Improving service part allocation at Liander's multi-echelon structure

Master thesis Industrial Engineering and Management

Diederick van den Berg
June 2013

Supervised by

Dr. P.C. Schuur	Utwente
Dr. M.C. van der Heijden	Utwente
Ir. M.X. de Vries	Liander

Preface

When I started this research, Liander had a clear long-term goal of where it was heading with its supply chain. They would, in the near future, gain control over the local warehouses and split the urgent order fulfilment from the other material deliveries. Although Liander had a clear view for the future, they did not have any methodology to allocate the service parts of urgent order fulfilment, let alone a methodology utilizing the multi-echelon structure they would manage in the near future. In order to provide Liander with an effective methodology such that they would realize their vision in a cost effective manner, this research was initiated.

This research is completed as a graduation requirement for my Master's degree in Industrial Engineering and Management at the University of Twente. The research is executed within Liander's Logistics department. During the research, two allocation models are developed using insights from recent literature and customized to Liander's situation. In addition to these two models, a simple allocation rule is developed which pursues the duplication of the most extensive allocation model. In the following nine chapters the entire research process is described.

It has been a very interesting period for me from which I learned much about service logistics. I appreciated that Liander gave me the opportunity to perform this assignment. More specifically, I would like to thank a number of people who contributed significantly to this research: Marnix de Vries, my external supervisor, who helped me quickly understand the situation at Liander; Jostijn Marsman, who gave me the opportunity to execute the research at Liander; Peter Schuur and Matthieu van der Heijden, my two supervisors from the University of Twente who contributed not only in assisting with quality report writing, but also by providing suggestions for the construction of the allocation models and simple allocation rule, which is the end result of this research. Also, I would like Elisa Alvarez, former Ph D student at the University of Twente, who contributed considerably by implementing suggestions of the theoretical model in a computer program; Paula Vergunst, who helped me find the right wording in the report; and finally, I thank my wife, Pauline van den Berg-Vergunst, who listened to all my stories.

Now, nine months after the start of this project, I hope and expect that this research will contribute to an improved service parts allocation, which in turn will lead to the purposes of Liander's vision for its supply chain: higher customer satisfaction and cost savings.

Apeldoorn, June 2013,

Diederick van den Berg

Management summary

Problem

Liander initiated a project, Standaard Op Maat (SOM), to reduce the total supply chain cost and increase customer satisfaction. A part of this project is the revision of the current urgent order fulfilment. The general impression is that a sizeable decrease in inventory costs is possible via proper alignment of the various stock locations for service parts, while maintaining a high service level. Therefore, we stated the following objective:

> Create a service parts allocation model for urgent orders, which determines the best storage network and order parameters in corresponding warehouses.

Current practice

Before developing a methodology, the current practice of urgent order fulfillment and its performance is analyzed. We identified the following key problem in the current practice:

The exclusive availability of local warehouses to Netcare mechanics employed for upkeep and failure work. The exclusive availability causes unnecessary emergency shipments from the central warehouse for all other types of work that are not upkeep or failure related, while the SKU may be available in a nearby local warehouse.

We defined the performance of the service part allocation for urgent orders by costs, as a result of Inventory value and emergency shipments, and the number of StoringsVerBruiksMinuten (SVBM) due to waiting time for spare parts. In order to approach the caused SVBM by Logistics, all expensive SKUs (>€ 100.-) are categorized according to their criticality on the gas or electricity network in the event of a failure. We refer to this value as SVBM'.

Models

In the developed model there are four different networks for storing a SKU (excluding van): 1) supplier only (consignment stock), 2) central warehouse only, 3) combination of central warehouse and MWs (manned warehouses) and 4) combination of central warehouse, MWs and UWs (unmanned warehouses). Networks 3 and 4 are configured in a two-echelon system, which enables us to utilize the risk pooling effect at the central warehouse. Besides choosing a network, the model also determines the optimal order parameters (minimizing costs while restricting the number of SVBM') at the corresponding warehouses. In order to find the best solution with the associated decision variables we apply column generation.

In order to measure the performance of order parameters in a network, the following assumptions are made: 1) There is no lateral resupply between local warehouses.,2) No orders are backordered; rather, an emergency order occurs from the next higher supply chain level, 3) Customer and replenishment order size is one.

The column generation model is used for SKUs valued above € 100.-. **The SKUs valued less than € 145.- are all locally stored, irrespective of other SKU's characteristics**, corresponding to networks 3 and 4. A second model is developed for the inexpensive SKUs below € 100.-, based on the assumption that these low valued service parts are always locally stored. This local model saves calculation time in the column generation model and enables the optimization of the replenishment size.

Performance of the models

With the column generation model, an annual improvement of €80,000 in costs, a reduction of 11.2%, is possible, or a reduction in caused SVBM' by 15, a reduction of 55.6%. We believe the actual improvement by the column generation model is higher due to the current backordering of tools. The computational time is around 5 minutes.

The column generation model applies an **increased fill rate differentiation at the warehouses.** For example, at the central warehouse a fill rate in network 2 is pursued between 63.9% and 100.0% rather than between 90% and 98% within regular deliveries. Furthermore, **local warehouses have an increased assortment, but lower stock levels.** Frequently low stock levels at the local warehouses are supported by a high fill rate at the central warehouse. This causes no delay in routine replenishment from the central warehouse and no required emergency shipments from the supplier.

For SKUs valued below € 100.-, we make an annual improvement of over € 250,000.- in costs, a reduction of 61%. On average, a fill rate of 99.7% is pursued for the cheap SKUs at the local warehouses.

Simple allocation rule

In addition to the two models, a simple allocation rule is developed in order to provide Liander with a way to practically approach the promising gains generated by the column generation model. The simple allocation rule consists of two allocation rules based on regression analysis, one which specifies which network to use and one to find the corresponding order parameters. We expect that an overly simplified use of this last rule results in a significant loss of performance, since simply applying average fill rates per network and supply chain level results in cost increase of € 143,000.-, 14%, while also the caused SVBM' triples.

Recommendations

We advise to **use the column generation model instead of the simple allocation rule for the SKUs valued above €100.-.** The results of the first method are not case specific and it outperforms the second method. With respect to **the SKUs valued below €100.-, the local model should be used,** as it easy to apply in practice and the improvement is sizeable. In the future, Liander should **measure the urgent demand at the local warehouses** by order line size, fill rate and date. This will improve the demand forecast and allow additional validation of the models. Furthermore, **the criticality of a failure per SKU should be added in the BIS** (Business Information System)**.** The criticality per SKU varies a lot, the inclusion of this data in BIS supports an improved service part allocation.

Further research

Additional improvement of the supply chain performance is possible by the a**lignment of stock for regular deliveries and urgent orders at the central warehouse.** In networks 2, 3, and 4, SKUs are centrally stored for urgent order fulfilment while these are already present for regular deliveries. **Interchangeability of SKUs** allows further reduction of costs.

The column generation model can be extended by the use of **lateral resupply between local warehouses,** as it results in additional risk pooling effect. Finally, the significant assumption that the customer order size per line is always one does not hold for all SKUs in the column generation model. By **incorporating the occurrence of large order line sizes,** the validity of the model can be improved.

Table of contents

1 Company description

In the framework of completing my Master of Industrial Engineering and Management, IEM, this report describes the graduate research study regarding improving the deployment of goods in benefit of urgent orders. Liander initiated a project to reduce their total supply chain cost, and the subject of this study, revision of Liander's current urgent order fulfilment, is a part of Liander's larger-scale cost reduction project. This report describes a methodology to determine the locations to store the SKUs and their corresponding order parameters. Before we elaborate on the subject, we will start with a description of the Liander company.

Chapter 1 describes the Liander company. Section 1.1 describes Liander's umbrella-company Alliander. Section 1.2 describes Liander's gas and electricity network. Section 1.3 describes the department Logistics, where the research was conducted.

1.1 The Alliander Company

Alliander N.V. is a utility company which has split from Nuon after the passing of the 'Wet Onafhankelijk Netbeheer' legislation in 2006. Alliander consists of 3 divisions: Liander, Liandon and Endinet. The company distributes electricity and gas to a third of the Netherlands. Alliander had 6,647 employees in 2011, €1.569 billion in turnover, and a resulting € 250 million net profit after taxes. The company vision is to be *involved* in regions which they connect, to be a *reliable* energy supplier through uninterrupted gas and electricity delivery and to provide the *best* service in the eyes of their customers. The three italics words are Alliander core values.

Liander is the most important part of Alliander and maintains and constructs the electricity and gas network in the following provinces: Gelderland, Noord-Holland, Friesland, Flevoland and Zuid-Holland. It has 3 million electricity customers and 2.1 million gas customers. Liander covers 90% of Alliander activities. Liandon maintains and constructs complex energy infrastructures within and outside The Netherlands. Endinet maintains electricity and gas network in parts of Noord-Brabant. Figure 1.1 displays the described service areas of Liander and Endinet.

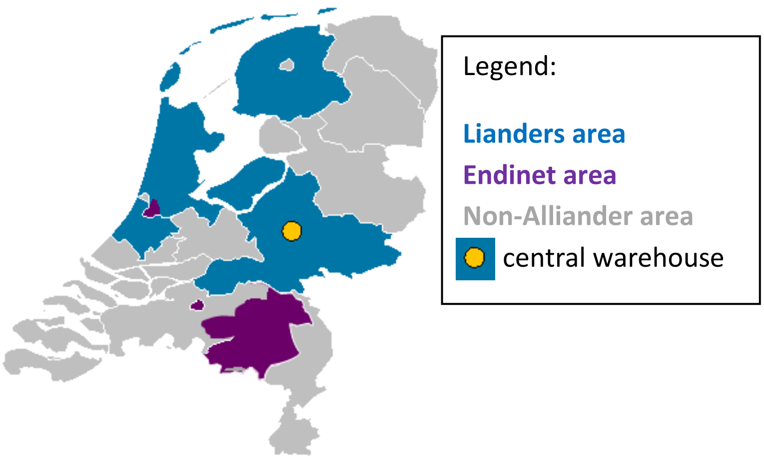

Figure 1.1 Area's where Liander, Endinet or Non-Alliander companies are active.

Source: Annual report 2011 Alliander

1.2 Liander's energy network

Figure 1.2 displays Liander's share, green plane, in maintaining and constructing the electricity network. See Appendix B for the gas network. In both energy types, Liander owns the network from the connection at people's home or company to upstream in the network, in the case of electricity until 50 kV upstream and in the case of gas until 8 bar upstream. Figure 1.2 shows that the high voltage cables (mostly situated above ground) and energy generation is not Liander's responsibility.

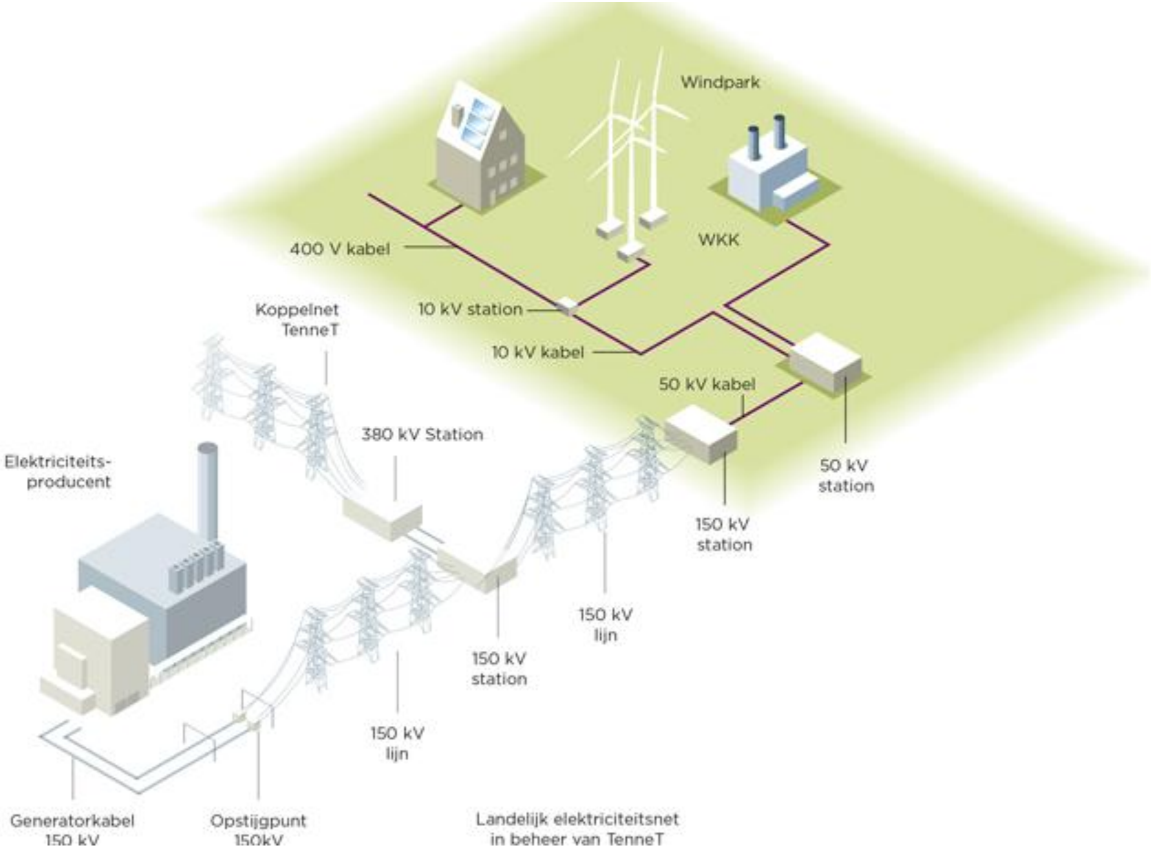

Figure 1.2 Illustration responsibility Liander (green plane) in electricity network.

Source: Annual report 2011 Alliander

As a result of these activities, Liander controls 87,483 km of electricity network and 42,460 km of gas network. Besides the transportation infrastructure, Liander also controls materials for converting gas pressure and electricity voltage and owns the gas and electricity meters in customer's homes and companies. The reliability of Liander's network is one of the best in the world, (Annual report, 2012).

The performance of the electricity network is measured by "StoringsVerBruiksMinuten" SVBM. This indicator measures the downtime of connections in minutes per year, corrected for the total number of connections, 3,000,000. Suppose a failure occurred where 10,000 households did not have electricity for one hour, this results in $\frac{10,000*60}{3,000,000} = 0.2$ SVBM. One minute of downtime at a household corresponds to € 1.-, therefore the prevention of one SVBM may cost € 3,000,000. Liander achieves on average a SVBM of between 20 and 30 minutes per year.

Liander operates in a regulated market, such that the government imposes maximum tariffs for energy distribution and connections. Furthermore, the government also regulates its activities.

1.3 Organisation of Logistics

The Logistics department within Liander is responsible for purchasing and distributing all the materials which are needed for Liander activities of constructing and maintaining the electricity and gas networks. The total annual turnover of these goods is about € 150 million. The Logistics department, including Liander's central warehouse, is located in Apeldoorn.

1.3.1 Organisational context logistics

Figure 1.3 represents the positions of the Logistics department within Liander and other relevant departments with respect to urgent orders.

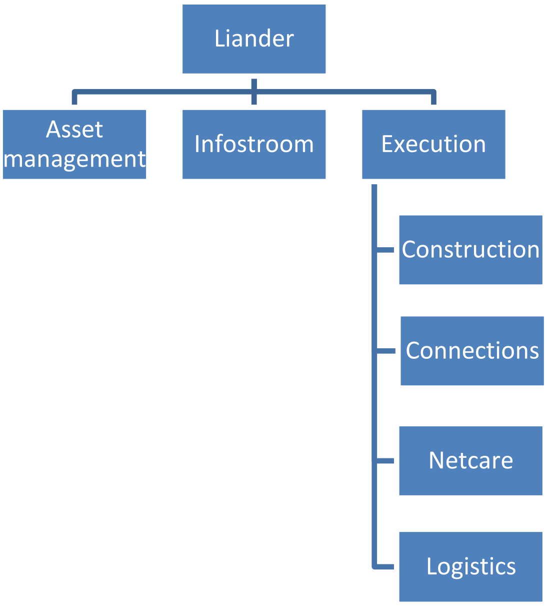

Figure 1.3 From urgent delivery perspective relevant departments for logistics

Asset management is responsible for keeping up and improving the energy network. It develops a future strategy for the network, converts it to a concrete plan and controls the planned activities.

Infostroom controls the traditional electricity meters for small consumption, households and small companies. Furthermore it plans to execute the installation of 5 million smart meters in the coming years.

Execution controls, constructs and maintains the infrastructure of the energy network and large consumption connections.

Construction manages the whole chain which realizes the infrastructure for energy distribution and large consumption connections.

Connections manages the whole chain which realizes the small consumption connections for households and small companies.

Netcare manages energy flows and realizes, controls and conserves the energy network and connections.

1.3.2 Customers

Logistics delivers materials to four different chains within Liander, which consist of both internal and external customers of Liander. Although the chains have different stakeholders, the physical receiver of the product is always a contractor or a mechanic of Netcare.

- **Large consumption connections** - This chain serves large energy consumers, such as businesses with their necessary connections, unit Construction.
- **Small consumption connections** - This chain is serving small energy consumers, like households and small businesses with their necessary connections, unit Connections and Infostroom.
- **Reconstruction and network** – This chain constructs and maintains the energy distribution network for current and future use, unit Construction.
- **Upkeep and failures** – This chain keeps the network in good condition, in order to prevent failures. It is additionally responsible for fixing power outages, unit Netcare.

1.3.3 Products

The main groups of products required for Liander’s activities are the following:

- Gas
  - Meters
  - Tubes
  - Control stations
  - Sockets
- Electricity
  - Meters
  - Cables
  - Transformers
  - Switchboards, Magnefix
  - Sockets
- Various materials
  - Mounting
  - Tools
  - Supportive materials

Within gas and electricity materials, a further distinction can be made. Electricity has materials for high, middle, and low voltage, while gas has different materials for low and high pressure. For most materials their function is clear. Sockets are used to connect tubes and cables. Supportive materials are a wide range of SKUs such as lubricant to insert a tube in a socket or even coffee for the mechanics’ well-being.

2 Research design

Chapter 2 gives an overview of how this research is structured. Section 2.1 describes the reason for research. Section 2.2 describes the problem of the research. Section 2.3 gives the purpose of the research. Section 2.4 provides the research questions. Section 2.5 details the research method.

2.1 Reason for research

Liander initiated a project, Standaard Op Maat(SOM), to reduce the total supply chain cost and increase customer satisfaction. A part of this project is the revision of the current urgent order fulfilment. Logistics will soon gain inventory control over the different local warehouses. The general impression is that a large decrease in inventory costs is possible by proper alignment of the various stock locations for urgent order fulfilment, all while maintaining a high service level. This report describes the background research and resultant proposed methodology to realize this ambition.

2.2 Problem identification

The supply chain of Alliander consists of suppliers, a central warehouse, about 27 local warehouses and about 1300 mechanic vans. Therefore Liander as an organisation can choose to deploy SKU at four supply chain levels: supplier, central warehouse, local warehouse or mechanic van. The central warehouse is controlled by department logistics. The local warehouses are currently locally controlled resulting in little mutual coordination between the central and local warehouses or among the local warehouses themselves.

Because of a long history of local energy distributor mergers which lead to the present structure of Liander, there is a geographical variety in the parts applied in the energy network. This variety causes local warehouses to have their own unique demand characteristics for urgent orders. While the demand characteristics at the central warehouse are known to the logistics department, those of the local warehouses remain unknown.

In order to enhance the performance of the inventory control, Alliander is intending to integrate the control of both the central and local warehouses in Apeldoorn as part of an improvement program, effectively revising the current distribution concept. The current key performance indicators of inventory control are fill rate per order line and inventory costs. Liander is interested in an inventory control system for urgent orders which creates an optimal balance in system control costs, system availability and inventory costs. These three factors should include the probable and impact that individual products have on the energy network. The tool should support inventory control by giving storing positions and ordering parameters per stock keeping unit (SKU). Urgent orders are used for fixing power outages and provide for the smooth execution of planned work when not all required materials are provided by regular delivery for any reason other than stock outages. The general impression is that too many products are held in the supply chain as a result of the overstock of individual SKUs at individual warehouses and the overlap of SKUs at different supply chain levels. This results in high obsolescence costs (R. Hermans). Furthermore, this tool must take into account the different demand characteristics at the local warehouse level.

2.3 Purpose

The purpose of the research is to:

Create a model for service part allocation for urgent demand, which determines the best storage location per article in the supply chain and the order parameters by which Liander can meet the required service level while minimizing its costs.

2.4 Research Questions

1) How does Liander currently handle the urgent orders from a logistical perspective? (Ch. 3)
   a) What is an urgent order?
   b) How does Liander currently deploy its inventory in benefit of urgent orders?
   c) Which performance indicators describe Liander's performance?
   d) What is the current performance on urgent order fulfillment?
2) Which deployment methodology from the literature best fits Liander performance ambition, vision and situation? (Ch.4)
3) How can the best deployment methodology be applied? (Ch.5)
4) What is the performance of the methodology under different system characteristics?(Ch. 6)
   a) How can the applied deployment methodology be verified and validated?
   b) Which parameters are adaptable?
   c) What improvement can be realized by the methodology?
5) To what extent can the methodology be reduced to a rule of thumb?
   a) How to choose storage locations?
   b) How to determine order parameters?

2.5 Research Method

By interviewing the persons from different departments concerned which urgent order fulfilment, studying available protocols and documents and visiting the different storage locations, the current process of urgent order fulfilment is identified (question 1). Figure 2.1 displays a model which is used to structure this research. It makes a distinction in three parts: control, process and performance. Although the model shows these as three distinct areas, there is in fact a mutual influence (free according to Trott, 2008). Furthermore, these areas are studied in light of four different hierarchical levels: strategic, tactical, operational off-line and operational on-line. This structure should give a clear picture of the current urgent order fulfilment (question 1).

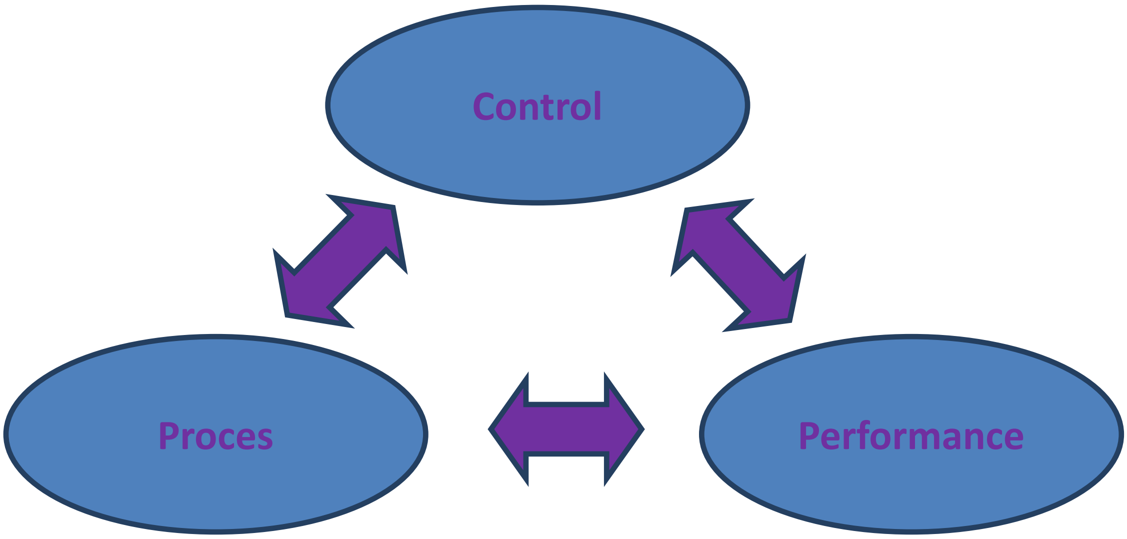

Figure 2.1 Model interaction business unit

Source: Free according to Trott, 2008

The demand for urgent orders at the local warehouses is not directly registered in SAP, the Business Information System (BIS) used at Alliander; however, the demand characteristics at the local warehouses of urgent deliveries can be deduced from the goods delivered from the central warehouse to the different local warehouses and the emergency deliveries from the central warehouse to the different regions. These data sets form the basis for demand characteristics estimation (question 1).

Liander has already made a detailed plan in consultation with the stakeholders named “standaard op maat” (SOM). This project plan contains among others, three distribution concepts: 1) standard work, 2) on order unless, and 3) urgent orders. The first is a material concept covering fast moving, cheap SKUs. The second covers expensive, slow moving SKUs. The last concept is the focus of this research, a vision for urgent order fulfilment. We will study these SOM documents especially with respect to urgent orders and we will also interview the initiators and project managers (question 1).

To bridge the gap between Liander’s current situation and its objective, a literature review is conducted in order to identify the possible deployment methodologies. In a meeting with the different internal stakeholders, we select an allocation methodology (question 2).

During the course of the research it must be determined if one methodology can return both the inventory location and the order parameters or if two separate methodologies, in hierarchical relation, must be developed to meet both needs. According to Francis and White (1974), there are basically two types of mathematical models, descriptive and prescriptive. The first is applied to describe the behaviour of a system. The second is used to prescribe which actions must be taken in order to reach an optimal solution to a present problem. Consequently, a prescriptive or normative model requires a measure of effectiveness in order to determine which solution is the best among the different alternatives. An allocation methodology is a normative model which therefore requires a measure of effectiveness. Therefore, the relevant costs are identified and related to the decisions to be made (question 3). Furthermore, we must determine during the course whether one solution approach fits all types of SKUs or if a separate model is required for inexpensive SKUs.

After the methodology is chosen and applied to Liander's situation, it is tested in order to identify how it will perform in practice. Also, the performance of the methodology will be measured under different system characteristics (question 4). Finally the methodology will be simplified down to some practical rules of thumb that might, depending on its performance, be effectively implemented. (question 5).

The actual replenishment of inventory needed for urgent orders, which can be done from the central warehouse, is outside the scope of this research.

3 Analysis: urgent order fulfilment

Chapter 3 describes the distribution organisation with emphasis on urgent order fulfilment divided into section 3.1 process, section 3.2 control, section 3.3 performance. Section 3.4 describes the vision of urgent order fulfilment. The conclusions are drawn in section 3.5.

3.1 Current distribution process

We will first describe urgent order fulfillment in detail, after which the whole supply chain will be briefly discussed.

3.1.1 Urgent orders

The definition of an urgent order depends at which supply chain level (supplier, central warehouse, local warehouse, or van) the SKU is subtracted. Based on the current practice, we describe it as the result of an unforeseen material need of a mechanic, which more specifically is a result of the following three events:

1. Item failure.
2. Delivery or ordering mistakes.
3. Deviation in working plan from actual situation encountered.

Delayed fulfillment in the first case may result in highly undesirable down time in parts of the energy network. It does not always result in downtime, as other parts of the network frequently take over the gas or electricity supply. In these cases the SKU is still urgently required but the consequences of delayed fulfillment are far less severe. Delayed fulfillment in the second and third cases results in idle time of mechanics. Imagine a team of mechanics who are starting to construct a new part in the energy network. When an order or delivery mistake has occurred, they cannot finish the job with the materials delivered. The same happens when mechanics have dug a hole to adapt the energy network and encounter a different situation than was assumed in the working plan, and consequently different SKUs are required to finish the job.

Delayed delivery can occur for two reasons: long lead time, when it must be delivered from a warehouse far away, or as a result of a stock outage. Stock outage in the case of a power outage is highly undesirable, as previously mentioned that 1 SVBM corresponds to €3,000,000.-.

In the case of an unforeseen material need, the mechanic will first check if his van covers the unforeseen material need, and if not, his corresponding local warehouse is consulted. In this way, most urgent orders are fulfilled. If neither of the lower supply chain levels can cover the unforeseen material need, the order is communicated to logistics. At the central warehouse, an order is defined as urgent if it is required in shorter time than is possible with regular deliveries, generally translating to less than 72 hours. Additionally, Logistics checks to determine if the need is large enough to handle it as an urgent order rather than a regular delivery. When Logistics concludes the need is large enough, the need is coupled to a delivery time window, which can be a 2, 24 or 48 hours delivery. Depending on the availability of the SKU in the central warehouse, the urgent order is delivered from either the central warehouse or the supplier. Figure 3.1 displays the different supply chain levels and shows the responsible organization for inventory levels in benefit of urgent order fulfillment. The three lines from the central warehouse correspond to the three different flows of urgent order fulfillment as a result of the three different time windows (2, 24, or 48 hours).

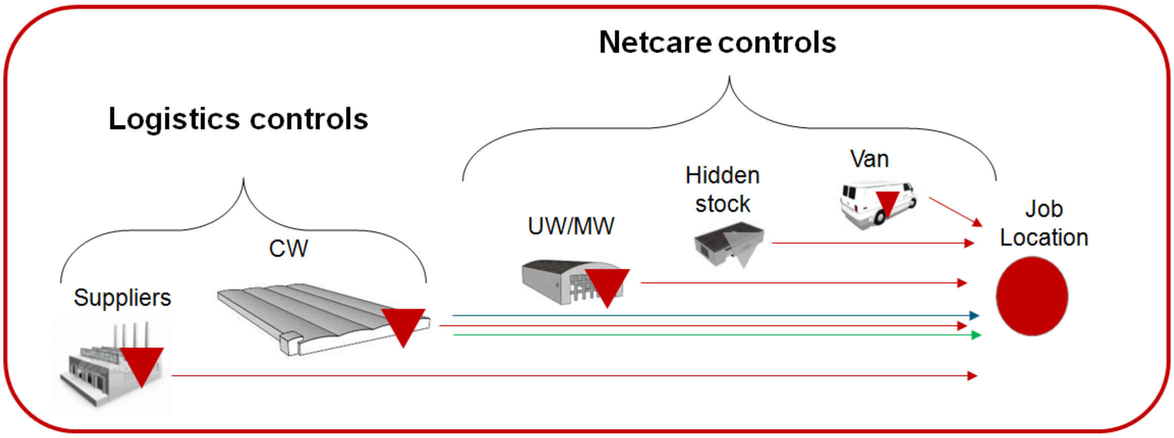

Legend
UW = Liander's Unmanned local warehouse, CW = central warehouse,
MW= Liander's Manned local warehouse

Figure 3.1 Different stock positions in urgent order fulfillment and corresponding controlling organization

Source: Liander SOM vision document

From Figure 3.1 follows that an urgent order is directly fulfilled from different supply chain levels. In general, the more upstream an urgent order is fulfilled, the more time is needed. And increased time is mainly a result of longer distances. The local warehouses, UW's and MW's, are discussed in section 3.1.2.

