27-1-2016

Dynamic consolidation decisions in a synchromodal environment

Improving the synchromodal control tower

Student:	P.J. Vinke
First supervisor:	Dr. Ir. M.R.K Mes
Second supervisor:	Dr. Ir. J.M.J. Schutten
External supervisor:	Ing. J.J. Tenhagen
External supervisor:	Ir. J. Grapendaal

Management summary

Optimisation of freight logistics is a long standing topic that currently still addressed. New planning methodologies are constantly being developed to make freight logistics more efficient. A recent trend in logistics is the shift from intermodal transport planning to synchromodal transport planning. In preceding research by Spikker (2014) the usefulness of synchromodal transport planning for Seacon has been shown.

This research continues on the aspect of synchromodal transport planning. We are interested in the effect that consolidation and equipment repositioning may have on synchromodal transport planning, where the bulk of the freight is containerised cargo. To this end we pose the following research question:

In what way and up to what degree is it possible to support real-time consolidation and equipment repositioning decisions in a synchromodal environment?

To answer this question we developed a heuristic that is tested through simulation. This heuristic covers three steps. First, orders are planned on an order-by-order basis. Second, the plan that results from the planning process is optimised through a simulated annealing procedure. Finally, a decision rule determines the amount of empty containers to be transported. To assess the performance of our heuristic, two benchmarks are defined. The first benchmark is a direct trucking benchmark, where all orders are trucked. The second benchmark is a planning heuristic that plans orders on an order-by-order basis and uses a decision rule for the determination of empty containers to be transported.

The proposed heuristic and benchmarks are tested on both a case network and generic networks. We distinguish between two cost structures: a fixed and a variable structure, where the fixed cost structure is considered the most interesting. In the comparison between our second benchmark and our optimisation heuristic in a network with a fixed cost structure our main findings are:

- Cost reductions varying between 11 to 17%.
- Utilisation rate increase on transport legs by 9 to 12%.
- Change in orders trucked varying between -3.5 to 95%.
- Reduction in empty containers moved varying between 38 to 47%.

For the case network, where we consider three different train schedules, we find that both zero and one intermediate stop (hop-on/hop-off terminal) per week are considered favourable. The key results for these two settings are:

- Expected utilisation rates of 80.4% and 84.8%.
- Expected number of orders trucked of 10.6 and 8.2 containers per day.
- Expected empty containers moved of 14.5 and 15 container per day.

The main recommendations that we give are the integration of the pre- and end haulage into the model to assess the interaction with the long haul transport. Further, we recommend to run a shadow project on a somewhat larger network, next to operational activities to assess the performance with real data.

Preface

***"No battle plan survives contact with the enemy"* (Helmuth von Moltke, no date)**

Just as no battle plan survives first contact with the enemy, so does a plan in the world of logistics not survive dynamical events. An event in the world of logistics usually alters the established transport plan. This can simply be an additional order added the plan or more drastically a complete revision of the plan. Several discussions with my supervisors helped me to present you this report in its current state.

When I first started this assignment, I was relatively clueless about the world of freight logistics. During my graduation time, I was able to experience some of the many aspects that freight logistics has to offer.

First and foremost I thank my supervisors for the assistance during my graduation time. Without their help this thesis would not be as it is now. Both Martijn and Marco for their no-nonsense feedback and the occasional laugh for minor idiotic mistakes. Further, both Joris and Jarco for their valuable contributions with regard to the contents.

Further I thank my colleagues from the business development department for the good time there. Most memorable are the constant stream of bad jokes and the weekly Schlagerstunde on Friday afternoons.

Finally and most important to me, I thank my family for their patience and support throughout the years.

Patrick Vinke

Wednesday, 27 January 2016

Table of Contents

List of figures

List of tables

1. Introduction

This chapter serves as introduction to the company and the problem. To start, Section 1.1 introduces Seacon logistics. We address the context of the problem and its core in Section 1.2. We set our scope in Section 1.3. Section 1.4 introduces the main research question and subdivides this into sub-questions. In Section 1.5, we give our methodology and the outline of this thesis.

1.1 Seacon

Seacon logistics is a leading logistics service provider (LSP) that operates at a global level. The company is founded in 1985 and has since then grown to be leading in multimodal chain management. Seacon operates as a third party logistics (3PL) provider in a constantly and rapidly changing environment. Goods have to be delivered in a rapid and timely manner against the lowest possible costs. At Seacon this is achieved by using intermodal transport. Intermodal transport involves transportation of goods in a standard transport unit, for example a container. This container is then transported through usage of at least two modes (truck, rail, barge, etc.) (Verweij, 2011).

As a 3PL provider, Seacon does not own assets in the form of trucks, barges, or trains. They do, however, own warehouses. As such, Seacon orchestrates the planning, contracts carriers to execute the actual transportation of goods, and has the ability to consolidate truckloads. Figure 1.1 gives a simple overview of the process that is mainly executed at the transport division of Seacon. As can be seen, the rough process consists of four steps: Incoming order, planning, transporting, and order receiving. A shipper sends an order to Seacon to request a transportation option to send goods from point A to point B. If the order gets accepted, Seacon communicates with carriers to transport the goods on modalities connecting A and B. From there, carriers transport the goods from A to B, where at point B the shipper receives the order.

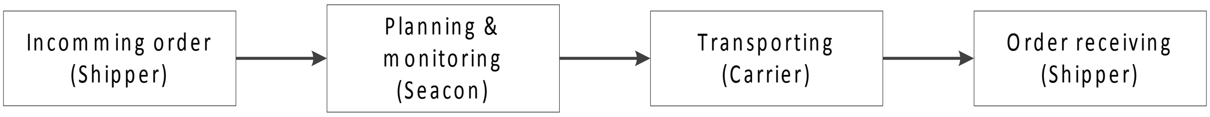

Figure 1.1: Rough process layout at Seacon.

In recent years, a new approach for planning and scheduling, with regards to intermodal transportation, received considerable attention. This approach is known as synchromodal transport planning. Synchromodal transport planning allows an LSP to schedule transport according to what currently is the best available transport mode. This can result in situations where on a given corridor at one moment truck transport is the best mode, whilst the next day transportations by rail or barge yields better results (Verweij, 2013). The key element here is that not the shipper or carrier but the LSP decides which mode to use, with decisions taken based on real time information.

The transportation market has a wide range of options that are considered. These options range from type of modality to type of container. On top of these choices, all kinds of restrictions and schedules with regards to pick-up and delivery times influence decision making. When combining these options and choices in a global network, decision making becomes a complex task. A project where a lot of these aspects come into play is a collaboration project at Seacon; we will go into more detail on this in Section 1.2.2.

1.2 Problem description

This section introduces the existing system, the problem context, and the core problem. The existing system is meant for comprehension of the problem context and its core. We give a rough sketch of the current situation in the problem context and the core problem is what our research aims to resolve.

1.2.1 Existing system

To understand the problem, a brief introduction to the existing system is needed. Within Seacon, the Control Tower (CT) department is responsible for the coordination of shipments and communication with external parties with regards to costs, progress, etc. To be able to do this, the department is assisted by a software program, the so called CT. Furthermore, this software is developed for intermodal transport planning. As Seacon aims to apply synchromodal transport planning, the CT recently has been adjusted for this. This is the so called Synchromodal Control Tower (SCT). In essence the SCT is software for synchromodal decision support, in this case specifically developed by Seacon in collaboration with the University of Twente and other partners. From here on we refer to the department as CT and the planning software as SCT.

In earlier research, Spikker (2014) compared three different planning strategies for Seacon: direct route planning, where shipments are sent from origin to destination by one mode (usually a truck), synchromodal planning, and consolidated synchromodal planning. The main findings were that both forms of synchromodal planning performed better than direct route planning. However, currently consolidation strategies are not applied structurally at the planning departments and consolidation decisions are not supported by the SCT. Therefore one of the recommendations of Spikker is to design a suitable consolidation strategy and implement this into the SCT.

The current way of working, at least as supported by the SCT, involves planning orders on a one-by-one basis. Transport orders arrive throughout the day and get planned for transport in that same sequence. When an order is planned, the planning is fixed and is not adjusted anymore. At this point it is possible to add new orders to the already planned shipments, which results in consolidation of orders. This consolidation can have positive effects on costs and carbon dioxide emissions. However, not in all cases it is possible to consolidate shipments. When it is not possible to consolidate the order with other orders, it is assigned to a new shipment. At this point it might be beneficial for other orders, which are already scheduled for transport, to be rescheduled to be transported with this newly planned shipment. This, however, is currently not supported.

A simple example to illustrate the issue with the one-by-one planning without considering replanning is the following. Consider three orders with the characteristics from Table 1-1 and two trains with characteristics from Table 1-2. We see that when we schedule the orders according to the one-by-one planning procedure, the first and second order get scheduled on train 1. The third order will then be scheduled on the second train without considering rescheduling to previous orders. A more logical schedule is to place all orders on the second train and therefore gain advantage by, for example, only having to pay fixed costs once, instead of twice.

Order	Time, order gets known	Due date
Order 1	2	7
Order 2	3	8
Order 3	4	9

Table 1-1: Set of orders with their characteristics

Train	Departure	Arrival	Capacity available for Seacon
Train 1	5	6	2
Train 2	6	7	3

Table 1-2: set of trains with their characteristics

1.2.2 Problem context

Seacon Logistics strives to play a leading role in multimodal chain management. To achieve and maintain this, Seacon keeps developing it services and capabilities. In recent years, intermodal transport planning is moving towards synchromodal transport planning. To stay competitive, Seacon is also adapting to this new planning method. One of the ways to achieve this competitiveness is through fully implementing the SCT. Currently, the SCT is not fully operational, it is still being improved and tested.

When considering consolidation, there are two distinguishable cases. First, consolidation can be done on goods level, where multiple less than truck loads (LTL) orders get consolidated into full truck loads (FTL) or near FTL, further referred to as LTL-consolidation. This results, amongst other benefits, in less containers to transport. Second, consolidation can be achieved on container level where orders consisting of one or multiple containers get consolidated in one transport batch and placed together on long-haul transport this phenomenon is called FTL-consolidation throughout this thesis. One of the results from this case of consolidation is economies of scale. This research focusses on the latter case. We elaborate on the decision for only considering FTL-consolidation in Section 1.4.

Collaboration project

For a collaborative logistics project with multiple shippers, executed at Seacon looks into possibilities of exploiting a rail service. For this project Seacon rents equipment. This equipment consists of containers and wagon sets. Wagon sets are a set of rail cars used for transporting, amongst others, containers by rail transport. The aim is to deploy these on the corridor Poland – Germany – United Kingdom. Naturally, Seacon strives to utilize the equipment optimally. Optimal utilisation here implies little unused capacity on the wagon sets and few empty container movements to meet the container demand at terminals. The latter is referred to as equipment repositioning.

For the mentioned corridor, it is known that a trade imbalance exists. This is shown by the fact that the westbound trade, from Poland to the United Kingdom, is larger than the eastbound trade. The rented containers used for the transport of goods are divided into two types of containers (Dry box and Reefers). Where the reefer containers are a more advanced type of container. The reefer containers are suitable for transportation of dry box goods as well. Dry box containers on the other hand are not suitable to transport reefer goods.

1.2.3 Core problem

The problem here is: *"How can the SCT be extended such that FTL-consolidation decisions are supported and the system is able to evaluate changes in orders already planned"* In other words, planning becomes integrated over several transport orders. The focus here lies on the dynamic decision model.

The aim is to develop a dynamic planning model where FTL-consolidation is optimized within acceptable computation time. An acceptable computation time is, as given by the CT department, execution within seconds. A favourable target is in less than one second. This is due to the fact that the system should not delay the order processing. Next to the first problem given, a second problem is required to be solved. The second problem has to do with the rented equipment which has to be positioned at the right place at the right time.

The first of these problems deals with freight consolidation and how to reschedule planned orders, whilst the second problem addresses the repositioning of empty containers. The insights gained from the literature are then used to create an algorithm which aids in decision making, on how to exploit a synchromodal network when long-haul consolidation and equipment repositioning are considered. Given a network, the algorithm should be able to compute which containers to send by which mode at what point in time, whilst being able to reschedule already planned orders.

The goal is to minimize the integral costs over a given planning horizon. This algorithm is tested on a case study. The case study involves the planning of shipments and empty containers to exploit the Poland – Germany – United Kingdom corridor.

1.3 Scoping

As with every real-life problem, accounting for every detail results in a too complex problem. To counter that, we focus on FTL consolidation on a railway (main-haul) corridor between Poland and the United Kingdom, with a possible 'hop-on-hop-off' rail terminal in Germany. The choice for railway transport comes from the fact that the case that we consider encompasses rail transport. Next to that, repositioning of equipment is addressed on this same corridor.

A reason for only selecting FTL trucking is the fact that for LTL consolidation, warehouses are needed at every terminal or hub. The corridor Poland – Germany – United Kingdom does not have warehouses at every terminal and therefore LTL consolidation is not a viable option for the time being. A direct result from only considering FTLs is the fact that consolidation is only considered on the long-haul as pre- and end-haul transport usually involves delivery by truck.

The case focusses on the transport corridor from terminals in Poland to terminals in the United Kingdom and vice versa. A corridor typically consists of multiple lanes that share a common origin and destination area. A lane is defined as the total transport route between a specific origin and destination (sender and receiver). Usually a lane is cut into three legs: pre-haul, main-haul, and end-haul. The pre- and end-haul is the transport between different shippers and terminals, this happens in a close area around the terminals, and is usually carried out by trucks. The main-haul is the freight transport between the terminals, this is usually performed by train or barge. Figure 1.2 graphically shows a basic transport lane. This scenario is how at the basis a transportation lane is partitioned (SteadieSeifi, Dellaert, Nuijten, van Woensel, & Raoufi, 2014)

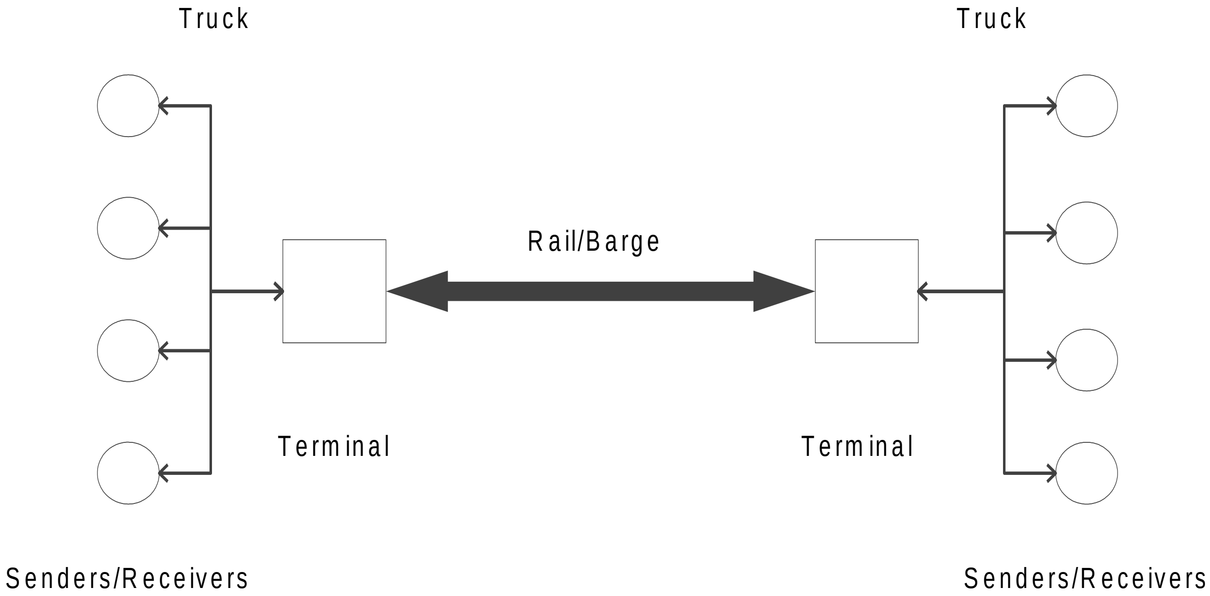

Figure 1.2: Basic transport lane

1.4 Research Questions

To reach the goal of developing a model for consolidation and equipment repositioning decisions, it is essential to gain insight into the concepts consolidation and equipment repositioning. As we are dealing with a synchromodal environment, we look into synchromodal planning as well. The main research question is as follows:

> ***In what way and up to what degree is it possible to support real-time consolidation and equipment repositioning decisions in a synchromodal environment?***

1.4.1 Sub-questions

To answer the main research question, this thesis is divided into six chapters. The first chapter serves as the introduction, the next four chapters focus on the many aspects surrounding the main research question with each of them having their own sub-questions. Finally, in the last chapter we answer the main research question.

