section 4.5. In this section a simulation optimization is provided in the form of a metaheuristic. And a metaheuristic needs an initial solution to start the optimization process. By analyzing the different models as initial solutions, it can be determined which model converges to the best value in the shortest time.

4.5. Simulation optimization model

In order to obtain the best solution for the situation at Nedap, a model should optimize the base stock levels for this situation. In this section it is explained how simulation is used to optimize the base stock levels, using simulated annealing.

4.5.1. Optimization heuristic (Simulated annealing)

Optimizing base stock levels by means of simulation comes down to evaluating as many different options as possible with a certain accuracy for a given time span. In this case we have ten decision variables (the base stock levels at component and half-fabric level) which can all obtain a large number of values. Since it is impossible to evaluate each individual solution due to the number of feasible solutions and the long running time, we use a simulated annealing.

4.5.2. Initial solution and neighbor generation

In Section 3.4.2 we found guidelines for determining the relevant simulated annealing settings. However, the best initial solution and how neighbors are generated depend on the situation. As explained above in Section 4.4.4, the best initial solution follows from the four models that are evaluated. For the generation of neighbor solutions, we only change one base stock level at a time, which is a fixed step size (the batch size). This has two main reasons: first of all, by incorporating a step size we shrink our solution space. Secondly, we do not lose an important part of the solution space. This is because we assume that only full batches are produced, hence when we have less than a single batch size on stock, we are never going to use those items. The unused items only increase the average inventory level without impacting the service level.

4.5.3. Objective function

In order to determine which solution is better, an objective function is formulated. The objective is two sided, a certain service level should be obtained with as little inventory as possible. Currently, we want to achieve a service level of at least 95%. Nevertheless, in Chapter 5 we look into the impact of the service level on the objective function, which means that different service levels are evaluated. When the service level is set, the goal is to maintain the lowest inventory possible while maintaining a service level of at least 95%. This average inventory over all the half-fabrics and components is taken. Moreover, if one of the service levels from the covered tags drops below the 95% mark, a huge penalty is given (in the form of a large increase in the objective value) in order to ensure that the minimum service level is reached.

Objective function:

$$penalty = \begin{cases} 1, \text{if } ServiceLevel_k < 0.95 \ \forall k \\ 0, \text{else} \end{cases}$$

$$Objective\ value = BigM * penalty + \frac{\sum_z \sum_w Inventory_{z,w}}{(MaxNumberComp + MaxNumberHF) * RunLength},$$

with $BigM$ at least larger than the second part of the objective.

4.5.4. Final settings

In Table 9, a summary can be found on the settings that have been selected for the simulated annealing, based on Brusco (2014).

Table 9. Simulated annealing settings

Setting	***Value***
Start temperature	1,000
End temperature	100
Cooling scheme	Exponential, factor: 0.97
Markov chain length	50
Neighbor creation	Increase/decrease one base stock level with batch size
Initial solution	Best initial model
Objective function	Minimize average inventory while all service levels $\geq$ 95%

4.5.5. Simulation settings

The simulated annealing algorithm evaluates many solutions (number of iterations * Markov length). If all of these values would be calculated by the simulation model with the original settings, this would take an immense amount of time. The time to obtain one objective value is already roughly 150 seconds. Therefore, we need to make sure that we can run this algorithm in a reasonable time, so hours rather than years. In order to do this without harming the performance too much, we can change some of the simulation settings in order to decrease the running time of the simulation. The consequence is that the value obtained is less accurate. This could mean that certain neighbor solutions get accepted while they might be worse, or the other way around. Nevertheless, we want to maintain a certain accuracy, and accept that one can never be a hundred percent sure in simulation. We acquired these lower running times by setting the number of replications from 400 to only 3. The analysis and explanation of these values can be found in Appendix A, which resulted in the following values:

Warm up period	15 weeks
Run length (excl. warm up)	500 weeks (10 ‘work’ years)
Number of replications	3 replications
Expected run time	1.03 seconds

4.6. Conclusion method

In this chapter we discussed the method we use to determine the best inventory parameters. First, we decided to create a simulation to model the current system at Nedap. This tactical-level simulation simulates the production steps together with the inventory levels of the SmartTags over a given time period. Given a set of base stock levels and other parameters, the simulation provides multiple KPIs in order to evaluate the performance of the base stock levels decision. In addition, some important assumptions have been made in Section 4.2 to model the simulation as accurate and close to reality as possible. Next, we evaluated the models which we derived from the literature, by using it as input for our simulation model. Besides the models we derived from literature, we also created our own optimization model. We applied SA with our simulation model to find the optimal base stock levels. The simulation model together with the models derived from literature and our own simulation provided enough information to optimize the base stock levels.

5. Results

This chapter presents the results obtained from the simulation. First, we discuss the different models from literature and our optimization model. Second, we evaluate the base stock levels and performance of the models. Next to that, we gain insights into the impact of the different variables by performing a sensitivity analysis. Finally, at the end of this chapter, we examine the impact of some of the assumptions in Section 4.2. Note that the values are altered because of confidentiality.

5.1. Evaluation of the models

In this section we evaluate six different models. First, we evaluate the C&M MTO and C&M ATO model strategy. Next, we discuss the results from the GSM. After that, two different set of base stock levels are also taken into account. The first set is the 'zero' model, in which we determine base stock levels without safety stock, with the goal to gain insights into the impact of safety stock. The second set of base stock levels is the 'Nedap' model, which is an estimation of what the base stock levels would look like for the current situation at Nedap, in which on average roughly a month of demand is kept as safety stock for the components. Finally, we assess how well our own optimization model performs with respect to the others.

5.1.1. Overview base stock levels

In Chapter 4, we explained how we determine the base stock levels from the different models. In this section we provide an overview of all the base stock levels that are obtained from these models.

Chopra & Meindl

The input and output of the models from Chopra & Meindl (2007) can be found in Table 10 and Table 11. Table 10 presents the values for the MTO strategy, in which base stock levels are only kept at component level. Table 11 provides the values for the ATO strategy, in which base stock levels are exclusively kept at half-fabric levels.

Table 10. I/O C&M MTO model (components)

	Components			
	Comp 1	**Comp 2**	**Comp 3**	**Comp 4**
Input (weekly)				
Demand (mean)	36	3	1	27
Demand (stdDev)	20	4	2	20
Lead time	6	6	6	6
Service level	85%	85%	85%	85%
Output				
Safety stock	55	8	4	50
Base stock	270	27	8	210

Table 11. I/O C&M ATO model (half-fabrics)

	Half-fabrics					
	HF 1	**HF 2**	**HF 3**	**HF 4**	**HF 5**	**HF 6**
Input (weekly)						
Demand (mean)	23	13	3	0	1	0
Demand (stdDev)	20	7	4	1	2	1
Cum. Lead time	7	7	7	7	7	7
Service level	75%	75%	75%	75%	75%	75%
Output						
Safety stock	86	30	14	0	6	0
Base stock	248	120	35	0	12	0

We initially aim to deliver 95% of the products on time, however, Chopra & Meindl (2007) use the cycle service level in their model. When we compute the service level with our simulation by evaluating the number of products that are delivered on time, we obtain values close to 100%. Therefore, we aim for a cycle service level which corresponds to our desired service level of 95%. This is done by running the simulation for different cycle service levels. It is shown in Table 12 which cycle service level came closest to our desired service level.

Table 12. Overview service levels Chopra & Meindl model

Cycle service level	Simulated service level (MTO)	Simulated service level (ATO)
75%	91.7%	99.2%
80%	94.6%	99.3%
85%	97.6%	99.5%
90%	99.7%	99.8%
95%	99.9%	99.9%

Although a cycle service level of 80% came closest to the 94.6% for the MTO model, an 85% cycle service level obtained a desired service level above the 95%. Therefore, we opt for a cycle service level of 85% for the MTO model. For the ATO model, all service levels are above the 95%, we select the cycle service level of 75%, since it is the closest to the desired number. Moreover, we decided not to go lower than a cycle service level of 75%, since the model gets close to becoming a MTO model when evaluated by the simulation. When the base stock levels for the half-fabrics are relatively low, the simulation model orders components so frequently (and in high numbers) that the inventory starts to pile up at component level, which resembles a MTO model.

Graves & Willems

For the GSM the values are put into one table, rather than the two separate tables for the MTO and ATO C&M models. The reason for this, is that the GSM evaluates the entire supply chain at once, which means that in addition to determining the safety stock level it also determines the placement. Furthermore, the original GSM would also incorporate the end products, but because the last production part has a lead time of zero weeks, there is no reason to stock items at that echelon level according to the GSM. In Table 13, the input and output of the GSM can be found. The input of the demand is the same for the GSM as for the C&M models. But instead of providing the lead times, the service times are shown, which are obtained from solving the non-linear model.

Table 13. I/O of the GSM

Input	Components				Half-fabrics					
	Comp 1	**Comp 2**	**Comp 3**	**Comp 4**	**HF 1**	**HF 2**	**HF 3**	**HF 4**	**HF 5**	**HF 6**
Demand (mean)	36	3	1	27	23	13	3	0	1	0
Demand (stdDev)	20	4	2	20	20	7	4	1	2	1
Service time	2	2	2	2	0	0	0	0	0	0
Service level	95%	95%	95%	95%	95%	95%	95%	95%	95%	95%
Output										
Safety stock	48	7	3	47	0	0	0	0	0	0
Base stock	264	26	8	208	0	0	0	0	0	0

Note: Service time $= s_z^{in} + PL_z - s_z^{out}$

We see from Table 13 that the GSM only places safety stock on component level. The service times for the components are two weeks, which means that safety stock should be implemented to cover for these two weeks of demand. When we compare these values to the C&M MTO model, we see that the base stock levels are slightly lower. That GSM provides lower values is not surprising considering the model's nature. Since GSM also takes the promised lead time to the customer into account, it suggests a shorter duration for which safety stocks should be kept. However, there are two reasons why this difference is not as big as one might expect: first, the service level of the C&M MTO model is 85% compared to the 95% of the GSM. Second, the lead/service time are in the square root of the equation of the safety stock, which means that the difference is not linear.

'zero' model & 'Nedap' model

In Table 14 and Table 15 an overview of the safety and base stock levels from the two additional models is provided. As can be seen from the tables, both models only incorporate base stock levels at component level.