In practice, there are also many "hidden stocks" at a variety of locations such as mechanic homes and garages, in order to quickly fulfill a urgent order, independent from Logistics and unregistered in the BIS. This situation is fostered by the fact that the inventory deployed at the local warehouses are only for failures and upkeep chain and not for the smooth execution of work within other chains. These alternative inventory locations are undesirable, because of two reasons. Firstly, these stocks are only known to a small number of mechanics, which consequently are the only ones who can use it in the case of an emergency. Secondly, if certain goods become legally prohibited for application in the energy network, there is no control to prevent these now illegal "hidden stock" goods from still being applied in the energy network.

3.1.2 Distribution concept

Figure 3.2 displays the current distribution concept and shows how its customers are served. In the framework there is no distinction between regular and urgent deliveries.

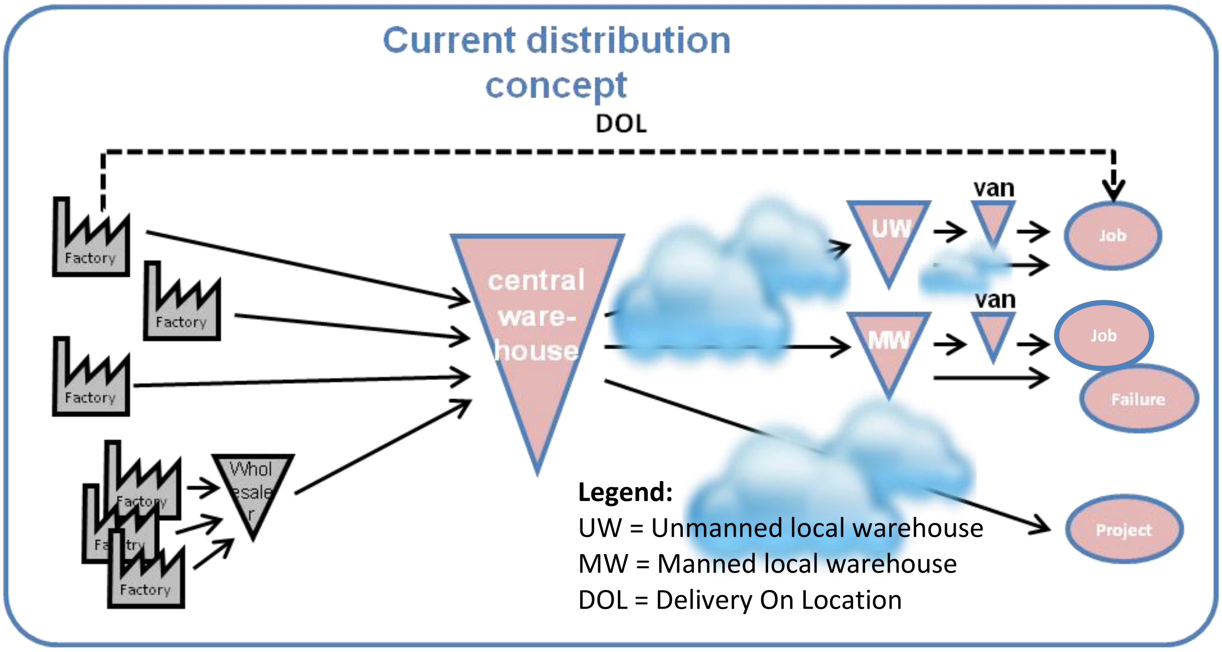

Figure 3.2 Current distribution concept, a four tier system, inventory levels down the stream unknown.

Source: Liander SOM vision document

As noted earlier, SKUs can be stored at four supply chain levels: supplier, central warehouse, local warehouse or van. The clouds represent the current invisibility of stock for Logistics when the SKUs are located at a local warehouse or van, since Netcare manages these stocks, and currently no stock information is exchanged between Netcare and Logistics. Logistics controls the central warehouse and deals with the suppliers. With the supplier, Logistics can make arrangements about delivery reliability and lead time.

Below the four supply chain levels are described, as well as the role in which they function within the supply chain.

Supplier

All suppliers combined deliver around 5400 different SKUs to Liander. There is a lot of variety from SKU to SKU in delivery times. Some deliver in days while others deliver in >2 months. This is mainly due to the fact that some products are produced on order. Consequently, to eliminate these long lead times Logistics can have consignment stock at the supplier. The supplier does not necessarily deliver to the central warehouse, especially if the delivery involves physically voluminous products such as cable drums, which are sometimes delivered on location (DOL).

Central warehouse

The central warehouse, located in Apeldoorn, promises to deliver within three days. Roughly half of the deliveries from the central warehouse are sent directly to a project, and the other half are sent to a local warehouse. Particularly the heavy, voluminous products, which are transported by a truck, are most often directly delivered to a job or project. Also, not all SKUs are stored in the central warehouse. The decision to deliver from stock or order is described in section 3.2.2. Deliveries to a project or job from the central warehouse are usually done by means of a container. Furthermore, the central warehouse processes a recycle and retour flow of goods, which are a result of leftovers when the work is done (not displayed in Figure 3.2). Urgent deliveries should always be directly delivered from the central warehouse to a job or a project.

Local warehouse

Local warehouses are organized in two ways. A distinction can be made between those that are consistently manned during working hours (OSP+ or LSP) and those that are unmanned (OSP-basic). We will refer to the first as Unmanned Warehouses (UW), and the second as Manned Warehouses (MW). The local warehouses are currently owned by Netcare. These internal local warehouses serve Netcare mechanics from a material perspective in three ways:

1. Urgent order fulfillment for failures and upkeep chain.
2. Transshipment of goods for other customer chains, described in section 1.3.2.
3. Replenishment of mechanics vans, which use the local warehouses as their base.

Note that only Netcare mechanics are allowed to subtract products for urgent order fulfillment and only within the failures and upkeep chain. Also, although the electricity and gas meters are fast moving products, they are only stored and issued at a MW, due to their susceptibility to fraud. Besides the meters, there is currently little difference in the product inventory of either a UW or a MW. Figure 3.3 displays the geographical distribution of the local warehouses.

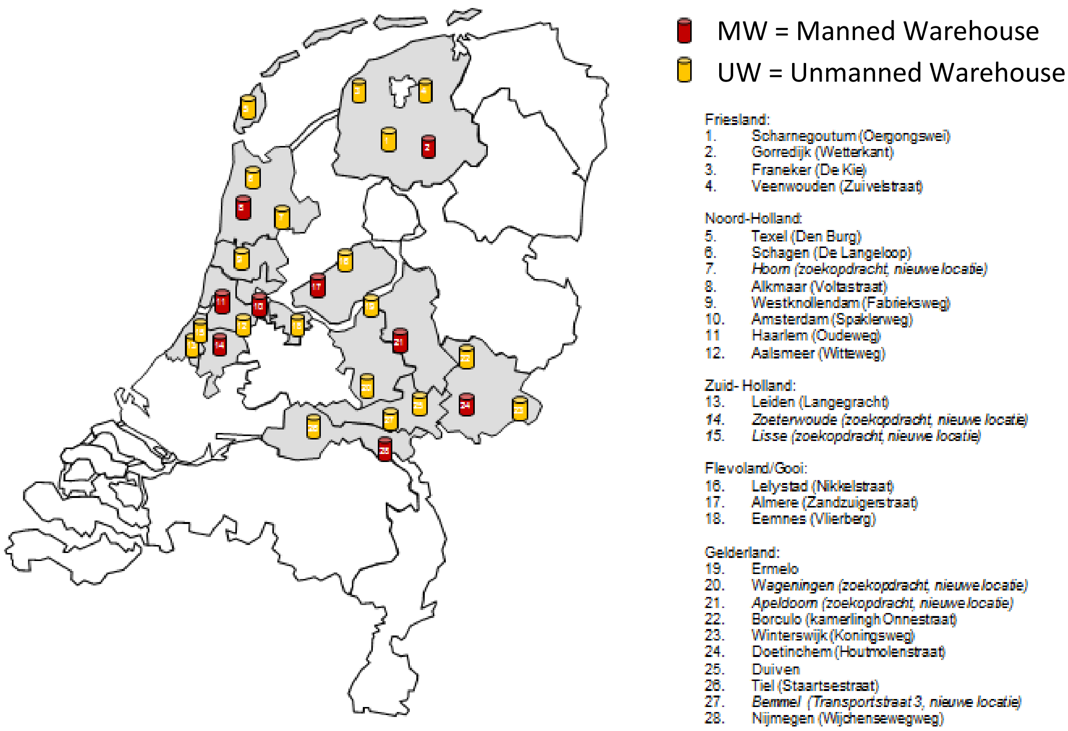

Figure 3.3 Locations of Liander's local, manned and unmanned warehouses.

Source: Liander, local warehouse policy

Note that in Figure 3.3 both the manned and unmanned local warehouses are equally distributed over Liander's service area.

Van

Because of the total weight of a van may not exceed 3,500 kg, the selection of SKUs stored in a van must be carefully determined. This selection is currently done by the mechanic, based on his experience. An empty van with storage installation weighs roughly 1,800 kg. About half of the weight allowance is fulfilled by necessary tools; consequently there is space left for (3,500-1,800)/2= 850 kg of SKUs. These figures are not entirely accurate due to the fact that the mechanic frequently transports products specifically needed for a project and also may be accompanied by one of his colleagues. Therefore, a small buffer is necessary and around 650 kg of SKUs is determined to be the maximum allowed. Some SKUs are too large to transport by van, but generally the weight restrictions are the dominant limitation. Generally, a mechanic is dedicated to one form of energy, either electricity or gas. Also, mechanics work within one of the different chains, which further determines the assortment the mechanic carries in his van. Groups of mechanics are coupled to a local warehouse, therefore the possible SKU assortment of the coupled vans reflects the inventory stocked in their assigned local warehouse.

Analysis of current distribution concept with respect to urgent order delivery

Figure 3.4 supports an analysis of the different identified issues surrounding urgent order fulfillment, displaying the relationships between different root causes leading to intermediate results and end problems.

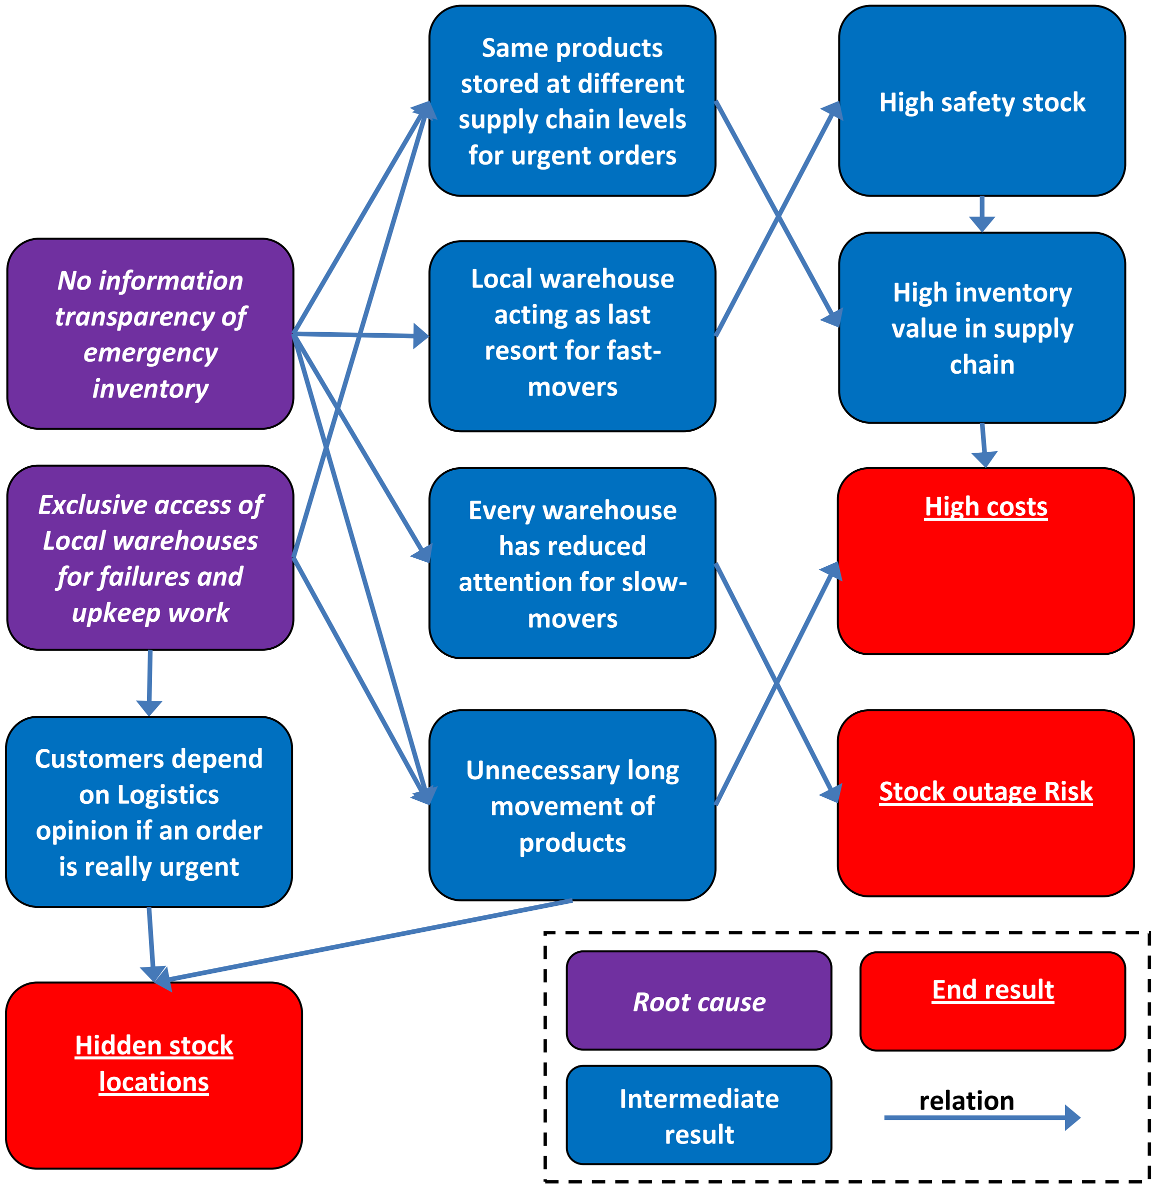

Figure 3.4 Interrelation of causes and problems in urgent order fulfillment

Source: interviews

Logistics has little information over emergency inventory as the stock becomes invisible when it is moved from the central warehouse to one of the local warehouses. This lack of information results in unnecessary movement of SKUs over long distances and therefore preventable idle time at projects. For example, if a socket is required in Amsterdam and out of stock, while a nearby local warehouse may have the socket available, this is unknown, and the material is consequently delivered from the central warehouse in Apeldoorn. This is furthermore deteriorated by the availability of the local warehouses exclusively for Netcare within failures and upkeep work. Consequently, when other customer chains release an urgent order, irrespective of demand location, it is delivered from the central warehouse or supplier. Other internal customers are therefore inclined to hold undesirable hidden stocks in order to quickly fulfill their needs, especially since they are dependent upon the logistics' opinion of whether or not their need is truly urgent.

The lack of information exchange in case of emergency inventory results in unnecessary high stock value in the whole chain, due to two reasons: 1. The local warehouses act as if their warehouse is the last resort for fast movers and hold high safety stock. 2. Identical SKUs are stored at two or more supply chain levels for urgent order fulfillment with no mutual alignment. Section 3.3.2 describes the numerical results of the inventory value. Furthermore, as there is no exchange of information between the supply chain levels, all levels might neglect certain slow movers, with adverse consequences.

Identical storage of SKUs at two or more supply chain levels for urgent order fulfillment is also a result of exclusive availability of local warehouses for failures and upkeep chain. This is due to the fact that an identical SKU can be requested in an urgent order by a different chain, but the chain is only allowed to subtract it from the central warehouse, while a nearby local warehouse also had the SKU stored.

3.1.3 SKUs' characteristics

Figure 3.5 displays the distribution of value of the roughly 5,400 SKUs of logistics. All of these SKUs should be considered for urgent orders. The assortment sometimes consists out of more than one indenture level. All SKUs can be requested by an urgent order.

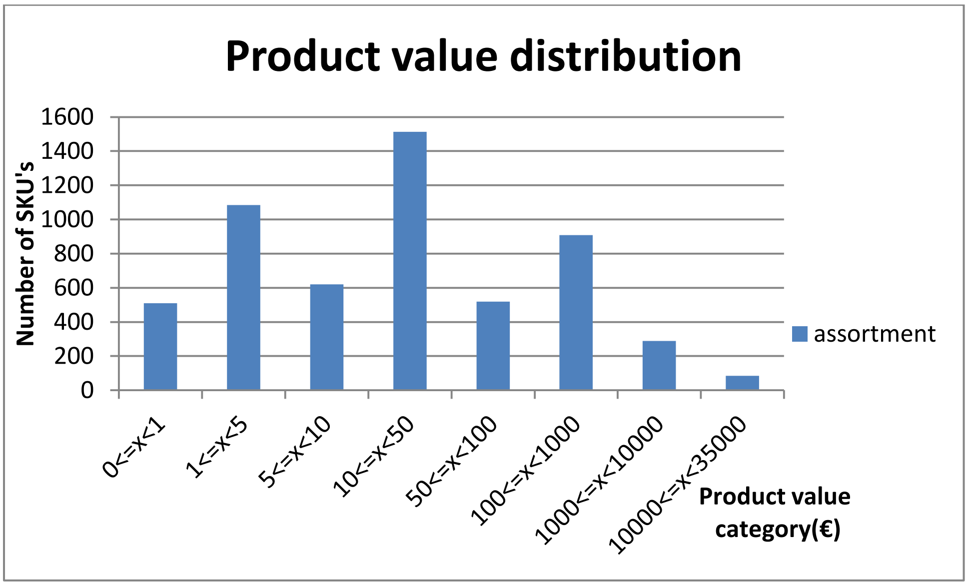

Figure 3.5 SKU value distribution of SKU assortment logistics

Source: SKU assortment in BIS, 2012.

From Figure 3.5 follows that most SKUs are relatively cheap.

The consequences of item failure corresponding to a specific SKU can vary a lot. Some failures only affect a correct measurement of energy consumption at one household, while another can cause a power outage for 10,000 households. Therefore, there is significant criticality variety in delayed urgent order fulfillment per SKU.

3.2 Control

Only the inventory levels of the central warehouse are managed by Logistics. The replenishment policy of the local warehouses is not supported by a clear methodology and is currently under the control of Netcare. The single local warehouse currently under the control of Logistics is also intuitively managed with respect to order parameters. Some local warehouses apply a critical level policy in benefit of urgent orders that are a result of power outages, due to the greater impact of a power outage in comparison to idle time of mechanics. The critical level is intuitively established, in general it is one item. Currently in the central warehouse no distinction is made between inventory for urgent orders or regular orders.

3.2.1 Inventory control policy

Inventory control determines when to reorder a product, based on the incoming orders. The four employees administer their own SKU group. Every night, BIS checks all stock levels. Based on comparison of the current inventory position and the reorder point, BIS produces an order line. The order lines of three nights are combined, resulting in an order to the concerned suppliers. All order lines are manually checked by the inventory controllers.

As has become clear, the replenishment at the central warehouse is organized by a (R,s,S)-system (Silver et al., 1998). A replenishment is released every R period (3 days), only when the inventory position of the SKU is lower than s, such that the inventory position is raised to the order-up-to level S. Sometimes the order quantity must be a multiple of a fixed quantity, because these SKUs are ordered per box or pallet. In such a case, the order quantity is rounded up to the closest number satisfying this requirement.

The replenishments at the local warehouses will be organized by a (s,Q) system (Silver et al., 1998). When the inventory position becomes s, a replenishment order of Q is released.

3.2.2 Service level

The service level is defined as a time-based order line fill rate, namely the fraction of order lines that can completely be filled from stock within initial promised delivery time. For example, when two order lines arrive, both sized 50, and corresponding delivery is 49 and 50, the performance will be 50% although 99% of ordered items are delivered.

The SKUs stored in the central warehouse are basically divided in two groups based on the number of picks: Demand driven SKUs, ordered at the supplier when customer orders (annually less than three picks), and forecast driven SKUs (three or more picks). Only the forecast driven SKUs with six or more picks have a safety stock calculated by an external consultancy company, Involvation. When fewer than six picks but more than two picks are conducted annually, safety stock is chosen intuitively, based on SKU cost, volume and average pick size. When two or fewer picks are conducted, no safety stock is applied, and “demand driven” implies. Within forecast driven SKUs, forecasted by Involvation, a subdivision is made based on turnover and number of picks per year. These groups have their own ambition level in service level, as displayed in Table 3.1.

Table 3.1 Fill rate per order line ambition per SKU group divided in number of picks and yearly turnover.

Yearly turnover/article		# picks A # ≥ 350	 B 35 ≤ # < 350	 C 6 ≤ # < 35
A	V ≥ € 15,000	98%	93%	90%
B	€ 1,500 ≤ V < € 15,000	98%	95%	93%
C	V < €1,500	98%	98%	95%

Source: Liander document of inventory controllers

The categorization which displays is well-considered. The annual turnover against the number of picks supports a high overall service level, while dedicating low inventory values because yearly turnover / number of picks, shows the value per pick, which is equal to the average value per order line. Consequently when a SKU is frequently picked, while yearly turnover is low, a high service level can reached with little inventory value. In this way, a smart differentiation can be made in service level ambition, such that overall ambition is reached with less inventory value needed, as opposed to pursuing the same service level per SKU.

Table 3.2 shows how the forecast driven SKUs, forecasted by Involvation are distributed over the different targets.

Table 3.2 Distribution of SKUs within different service levels target corresponding to table 3.1

Yearly turnover/article		# picks A # ≥ 350	 B 35 ≤ # < 350	 C 6 ≤ # < 35
A	V ≥ € 15,000	93	339	128
B	€ 1,500 ≤ V < € 15,000	27	552	559
C	V < €1,500	2	307	1135

Source: SKU assortment in BIS 2012

As is clear from Table 3.1 and Table 3.2 there is a large number of SKUs where a low service level, <95%, is pursued: 559+339+128=1026 in the central warehouse. Table 3.3 displays the number of SKUs corresponding to the number of picks, for five or fewer picks.

Table 3.3 Distribution of demand driven SKUs corresponding to # picks

# picks	0	1	2	3	4	5
Quantity	611	397	292	227	191	177

Source: SKU assortment in BIS 2012

Table 3.3 shows there is a large number of SKUs where no safety stock at all is applied in the central warehouse, 24 % of the total assortment.

Realized fill rate at central warehouse

Table 3.4 displays the performance on fill rate per order line from the central warehouse. Besides the normal order lines, these scores also describe the urgent orders. Unfortunately, this performance is only measured for deliveries to contractors. The deliveries to contractors describe only 780/5400=14% of the SKUs, while at the same time, within this 14% of SKUs, deliveries also occur to other customers. Consequently a limited view is available.

Table 3.4 Performance on order line fill rate per month from central warehouse, including urgent orders.

	Performance
January	93.54%
February	95.81%
March	96.88%
April	95.39%
May	95.77%
June	91.70%
July	91.19%
September	97.45%
Average	94.54%

Source: Outcomes reportage over 01-01-2012 to-30-09-2012, BIS based.

Table 3.4 shows that the ambition level of an overall 95% fill rate per order line is nearly reached, as the average performance over 9 months is only 95-94.54=0.46% lower. Within these 780 SKUs, 31 SKUs have no safety stock, 4% of assortment. Recall that in the total assortment, 24% have no safety stock. Therefore, it is likely that the fill rate for urgent orders at the central warehouse is even less than 95%.

3.3 Supply chain performance

In order to determine the performance of urgent order fulfillment, the whole supply chain should be evaluated, since all echelon levels contribute to the fulfillment of urgent orders.

3.3.1 Performance indicators

From the supply chain perspective, three performance indicators are relevant for urgent order fulfillment:

1. Inventory value in the supply chain in benefit of urgent order fulfillment
2. Transport costs.
3. The caused number of StoringsVerBruiksMinuten (SVBM) due to waiting time for spare parts.

Emergency shipment costs occur when the order is delivered from the supplier or central warehouse. These shipments interrupt regular processes and a courier is required. Only the additional costs of an emergency shipment from the supplier or central warehouse are necessary to compare to an emergency shipment from a local warehouse.

Approach of caused SVBM

The caused number of SVBM due to waiting time for a spare part is more difficult to measure. There is no information available to measure the exact SVBM that results from stocking decisions. Furthermore, as described in section 3.1.3, possible caused SVBM is variable from SKU to SKU, when waiting for it one hour. In order to estimate the caused SVBM due to waiting time for a spare part, we assume a certain number of caused SVBM per hour waited for a specific SKU. This number of caused SVBM per SKU, per hour is based on a categorization made for all SKUs valued more than €100.- with different material experts. As a result, the caused SVBM in our measurement becomes:

Where:

q_i Average number of caused SVBM when the delivery for SKU i takes one hour (SVBM/hour).
D_i Average delivery time of SKU i (hour).
O_i Yearly occurrence of an urgent order for SKU i in one year.

$$Caused\ SVBM = \sum_{i \epsilon All\ SKUs} q_i * D_i * O_i$$

Note that this measurement of caused SVBM due to logistical decisions is different from the actual SVBM caused by Liander as a whole. Therefore, we refer to this measure from now on as SVBM'. We choose to approach the waiting time for a spare part by considering which warehouse delivered it. When a local warehouse delivered it to its own region, the absolute minimum waiting time was spent. However, when the central warehouse or supplier delivered the spare part, less waiting time was possible, since a local warehouse is generally closer to the demand location. Consequently, the waiting time for a delivery from the central warehouse or supplier is calculated by the central warehouse or supplier delivery time minus delivery time from the nearest local warehouse. All deliveries from the local warehouses are regarded as deliveries to their own region, so no less waiting time was possible and consequently these deliveries do not contribute to the caused SVBM.

3.3.2 Actual performance

According to a service desk employee stock outages in the case of an urgent order at the central warehouse have occurred in 10 years time fewer than 10 times. In practice, we noticed this is due to the interchangeability of SKUs. For example, a specific transformer was ordered, but as it was not available, another transformer was delivered from the central warehouse. Another reason for this low figure is the fact that some urgent orders were judged as not urgent enough to arrange a delivery from the supplier. For these reasons we present two measurements in this section to describe the caused SVBM and the transport costs: the first in which we assume a fill rate at the central warehouse of 100%, reached by interchangeability and the ignoring of certain urgent orders, and the second in which we assume a fill rate at the central warehouse of 95%. This assumption is based on an observed fill rate of 95% within deliveries to contractors between 01-01-2012 and 30-09-2012 obtained from the BIS (see Table 3.4).

Inventory value

Due to a lack of registration, it is not clear what is stored at all the local warehouses, let alone the value of the hidden stock. The inventory value from one local warehouse in Nijmegen, however, is known, as it is currently under control of Logistics. This value is currently comparable to all of the local warehouse, as currently no distinction is made between assortments of manned or unmanned local warehouses.

The value is € 167,934.-, which extrapolated on the relative turnover for spare parts at Nijmegen, (1/18.8), becomes € 3,157,159.-. For the products we optimize in our column generation model in chapter 5, these values become €78,917.- and €1,483,640.- respectively. Besides emergency inventory value in the local warehouse, there is also emergency stock available in the central warehouse. Unfortunately, there is not a clear distinction between inventory value for regular deliveries and emergency deliveries. Therefore, after careful consideration with supply chain expert, M.X. de Vries, we assume 10% of central warehouse inventory value to be dedicated to emergency deliveries. As the current value consists of € 21,000,000,- the emergency value becomes: € 2,100,000.-. From this € 2,100,000.- value at the local warehouses, we apply the same proportion of value to the products in our column generation model at the central warehouse. Table 3.5 summarizes the mentioned values.

Table 3.5 Emergency inventory value and costs, holding costs = 25% of inventory value

	Local value	Central value	Total value	Total costs
All SKUs	€ 3,157,159.-	€ 2,100,000.-	€ 5,257,159.-	€ 1,314,290.-
Modeled SKUs	€ 1,483,640.-	€ 987,140.-	€ 2,470,780.-	€ 617,695.-

Source: Logistics order parameters Nijmegen, BIS

The total values which Table 3.5 displays are in practice even higher, as hidden stock points are also known to exist. We choose to ignore these stock points in the performance evaluation. Furthermore, it is interesting to see that there is a relative high inventory value kept locally rather than centrally. Perhaps the local warehouses are inclined to hold high safety stocks due to the central warehouse fill rate of 95%.