Literature review

A central theme in this thesis is synchromodality. Before we focus on consolidation and repositioning of empty containers, synchromodality is explained. To do this, the following questions are answered.

1. What is synchromodal transport planning?

2. What has to be considered within a transport network for consolidation of container transport?
3. What has to be considered within a transport network for repositioning of empty containers?
4. Which type of models are used for planning, regardless of the transport planning environment?
5. Which methods are generally applied in solving the models found in literature?

Case & data analysis

To test our algorithm on the case study, we have to gain insight into the transport corridor. The following questions have to be answered to gain this insight.

6. What are the characteristics of the United Kingdom – Germany – Poland corridor?
7. Which data is available for the case?

Solution design

After conducting the literature study and the case & data analysis, we are able to build a planning algorithm. Here, the following question is answered.

8. What is an appropriate algorithm?

Simulation

To assess the performance of the proposed algorithm, we decide to simulate the transport process using discrete event simulation and apply it to the case. To gain managerial insights, the following questions have to be answered.

9. What performance can be expected from exploiting the United Kingdom – Germany – Poland corridor?
10. What are the results of the proposed algorithm?

1.5 Methodology and outline

As we mentioned in Section 1.4 we perform a literature study with regards to consolidation and repositioning of equipment. The focus is on planning models and methods for these two subjects.

After conducting the literature study, the case and its data are analysed. Partly this is achieved by evaluating an order set corresponding to the case data. However, not everything can be retrieved from data, because, not all data is known on the corridor Poland – Germany – United Kingdom and most of the data are aggregated. Therefore, we make assumption. These assumptions are verified by consulting experts to see if they reflect reality. When this is done, an algorithm can be constructed.

With the results from the literature study and the analysis we construct an algorithm. This algorithm is able to solve the consolidation and repositioning problems. To verify this, we apply the algorithm to the third sub-problem, the case. A good way to achieve this verification is through performing a simulation study.

In a simulation study, reality is represented through a model, where some of the aspects of reality are simplified. This model needs to be validated. This can be achieved in multiple ways. Common methods for validation are comparing the results of the model to the existing system and using expert opinions.

After the model is validated, the experimental design is constructed and simulation runs are made. The output from these runs are evaluated and based on these results, advice is given on the best to use policies with regards to consolidation and repositioning. The methodology applied is based on the procedures for a sound simulation study (Law, 2007). Figure 1.3 gives a schematic overview of the complete methodology used.

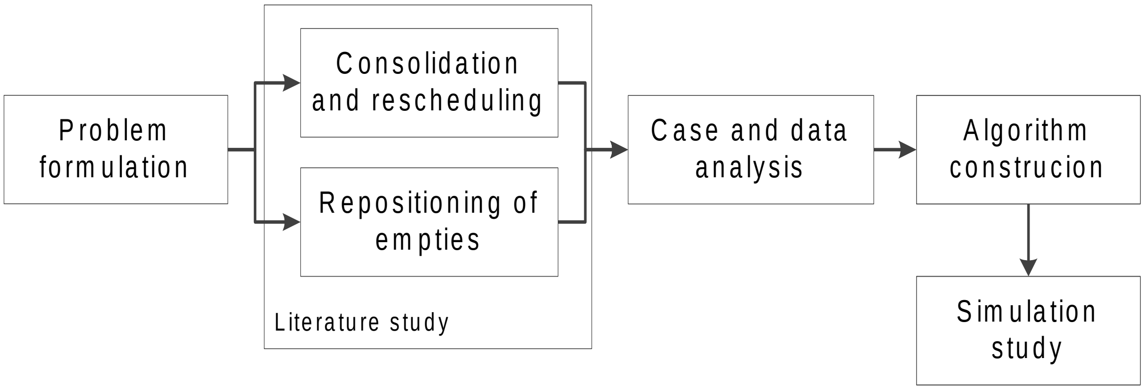

Figure 1.3: Methodological outlining

To finalise this chapter, we give an outline of this thesis in this section. In Chapter 2 the literature study is conducted, here we answer questions 1-5. To complement the answers found in literature, analysis of the case is performed in Chapter 3, sub questions 6 and 7 are answered here. Chapter 4 is about the solution design, here we answer question 8 as well, which we simulate in Chapter 5 where questions 9 and 10 are answered. Thereafter, in Chapter 6, the conclusions of this result are given, to end with a discussion and recommendations for further research in Chapter 7. Figure 1.4 gives a graphical representation of the outline.

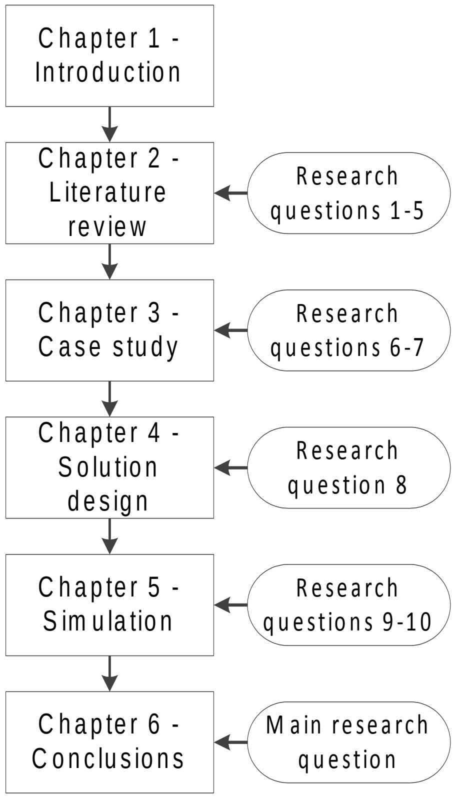

Figure 1.4: Outline of thesis.

2. Literature review

In this chapter we perform a literature study and answer the questions related to this. Section 2.1 focusses on synchromodal transport planning. Section 2.2 then focusses on freight consolidation whilst Section 2.3 focusses on the repositioning of empties. In Section 2.4 we address planning in a general setting for intermodal operators. Next, in Section 2.5 we address models found in literature. We present applied solution methods in Section 2.6. In Section 2.7 we give an overview of the models from Section 2.5, what they address and how they are solved. We conclude this chapter in Section 2.8

2.1 Synchromodality

In Chapter 1, the switch from intermodal to synchromodal transport planning has been addressed briefly. Intermodal transport planning is what currently is applied at Seacon Logistics. In this section we briefly address several transport planning concepts before going into more detail on synchromodal transport planning.

2.1.1 Synchromodal transport planning

In the past, several transport planning methods have been introduced. All of these methods aim to plan transport to get goods from A to B. The difference however comes from the way the goods are transported, how the transport modes are selected, and how the goods are packed. Verweij (2011) identifies five different transport planning types: unimodal, multimodal, intermodal, co-modal, and synchromodal.

Unimodal transport encompasses transportation from A to B where only one transport mode is used. In most cases this involves transportation by truck or van. The goods can be packed in standard units, for example a sea container, but this is not necessary. Next to unimodal transport, there is multimodal transport. Multimodal transport planning uses multiple modes of transport as opposed to unimodal transport planning (Verweij, 2011).

Intermodal transport is the same as multimodal transport in the sense that transport takes place through multiple modalities. The key difference here is that transport takes place in standard units (Verweij, 2011). Usage of standard units has the advantage that it is easier to switch between modes.

Co-modal transport makes use of the above three transport planning concepts. It combines them and aims to utilise resources in an optimal and sustainable way. Very similar to co-modal transport is synchromodal transport, which also strives for optimal utilisation. The difference here being that synchromodal transport puts emphasis on the total network instead of separate chains within the network (Verweij, 2011).

Besides the emphasis on the network there are some other details that distinguish synchromodal from the other transportation modes. Opposed to intermodal transport, where the mode choices have been predetermined, the mode choices happen based on real time information. These choices can be based on the current status of the different transport routes. Next to flexible decision making based on real time information, it is not the shipper that decides which transport modes are used but the logistics service provider. This is also known as A-modal or free mode booking

(Stuurgroep synchromodaliteit, 2013). To show some key differences, Figure 2.1 gives a graphical representation of intermodal, co-modal, and synchromodal.

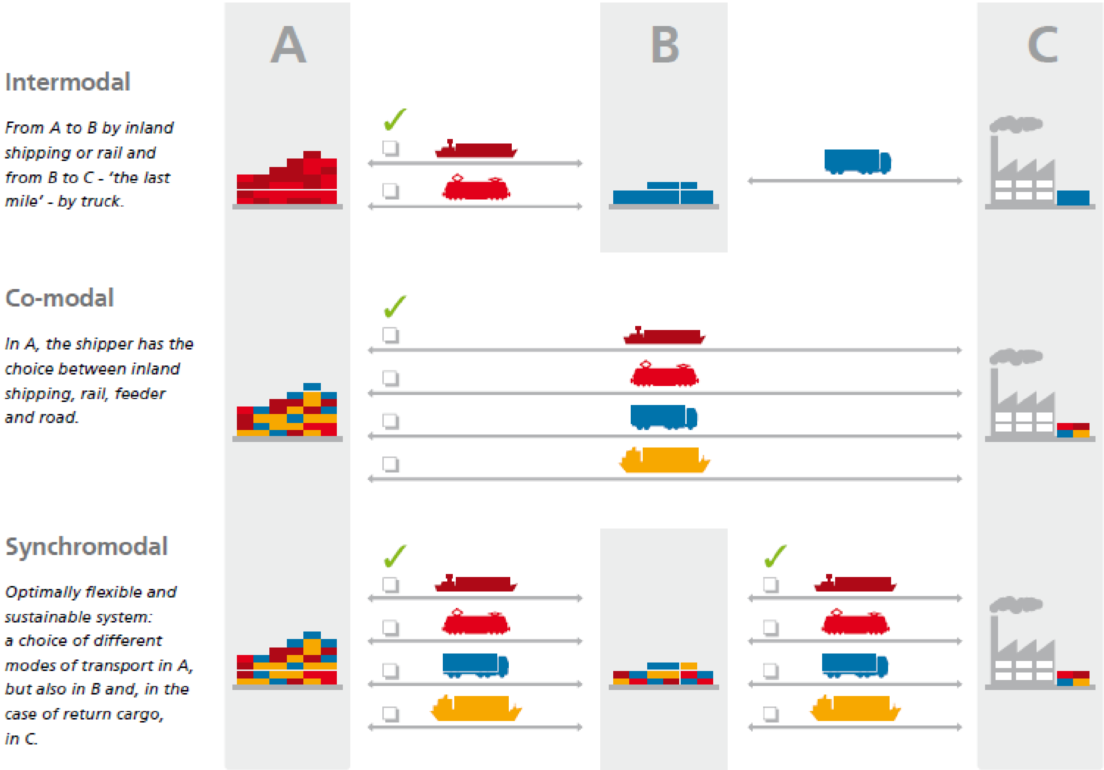

Figure 2.1: Transportation concepts (European container terminal, 2011).

In the literature, several definitions and descriptions for synchromodal transport planning are given. We believe that the description given by DINALOG is the most complete. For completeness we give several definitions found in literature.

Behdani, Fan, Wiegmans, & Zuidwijk (2014) define synchromodal transport planning as: *"An integrated view of planning and usage of different transport modes to provide flexibility in handling transport demand."* They further state that this is a general agreement on what the concept of synchromodal transport planning encompasses.

"Intermodal planning with the possibility of real-time switching between the modes" is the definition used in Van Riessen, Negenborn, Dekker, & Lodewijks (2013). The focus here lies on the real-time decisions with regards to the transport modes.

DINALOG (2015) describe synchromodal transport as an agreement between shipper and LSP on the delivery of products at specified costs, quality, and sustainability but leaves the mode choice free for the LSP to decide. Besides flexibility, the points of costs, quality and sustainability are specifically emphasised.

Another description found for synchromodality is: the optimal, flexible and sustainable usage of different transport modalities in a network under supervision of a logistics service provider, such that the customer receives an integrated transport solution. (Stuurgroep synchromodaliteit, 2015). Optimal in this case is left open for interpretation, but can be translated to: achieving lowest possible costs, fast throughput times, or high equipment utilistion.

Roth, Klarmann, & Franczyk (2013) describe synchromodality, when compared to multimodality, as: Switching between different forms of transport, but within a strategy of a more timely, efficient and environmentally friendly distribution. Another difference is that synchromodal logistics minimize buffer times, support bundling and allow quick changes between the different modes of transportation.

"Making optimal use of all modes of transport and available capacity, at all times, as an integrated transport solution" is the definition used by Ham (2012). Again, optimal is not given a clear definition and is left open for interpretation.

As we can see, there is no uniform definition for synchromodal transport planning, but all of the definitions have at least some of the aspects in common.

Behdani (2013) presents five key aspects for synchromodal transport planning. These are mode-free (or A-modal) booking, joint planning and coordination, bundling, flexibility, and visibility. Joint planning and coordination implies collaboration across a network and not simply focussing on the individual chains. Bundling is the combination of shipments into one transport batch. Management of the different modes and switching accordingly between them, defines the flexibility. Finally, visibility translates to situational awareness and the requirement of information sharing to make synchromodal transport planning viable.

Ham (2012) identifies the following characteristics:

- Mode free booking
- Dynamic planning of transportation
- Real-time switching between modes
- Decision making based on network utilization
- Combining transport flows
- Cooperation between actors in the transportation chain
- Information availability and visibility among actors.

These seven points are in essence a subdivision of the five key aspects identified by Behdani (2013).

From the definitions found in literature, all seem to agree that a certain emphasize is given with respect to free mode choice. Further characteristics that appear in the definitions are sustainability, costs, and quality. Sustainability can be expressed in, for example: reduction of carbon dioxide emissions. Costs is self-explanatory. Finally, quality can be measured in several ways, for example: throughput time or service level agreements.

2.2 Consolidation

Consolidation in a logistics contexts is the act of combining multiple orders or shipments into one batch. Reasons for consolidation can vary from economies of scale to sustainability improvement. Before we focus on consolidation characteristics we introduce general characteristics for transport networks in Section 2.2.1. Then, in Section 2.2.2 we address consolidation, followed by KPIs in Section 2.2.3.