Table 14. Output zero safety stock

Zero safety stock	Components				Half-fabrics					
	Comp 1	**Comp 2**	**Comp 3**	**Comp 4**	**HF 1**	**HF 2**	**HF 3**	**HF 4**	**HF 5**	**HF 6**
Safety stock	0	0	0	0	0	0	0	0	0	0
Base stock	217	18	4	161	0	0	0	0	0	0

The base stock levels in Table 14 are only the demand during lead time, since there is no safety stock. This model shows how the base stock levels look when no safety stocks are used.

Table 15. Output Nedap's current stock levels

Nedap's current stock levels	Components				Half-fabrics					
	Comp 1	**Comp 2**	**Comp 3**	**Comp 4**	**HF 1**	**HF 2**	**HF 3**	**HF 4**	**HF 5**	**HF 6**
Safety stock	145	5	2	27	0	0	0	0	0	0
Base stock	361	24	7	188	0	0	0	0	0	0

The safety and base stock levels in Table 15 are an estimation of how Nedap's base stock levels would look like if they would have used a base stock policy. As explained in Section 2.4.2 they currently keep a month of expected demand as safety stock, this corresponds to the values in Table 15. The base stock levels are significantly higher than both C&M models, the GSM, and the 'zero' model.

Simulated annealing

For the simulated annealing algorithm, we already determined most parameters. These are determined based on Brusco (2014), which provided guidelines for optimal simulated annealing settings. In this section we evaluate these parameters and see how they perform. First, we evaluate which models perform best as initial model. After that we extend the running time of the algorithm by changing the parameters and seeing if the difference in performance makes up for the increase in running time.

Initial model

As explained in Section 3.4.2, simulated annealing needs an initial solution in order to provide solutions. The choice of an initial solution can have a significant impact on the performance and running time of the algorithm. Initial solutions that are 'far away' in the solution space from the global optimum require more time to converge to good solutions. Therefore, they are less efficient than initial solutions which are already closer to good solutions. We already have different ways to generate initial solutions, namely the five different models to determine the base stock levels. In this section we evaluate these models, by performing the simulated annealing algorithm with the same settings obtained from Brusco (2014). In Table 16 the performance of these initial models are provided. In Table 17, we can see the base stock levels after running the algorithm.

Table 16. Performance models after SA (with the settings from literature)

Initial Model	Average inventory	(rank)	Service level	(rank)
SA (C&M (MTO))	95	(5)	98.6%	(1)
SA (C&M (ATO))	65	(2)	92.6%	(2)
SA (GSM)	76	(3)	92.4%	(3)
SA (Zero)	39	(1)	88.3%	(5)
SA (Nedap)	91	(4)	90.5%	(4)

Table 17. Base stock levels after SA (with the settings from literature)

Base stock levels (after SA)	Components				Half-fabrics						Total base stock
	Comp 1	Comp 2	Comp 3	Comp 4	HF 1	HF 2	HF 3	HF 4	HF 5	HF 6	
SA (C&M (MTO))	261	24	12	199	22	12	2	7	0	2	541
SA (C&M (ATO))	53	19	0	10	252	74	39	40	33	17	537
SA (GSM)	257	24	7	182	2	4	2	7	5	5	495
SA (Zero)	201	21	2	151	21	7	2	5	12	2	424
SA (Nedap)	345	21	45	132	12	10	41	2	31	0	639

The overview of the performances provided by Table 16 shows that choosing the C&M MTO model as initial model is the only option that is able to achieve the minimum service level of 95%. However, this is achieved with the highest average inventory of all the models. It still is the best initial model, because the neighbor solutions are likely to also be solutions with a service level above 95%. The neighbor solutions are important, since they determine the way to the best solution. Nevertheless, it becomes clear that all models differ quite a lot from each other in terms of performance (Table 16) and base stock level (Table 17). These differences mean that the solutions did not converge to a(n) (local) optimum, which tells us that the simulated annealing settings are not optimal for this case.

Parameter fine tuning

In order to achieve a(n) (local) optimum, we extend the current simulated annealing parameters. Next to the general simulated annealing parameters, we also include the number of replications the simulation performs in order to determine the objective value of the neighbor solution. This parameter is not a simulated annealing parameter, but it does determine the accuracy of the objective function from the neighbor solution. Increasing the accuracy of the neighbor's objective increases the chance of achieving a solution with a minimum service level of 95%. In Table 18 three different settings are evaluated in order to determine the best simulated annealing parameters. These runs use the C&M MTO model as an initial solution, since it was the only model that was able to obtain a service level above the 95%.

Table 18. Tuning the simulated annealing parameters (initial model: C&M MTO)

	Start temp	End temp	Decrease factor	Markov length	Replications	Average inventory	Average service level	Run time (sec)
Setting 1	1,000	100	0.97	50	3	95	98.6%	1,816
Setting 2	1,000	10	0.98	75	5	71	95.9%	12,683
Setting 3	1,000	1	0.98	100	5	71	95.9%	31,968

From Table 18 we observe that running the simulated annealing algorithm more intensely, and thus longer, results in better objective values. However, the second change in parameters (from Setting 2 to Setting 3) did not result in an improvement. This could be because the model got stuck in a local or global optimum. We have to take into account that this problem is a stochastic problem, and the algorithm is based on random values. Therefore, it is difficult to talk about optimal values. Nevertheless, it seems that to a certain extent, the algorithm is able to find improvements. So, in conclusion, running the algorithm for more iterations results in better values. However, the algorithm has its limits and for this case it seems that the limit is found at running the algorithm for three and a half hours. Next to that, we observed that the hyperparameters provided by Brusco (2014) did not provide the optimal values. The explanation for this lies in the different problem sizes, larger problems have more different values to evaluate, which requires more iterations to evaluate different solutions.

Final base stock levels

After performing the simulated annealing algorithm with the best parameters, we obtained the base stock levels presented in Table 19.

Table 19. Base stock levels simulated annealing

Simulated annealing	Components				Half-fabrics					
	Comp 1	Comp 2	Comp 3	Comp 4	HF 1	HF 2	HF 3	HF 4	HF 5	HF 6
Base stock	84	21	9	182	194	93	9	43	29	9

5.1.2. *Evaluation base stock levels*

In Table 20, a summary of all the base stock levels is provided. All the numbers are rounded to the nearest number that is dividable by the batch size, see Section 4.2.6 for the assumption that explains this action. From this we can see that the simulated annealing is the only model that includes base stock levels on component and half-fabric level. Moreover, all other models only include base stock levels on component level, except the C&M ATO model, which only includes base stock levels on half-fabric level.

If we take a look at the total base stock level, we can see that the zero-safety stock model has the lowest total base stock. And the simulated annealing has by far the highest total base stock. Furthermore, it is interesting to see that the C&M ATO model incorporates 160 less total base stock than the C&M MTO model. Due to the risk pooling effect, one expects the total base stock of the C&M MTO model to be lower. However, the reason that the total base stock of the C&M MTO model is not lower, is that all the half-fabrics that start with (HF 1, HF 3, HF 5) need two components, the 'ID' and one other components. Also, the difference in cycle service level (75% for ATO vs 85% for MTO) plays a role in this difference. However, if we take a cycle service level of 95% for both, the difference is still 100, these base stock levels can be found in Appendix C: base stock levels C&M models. Although the base stock levels for the MTO model are higher compared to the ATO model, the safety stock for the MTO are lower compared to the ATO model.

Table 20. Overview base stock levels: comparison of the output from the different models

	Components				Half-fabrics						Total
Base stock levels	**Comp 1**	**Comp 2**	**Comp 3**	**Comp 4**	**HF 1**	**HF 2**	**HF 3**	**HF 4**	**HF 5**	**HF 6**	**base stock**
C&M (MTO)	270	27	8	210	0	0	0	0	0	0	515
C&M (ATO)	0	0	0	0	248	120	35	0	12	0	415
GSM	264	26	8	208	0	0	0	0	0	0	506
Zero	217	18	4	161	0	0	0	0	0	0	400
Nedap	361	24	7	188	0	0	0	0	0	0	580
SA	84	21	9	182	194	93	9	43	29	9	673

5.1.3. *Evaluation of the performances*

Based on the base stock levels from Table 20, we calculated the average inventory and service level with our simulation. The results of these simulations are shown in Table 21.

Table 21. Performance different models

Model	**Average inventory**	**(rank)**	**Service level**	**(rank)**
C&M (MTO)	97	(4)	97.6%	(2)
C&M (ATO)	295	(6)	99.2%	(1)
GSM	94	(3)	96.9%	(3)
Zero	41	(1)	79.8%	(6)
Nedap	103	(5)	88.4%	(5)
SA	71	(2)	96.3%	(4)

First of all, when we compare Nedap's current solution to the best solution (simulated annealing), we see that the average inventory decreases with 30.7%. In addition, the service level is 96.3% rather than the 88.4% Nedap would have obtained. Thus, the simulated annealing solution provides better average inventory values and service level, making it the superior solution on both fronts. In addition, both the C&M MTO model and GSM outperformed Nedap's current model, with 6.0% and 9.0% lower average inventory respectively. Not only did both models perform better in terms of average inventory, both models also obtained a service level above 95%.

Next to that, the lowest average inventory in Table 21 is the zero-safety stock configuration. However, the service level is below 80%, which is far off the 95% target service level. The second lowest average inventory is the simulated annealing, which is surprising considering its total base stock level is by far the largest. The reason why high total base stock levels do not always result in high average inventory levels, can be explained by the deviation of the base stock levels over the different components and half-fabrics. For instance, if we take Nedap's current model, we notice a high base stock level for Comp 1 (361). When we compare this number to the base stock level for the same component provided by simulated annealing (SA model), we perceive a much lower value (84). However, the part of inventory that is nonmoving during each cycle (on average), is noticeably different for the two models. The part of inventory that is not moving in case of the Nedap model is extremely high. For the SA model this is different, because the inventory in the SA model is used more effectively. There is little nonmoving inventory, resulting in low average inventory levels. Furthermore, Nedap's model does not include base stock levels on half-fabric level, but inventory arises as work in progress inventory. For this reason, a higher average inventory is obtained with respect to the base stock level. In Figure 11 an illustrative example of these effects is provided. From this figure it becomes clear that Nedap's model has significantly more non-moving inventory, which results in higher average inventory, despite the lower total base stock level.