Transport costs

We assume the additional costs of a delivery from the central warehouse to cost €80,-, and when we include deliveries from the supplier, it costs € 205. Over all products, 8,040 emergencies occurred in 102 weeks, and 1,607 emergency shipments occurred for the modeled products. Table 3.6 displays the *yearly* occurrence per item and per order and the corresponding costs. Occurrence on item and order basis differs because numerous urgent order lines consist of order line sizes > 1. Furthermore, Table 3.6 distinguishes between the modeled SKUs and all SKUs.

Table 3.6 Yearly 'Transport' costs for all SKUs and SKUs in model

	Yearly occurrence	Transports costs Supplier excluded	Transport costs Supplier included
All SKUs/order	4,099	€ 327,920.-	€ 538,037.-
All SKUs/item	24,212	€ 1,936,960.-	€ 3,177,851.-
Modeled SKUs/order	819	€ 65,520.-	€ 70,639.-
Modeled SKUs/item	1,626	€ 130,080.-	€ 140,243.-

Source: Emergency shipments BIS 2012 3-1-2011 until 13-12-2012

In Table 3.6 we observe that on a per order basis, only 20%, 819/4099, of the emergency shipments from the central warehouse consist of the modeled SKUs. Recall that the SKUs in the model are roughly 80% of the SKUs above € 100,-. Therefore, it becomes clear that cheap items are frequently delivered from the central warehouse with an emergency shipment. This is suboptimal as the resulting high transport costs of the cheap SKUs can easily be prevented by a small increase in holding costs In the local warhouses. Furthermore, it appears that the modeled SKUs are more frequently ordered with order line size of 1 compared to the other SKUs, since the ratio between the per-order and per-item measurements is 2 (1626/819) for the modeled SKUs, compared to the ratio of 6 for all SKUs (24212/4099).

Caused SVBM'

We assume that the additional waiting time caused by a delivery from the central warehouse in comparison to a delivery from a local warehouse to be 1.67 hours, and between the supplier and a local warehouse to be 71.67 hours. These values are determined via consultation with supply chain experts, J. B. M. Delleman and M.X. de Vries. The relative long time for a delivery from the supplier is based on the fact that it is unprepared for an emergency order. The caused SVBM' can only be measured over a small range of products as we only have the criticality of waiting time per SKU for the limited number of modeled SKUs listed in chapter 6. We approach the caused SVBM' for the Nijmegen region and Liander's total area in four ways, again at the per-order and per-item levels. Furthermore, we include a correction for unnecessary shipments from the central warehouse as these unnecessary shipments, causing preventable waiting time, cannot be prevented by improved stocking decisions. Recall from the discussion following Figure 3.4 that for various reasons, while a SKU was available in a nearby local warehouse, emergency shipment from the central warehouse occurs instead.

We corrected for these unnecessary shipments from the central warehouse by excluding those SKU deliveries to the Nijmegen region which were stored at the local warehouse of Nijmegen. This measurement can only be done for Nijmegen, as the inventory at Nijmegen is known. We therefore applied these values by extrapolation to Liander's total area. Note that in this correction we assume a local fill rate of 100% for the Nijmegen assortment. Therefore, the corrected values can be interpreted as a lower bound. Table 3.7 displays the SVBM' for Nijmegen and Liander's total area in four different estimation approaches. Furthermore, Table 3.7 distinguishes between the inclusion and exclusion of supplier deliveries.

Table 3.7 Caused SVBM' for modeled SKUs on order and item basis and corrected for central coordination or no central coordination for Nijmegen and Liander's area.

Yearly Caused SVBM	Nijmegen Excl supplier	Liander Area Excl supplier	Nijmegen Incl supplier	Liander Area Incl supplier
CW deliveries/order	0.93	13.93	2.91	43.83
CW deliveries /item	1.55	24.58	4.89	77.32
Corrected CW deliveries/order	0.34	*5.09**	1.06	*15.97**
Corrected CW deliveries/item	0.45	*7.14**	1.42	*22.45**

*=estimated value

Source: Emergency shipments BIS 2012 3-1-2011 until 13-12-2012 and expert opinion.

We conclude from Table 3.7 that significant improvement is possible in caused SVBM' by the simple prevention of unnecessary emergency shipments from the central warehouse, since there is a big difference between the corrected and uncorrected deliveries from the central warehouse on a per item basis (71%). The logic of correction for unnecessary deliveries from the central warehouse could also be applied to transport costs on an item level, as it reduces the costs by 28% for the modeled SKUs. The difference in improvement between transport costs and caused SVBM' is that a large portion of the unnecessary transportations were high criticality SKUs.

The displayed SVBM´ values as a result of Logistics performance in Table 3.7 should be interpreted with caution as these values ignore the following:

1. Item failure does not necessarily result in downtime, as other parts of the energy network take over its function.
2. When an SKU is ordered, it is not always required immediately, since there is frequently a setup time. For example, the defect item must be removed first.
3. The calculations above ignore the effect of vans' inventories on the caused number of SVBM'.
4. Down time of gas is not registered in terms of SVBM.

3.4 Vision

In order to improve material supply reliability and save costs, an improvement project is begun "standard op maat", SOM. This project splits the current material delivery from one, as Figure 3.2 displays, into three concepts of material supply: "Standard materials", "on order unless" and "urgent order". The first two concepts deliver the regular orders. "Standard materials" delivers the cheap items from stock within 24 hours, while "on order unless" delivers the expensive items within two weeks, whereby the SKUs are stocked at the supplier as much as possible. For an extensive discussion, see Appendix B. The stock levels in the supply chain in benefit of urgent order fulfillment are also replenished according to these two concepts. Practically it means that with a certain fill rate per order line, the stock levels of emergency inventory are replenished within 24 hours by "Standard work" or are replenished within two weeks by "on order unless" (see Figure 3.6).

3.4.1 Urgent order vision

The materials concept for urgent orders (Figure 3.6) fulfill these orders for same reasons as in the current supply chain:

1. Item failure
2. Delivery and ordering mistakes
3. Deviations in working plan from actual situation encountered.

Based on influencing factors, such as volume, weight, value, impact on network and failure probability, a storing decision must be made with respect to both location and order parameters.

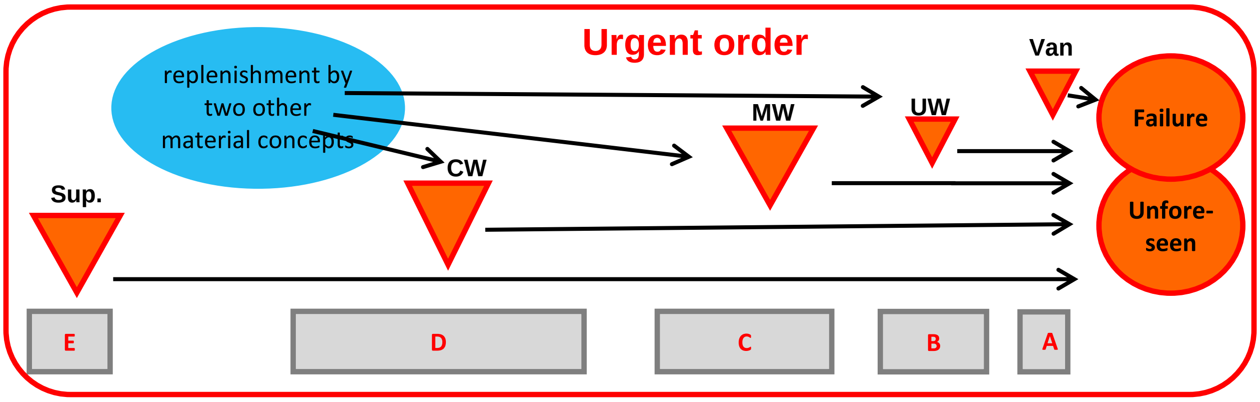

A, B, C, D, E = SKU assortments of spare parts

Figure 3.6 Urgent order distribution concept

Source: Liander SOM vision document

The bars in Figure 3.6 beneath the different supply chain levels represent the number of SKUs in stock corresponding to that supply chain level. The envisioned new urgent order concept excludes the same SKUs stored at different supply chain levels. For example, if a certain cable is stored at the central warehouse, it will not also be stored in a van or at a local warehouse or supplier, as it is believed that it prevents redundancy in inventory value in the supply chain. The exclusivity of SKUs stored at different supply chain levels does not hold true for MWs and UWs. With respect to these two levels, the width of a MW is minimally the width of an UW. Consequently, when we define the displayed letters as subset of SKUs stored, the following is true: $B \subset C$, and the following formula is true for every other combination of subsets: $E \cap D = \emptyset$, with respect to SKUs stored for urgent order fulfillment.

Besides the exclusiveness of SKUs stored at different supply chain levels, the question remains for supply chain levels that represent more than one warehouse (MWs and UWs), if they all have the same SKUs stored. Likely, inventory value can be saved by not storing all SKUs at an individual UW corresponding to the assortment of this supply chain level, since the corresponding region hardly ever launches an urgent order for this SKU. When it does happen, it can easily by fulfilled by an MW or the central warehouse.

In the new situation, other customer chains are also allowed to subtract SKUs from the local warehouses. The demand, with respect to urgent orders, will be registered at the local warehouses. The vans, coupled to a local warehouse, will not use the local warehouse inventory to replenish their vans, but use the local warehouse as a transshipment location for replenishments, which come from the central warehouse.

3.5 Conclusions

In chapter 3 we answered the first research question: *How does Liander currently handle the urgent orders from a logistical perspective?* We saw that an urgent order is a result of an unforeseen material need of a mechanic, which cannot be delivered by a regular delivery. Currently there is not a clear methodology supporting the service parts allocation for urgent orders. Furthermore, the performance is exacerbated by the lack of information exchange regarding stock levels between the warehouses. Liander envisions a situation where every SKU is stored maximally at a single supply chain level (except for the combination of central warehouse and local warehouses), and where Logistics coordinates the stock levels at the different supply chain levels. We defined the performance of the service part allocation for urgent orders by costs, as a result of inventory value and transportation, and the number of StoringsVerBruiksMinuten (SVBM) due to waiting time for spare parts. In order to approach the caused SVBM by Logistics, all expensive SKUs (>€ 100.-) are categorized according to their impact on the gas or electricity network in the event of a failure. Table 3.8 displays the estimated performance of the current service part allocation for 80% of the SKUs valued above €100.-.

Table 3.8 Estimated current performance

Performance indicator:	Score
Caused SVBM’	27.94
Transport costs	€ 104,901.-
Inventory costs	€ 617,695.-
Total costs	€ 722,596.-

Source: Direct and indirect deliveries from central warehouse, BIS, 1-1-2011 until 13-12-2012.

4 Literature

Chapter 4 reviews the relevant literature for a model to solve the problem presented in section 3.4.1. Before we elaborate on the literature, section 4.1 specifies the optimization problem. Section 4.2 describes the column generation procedure. Section 4.3 describes the evaluation literature with respect to both single and two-echelon inventory systems. Section 4.4 describes how to find an optimal stock allocation in a two-echelon inventory system. Section 4.5 reviews the literature regarding the regression analysis required to obtain a simple allocation rule.

4.1 Optimization problem

From section 3.4.1. follows that a combined optimization of storage location and order parameters(or stock levels) is required per SKU, such that an aggregated optimal balance is reached across all SKUs in:

1. Caused SVBM'.
2. Inventory value in the supply chain in benefit of urgent order fulfillment.
3. Transport costs in case of an urgent order.

From section 3.4.1. follows that a SKU is delivered and stored at a single supply chain level. We choose to allow two supply chain levels, in case the SKUs are stored at the local warehouses (UW or MW). In this case, stock at the central warehouse allows us to reap the benefits of risk pooling. Less stock is then locally required for its demand region, since it can be quickly replenished from the central warehouse and, if necessary, directly be delivered to the demand location from the same. From now on we do not use the term storage location(s), but network, as different warehouses at different supply chain levels collaborate for the urgent order fulfillment of each individual SKU. We choose to simplify the problem by excluding the supply chain level van. From section 3.4.1. we deduce that a SKU can then be placed in four different networks:

1. Supplier only - consignment stock
2. Central warehouse only.
3. Combination of central warehouse and MWs.
4. Combination of central warehouse, MWs and UWs.

A delivery policy is a decision regarding in which network and in which quantities (the stock levels) over the available warehouses a SKU is stored. The performance of a delivery policy is expressed in costs and waiting time. Costs are a result of inventory holding costs and transport costs. Waiting time is a result of transport time between demand region and delivery location multiplied by total demand. Note that from the waiting time of a SKU, the caused SVBM' can easily be deduced by including a parameter per SKU indicating the caused SVBM' per unit waiting time.

Per SKU, there are almost an infinite number of delivery policies possible (different networks and stock level combinations), although obviously each SKU requires only a single policy. The delivery policies must *mutually* result in a certain value for caused SVBM', while minimizing the costs. This dependency between the SKUs results in a very large solution space, as individual SKUs cannot simply be optimized separately.

The decision on stock levels influences the fill rates at the different warehouses. The relationship between stock levels and fill rates is non-linear. From the reached fill rates at the different warehouses, the waiting time and transport costs can easily be deduced. Consequently, the decision on stock levels influences performance indicators 1 and 3 in a non-linear manner. In the literature, an approach based upon column generation and decomposition has been previously used in order to solve spare parts optimization problems with aggregate waiting time restrictions over all items. (Kranenburg and Van Houtum, 2007 and 2008, Wong et al., 2007, and Alvarez et al., 2012b). This problem is comparable to our problem, all though we do not have a waiting time restriction over all items. However, such a restriction can easily be transformed to a max caused SVBM' restriction over all items by including a parameter indicating the severity of the waiting time per SKU. The column generation approach is suitable for many different allocation strategies, as long as a waiting time can be determined (Alvarez et al., 2012a). In the literature, multiple strategies are mentioned, such as critical policy level. Therefore, we can use this advantage of column generation to determine the expected waiting time of a SKU when placed in different networks.

4.2 Column generation

Column generation allows us to make a decomposition of the problem per SKU and allows us to include nonlinear aspects in the problem. Furthermore, it is a robust method, as adaptations can easily be made in both the definition of and the evaluation criteria of a delivery policy.

This approach consists of an ILP problem, integer linear programming, which chooses a delivery policy per SKU, such that the max caused SVBM' restriction is met while minimizing the costs. Figure 4.1 is a graphical representation of the column generation approach. This method solves the problem, as described above, by two sub-methods, namely the LP-relaxation, A, and the column generator, B. The LP-model can be solved after the creation of an initial set of policies that make a feasible solution possible, step 1. After it's solved, the LP-model returns shadow prices to the column generator, step 2. The column generator adds an unconsidered delivery policy per SKU to the LP-relaxation. It then selects the unconsidered delivery policy with minimal reduced cost, step 3. The determination of the reduced costs of a delivery policy requires the shadow prices of the constraints of the LP-relaxation. The addition of new delivery policies with negative reduced costs to the policy set will change the shadow prices when the LP model is resolved. Consequently, a number of iterations are required to converge for the optimal LP-relaxation solution. The iterations are stopped when we cannot find any more new delivery policies which improve the LP-relaxation, i.e. no policies with negative reduced cost can be found. When no new policies can be found, the ILP is solved with all added policies, step 4, which returns the final solution.

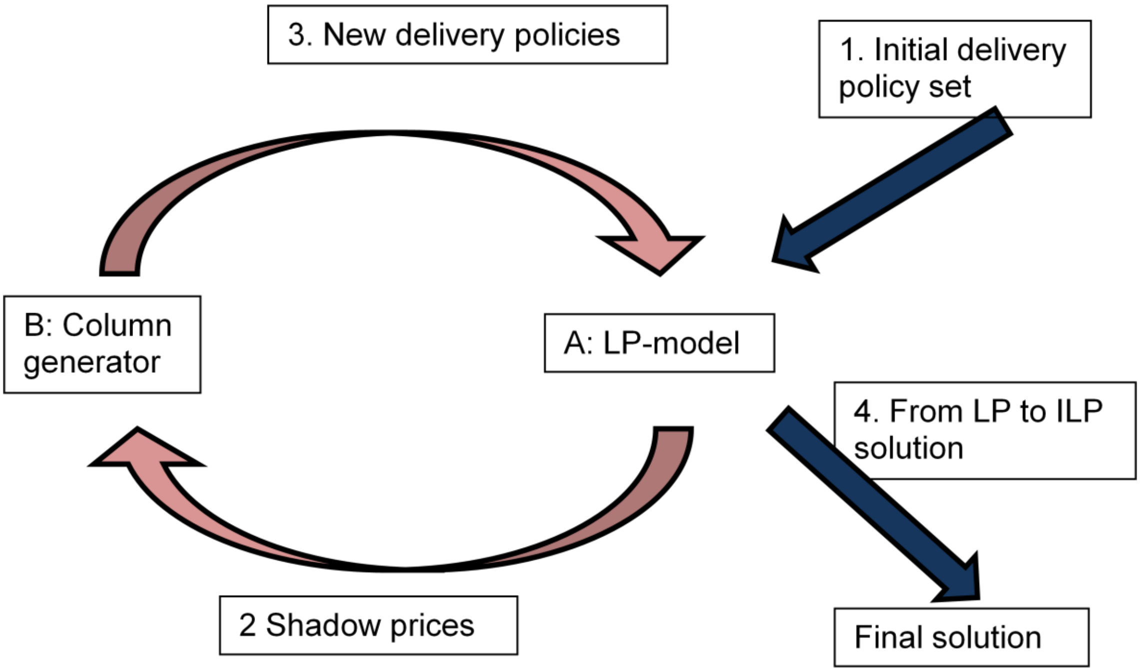

Figure 4.1 Graphical represenation of column generation solution approach

The explanation above of the solution approach as illustrated in Figure 4.1 implies that the column generator, based on certain criteria, is able to find the best stock level combination over the available warehouses, which in turn implies that the performance of a certain stock allocation in a network can therefore be evaluated. Section 4.3 reviews the literature regarding how to evaluate a certain stock allocation in a network and section 4.4 regarding how to find the optimal stock level combination.

4.3 Evaluation of a Network

In order to find a suitable method from the literature to evaluate a network we must first outline a number of assumptions regarding Liander urgent order delivery, described in section 4.3.1.

4.3.1 Assumptions on reality

In order to evaluate the chosen networks we make the following assumptions on reality:

- Demand is Poisson distributed from the regions corresponding to the MWs and UWs.
- A (s-1,s) policy is used at every warehouse.
- Demand that is not fulfilled through the chosen network is not backordered, but a “lost sale” occurs, which in our case refers to an emergency delivery from an infinite source, either contracted or received from another supplier by other means.
- No lateral resupply occurs between the local warehouses, MWs and UWs.
- Delivery from different supply chain levels imply different delivery times to the demand regions.

4.3.2 Evaluation method network 3 and 4

From section 4.3.1 follows that there are exist three possible supply chain levels for an urgent delivery when a SKU is placed in network 3 or 4, namely the local warehouse, the central warehouse, and the infinite supplier. The frequently used Multi-Echelon Technique for Repairable Item Control (METRIC) approach is inapplicable to evaluate a stock allocation, as it assumes all orders are backordered (Sherbrooke, 2004), while Liander’s urgent orders rather result in lost sales.

Consequently, when the local warehouse does not have the item in stock when ordered, an emergency delivery occurs from the central warehouse, and when the central warehouse is also out of stock, an emergency delivery occur from the supplier. The inadequacy of METRIC in this situation is unfortunate, as it not only offers an evaluation of a specific stock allocation, but also a method to find the optimal stock level combination.

Muckstadt and Thomas (1980) elaborated on the work of Sherbrooke (1968), such that the evaluation of networks with emergency deliveries from the central warehouse and the central repair facility is also possible. This model matches our need precisely. Özkan et al. (2011), improved the work of Muckstadt and Thomas (see section below). Appendix D describes the iterative calculations of the model of Özkan et al. in detail. Via these calculations we obtain in a 2-echelon network per local warehouse, and the fraction of demand fulfilled from the three different supply chain levels. Figure 4.2 displays a graphical representation of such a network.

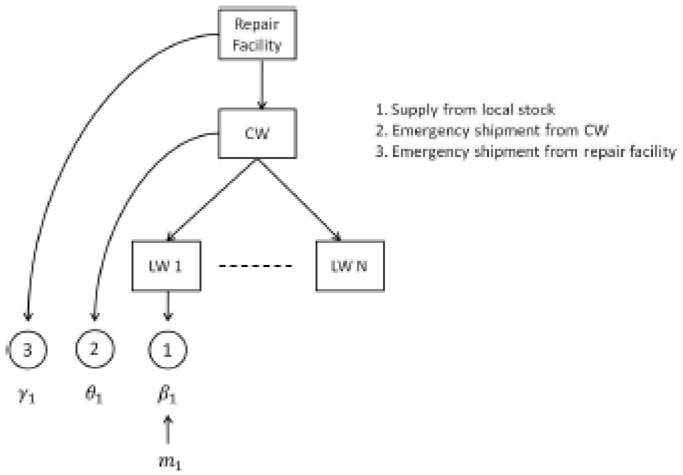

Figure 4.2 Graphical representation of 2-echelon model with lost sales.

Source: Özkan et al., 2011

As can be seen in Figure 4.2, two types of flow occur from the supplier/repair facility and central warehouse, namely emergency and replenishment deliveries.

Model Özkan et al. (2011) versus Muckstadt and Thomas (1980):

The model of Özkan et al. improves the method of Muckstadt and Thomas as follows:

1. It takes into account that a fraction of demand at the central warehouse is lost, due to emergency deliveries from the external supplier.
2. Muckstadt and Thomas (1980), in contrast to Özkan et al. (2011), ignore the dependency between the inventory levels of the central warehouse and the local warehouses. This dependency becomes especially strong when we have a limited number of local warehouses.

With respect to maximum and average absolute differences between the values of reality, obtained by simulation, the method of Özkan et al. outperforms that of Muckstadt and Thomas for al performance measures. Furthermore, the method of Muckstadt and Thomas requires the inventory level of all warehouses to exceed the mean demand during replenishment lead time of the corresponding warehouse, which is not true of the Özkan et al. method.

Complicating assumptions model Özkan et al. (2011)

- The model assumes the replenishment lead time of the central warehouse is exponentially distributed with mean t_0. Alfredson and Verrijdt (1999) show that the outcomes of a system with slow movers is rather insensitive to the probability distribution of the central warehouse replenishment lead time.
- The method consists of iterative calculations, which may lead to long computational time and no convergence.
    - The authors experimented with different problem instances in their numerical study, and obtained convergence for all instances, although they do not have theoretical results on the convergence of this algorithm.
    - The average computation time of the problem instances, N= 2, 4, 10, 20, measured by the authors on an Intel core2 DUO 3 GHz, resulted in an average computation time of less than 4 milliseconds. It is evident then that it is possible to quickly evaluate high numbers of delivery scenarios.

4.3.3 Evaluation method network 1 and 2

The evaluation of network 1 and 2 is much easier to complete, as the network consists of only two warehouses, namely a central warehouse and a infinite supplier (or in the case of network 1, a supplier location and an infinite supplier). Although named differently, both networks function exactly the same. Furthermore, network 2 is actually a special case of networks 3 and 4, one in which no stock is stored locally. However, the evaluations of network 1 and 2 do not require the method of Özkan et al. (2011), but can be evaluated easier and more precisely by the Erlang loss Formula (Ross, 2007), equation 9.

4.4 Determine optimal stock combination

When we add a new column for a SKU, it must be that with minimal reduced costs. Section 4.4.1 describes briefly how the minimal reduced costs formula can be obtained from the shadow price results of the LP-relaxation solution.

4.4.1 Criterion stock allocation: reduced cost

The shadow price of a constraint is the amount by which the objective value is improved when the right hand side of the corresponding constraint is increased by one, assuming the same basis remains optimal (Winston, 2004). Improvement, in our minimization problem, means a decrease in the objective value.

The reduced costs of a new variable are calculated by the following formula, where Δ is the shadow price of the primal problem constraint, $\underline{c}$ is the cost parameter of the new variable in the primal objective function, and $\underline{w}$ its left hand side parameter, multiplying the decision variable in the constraint.

$$reduced\ costs\ = \underline{c} - \Delta \underline{w}$$

Note that as we determine a new variable, the corresponding parameters in the primal becomes the decision variables in the column generator (Manthey, 2011). This is a very simple representation of calculating the reduced costs. In appendix C the logic is explained in the context of a cutting stock problem.

4.4.2 Heuristic: finding optimal stock level combination

The 2-echelon networks, network 3 and 4, have many different inventory level possibilities. Suppose that in network 4 the reasonable range in inventory levels per local warehouse is three, then the number of possibilities becomes $3^{27} \approx 7.6*10^{12}$. In network 3, the number of possibilities within the same reasonable inventory level are significantly less: $3^{9} \approx 1.97*10^{4}$, but in both situations the number of possibilities are far too much to consider all explicitly.

From the calculations follows that the reasonable inventory level range per warehouse should become as small as possible. This can be done to set a lower and upper bound on the inventory level per warehouse. We make use of the four observations from the article of Alvarez et al. (2012a). These observations enable us to find upper and lower bounds on the inventory levels. The first three observations are aligned with our problem context (see appendix E for details). The article also shows why an increase of the stock level at one local warehouse causes an increase in waiting time at the other local warehouses. Lastly, Alvarez et al. concludes emperically between what stock level the optimal stock level of a local warehouse would be, given a proposed stock level of the central warehouse (see appendix E for details). Although the authors do not guarantee this empirical conclusion always holds, they did test its efficacy across scenarios of varied stock quantities (20 and 50 items) and varied number of local warehouses (8 and 16), in which all cases yielded the same solution by complete enumeration (Alvarez et al., 2012a). We assume the empirical conclusions Alvarez et al. draw also hold in our problem context, as our two- echelon inventory systems are comparable.

4.5 Regression analysis

Out of the presented literature we are able to formulate a model which makes a combined network decision for both storage and corresponding stock levels. However, Liander expressed the need for a simple allocation rule which yields the optimal storage network and stock levels based upon a number of SKU characteristics. Therefore, as we have already obtained the optimal decision by the column generation model, we can now attempt to forecasts the results by regression.

4.5.1 Network choice: Ordinal logistic model

Although the real distance between the different categories and networks is unknown, they can be ordered from centralized to decentralized, allowing them to be scaled in ordinal sequence. Ordinal regression enables a network forecast, based upon a number of predictors. It is also possible to ignore the ordering in the dependent variable and apply a multinomial logit model. However, the practical application of a multinomial logit model is more cumbersome, making ordinal regression the best choice for this study.

There are different link functions possible between the dependent variable and the actual network number. We select the logit function (O'Halloran, 2005):

$$\ln\left(\frac{P(event)}{1 - P(event)}\right) = \varphi_0 + \varphi_1 X_1 + \varphi_2 X_2 + \ldots + \varphi_n X_n$$

The formula to the left of the equality sign is defined as logit. It is the natural logarithm of the odds that an event will occur. The odds are equal to the probability that an event will occur, divided by the probability that the event will not occur. The equation above only allows a binary dependent variable, although more than two networks options exist in our situation,. The forecast of multiple options can be done by multiple logits in the form of P(network # = 1 or 2 or 3)/P(network # =4)=τ_3

The ordinal logistic model for a single variable becomes the following equation, where j is the number of categories minus 1 (O'Halloran, 2005):

$$\ln(\tau_j) = \omega_j - \varphi_1 X_1 + \varphi_2 X_2 + \ldots + \varphi_n X_n$$

The definition of the ordinal logistic model implies that every logit has its own ω_j but with the same φ coefficients. This assumes that the influence of the independent variables is identical for the different logit functions. This assumption must be verified by the test of parallel lines. Furthermore, there are additional tests necessary for comparing the observed and the expected counts (e.g., the Pearson residual). These tests are not applicable in cases where the independent or dependent variable have many possible values or are continuous (M.J. Norusis, 2005). In our case, we expect continuous predictors, such as demand, which render these tests inapplicable.

From the estimated ω's and φ's, the probability of the optimal network number becomes:

$$P(network = 3) = P(network \le 3) - P(network \le 2)$$

where,

$$P(network \le j) = \frac{1}{1 + e^{-(\omega_j - \varphi_1 X_1 + \varphi_2 X_2 + \ldots + \varphi_n X_n)}}$$

4.5.2 Stock levels: Log-linear model

In order to predict the stock levels, we use the average waiting time for a SKU. From the forecasted waiting time, the stock levels in the corresponding warehouses can be deduced. In order to predict the waiting time from a number of characteristics of a SKU, we use a log linear model. A log linear model is required due to the lack of a linear relationship between waiting time and the predictors. By logarithmically transforming the dependent variable, the effective relationship remains non-linear, while a linear model can still be applied (Benoit, 2011). A log linear model:

$$\ln Y_i = \propto + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_k X_k + \epsilon_i$$

5 Optimization model

Chapter 5 describes the application of the literature presented in chapter 4 to Lianders problem as specified in section 4.1. Section 5.1 describes the ILP-problem. Section 5.2 describes how we determine the relevant performance indicators of a stock allocation of a SKU. Section 5.3 describes the solution approach and how a new delivery policy is derived. Lastly, section 5.4 describes how the inexpensive SKU are allocated.

5.1 ILP-model

Recall that the proposed model should pursue an optimal balance between caused “StoringsVerBruiksMinuten” (SVBM’) and costs. We pursue a limited number of caused SVBM’ by preventing the combined delivery policies from exceeding a certain value for it. In order to apply such a restriction we need one parameter per SKU in the ILP problem, which specifies how much SVBM’ is caused per hour delivery time. This value depends upon the probability of an urgent demand of this SKU to be used for fixing a failure, and the average number of household connections down in case of a failure due to the SKU.