2.2.1 Characteristics

For each order in a network, we identify several characteristics. Orders have an announcement time. The announcement time is the moment where an order gets known at the 3PL provider. Next, each order has an earliest pickup and latest delivery time (or a pickup and delivery time window). These pickup and delivery times are the earliest time that an order can be picked up at a shipper and the latest moment in time that it has to be delivered at the customer. Finally, each order has an origin, destination, volume, weight, and price.

Besides orders that have their own characteristics, the transport modes have their own characteristics as well. For the long-haul transport we consider three different modes: truck, ship (barge, short sea, and deep sea), and train. Arguably airfreight forms another mode but is not considered in this thesis. Characteristics that can be identified here are: capacity, transportation time (or speed), origin, fixed and variable costs, destination, departure and arrival time. Table 2-1 gives an overview of the characteristics for four common freight transportation modes considered in (Christiansen, Fagerholt, & Ronen, 2004).

	Mode		
Operational characteristics	Ship	Truck	Train
Fleet variety (physical and economic)	Large	Small	Small
Power unit is an integral part of the transportation unit	Yes	Often	No
Transportation unit size	Fixed	Usually fixed	Variable
Operating around the clock	Usually	Seldom	Usually
Trip (or voyage length)	Days or weeks	Hours or days	Days
Operational uncertainty	Larger	Smaller	Smaller
Right of way	Shared	Shared	Dedicated
Pays port fees	Yes	No	No
Route tolls	Possible	Possible	Possible
Destination change while underway	Possible	No	No
Port period spans multiple operational time windows	Yes	No	Yes
Vessel-port compatibility depends on load weight	Yes	No	No
Multiple products shipped together	Yes	Yes	Yes
Returns to origin	No	Yes	No

Table 2-1: Operational characteristics of freight transportation modes.

2.2.2 Decisions

An important question when considering consolidation strategies is where consolidation should take place for particular orders (Campbell, 1990). In the work of

Campbell the focus is on combining LTL shipments. However, given a network with long-haul transport, we believe that concept of combining LTL shipments is applicable for FTL shipments as well.

Higginson & Bookbinder (1994) give a generalized overview of management decisions that have to be made for shipment consolidations. The general decisions that are made are:

- What is consolidated?
- When are consolidated orders released?
- Where will consolidation be done?
- Who will perform the consolidation?
- How will the consolidation be carried out?

After the introduction of the general decisions Higginson & Bookbinder (1994) focus on the “when” question. They identify three shipment release policies: a quantity policy, a time policy, and a hybrid policy. When applying a quantity policy, consolidated shipments wait until a threshold quantity is reached before execution of the transport. When applying a time policy, consolidated shipments wait until a threshold time is reached before execution of the transport. The final policy is a hybrid policy. This policy has both a time and a quantity threshold. The first of the two thresholds that is reached triggers the execution of the transport. (Higginson & Bookbinder, 1994). The trade-off is the choice between time delay and utilization rate. With a higher order arrival rate, a quantity policy becomes more attractive.

Table 2-2 gives an overview of policies, based on the five general decisions presented in (Higginson & Bookbinder, 1994) and in which articles these policies are presented. We refer to these five general decisions as the what, when, where, who, and how questions. Naturally, this list is extendable, but it should give an overview of what to consider in decision making for intermodal freight transportation. Bontekoning, Macharis, & Trip (2004) note that for intermodal transport, specifically rail transport, fixed schedules are used as opposed to schedules in classic rail transport. In the classic case, trains depart when they are fully loaded. Because we are dealing with intermodal transport, a quantity policy is no option because trains depart at their designated time. Arguably, a hybrid policy is possible if multiple transport shipments are performed per day or a train can depart at multiple points in time. Otherwise a time policy is the only option.

Newman & Yano (2000) identify three factors in intermodal terminal networks. First, due to the structure of these networks relative few terminals are used contrary to conventional terminal networks. This in turn can lead to economics of scale. Second, due to distances only few stops are needed per journey, eliminating the need to consider blocks, a group or railcars sharing the same segment of the journey. Finally, and most relevant, shorter lead times are promised for intermodal freight. This in turn requires good management of scheduling transport modes. In conventional operations, freight may wait some time until enough is accumulated for a shipment. This however is less likely realistic with intermodal operations where there is less time available for waiting because of the shorter lead times.

In more recent research concerning consolidation, Rivera & Mes (2015) state: *"The combination of destinations of the orders consolidated determines last-mile costs."* The challenge is to decide which of the known and released for transport freights to transport on the long-haul vehicle. Given a set of orders, which orders should be selected for transport, which orders should be delayed for later shipment, and which orders need to be delivered by an alternative mode, e.g., a truck for direct shipment.

Hall (1987) looks at the "Where" question for freight consolidation. He distinguishes between three major different locations for consolidation: inventory, vehicle, and terminal consolidation. Inventory consolidation is what happens if, for example a production company lets products accumulate in inventory. Vehicle consolidation happens when a vehicle performs a roundtrip from a terminal to several customers to collect goods. This is a classic example of the well-known vehicle routing problem (VRP). Finally, terminal consolidation encompasses collecting the goods at a terminal where the goods then get sent as one transport batch to their (intermediate) destination terminal. As we are only considering the long-haul transport between terminals, it is not relevant to look into further detail at the "where" question.

The "Who" question identified by Higginson & Bookbinder (1994) is closely related to the policy chosen for where to consolidate. If an inventory policy is chosen, it is logical that the shipper arranges the consolidation whilst for a vehicle policy the carrier is a logical option. The terminal policy, bringing all the containers together for further transport, would be performed by the 3PL.

The "How" question, focusses on the specific techniques used for consolidation freight. Janic, Regglani, & Nijkamp (1999) introduce four different models for network layout: point-to-point, trunk line with collecting/distribution forks, hub-and-spoke, and line models. These four models are identified by Kreutzberger (1999) as well, he also mentions trunk-and-feeder networks. Figure 2.2 gives graphical representation of these network configurations.

In a point-to-point network configuration every origin-destination pair forms a single leg. In a trunk-line with collection and distribution forks long haul transport takes place between a select few main terminals. These terminals are connected to several feeder terminals that send and receive goods to and from the main terminals, the so called feeding. The actual long haul transport takes place between the select few terminals. The hub-and-spoke configuration uses a central hub, where goods from several origins are collected, sorted, and distributed to their destination. A line configuration sends a shipment across several terminals, loading and unloading at each of them. Finally, the trunk-and-feeder network combines the aspects of a line and the trunk-line configurations.

In more recent research, Kreutzberger (2010) compared the performance of the five different bundling networks. His main findings are that line, trunk collection and distribution, and trunk-and-feeder networks perform best with regards to costs when annual freight volumes are below 14,000 load units. At that point these network configurations get outperformed by hub-and-spoke networks. This remains till up to about 175,000 load units, at that point direct line networks is the choice of network for attaining lowest costs.

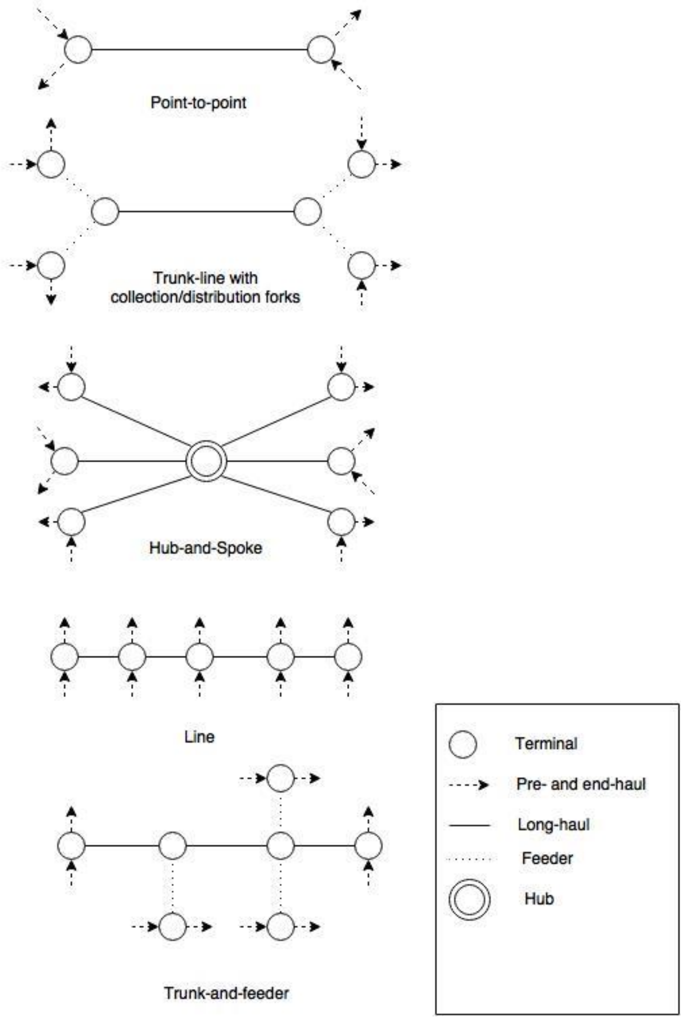

Figure 2.2: Types of network layout (Kreutzberger, 1999)

General decision	Decisions policy aspects	Articles
What?	Consolidate, delay, ship individual	Newman & Yano (2000) Rivera & Mes (2015)
When?	Triggers for dispatching	Higginson & Bookbinder (1994) Bontekoning et al. (2004)
Where?	Inventory, vehicle, or terminal consolidation	Hall (1987)
Who?	Shipper, carrier, third party, etc.	No articles found
How?	Specific consolidation techniques	Janic et al. (1999) Kreutzberger (1999)

Table 2-2: General decisions overview

2.2.3 Key performance indicators

As with every process improvement, we like to be able to say something about the actual improvement. To assess improvement we need to look at KPIs. Naturally cost or profit is one of the main performance indicators for a company LSPs are not excluded from this. However shippers would rather see high quality and low prices

(Lai, Ngai, & Cheng, 2004). This results in LSPs having to focus on more than just costs as performance indicator.

Krauth, Moonen, Popova, & Schut (2005) identify two major perspectives in literature when it comes down to KPIs for 3PL providers. The first group distinguishes between KPIs on individual firm performance versus a supply chain wide perspective. The second group distinguishes between KPIs on strategic and operational level. Further, they identify four stakeholders within the logistics: management, employees, customers, and society. Each of these stakeholders has their own desired performance indicators. We give an overview of some common indicators per stakeholder, logically this list can be extended.

- Management
    - Profit increase
    - Maximising utilisation
- Employees
    - Work conditions
    - Labour efficiency
- Customers
    - On time delivery
    - Low prices
- Society
    - Reduction of greenhouse gasses.
    - Reduction of overexploitation of roads

2.3 Equipment repositioning

This section covers the repositioning of empty containers. We look at characteristics and decisions.

In current logistic supply chains, a trade imbalance with regards to containers is noticeable. This creates the need for good empty container management, which involves repositioning of the empty containers (Theofanis & Boile, 2009). To further enhance the scope of the problem we look at the results from the first quartile of 2015 in the Port of Rotterdam. During this period a total of 1.64 million Twenty-foot equivalent unit (TEU) containers is imported. At the same time 1.49 million TEU is exported. This results in a difference of 150,000 TEU for the first quartile (Port of Rotterdam Authority, 2015).

Due to this trade imbalance it is possible that at certain locations an overflow of containers emerge, whilst other locations have a lack of containers. The lack of empty containers results in the need for a flow of empty containers. This imbalance forms a problem for carriers who usually perform the transport. For them the movement of empty containers does not yield revenues and is therefore unwanted.

Transport modes arriving at a terminal, usually carry one or multiple containers. These containers are either loaded or empty. Loaded containers are then transported onwards to their final destination, after some terminal handling. Empty containers on the other hand have two possibilities, they can remain at the terminal or be dispatched to a location where they might be needed for picking up new loads (Crainic, Gendreau, & Dejax, 1993).

After unloading at the destination, the container can return empty or it is sent to a shipper to pick up a new load. This is part of the pre- and end-haul optimization which is an interesting topic as well. However, this is not within the scope of this research but is covered in our recommendations in Chapter 6.

When considering the repositioning of empty containers, the major decision to take is thus whether to keep an empty container at a terminal for potential freight in the future or move the container empty to a different terminal where there is certain demand. A risk that occurs when this process is not correctly managed is that a shortage of containers at a terminal can occur. This shortage results in not being able to pick up new freights and therefore having to rent containers from an external party, decline an order, apply non-containerised transport, or delay transportation of the order. Higher total transport costs are usually the consequence of this.

2.4 Planning in a general setting

Increasing utilization rate of fleets can have a significant financial benefit. Another positive result of increasing fleet utilization is reduced emissions because of reductions in transport operations, e.g., less trucks on the road when trucks are filled for 80% instead of 50%. This is then an improvement to the sustainability of a company, which can be measured in terms of carbon dioxide emissions (Christiansen et al., 2004). To achieve good utilization, we have to look beyond the daily operational planning and look at the tactical planning level.

When considering planning at the tactical level in a freight transportation environment, Crainic & Laporte (1997) make a distinction between long and short distance transportation. The long distance transport is focussed on, amongst others, rail transport. Whilst the short distance focusses on pick-up and delivery. The former group is often referred to as service network design (SND), whilst the latter often gets addressed as a vehicle routing problem (VRP).

According to Crainic & Rousseau (1986) tactical planning for freight transportation has a main focus on three major problems, namely:

- Service network design
- Traffic routing design
- Terminal policies

These problems are related to each other, addressing them has impact on the whole network. SND is concerned with the type and level of service: which modes are offered, which routes are exploited, and with what frequency is transport offered. Traffic routing design focusses on the movement of modes through a network. It aims to determine the routes, terminals to visit, and the volume on that route, for each traffic class (origin-destination-commodity combination). Finally, the terminal policies aim to construct the strategy for classification or consolidation of traffic. Specifically, this encompasses modal shifts and assignment of traffic to specific carriers.

Caris, Macharis, & Janssens (2008) give an overview of different problems in intermodal freight transportation planning, see Table 2-3. They distinguish between four different decision makers: drayage operator, terminal operator, network operator, and intermodal operator. Further, they divide the problems over the three planning layers: strategic, tactical, and operational level. Strategic planning models deal with long term decisions. At the tactical level medium term decisions are covered, and finally at the operational planning level, models consider the short term decisions (Andersen, Crainic, & Christiansen, 2009a).

Decision maker	Time horizon		
	Strategic	**Tactical**	**Operational**
Drayage operator	Co-operation between drayage companies Truck and chassis fleet size	Allocation of shippers and receiver locations to a terminal Pricing strategies	Vehicle routing Redistribution of trailer chassis and load units
Terminal operator	Terminal design	Capacity levels of equipment and labour Redesign of operational routines and layout structures	Resource allocation Scheduling of jobs
Network operator	Infrastructure network configuration Location of terminals	Configuration consolidation network Production model Pricing strategy	Load order of trains Redistribution of railcars, barges and load units
Intermodal operator	n.a.	n.a.	Routing and repositioning

Table 2-3: Problem types in transportation networks (Caris et al, 2008).