Note: ID = Identification (component)

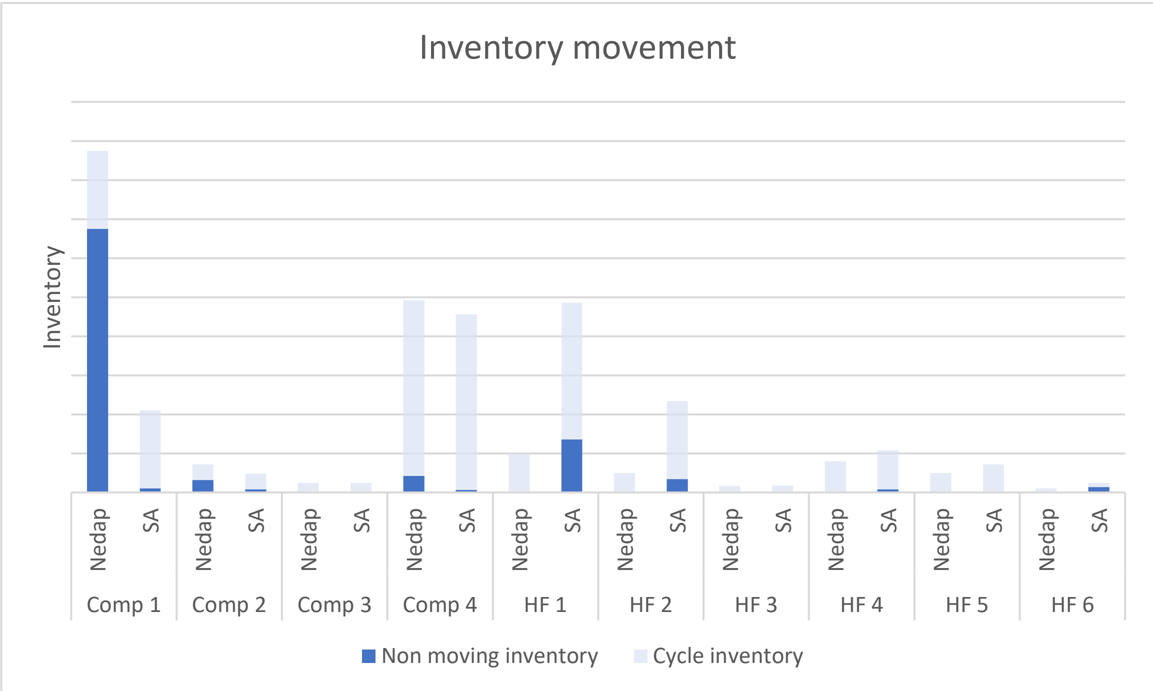

Figure 11. Illustrative example inventory movement

This decrease in average base stock level is the result from understanding exactly where and how many items are needed in the supply chain, this effect becomes especially clear when we look at the first component (Comp 1). Keeping items upstream the supply chain is beneficial due to the risk pooling effect. On the contrary, stocking items downstream results in reducing the risk of waiting for other orders to be produced first. The SA model is able to find the balance between the two and thus provides lower average inventory values, despite the high(er) total base stock level.

5.2. Sensitivity analysis

In the previous sections we found the best base stock levels simulated annealing was able to provide. However, this was solely based on the average inventory given a service level (of 95%). In order to understand the impact of the solution, we evaluate the effect of the input on the output by means of a sensitivity analysis. We use the best base stock levels found by the simulated annealing algorithm for the sensitivity analysis, which tells us how the solution performs under different circumstances.

Moreover, the following input parameters are considered in the analysis: lead time, demand, production capacity, batch size, and minimum order quantity. For the output we are not only interested the average inventory. In order to gain more insights into the performance of the best solution, we also discuss other KPIs. First, we evaluate the service level per end product, which gives us information about how the solution performs per end product. Next, the average order size and frequency per component are discussed, with the goal to examine the impact of base stock levels on the order patterns. Furthermore, we keep track of the utilization of the production (both parts), this tells us how 'full' the system is. Finally, we consider the percentage of time we do not have any components on stock, per component.

Lead time

For the sensitivity analysis of the lead time, we evaluate lead times between four and eight weeks, with steps of one week. Below in Figure 12, the impact of these component lead times on the average inventory and service level is shown.

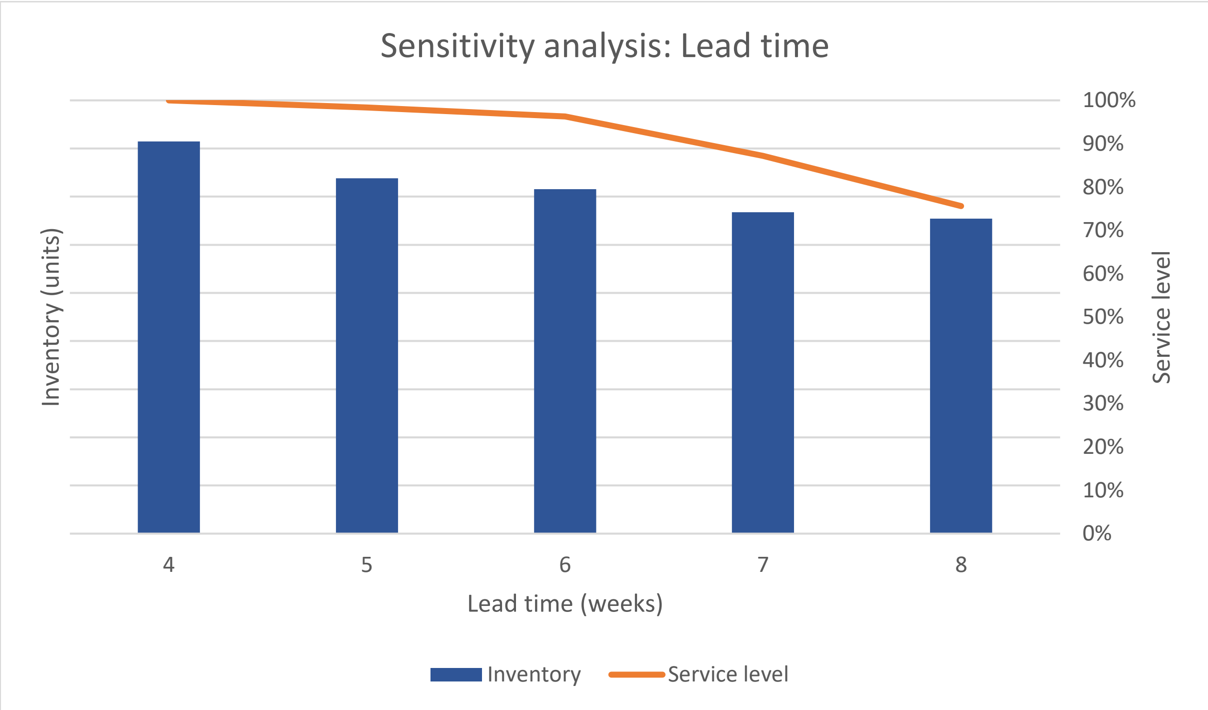

Figure 12. Sensitivity analysis: Lead time

From Figure 12 we conclude that both the average inventory and the service level decrease when the component lead time increases. When components arrive later, we have a higher chance of stockouts due to the increased waiting time of our components. These stockouts result in lower inventory levels and lower service levels. Currently, we assume the lead time to be six weeks. However, from the figure it becomes clear that if the lead time increases by a week, the service level drops by a significant amount. This emphasizes the impact of the accuracy of the lead time. Furthermore, it is in practice unlikely that suppliers deliver sooner than promised, which would decrease the lead time. However, if the lead time does decrease, we end up with more inventory. We notice an increase in average inventory of 3% and 14% for the lead times of five and four weeks respectively. This increase in inventory means more cost and a higher service level. However, the service level was already above the target level of 95%, therefore this increase in service level might not be worth the extra cost. In addition, the production utilization of the first production part and the percentage of no component on stock explain the cause of the decrease in service level. We see that for the original scenario, the percentage of no components on stock is 18% on average, with a production utilization (first part) of 64.2%. These numbers already show that the service level is limited by the availability of components rather than the production capacity. If we increase the component lead time to 8 weeks, we notice an even higher zero component percentages of 25% on average. The production utilization however, does not change at all.

We perceived that if we use the current base stock levels and the lead time is 7 weeks (for the components) it results in a service level under the target of 95%. Therefore, we reoptimize the base stock levels for a lead time of 7 weeks, to see what the difference in base stock levels and performance is. In Table 22 an overview of the base stock levels for a lead time of 7 weeks is provided.

Table 22. Optimized base stock levels (lead time = 7 weeks)

Base stock levels	Components				Half-fabrics						Total base stock
	Comp 1	Comp 2	Comp 3	Comp 4	HF 1	HF 2	HF 3	HF 4	HF 5	HF 6	
SA (lead time 7)	120	31	36	163	216	91	14	14	48	5	738

When we compare the base stock levels for 6 and 7 weeks of lead time, the total base stock increases by approximately 8%. If we then evaluate the performance with our simulation, we obtain an average inventory of 71, with a service level of 98.1%. Not only is the service level above the target of 95%, the average inventory is similar to the average inventory of the original scenario. Therefore, we obtain similar performance if we reoptimize our base stock levels for the alternative parameters.

Demand

In order to determine the effect of the demand, we analyze the interval between 20% less demand and 20% more demand (than currently assumed), with an interval of 10%. In addition, we also include a scenario of 50% more demand. In Figure 13, the impact of the demand on the average inventory and service level is shown.