5.1.1 Primal

sets:

i SKU or item, $i \in \{1,\dots,I\}$.

p Delivery policy, $p \in B_i$, where B_i is the set of delivery policies for SKU i.

Decision variables:

Y_{ip} {0,1} 1 if policy $p \in B_i$ is applied for SKU i and 0 otherwise.

Parameters:

c_{ip} Total cost per year when policy $p \in B_i$ is used for SKU i.

w_{ip} Total waiting time per year for SKU i when policy $p \in B_i$ is used in hours.

q_i Average number of caused SVBM when the delivery for SKU i takes one hour.

mx Maximum number of caused SVBM’.

Objective:

(1)

$$minimize: \sum_{i=1}^{I} \sum_{p \in B_i} Y_{ip} * c_{ip}$$

Constraints:

(2)

$$\sum_{i=1}^{I}\sum_{p\in B_i} Y_{ip} * w_{ip} * q_i \leq mx$$

(3)

$$\sum_{p\in B_i} Y_{ip} = 1 \qquad \forall i$$

(4)

$$Y_{ip} \in \{0,1\} \qquad \forall i, p \in B_i$$

Parameters q_i and mx are given by Liander. We derive the cost and waiting time parameters from the results of a proposed delivery policy. The decision variables within a policy are namely which network is being analyzed, and what the stock levels are at the warehouses in the corresponding network. Section 5.2 describes the determination of costs and waiting time of a delivery policy.

5.2 Evaluation delivery policy

Section 4.3 specified which methods will be used to obtain the fill rates within the inventory system resulting from a chosen stock level. Section 5.2 further specifies how the networks are defined and how the fill rates will be transferred to costs, c_{ip}, and waiting time, w_{ip}. For convenience's sake, we omit the suffix for SKUs i, when we describe the network models.

For costs resulting from transport, we make the assumption that delivery costs from a local warehouse and its replenishments are excluded from the costs calculations. Therefore, the costs for emergency shipments from the supplier or central warehouse are additional costs, compared to a routine delivery from a local warehouse and its replenishment.

Notation

sets

g Delivery network g∊{1,2,3,4}
r demand regions, r∊ { 1,…,R},
n, k warehouses(-1=infinite supplier, -1'=supplier, 0=central warehouse, {1,…,N}= local warehouses,). n∊L_g where L_g is the set of local warehouse in network g, {1,...,Z}=MW and {Z+1,…,N}=UW.

Parameters

t_n Planned replenishment time in weeks of warehouse n∊{0,…,N}.
D_{nr} Delivery time in hours from warehouse n∊{-1,…,N} to region r.
CW Cost of an emergency shipment from the central warehouse.
CS Cost of an emergency shipment from the supplier.
m_r The expected weekly demand rate at region r.
m $= \sum_{r=1}^{R} m_r$, total demand.
h Holding costs.

U_{nrg} {0,1} 1 if region r is linked to local warehouse n else 0 in network g ϵ{2,3,4}.
DS Relative size of urgent order demand of Liander at the supplier.

Performance indicators

β_n Fraction of demand satisfied by local warehouse nϵ{1,…,N} in network (fill rate).
θ_n Fraction of demand satisfied by central warehouse at warehouse nϵ{1,…,N}.
Υ_n Fraction of demand satisfied by supplier at warehouse nϵ{1,…,N}.
s_n Stock level at warehouse nϵ{0,…,N}.
c_g Total annual cost of a delivery policy for network g ϵ{1,2,3,4}, corresponding to the c_{ip} in the ILP-model, see section 5.1.
w_g Total annual waiting time of a delivery policy for network g ϵ{2,3,4}, corresponding to the w_{ip} in the ILP-model, see section 5.1.

Auxiliary variables

LT_n The mean of the realized lead time for warehouse nϵ{1,…,N} in hours.
DL_0 Delay at central warehouse until *replenishment* order is fulfilled in weeks.
β_0 Fill rate at central warehouse in network.

The model we apply for network 3 and 4 assumes demand is fulfilled at the local warehouses, although in our context the demand is met in the regions. This difference is unimportant, however, since irrespective of where demand exactly occurs (local warehouse or region), the fractions from where the demand is met remains the same. We pursue a uniform formulation of the different network costs and waiting time. Therefore, although no local warehouses are used in network 2, we still require them to obtain costs and waiting time:

$$|L_2| = N$$

By U_{nrg} we link each region to its corresponding local warehouse. Note that the following holds:

$$\sum\nolimits_{n=1}^{|L_g|} U_{nrg} = 1 \qquad \forall\, r \in \{1, \dots, R\}\ and\ g \in \{2,3,4\}$$

When we have obtained β_n,θ_n and Υ_n in a network, given a certain stock deployment, we can easily obtain costs and waiting time, which are now suffixed by g (network). The total annual costs are determined by:

(5)

$$c_g = \sum_{n=0}^{|L_g|} hs_n + \sum_{r=1}^{R} 52 * m_r U_{nrg}\,(\theta_n CW + \gamma_n CS) \qquad \forall\, g \in \{2,3,4\}$$

Note that in network 2 all s_n {1,…,N} are zero. Besides charging costs as a result of holding cost we also charge transports costs for every emergency shipment from the central warehouse and supplier, as it interrupts regular processes and requires fast transportation over relatively long distances. As mentioned before, these costs only reflect the additional costs of a delivery from a local warehouse and its replenishment.

The total annual waiting time is determined as follows:

(6)

$$w_g = 52 * \sum_{n \in L_g} \sum_{r=1}^{R} m_r U_{nrg} (\beta_n D_{nr} + \theta_n D_{0r} + \gamma_n D_{-1r}) \qquad \forall\, g \in \{2,3,4\}$$

Note that the following holds in equation 6, since every demand is met:

$$\gamma_n + \theta_n + \beta_n = 1 \qquad\qquad \forall\, n$$

5.2.1 Supplier – Network 1

Network supplier is a special case which requires a different cost and waiting time determination. The delivery times D_{-1r} will not apply, because when we make prior arrangement with the supplier, it will deliver more quickly than D_{-1r}. Moreover, it will also charge for the consignment stock, in order to deliver quickly. We have yet to consider a stock outage, since it is unrealistic to expect the supplier to achieve a fill rate of 100%. The consignment stock for an emergency shipment from the supplier are lower, due to the risk pooling effect, as Liander only contracts suppliers whose yearly turnover in sales is less than 50% dependent on Liander. We introduce a new parameter DS which denotes which fraction of urgent demand at the supplier origins from Liander. Note that when DS decreases, the portfolio effect increases. Furthermore, we assume that the supplier achieves a certain fill rate of X<1. Waiting time and costs become:

(7)

$$w_1 = 52 * \sum_{r=1}^{R} m_r \,(X D_{-1'r} + (1 - X) D_{-1r})$$

(8)

$$c_1 = s_{-1} * DS * h + 52 * m * CS$$

s_{-1} follows from the Erlang loss formula where the demand size becomes m/DS and the lead time is equal to the lead time of the central warehouse. Erlang loss formula:

where:

$\rho_n = \sum_{r=1}^{R} m_r\, U_{nrg} t_n$ product of demand at warehouse n and its replenishment time.

(9)

$$ERL(s_n, \rho_n) = \frac{\frac{\rho_n^{\,s_n}}{s_n!}}{\sum_{i=0}^{s_n} \frac{\rho_n^{\,i}}{i!}}$$

s_{-1} becomes the minimal stock level that results in a fill rate of at least X. See Figure 5.1 for a graphical representation.

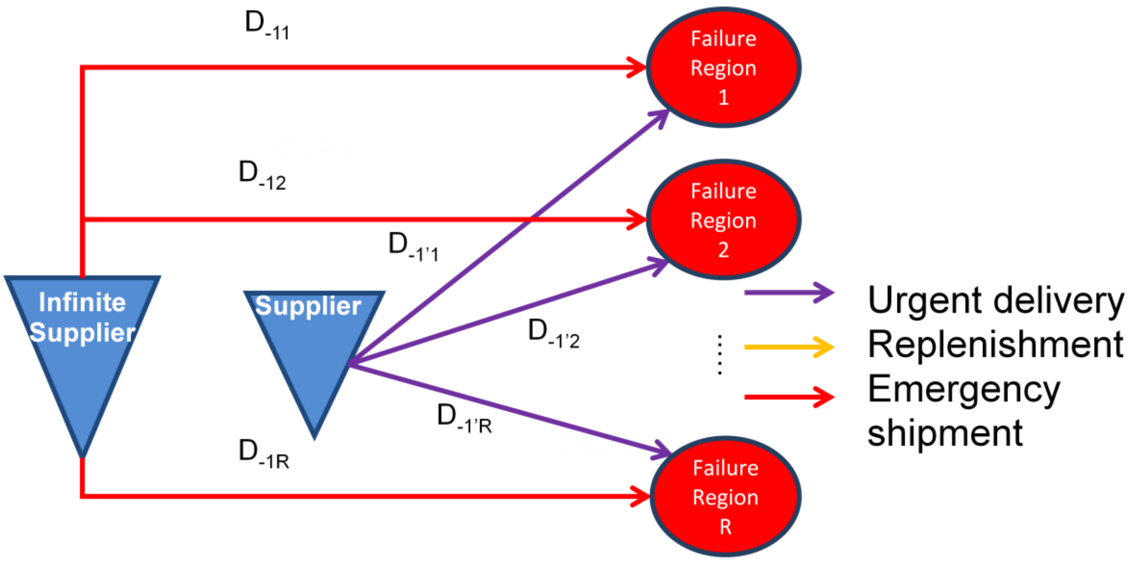

Figure 5.1 Graphical representation network 1, only supplier delivers the SKU.

Within Figure 5.1 represents the infinite supplier a variety of solutions to fulfill the urgent demand. The Infinite supplier can include, for example, deliveries by other suppliers of the SKU, deliveries by a colleague utility company performing the same activities in other parts of the Netherlands, or even the redirection of electricity supply through a generator.

5.2.2 Central warehouse – Network 2

Network 2 is exactly the same as network 1, other than that we can now determine the fraction of demand satisfied by the central warehouse, θ_n, by the stock level s_0. The determination of θ_n is not labour-intensive as inventory is only kept at the central warehouse. It is a single echelon supply chain with one warehouse, where demand arrives according to a Poisson distribution, the sum of the region's arrival rates. To evaluate this delivery network with a certain stock level, we can use the Erlang loss formula, equation 9, as stock outages are served by emergency shipments from the infinite supplier. See Figure 5.2 for a graphical illustration of this network.

In network 2 we obtain the fraction of demand satisfied from the central warehouse, θ_n, and supplier, γ_n, from the results of the Erlang loss formula as follows:

$$\theta_n = 1 - ERL(s_0, \rho_0) \qquad \forall\, n \in L_2$$

$$\gamma_n = ERL(s_0, \rho_0) \qquad \forall\, n \in L_2$$

Since we do not have local warehouses, $\beta_n = 0 \;\forall\, n \in L_2$.

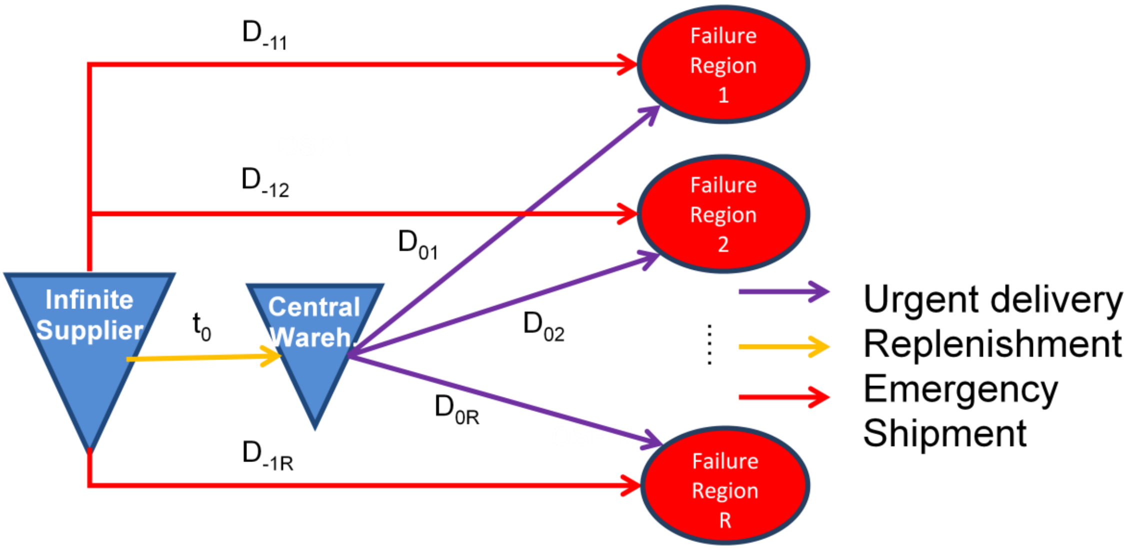

Figure 5.2 Graphical representation of network 2; only central warehouse delivers the SKU.

5.2.3 Central warehouse and UWs and MWs – Network 3 and 4

From section 4.3.2 follows we use the method of Özkan et al. (2011), to evaluate a stock allocation in networks 3 and 4. Appendix D describes the iterative calculations of the model of Özkan et al. in detail. By these calculations we obtain in a 2-echelon network per local warehouse β_n, θ_n and Υ_n. To clarify the operation of the 2-echelon network, we explain which three things can happen when region r releases an urgent order:

1. When corresponding local warehouse n has the SKU in stock, the item is delivered from this warehouse with delivery time D_{nr}. The local warehouse in turn places a replenishment order at the central warehouse, and the central warehouse places a replenishment order at the supplier. No transport costs occur.

2. When corresponding local warehouse n does not have the SKU on stock, but the central warehouse does, the item is delivered from the central warehouse via an emergency delivery, with delivery time D_{0r}. The central warehouse places a replenishment order at the supplier and CW transport costs occur.

3. When neither the corresponding local warehouse n nor the central warehouse have the SKU in stock, the item is delivered from the supplier via an emergency delivery, with delivery time D_{-1r}, and CS transport costs occur.

Recall that network 4 has N local warehouses and network 3 has Z local warehouses. We will apply this 2-echelon network to network 4 first (see Figure 5.3). In this network, every local warehouse n has one region r, $\sum_{r=1}^{R} U_{nr4} = 1 \quad \forall\, n \in \{1, \dots, N\}$. The local warehouses are both MWs and UWs. Network 3 is a copy of network 4, but differs in the number of local warehouses used, all of which are MWs. Therefore, the MWs in this network also serve the regions of the UWs, resulting in $\sum_{r=1}^{R} U_{nr3} \geq 1 \quad \forall\, n \in \{1, \dots, Z\}$, with higher delivery times to the regions, D_{nr}. Figure 5.4 displays a graphical representation of network 3.

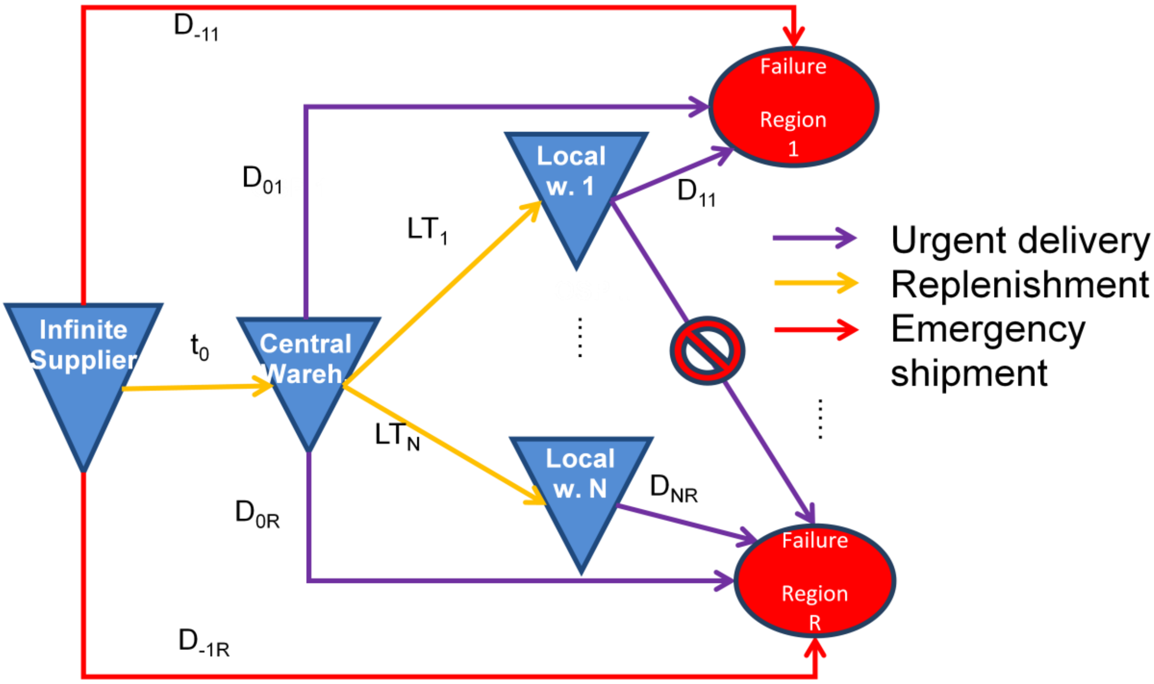

Figure 5.3 Graphical representation of network 4, CW with both UWs and MWs, $\sum_{r=1}^{R} U_{nr4} = 1 \quad \forall\, n \in \{1, \dots, N\}$.

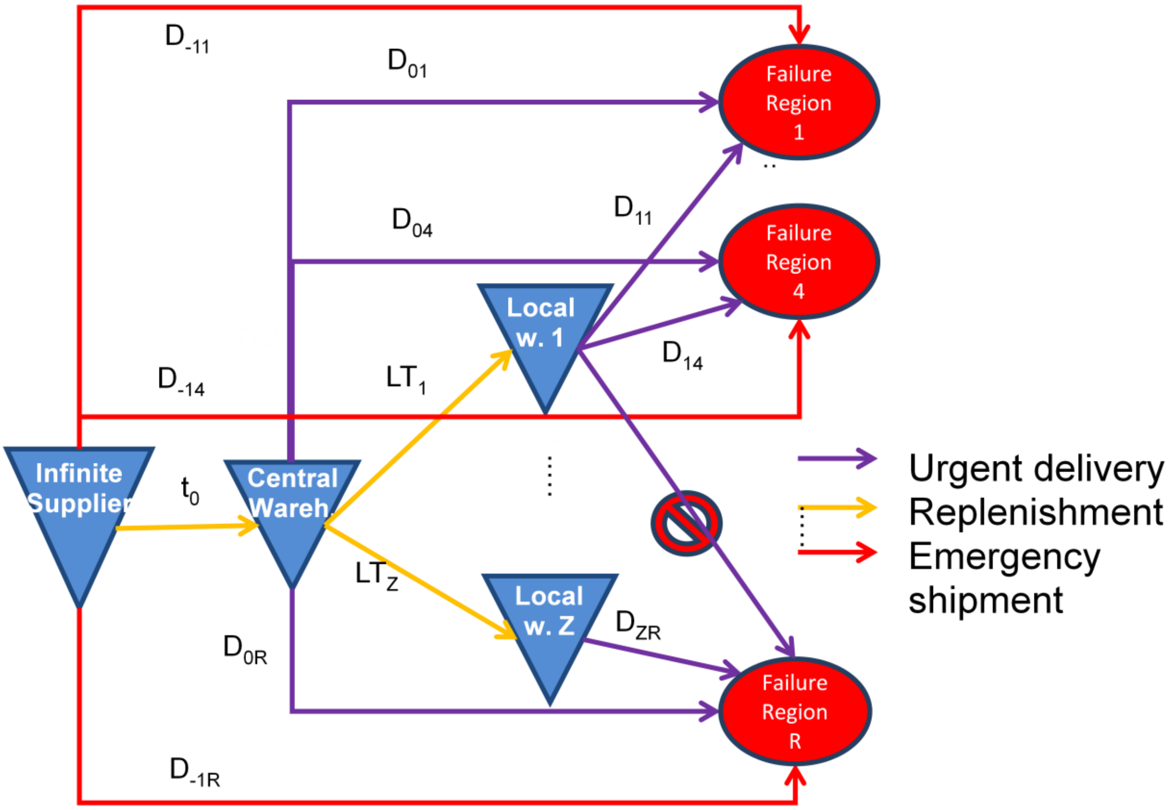

Figure 5.4 Graphical representation of network 3, CW with only MWs, $\sum_{r=1}^{R} U_{nr3} \geq 1 \quad \forall\, n \in \{1, \dots, Z\}$.

Note that in Figure 5.3 and Figure 5.4 we do not apply t_n as the replenishment lead time for the local warehouses. This is due to the fact that the lead time can increase if the central warehouse has a stock outage. Consequently:

$$LT_n = t_n + DL_0$$

Recall that DL_0 represents delay at the central warehouse until a replenishment order is fulfilled. Furthermore, we do not allow lateral transshipment which is illustrated by the forbidden sign through a lateral supply.

5.3 Solution approach

In order to limit calculation time of the whole method, we will limit the number of SKUs for evaluation, and will only consider the SKUs which constitute significant monetary value, >€100,-. The allocation of the low valued SKUs will be discussed in section 5.4.

5.3.1 Create initial primal solution

The shadow prices from the primal relaxation are required in order to obtain a new variable or delivery policy. The primal therefore needs initial delivery policies for all SKUs, such that a feasible solution is possible. The better the initial primal solution is, the better the shadow prices will represent the optimal balance in waiting time and costs, which in turn will likely result in fewer iteration steps. From sections 0, 5.2.1, and 0 follows that waiting time and costs are easily obtained when a SKU is placed in network 1 or 2. For all SKUs, therefore, we add the relevant delivery policies of these two networks.

In network 2, this is done as follows. Suppose for SKU i, stock level S_{i0}, results in θ_{in}≥1- ε, then we add delivery policies for SKU i with the following s_{i0} stock levels.

$$s_{i0} \in \{1, \dots, S_{i0}\}$$

Networks 1 and 2 represent the most centralized stock allocations, resulting in high waiting times. As a result, delivery policies from these networks alone will not likely result in a feasible LP-relaxation solution, as the maximum caused SVBM' restriction is not held (equation 2). Therefore, we also add one network 4 option per SKU, with zero stock at the central warehouse and the minimal local warehouse stock level of s_{in} resulting in β_{in}≥1- ε. Note that network 4 is the most decentralized option, resulting it the least waiting time for a delivery policy with $\beta_{in} = 1 \quad \forall\, n \in \{1, \dots, N\}$. Fill rate as a result of s_{in} with s_{i0}=0 can be calculated by the Erlang loss formula (equation 9), where the replenishment time becomes $t_n + t_0$ (Alvarez et al., 2012a).

As has become clear, a certain value of MX is unpreventable, since the most decentralized allocation still results in waiting time and corresponding SVBM'. Therefore, we first calculate the unpreventable number of SVBM' by the resulting waiting time within the above described delivery option in network 4. We then add this value to the maximum caused SVBM' value. The maximum caused SVBM' can be specified by management. Therefore, as long as the management does not indicate a negative value for maximum caused SVBM', a feasible solution can be reached.

5.3.2 Selection of a new column

We apply the logic of section 4.4.1 to our problem, extending the notation as follows:

A	The shadow price of SVBM' constraint 1 (equation 2).
J_i	The shadow price of policy constraint 2 for SKU i (equation 3).

These shadow prices are generated when we solve the above formulated LP-relaxation of section 5.1.1 with a set of delivery policies. The reduced cost of a new delivery policy p for SKU i becomes:

(10)

$$reduced\ cost\ policy_p = c_{ip} - J_i - w_{ip} * q_i * A$$

The interpretation of the formula above is as follows. c_{ip} and w_{ip} respectively represent the cost and waiting time of a new policy. The shadow prices describe the improvement of the objective value when the right hand side of the corresponding constraint is increased by 1. This definition only applies when the same basis remain optimal. Therefore, $A \le 0$ and $J_i \ge 0$. The improvement in the objective value as a result of less waiting time should be corrected for q_i, which transforms the waiting time to caused SVBM' in the primal SVBM constraint (equation 2).

As long as we can find a policy for one SKU which returns negative reduced costs, improvement of the primal problem is still possible. Consequently, when no policies with negative reduced costs exist for any SKU, the optimal solution to the LP-relaxation is found.

Every time a new policy p is added to the set of policies B_i of SKU i, the solution to the LP-relaxation and the corresponding shadow prices will be different. The addition of a new policy for one item will not drastically change the shadow prices corresponding to the other SKUs. Therefore, in every iteration run (recall that one iteration run corresponds to steps 2 and 3 in Figure 4.1 from section 4.2), we choose to add a policy for every SKU, given the fact that a negative reduced cost policy exists. In this manner, we save computation time by limiting the number of required LP-relaxation solutions.

5.3.3 Reducing numerous delivery policies

From section 5.3.1 follows that the delivery policies from networks 1 and 2 do not need to be considered by the column generator, as these are all already part of the initial policy set. For networks 3 and 4, section 4.4 and appendix E describe the theory from Alvarez et al. (2012a), where the authors make four observations and a empirical observation which enable us to quickly find the optimal stock allocation in a 2-echelon inventory system. The formulas from Alvarez et al. (2012a) in appendix E are already tailored to our problem context

In addition to these four observations and the empirical conclusions, we add two further limitations to a stock allocation in networks 3 and 4, which are required by Liander, M.X. de Vries. These will be discussed below. A welcome advantage of these additional restrictions is that they will speed up the selection of a delivery policy. Before we elaborate on these additional restrictions, we extend our notation by:

s_{in}^{LB} Lower Bound for stock level of SKU i at warehouse n.
s_{in}^{UB} Upper Bound for stock level of SKU i at warehouse n.

Obviously, we do not store units at warehouses corresponding to a region or regions where no demand occurs, resulting in the following restriction:

$$s_{in}^{UB} = s_{in}^{LB} = 0 \qquad \forall \sum_{r=1}^{R} m_{ir} U_{inr} = 0 \ and\ n \in \{1, \dots, |L_g|\}$$

Those warehouses without demand for that specific SKU do not require further consideration.

Restriction 1: Uniformity of assortment.

For uniformity of assortment, we require in networks 3 and 4 that local warehouses have at least one in stock for every item of which demand occurs in the corresponding region(s). As a result we have a lower bound of 1 for the inventory level of local warehouses, s_{in}^{LB}.

$$s_{in}^{LB} = 1 \qquad \forall \sum_{r=1}^{R} m_{ir} U_{inr} > 0 \; and \; n \in \{1, \dots, |L_g|\}$$

Restriction 2: Central stock always available.

In extreme cases, demand may occur at only one or two local warehouses, according to the forecast. As a result, the model will obtain little risk pooling by storing items at the central warehouse and instead allocate all stock to the corresponding local warehouses. However, as we have a relatively short period of only 2 to 3 years for the forecast of demand for spare parts, more regions might have demand which has not yet been registered. In order to quickly deliver to these "no demand regions", we restrict the heuristic stock allocation to store at least one SKU centrally.

$$s_{i0}^{LB} = 1$$

5.3.4 Steps of column generation procedure

This section summarizes the steps which are performed by the column generator, which will generally obtain the optimal stock allocation in terms of reduced costs. These steps are mostly deduced from Alvarez et al. (2012a). At every step, we refer to which observation or restriction the step is deduced from. Recall that Appendix E describes the observations. We start with network 3, which will result in the quickest optimal stock allocation, as it has fewest local warehouses. Its corresponding reduced costs help to generate tighter stock levels bounds for network 4, which in turn will result in less calculation time.

1. Exclude all local warehouse with no demand $s^{LB}_{in} = s^{UB}_{in} = 0$.
2. Warehouses with demand $s^{LB}_{in} = 1$ (Restriction 1 & 2) n∊ {0,..., $|L_g|$}, n≠excluded.
3. Determine s^{UB}_{i0} by observation 1
4. Determine s^{UB}_{in} by observation 3 n∊ {1,..., $|L_g|$}, n≠excluded.
5. For each s_{i0} ∊ { s^{LB}_{i0},..., s^{UB}_{i0}} do
    a. Determine if the tighter lower bound on the reduced cost by observation 2 is lower than the best found so far RED*(i), if not considering next s_{i0} stock level.
    b. By observation 4, tighten upper bound of $s^{UB}_{in}(s_{i0})$ n∊ {1,..., $|L_g|$}, n≠excluded.
    c. Determine $s'_{in}(s_{i0})$ and $s''_{in}(s_{i0})$ n∊ {1,..., $|L_g|$}, n≠excluded, by empirical conclusion.
    d. Enumerate all possibilities for the local warehouse $s'_{in}(s_{i0}) \le s_{in} \le s''_{in}(s_{i0})$, n∊ {1,..., $|L_g|$}, n≠excluded, in order to find minimal RED*(i), reduced costs.
    e. If the solution is best so far, store solution and its reduced costs RED*(i). If further increase of s_{i0} cause the lower bound in observation 1 to exceed RED*(i) exit for loop.
6. Continue for network 4.

The actual algorithm is a little different, as step 5 requires extensive calculation time when s^{UB}_{i0} is large (> 15). In these cases, we choose 5 stock levels between s^{LB}_{i0} and s^{UB}_{i0} and check if this results in a delivery policy with negative reduced costs. If this is not the case, all s_{i0} stock levels will still be considered. This saves computation time, since in the first number of iteration steps, the shadow prices do not well reflect the optimal balance between costs and waiting time for a SKU. The added value of obtaining the optimal delivery policy with these shadow prices is limited, as it likely is not the optimal one in the final solution. Consequently, as long as we find a policy for each SKU which improves the LP-relaxation solution, it is fine, as it will improve the shadow prices.