As the drayage operator is concerned with pre- and end-haul, we do not look further into this. The terminal operator planning problems aim at optimal exploitation of the terminal and therefore is not considered either in this research. Further, the strategic level deals with long term decision making and therefore we do not consider it.

This leaves the network and intermodal operator for tactical and operational planning. For the tactical planning level this then leaves planning problems for configuration consolidation networks, production models, and pricing strategies. For the operational planning models the remaining subjects are load order of trains, redistribution of railcars, barges, and load units, and routing and repositioning (Caris et al., 2008).

An important aspect to logistics is the unknown demand across a time horizon. Unknown demand induces uncertainties into the planning, in literature we find two ways to deal with uncertainties in demand.

Pedersen & Crainic (2007) note that some quantitative knowledge for demand levels might be known by operators, but this is usually on an aggregated level. Therefore, they assume that a realistic approach can be attained through aggregation of customers in several zones. When we consider, for example, a continental train that runs between rail terminals, the demand can be forecasted on a regional level. Figure 2.3 shows two networks, on the left a conventional network where customers are represented individually, and on the right a network where customers in a close proximity to each other are aggregated into zones.

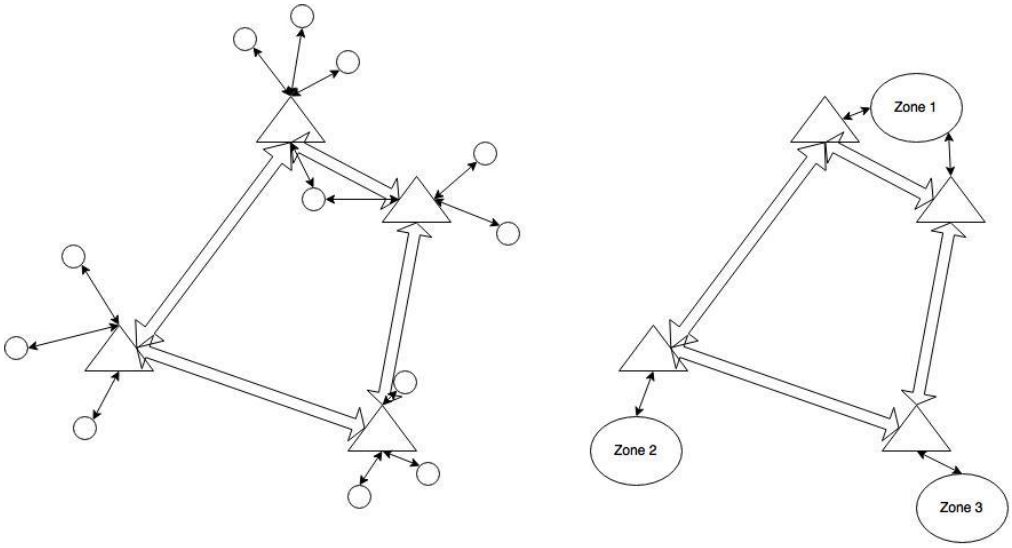

Figure 2.3: On the left a network with individual customers, and on the right a network with aggregated customers (Pedersen & Crainic, 2007).

Another practice that occurs in planning operations is ignoring the uncertainties which are present in actual operations. When uncertainties do get addressed, it is usually given as a two-phase procedure. First, a base planning is made for a standard operation period, e.g., a week. Second, during execution of the planning it is adjusted according to the circumstances happening (Lium, Crainic, & Wallace, 2009).

When considering static models, demand is usually either assumed to be routine (equal every day, week, or month) or variable throughout time, which happens when some of the shippers have weekly orders, whilst others only have orders recurring every two to three weeks. In the first case, the solution will be a recurring schedule until major changes occur in the demand pattern. The second case will require a solution that spans across a longer horizon and usually covers a demand cycle (Braekers, Caris, & Janssens, 2013).

As we indicated, the unknown demand occurs over a time horizon. To capture the time horizon of network problems, a time-space network can be applied. Figure 2.4 shows a time space network for three terminals across a five-day period. Every terminal is represented for every day on the time horizon. Arcs between two nodes representing the same terminal, imply that a vehicle is waiting for a given day. Whilst arcs between two different terminals imply movement of a vehicle (Pedersen, Crainic, & Madsen, 2009).

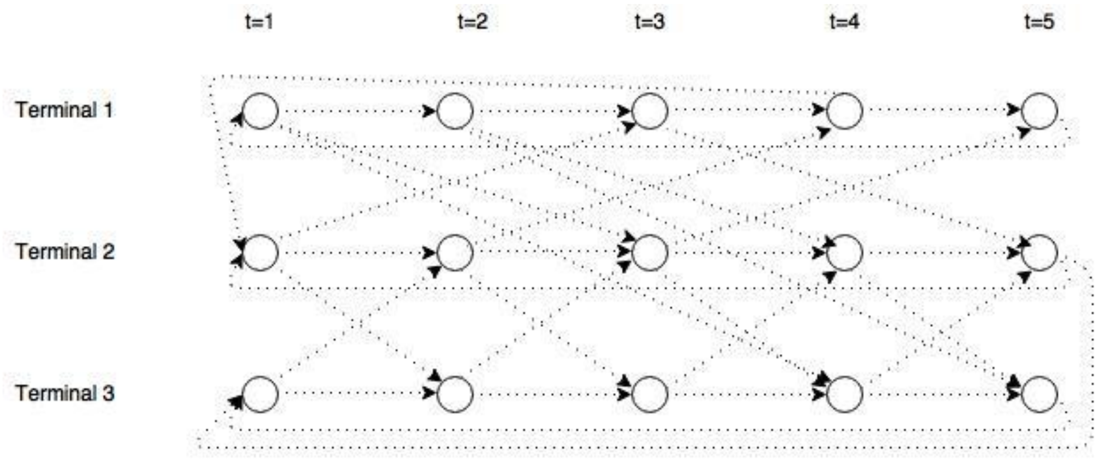

Figure 2.4: Time space network with three terminals and a five day horizon

Real-time-oriented control systems have to be able to deliver feasible and efficient execution of transportation processes. This in turn suggests that, after starting of the transportation process, the control system needs to cope with occurring disturbances in the network (Bock, 2010).

Models used for planning in the current logistic environment have to deal with a lot of uncertainty. For example new inbound orders, delayed orders or equipment breakdowns. This leads to the point that models should be flexible, stable, and robust. Flexible here means easy adjustable. Small changes to the planning should only have little effect on the planning, also known as stable. Finally, robust indicates that the overall results need to stay acceptable even when a large number of uncertain events take place (Mes, van der Heijden, & van Harten, 2007).

2.5 Models

In Section 2.4.1 we have seen that a good and common way to model consolidation policies is with SND models. Therefore we go into further detail on these types of models. Because network flow problem (NFP) models are, in essence, a simplified SND due to leaving out the fixed costs, we cover them alongside the SND models.

An interesting point described by Pedersen & Crainic (2007) is that intermodal freight trains become a viable option when a minimum utilization of 90% reached. Other than that, intermodal freight is generally viable for distances over 400 km. When distances become shorter, freight trains start getting outperformed by direct shipping, e.g., delivery by truck. As we are dealing with freight trains in our case, see Section 1.4, it is interesting to see if the statement of 90% utilization holds. Managing the order planning and consolidation of orders in a good fashion can result in attaining this 90%.

2.5.1 Service network design

SND models are used for a broad spectrum of planning problems. One of the more common fields is the tactical planning of consolidation based modal and multimodal operations (Crainic, Hewitt, Toulouse, & Vu, 2014).

SND models in general address: service selection, scheduling, blocking, train makeup, and freight routing. Service selection encompasses train route and type selection. Accompanying the selected service, a schedule is created for a certain period which is then repeated for a given planning horizon. Blocking is the act of determining the blocks to be made out of train cars which are then assigned to specific trains. Finally, freight routing specifies for each demand the route to take to transport the train from its origin to destination. SND models get addressed in, amongst others, (Song & Dong, 2013), (Zhu, Crainic, & Gendreau, 2014), and (Meng & Wang, 2011).

In SND models, goods have to be transported between two nodes. We addressed the time-space network in Section 2.4.1, this network represents services in time between nodes. When the time dimension is an important aspect of the problem, a time-space network is a logical choice. In fact, the combination of SND with time-space network is found in several works, for example, (Andersen et al., 2009a) and (Pedersen et al., 2009).

An interesting point, identified in an overview of tactical models by SteadieSefie et al. (2014), is that SND models, in special the dynamic variants, are capable of dealing with empty flows as well. Meng & Wang (2011) note that empty container repositioning should be considered simultaneously with loaded container movement when planning on a medium-term horizon.

Newman & Yano (2000) develop a model which determines schedules for both direct and indirect shipping and allocating containers to them. This is done for the long-haul rail transport part of the intermodal transport trip. The indirect shipping implies an intermediate stop at a hub for rearranging train cars. They address a short-term, finite-horizon, discrete-time scheduling problem for the long-haul session of intermodal freight transport. The set of commodities is reduced to one, and a time dimension is added. Further, the model keeps track of the amount of containers at each location. Some issues that can be identified are: the model does not deal with in- and outbound containers at the hub locations, there is no alternative transport available, and containers do not get assigned specifically to a train but instead just an indication of the amount of containers using an arc is given. The last issue is common for all networks identified. Figure 2.5 gives the representation of a two day horizon with two origins, two destinations, and one hub, which is used in (Newman & Yano, 2000).

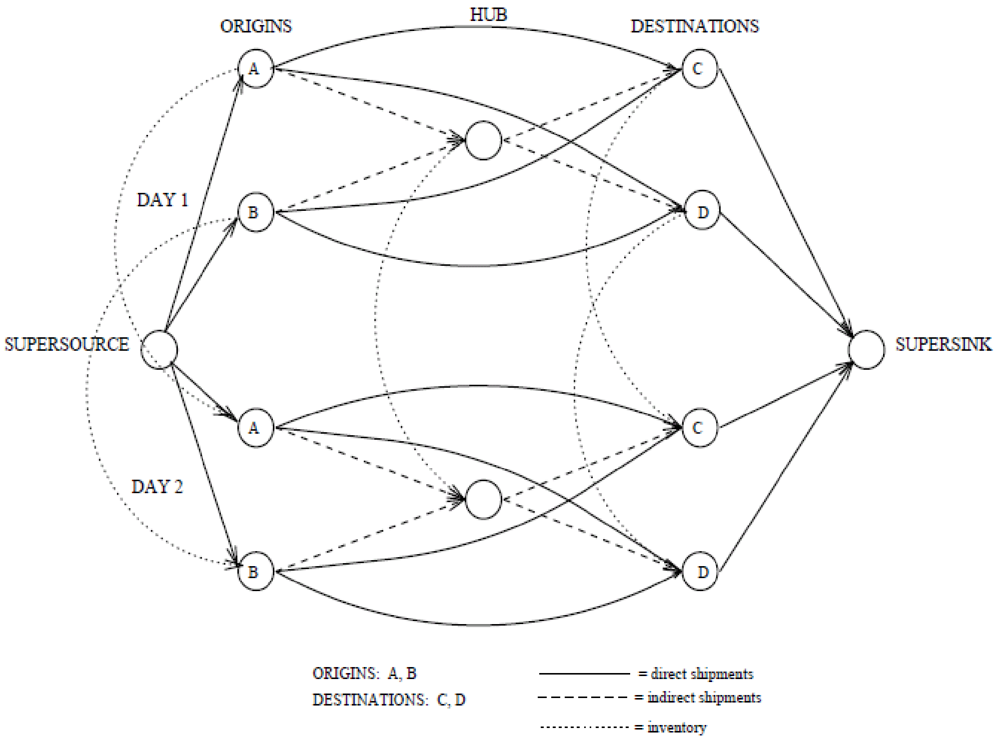

Figure 2.5: two origin, two destination, one hub multi commodity network (Newman & Yano, 2000)

2.5.2 Markov models

A way to account for the uncertainties in demand is through a Markov model. In a Markov model an optimal solution can be found by solving a recursive formula. A way to do this is, for example, considering a time period of five days and start solving (take a decision) on the last day of the horizon. The idea is, because it is the last day, the decision is easy. Then moving one step earlier into the horizon solving the problem for that day, with the knowledge of possibilities on later days.

Rivera & Mes (2015) propose a Markov model for solving dynamic multi-period freight consolidation problems. The model is applied to long-haul freight transport with uncertain demand. The aim here is to minimize the end-haul costs by selecting freights for consolidation such that only a few terminals have to be visited instead of all. Downside to the model is the fact that it is suitable for small problem instances only, again due to the nature of real life sized network problems.

2.5.3 Inventory control models

Besides the SND models found in literature, another approach for the handling of empty containers is done through inventory control policy models. Yun, Lee, & Choi (2011) consider a model for repositioning and leasing of empty containers under probabilistic demand and supply. Dang, Yun, & Kopfer (2012) present a model that combines four policies. The four policies serve different hierarchical levels within the network. First a coordinated overseas positioning policy, which determines ordering times and quantities based on the inventory position of the total network. The next policy is the distribution policy, which deals with the distribution of empty containers from the different terminals to the depots. The inland positioning policy covers interdepot positioning of containers. Finally, the leasing policy covers for the case when the other policies are not able to fulfil the demand of empty containers. This policy covers the container demand that cannot be met through the other three policies. This is done by assuming that additional containers can be leased instantaneous against additional costs. Figure 2.6 gives a graphical representation of the relation of the different policies for a network with one terminal and two depots.

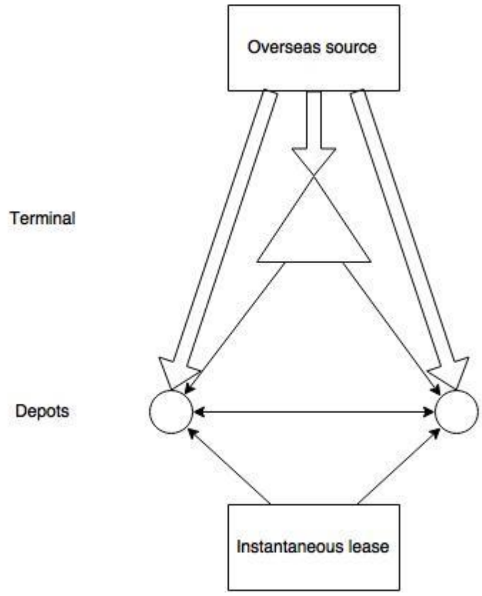

Figure 2.6: Graphical representation of the different inventory policies (Dang, Yun, & Kopfer, 2012).

2.6 Algorithms

At the tactical level SteadieSeifi et al. (2014) subdivide the solution methods for tactical planning problems under six different types: exact, approximation, heuristics, meta heuristics, hybrid heuristics, and others. Arguably approximations and heuristics are the same thing as both aim to find an as good as possible solution.

Among the exact solution methods they mention Branch & Bound, Branch & Price, and Decomposition. The aim of these is, as the solution type suggests, to solve the model to optimality. A major issue with this is the fact that with an increasing instance size, the solution time required might grow in an exponential fashion.

Feeding a solution algorithm with data that represents reality most likely results in a solution time far exceeding reasonable boundaries. Heuristics are often successful in finding a good solution (close to optimal) for solution methods that otherwise take an excessive amount of time to solve.