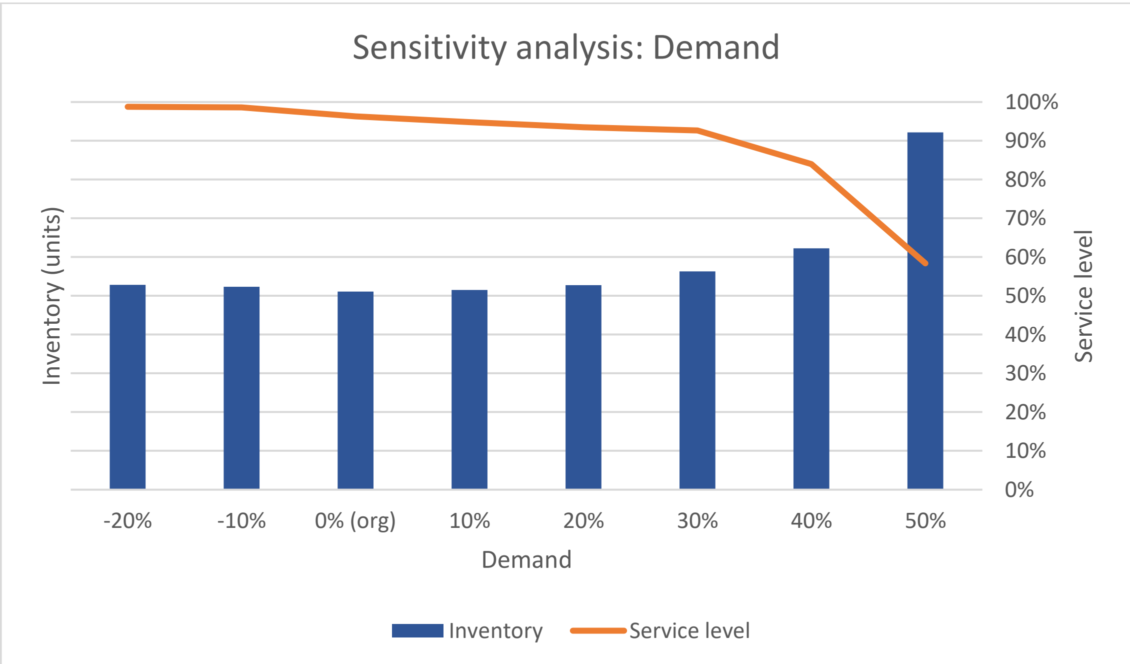

Figure 13. Sensitivity analysis: Demand

From Figure 13, we are not able to see a large difference in both the service level and inventory for the first five scenarios. The first four scenarios all obtain a service level of at least this 95% mark. Moreover, the service level reaches 93% with 20% more demand, which is just below the target service level of 95%. The average inventory levels for the first five scenarios are also not quite different. However, when we look at the last three scenarios, we see that the average inventory increases heavily and the service level drops. These values can be explained by the production capacity. In the first five scenarios the production capacity was enough to handle the demand. Hence, no large difference in service level and average inventory were noticeable. An increase in demand of 30% results already in a small change in both the average inventory and service level. However, in case of an increase of 40% demand, we notice that production capacity becomes the bottleneck, resulting in significantly higher inventories and lower service levels. The average inventory increases because components remain longer on stock, since they cannot be used when the production capacity is reached. If demand rises, we order additional components to fulfill this increased demand. Nevertheless, if the production capacity cannot handle this demand the inventory piles up. For the same reason the service level decreases, because end products cannot be produced due to production capacity limitations. This is confirmed by the production utilization of the first production part, which is 64.2% for the original demand scenario and 96.2% for the 50% more demand scenario. From this we conclude that, at this level of demand, the production capacity becomes the bottleneck rather than the component availability.

Production capacity

For the sensitivity analysis of the production capacity we only evaluate the production capacity of the first production part. We change the production capacity of the second part with the same margin as the first production part. However, the second production part has a lower utilization than the first part, which means that the first production part is more likely to become the bottleneck. Therefore, we evaluate the following five production capacities (next to the original capacity) for our analysis of the first production part: -30%, -20%, -10%, +10%, and +20%. The results can be found in Figure 14.

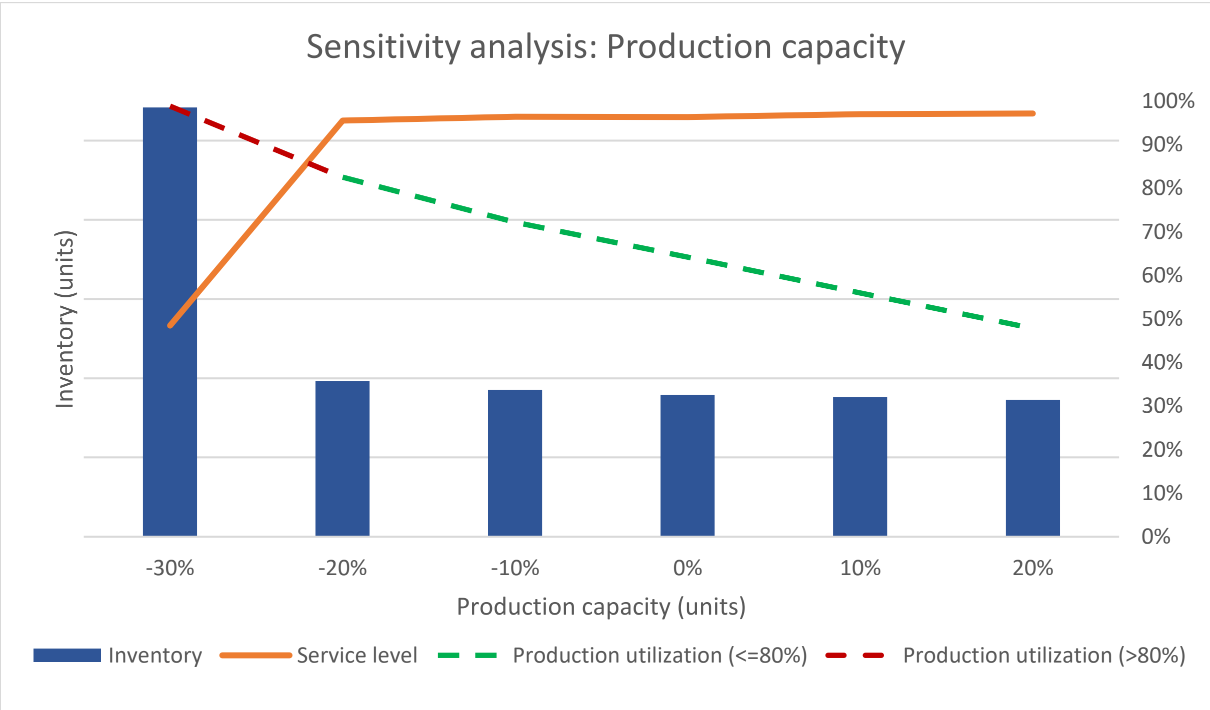

Figure 14. Sensitivity analysis: Production capacity

Figure 14 shows us that the service level barely changes and the average inventory decrease if the capacity does not drop by 20%. Decreasing the production capacity means that less components can be used per week, which means that components remain on average longer on stock if the capacity is reached. Because the components remain longer on stock, the average inventory increases. However, the service level is not affected by it, because the production capacity is not the bottleneck for the service level. Some end products might take longer to be delivered, but if it is within the five weeks promised lead time, the service level does not get effected. However, we notice that production capacity does become the limiting factor when the production capacity reduces by 30%. The average inventory becomes three times higher compared to the other experiments and the service level drops below 50%. Moreover, by looking at the production utilization we see that the -20% production capacity is already close to becoming the bottleneck, since the first production part obtains a utilization of 82.5%. This utilization indicates that production capacity of -20% of the total is just enough to keep the system from overloading.

Batch size

The batch size is the next input parameter which is included in the sensitivity analysis. For this we take 5 different batch sizes, these are scenarios 1 to 5. Currently, scenario 5 is the largest batch size Nedap is able to work with, due to production limitations. In Figure 15, the performance of the different batch sizes is presented.

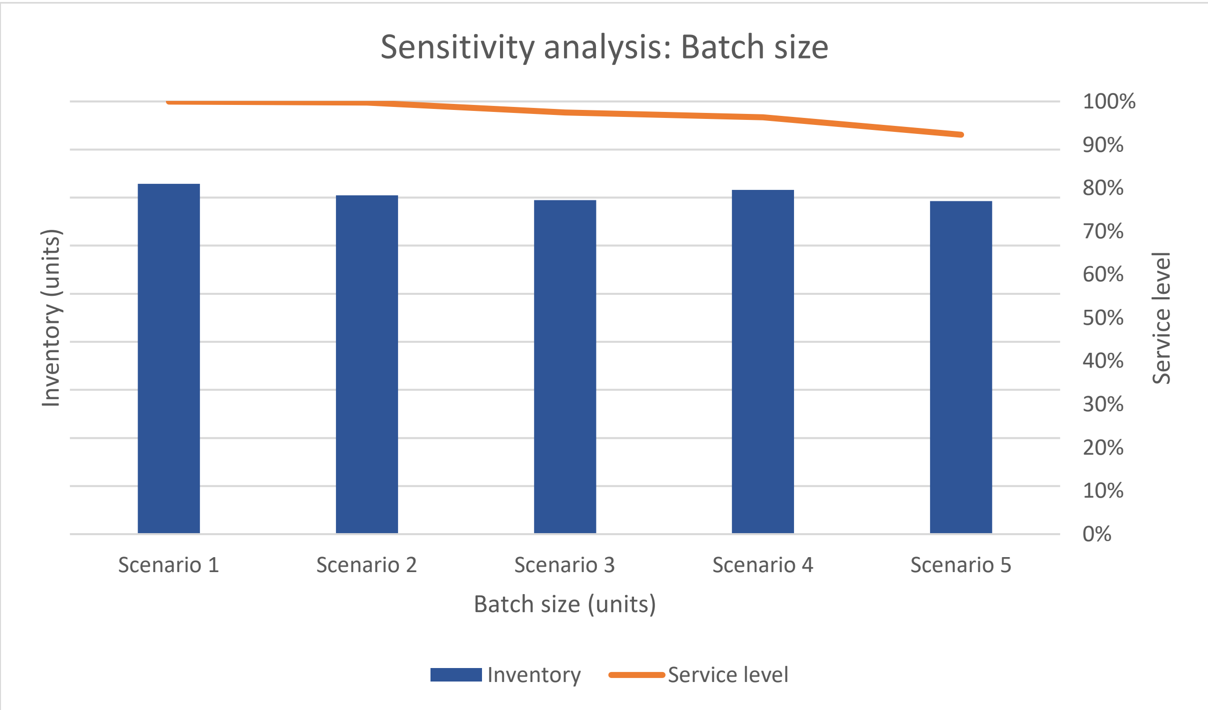

Figure 15. Sensitivity analysis: Batch size

Figure 15 shows that the batch size effects the service level more than the average inventory. The main advantage of decreasing the batch size is the increase in flexibility. A lower batch size means that we produce less products that we do not immediately need. This increases flexibility, resulting in a higher service level. Furthermore, there is no clear relation between the average inventory and the batch size. The base stock levels and MOQ remain the same, leading us to order the same quantity and maintain the same level. Although, we might place an order sooner because we produced more due to a higher batch size, the order size remains unchanged. Hence, the batch size does not have a clear impact on the average inventory. KPIs which do show different outcomes for different batch sizes are the production utilization of the first production part and the percentage of zero components on stock. The production utilization and percentage zero components are 62.6% and 10.6% respectively for a first (most flexible) scenario. These values are better compared to the 66.5% and 19.8% for the larger batch size of scenario 5. This difference is solely obtained by being able to produce the exact amount we need.