From fractional LP-relaxation to an integer ILP solution

It is likely that the LP-relaxation solution is fractional; i.e., more than one policy is used for one SKU. Per SKU, we have numerous policies from both the initial set and those added during the column generation procedure. These numerous policies are no problem for the LP-relaxation, but might result in long solving time in the ILP-problem. Therefore, we delete all dominated policies before solving the ILP-problem, namely those policies with both higher costs and waiting times than at least one other policy in the set.

5.4 Low value SKUs

There are 4249 SKUs valued at less than € 100.-, (BIS, 2012). Optimization of all these items in the presented model would cause long computational time, and furthermore, the structure for these SKUs in the optimal solution is already clear in advance. Due to the low item value, a high fill rate for the cheap items at all the local warehouses (network 4) is relatively cheap to realize. This will prevent both transport costs and caused SVBM'.

Additionally, the presented model assumes a reorder quantity of one. This assumption does not hold for low value items, as it was based on the fact that holding cost for high value items are large, compared to the order costs. Recall that the standard EOQ formula balances order and holding costs. Obviously, the holding cost for cheap items are low, resulting in an EOQ>>1. Consequently, we first calculate the reorder quantity Q by the standard EOQ formula, and then the fill rate for reorder point s is determined by:

(11)

$$fill\ rate\ local\ warehouse\ n = \frac{1}{Q}\sum_{y=s+1}^{s+Q} F(y-1) = \beta_n$$

where F(y) is the Poisson distribution. This formula is based upon the fact that the inventory position follows a simple continuous time Markov chain, a birth and death process. As a result, the inventory position has a uniform distribution on the integers {s+1, s+2,…,s+Q}, (Axsäter, 2006).

In order to approach the reorder point of a low valued SKU, the resulting costs and waiting time of a reorder point are calculated per local warehouse by equations 5 and 6, assuming that the central warehouse reaches a fill rate of 100% for emergency shipments. Replenishment lead time, required to determine F(y), becomes $\beta_0 * t_n + (1-\beta_0) * (t_n + t_0)$, assuming a certain value for β_0, the fill rate of the central warehouse.

The resulting waiting time and costs can be plugged into equation 10, ignoring policy shadow price J_i. The shadow price of the SVBM constraint, A (see equation 2), is used from the last column generation step, representing the optimal balance in waiting time and costs. We apply the reorder point which results in the lowest reduced costs, and as this function is convex, the lowest reduced cost is quickly found.

Although a Poisson distribution is incorrect for a context of emergency shipments, we nonetheless apply it, for the following reason. In practice, the fill rate at the local warehouses will be higher, due to lost sales, which are assumed to be backordered in the Poisson distribution. Consequently, the performance at the local warehouses is actually underestimated with the Poisson distribution. Furthermore, we have to consider that the outcome difference for high fill rates are minimal between Erlang loss and the Poisson distribution.

6 Solution test

Chapter 6 describes the performance of the model as presented in chapter 5. Section 6.1 describes the verification of the model. Section 6.2 describes how we set the parameters values. Section 6.3 describes the validation of the model. Section 6.4 describes the sensitivity analysis. Section 6.6 draws the conclusions.

6.1 Verification

"Verification is concerned with determining whether the conceptual simulation model has been correctly translated into a computer 'program,'" (Law and Kelton, 2000). The authors describe several techniques to verify the program. This section describes how the model of chapter 5 is verified, applying these techniques.

The model uses different building blocks which transfer information to one another. The three basic building blocks are:

1. LP-model.
2. Column generator.
3. Evaluation of Stock deployment by method (Özkan et al., 2011).

These three blocks and the transfer of information between the building blocks could cause deviation from the paper model. We discuss the verification of these four sources below.

LP-model

Although we did not run the same LP problem in a different program and compare the results, we are convinced that no deviations are caused by the LP-model. The solving method used, CPLEX, is frequently used in different settings, so we have no reason to doubt this part of the LP-model. Problems could occur, however, if we constructed the wrong LP-problem. To prevent such an error, we have printed the LP-problem as solved by CPLEX multiple times, to verify if the LP-relaxation was correctly constructed, corresponding to the described ILP model in section 5.1.1. We observed no deviations (for an extensive discussion, see Appendix F).

Column generator

The column generator quickly rules out large numbers of possible stock deployment combinations, based on the observations and empirical conclusions mentioned in Appendix E and described in section 4.4. For three products in a supply chain with 7 local warehouses, we calculated every reasonable inventory combination. Reasonable combinations are defined as every combination possible within the bounds of all warehouse stock levels. The upper bounds for the local warehouse are given by observation 4 of Appendix E. The lower bound is 0 if no demand occurs, and 1 if demand does occur at a local warehouse. The bounds on the stock level for the central warehouse are given by observation 2. These boundaries allow a large number of combinations. Therefore we chose to limit the supply chain to 8 warehouses. For none of the three SKUs was a better solution found by this enumeration procedure as compared to the column generator when it considered all central warehouse stock levels. E.M. Alvarez reported that in some specific instances a better allocation could possibly be found by enumeration, when compared to the method described in section 4.4. However, even when a better solution was found, the difference was minimal. As noted then, the programmed column generator almost always finds the best new column, and when it does not, the deviation from the optimum is minor.

Evaluation by Özkan et al.

The functioning of the method of Özkan et al. (2011), is tested by running around five problem instances from their article, and comparing our results with their results. These turned out to be exactly equal, and the programmed evaluation of a stock deployment was determined to function as required.

When running the actual problem instance, the method of Özkan et al. did not always converge, and consequently, the program got stuck. We included an escape method and printed the stock allocations which caused non-convergence to a text file. Recall from Appendix E that the method of Özkan et al. is frequently used to establish bounds on stock levels. Based on the produced text file, we deduce that most of the freeze occurrences are happening during the procedures which aim to create these bounds. These procedures frequently require the evaluation of extreme allocations, e.g., high stock at one local warehouse, while other local warehouse stock levels are low. This inconvenience does not affect the quality of the final solution, but it does cause increased computation time, since looser bounds result in more stock allocations that require explicit evaluation. Certain evaluations of stock allocations which were required for the actual performance but not for the bounds creation, might also be neglected, this can obviously affect the solution quality. Regardless of this threat, the deviation from the optimum will be minimal, since an escape is registered for only five SKUs.

Comparison model Alvarez et al.

In order to test the functioning of the total model, we have made a comparison on a problem instance in order to see if both of the models yield the same solution. We expect only the LP solution to be the same, not necessarily the ILP solution, as we have different ways of establishing an initial policy set and a different column generator procedure. Appendix I describes this comparison in detail, and how we overcome the model differences and the outcomes on a tested problem instance for both models. From Appendix I follows that besides a negligible deviation, the models yield exactly the same solution. Therefore, we conclude that our model functions correctly for network 4. Note that we could only test one network, since the model of Alvarez et al. (2012a) uses only one network. However, as network 3 is highly comparable to network 4, the column generation will also function correctly for network 3. Since it functions correctly for network 4, it is highly likely that given correct values for waiting time and costs, it also functions correctly for the four networks combined.

Besides these efforts, we debugged almost every small part of the program. We tested the accuracy of the program's results by seeing how well its resultant values compared to outcomes derived by different means. Finally, numerous intermediate results of the algorithm were printed in various text files, and no deviations from the expectations were noted, other than in instances of inaccurate expectations.

To finalize the verification of the model, we verify the determination of costs and waiting time for all networks, given certain fill rates. The costs and waiting time determination in network 4 should also be verified, as we have included our costs and waiting time determination in the model of Alvarez et al. (2012a). Finally, we should verify the determination of both the risk pooling effect and the central warehouse fill rate in networks 1 and 2. See Appendix G, under heading *Determining cost and waiting time*. Based on the discussion above and Appendices G and I, we determine the programmed model to most likely coincide with the paper model.

6.2 Parameter value selection

Numerous parameter values must be chosen, such that the model approaches the situation of Liander. In a dedicated meeting with the managers of Supply Chain Management and Service & Planning and a supply chain consultant, the parameter values were established as follows

t_n Planned replenishment time in weeks of warehouse n∈{0,…,N}.

The actual applied inventory policy at the warehouses is (R,s,Q), which is different than a (s-1,s) inventory system. The difference between continuous review versus periodic review with replenishment opportunities ever R periods, we choose to incorporate in variable t_n. For simplicity, we assume Q =1 for every SKU under consideration. We deduce t_0 from the promised lead time of the supplier available per SKU in the BIS. These values reflect the lead time of the supplier with the longest lead time. Therefore, we neglect the review period and lead time delays in t_0 and t_n becomes 3 days for n∈{1,…,N}, which is the sum of the review period at the local warehouse + lead time from the central warehouse.

D_{nr} Delivery time in hours from warehouse n∈{-1,…,N} to region r.

There are basically five different situations: For delivery from the local warehouse to its own region, we choose 20 minutes. For delivery from a manned warehouse to a region of an unmanned warehouse, we choose 45 minutes. For a delivery from the central warehouse to a region, depending on distance, we choose a value between 1.5 and 2.5 hours. For delivery from the supplier when contracted for emergency shipments, we choose 24 hours, and delivery from the supplier when not contracted for emergency shipments, 72 hours. Note that from section 5.2.1, the supplier in the last case has a broad definition, which makes delivery within 72 hours possible.

CW, CS Cost of a shipment from respectively the central warehouse and supplier.

This cost parameter only reflects the additional costs made in comparison to an emergency delivery from a local warehouse and its replenishment, €45.-. These costs are a combination of courier costs, € 100, and the interruption of regular picking and planning processes, €25, at Logistics. Therefore, CW becomes € 80,- and CS is increased from € 125.- to € 205.- as a result of generally longer transport distances and the additional organizational efforts required.

m_r The expected weekly demand rate at region r.

Deduced from the emergency deliveries from the central warehouse to a region and replenishments from the central warehouse to the local warehouses.

h_n Holding costs at warehouse n per year n∈{-1,…,N}.

At every warehouse, we choose this figure to reflect 25% of the purchase value of products, reflecting capital, storage and risk costs.

U_{nrg} {0,1} 1 if region r is linked to local warehouse n else 0 in network g ∈{2,3,4}.

In networks 2 and 4, this is straightforward. In the case of network 3, it is derived from the division of the energy network into districts. These are available from Liander OSP-beleid (2012).

q_i Average number of caused SVBM when SKU i is missing for an hour.

As has become clear from section 3.1.1, there are basically two different consequences of waiting time: idle time of mechanics and downtime of connections. The expected number of affected connections in case of an SKU failure differs enormously. For example, a failure in an electricity meter affects a single household, while a high capacity transformer can affect 10,000 households. Therefore, we have asked different material specialist to categorize all SKUs above €100.- into four categories. Table 6.1 displays the number of connections down due to a failure that corresponds to a SKU category.

Table 6.1 Number of expected connections down corresponding to the criticality category in case of a failure.

Category	# connections down
1	10,000
2	1,000
3	50
4	1

We choose to rate idle time of a mechanic equal to downtime of one connection, category 4. By a percentage, urgent order due to a failure q_i can be estimated. We estimate it at 10%. Suppose a SKU is in category 3, then q_i becomes:

$$\frac{(0{,}1*1000+(1-0.1)*1)*60}{3{,}000{,}000} = 0.002018$$

Sometimes we have different category estimations for one SKU; in these cases we apply linear interpolation to find q_i.

DS relative demand size of Liander at the supplier &
X Fillrate of the supplier in network 1.

We estimate Liander's relative demand size for urgent orders at the supplier to be 25%, and the supplier should reach a fill rate of 98%. 25% is a conservative estimation since Liander's suppliers frequently deliver to numerous other network administrators in Europe, or even to other industries. Consequently, Liander's average relative demand size for urgent orders is less than 25% and consequently, the actual risk pooling effect will be higher. This is a deliberate decision, since the added value of outsourcing emergency shipments to a supplier should be high in order to make such a decision worthwhile.

6.3 Validation

"Validation is the *process* of determining whether a simulation model (as opposed to the computer program) is an accurate representation of the system, for the particular objectives of the study" (Law and Kelton, 2000). Law and Kelton describe several techniques for increasing the validity. This section describes which of these techniques we used and how we applied them.

6.3.1 Model determination

Throughout the project there was regular contact with subject matter experts. These included the central and local warehouse managers, energy network engineers, supply chain experts, inventory controllers, data analysts and management team. These contacts had four purposes in the beginning of the project:

1. Establishing objectives of the study
2. Establishing the performance measures
3. Making the right simplifying assumptions
4. Establishing the right sources of data.

We especially emphasized the last three purposes by arranging meetings with different stakeholders and performing a structured walk-through. In these meetings, we discussed the performance measures, the most important simplifying assumptions and the sources of important or controversial information. With respect to the last two issues, we discussed the difficulty of establishing correct demand forecasts per region and SKU, due to a lack of registration of demand at the local warehouses. Furthermore, we identified the subjectivity in establishing the correct criticality of waiting time per SKU (categorization of SKUs by material experts). Finally, we saw that customer demand size per order line is not always one, although for the modeled products it frequently holds.

6.3.2 Model outcomes

We discussed the model outcomes of small problem instances with the supply chain expert, M.X. de Vries, and concluded the model outcomes matched our expectations. This conclusion was confirmed when we changed the input data and compared the results of these tests with previous model runs. After this validation of small problem instances, we ran the actual problem with the real data. The discussion of these results are divided into three parts: a discussion of the results with respect to one local warehouse Nijmegen, the overall results, and a comparison between the performance of the current and the model allocation. The figures of the model outcomes are based on an mx value , restricting the number of SVBM' to 2.01 unless otherwise stated. Section 6.4 discusses why.

Local warehouse, Nijmegen

We invited one of the local warehouse managers, A.G.M. Marneef from local warehouse in Nijmegen, to compare the model inventory parameters with the current inventory parameters. Table 6.2 displays the inventory value based on the current order parameters versus column generation model order parameters.

Table 6.2 Comparison inventory value Nijmegen between current allocation versus model allocation, with mx 2.01.

	# SKUs	Value	Model value	Tools
Current Nijmegen	71	€ 78,916.84	€ 24,682.22	€ 0.-
Model Nijmegen	229	€ 128,590.86	-	€ 43,534.10(96 SKUs)

Source: Logistics order parameters Nijmegen, BIS (2013).

We conclude the following from Table 6.2 and the comparison of the order parameters made with the local warehouse manager:

- The number of different SKUs stored at Nijmegen is increased from 71 to 229.
- Only 2 of the current 71 SKUs were removed from the assortment at Nijmegen
- With respect to these 71 SKUs, the model allocated less than one third of the current value, € 24,682.22 versus € 78,916.84.
- Although the model almost triples the assortment width at Nijmegen, the inventory value only increases by 64%. Consequently the average value stored per SKU is lowered.
- The SKUs parameters correspond well with A.G.M. Marneef's perception about the different slow and fast moving SKUs.
- In the model, tools constitute a large part of both the model inventory value at Nijmegen, 34%, and SKU width, 42%.
- Some SKUs have a reorder point of one, while they are mostly demanded in quantities of two. Unfortunately this data per SKU is unavailable.

Overall model outcomes

Table 6.3 displays the range of the different fill rates pursued per network and per supply chain level.

Table 6.3 Fill rates per network and supply chain level, with mx 2.01.

	Min	Max	Average
n2, CW	63.97%	99.99%	95.27%
n3, CW,	39.27%	99.89%	91.25%
n3, MW	90.27%	100.00%	99.89%
n4, CW	42.34%	99.97%	94.74%
n4, UW/MW	85.03%	100.00%	99.16%

As Table 6.3 displays, the ranges of the fill rates at the central warehouse in the different networks vary a lot more than the currently pursued range of fill rates at the central warehouse, 90%-98% (see Table 3.1). For example, in network 2 we observe a minimum fill rate of 63.97% and a maximum fill rate of 99.99% at the central warehouse. The increased fill rate range at the central warehouse makes sense, because besides demand and price per SKU, the criticality per SKU also reinforces increased fill rate differentiation. In addition to these three factors, the risk pooling effect can cause high fill rates in networks 3 and 4. By pursuing a high fill rate centrally, a high local fill rate can be achieved with low stock levels. Note that in this case, the replenishment from the central warehouse is seldom delayed. Moreover, if a stock outage occurs locally in the event of an urgent order, most likely the central warehouse can deliver it with a emergency shipment. At the local warehouses, high fill rates are especially observed for highly critical SKUs.

Table 6.4 displays the averages of the SKU characteristics of criticality, price and demand, per network.

Table 6.4 Per network, the averages over the SKUs characteristics: criticality, demand, and price, with mx 2.01.

SKU characteristic: Network	Criticality (SVBM'/hour)	Price (€)	Demand (#/week)
1 (34 SKUs)	0.00006	€ 4,214	0.01
2 (134 SKUs)	0.00187	€ 2,257	0.05
3 (261 SKUs)	0.00386	€ 636	0.19
4 (348 SKUs)	0.01086	€ 485	0.58

In Table 6.4, we observe an uninterrupted increase in the averages of criticality and demand when the network number increases. The same holds for price in the opposite direction. This is in line with our expectations as higher demand and criticality drive the products to a local storage location, corresponding to a higher numbered network, as this prevents SVBM' and transport costs. Conversely, a higher price drives the products to a central storage location, corresponding to a low numbered network, as this prevents high holding costs.

When observing the different order parameters at the warehouse we observed numerous occurrences of zero stock at local warehouses when the model applied local network 3 or 4 for a SKU, as a result of no demand. Due to a relatively short demand forecast measurement (two to three years), a demand for a specific SKU might not have occurred, while in reality there might be demand. This becomes especially problematic when only one region has a forecasted demand for a SKU. In these cases, the model will only allocate demand at the corresponding local warehouse, since no risk pooling effect can be achieved by allocation stock centrally. When another region releases an urgent order for this SKU, an undesirable emergency shipment from the supplier is required. We solved this by restricting the model in case of a SKU allocation in network 3 or 4, to have at least one item centrally stored, as previously discussed in section 5.3.3.

Furthermore, sometimes SKUs which cannot transported by van due to their weight and/or volume, were stored locally. This allocation is acceptable, since in such cases Liander might arrange a special transport contract with a nearby transport company.

Comparison of current and model performance

As has become clear from section 3.3.2, the urgent order size is sometimes larger than a single item of a SKU. Although orders greater than one occur less frequently for the modeled SKUs, in reality these orders still constitute a significant part. For this reason, an exact comparison between the model outcomes and the current performance is difficult. Note that when the order size becomes greater than one, lower transport costs and less SVBM' are caused. Therefore we have also converted the performance values in section 3.3.2 to the item level. Even this is not a completely fair comparison, since when the urgent order size is *always* one, less inventory value is required in comparison to a situation where the order size is *sometimes* greater than one.

We choose to use the performance outcomes when supplier shipments occur. In practice it does hardly happen, due to the interchangeability of items, lateral resupply, and backordering. Furthermore, we assume 90% of the potential improvement between the corrected and uncorrected SVBM, and the same for transport costs, can be reached by central coordination. Recall from section 3.3.2 that the corrected measurement, in contrast to the uncorrected measurement, excludes all emergency SKU deliveries from the central warehouse, which were actually locally stored. Table 3.8 displays the current performance based upon the mentioned choices. The current performance is plotted in Figure 6.1, which displays the performance of the column generation model for different values of SVBM'.

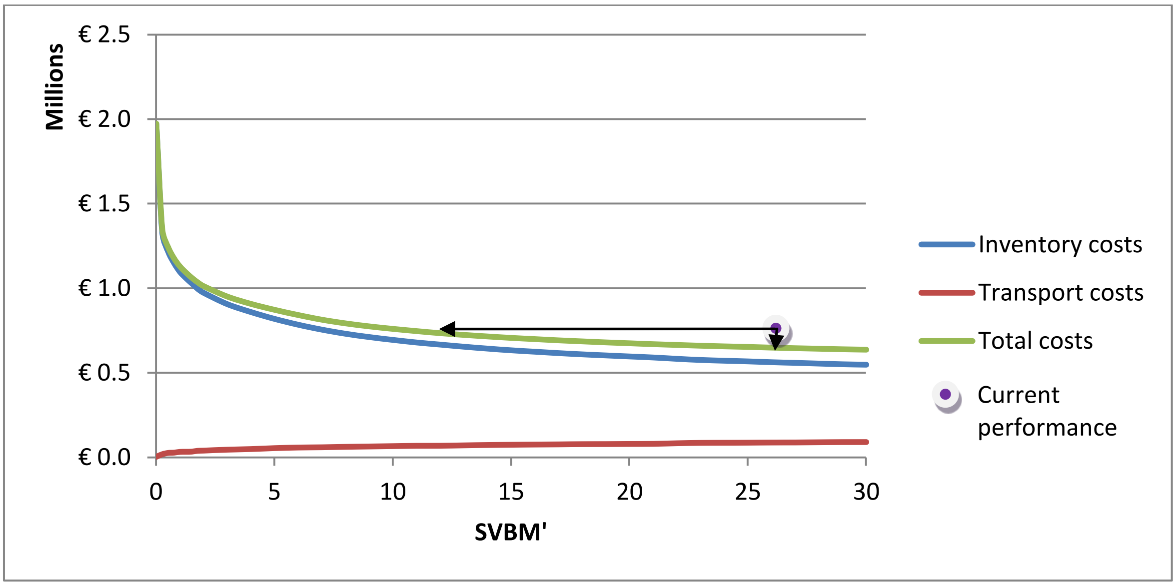

Figure 6.1 Comparison between current and model performance, ratio failure 10%

From Figure 6.1 follows that for the same costs, an improvement in caused SVBM' of 15 is possible, translating to a decrease in annual costs of around € 80,000.-. This improvement is lower than expected, taking into account the improvement potential in respect of low valued SKUs, which section 6.5 discusses. The relatively good performance of the current allocation is due to the fact that certain less severe urgent orders are backordered at the local warehouses, e.g., tools. Furthermore, the interchangeability of spare parts at the local warehouses could also play a role (A.G.M. Marneef) and lateral resupply between local warehouse (R. Hermans). Consequently, the efficient frontier could be lowered by incorporating these options in the column generation model and thereby further increasing the improvement. In the modeling, we have ignored the fact that many SKUs are present in the central warehouse for regular deliveries. By smartly integrating the urgent order and regular delivery stock, additional risk pooling is possible, which will further improve the model performance.

Table 6.5 displays a detailed comparison of the different costs items between the current and model performance. The column generation model is run by restricting the SVBM' to the exact one measured in the current situation, 27.94.

Table 6.5 Comparison in costs between the current and model allocation, mx=27.94.

Item:	Current	Model
CW inventory	€ 370,910.-	€ 315,817.-
UW+MW inventory	€ 246,785.-	€ 201,421.-
Supplier inventory	€ 0.-	€36,986.-
Transport	€ 104,901.-	€ 88,302.-
Total	€ 722,596.-	€ 642,526.-

We do not observe great deviations in any cost item between the current and model performance. All items are slightly lower by the model, and only supplier costs are different, since we have excluded that value in the current performance measurement.

We conclude that the output of the model meets our expectations, although for some items the assumption of one SKU per order line is not valid. Furthermore, we can doubt the criticality data gathered from different material experts with respect to tools. Are they truly so important that they should be stored locally and shouldn't the holding costs be increased for these SKUs, due to theft risks?

In the future, the actual demand will be measured at the local warehouses. When the stock outage moments are also included in this measurement, a comparison can be made with the forecasted fill rate of the model and the actual fill rate. Furthermore, the customer order size should be registered in order to determine this value per SKU and be able to incorporate it in the model.

6.4 Sensitivity analysis

Before elaborating on the sensitivity analysis, two different performance measurements of the column generation are discussed: the integrality gap and computational time.

Alternative performance measurements

The integrality gap defines how much the results of the ILP solution deviate from the LP-relaxation solution (Manthey, 2011). We observe a very small gap of 0.0027%, for mx=2.01 and ratio failure 0.10, which means the ILP solution is very close to the optimum, as the solution to the LP-relaxation is a lower bound for the optimal solution.

Computational time of the column generation model is around 5 minutes, optimizing 777 SKUs. We observed that the computational time is longer for low values of allowed SVBM’ in comparison to high values of allowed SVBM’. We never observed computational times above 7 minutes.

Varying allowed SVBM’ and ratio failure

The parameter limiting the maximum number of SVBM’ is not discussed in section 6.2. This is one of the most important parameters, as it heavily influences the outcome of the model. Liander does not have a clear idea of what this value should be in order to have the desired balance between costs and SVBM’. Moreover, since the SVBM’ calculated by the model does not coincide with the SVBM caused in reality. Therefore, we determine different model outcomes for different maximum SVBM’, corresponding to the value mx in equation 2.

Also, Liander is not sure what the actual balance is between an urgent order due to a failure and an urgent order for different reasons. During an expert meeting this value was estimated to be 10%, displayed by Figure 6.2. Figure 6.3 and Figure 6.4 display the results for possibilities of 5% and 20%.

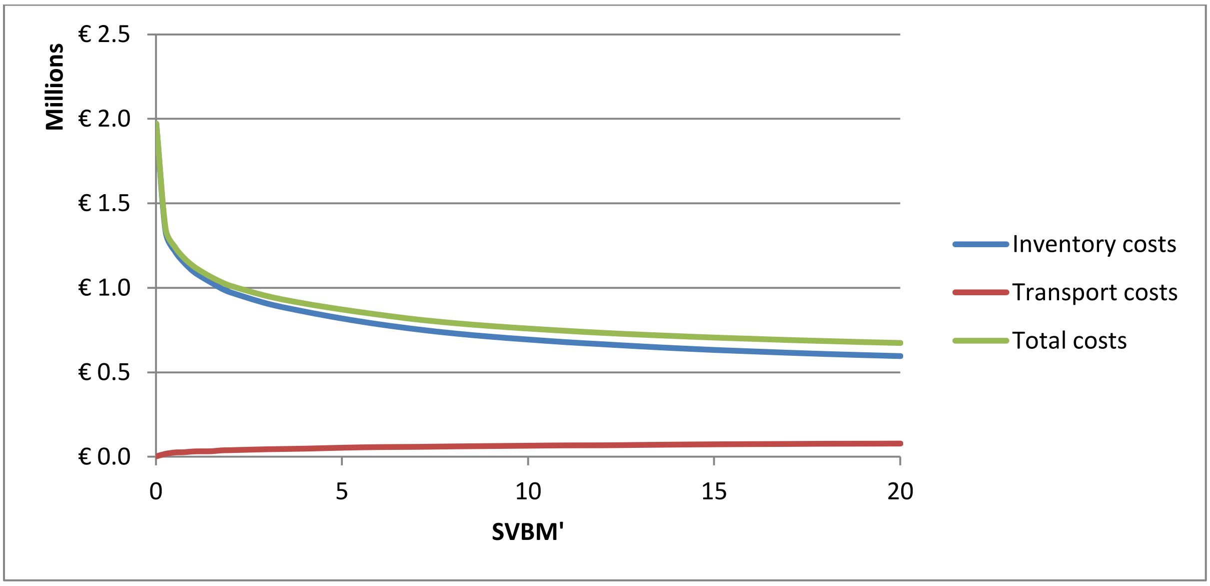

Figure 6.2 Relation between transport and inventory costs and maximum allowed SVBM, ratio failure 10%.

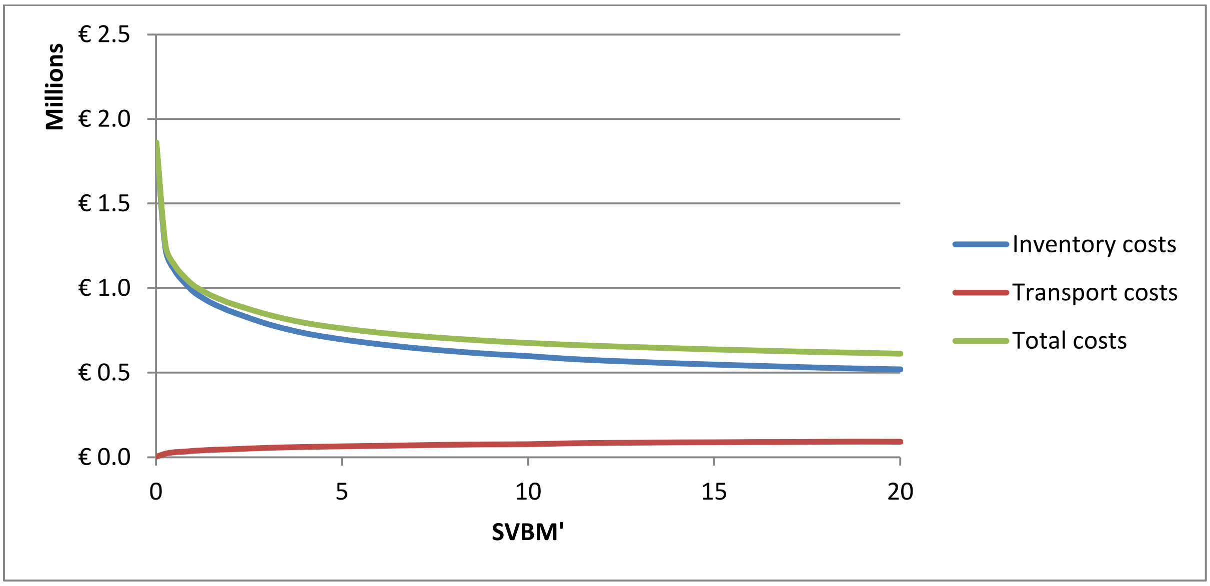

Figure 6.3 Relation between transport and inventory costs and maximum allowed SVBM, ratio failure 5%.