Heuristics can be classified as constructive and local search techniques. In a constructive heuristic, the aim is to find a good solution, which is constructed from scratch (Blum & Roli, 2003). Local search techniques are used to improve an initial solution by making small changes. Heuristics are devised for a particular problem. The risk with local search techniques is the fact that they can get stuck in a local optimum. An answer to this problem are the so called meta heuristics. These meta heuristics can be applied, independent of the problem type, and can guide a heuristic to step out of a local optimum (Blum & Roli, 2003). Some well-known meta heuristics are tabu-search and simulated annealing, we explain these in Section 2.6.1.

Hybrid heuristics combine elements of several meta heuristics into one heuristic. As this requires identification of more meta heuristics and their characteristics before we are able to apply this, we do not go into further detail on this.

In the last category SteadieSeifi et al. (2014) identify, amongst others, simulation. Simulation aims to capture the reality in a model, based on some simplifying assumptions. The simulation model is then used to gain managerial insights for decision making. Simulation is extensively addressed in Chapter 5.

In the remainder of this section we address several solution methods. As we mentioned, Section 2.6.1 covers some well-known meta heuristics. Section 2.6.2 addresses approximate dynamic programming. Next, in Section 2.6.3 we introduce linear programming. Finally, we give a short introduction to simulation in Section 2.6.4.

2.6.1 Meta heuristics

Tabu search generally consists of two phases: an initialization phase and an iteration phase. The initialization phase gives an initial solution for the tabu search heuristic to use. This initial solution often stems from a heuristic procedure. After an initial solution is created the iteration phase can be performed. In the iteration phase all neighbourhood solutions of the current solution are considered. From these neighbourhood solutions the best solution that currently is not on the tabu-list is selected as the current solution. The tabu-list contains recently made changes. The new solution does not have to be better than the previous solution. The change in the solution is then added to the tabu-list, e.g., a swap of two orders. A solution drops from the list after a predetermined amount of new neighbourhood solutions are added to the list. This iteration phase keeps repeating itself until either one of two stopping criteria is reached. The first criteria is the total amount of iterations performed, whilst the second one occurs when a predetermined maximum solution time has expired. During the whole execution, the best found solution is remembered.

Simulated annealing is another meta heuristic capable of improving an initial solution. Just as tabu search, simulated annealing consists of an initialization and iteration phase. In the initialization phase a random solution to the problem is created. Next, during the iteration phase a small adjustment is made to the current solution such. This results in a nieghboursolution. At this point there are two possibilities, either the solution is better and gets accepted, or it is worse. In the latter case the new solution gets accepted with a certain probability. Throughout the execution of simulated annealing it becomes more unlikely that worse solutions get accepted.

Both tabu search and simulated annealing are so called meta-heuristics. They aim to find a good solution in a short period of time. A major advantage of these two meta heuristics is the fact that they are able to escape from local optimum solutions. A local optimum is a solution from where all neighbour solutions are worse than or equal to the current solution.

2.6.2 Approximate dynamic programming

A downside to most (meta) heuristics and mathematical programming techniques is that they have issues with dealing with stochastic settings (Rivera & Mes, 2015). Dynamic programming (DP) can handle this, but does not scale well. Another solution is approximate dynamic programming (ADP). ADP aims to solve the three curses of dimensionality found in DP. These three curses are the sizes of: state space, outcome space, and action space (feasible region) (Powell, 2007).

The algorithmic strategy for ADP is based on stepping forward in time. This is contrary to DP that typically steps backwards in time. DP assumes that the value function at a later point on the time horizon is known and computes the value function for an earlier point in time. ADP on the other hand learns new information by stepping forward in time where it simulates transitions. Several iterations of this forward pass are performed throughout the horizon.

2.6.3 Linear programming

A common method for solving problems to optimality is Linear Programming (LP). LP models aim to minimize or maximize an objective, often costs or profits. This is done by formulating an objective function, based on a set of decision variables. Decision variables give the decisions that have to be made, e.g., how much money to invest or how many products to manufacture. The value of the decision variables in the objective function are bounded by a set of constraints. Constraints can, for example, state that no more than twenty units of product 1 can be produced each week. In a standard LP, the decision variables can take any real number. This is usually restricted by sign restrictions. These restrictions state, for example that a decision variable has to be greater than or equal to 0. This can be taken further by restricting variables to only take integer values, creating so called integer linear programs (ILP) or mixed integer linear programs (MILP). Most problems require either ILPs or MILPs. ILPs and MILPs are of NP-hard resulting in exponential increase computation time.

Small LP models, with few variables, can often easily be solved by hand. However, it is not uncommon for LP models to contain hundreds of variables, solving such a model by hand becomes cumbersome, if not an impossible task. To counter this, an LP solver, for example CPLEX, is often applied. An advantage of using LP models is that it is possible to find an optimal solution.

2.6.4 Simulation

Several models that SteadieSeifi et al. (2014) identified get “solved” through usage of simulation. Simulation can be used to simulate several scenarios, these scenarios can be different decision policies. By simulating these different policies, their performance can be compared. This can then be used to gain managerial insights into which policy performs the best under the given conditions.

Next to that, simulation can be applied when decision makers want to get insight into performance of a new, yet to implement scenario. This scenario can be anything: a decision policy, a new warehouse, etc. Usually this is applied when the costs for actual testing are too costly and the effects of the scenario are not known upfront, e.g., expected performance of a warehouse that has not been built yet.

2.7 Framework for articles

In Table 2-4 we show an overview of types of models we found in literature, we indicate whether they deal with consolidation, empty flows, or both. Further, we present the objective, special considerations, and the specific solution method applied.

Article by Author(s)	**Model**	**Consolidation**	**Empty Flows**	**Objective**	**Special consideration**	**Solution Method**
Newman & Yano (2000)	NFP	Yes	No	Minimize fixed, holding, handling, and variable transportation costs	Tracking of container inventory	Decomposition algorithm
Lin & Chen (2004)	SND	Yes	No	Minimize fixed and variable transportation costs	Time-constraints	Enumeration algorithm
Pedersen & Crainic (2007)	SND	Yes	Yes	Minimize fixed and variable transportation costs	Value for time-costs	ILP model with linear solver
Pedersen et al. (2009)	SND	Yes	Yes	Minimize fixed and variable transportation costs	Design balance constraints	ILP model with Meta heuristics (Tabu search)
Lium et al. (2009)	SND	Yes	Yes	Minimize costs for operating vehicles	Non-linear, no costs for moving goods, LTL trucking	Monte-Carlo like simulation, scenario trees
Andersen et al. (2009a)	SND	Yes	Yes	Minimize fixed and variable transportation costs	Asset management	ILP model with linear solver, relaxation
Yun et al. (2011)	Inventory control	No	Yes	Determine inventory control policies, to minimize costs for holding, leasing and ordering	-	Simulation
Meng & Wang (2011)	NFP	Yes	Yes	Minimize total operating costs	Liner shipping	ILP model with linear solver
Dang et al. (2012)	Inventory control	No	Yes	Determine inventory positioning policies, minimizing expected total costs	-	Simulation
Braekers et al. (2013)	SND	Yes	Yes	Maximize profit for transporting loaded and empty containers	-	ILP model with linear solver
Song & Dong (2013)	SND	Yes	Yes	Minimize total operating costs	Liner shipping	Decomposition algorithm and heuristic method
Crainic et al. (2014)	SND	Yes	Yes	Minimize fixed and variable transportation cost and costs for using resources	Resource management	ILP model with linear solver
Rivera & Mes (2015)	Markov model	Yes	No	Find the best policy to minimize expected costs	Alternative transport option	Approximate dynamic programming

Table 2-4: article, subject, solution method matrix

From literature it becomes apparent that in most cases, network problems for transportation and scheduling are modelled through SND. All of the models determine the number of containers transported per arc, whilst not specifically assigning a container to a transport mode. When we look at models dealing with both consolidation and empty flows, the models proposed in, amongst others, (Andersen et al., 2009a) and (Pedersen & Crainic, 2007) are the most promising. In these two works, the network models are solved through formulating the network as an ILP.

2.8 Conclusions

From the literature on synchromodal planning we conclude that synchromodality is the act of planning and scheduling of intermodal transportation in a flexible fashion based on real-time information, where main drivers are costs, throughput time, and sustainability. Furthermore, besides the possibility of transportation by truck additional transport options should be available, e.g., barges or trucks.

When considering consolidation, the utilisation rate should be as high as possible. A breakeven point for intermodal freight trains, found in literature and confirmed in practice, is a utilization rate higher than 90%. When loading containers onto a transport mode the number of containers loaded onto the mode should be maximised whilst not exceeding the maximum allowable weight. We should strive for a utilization of at least 90% of maximum container capacity.

The models that we found in literature, did not cover replanning of orders. If the models that we found have to perform replanning, it implies that they have to be fully executed again. This is not a desirable situation as we want to only replan the part of the planning that is influenced by a new order. Even though the models found in literature are not entirely suitable, concepts of the SND models (Andersen et al., 2009a) and (Pedersen & Crainic, 2007) are suitable

As one of the requirements for the algorithm is to get solutions in a reasonable amount of time, a couple of seconds, an exact solution is not an option. Therefore we construct an algorithm tailored for Seacon.

In Chapter 3 we focus on the case and data analysis, the results are used to validate the model and to perform experiments with scenarios derived from the case study.

3. Case study

In this chapter we focus on the analysis of the case and relevant data. We answer sub-questions related to the case study. In Section 3.1 we focus on the characteristics of the case. In Section 3.2 we analyse and quantify the data set corresponding to the east-west corridor Poland – Germany – United Kingdom. Finally, we conclude this chapter in Section 3.3.

3.1 Case study

Before we are able test the performance of our planning heuristic on the case, we analyse the important aspects that are considered at Seacon for this case. We give a description of the case in Section 3.2.1. Section 3.2.2 introduces the equipment. Finally, Section 3.2.3 gives the train schedules.

3.1.1 Case description

Currently Seacon is working on setting up a rail transport service between the United Kingdom and Poland. This rail transport service is positioned between different terminals in UK and Poland and with an intermediate hub in Germany. Figure 3.1 graphically represents part of this network. The case is focussed on the terminals 1 to 6 together with the corresponding connections. Furthermore this network shows potential feeder terminals and a short sea trip. These options can be considered when the initial project is a proven concept.

As we address in further detail in Section 3.2.1, there exists an imbalance in the network. The number of orders transported from Poland to the United Kingdom is larger than the return flow of orders.

Seacon wants to gain insight into this network such that they are able to exploit this network to its fullest potential, against lowest possible costs whilst meeting all of the delivery due dates for containers. To do this, several decisions have to be considered, their effects analysed, and eventually combined into an optimal decision policy for the assignment of containers to shipments.

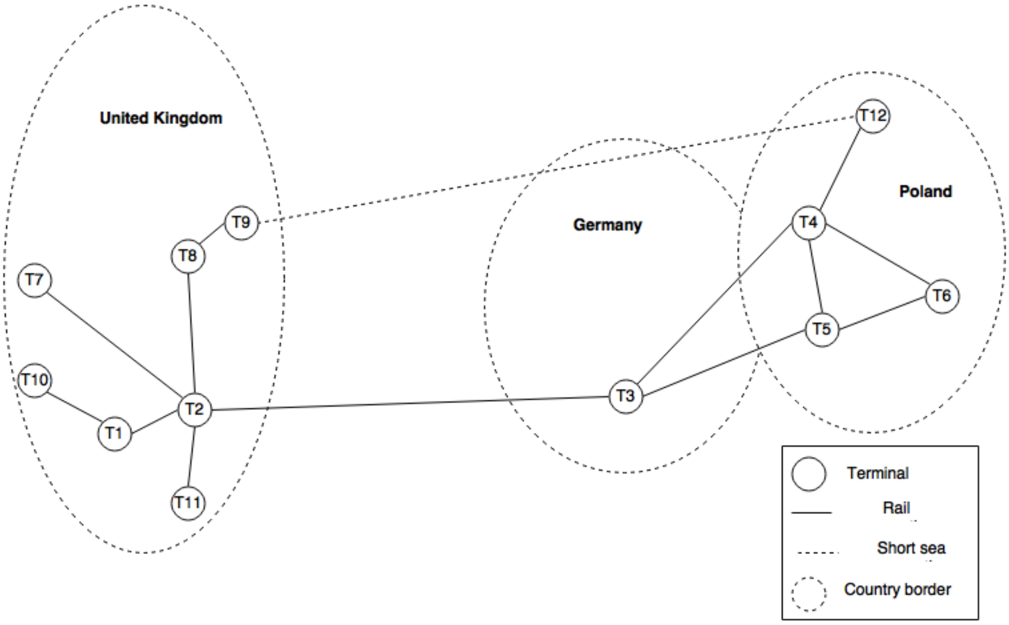

Figure 3.1: network

3.1.2 Equipment

The containers that we consider in this case are two different types of 45 ft. containers: a 45 ft. dry box and a 45 ft. reefer. The main difference between these two types of containers is that the reefer container is an insulated container with climate control, meaning that it is suitable for transportation of perishable goods such as fruits and vegetables. This in turn means that goods that can be transported in a dry box, can also be transported in a reefer. However, this does not work vice versa. Naturally, as the reefer container is an advanced type of container, the costs for using it are higher than for the dry box container. Table 3-1 gives an overview of specifications per container type.

Specification	**45 ft. dry box**	**45 ft. reefer**
Transportable goods	Limited	All
Maintenance costs	Low	High
Dry weight	4,260 kg[1]	6,995 kg[1]
Max. Gross weight	34,000 kg[1]	34,000 kg[1]

Table 3-1: Specification difference between dry box and reefer containers.

Further, for exploitation of the network, Seacon uses flatcar wagon sets. A wagon set consists of a number of flatcars, each flatcar is approximately 90 ft. long, giving room to two 45 ft. containers. A full wagon set consists of 20 flatcars, bringing capacity to a total of 40 containers.

A total of 2 wagon sets and 240 containers are rented. Besides the maximum number of containers, we also have to deal with a maximum allowed weight which forms the second capacity restriction. Depending on the locomotive used, this is different. In our case the locomotives that are intended to be used, are allowed to pull 1,600 tons as a maximum weight. The wagon set is assumed to weigh 571 tons. This leaves a total of 1,029 tons for freight and container transport.

3.1.3 Train schedules

For our case trains are running between the United Kingdom and Poland. As we already found in the literature, it is common for intermodal networks to have fixed train schedules, see Section 2.2.2. For our case this is also true. Table 3-2 gives the train schedule used for the case, which is currently a fictive schedule. The schedule is based on a schedule using the same route but without an intermediate stop. From the schedule it becomes clear that a roundtrip starting in Poland, going to the United Kingdom and returning takes approximately 5 days.

Important details in the schedule are the locomotive changes and the arrival times at border crossings. The locomotive changes are required due to the many different track settings between countries. A time slot for crossing is booked and when this is not made in time, delays of up to 24 hours are possible. The four different times mentioned in the table (arrival, slot time begin, slot time end, and departure) represent the estimated time of arrival, the start and end time for the loading and unloading of the train, and the estimated time of departure.