Minimum Order Quantity (MOQ)
Finally, we analyze the impact of the MOQ. The MOQ is determined by Nedap's suppliers, which reduces the order flexibility. In order to gain insights on how much impact the MOQ constraint has, we evaluate the solution without the MOQ constraint.

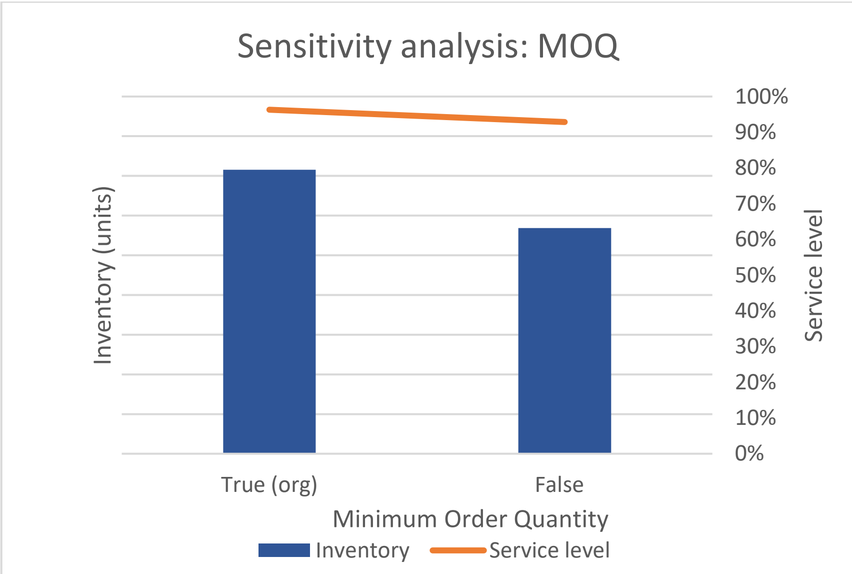

Figure 16. Sensitivity analysis: MOQ

Figure 16 shows that the MOQ has significant impact on the performance of the solution. First of all, the average inventory decreases by 20.6%, while the service level decreases from 96% to 93%. The explanation for this effect is similar to the effect we saw with the batch sizes. By removing the MOQ constraint we are more flexible, which means that we only order what we need. This results in a large decrease in average inventory units (20.6%), but also in less service level. One advantage of ordering more than initially required, is that sometimes we require more than expected. Because of this effect we obtain a slightly higher service level with the MOQ constraint. This effect is confirmed by the ordering patters. We see that we order three times more often without the MOQ constraint. In addition, we see that our average order is also significantly lower with 18 units instead of 57 (almost a third of the total).

When we do not incorporate the MOQ, the service level falls slightly below the target of 95%, since the base stock levels are determined based on including an MOQ. Optimizing the base stock levels without an MOQ results in the values provided in Table 23.

Table 23. Optimized base stock levels (no MOQ)

Base stock levels	Components				Half-fabrics						Total base stock
	Comp 1	Comp 2	Comp 3	Comp 4	HF 1	HF 2	HF 3	HF 4	HF 5	HF 6	
SA (no MOQ)	84	24	9	182	194	93	19	43	29	9	686

The new base stock levels are unchanged apart from comp 2 and HF 3. When we take a closer look at the KPIs from the original scenario, we notice a service level of approximately 63% for the HF 3. Hence, it makes sense that this product with its related components are the two items that have an increased base stock level. These extra 13 base stock units (3 of Comp 2, and 10 of HF 3), result in an average inventory of 59 units and a service level 96.1%. This means that the target service level of 95% is reached and we are still able to maintain a 17.5% lower average inventory than the original scenario with an MOQ.

5.3. Assumption validation

In this section we evaluate the impact of two assumptions we made for our simulation. The first assumption is about not including safety stock / base stock levels at end product level. Therefore, we adapt our simulation to include the option to keep inventory on end product level. The other assumption is about the order in which batches are planned. Currently, we decide which batch is produced based on a fixed order, sorted from most to least expected demand. We evaluate to what extent this decision changes the objective.

End products on stock

In order to include the option to include inventory on end products level, the simulation is slightly altered. In addition to evaluating if it is possible to create additional batches to obtain half-fabric inventory (inventory orders), we now also evaluate if it is possible to cover half-fabrics already to create inventory for the end products. We previously made an assumption that we do not make a distinction between the different covers, which makes it difficult to decide for which end products we maintain inventory. Therefore, we only place inventory for the ‘runners’ for which we know there is plenty demand. These ‘runners’ include the covered HF 1 and HF 2 tag. For these two items we are keeping end product inventory with a maximum of half a week of demand. We only keep a maximum of half a week of demand, since this limits the risk of keeping stock on that level but shows the possible gain of keeping stock.

After performing the simulation with the addition of keeping inventory at end product level, we saw a small decrease in both average inventory and service level. The average inventory we obtained was 71 with a service level of 96.0%. The average inventory does not include the inventory of the end products. However, since a small part of the inventory is now shifted to the end product inventory the average inventory is slightly lower. Furthermore, the main benefit when placing inventory on end product level is that we do not have to cover (second production part) the tags. Nevertheless, the production capacity for covering a tag is never the limiting factor in the current setting. Therefore, including inventory on end product level did not make a significantly large difference.

Production planning order

We also analyzed the impact of altering the order in which products are planned. Currently the products with the highest expected demand are planned first. For this analysis we plan the products with the least expected demand first. By changing the order in which orders are planned, we evaluate the impact the planning order has on the performance of the solution.

The new planning order resulted in an average inventory of 72 and a service level 95.5%, which is a slightly worse solution than the original planning order. The reason why this solution is slightly worse, is because the products with a higher expected demand are less likely to be planned quickly, since the other products are prioritized. Because these products need to wait longer to be produced, the relative components spend more time on the shelf, which results in a higher average inventory. Moreover, when these products are planned too late, this results in a lower service level as well. The additional delay of the 'high demand' products as a result of the worse planning order, impacts the objective slightly. However, this impact could be much worse if the production would form the bottleneck in this setting. When the production is the bottleneck, the decisions regarding the planning order carry greater consequences.

5.4. Conclusion results

First, we obtained the base stock levels from the different models. After fine tuning the service level for the model provided by Chopra & Meindl (2007) and the simulated annealing parameters, the simulated annealing model was able to obtain the best performance. With an average inventory of 71 and a service level of 96.3%, it achieved the lowest average inventory while maintaining a service level above the 95% mark. In addition, our 'initial model' experiment showed that the C&M MTO model was able to obtain the best value (only service level above 95%) after running the simulated annealing model for a fixed time. Next to that, we noticed that the total base stock level of the best performing model (simulated annealing) was significantly higher than the other models. While the total base stock level was higher, it was still able to outperform the other models in terms of average inventory. We learned that the placement of the base stock levels plays an enormous role in the performance of the model. Besides the performances we acquired, we also gained more insights in the effect of the input and output parameters with a sensitivity analysis. From this sensitivity analysis we obtained many insights in the following input parameters: lead time, demand, production capacity, batch size, and MOQ. Three important conclusions can be drawn from this sensitivity analysis. First, the lead time has a large impact on the performance, regarding both the inventory and the service time. Second, in the current setting we are not bounded by production capacity, but rather by the availability of components. Third, the MOQ constraint has a significant impact on the average inventory. In addition to conducting these analyses, we also reoptimized the base stock levels for two different scenarios. The first scenario involved a lead time of 7 weeks, while the second scenario excluded the MOQs. These re-optimizations showed where and how much to change the base stock levels in order to obtain an appropriate service level and average inventory. Finally, we evaluated the assumptions we made about the inventory placement on end product level and the order in which products are planned. We showed that both assumptions have minor impact on the model.

6. Conclusion and discussion

This chapter concludes this thesis. In this Section 6.1 we conclude the thesis by answering the research question: *"How can the inventory be optimally managed for the production process of the SmartTags?"*. In addition, we provide other findings from the thesis. Next to that, in Section 6.2 we discuss the limitations of the research. Furthermore, in Section 6.3 we provide recommendations to Nedap. Finally, in Section 6.4 we discuss the opportunities for future research.

6.1. Conclusion

We developed a simulation model to evaluate the impact of the base stock levels on the average inventory and service level in a multi-echelon multi-item supply chain. By implementing SA to optimize the base stock levels, we found a set of base stock levels which provided great results for Nedap. If we compare these results to Nedap's current inventory policy, we obtain a decrease of approximately 30.7% average inventory while increasing the service level from 88.4% to 96.3%. This improvement is realized by understanding the requirements of each item on every echelon level. By utilizing the entire supply chain, rather than only the component level, much performance has been gained.

From the literature we found that both Chopra & Meindl (2007) and Graves & Willems (2003) provided a model to calculate safety stocks and thus base stock level. The C&M MTO model was able to find a 6.0% reduction in average inventory compared to Nedap's current policy, after fine tuning the cycle service level. In addition, the GSM reduced the average inventory by 9.0% compared to Nedap's policy. Both models are superior to Nedap's in terms of average inventory, while obtaining a service level above the target of 95%. However, the reason why these models did not perform as great as our SA algorithm, is the placement of the safety stocks. No other model, except the SA, suggested placing stocks on more than one echelon level, which caused higher average inventory values. Moreover, the research on the impact of production capacity on inventory optimization confirms the differences we see between the SA model and the others. From evaluating the models that suggest only placing inventory at component level, such as C&M MTO and GSM, it becomes clear that the production constraints are not taken into consideration. Our optimization model does include the capacity constraints, which is why our model suggests placing inventory on both echelon levels, therefore benefitting from both risk pooling effect (component level) and not waiting for production (half-fabric level). Therefore, it is clear that our optimization model effectively takes these production capacity and planning into consideration. This shows the advantages our model has over the models discussed in Chopra & Meindl (2007) and Graves & Willems (2003) that do not incorporate these production constraints.

Furthermore, the sensitivity analysis provided interesting insights about the different input parameters and behavior of Nedap's current system. First of all, we observed a large impact of the lead time on both the service level and the average inventory. Increasing the lead time by one week resulted in a service level of 87%, increasing it by one more week resulted in a service level of 76%. The average inventory did drop as well, however, this is solely because the increase in stockout due to the increased lead time. Therefore, it can be concluded that the availability of the recourses has a large impact on the objective value.