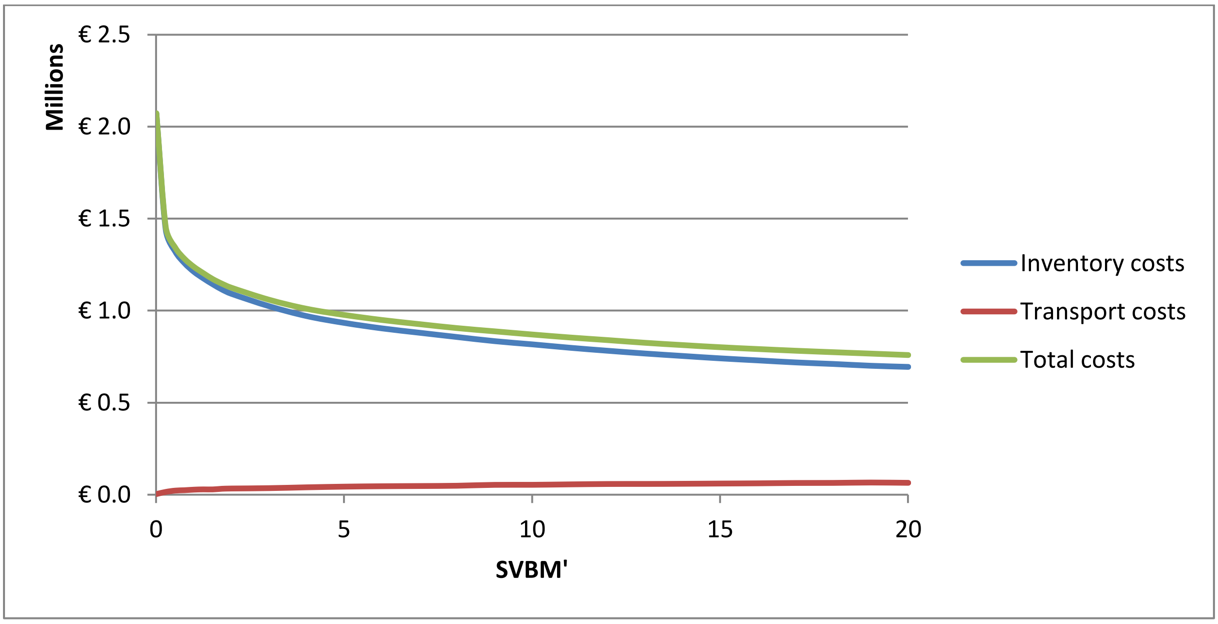

Figure 6.4 Relation between transport and inventory costs and maximum allowed SVBM, ratio failure 20%.

Although Figures 6.1, 6.2 and 6.3 do not display the results for large values of SVBM´, these results indicate that regardless of the ratio, the solutions converge to the same solution. This is logical, since as the allowed SVBM' increases in size, it no longer influences the solution. For very large allowed SVBM', the solution becomes the optimal balance between inventory and holding costs. Note that the ratio per SKU only influences how much SVBM is caused by one hour of waiting.

With supply chain expert M.X. de Vries, we concluded from Figures 6.1, 6.2 and 6.3 that changes in the ratio only causes the graph line to move slightly higher or lower compared to other ratio values. Therefore, we simply choose the initial failure ratio of 10%. Furthermore, the optimal value for maximum caused SVBM' is set at 2.01, since from that value on, the costs are not very sensitive to the maximum caused SVBM' value. The outcomes for this value will be used in the regression analysis, which is described in sections 7.2 and 7.3. Appendix J describes the column generation model outcomes in terms of network choice. Table 6.6 displays the allocation results for mx is 2.01.

Table 6.6 Comparison of current and model allocation, mx=2.01.

Item:	Current	Model
CW inventory	€ 370,910.-	€ 399,918.00
UW+MW inventory	€ 246,785.-	€ 547,758.00
Supplier inventory	€ 0.-	€ 24,811.00
Transport	€ 104,901.-	€ 39,354.00
Total	€ 722,596.-	€ 1,011,841
SVBM’	27.94	2.01

The allocation in Table 6.6 increases the current costs by €290,000.-, an increase of 40%, while lowering the SVBM’ by nearly 26, a reduction of 93%.

6.5 Low valued SKUs

Table 6.7 displays the model results for the SKUs for which we could establish the required information: lead time from supplier, current order parameters in Nijmegen, demand and price. Recall that low valued SKUs have an item price of less than € 100.-. These results are based on data from the local warehouse of Nijmegen and extrapolated to all local warehouses, based on turnover.

Table 6.7 Comparison of current and model performance of inventory costs and value for Nijmegen and all local warehouses

Inventory	Current situation	Model performance	Improvement
value Nijmegen	€ 89,017.-	€ 34,811.-	€ 54,206.-
cost Nijmegen	€ 22,245.-	€ 8,703.-	€ 13,542.-
value All local warehouses	€ 1,664,618.-*	€ 650,966.-*	€ 1,013,652.-
cost All local warehouses	€ 415,982.-*	€ 162,746.-*	€ 253,235.-

Source: Assortment Nijmegen in BIS and model outcomes, * by extrapolation

From Table 6.7 follows that a huge improvement is possible for cheap items by optimizing the order parameters, namely an improvement of € 253,234.- annually. However, we should note that certain SKUs are ordered by two at a time. For these items, a higher reorder point could be needed to reach the same performance. Unfortunately this information is unavailable per SKU. In the comparison we ignore transport costs, € 804.-, and caused SVBM’, 0.00186, as these are very low in the model. Due to the high current stock, these values will also be low in the current situation.

6.6 Conclusions

In chapter 6 we answered the fourth research question: *What is the performance of the methodology under different system characteristics?* Before answering this question we verified and validated the developed column generation model. From the validation we conclude:

- The validity of the model can be increased by:
    - Improving the measurement of SKU criticality, especially tools.
    - Incorporating the occurrence of order lines with size greater than one.
    - Implementing measurements of the fill rate at the local warehouses and comparing these with the forecasted fill rate by the model.

- The performance of the column generation model can further be improved by incorporating:
  - Lateral resupply between local warehouses.
  - Interchangeability of SKUs at all warehouses.
  - The option of backordering certain SKUs.
  - Integrating stock of regular deliveries and urgent orders at the central warehouse.

From the comparison of the current and model performance, we conclude that Liander can annually save €80.000, a reduction of 11%, on SKUs valued at less than €100.-, or choose to reduce the caused SVBM' value by 15, a reduction of 55%. From the sensitivity analysis of costs versus caused SVBM', we choose to restrict the column generation model to a max caused SVBM' value of 2.01. This corresponds to a cost increase of 40%, while lowering the SVBM' by 93%.

With respect to the inexpensive SKUs, < €100.-, Liander can annually save € 253,235.-.

7 Simple allocation rule

Chapter 7 describes how the obtained results of the column generation model can be duplicated by a simple allocation rule. Section 7.1 investigates the possibilities of regression for network choice. Section 7.2 describes a simple allocation rule based on ordered logistic regression. Section 7.3 describes how to obtain a stock allocation in a network. Section 7.4 compares the performance of the allocation obtained by the column generation model with the simple allocation rule. Section 7.5 draws the conclusions.

7.1 Graphical representation network selection

Liander is primarily interested in a simple rule which specifies which network to choose for a SKU. Many parameter values have influence, but it is impossible to involve all in a simple allocation rule. Therefore, we propose to take only a few product-dependent characteristics into account.

We identify three product parameters in the discussion in Appendix G which mainly influence the network decision: demand size, criticality and holding costs. The higher the first two factors are, the higher the network number will be. The higher h is, the lower the network number will be. Figure 7.1 displays the decision of the model with respect to network selection plotted against item price and demand*criticality. Note that the axes of the graph in Figure 7.1 are shown on a logarithmic scale.

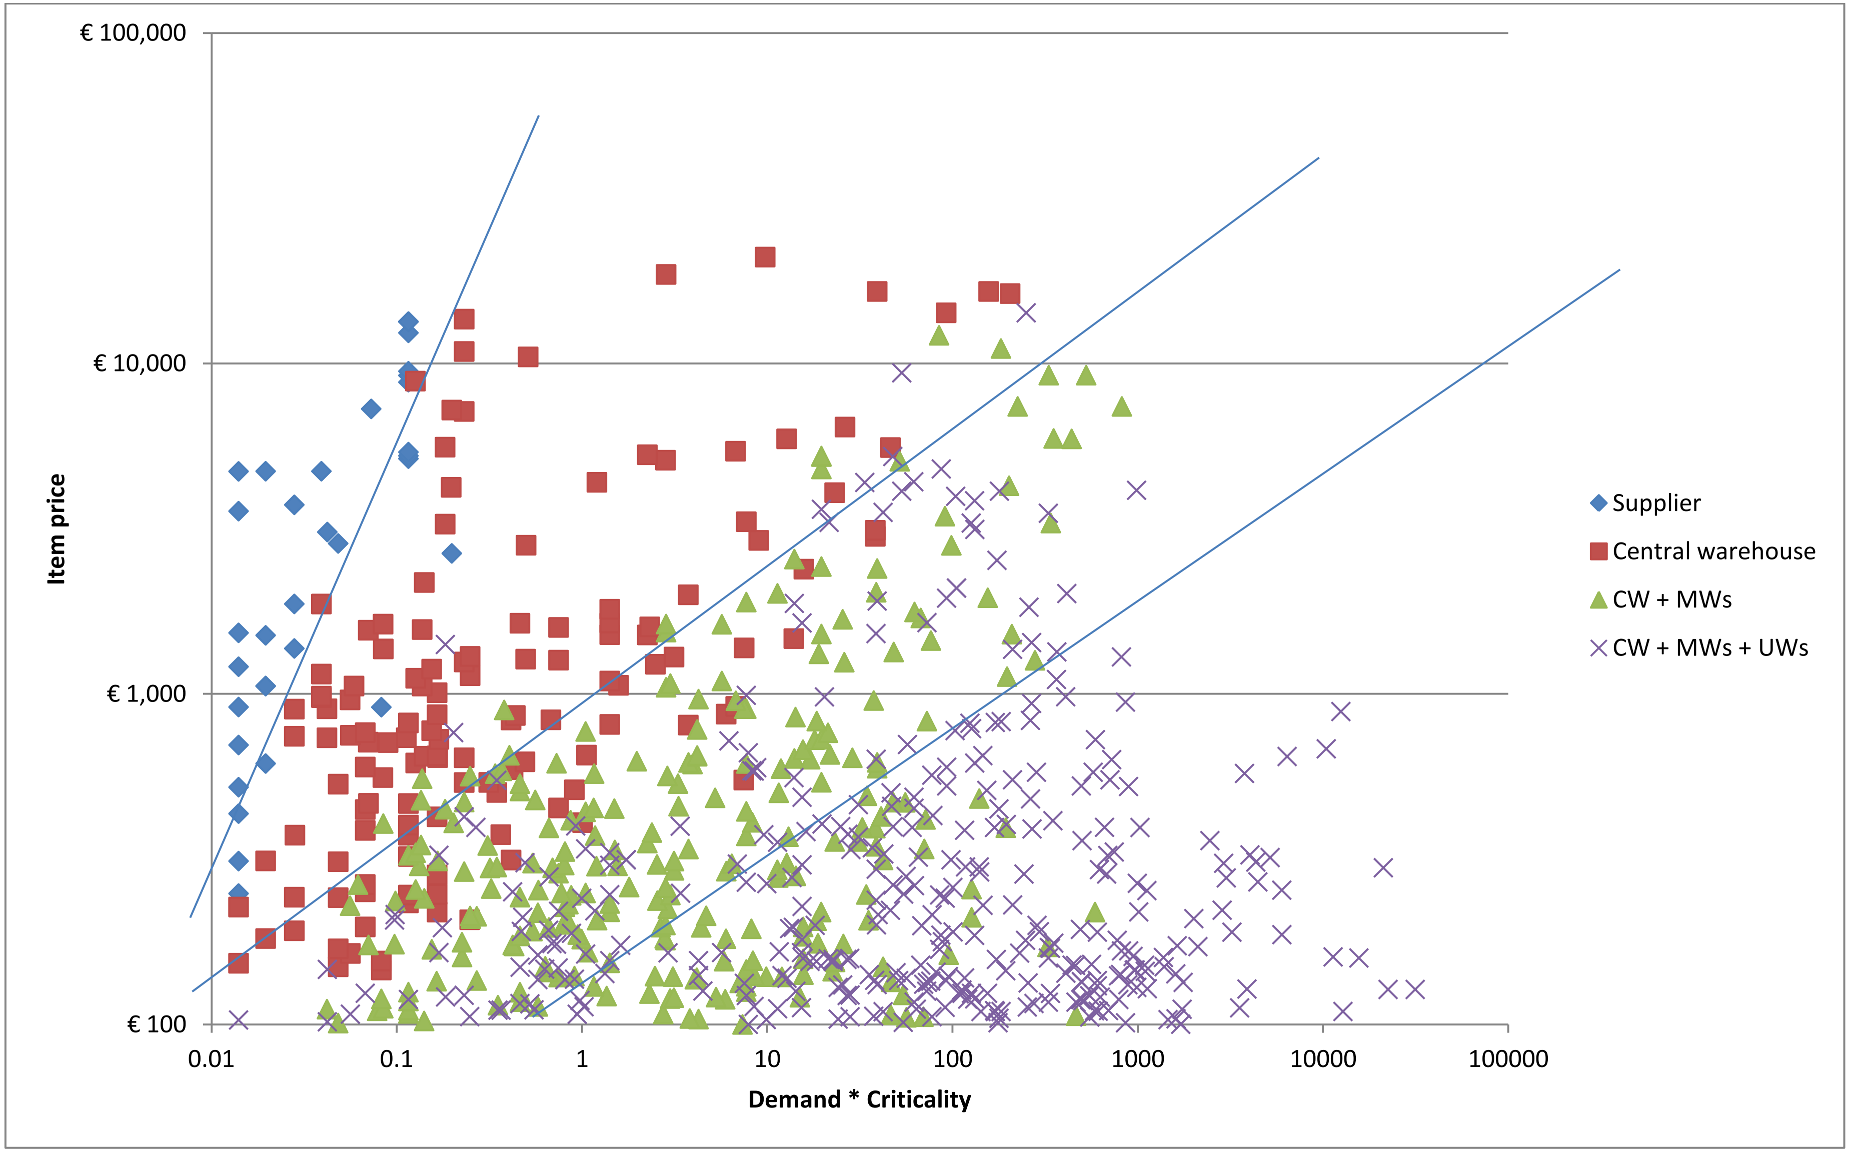

Figure 7.1 Network selection by SKU price and demand*criticality. These values are based on a mx of 2.01 and a failure ratio of 0.1.

From Figure 7.1 we conclude that based on the chosen SKU characteristics, likely a good approximation can be made which network is optimal, although a lot of overlap can occur for networks 3 and 4. This might be due to that in certain demand situations, the choice between networks 3 and 4 makes almost no difference. Suppose that only regions corresponding to manned local warehouses demand a certain SKU. The solution to networks 3 and 4 will be exactly the same. Additionally, when relatively little demand occur from regions of unmanned local warehouses, the solution to networks 3 and 4 will still differ minimally.

7.2 Ordinal logistic model: network selection

The intuitive selection of the parameters, or predictors, in section 7.1 can be improved by the ordinal regression as described in section 4.5.1. The output generated by SPSS returns which predictors significantly influence the outcome. Furthermore, we can include interactions terms of predictors. For example, we expect that interaction between demand and criticality significantly influences the network selection, as these two terms combined influence the SVBM constraint (equation 2).

From including various different predictors per SKU, such as lead time, sum of demand at manned local warehouses and sum of demand at unmanned local warehouses, we identify four significant predictors: total demand, criticality, SKU price and supplier's lead time. Although lead time has the poorest significance level, when we simply compare the forecast outcome with the actual decision of the column model, the results of the forecast *with* lead time is poorer (56.4% correct) than *without* lead time (56.6% correct). Consequently we choose to exclude lead time, since in practice it requires additional input for Liander to determine the network number, while the benefit is unclear. In contrast to our expectations, no interaction terms were significant.

Table 7.1 displays the estimates of the significant parameters. Appendix H displays the total regression analysis output.

Table 7.1 Parameter estimates of ordinal logistic model based on a model run with ratio failure is 0.1 and mx is 2.1 with corresponding symbols and odds changes.

		estimate	symbol	Δodds
Threshold	Network 1	-3.252	ω_1	-
	Network 2	-.734	ω_2	-
	Network 3	1.471	ω_3	-
Location	Price/1000	-0.410	φ_1	0.66
	Demand/week	3.186	φ_2	24.19
	Criticality	178.684	φ_3	3.99E+77

Source: Output regression analysis SPSS

As section 4.5.1 discusses from the predictors displayed in Table 7.1. The probability that a network is optimal can be deduced from:

$$P(network \leq j) = \frac{1}{1 + e^{-(\omega_j - \varphi_1 X_1 + \varphi_2 X_2 + \ldots + \varphi_n X_n)}}$$

From Table 7.1 follows that one unit of increase of the SKU price, corresponding to an increase in price by € 1000,-, decreases the odds of entering a higher numbered network by 0.66. While an increase of one unit of demand/week increases the odds of entering a higher numbered network with 24.19. The change in odds for criticality is very large. Note that in practice these values and their corresponding variation is small.

Performance of forecast

Table 7.2 displays the absolute deviation between the forecast and column generator result with respect to the network decision, recall these are numbered from one to four so that a maximum deviation of three is possible. The network decision of the forecast is simply based on the highest forecasted network probability.

Table 7.2 Absolute deviation between forecast and column generator result with respect to the network decision

Absolute Deviation	Occurrence	%
0	440	56,6%
1	322	41,4%
2	15	1,9%
3	0	0,0%

From Table 7.2 we conclude the regression results never propose the total opposite of the column generation results, as a deviation of three never occurs. A deviation of two is also seldom. The correct forecast (or a deviation of one) occurs for 98% of the SKUs. If the forecast is simply a choice between local storage, corresponding to networks 3 and 4, or central storage, corresponding to networks 1 and 2, the regression results forecasts 83.7% correctly.

Verification of regression outputs

SPSS returns on command various test results of the regression analysis, which are displayed in Appendix H. The ordinal logistic model assumes that the influence of the independent variables is identical for the different logit functions. This assumption is verified by the test of parallel lines, but unfortunately the model is unable to test it with the presented predictors. However, we did run many regression analyses, and did not observe a problem in any of them, since the observed significance levels were always above 0.05. This indicates the assumption of parallel lines holds.

Appendix H displays the table Model Fitting Information. This tests the null hypothesis that all of the regression coefficients are equal to zero. That is, it tests if there is no effect of the predictor variables. From these results we can reject the null hypothesis.

As section 4.5.1 discusses, the other output of the regression output in Appendix H is inapplicable, due to the use of continuous predictors.

In order to verify the probability calculations for the different networks, we calculated by hand the different probabilities based on the SKU characteristics and the estimates of φ's and ω's from SPSS. We compared these values with the probability calculations from SPSS and identified the results to be exactly the same. Based on the discussion above and Appendix H, we conclude the forecast model to coincide with the paper model and it is applicable to support Liander in a network decision.

7.3 Inventory level

Besides a decision regarding which network the SKU should be placed in, a reorder point also must be chosen such that it approaches the warehouse fill rates chosen by the model. In section 6.3.2

Table 6.3 displays the average fill rates within the different networks at the different supply chain levels. Furthermore the ranges of these fill rates are displayed over the different SKUs.

Especially for network 2, we observed significant differences in pursued fill rate at the delivery warehouse, from 0.64 to 1.00. Besides a selection of the stock level at the local warehouses, a selection of stock for the central warehouse is also required in networks 3 and 4. Again, we observe great variation in fill rate for networks 3 and 4 at the central warehouse, from 0.39 to 1.00 and from 0.42 to 1.00, respectively.

When we simply apply the average fill rates displayed in Table 6.3, it results, as expected, in very poor outcomes. The application of these percentages is done by choosing s such that it minimally results in the displayed fractions. The fill rates at the central warehouse in network 2 and at the local warehouses in network 3 and 4 are approached by the Erlang loss formula, while the fill rate at the central warehouse in network 3 and 4 are approached by the Poisson distribution. It results in the tripling of the caused SVBM´, from 2.01 to 6.57, while the costs also increase from €1,011,841.- to €1,154,762.-

An improvement could be made possible by differentiating the various fractions based upon specific SKU characteristics. This results in five tables. The application of all these tables becomes cumbersome, which contradicts Liander's intention of a simple allocation rule.

Therefore, we changed our approach in an attempt to forecast the average waiting time per SKU. This forecast gives an indication which fill rates must be pursued at the different supply chain levels. This can easily be done, since in the network selection of various SKU characteristics are already used as input. In order to apply a linear regression, we transform the average waiting time for a SKU by the natural logarithm so that a linear regression model is possible:

$$\ln Y_i = \propto + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_k X_k + \epsilon_i$$

Note that in our case a lower value for average waiting than a local delivery in network 4 is impossible. Recall that in this case D_{nr}=0.33 Therefore, we subtract this value from the average waiting from the model before we apply linear regression. Otherwise, a lower value for $\ln Y_i$ of -1.11 would be impossible.

We choose to include exactly the same SKU characteristics which are used in the network selection regression model: Demand, price, criticality and Demand*criticality. Furthermore, we include the network number in the regression model and treat it as a continuous variable. Table 7.3 displays the log linear regression results of the waiting time in hours.

Table 7.3 Log linear regression parameter estimates in order to determine average waiting time in hours.

Parameter	Symbol	Coefficients	e^Coefficients
Intersection	$\propto$	3.52	33.84
Demand/week	β_1	-0.28	0.75
Price/1000	β_2	0.25	1.28
Criticality	β_3	-145.44	6.86E-64
# network	β_4	-1.56	0.21

Source: Output regression analysis Excel

One unit of increase in X_i results in a multiplication of Y_i, the expected average waiting time of SKU i, by e^{β_i}. From Table 7.3 we conclude that each unit of increase of demand/week, criticality and network # decreases the expected average waiting time, while only price increase the waiting time when the corresponding X value increases. The outcomes on the β_i's are in line with our expectations. A higher network number means a more local delivery, resulting is less waiting time. A SKU with higher demand will have more impact on transport costs and on the SVBM constraint. As a result, the column generator from transport costs perspective will aim for a high fill rate locally, as it prevents transport costs, which in turn results in less waiting time. Moreover, the column generator will also aim, for a low waiting time (from the perspective of the SVBM constraint, equation 2), as high demand SKU heavily influences the SVBM constraint. The same logic of the SVBM constraint can be applied on the criticality e^{β_i} value. Note that all the criticality values, q_i's, for all SKUs are very small, resulting in an extremely small value for corresponding e^{β_i}. As expected, a higher SKU price results in additional waiting time, as it is more expensive to reach a high fill rate, resulting in higher average waiting times.

7.4 Comparison simple allocation rule and column generation model

The obtained regression results are case specific, which means that as one of the parameter values changes in the column generation model, the parameters obtained from the regression model are no longer based on the optimal solution. It should be noted that there is variety in the sensitivity of the optimal solution as a result of changes in one of the column generation model parameters.

From section 7.3 follows that when the service part allocation is based on the optimal network choice in combination with the average fill rates per network and supply chain level, it results in a cost increase of 14 % and a SVBM' increase of 227%. It shows that the order parameters per SKU need to be smartly deduced from the waiting time forecast in order to prevent a drastic loss in performance. We expect this order parameter deduction per SKU will be a time-consuming job. Also hardly any less data per SKU is required for the simple allocation rule in comparison to the column generation model. We only save the gathering data with respect to lead time from the supplier and the demand per region. Note that total demand per SKU is still required.

Consequently, we expect hardly any labor time is saved by the application of the simple allocation rule in comparison to the column generation model, while at the same time the performance will certainly less for four reasons:

1. Deviation from optimum in regression results network choice.
2. Deviation from optimum in regression results average waiting time.
3. Deviation from optimum in transforming waiting time to stock allocation.
4. Over time, the regression results lose quality, as the optimal solution changes due to parameter changes

7.5 Conclusions

We conclude that it is possible to give a reasonable forecast of the optimal network for storing a SKU, based upon a number of its characteristics. We have obtained a verified model which can forecast the optimal network based on ordered logistic regression. It forecasts 56.6% of our problem instances (777 SKUs) correctly and for an additional 41.6%, it forecasts just one network rank above or below the optimal. Besides the forecast of a network, we need to forecast the stock levels in corresponding warehouses. Simply applying average fill rates per network and supply chain level results in suboptimal results. Therefore, we apply a log-linear model which forecasts the expected average waiting based upon a number of SKU characteristics. From these average waiting times, Liander can try to deduce an optimal stock allocation. We expect this deduction will be time-consuming. We predict that an overly simplified approach to this last step results in disappointing results.

Overall, we expect a loss in the performance of the simple allocation rule versus the column generation model due to:

1. Deviation from optimum in regression results network choice.
2. Deviation from optimum in regression results average waiting time.
3. Deviation from optimum in transforming waiting time to stock allocation.
4. Over time, the regression results lose quality as the optimal solution changes due to parameter changes.

8 Implementation

Section 8.1 describes the application of the different models. Section 8.2 discusses the application of the simple allocation rule. Section 8.3 draws the conclusions.

8.1 Application of the models

Liander can choose to directly use the order parameters returned by the models, column generation model and model for low valued SKUS. Besides this option, Liander can also make use of the simple allocation rule to replace the column generation model. Section 8.1.1 describes the interpretation of the model outcomes.

If Liander decides to apply the column generation model instead of the simple allocation rule, we will train a Liander employee to practically apply it. From section 6.5 follows that an annual reduction in costs of € 250,000.- can be easily reached for the inexpensive SKUs (<€ 100.-) by the corresponding model. Therefore, we develop an easy-to-use version of the Delphi program, and we will help Liander employees to practically apply it.

8.1.1 Interpretation model outcomes

Column generation model

The reorder points returned by the column generator and the simple model for the cheap items should be interpreted carefully. The column model returns s which is the order-up-to level in an (s-1,s) inventory system. Therefore, as the central warehouse uses a (R,s,S) system, in terms of column model output, it becomes: (R,S-1,S). The reorder point becomes one item less than the returned s of the column model. At the local warehouses an (s,Q) inventory system is used. In terms of output of the column model, it becomes: (s-1,1).

Low valued SKUs model

The reorder point, s, and order quantity, Q, returned by the model can easily be applied to the local warehouses' inventory systems. As discussed above, it is a (s,Q) inventory system, where the s and Q corresponds with the s and Q returned by the model.

8.2 Applying simple allocation rule

Note that the simple allocation rule is only developed for the expensive SKUs (>€ 100.-). It duplicates the results of the column generation model.

Network choice

From the ordinal logistic regression model we obtain four probabilities referring to the four different networks, based on the following SKU characteristics: demand, price, and criticality (see section 7.2). Liander can simply choose to apply the network which returns the highest probability. Liander can also carefully consider the returned probability and as network 2, for example, is the second best while network 3 is best, since Liander knows the demand is overestimated, it is smart to apply network 2, as less demand lowers the odds for a higher network number (see Table 7.1, ΔOdds column). Also, based on other characteristics not applied in the regression model, it can be smart to overrule the network with the highest probability if the difference from the best to the second or third best is relatively small.

Average waiting time

The results, which section 7.3 describe, of the log-linear regression forecast of waiting time is less intuitive to apply to networks 3 and 4. From these results we must obtain the optimal stock allocation levels of the corresponding warehouses. This is difficult, since in certain circumstances such as with many demand regions, it can be smart to allocate much stock centrally in order to harvest the risk pooling effect. In contrast, in other circumstances such as few demand regions, not much risk pooling effect can be reached by allocating stock centrally. Moreover, section 7.3 shows that a simple approach results in disappointing results for the expensive items, even when we obtain the optimal network for all SKUs. It corresponds to a costs increase of 15% while the SVBM' also increases by 226%.

8.3 Conclusions

The order parameters generated by the models should be interpreted with care, as these are different. Liander employees will be trained in the application of the model(s) if Liander decides to apply them. An overly simple transformation of the waiting time forecast, obtained by the simple allocation rule, to order parameters will result in a bad performance.

9 Conclusions and Recommendations

9.1 Conclusions

1) How does Liander currently handle urgent orders from a logistical perspective?

An urgent order is a result of an unforeseen material need of a mechanic, which cannot be delivered be a regular delivery. Currently, there is not a clear methodology of service parts allocation with respect to urgent orders. This problem is exacerbated by no central coordination, since both Netcare and Logistics control the stock allocation at different supply chain levels. The performance of the service part allocation for urgent orders is defined by costs, as a result of inventory value and emergency shipments, and the number of StoringsVerBruiksMinuten (SVBM) due to the waiting time for spare parts. Table 9.1 displays the current performance for SKUs valued above € 100.-.

Table 9.1 Estimated current performance for SKUs>€100.-

Performance indicator:	score
Caused SVBM'	27.94
Transport costs	€ 104,901.-
Inventory costs	€ 585,778.-
Total costs	€ 690,679.-

2) Which deployment methodology from the literature best fits Liander's performance ambition, vision and situation?

We obtained the column generation approach from the literature, which has been previously applied in the context of spare parts optimization with aggregate waiting time restrictions over all items. Although we do not have a waiting time restriction, it can easily be transformed to a maximum caused SVBM´ constraint by including a criticality measure per SKU. In order to determine the fill rates in a two-echelon network, with emergency shipments, at the local warehouses as a result of a stock allocation for a SKU, the method of Özkan et al. (2011) is used. In order to determine the optimal stock allocation in a two-echelon network, the observations and empirical conclusions from Alvarez et al. (2012a) are used.