[1] Source: (Unit 45, 2015)

Roundtrip schedule	Day 1	Day 2	Day 3	Day 4	Day 5	Special event
Slottime begin terminal 5	09:00					
Slottime end terminal 5	14:00					
Departure terminal 5	14:30					
Arrival border crossing 1	23:30					Locomotive change
Departure border crossing 1		00:00				
Arrival terminal 3		12:00				
Slottime begin terminal 3		13:00				
Slottime end terminal 3		15:00				
Departure terminal 3		16:00				
Arrival border crossing 2		19:00				Locomotive change
Departure border crossing 2		20:00				
Arrival border crossing 3			02:00			Locomotive change
Departure border crossing 3			03:30			
Arrival terminal 2			06:00			
Slottime begin terminal 2			06:00			
Slottime end terminal 2			12:00			
Departure terminal 2			23:00			
Arrival border crossing 3				01:30		Locomotive change
Departure bordercrossing 3				06:30		
Arrival border crossing 2				12:30		Locomotive change
Departure border crossing 2				13:30		
Arrival terminal 3				16:30		
Slottime begin terminal 3				17:30		
Slottime end terminal 3				19:30		
Departure terminal 3				20:30		
Arrival border crossing 1					08:30	Locomotive change
Departure border crossing 1					09:00	
Arrival Terminal 5					16:00	
Slottime begin terminal 5					16:00	
Slottime end terminal 5					21:00	

Table 3-2: Event and time table.

3.2 Data analysis

This section focusses on the data that corresponds to the case and a generic network. A generic network is a random generated network that shares aspects with the case. We do this to gain more insight into the performance of our planning heuristic because the case network alone is too specific to base conclusions on. Section 3.2.1 examines how the orders are divided across the network, Section 3.2.2 then focusses on the arrival process of orders. The weight and lead time characteristics are analysed in Sections 3.2.3 and 3.2.4 respectively. In Section 3.2.5 we examine the price structure corresponding to the network. In Section 3.2.6 we look at inter-terminal truck times. Finally, we perform experiments on generic networks. Data covered in the before mentioned sections does not encompass all insight needed for a generic network. Therefore we address this in Section 3.2.7.

3.2.1 Data set

For the case we use the order dataset on the east-west corridor Poland - Germany - United Kingdom from the year 2014. A total of 81 origin-destination lanes are identified for that year, each of these consisting of one or multiple orders. A lane corresponds to the total transport route from sender to receiver. The total number of orders identified is 11,024, where each order encompasses one container. These orders are assigned to six different terminals spread across the three countries. Figure 3.2 shows how the 11,024 orders are divided per country. This is done for both in- and outbound orders (receiving and sending markets). The figure clearly shows a trade imbalance. Figure 3.3 shows the in- and outbound orders on a terminal level, from this we can see that terminals 1, 2, 5, and 6 are the major contributors to the imbalance.

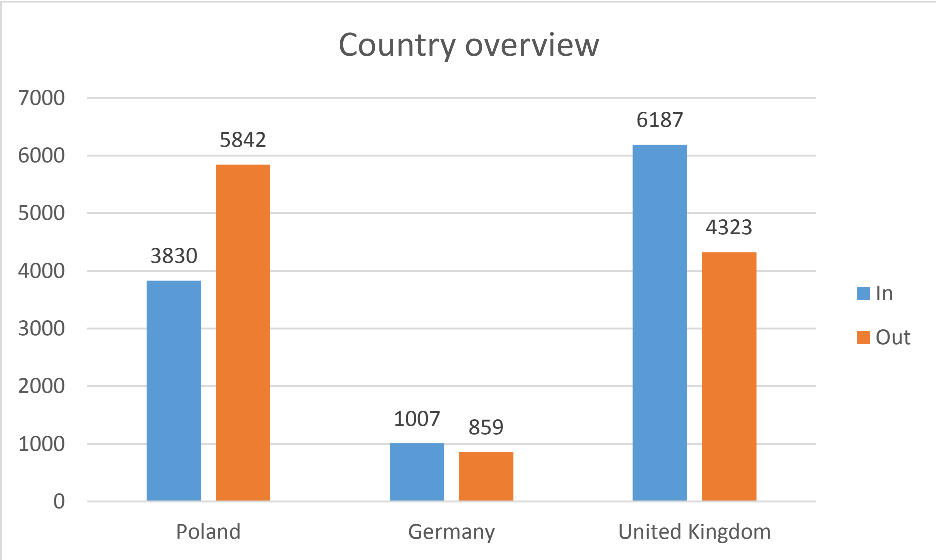

Figure 3.2: Overview of in- and outbound containers per country, in 2014.

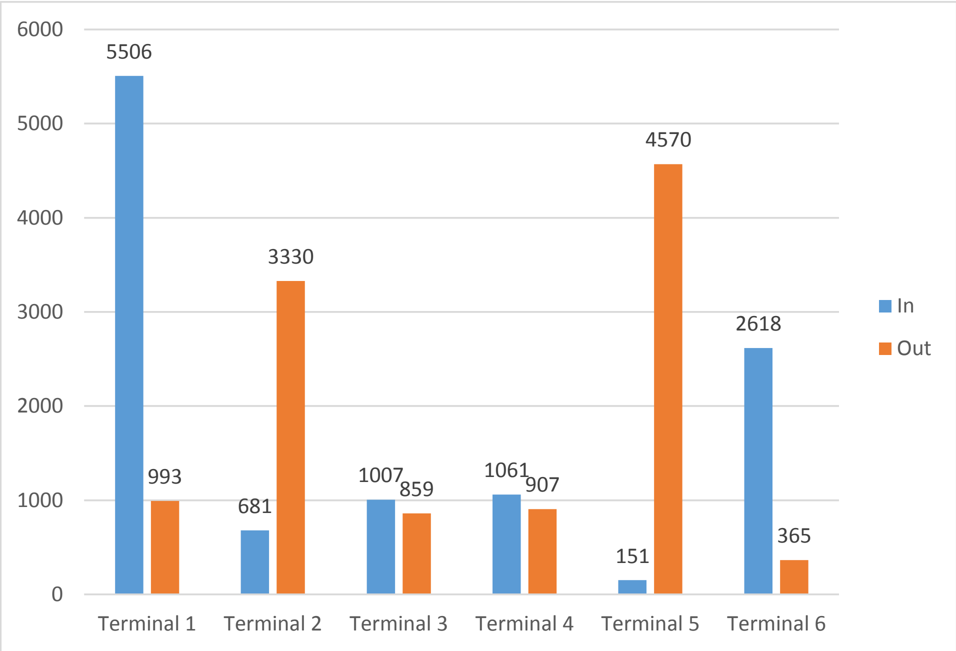

Figure 3.3: overview of in and outbound containers per terminal in 2014

3.2.2 Arrival process

Figure 3.4 gives an overview of how the orders are divided over the months. We notice a clear distinction in number of orders per month in the periods December until June and July till November. In the former group, each month has a total number of orders less than 900, whilst the latter group each month has total orders over 900. This difference is significant, see Appendix B. The different seasons are further referred to as low and high season. The major reason that can be identified for this difference in seasons comes from the fact that goods are getting stocked for the Christmas holiday season. This starts already in July.

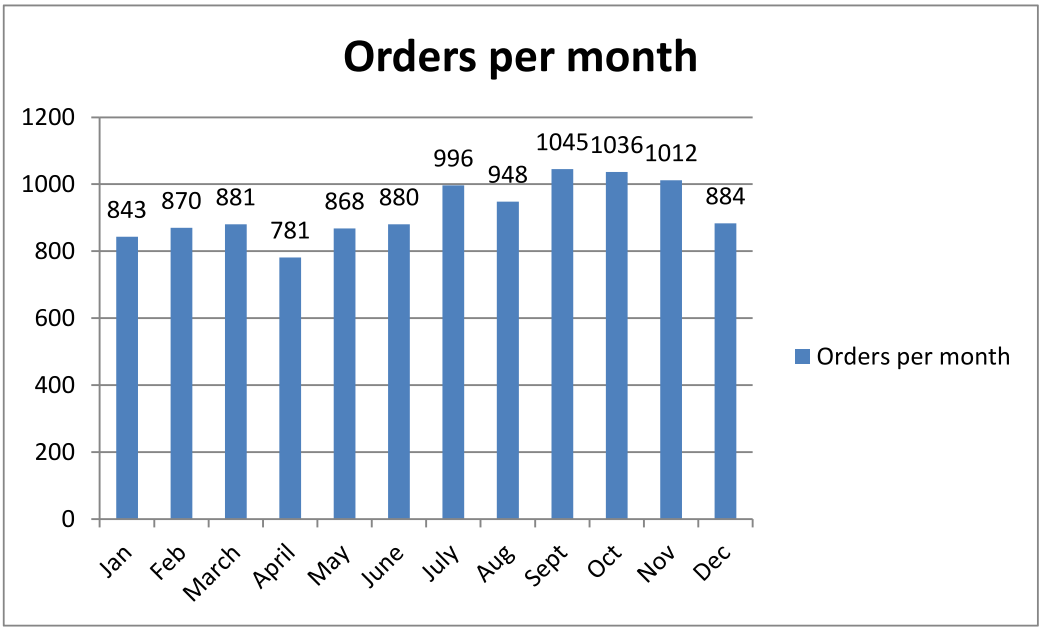

Figure 3.4: Orders per month

Now that two different seasons are established, we determine the arrival processes for each of the seasons. We want these processes for our experimental settings in our simulation study.

We assume that inter arrival time between two orders are independent. The average number of order arriving across the entire network, per month for the low and high season are 858.1 and 1007.6. The arrival of orders are stochastic events. We test whether a Poisson process is a valid representation for the order arrival process. We find no statistical evidence against this, see Appendix B for details.

3.2.3 Door-to-door lead times

The door-to-door lead time of an order is the time from the moment an order can be picked up at the shipper until its latest allowed point in time that it can be delivered at the receiver. According to the experts, the door-to-door lead time for orders is dependent on the season. In the low season a smaller door-to-door lead time is expected than in the high season. This is logical due to the fact that in the low season carriers have more available resources to transport goods and thus there is more room available for competition. Table 3-3 gives an overview of demanded door-to-door lead times during low and high season.

	Min (days)	Max (days)	Average (days)
High season	7	9	8.5
Low season	7	8	7.5

Table 3-3: demanded door-to-door lead times.

In the absence of data it is difficult to determine a good distribution for the door-to-door lead times. To be able to fit a distribution to the data in Table 3-4, we apply a heuristic procedure for selecting a distribution in the absence of data. The heuristic is described by Law (2007), see Appendix C. For the door-to-door lead-time, we use a triangular distributions with parameters corresponding to the min, max, and average shown in Table 3-4.

Because the door-to-door lead times include pre- and end-haulage, we need to determine the lead time for the long-haul transport. Often the origin and destination terminals are based on the distance to the pickup and delivery points. This is to reduce pre- and end-haul distances. Table 3-4 summarizes the data with regard to pre- and end-haul. At Seacon the average speed that is assumed for pre- and end-haul trucking is 60 km/h. Therefore we assume the same average trucking speed.

	Min	Max	Average
Distance (km)	1	375	82.38
Time (hours)	0.02	6.25	1.37

Table 3-4: Pre- and end- haul distance and time.

For the pre- and end-haul times we fit a Weibull distribution with parameters α=1.1436 and β=1.4873, see Appendix D for the data fitting.

3.2.4 Order weight

In an intermodal network, different types of containers are transported, one of the characteristics that we identified is the weight of containers. The weight can vary due to multiple factors. These factors are, but not limited to, size, type, and load of the container. Because weight restrictions are in place, insight in to the weight distribution of containers is needed.

The weight distribution is determined by consulting experts. Table 3-5 gives the overview of the loads. There is a noticeable difference between refrigerated goods (requiring a reefer container) and non-refrigerated goods (can be loaded in either type of container). Further we notice that the maximum weight of 34 tons (container and freight) is never reached.

	Min (tons)	Max (tons)	Average (tons)
Refrigerated goods	19	22	21
Non-refrigerated goods	18	25	24

Table 3-5: Weight spread for loads.

Just as with the lead time distribution, we only have a range for the weight distribution. Following the same procedure as with the lead time, we decide on triangular distributions for the weight distributions. The parameters correspond with Table 3-6.

3.2.5 Price structure

For the case network that we consider, several sources that contribute to the total costs can be identified. We subdivide these costs into transport modes, containers, and additional costs.

We identify three different modes and three different cost aspects. These modes are: truck, barge, and train. However, for our case we do not consider barge transport and therefore leave these costs out. The aspects are fixed tariff, variable tariff, and cost for using a leg. Table G-1 in Appendix G gives an overview of the costs per for both truck and train.

The prices for a trips between terminal 2 and either terminal 4 or terminal 5 are not based on a price per leg. In essence if a train departs from terminal 4 or terminal 5, it returns there as well, creating a round trip. If the train departs the whole trip is payed for. As we have two departure points, we also have two different tariffs. The departure tariffs for terminal 4 and terminal 5 are shown in Table G-2 in Appendix G.

For the containers we look at the two types we discussed in Section 3.2.2. For these containers we identify two cost aspects. Table 3-6 gives an overview of the aspects the corresponding costs. It is noteworthy that, in accordance with terminal operators, holdings costs are paid for storage after a fixed amount of days has passed.

	Variable tariff (per transport)	**Holding costs at terminal (per day)**
Dry box	€30	€5.50 to €11.50
Reefer	€250	€5.50 to €11.50

Table 3-6: Costs for containers

The additional costs that we consider in a network are the terminal handling costs. For terminal handling costs there are two different types identifiable. The first type is handling of a container that is transferred from rail-to-rail, where a container is transferred from one train to another. The second type is the transferring of a container from rail-to-truck. This consists of transferring a container from the train to a truck and vice versa. This container can either be loaded or empty. Table G-3 in Appendix G gives the overview of these costs.

3.2.6 Distance between terminals

Next to train transport there is the option to transport goods by truck. For this a fixed charge with a variable km tariff is applicable. Table 3-7 gives the inter-terminal distances. Table 3-8 gives the corresponding times required for these distances. For the pre- and end-haul transport an average speed of 60 km/h is assumed. This is due to the fact that pre- and end-haulage takes place in more rural areas. For the long-haul transport an average speed of 68 km/h is assumed as most of the transport occurs on the highway.

	Terminal 1	Terminal 2	Terminal 3	Terminal 4	Terminal 5	Terminal 6
Terminal 1	0					
Terminal 2	157	0				
Terminal 3	727	578	0			
Terminal 4	1483	1328	756	0		
Terminal 5	1497	1351	779	182	0	
Terminal 6	1758	1602	1031	292	359	0

Table 3-7: Inter-terminal distances in km.

	Terminal 1	Terminal 2	Terminal 3	Terminal 4	Terminal 5	Terminal 6
Terminal 1	0					
Terminal 2	2.3	0				
Terminal 3	10.7	8.5	0			
Terminal 4	21.8	19.5	11.1	0		
Terminal 5	22.0	19.9	11.5	2.7	0	
Terminal 6	25.9	23.6	15.2	4.3	5.3	0

Table 3-8: Inter-terminal time table in hours.

It is important to note that the truck times in Table 3-8 only consists of *driving* times. Regulations with regard to truck driving state that truck drivers are allowed to drive 9 hours a day, and only 4.5 hours without a break of 45 minutes (European Parliament, 2006). If the total driving time exceeds 9 hours a longer break is required, the rest period should then at least be 11 hours. Adding this information gives the updated timetable shown in Table 3-9. The times presented here are the *minimum driving* times required. The latest possible dispatch time, while ensuring that container arrives before its due date, is then given by deducting the values from the due date.