Next to that, the sensitivity analysis on the production capacity and demand showed that the impact of the capacity is not the same as component availability. When we decrease the production capacity by 20% or increase the demand by the same amount, the objective value does not change that much. When this percentage gets closer to 50%, we see that the system gets overloaded and then the production capacity does become the limiting factor. However, in the current system we see that the availability of resources is limiting the performance more than the production capacity. Finally, we observed that excluding the MOQ of components leads to improved performance and reduced base stock levels. This outcome is expected, since MOQs often result in ordering quantities that exceed actual requirements.

6.2. Limitations

In this section we discuss the two main limitations we encountered during this research.

Stochastic lead time

The first limitation is the stochasticity of the lead time. We know that in practice the lead time is not deterministic, however since we do not have accurate information about the lead time and the simulation would include even more stochasticity, we decided to apply deterministic lead times. One more reason for using deterministic lead times was to be able to compare our model to the models derived from literature, which also use deterministic lead times. However, we did incorporate the lead time in the sensitivity analysis and evaluated the effect on the solution for different lead times. This showed that altering the lead time has a significant impact on the solution. Therefore, not including stochastic lead time is a limitation of this research.

Order generation

The second limitation is regarding the way orders are treated in the simulation. In practice, orders can arrive every day of the week and consist of multiple product types per order. We assumed that orders consist of one type of product and always arrive at the start of the week. The simulation evaluates how many types of a certain tag needs to be produced that week, but this can change in practice during the week depending on the incoming orders from the customer. Although this current way of simulating demand is not how it is currently going, it still provides enough insights to evaluate the demand on a tactical level. Nevertheless, it is a limitation since the simulated demand is not presented in the same way as the actual demand.

6.3. Recommendations

In this section we provided recommendations to Nedap.

Incorporating base stock levels

In order for the simulation to have any effect on the future performance, the values need to be implemented. Therefore, these base stock levels should be put into the ERP system. Moreover, because Nedap has access to the simulation model, they can use it to recalculate the base stock levels when parameters have changed. Therefore, we encourage Nedap to use the simulation model and update the values after an event of change. The importance of updating the input parameter is emphasized by the results we obtained from the sensitivity analysis.

Nedap needs to change their inventory settings in their ERP system and reconsider certain operation decisions to make sure inventory on half-fabric level is accurately monitored and placed. Apart from these measures, there are no actions of substantial impact expected to be taken.

KPI monitoring

Besides using the current base stock levels and updating them, we recommend Nedap to gain more insights into their own KPIs and parameters. When they obtain data about their performance and other parameters, they are able to utilize the simulation model even more. When more accurate data is put into the model, more accurate data comes out. In addition, tracking the KPIs will give Nedap the opportunity to evaluate their performance, but also the performance compared to the expected performance form the simulation model. Tracking the KPIs indicates whether the input parameters need to be altered due to unforeseen behavior.

We have already seen how important the accuracy of lead time is, which asks for accurate monitoring of the lead time from the suppliers. In addition, it is crucial to evaluate the simulation by monitoring the service level and the average inventory, this shows the accuracy of the model itself. Finally, it is crucial to continue actively monitoring KPIs that are already available such as: production failure rate, demand/sales, and order sizes. It is important that these values are being monitored since they keep the simulation model up to date.

Sensitivity analysis
From the sensitivity analysis in Section 5.2, we already gained a lot of interesting insights. Therefore, we recommend Nedap to create insights on their own as well. By using the simulation model to create these sensitivity analyses, Nedap can create many interesting quantitative insights, which can be used as a foundation for both strategic and tactical decisions. An example in which a sensitivity analysis could be used as a quantitative foundation for a strategic decision is a production scale up. A sensitivity analysis can provide insights in when and how much production needs to scale up in order to fulfill demand. Another application would be to evaluate the effect of the lead time Nedap promises their customers. A sensitivity analysis shows Nedap if it is possible to lower the promised lead time to their customer, if needed.

Volatile lead time
The final recommendation is about determining the lead time. From the sensitivity analysis of the lead time, we saw how lead time can impact the objective values. By creating agreements with suppliers on lead times and service levels, Nedap makes sure that the lead time from the simulation is promised. Therefore, achieving this lead time results in the optimal usage of the base stock levels provided by the simulation.

6.4. Future research
This section shows the opportunities for future research. We provide four possibilities for future research.

Optimization algorithms
We proposed a model that uses SA to optimize the base stock levels in a multi-echelon multi-item inventory production system. However, there are more algorithms available to approach such a problem. The literature provides us already with alternatives for similar problems such as machine learning algorithms and other (meta)heuristics (Pirhooshyaran & Snyder, 2020; Taleizadeh et al., 2016; Zhao & Sun, 2010). It is interesting to see how these other algorithms perform in terms of running time and objective value when compared to the SA algorithm for the current system.

Stochastic lead time
Secondly, incorporating stochastic lead time within the model is a possible improvement to our model. Including these stochastic values to evaluate the impact of stochastic lead times is another interesting possibility for future research. When there is accurate information available about the distribution of the lead time, including this lead time makes the model more realistic. Including stochastic lead times within a simulation model should be straightforward. However, formulating stochastic lead times within an exact model is likely to be unattainable.

Demand distribution

Next to that, we assumed normally distributed demand for our model. In our case normally distributed demand matches the historical demand, however this does not have to be the case for other systems. Therefore, exploring different distributions for demand is an addition to our model. Including different distributions makes the model more applicable for other systems.

Alternative production planning

Finally, we focus on the base stock levels in an inventory production system, however, we assume all production related parameters to be fixed. For instance, we do not focus on the production planning and thus the order in which production is planned. We evaluated an alternative production planning in Section 5.3, but we do not have a lot of information on what the impact is of this production planning and which production planning is optimal for our inventory production system. Focusing on the production planning and its effect on the inventory parameters is interesting material for future research.

References

Benjaafar, S., Cooper, W. L., & Kim, J. S. (2005). On the benefits of pooling in production-inventory systems. *Management Science*, *51*(4), 548–565. https://doi.org/10.1287/MNSC.1040.0303

Brusco, M. J. (2014). A comparison of simulated annealing algorithms for variable selection in principal component analysis and discriminant analysis. *Computational Statistics and Data Analysis*, *77*, 38–53. https://doi.org/10.1016/J.CSDA.2014.03.001

Bylka, S., & Rempala, R. (2007). On a Nonlinear Production-Inventory Problem. *Advances in Information Processing and Protection*, 447–457. https://doi.org/10.1007/978-0-387-73137-7_40

Chopra, S., & Meindl, P. (2001). *Supply Chain Management: Strategy, Planning, and Operation*. Prentice Hall. https://www.scholars.northwestern.edu/en/publications/supply-chain-management-strategy-planning-and-operation

Chopra, S., & Meindl, P. (2007). *SUPPLY CHAIN MANAGEMENT* (3rd ed.). Pearson.

Du, J., Jiao, Y. Y., & Jiao, J. (2005). Integrated BOM and routing generator for variety synchronization in assembly-to-order production. *Journal of Manufacturing Technology Management*, *16*(2), 233–243. https://doi.org/10.1108/17410380510576859/FULL/XML

Eppen, G. D. (1979). Effects of Centralization on Expected Costs in a Multi-Location Newsboy Problem. *Management Science*, *25*(5), 498–501. https://doi.org/10.1287/mnsc.25.5.498

Ghrayeb, O., Phojanamongkolkij, N., & Tan, B. A. (2009). A hybrid push/pull system in assemble-to-order manufacturing environment. *Journal of Intelligent Manufacturing*, *20*(4), 379–387. https://doi.org/10.1007/S10845-008-0112-6/METRICS

Grasso, E. T., & Taylor, B. W. (2007). A simulation-based experimental investigation of supply/timing uncertainty in MRP systems. *Http://Dx.Doi.Org/10.1080/00207548408942468*, *22*(3), 485–497. https://doi.org/10.1080/00207548408942468

Graves, S. C., & Willems, S. P. (2003). Supply Chain Design: Safety Stock Placement and Supply Chain Configuration. *Handbooks in Operations Research and Management Science*, *11*(C), 95–132. https://doi.org/10.1016/S0927-0507(03)11003-1

Haddock, J., & Mittenthal, J. (1992). Simulation optimization using simulated annealing. *Computers & Industrial Engineering*, *22*(4), 387–395. https://doi.org/10.1016/0360-8352(92)90014-B

Hillier, M. S. (1999). Product commonality in multiple-period, make-to-stock systems. *Naval Research Logistics*, *46*(6), 737–751. https://doi.org/10.1002/(SICI)1520-6750(199909)46:6<737::AID-NAV9>3.0.CO;2-M

Iannone, R., Lovaglio, C., & Riemma, S. (2022). *Determining safety stock with backlogging and delivery spare time*. 404–413. https://art.torvergata.it/handle/2108/46707

Jung, J. Y., Blau, G., Pekny, J. F., Reklaitis, G. V., & Eversdyk, D. (2004). A simulation based optimization approach to supply chain management under demand uncertainty. *Computers and Chemical Engineering*, *28*(10), 2087–2106. https://doi.org/10.1016/J.COMPCHEMENG.2004.06.006

Kapuściński, R., & Tayur, S. (1998). A Capacitated Production-Inventory Model with Periodic Demand. *Https://Doi.Org/10.1287/Opre.46.6.899*, *46*(6), 899–911. https://doi.org/10.1287/OPRE.46.6.899

Köber, J., & Heinecke, G. (2012). Hybrid Production Strategy Between Make-to-Order and Make-to-Stock – A Case Study at a Manufacturer of Agricultural Machinery with Volatile and Seasonal Demand. *Procedia CIRP*, *3*(1), 453–458. https://doi.org/10.1016/J.PROCIR.2012.07.078

Kumar, K., & Aouam, T. (2018). Integrated lot sizing and safety stock placement in a network of production facilities. *International Journal of Production Economics*, *195*, 74–95. https://doi.org/10.1016/J.IJPE.2017.10.006

La Fata, C. M., & Passannanti, G. (2017). A simulated annealing-based approach for the joint optimization of production/inventory and preventive maintenance policies. *International Journal of Advanced Manufacturing Technology*, *91*(9–12), 3899–3909. https://doi.org/10.1007/S00170-017-0053-3/METRICS

Law, A. M. (2014). *Simulation Modeling and Analysis (Mcgraw-hill Series in Industrial Engineering and Management)*. 776.