3) How can the best deployment methodology be applied? (Ch.5)

The option of a delivery from a mechanic van is excluded from the column generation model. Therefore, we have four network options for urgent order delivery:

1. Supplier only - consignment stock
2. Central warehouse only.
3. Combination of central warehouse and MWs.
4. Combination of central warehouse, MWs and UWs.

Every network has an infinite delivery option, which consists of all possible suppliers for a SKU or Liander colleague companies. The column generation model is only applied on the expensive SKUs (>€ 100.-). For the inexpensive SKUs (<€100.-), a separate model is developed, assuming these items are by definition stored in network 4. In this model, the reorder quantity, in addition to the reorder point is optimized.

4) What is the performance of the methodology under different system characteristics?(Ch. 6)

Verification is performed by aligning the Alvarez model with our model and duplicating the results of the same problem instance. The model of Özkan et al. is validated by duplicating the results of given problem instances in their article. We varied the values of maximum caused SVBM' and the ratio between urgent orders as a result of item failure and other reasons. For the expensive SKUs (SKUs >€100) in terms of costs, an improvement of € 80,000 (11%) can be realized annually, or an annual reduction in SVBM' of 15 (55.6%). For Liander, an SVBM' value of 2.01 reflects the best balance between costs and SVBM'. This corresponds to a cost increase of 40%, while lowering the SVBM' by 93%. With respect to the low valued SKUs, a reduction in annual cost of € 253,235.- is possible.

5) How can the methodology be reduced to a rule of thumb?

The results of the column generation model are duplicated by two allocation rules. The first specifies which network to use, while from the second the order parameters can be deduced. The network allocation rule is based on ordered logistic regression analysis, and a forecast is made of the optimal network, based upon a number of SKU characteristics. In our problem instance, it forecasted 56.6% correct, and for an additional 41.6% it forecasted a network just one rank above or below the optimal. The second allocation rule does not directly forecast the order parameters, but forecasts the average waiting time for a SKU. From this average waiting time, a corresponding stock allocation can be found. We expect this deduction will be time-consuming and that an overly simplified approach to this last step results in disappointing results.

9.2 Recommendations

From our conducted research we make the following recommendations:

1) Incorporate lateral resupply between local warehouses and the interchangeability of SKUs in the allocation models. Lateral resupply between local warehouses creates additional risk pooling effect, which enables the creation of a stock allocation which causes the same SVBM' at less cost. Also, the interchangeability of SKUs allows inventory cost reduction while providing the same service level.

2) Integrate the stock of urgent orders with the stock for regular deliveries at the central warehouse. In networks 2, 3 and 4, SKUs are set aside at the central warehouse specifically for urgent order fulfilment, while these same SKUs are already stocked centrally for regular deliveries. By integrating these stocks, additional risk pooling effect is possible.

3) Improve criticality data of SKUs. Some SKUs that are classified as highly critical are backordered at the local warehouse. Apparently these SKUs are not as critical as reported. By an improved measurement, the service part allocation can be further aligned with Liander's need for an effective urgent order fulfilment.

4) Incorporate the occurrence of order line sizes greater than one. For some SKUs, especially the inexpensive SKUs (< € 100.-), the customer order line size is not always one, which contradicts the assumption of both allocation models. When the customer order line size is stable but not equal to one, the unit can simply be adapted. This is applicable in both models.

5) Measure demand by date, order line size and fill rate at the local warehouses. The demand measurement allows for an improved demand forecast in the future, while also making the application of the former recommendation easier. Measurement of the fill rate will boost trust between Netcare and Logistics, preventing hidden stock locations. Furthermore, it enables additional validation of the allocation models.

6) Use the column generator model instead of the simple allocation rule. A stock allocation of the simple allocation rule will deviate from an allocation of the column generator model for several reasons. Furthermore, over time as parameters change, the results of the simple allocation rule will deviate from the results of the model. Therefore, as Liander deems the simplicity of the simple allocation rule to outweigh the loss in performance, the simple allocation rule requires an annual update by the column generation model.

9.3 Research limitations

We ignored the effect of stock in the mechanic vans in all models. In the column generation model, we assumed replenishment order quantity at all warehouses to be one. In practice, this will not hold for relatively inexpensive SKUs (> € 100.-).-, which are fast-moving. We also assumed a customer order size of one in both models, which will not always hold in practice.

Additionally, in practice, some urgent orders are backordered, fulfilled by lateral transhipment or fulfilled by interchangeability between SKUs. These options are excluded from the column generation model.

9.4 Further research

Liander can likely further improve the stock allocation of the models by:

- Integration of mechanic vans' inventories in the allocation models.
- Integration of inventory for regular deliveries and urgent orders.
- Utilizing the interchangeability of certain SKUs.
- Application of a critical policy level in the local warehouses for urgent orders as a result of item failure.
- Include the option of lateral transhipment between the local warehouses.
- Improvement of waiting time criticality per SKU:
    - The probabilities that the whole energy distribution network can overtake its function vary from SKU to SKU.
    - Include differences of downtime between gas and electricity outages.
    - Criticality data with respect to tools seems to be overestimated.

The validity of the allocation models can be improved by:

- Incorporate that customer order size is not always one.
- Replenishment order size is not always one; this holds only for the column generation model.
- Registration of urgent order demand at the local warehouses enables a comparison between the expected and registered fill rate.

10 Abbreviations

SVBM	Corrected downtime of electricity network (StoringsVerBruiksMinuten)
SVBM'	Alternative measure for 'StoringsVerBruiksMinuten' used in the modelling and performance measurement
CW	Central warehouse
DOL	Delivery on location
MW	Manned local warehouse(also referred to by OSP+, Operationeel steunpunt)
UW	Unmanned local warehouse(also referred to by OSP, Operationeel steunpunt)

11 References

- Alfredsson, P., and Verrijdt, J. (1999), Modeling emergency supply flexibility in a two echelon inventory system. Management Science, 45:1416–1431.
- Alvarez, E.M., Van der Heijden, M.C., van Zijm, W.H.M. (2012a), Service differentiation in spare parts supply through dedicated stocks, Beta Working Paper series 373.
- Alvarez, E.M., Van der Heijden, M.C., van Zijm, W.H.M. (2012b), The selective use of emergency shipments for service contracts differentiation.
- Axsäter, S. (2006), Inventory control, Lund(Sweden), second edition, Springer science + Business Media, LLC.
- Benoit, K., (2011). Linear Regression Models with Logarithmic transformations, London School of Economics.
- Kranenburg, A.A., and Van Houtum, G.J. (2007), Effect of commonality on spare parts provisioning costs for capital goods. International Journal of Production Economics, 108, 221-227.
- Kranenburg, A.A., and Van Houtum, G.J. (2008), Service differentiation in spare parts inventory management. Journal of the Operational Research Society, 59, 946-955.
- Law, A.M., Kelton, W.D.(2000). *Simulation Modeling & Analysis*, third edition, McGraw-Hill
- Manthey, B. (2011), Optimization modeling. Reader, University of Twente.
- Muckstadt, J. A., and Thomas, L. J., (1980), Are multi-echelon inventory methods worth implementing in systems with low-demand-rate-items? Management Science, 26:483–494.
- O'Halloran, S. (2005). Econometrics for Sustainable Development II, sheets, Columbia University.
- OSP-beleid Liander 2.0, (2012)
- Özkan, E., van Houtum, G.J., and Serin, Y. (2011), A new Approximate Evaluation Method for Two-Echelon Inventory systems with Emvergency Shipments, Beta Working Paper series 363.
- Ross, S.M. (2007), Introduction to probability models, ninth edition, Academic press.
- Sherbrooke, C. C. (1968), METRIC: A multi-echelon technique for recoverable item control. Operations Research, 16:122–141.
- Sherbrooke, C.C. (2004), Optimal inventory modelling of systems: Multi-echelon techniques. Kluwer Academic Publishers, New York. 2nd edition.
- Silver, E.A., Pike, F.P., Peterson, R. (1998), Inventory management and Production Planning and Scheduling, Third edition, Wiley
- Trott, P. (2008). Innovation Management and New Product Development, Essex, Pearson Education Limited.
- Winston, W.L., (2004). Operations Research, Belmont, fourth edition, Thomson Learning.

Appendices

A. Lianders energy network: gas

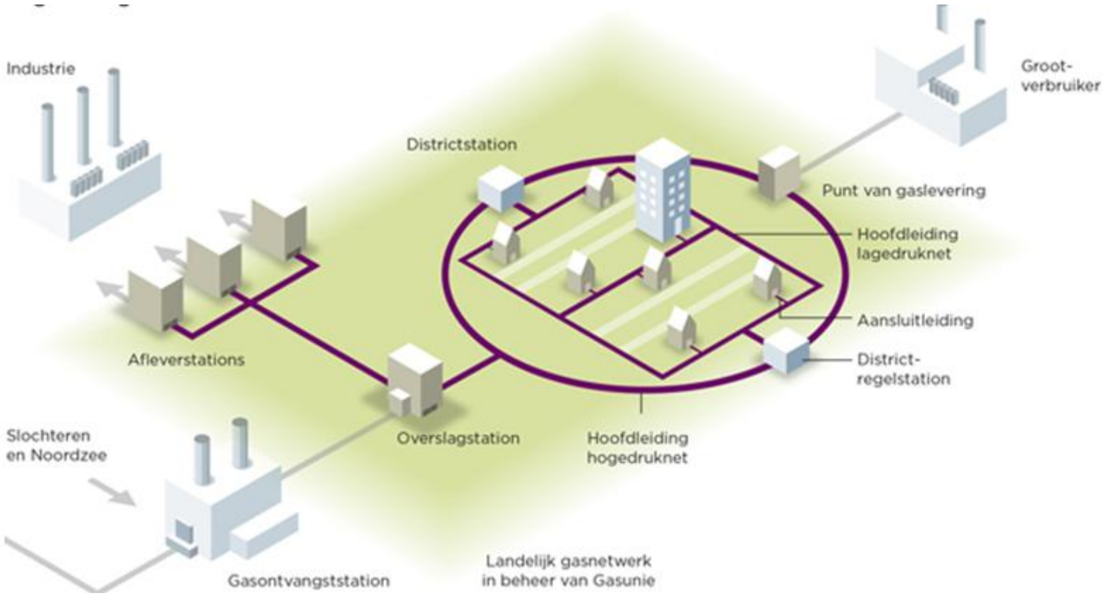

Illustration A-1 Responsibility Liander (green plane) in gas network

Source: Annual report 2011 Alliander

B. Distribution concept

Standard work

The first distribution covers the standard materials, these SKUs are cheap, standard and fast moving with little variation, see Illustration B-1.

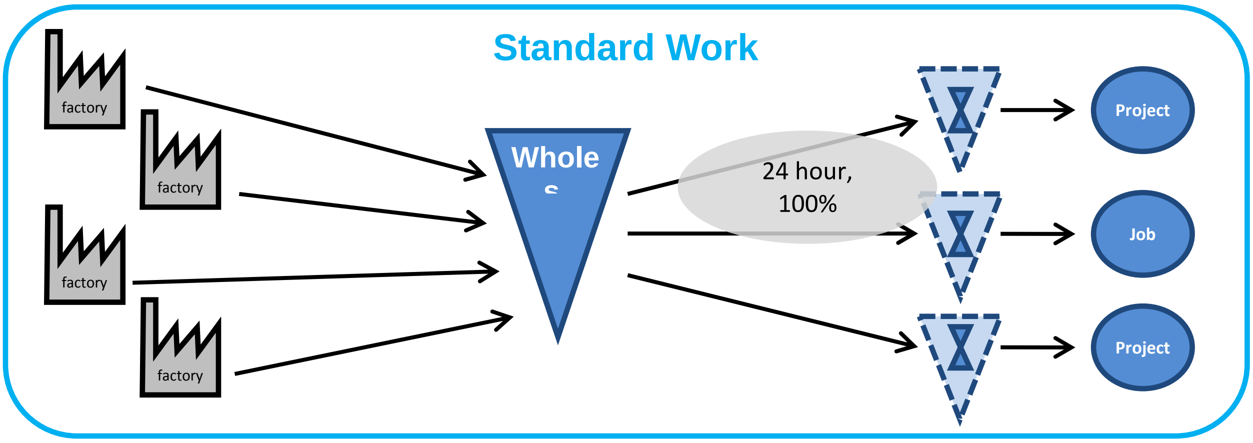

Illustration B-1 Standard work distributin concept

Source: Liander SOM vision document

In order to execute this function effective as possible a Wholesaler concept is chosen. It stores all material and delivers from inventory to transshipment locations at the local warehouses or directly to a project. The Wholesaler needs to deliver the total order within 24 hours. Who in the end wholesaler will be is still under consideration, it can be Liander or an external party. This concept requires less inventory local, since local material needs can quickly be delivered from central wholesaler.

On order unless

The second distribution concept, on order unless, are in general non-standard SKUs, expensive, slow moving SKUs with high variation, see Illustration B-2.

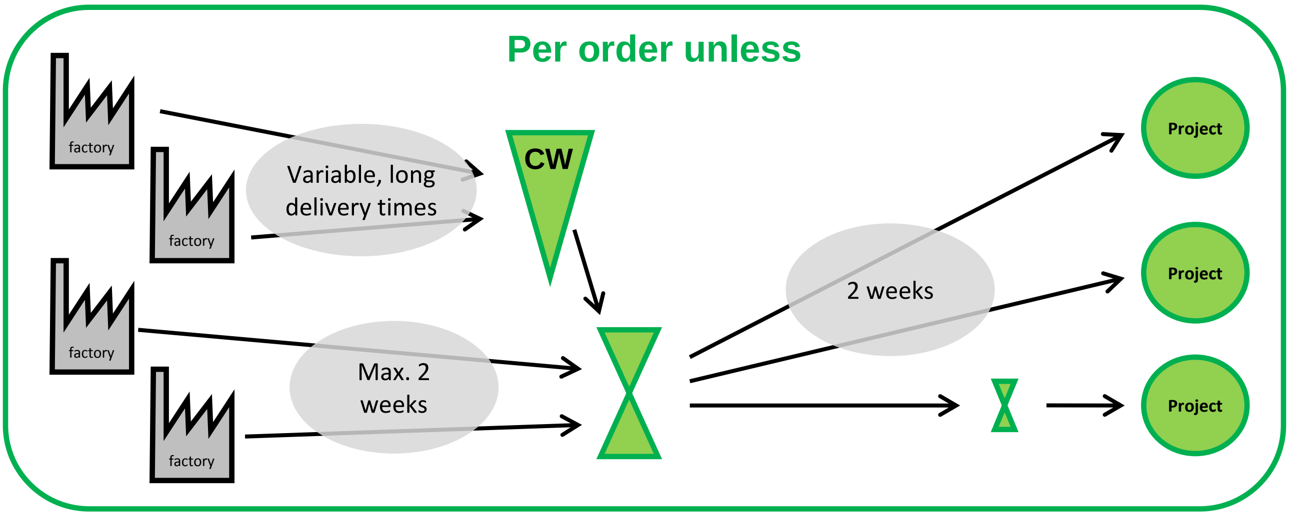

Illustration B-2 On order unless distribution concept

Source: Liander SOM vision document

The central warehouse commits to deliver these SKUs in two weeks, consequently these SKUs are only stored in the central warehouse if the lead times from suppliers are more than two weeks. If the lead times from the supplier are less than two weeks the central warehouse is used as a transshipment location. From the central warehouse the SKUs can be transported directly to a project or again be transferred at a local warehouse before reaching a project. In this way less expensive SKUs have to stored resulting in less costs.

C. Column generation

By the cutting stock problem the logic of column generation is easily understood. Suppose a factory produces standard raws with width W and different customers demand different raw widths, d_w, w is a certain width. Z is the number of different patterns demanded. A pattern, p, is a way of cutting a standard raw, where p_w is the number of raws in pattern p with width w, such that $\sum_{w=1}^{Z} p_w w \leq W$. Suppose we start with an initial set of patterns P', since there are many options. P denotes the set of all patterns consequently $P' \subseteq P$.

$$minimize: \sum_{p \in P'} x_p$$

$$s.t.: \sum_{p \in P'} x_p p_w \geq d_w \qquad \forall w \in \{1, \dots, z\}$$

$$x_p \in \mathbb{N} \qquad \forall p \in P'$$

Suppose we have solved the LP-relaxation of the problem above. The new pattern to add should have minimal reduced costs. The shadow prices of the demand constraint above are given by y_w. The reduced cost of a new pattern become: where c corresponds to the costs in the objective value of the a new pattern, which is non-basic.

$$c - \sum_{w=1}^{W} y_w p_w$$

Note that in this problem c=1 and p_w becomes a variable, consequently when we optimize the following ILP we find the pattern with minimal reduced costs:

$$minimize: 1 - \sum_{w=1}^{W} y_w p_w$$

$$s.t.: \sum_{w=1}^{W} p_w w \leq W$$

Note that instead of finding a new variable by linear programming we can use all kind of algorithms to find a new variable as long it creates a feasible pattern. In our method we also don't use linear programming to generate a new variable.

D. Two-echelon inventory systems with emergency shipments Model description

Appendix D describes the model developed by Özkan et al.(2011).

Notation:

sets

n Locations(0=central warehouse, {1,…,N}= local warehouses)

Parameters

t_n Planned replenishment time of warehouse n.
EC_n Emergency delivery from central warehouse to local warehouse n.
ES_n Emergency delivery from supplier to local warehouse n.
m_n The expected daily demand rate at warehouse n, $m_0 = \sum_{n=1}^{N} m_n$

variables

β_n Steady state fraction of demand satisfied by local warehouse n(fill rate)
θ_n Steady state fraction of demand satisfied by central warehouse
Υ_n Steady state fraction of demand satisfied by supplier
s_n Stock level at warehouse n
IL_n inventory level of warehouse n.
BC_0 Number of backorders at central warehouse.
LT_n The mean of the realized replenishment lead time for local warehouse n.
DL_0 Delay at central warehouse until *replenishment* order is fulfilled in weeks.

From β_n,θ_n and Υ_n easily al kind of performance measurements can be obtained like average waiting time or costs. For a graphical representation of the described 2-echelon model see Illustration D-1. Note that by definition(1):

$$\theta_n + \gamma_n + \beta_n = 1 \qquad \forall n$$

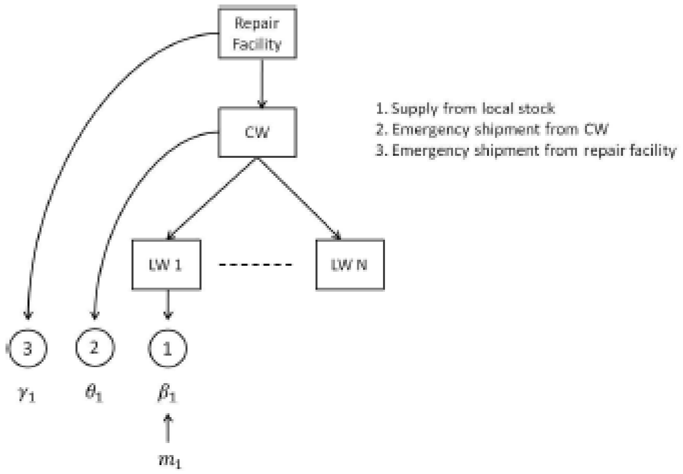

Illustration D-1 Representation of 2-echelon model with lost sales.

Approximation method

The method iteratively determines the fill rates at the local warehouse, β_n, and the expected delay at the central warehouse, DL_0. The calculations determine the fill rates at the local warehouses given a certain delay at the central warehouse and proceed with calculating the delay at the central warehouse given a certain β_n.

Calculating fill rates local warehouses

The lead time at local warehouse n is the result of the planned lead time and the delay at the central warehouse(2):

$$LT_n = t_n + DL_0$$

Note that the higher the base stock level at the central warehouse, the lower the value of DL_0. Demand at the local warehouses are not backordered, since it will be fulfilled from central warehouse or supplier. Consequently we can model the local warehouse by an Erlang loss system, M|G|c|c queue. Each item of a SKU acts as a server that is on average busy for LT_n time units when it serves an order. The system has S_n servers, arrival rate m_n and mean service time LT_n. As a result the percentage of accepted customers in the Erlang loss system is equivalent to β_n

Where:

P_n m_n*LT_n, product of arrival rate and the mean of realized lead time(3).

$$\beta_n = 1 - ERL(s_n, \rho_n) = 1 - \frac{\frac{{\rho_n}^{s_n}}{s_n!}}{\sum_{i=0}^{s_n} \frac{{\rho_n}^i}{i!}}$$

Calculating expected delay at central warehouse

Given certain fill rates at the local warehouses, the inventory level at the central warehouse is modeled by a birth-death process, i.e. a continuous-time Markov process with states $x \leq S_0$. At the central warehouse per local warehouse two demand streams occur Poisson distributed, replenishment orders($\beta_n m_n$) and emergency shipments($(1-\beta_n)m_n$). Only the first can be backordered, the second only occur when the central warehouse has at least on item on stock. As a result the demand at the central warehouse when stock is positive:

$$\sum_{n=1}^{N} m_n\beta_n + m_n(1-\beta_n) = \sum_{n=1}^{N} m_n = m_0$$

When stock level at central warehouse zero or negative(4):

$$m'_0 = \sum_{n=1}^{N} m_n\beta_n$$

The deterministic lead time to the central warehouse t_0 is replaced by an exponential distribution with rate $\mu_0 = \frac{1}{t_0}$. The authors show that the outcomes are rather insensitive to the lead time distribution. This enables the modeling of the inventory level by a birth-death process. The birth-death process state space in truncated, since backorders can only occur as a result of replenishment orders. The total number replenishment order from the local warehouse cannot exceed $\sum_{n=1}^{N} s_n = \bar{S}$, consequently the states x are truncated with $S_0 \leq x \leq -\bar{S}$. Illustration D-2 displays this birth-death process.

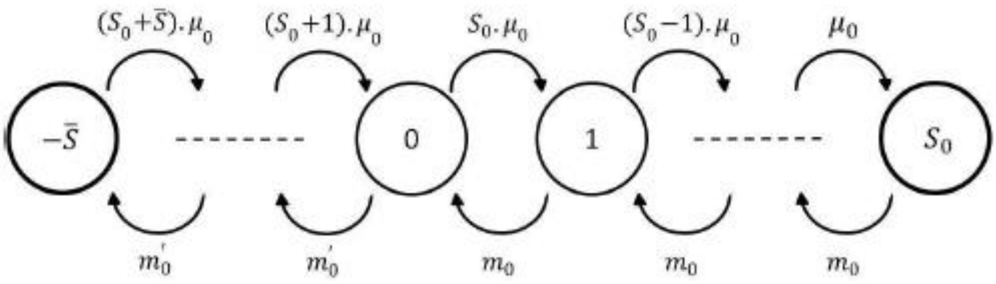

Illustration D-2 Graphical representation of the birth-death process for the inventory level at the central warehouse.

From the steady state distribution ,π_x, of the birth-death process follows the mean delay, DL_0. The steady state probabilities follows from(5):

$$\pi_x \begin{cases} \dfrac{m'_0}{(s_0 - x)\mu_0}\pi_{x+1}, & -\bar{S} \leq x < 0 \\ \dfrac{m_0}{(s_0 - x)\mu_0}\pi_{x+1}, & 0 \leq x < s_0 \end{cases}$$

The mean number of backorders, BC_0, is given by equation(6):

$$BC_0 = \sum_{x=-\overline{S}}^{-1} (-x)\pi_x$$

From Little's law follows, note that the arrival rate of replenishment orders is m'_0 (7):

$$DL_0 = \frac{BC_0}{m'_0}$$

The iterative procedure in order to determine the fill rates, β_n, n∊K and mean delay, DL_0, becomes:

Step 1: $DL_0=0$
Step 2: determine β_n via (2) and (3).
Step 3: Determine DL_0 via (4), (5), (6), and (7).
Step 4: Repeat step 2 and 3 while DL_0 change is more than ε

The authors experimented with different problem instances and obtained for all problem instances in their numerical study convergence, although they don't have theoretical results on convergence of this algorithm.

Calculating fraction of demand directly fulfilled from central warehouse and supplier

In order to determine these steady state fraction θ_n and Υ_n, the central warehouse and the supplier respectively. The authors introduce a random variable IL_n denoting the inventory level of the central and local warehouses. θ_n becomes:

$$\theta_n = \boldsymbol{P}(IL_0 > 0, IL_n = 0) = \boldsymbol{P}(IL_n = 0|IL_0 > 0)\boldsymbol{P}(IL_0 > 0)$$

The probability of a positive stock level of the central warehouse is determined by the birth-death process:

$$\boldsymbol{P}(IL_0 > 0) = \beta_0 = \sum_{x=1}^{s_0} \pi_x$$

The conditional probability $\boldsymbol{P}(IL_n = 0|IL_0 > 0)$, the authors propose to pretend the central warehouse has a positive inventory level for a long period. Then this probability can be modeled by a Erlang loss system, with mean service time t_n:

$$P(IL_n = 0|IL_0 > 0) \approx ERL(s_n, m_n t_n)$$

Consequently θ_n becomes(8):

$$\theta_n = \beta_0 ERL(s_n, m_n t_n)$$

Υ_n can be determined by (1), since β_n and θ_n are known.

E. Finding optimal solution of inventory allocation in local network

These observations are deducted from the article by Alvarez et al. (2012a) and enable us to find the optimal stock allocation in a 2-echelon network.

A complication with the local delivery network with emergency shipments is that the waiting time at local warehouse n depends on the stock levels at the other local warehouses. When the central warehouse stock level is positive, $s_{i0}>0$, an increase in stock level at local warehouse n will reduce its waiting time at the expense of the other local warehouse waiting times. As we consider the following case, where we increase s_{in} from 0 to 1, while other local warehouses have a stock level of zero. It means the regions corresponding to the local warehouses with stock level 0 are delivered with emergency deliveries from the central warehouse and the supplier. By a Markov chain we can model the stock level of the central warehouse. Only when local warehouse n changes its stock level from 0 to 1 we can reach a negative stock level at the central warehouse, as a result the steady state probability of a positive stock level at the central warehouse decreases. Hence an increase of s_{in} causes an increased waiting time at the other local warehouse, since more often deliveries needs to come from the supplier, instead of the central warehouse.

Observation 1 Upper bond on s_{i0}^{UB}

We define $red(i)^{LB}(s_{i0})$ as the minimal reduced costs for a given s_{i0}, stock level of central warehouse.

Recall that the reduced costs are given by:

$$reduced\ cost\ policy_p = c_{pi} - J_i - w_{pi} * q_i * A$$

Where A(≤0) is shadow price of the SVBM constraint and J_i(≥0) the shadow price of constraint requiring every item has one policy.

From these results we can obtain $red(i)^{LB}(s_{i0})$, a lower bound on the reduced cost given a stock level s_{i0} of item i at the central warehouse. For notation simplicity we omit suffix p and g, the reduced costs will be at least:

$$h_{i0}s_{i0} + \sum_{n=1}^{N} s_{in}^{LB}\, h_{in} - J_i - 52 * \sum_{n=1}^{N}\sum_{r=1}^{R} m_{ir}U_{inr}D_{nr} * q_i * A$$

Note that the smallest waiting time and no transportation cost is possible by transporting all items from the local warehouses, implying β_{in}=1, since $D_{nr} < D_{0r} < D_{-1r}$ for n∈{1,…,N}, r∈{1,…,R} driving times from local warehouse, central warehouse supplier respectively to a region. Consequently this is a lower bound on the reduced costs because the reduced costs will always be higher as a result of increased inventory level at the local warehouses, increased waiting time and transport costs as normally β_{in} <1. Now we can find a maximum on the inventory level of the central warehouse, s_{i0}^{max} , by finding the smallest s_{i0} where the following holds:

$$h_{i0}s_{i0} + \sum_{n=1}^{Z} s_{in}^{LB}\, h_{in} - J_i - 52 * \sum_{n=1}^{N}\sum_{r=1}^{R} m_{ir}U_{inr}D_{nr} * q_i * A > 0$$

Note adding policies with nonnegative reduced costs will not improve the LP-relaxation.

Observation 2 – Tighter upper bound on s_{i0} by considering s_{in}

In order to create a tighter upper bound on s_{i0} we need a tighter lower bound on the reduced costs as used in observation 1. This is done by considering the inventory levels of the local warehouses. We define $RED_{in}(s_{i0,}s_{in})$ as the related reduced cost to local warehouse n with stock level s_{in} when the central warehouse stock level is s_{i0}. For simplicity we omit parameter U_{nr} and the required summation which links a local warehouse to a region.

$$RED_{in}(s_{i0}, s_{in}) = h_{in}s_{in} - 52m_{ir} * ((\beta_{in}D_{nr} + \theta_{in}D_{0r} + \gamma_{in}D_{-1r}) * q_i * A - \theta_{in}CW + \gamma_{in}CS)$$

In comparison to the lower bound on the reduced costs formulated above also transport costs are included in this formula. Transport costs occur when items are delivered to a region directly from the central warehouse or supplier. Due to the same reason increased waiting time occur since delivery times from the central warehouse and supplier are longer than from the local warehouse. Waiting time is multiplied by the shadow price of the SVBM constraint, A. Furthermore q_i is included since this is a LHS parameter in the SVBM constraint. J_i, shadow price of constraint requiring one policy for an item, is excluded as this formula only considers the reduced costs related to local warehouse n.

$Red(i)^{LB}(s_{i0})$ becomes:

$$h_{i0}s_{i0} - J_i - \sum_{n=1}^{N} \min_{s_{in} \in [s_{in}^{LB}, \dots, s_{in}^{UB}]} RED_{in}(s_{i0}, s_{in})$$

S^{UB}_{in} follows from observation 3. In order to evaluate $RED_{in}(s_{i0}, s_{in})$ we need to specify the stock levels of the other local warehouses. We put the stock level of the other local warehouses on s_{in}^{min}, as a result we will obtain the lowest waiting time and transportation costs at local warehouse n and consequently the lowest reduced costs.