	Terminal 1	Terminal 2	Terminal 3	Terminal 4	Terminal 5	Terminal 6
Terminal 1	0					
Terminal 2	2.3	0				
Terminal 3	21.7	9.3	0			
Terminal 4	43.8	41.5	22.1	0		
Terminal 5	44.0	41.9	22.5	2.7	0	
Terminal 6	48.6	46.3	26.9	4.3	6.0	0

Table 3-9: Inter-terminal time table with breaks in hours.

3.2.7 Generic network data

To be able to create generic networks (which we introduced in Section 3.2), the before mentioned data is not sufficient. More specific, information with regards to costs per km for truck and barge are needed. Van Heeswijk, Mes, Schutten, & Zijm (2014) and Spikker (2014) apply normalized costs for transportation of freight on the three modes truck, train, and barge. Contrary to the case network we do look at barge transportation in the generic networks. We incorporate barges due to the fact that in generic networks transportation by barges does occur. Table 3-10 gives these normalized transport rates.

Transport mode	Transport rate (per container)
Truck	1.00 per km
Train	0.51 per km
Barge	0.33 per km

Table 3-10: Normalized transport rates.

Further, the transhipment costs between three different modes is not addressed for generic networks in Section 3.2.5. Spikker (2014) found that the transhipment from barge or train to any of the modalities besides truck have equal rates. Further, transhipment from and to truck from any of the modalities also holds the same rates. Based on the costs found in Section 3.2.5 we present the transhipment costs in Table G-4 in Appendix G.

Although reduction in carbon dioxide emissions is not the main criterion on which we optimize, it can predominate decision making in the case that both costs and throughput times are minimized. Den Boer, Otten, & Van Essen (2011) performed a study in which various transport modes are compared on an EU scale. One of the parameters in this study is carbon dioxide emissions. Table 3-11 gives the values for carbon dioxide emissions for transportation of a container. In the case of truck transport, we distinguish between a loaded and an empty container.

	Loaded (gram CO2/km)
Truck (loaded)	819
Truck (empty)	558
Train	322.8
Barge	365.9

Table 3-11: Carbon dioxide emissions in gram/km per 2TEU container (den Boer, Otten, & van Essen, 2011).

3.3 Conclusions

In this chapter we described the case and quantified the data, with a focus is on transport on the east-west corridor Poland – Germany – United Kingdom. We considered the orders for a whole year, giving a total of 11,024 orders. While analysing the order set, several issues arose. First, there is an imbalance in transport between east and west. Terminals 1, 2, 5, and 6 are the major contributors for this imbalance. Going into more detail, the imbalance of container type becomes clear. Transport to the United Kingdom is mainly done in reefers, whilst transport to Poland is mainly done in dry boxes. One of the ways this imbalance is addressed is through the allocation of different rates for east and west bound transport.

For transportation we provided the cost structures, with train lanes having a fixed rate whilst truck lanes having a variable rate. Orders are defined through their characteristics: arrival rate, origin, destination, lead times, and weight. With regards to arrival rate, we identified two different seasons, a high and a low season. Finally, for generic networks we defined cost rates for the transportation modes and terminal handling charges as well as emission rates for the transportation of 2TEU containers.

4. Planning model

This chapter explains our modelling choices and the model design. In Section 4.1 we cover assumptions for the planning model. Next, Section 4.2 provides a detailed description of the problem. We introduce our planning heuristic in Section 4.3. Finally, we conclude this chapter in Section 4.4.

4.1 Assumptions

In this section we cover our assumptions for the model. Some of the assumptions are required to cover for aspects that do not fall within the scope (e.g. pre- and end-haul transport) of our research but are essential. Others serve to deal with issues that would cause for too much detail in our model. The assumptions that we identify are:

1. An unlimited number of trucks is available.
2. Terminals are open 24/7.
3. The tractionair is responsible for the actual transportation. Therefore we use static travel times.
4. Containers that have been delivered to their destination terminal, become available again after a fixed period.
5. There are no weight restrictions on transport modes.
6. Loading and unloading of containers on the different transport modes happen instantaneously.

Assumption 1 ensures that every order is deliverable through direct trucking if no intermodal options are available or intermodal transport results in delivery after the due date. Assumption 4 is in place to cover for the time period that is needed for the pre- and end-haul before containers become available again. Because our model does not cover the pre- and end haulage during which containers are unavailable for long haul transport.

4.2 Problem description

This section covers a detailed description of our problem. Section 4.1.1 gives an explanation on why the problem is relatively difficult. Section 4.1.2 covers the indices that we use. Next, Section 4.1.3 introduces the decision variables applicable in our problem. The parameters used are then introduced in Section 4.1.4. Section 4.1.5 gives our objective function and finally, in Section 4.1.6, we present the constraints for our problem.

4.2.1 Difficulty

The problem that we aim to solve is relatively difficult. It is hard to align in- and export loads across a network in such a way that there is always equipment available to execute transport within allowable time whilst minimizing the total costs incurred (or maximizing profit). This stems from a multitude of factors: uncertainty and imbalance in demand, limited equipment, uncertainty in transport times, and mode (or service) choices. To illustrate the mode choices, consider the network from Figure 4.1. In this figure, there are three terminals and two transport modes. If all route possibilities from A to B are allowed, there are six different routes, see Table 4-1. The route set in the example is extendable to multiple ways. Furthermore, we could include additional transport modes, e.g., a barge. Therefore we define a route as the transport from start to end terminal using one or multiple legs. It is possible that between each terminal pair multiple legs exist, the difference between these legs is then the transport mode used.

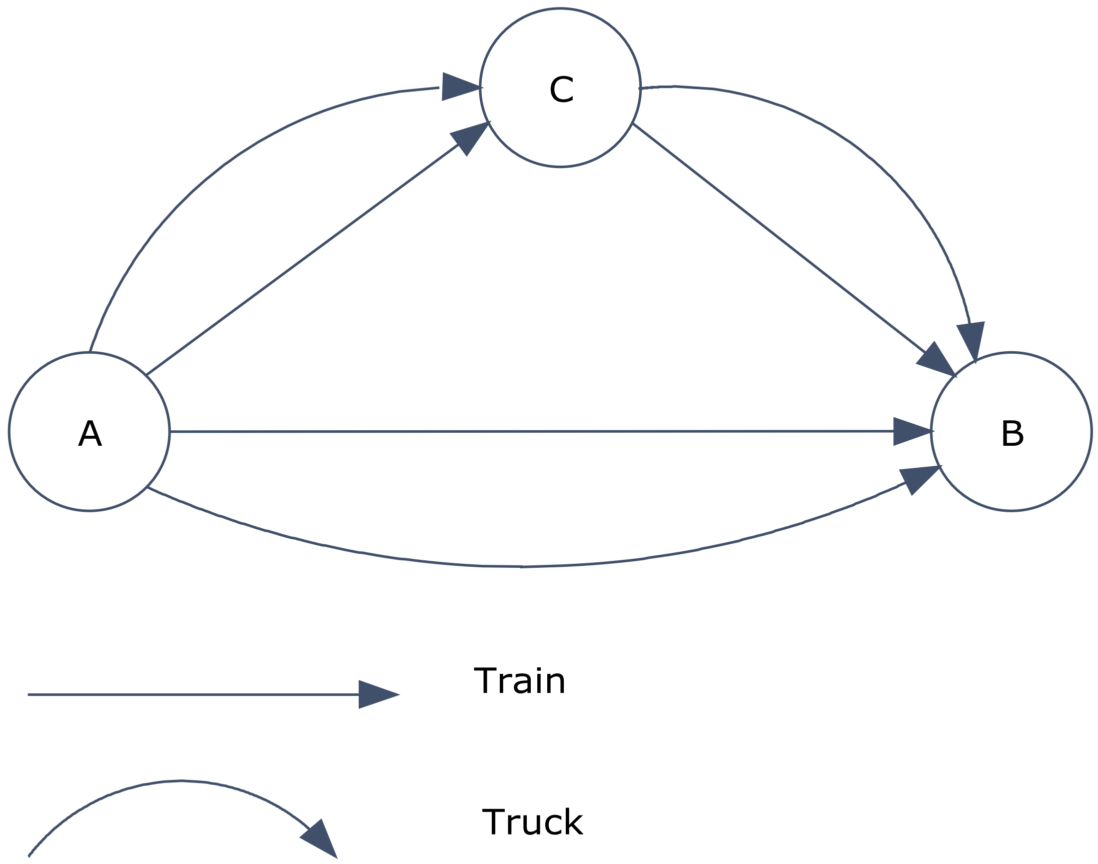

Figure 4.1: Three terminal, two mode network with possible routes.

	Leg 1	Leg 2
A to B	Train	Not applicable
A to B	Truck	Not applicable
A to B via C	Train	Train
A to B via C	Train	Truck
A to B via C	Truck	Train
A to B via C	Truck	Truck

Table 4-1: Possible routes from A to B.

As established in Chapter 2, service choices are one of the aspects in service network design (SND) models. We found several SND models dealing with asset balancing. These two aspects occur in our research as well. Therefore we believe that an SND model is a good representation of our problem. Most of the SND models in literature are formulated as an Integer linear programming (ILP) model, see Appendix A for an example. A common practice with regards to solving SND models with ILPs is representing the network as a time-space network. A problem that the models do not cover is replanning of orders. Replanning of orders is hard, if not impossible, to realize in an ILP. Replanning would imply that the ILP has to be solved entirely again upon replanning. Therefore, we resort to heuristic procedures.

4.2.2 Indices

a : $containers$ $\{1,2,..\}$

b : $orders$ $\{1,2,..\}$

i : $terminals$ $\{1,2,..\}$

k : $routes$ $\{1,2,..\}$

l : $legs$ $\{1,2,..\}$

m, n : $modes$ $(train, barge, truck)\{1,2,3\}$

4.2.3 Decision variables

In our problem we define four different decision variables.

$$X_{bka} = \begin{cases} 1 : order\ b\ is\ assigned\ to\ intermodal\ route\ k\ in\ container\ a \\ 0 : else \end{cases}$$

$$Y_{ka} = \begin{cases} 1 : empty\ container\ a\ is\ assigned\ to\ intermodal\ \ route\ k \\ 0 : else \end{cases}$$

$$Z_{ba} = \begin{cases} 1 : order\ b\ is\ assigned\ to\ a\ direct\ truck\ in\ container\ a \\ 0 : else \end{cases}$$

$$Z_{b} = \begin{cases} 1 : order\ b\ is\ assigned\ to\ a\ direct\ truck\ without\ container \\ 0 : else \end{cases}$$

Each order is assigned to either an intermodal route or a direct truck but never both. Therefore for each order either Z is equal to 1 or X is equal to 1. Intermodal routes are all transport routes with the exclusion of direct trucking between origin and destination.

4.2.4 Parameters

We identify the following parameters.

c_k^r : $costs\ for\ using\ route\ k, including\ terminal\ handling\ costs$

c_i^h : $costs\ per\ day\ for\ holding\ container\ at\ terminal\ i$

c_b^d : $costs\ for\ delivering\ order\ b\ after\ its\ due\ date$

c_a^c : $costs\ for\ using\ container\ a\ for\ transport$

o_b : $origin\ of\ order\ b$

d_b : $destination\ of\ order\ b$

t_b^o : $type\ of\ order\ b\ (reefer\ or\ dry\ box)$

t_a^c : $type\ of\ container\ a\ (reefer\ or\ dry\ box)$

dd_b: $due\ date\ of\ order\ b$

Cap_{ml} : $capacity\ of\ mode\ m\ on\ leg\ l$

$$r_{klm} : \begin{cases} 1 :\ if\ route\ k\ has\ mode\ m\ on\ leg\ l \\ 0 :\ else \end{cases}$$

4.2.5 Objective function

The objective is to minimize the total costs over the planning horizon. The costs consist of holding costs for containers at terminals, costs for using a route (this includes terminal handling costs), costs for using containers, and penalties for delivery after the due date.

4.2.6 Constraints

Most of the constraints that are applicable can be derived from a standard ILP for SND, see Appendix A. For the problem that we consider, we identify the following additional constraints.

- An order can only be loaded in a compatible container.
- Each order can only be assigned to one route (intermodal route, direct truck with or without container) at one departure time.
- Container balance constraints: a minimum level of containers should be available at a terminal. This includes the physical stock of containers at the terminal and the containers that are currently on their way to the terminal.

4.3 Planning heuristic

In this section we cover our solution approach to the problem. In Section 4.3.1 we elaborate on setting benchmarks. We propose our planning heuristic in Section 4.3.2. The benchmarks and planning heuristic are compared through simulation. We address simulation in Chapter 5.

4.3.1 Benchmarks

To assess the performance of our heuristic, we use two benchmarks: a direct trucking benchmark and a "best" route with empty container repositioning benchmark.

Direct trucking

The direct trucking benchmark simply sends every order from its origin to its destination terminal by truck. As one of the assumptions is that truck transport is always available, every order gets delivered. This benchmark does not consider consolidation or repositioning of empties. Containers are used if they are available, if no container is available, trailer trucking is used.

Best route with container repositioning

The best route benchmark assigns, as its name suggests, the "best" available route to every order. Best is here defined as the route that is able to deliver the order before the due date against lowest costs at the latest possible departure time. If there are no possible routes available that can meet the due date criteria, direct trucking is performed. Further, if there are no containers available at the start terminal, direct trucking is performed as well. No replanning of orders (to search for better consolidation options) or repositioning of empties is applied. By setting a benchmark that does not consider consolidation, replanning, and repositioning of empties, we can assess how good our heuristic that does cover these three points performs.

Figure 4.2 shows the flowchart for the best route benchmark. When an order arrives, the first procedure is to check if a container is available. The corresponding procedure is a sub-procedure, which we show in Figure 4.3. The return value from the sub-procedure is whether or not a container is available. If no container is available, the order is trucked directly to its destination. Otherwise, out of the available routes between origin and destination, a sub-selection is made. The criterion here is whether the route is fast enough to deliver before the due-date. If no route is available, the order still gets trucked directly but with a container. If there are routes available, the next selection criterion is based on costs. If multiple cheapest routes are available, the latest departing route in the list of remaining routes is selected. This choice is done to enhance the consolidation possibilities. Otherwise the cheapest route is selected.

Empty container repositioning is a research field on its own and has a lot of potential for optimisation. Because it is a research field on its own, we apply a simple decision rule rather than actually optimising the empty flows. As we address in Chapter 5, we divide the networks in two areas. We use these two areas (for example, east and west side) to study the effect of imbalance. One side receives a higher amount of orders than the other side. For the repositioning, the number of orders going to each area are tracked until a point in time where empty containers need to be repositioned. Repositioning can be done once a day, once a week or at any other time interval.

The difference in the number of orders going to either side determines the maximum amount of containers to send to the side with a container deficit. From the terminals on the surplus side of the network, the terminal with the highest number of available containers is selected. For the deficit side, the terminal with the lowest number of available containers is selected. Next for all the routes running between these two terminals, we select the cheapest available route that is already in use. If no capacity remains on a route or if there is no active route between the two terminals the next terminal with the highest surplus is selected. This is repeated until either the imbalance for that period is brought down to zero, or no more routes are available, see Figure 4.4 for the flowchart showing the this procedure.