Mason-Jones, R., & Towill, D. R. (1999). Using the Information Decoupling Point to Improve Supply Chain Performance. *The International Journal of Logistics Management*, *10*(2), 13–26. https://doi.org/10.1108/09574099910805969/FULL/XML

Meisel, F., & Bierwirth, C. (2014). The design of Make-to-Order supply networks under uncertainties using simulation and optimisation. *Http://Dx.Doi.Org/10.1080/00207543.2014.904533*, *52*(22), 6590–6607. https://doi.org/10.1080/00207543.2014.904533

Nel, J. D., & Badenhorst-Weiss, J. A. (2010). Supply chain design: Some critical questions. *Journal of Transport and Supply Chain Management*, *4*(1). https://doi.org/10.4102/JTSCM.V4I1.68

Olhager, J. (2010). The role of the customer order decoupling point in production and supply chain management. *Computers in Industry*, *61*(9), 863–868. https://doi.org/10.1016/J.COMPIND.2010.07.011

Olhager, J., & Prajogo, D. I. (2012). The impact of manufacturing and supply chain improvement initiatives: A survey comparing make-to-order and make-to-stock firms. *Omega*, *40*(2), 159–165. https://doi.org/10.1016/J.OMEGA.2011.05.001

Peeters, K., & van Ooijen, H. (2020). Hybrid make-to-stock and make-to-order systems: a taxonomic review. *Https://Doi.Org/10.1080/00207543.2020.1778204*, *58*(15), 4659–4688. https://doi.org/10.1080/00207543.2020.1778204

Pirhooshyaran, M., & Snyder, L. V. (2020). *Simultaneous Decision Making for Stochastic Multi-echelon Inventory Optimization with Deep Neural Networks as Decision Makers*. https://arxiv.org/abs/2006.05608v2

Rajagopalan, S. (2002). Make to Order or Make to Stock: Model and Application. *Https://Doi.Org/10.1287/Mnsc.48.2.241.255*, *48*(2), 241–256. https://doi.org/10.1287/MNSC.48.2.241.255

Sitompul, C., Aghezzaf, E. H., Dullaert, W., & Van Landeghem, H. (2008). Safety stock placement problem in capacitated supply chains. *International Journal of Production Research*, *46*(17), 4709–4727. https://doi.org/10.1080/00207540701278406

Sobel, M. J. (2008). Risk pooling. *International Series in Operations Research and Management Science*, *115*, 155–174. https://doi.org/10.1007/978-0-387-73699-0_9/COVER

Sun, X. Y., Ji, P., Sun, L. Y., & Wang, Y. L. (2008). Positioning multiple decoupling points in a supply network. *International Journal of Production Economics*, *113*(2), 943–956. https://doi.org/10.1016/J.IJPE.2007.11.012

Taleizadeh, A., Salehi, H., & San-José, L. A. (2016). An Inventory-production Model for Supply Chain under Partial Backordering and Disruption. *Iranian Journal of Operations Research*, *7*(1), 18–29. https://doi.org/10.29252/iors.7.1.18

van Donk, D. P. (2001). Make to stock or make to order: The decoupling point in the food processing industries. *International Journal of Production Economics*, *69*(3), 297–306. https://doi.org/10.1016/S0925-5273(00)00035-9

Zaerpour, N., Rabbani, M., Gharehgozli, A. H., & Tavakkoli-Moghaddam, R. (2008). Make-to-order or make-to-stock decision by a novel hybrid approach. *Advanced Engineering Informatics*, *22*(2), 186–201. https://doi.org/10.1016/J.AEI.2007.10.002

Zhao, G., & Sun, R. (2010). *Application of multi-agent Reinforcement Learning to supply chain ordering management*. 3830–3834. https://doi.org/10.1109/ICNC.2010.5582551

Zhou, A. H., Zhu, L. P., Hu, B., Deng, S., Song, Y., Qiu, H., & Pan, S. (2018). Traveling-salesman-problem algorithm based on simulated annealing and gene-expression programming. *Information (Switzerland)*, *10*(1). https://doi.org/10.3390/INFO10010007

Appendices

Appendix A: simulation settings

Warm up period

In order to determine the warmup period, 100 weeks have been run. Once with (relatively) low input variables and once with (relatively) high input variables. Since the goal is to see when the system approaches a steady state, two different scenarios have been selected to get a better understanding of when this happens. In Figure 17 and Figure 18, the warmup period is presented, by mapping the average inventory per week. In both figures it can be seen that the weeks up until week 15 look slightly different from the steady state that is reached after week 15. That is why we determine that after 15 weeks, the system should be in its steady state. Hence, we obtain a warm-up period of 15 weeks.

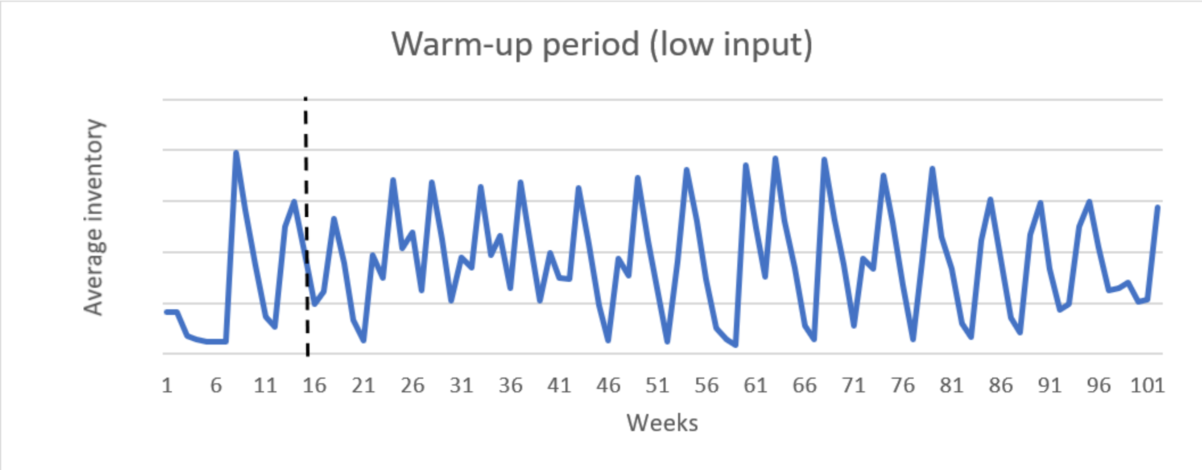

Figure 17. Warm up period, low input

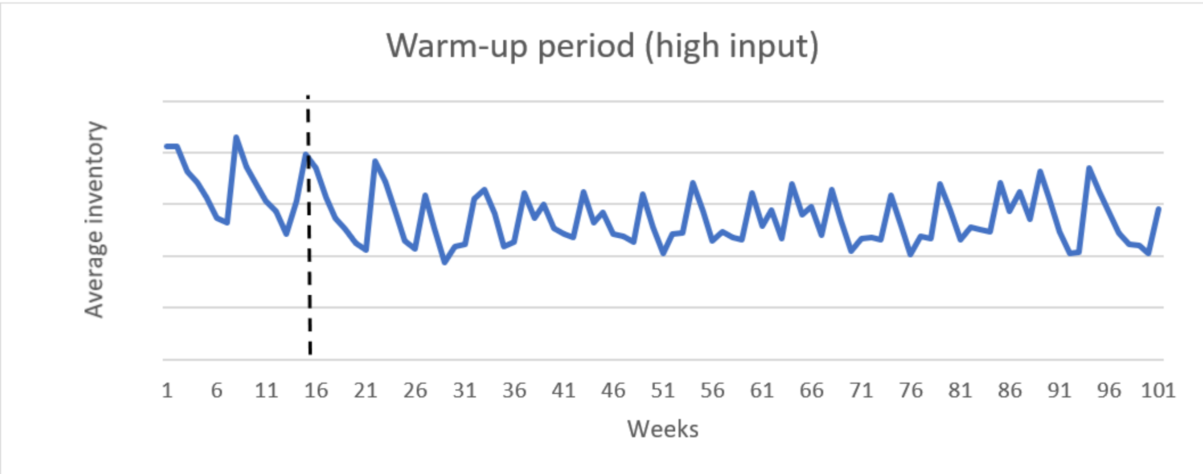

Figure 18. Warm up period, high input

Run length and replications

After the warmup period has been determined, the run length and number of replications should be determined in order to predict the accuracy. This is done by the sequential procedure proposed by (Law, 2014). This procedure is performed for both the evaluation and the optimization simulation. The main difference between these simulations is the required accuracy. For the (more accurate and slower) evaluation simulation, a relative error of 1% with a 99% confidence level is the target. For the (less accurate and faster) optimization simulation, a relative error of 10% with a 90% confidence interval should be enough to give an indication whether one solution performance better than another. And more importantly, with these values we can achieve running times which are far more reasonable. In Figure 19 and Figure 20 it is shown how many replications are needed in order to achieve these values for the given run length.

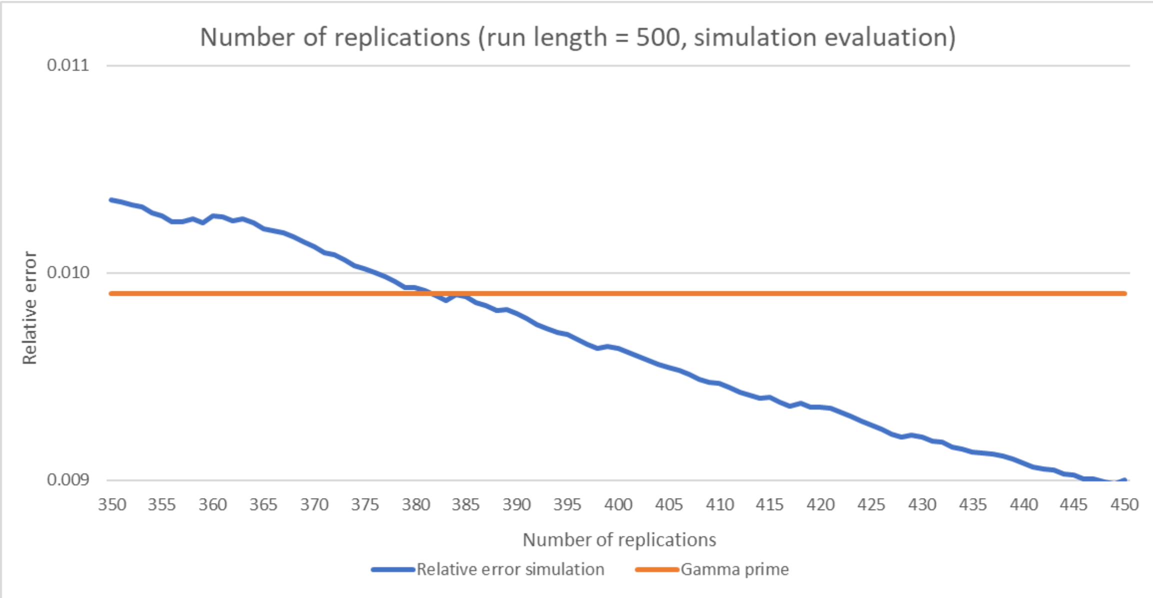

Figure 19. Simulation settings (simulation evaluation)