Observation 3 – upper bound on s_{in}, nϵ {1,...,N}

An increase of s_{in}, nϵ {1,...,N} can only benefit region r corresponding to local warehouse n. Consequently we can find s_{in}^{max} as additional inventory holding costs, h_{in}, outweighs the maximum reduction in waiting time and transport costs at local warehouse n. s_{in}^{UB} is the smallest s_{in}, where the following holds:

$$h_{in} > -52m_{ir} * \left(\left(\theta_{in}(D_{0r} - D_{nr}) + \gamma_{in}(D_{-1r} - D_{nr}) \right) * q_i * A - \theta_{in}CW + \gamma_{in}CS \right) \forall\, n \in \{1, \dots, N\}$$

To ensure s_{in}^{UB} nϵ {1,...,N} is large enough we need an upper bound on the waiting time and transport costs. This we do by taking an upper bound on γ_{in} by assuming demand can only be met from on hand stock at local warehouse n or from the supplier, implying $\theta_{in} = 0$. This can be modelled by an Erlang loss system, where the resupply time becomes replenishment time of central warehouse, t_0, and local warehouse, t_n, combined.

$$\beta_{in} = 1 - ERL(s_{in}, m_{ir}(t_0 + t_n)) \qquad \forall\, n \in \{1, \dots, N\}$$

$$\gamma_{in} = 1 - \beta_{in} \qquad \forall\, n \in \{1, \dots, N\}$$

Observation 4 – Given s_{i0} a tighter upper bound on s_{in} nϵ {1,...,N}, $s_{in}^{UB}(s_{i0})$

We apply the same reasoning as in observation 3, we determine $s_{in}^{UB}(s_{i0})$ when the holding costs exceeds, as a result of increasing s_{in}, the waiting time and transport costs reduction at warehouse n. we now determine more accurate the waiting time compared to observation 3, as we now also consider the stock levels at the other local warehouses. We put them all on s_{ik}^{UB} $k \neq n$ and determine the actual waiting time and transport costs. The waiting time and transport costs will be the largest at warehouse n when all other local warehouses have the highest stock level.

Alvarez et al concludes empirically that the optimal value of s_{in} given a certain s_{i0}, denoted by $s^{opt}_{in}(s_{i0})$ lies between two values $s'_{in}(s_{i0})$ and $s''_{in}(s_{i0})$. $s'_{in}(s_{i0})$ is found by $\min_{s_{in} \in [s_{in}^{LB}(s_{i0}),\dots,s_{in}^{UB}(s_{i0})]} RED_{in}(s_{i0}, s_{in})$ while the stock levels of the other local warehouse are put to their maximum. $s''_{in}(s_{i0})$ is found in the same way but the stock levels of the other local warehouses are put to their minimum. This procedure gives us new values on $s_{in}^{min}(s_{i0})$ and $s_{in}^{max}(s_{i0})$, which we can reuse to find again $s_{in}^{LB}(s_{i0})$ and $s_{in}^{UB}(s_{i0})$. When these bounds stabilize after a number of iterations or $s'_{in}(s_{i0}) = s''_{in}(s_{i0})$ we stop.

F. LP-model – verification

Illustration F-1 displays a LP-model based on fictitious data, four SKUs, from the programmed algorithm in Delphi. This should match the described model in section 5.1.1.

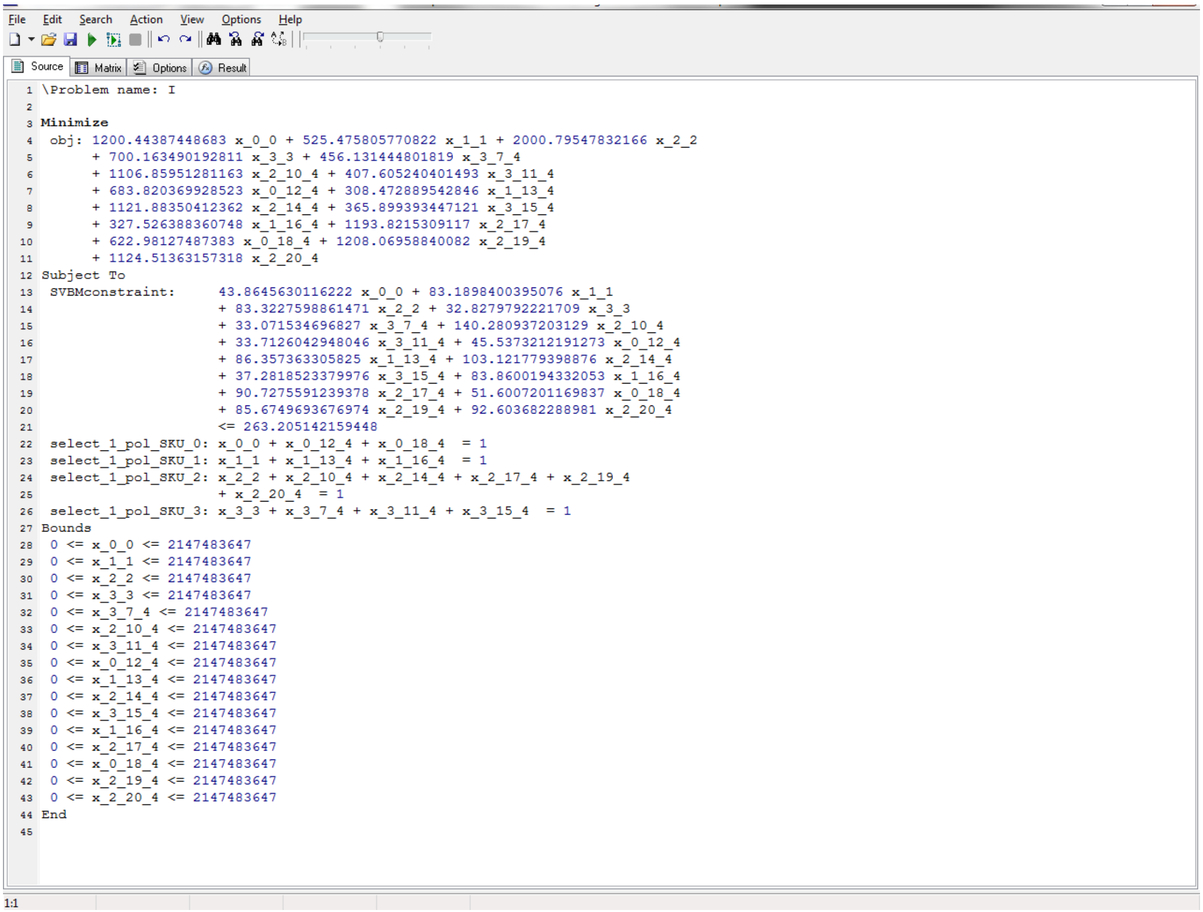

```
\Problem name: I

Minimize
 obj: 1200.44387448683 x_0_0 + 525.475805770822 x_1_1 + 2000.79547832166 x_2_2
      + 700.163490192811 x_3_3 + 456.131444801819 x_3_7_4
      + 1106.85951281163 x_2_10_4 + 407.605240401493 x_3_11_4
      + 683.820369928523 x_0_12_4 + 308.472889542846 x_1_13_4
      + 1121.88350412362 x_2_14_4 + 365.899393447121 x_3_15_4
      + 327.526388360748 x_1_16_4 + 1193.8215309117 x_2_17_4
      + 622.98127487383 x_0_18_4 + 1208.06958840082 x_2_19_4
      + 1124.51363157318 x_2_20_4
Subject To
 SVBMconstraint:     43.8645630116222 x_0_0 + 83.1898400395076 x_1_1
                     + 83.3227598861471 x_2_2 + 32.8279792221709 x_3_3
                     + 33.071534696827 x_3_7_4 + 140.280937203129 x_2_10_4
                     + 33.7126042948046 x_3_11_4 + 45.5373212191273 x_0_12_4
                     + 86.357363305825 x_1_13_4 + 103.121779398876 x_2_14_4
                     + 37.2818523379976 x_3_15_4 + 83.8600194332053 x_1_16_4
                     + 90.7275591239378 x_2_17_4 + 51.6007201169837 x_0_18_4
                     + 85.6749693676974 x_2_19_4 + 92.603682288981 x_2_20_4
                     <= 263.205142159448
 select_1_pol_SKU_0: x_0_0 + x_0_12_4 + x_0_18_4  = 1
 select_1_pol_SKU_1: x_1_1 + x_1_13_4 + x_1_16_4  = 1
 select_1_pol_SKU_2: x_2_2 + x_2_10_4 + x_2_14_4 + x_2_17_4 + x_2_19_4
                     + x_2_20_4  = 1
 select_1_pol_SKU_3: x_3_3 + x_3_7_4 + x_3_11_4 + x_3_15_4  = 1
Bounds
 0 <= x_0_0 <= 2147483647
 0 <= x_1_1 <= 2147483647
 0 <= x_2_2 <= 2147483647
 0 <= x_3_3 <= 2147483647
 0 <= x_3_7_4 <= 2147483647
 0 <= x_2_10_4 <= 2147483647
 0 <= x_3_11_4 <= 2147483647
 0 <= x_0_12_4 <= 2147483647
 0 <= x_1_13_4 <= 2147483647
 0 <= x_2_14_4 <= 2147483647
 0 <= x_3_15_4 <= 2147483647
 0 <= x_1_16_4 <= 2147483647
 0 <= x_2_17_4 <= 2147483647
 0 <= x_0_18_4 <= 2147483647
 0 <= x_2_19_4 <= 2147483647
 0 <= x_2_20_4 <= 2147483647
End
```

Illustration F-1 Example of produced LP-model by Delphi and CPLEX.

In Illustration F-1 we observe from top to bottom in the following sequence:

1. objective function – equation 1
2. SVBM constraint – equation 2
3. policy constraint – equation 3
4. the boundaries of the variables.

The variables corresponds with the policies. In the objective functions the variables are multiplied by the costs. In the waiting time/SVBM constraint all variables multiplied by the corresponding waiting time should combined result in less than 263.205, the max SVBM value. Finally every SKU should have one policy, therefore we observe for every SKU a specific policy constraint. The first displayed number per variable corresponds to the SKU number, these are consequently all equal per policy constraint. The boundaries do not correspond with the described model in section 5.1.1 as the model restricts the upper bounds of the variables to one while displayed model in Illustration F-1 allows variable values equal to $2.14*10^{9}$. All though this is huge difference the final result will be the same since variables in both model may not be negative and are all mentioned in one policy constraint restricting the sum of the variable values to 1, consequently higher variable values than 1 are in both model impossible. Consequently the displayed model in Illustration F-1 corresponds to the described model in section 5.1.1.

G. Toy problem

In order to show the operation of the presented model in chapter 0 we apply the model on a toy problem. The toy problem consist out of 7 regions, 4 unmanned warehouses and 3 manned warehouses. Every region has one local warehouse, manned or unmanned, MW's and UW's. Furthermore there is one central warehouse and a supplier. The 7 products needs to be allocated such that the 7 products combined does not cause more than a certain number of "StoringsVerbruiksMinuten", SVBM, while the costs are minimized. Matrix G-1 displays the demand characteristics per warehouse per product in the local supply chain, network 4. Furthermore displays Matrix G-1 the price and the waiting time criticality category per product and the direct delivery times from the central warehouse to the local warehouses corresponding regions.

Matrix G-1 SKU characteristics per warehouse: Demand, weekly arrival rates, per warehouse per product in the local supply chain(network 4) and price € and criticality category per product. Supply chain characteristics: direct delivery time from the central warehouse to the local warehouse corresponding region, D_{0r}.

Region/warehouse: Products:	1	2	3	4	5	6	7	price	Criticality category
Fuses	1.5	1.5	1.5	1.5	1.5	1.5	1.5	100	2
Joints	0.2	0.2	0.2	0.2	0.2	0.2	0.2	200	2
Elec meter s consump	2	2	2	2	2	2	2	300	3
Elec meter xl consump	0.2	0.2	0.2	0.2	0.2	0.2	0.2	600	3
Transformator	0.02	0.02	0.02	0.02	0.02	0.02	0.02	8000	2
Hermetic seal	0.08	0.08	0.08	0.08	0.08	0.08	0.08	1000	2
gas control station	0.01	0.01	0.01	0.01	0.01	0.01	0.01	15000	3
Delivery time from CW, D_{0r}	2.5	1.5	2.5	2.5	2.5	2	1.5	-	-

Criticality category 1 corresponds with 10 households averagely affected and 2 with averagely 100 households affected during a failure. For simplicity we assume all urgent orders are due to failures. Consequently results one hour of waiting time for a product in category 1 in $q_i = \frac{60*10}{3,000,000} = 0.0002$ SVBM and category 2 in $q_i = \frac{60*100}{3,000,000} = 0.002$ SVBM. q_i is used in the SVBM constraint of the LP problem. We restrict the caused SVBM to 0.1. Consequently mx, the RHS in the SVBM constraint becomes: the minimum number of SVBM possible in network 4 + 0.1. Note that in the solution space a fill rate of 1 at the local warehouses in network 4 results in the minimal SVBM possible for a product.

Matrix G-2 displays the demand characteristics per MWs in network 3, the regional supply chain.

Matrix G-2 Demand per manned warehouse, mw, in the regional supply chain(network 3).

manned warehouse: Products:	1	2	3
Fuses	3	6	1.5
Joints	0.4	0.8	0.2
Elec meter s consump	4	8	2
Elec meter xl consump	0.4	0.8	0.2
Transformator	0.04	0.08	0.02
Hermetic seal	0.16	0.32	0.08
gas control station	0.02	0.04	0.01

From Matrix G-2 follows that MW 1 serves two regions, MW 2 serves four regions and MW 3 serves one region in network 3. The direct delivery time form local warehouse n to region r, D_{nr}, in network 4 is always 20 minutes for every warehouse. D_{nr} in network 3 is 20 minutes in case when the warehouse serves its own region and 45 minutes when it servers a nearby region. Consequently on average the delivery time in network 3 becomes for MW 1: $\frac{20+45}{2} = 32.5$ minutes, MW 2: $\frac{20+45+45+45}{4} = 38.75$ minutes and MW 3: $\frac{20}{1} = 20$ minutes.

The remaining parameters values for both models are:

t_0 = 2 weeks, planned replenishment time for the central warehouse
t_n =0.29 weeks≈2 days, planned replenishment time for the local warehouses.
D_{-1r} =72 hours, direct delivery time from supplier to region r.
CW = € 50,- Cost of a direct delivery from the central warehouse.
CS = € 200,- Cost of a direct delivery from the supplier.
h_n = 0.25 of product price n∈{-1,…,n}.

In case we allocate the products at the supplier the additional required parameter values are:

$D_{-1'r}$ =24 hours, delivery time from the supplier to region r.
Safety factor = 1
Supplier's fill rate = 0.95.

Matrix G-3 displays the end solution. It specifies per product the optimal policy: network number, stock levels and expected yearly costs in € and waiting time in hours.

Matrix G-3 End solution per SKU: network number, costs €, waiting time hours and stock levels.

Item	Network number	Costs (€)	Waiting time (hours)	s_0	s_1	s_2	s_3	s_4	s_5	s_6	s_7
Fuses	4	1628.42	181	30	5	5	5	5	5	5	5
Joints	4	1057.47	25	7	2	2	2	2	2	2	2
Elec meter s consump	4	5084.34	259	38	4	4	4	4	4	4	4
Elec meter xl consump	3	1841.25	53	5	2	3	2				
Transformer	2	4396.45	31	2							
Hermetic seal	3	2261.15	18	4	2	2	1				
gas control station	1	2978.00	96								

When we analyze the solution we see that the relative cheap items which are frequently demanded with significant criticality are placed in network 4, see Matrix G-3 first three items. The electricity meter small consumption is an exception as it is the most expensive of these three and has small criticality. This product is discussed in detail below. Items which are more costly and less frequently demanded are placed in network 3. The most costly items are placed only at the central warehouse or at the supplier, corresponding to network 2 and 1. Network 1 and 2 allows relative cheap policies at the expense of high waiting time per order, especially when the product is placed at the supplier. These options are optimal for the transformer and gas control station due to high item costs. When a product is rarely demanded as is the case for these two products it still results in little SVBM, while the waiting time per order is relative high.

Note that in network 3 little stock is kept at local warehouse 3, s_3, in comparison to local warehouse 2, s_2, this is due to the fact that demand size of local warehouse 2 is 4 times bigger as the demand size at local warehouse 3.

Determining costs and waiting time

In order to give additional insight and to verify cost and waiting time determination we discuss for the third product, electricity meters small consumption, and seventh product, gas control station, how the costs and waiting times are determined. We start with electricity meter small consumption, which is placed in network 4. In our problem context every local warehouse has the same demand in network 4 resulting is the same stock levels at the local warehouses. Consequently the fractions from where the demand is fulfilled will be the same for every local warehouse, these are for the electricity meter small consumption:

β_n =0.9972 Fraction of demand satisfied by local warehouse n.
θ_n =0.0025 Fraction of demand satisfied by central warehouse.
Υ_n =0.0003 Fraction of demand satisfied by lost sales.

These fractions follows from the model of Özkan et al., (2011). From these fractions the costs and waiting time can be determined. Recall yearly costs are given by:

$$c_g = \sum_{n=0}^{|L_g|} h_n s_n + \sum_{r=1}^{R} 52 * m_r U_{nrg} \left(\theta_n CW + \gamma_n CS\right) \qquad \forall\, g \in \{2,3,4\}$$

Holding costs per item become: $h_n = 300 * 0.25 = 75$

$$c_4 = \sum_{n=0}^{7} (75 * s_n + 52 * \ 2 * (0.0025 * 50 + 0.0003 * 200))$$

$$\approx 4950 + 135 \approx 5085$$

€ 5085,- is not exactly equal to € 5084.34 due to the loss of significance in the different fractions. Recall yearly waiting time is given by:

$$w_g = 52 * \sum_{n \in L_g} \sum_{r=1}^{R} m_r U_{nrg} (\beta_n D_{nr} + \theta_n D_{0r} + \gamma_n D_{-1r}) \qquad \forall\ g \in \{2,3,4\}$$

On average D_{0r}, displayed in Matrix G-1, is for every region: $D_0 = \frac{2.5*4+1.5*2+2}{7} = 2.14$. Note this way of calculating is only valid when every region has the same demand size.

$$w_4 = 52 * \sum_{n=1}^{7} 2(0.9972 * .33 + 0.0025 * 2.14 + 0.0003 * 72) \approx 259\ hours$$

After rounding the yearly waiting time becomes equal as in Matrix G-3 to 259 hours waiting time in a year.

The gas control station is allocated in network 1 in this case the costs and waiting time are calculated by:

$$c_1 = s_{-1} * DS * h_0 + 52 * m * CS)$$

Where s_{-1} follows from the Erlang loss formula. Where s_{-1} is the lowest value yielding:

$$1 - ERL\left(s_{-1}, \frac{m}{DS} t_0\right) = \ 1 - ERL\left(s_{-1}, \frac{0.07}{0.2} 2\right) > 0.95 \rightarrow \ s_{-1}{=}4$$

$$c_1 = 3 * 0.2 * 0.25 * 15000 + 52 * .07 * 200 = 2978$$

Yearly waiting time in hours in network 1 is given by:

$$w_1 = 52 * \sum_{r=1}^{R} m_r \left(X D_{-1'r} + (1 - X) D_{-1r}\right)$$

$$= 52 * \sum_{r=1}^{R} 0.01\,(0.95 * 24 + 0.05 * 72)$$

$$= 52 * 0.07 * 26.4 \approx 96$$

Actually we should also discuss the determination of the fill rate when a product is placed in network 2, this is done manually with correct outcomes. Recall that the fill rate determination in network 2 is highly comparable to network 1.

Column generation Toy problem

Recall that the reduced costs are given by, where J_i>0 and A<0:

$$c_{ip} - J_i - w_{ip} * q_i * A$$

"A" represents the shadow prices of the constraint restricting the number of SVBM, equation 2, and J_i represents the shadow price of the constraint restricting every product to have one policy, equation 3.

Matrix G-4 displays the intermediate results of the column generation procedure for electricity meter small consumption. From Matrix G-4 we can also find the corresponding shadow prices when the above optimal stock allocation was calculated by these results we can reproduce the found reduced cost for the electricity meter small consumption. See Iteration step 4, A=-53248 and J_i=7862.

$$5084.34 - 7861.68 - 258.69 * 0.002 * -53248 \approx -22.39$$

Nearly equal to the displayed 22.36 in Matrix G-4 iteration 4.

Matrix G-4 Progress column generation for electricity meter small consumption, iteration number, stock levels, reduced cost new variable, shadow prices.

Iteration	s_0	s_1	s_2	s_3	s_4	s_5	s_6	s_7	Reduced Costs	Shadow price SVBM, A	Shadow price Item, J_i
1	37	4	4	4	4	4	4	4	-2650.33	-45786	10115
2	33	7	10	5					-45.32	-25609	6391
3	40	4	4	4	4	4	4	4	-18.44	-63479	8406
4	38	4	4	4	4	4	4	4	-22.36	-53248	7862
5	-	-	-	-	-	-	-	-	-	-47109	7522
6	-	-	-	-	-	-	-	-	-	-47109	7522

In iteration step 2 we observe a relative big A(<0) and relative small J_i(>0). As a result waiting time is not severely punished and the procedure finds the best allocation in network 3. While in all other steps network 4 is the best. Recall that in general network 3 is cheaper as network 4, but causes more waiting time since we have less local warehouses. In iteration step 3 the opposites happen when A is relative small and J_i is relative big. The procedure chooses network 4 with compared to other network 4 options high stock levels, see s_0.

In the last iteration steps the shadow prices stabilizes and no more variables with reduced costs can be found for the product under consideration. Note that in the last iteration step never new variables are found otherwise a new iteration step is required.

When we observe the chosen networks of the added policies during the column generation only shifts occur in the chosen networks for the electricity meter small consumption and the Hermetic seal. In case of the electricity meter small consumption it is logical compared to the first two products as these are more cheap and have more waiting time criticality. The hermetic seal is although not frequently demanded and relative expensive still interesting to place in network 4 due to the high criticality of waiting time. Finally the first two factors overrule the criticality factor for the hermetic seal as it is finally placed in network 3.

H. Ordinal regression output

Case Processing Summary

		N	Marginal Percentage
Network	1	34	4.4%
	2	134	17.2%
	3	261	33.6%
	4	348	44.8%
Valid		777	100,0%
Missing		0	
Total		777	

Model Fitting Information

Model	-2 Log Likelihood	Chi-Square	df	Sig.
Intercept Only	1812.328			
Final	1359.142	453,186	3	,000

Link function: Logit.

Goodness-of-Fit

	Chi-Square	Df	Sig.
Pearson	1804.714	2307	1.000
Deviance	1359.142	2307	1.000

Link function: Logit.

Pseudo R-Square

Cox and Snell	.442
Nagelkerke	.489
McFadden	.250

Link function: Logit.

Parameter Estimates

		Estimate	Std. Error	Wald	df	Sig.	95% Confidence Interval	
							Lower Bound	Upperbound
Threshold	[network = 1]	-3.252	.243	178.929	1	.000	-3.728	-2.775
	[network = 2]	-.734	.118	39.004	1	.000	-.965	-.504
	[network = 3]	1.471	.130	128.296	1	.000	1.216	1.726
Location	price	-.410	.039	110.615	1	.000	-.486	-.334
	tot_dem	3.186	.482	43.747	1	.000	2.242	4.130
	criticality_alt	178.684	13.597	172.685	1	.000	152.033	205.335

Link function: Logit.

Test of Parallel Lines[a]

Model	-2 Log Likelihood	Chi-Square	df	Sig.
Null Hypothesis	1359.142			
General	1361.818[b]	.[c]	6	.

The null hypothesis states that the location parameters (slope coefficients) are the same across response categories.[a]

a. Link function: Logit.

b. The log-likelihood value cannot be further increased after maximum number of step-halving.

c. The log-likelihood value of the general model is smaller than that of the null model. This is because convergence cannot be attained or ascertained in estimating the general model. Therefore, the test of parallel lines cannot be performed.

I. Verification model comparison Alvarez et al.

Before we can compare the outcomes of the model of Alvarez et al. (2012a) with our own developed model, we need to make adoptions in both models to overcome the differences. The section below describes this process.

Overcome model differences

Alvarez et al. use differently the fill rates outcomes of the warehouses in order to determine costs and waiting time. Alvarez et al. use two fractions to calculate these, instead of three in our model. The two fractions are defined and used as follows: The fraction of demand fulfilled by the supplier and the fraction of demand fulfilled by the regular system. The regular system is the central and local warehouse combined or in her context depot and customer location. As a result she doesn't make a distinction between an emergency delivery and a regular replenishment delivery from the central warehouse, furthermore no delivery time occur from the local warehouse or customer location. Therefore we put in our model t_n, planned replenishment time from the central warehouse to the local warehouses equal to D_{0r}, direct delivery time from central warehouse to region r. Furthermore we put D_{nr}, transport time from local warehouse n to region r, to zero and we put CW, cost of a direct shipment from the central warehouse to a region, to zero.

Besides these parameter adaptions we copied our method in respect of the cost and waiting time determination to Alvarez method, since no cost and waiting time determination was included for Özkan's method(2011). When we changed the determination of the waiting time and costs the column generator procedure did not function anymore as E.M. Alvarez method was based on the original cost and waiting time determination of different network evaluation methods. We solved this by lowering and increasing respectively the lower and upper bound of the stock levels to evaluate explicitly. As these bounds where also wrongly established due to the different way of calculating the waiting time.

The LP-model is differently as Alvarez restricts the average waiting per warehouse to a certain value, while we restrict the total waiting of all warehouses combined to a certain value. In this respect the model of Alvarez is more restrictive. We solved this by first run our own model and copy the corresponding local warehouse average waiting times of the solution to the model of Alvarez.

In our model we use different units of time, namely hours and weeks, while Alvarez applies one unit. Therefore careful attention was required to create the same problem instance for both models. In our model we do not allow a local stock of zero when demand occur, therefore the fictitious problem instance was created such that a zero stock allocation at a demand location was suboptimal. Finally we changed our model to only allow policies which uses network 4 and putting all q_i, waiting time criticality per SKU, see equation 2, to one.

Problem instance

Matrix I-1 displays a fictitious problem instance with four products and 3 regions. This problem we solved with both models.

Matrix I-1 Ficticious problem instance applied in both models.

Product\region	1	2	3	Price
1	0.1	0.4	0.5	300
2	0.2	0.8	0.9	100
3	0.3	0.4	1.2	400
4	0.05	0.6	0.1	200

Matrix I-2 and Matrix I-3 display the outcomes to the problem instance by the model of Alvarez et al. and the developed column generation model respectively.

Matrix I-2 Outcomes of LP-relaxation to problem instance of Table 6.1 by model Alvarez, solution value: 2262.73472214.

itemid	Variable Value	Costs	S0	S1	S2	S3
1	1.00000000	608.56	3	1	2	2
2	0.99999993	279.87	6	1	2	2
2	0.00000007	282.55	5	2	2	2
3	0.92776110	1023.08	6	1	1	2
3	0.07223890	977.27	5	1	1	2
4	1.00000000	354.54	3	1	2	1

Matrix I-3 Outcome of LP-relaxation to problem instance of Table 6.1 by model Liander, solution value: 2262.73473017.

itemid	Variable Value	Costs	S0	S1	S2	S3
1	1.00000000	608.56	3	1	2	2
2	1.00000000	279.87	6	1	2	2
3	0.92776121	1023.08	6	1	1	2
3	0.07223879	977.27	5	1	1	2
4	1.00000000	354.54	3	1	2	1

Both solutions do not exactly match due to the fact that the model of E.M. Alvarez applies an additional policy for SKU 2. This variable value is very low 0.00000007 and results in a comparable sized improvement of €0.00001. Consequently we see the solution as equal. When we changed the problem instance in respect of replenishment times and or emergency shipment time we observed the same occurrence of additional policies with very low corresponding variable values. This results in a slightly better objective value for Alvarez model. This is likely due to an acceptance threshold used in the column generation model or CPLEX.

Besides the small mentioned deviations the models yield exactly the same solution, see costs and the different stock levels. Therefore we conclude that our model functions correctly for network 4.

J. Column generation model outcomes

Matrix J-1 displays the results of the column generation model in terms of number of SKUs allocated in the different networks possible.

Matrix J-1 SKU allocation in terms of network number for different values of max SVBM', failure ratio 0.1

SVBM'	n1	n2	n3	n4
0,01	0	16	218	543
1,01	34	107	256	380
2,01	34	134	261	348
3,01	35	150	268	324
4,01	39	156	275	307
5,01	39	168	277	293
6,01	41	173	278	285
7,01	41	179	287	270
8,01	44	185	293	255
9,01	44	191	294	248
10,01	46	194	296	241
11,01	47	198	296	236
12,01	46	202	298	231
13,01	50	204	297	226
14,01	50	205	300	222
15,01	51	208	302	216
16,01	54	208	301	214
17,01	52	216	297	212
18,01	53	220	297	207
19,01	53	220	299	205
20,01	53	221	299	204
21,01	54	221	303	199
22,01	53	228	300	196
23,01	55	234	294	194
24,01	54	238	296	189
25,01	58	235	295	189
26,01	56	239	296	186
27,01	57	238	302	180
28,01	56	240	307	174
29,01	56	241	307	173
30,01	58	239	310	170

In line with our expectations we observe in Matrix J-1 more items stored in decentral networks as the allowed numer of caused SVBM' increases. Recall network 1 is most centralized network and network 4 most decentralized network.