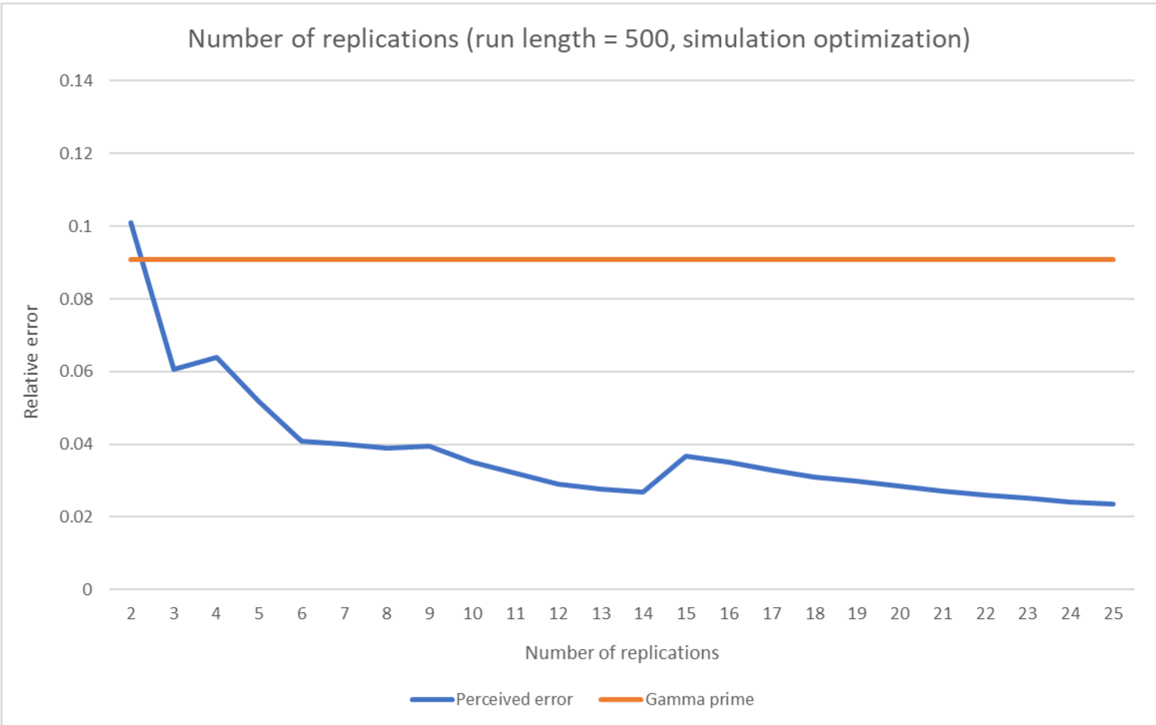

Figure 20. Simulation settings (simulation optimization)

From these figures the values in Table 24 can be achieved. The warmup period is the same for both simulations, since that is the time needed for the simulation to get into the steady state. Next, the run length and number of replications together result in the accuracy, which is why it is decided to set the run length to the same value, namely 500 weeks. This corresponds to ten years, since the production runs 50 weeks a year. Finally, the number of replications is where one can see the difference between the two. However, this higher accuracy comes at a higher computation time. The hardware used was a windows 10 laptop using a 2.4GHz Intel Core i5-1135G7 processor with 16.0 GB of RAM.

Table 24. Simulation settings

	Evaluation simulation	Optimization simulation (SA)
Warm up	15 weeks	15 weeks
Run length (excl warm up)	500 weeks	500 weeks
# of replications	400	3
Computation time	146.05 seconds	1.03 seconds

Appendix B: flowchart simulation

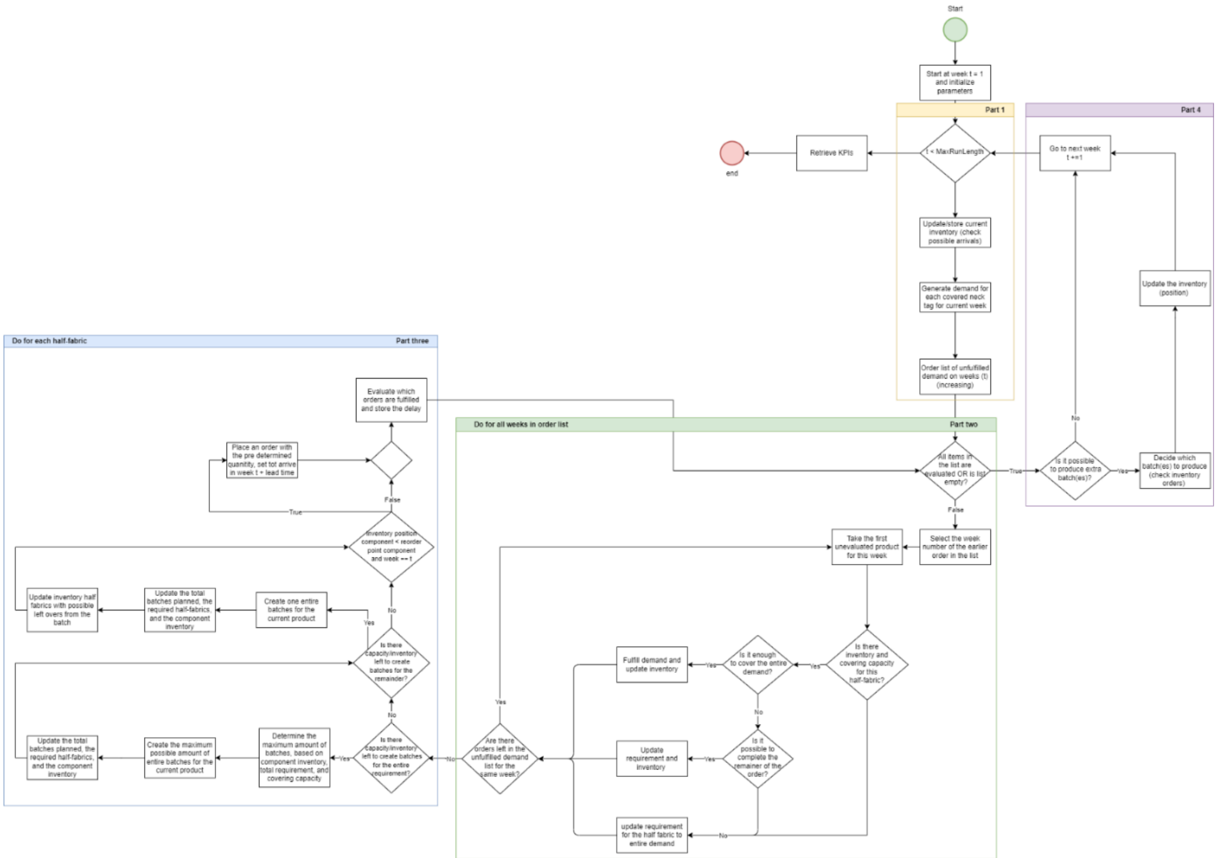

Figure 21. Flowchart simulation model

Appendix C: base stock levels C&M models

Table 25. Base stock levels C&M models with 95% cycle service level

Base stock levels	Components				Half-fabrics						Total base stock
	Comp 1	Comp 2	Comp 3	Comp 4	HF 1	HF 2	HF 3	HF 4	HF 5	HF 6	
Chopra & Meindl (MTO)	276	29	9	209	0	0	0	0	0	0	523
Chopra & Meindl (ATO)	0	0	0	0	247	120	36	2	12	2	419

Appendix D: Q-Q plot demand

In order to evaluate if the demand is even close to being normally distributed, we create a Q-Q plot of the total weekly demand, compared to the normally distributed demand. We got rid of the first and last 5% demand values, since they are obvious outlier. In Figure 22 it is shown that the demand follows the 'y=x' line in the middle part of the graph. The first and last part show small deviations from the line. From this we can see that the normal distribution is a decent fit, however it is not perfect, especially in both tails we see more deviation from the line.

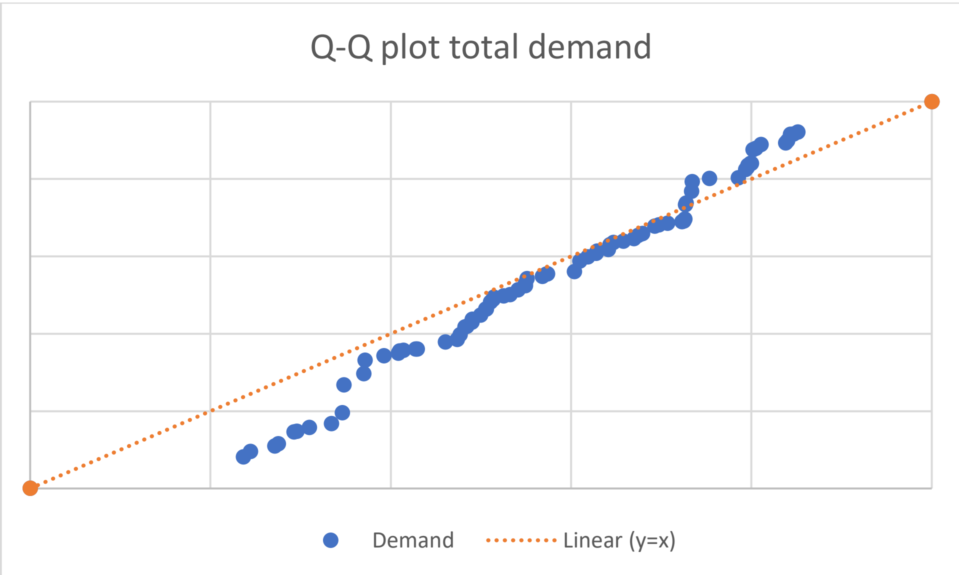

Figure 22. Q-Q plot total demand