The decision policy for repositioning of empties has its pros and cons. First and foremost this method aims to keep the number of containers balanced at either side of the network. Next, it makes use of routes already in use. Therefore, no additional costs are made when working with fixed charges for using transport. Downside to this however is that it is possible to match two terminals that are relatively far away from each other, whilst a terminal that is relatively close would be good enough. Next, it is possible that on a given day we send an empty container from A to B and the next day we send an empty container from B to A. Finally, the method does not deal with imbalance within one side of the network that occurs for example in the case. If we look back at the data from Chapter 3 we see for example that a lot of orders get sent to terminal 1 and a lot of orders get sent from terminal 2. Meaning that empty containers pile up at terminal 1, whereas they are needed at terminal 2. For the case network the method for empty repositioning is slightly extended by also repositioning the containers at one side of the network. This repositioning happens in the same fashion as the repositioning between the two areas (start with highest surplus and highest deficit).

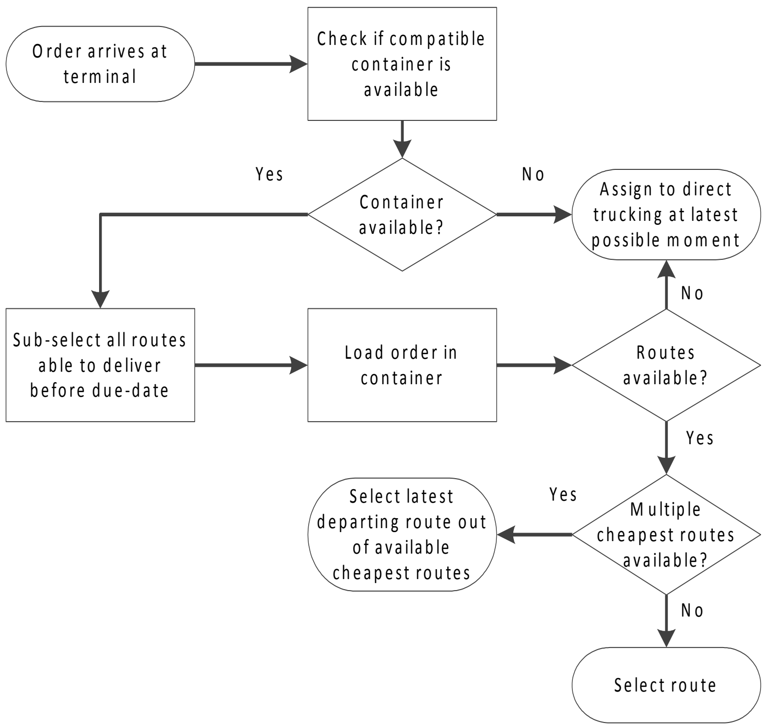

Figure 4.2: Best route benchmark

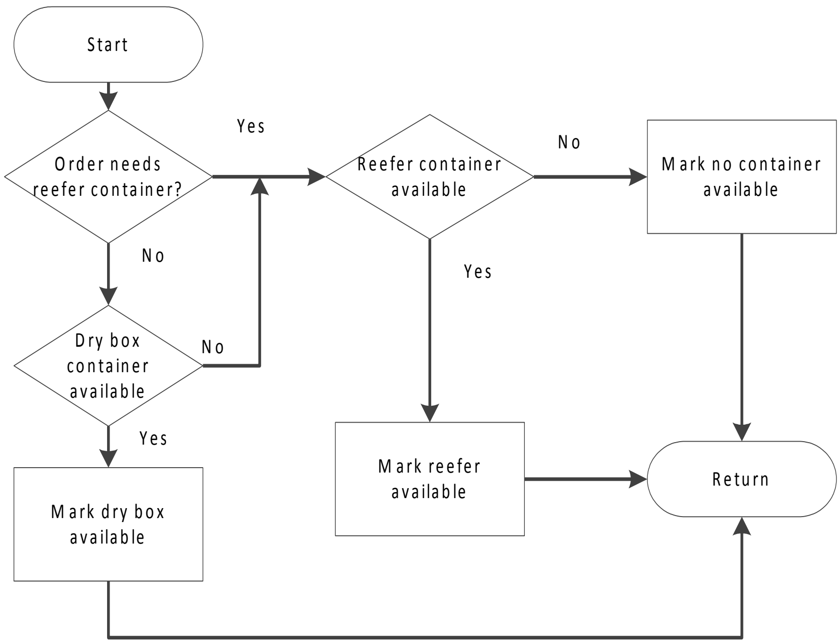

Figure 4.3: Check for available containers

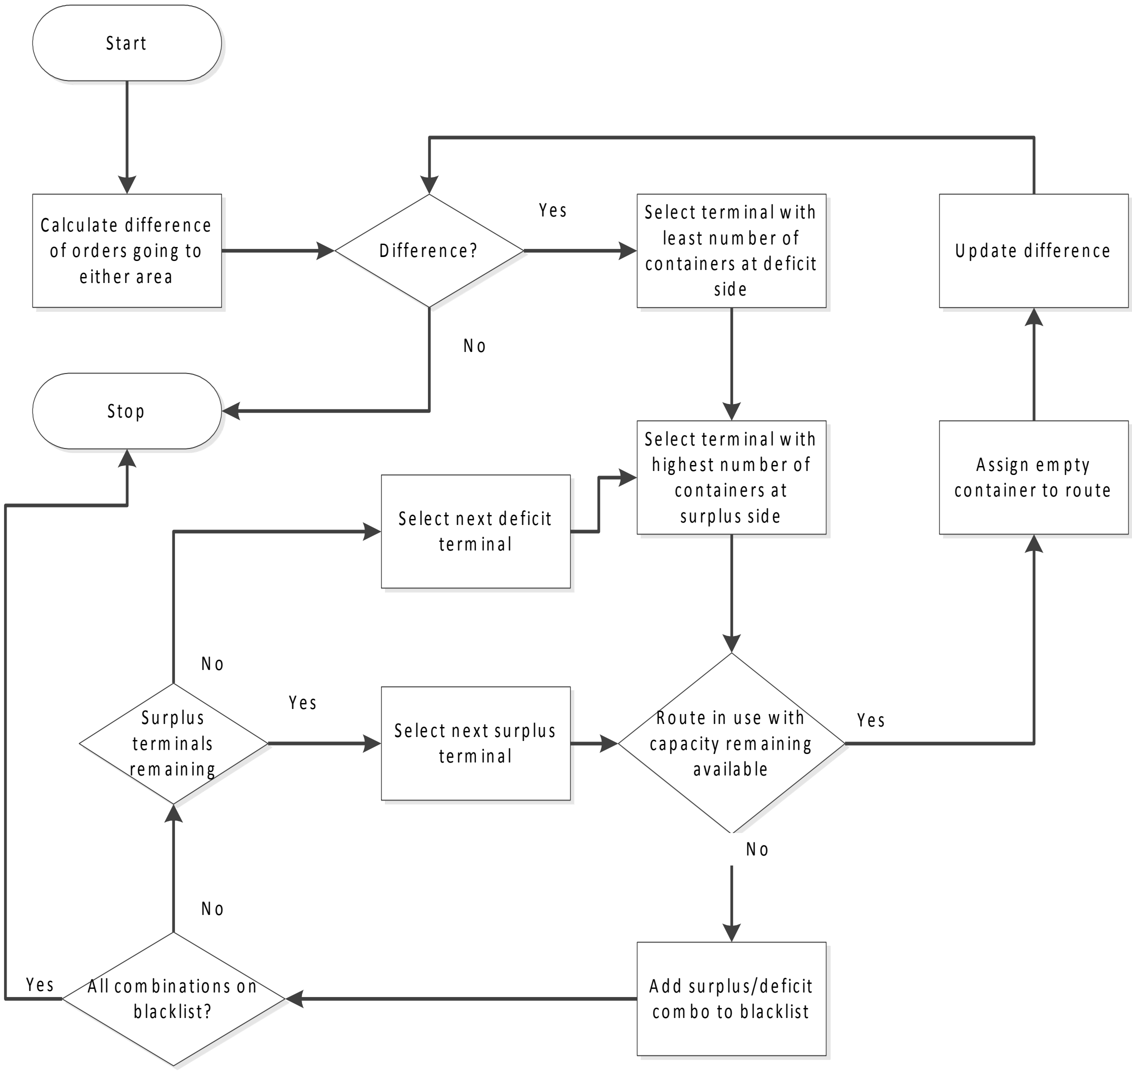

Figure 4.4: Flowchart for empty container repositioning.

4.3.2 Optimisation heuristic

Before we go into detail about our heuristic, we briefly describe the general idea.

Description

In essence, there are three steps that have to be performed. Initial order planning, optimisation, and determining the empty flow. Once a transport request arrives, it has to be planned for transport. This is solved through selecting the best available route as explained in Section 4.3.1. The initial order planning occurs throughout the day, each time that a transport request arrives it is immediately planned according to the method shown in Figure 4.2.

Because it is possible that the initial plan is not optimal, we want to check whether it is possible to create a better plan through an optimisation heuristic. We look at the effect on costs when swapping and moving orders in the schedule. These swaps and movements can be performed in multiple ways, the details follow later in this section. With regard to costs we identify several ways of allocating costs to transport modes: costs for each container placed on a transport mode, costs per transport mode regardless of the number of containers, and a mixed variant.

Finally, after the plan for loaded containers is established, empty containers are (if needed) added to the planned transport routes. This is to counter the imbalance that is present in most cases. Because empty repositioning is a problem on its own we create a decision policy to allow for empty repositioning.

Order planning

The heuristic that we propose uses as input a set of possible routes per origin-destination pair. We choose to set a maximum number of routes per origin-destination pair, to ensure that no routes that are not logical, reasonable, or feasible are used. A route that is not logical, is for example a route that takes a detour exceeding a distance threshold. A route that is not reasonable is for example a route that is more expensive than direct trucking. Finally, a route that is not feasible is for example a route that takes longer than the maximum lead time of orders. As we mentioned earlier on our heuristic uses predetermined routes as input. These routes are selected using a k-shortest path algorithm. Although it is possible to perform the three before mentioned check while performing the k-shortest path algorithm, we do not do this because we assume that due to the small size of the networks we are already pick out most if not all of the valid routes. As we defined in Section 4.1, a route between origin and destination terminal pair consists of one or multiple legs where transportation takes place at a designated time. Between two terminals multiple legs can exist, the difference being different modes.

To give an example, consider a network with 5 terminals and we allow for a maximum of 5 routes per origin destination pair. We then have a total maximum route-set of 100 routes. Routes that are too expensive (exceeding direct truck costs between the origin and destination) etc. are not included in the route-set meaning that it is possible to have only 1 possible route between an origin and destination.

Optimisation

As we initially plan our orders when they arrive (order for order), we are not able to guarantee that the initial plan is optimal. The optimisation procedure has the freedom to replan orders such that consolidation is encouraged. When our heuristic optimises the current transport plan it considers all possible orders that are still located at their origin terminal.

The first step is optimisation of the established plan. Optimisation is done through simulated annealing. We choose simulated annealing (SA) over tabu search due to the fact that simulated annealing selects a neighbour solution and either accepts or rejects it. Tabu search evaluates all neighbours before accepting a neighbour solution as the new solution. Evaluating all neighbour solutions has a negative influence on the computation times.

As we already stated in Chapter 2, simulated annealing uses the concept of neighbour solutions. Selecting neighbour solutions can be done in a variety of ways. We consider the following four options for creating a neighbour solution.

1) Swapping an order from an intermodal route with an order on a direct trucking route.
2) Moving an order from a direct trucking route to an intermodal route that has capacity remaining.

3) Moving an order from an intermodal route to a direct trucking route.
4) Moving an order from an intermodal route to another intermodal route that has capacity remaining.

Each of the options has an equal probability of being chosen, where a random order or random orders are selected for a move or a swap. The neighbour structures are depicted in Figure 4.6.

Option 1 swaps an order from a truck with an intermodal route, if the truck order does not have a container assigned to it, the intermodal order needs to have a compatible container for the truck order. For option 2 a compatible empty container has to be available at the origin terminal otherwise the move creates an infeasible schedule. The need for an empty container only holds if the direct truck order does not yet have a container assigned to it. Option 3 frees a container for other orders to be used. The order is trucked at the latest possible time as direct trucking is still unwanted. Finally, option 4 allows orders to be moved to different intermodal routes and/or different departure times. It should be noted that when moving or swapping an order to an intermodal route it is possible to switch to any route, even routes that are not yet in use.

Depending on the neighbour structure selected the procedure of selecting containers to swap or move differs. In case of neighbour structure 1 the heuristic first selects a random order that is planned for direct trucking. Next, an order from the same origin terminal is selected. Depending on whether the direct truck order already has a container assigned to it or not a check is performed to see if a container is available. Further, a check is performed whether there is an intermodal route with a lead time sufficiently small to deliver the order before its due date.

For the second neighbour structure the heuristic again selects a random order that is planned for direct trucking. A check is performed whether a container is available. Next, a check is performed whether there is an intermodal route with capacity remaining. If the checks are passed a swap is considered valid.

The third neighbour structure first selects a random terminal that has orders waiting to be transported. Next, the heuristic selects an order at random which it then replans from an intermodal route to a direct truck route.

The fourth swap selects an order in the same fashion as the third neighbour structure. A check is performed whether there is any route available with a lead time sufficiently small to deliver the order before its due date.

An order that has already partially transported is never selected by any of the four methods. Further, even if a swap or move is valid it still depends on the contribution of the swap or move whether it gets accepted.

The main reason for choosing to swap between routes instead of legs stems from the fact that if we only change one leg, we still have to evaluate all the legs on which the current order is planned. For example, consider the planning of a random order in Figure 4.5a. If we want to replan the third leg to an earlier point in time on the same leg we would have to consider the replanning of legs 1 and 2 as well. The same idea holds for planning an order on a different leg. Consider the planning of the order in

Figure 4.5b. If we are replanning the first leg to the alternative available leg (with a longer duration), we would have to re-evaluate the second and third leg as well. In short, replanning of a leg would require the evaluation of the other legs on which an order is planned as well. Because we have to consider every leg anyway we decide to replan a complete route rather than a leg. Given that every origin destination pair should have multiple alternative routes, enough freedom should be available.

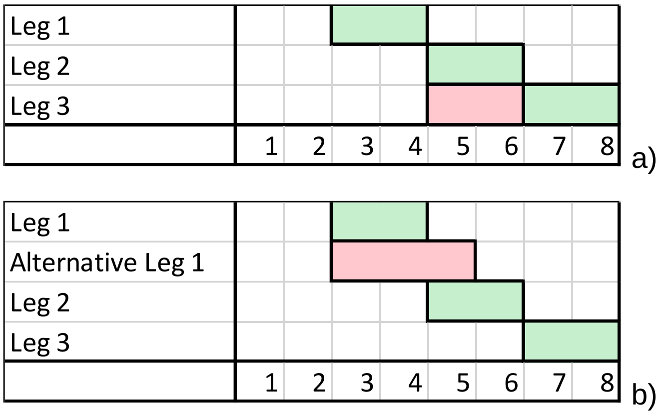

Figure 4.5: Replanning of a random order.

As we plan our orders on a one by one basis, we have an initial plan before the SA procedure starts. During the execution of SA it is possible to replan every order that has not started on its transport journey. In practice it is hard to alter the transport planning of an order when it is already on its way.

Finally, our choice of 4 moves/swaps do not cover all possible swaps and moves that are possible. For example:

- Moving multiple orders from one intermodal route to another intermodal route.
- Taking orders from multiple intermodal routes and place them on one different intermodal route.

We did not include these swaps into our procedure as we expect them to have a major impact on computation times.

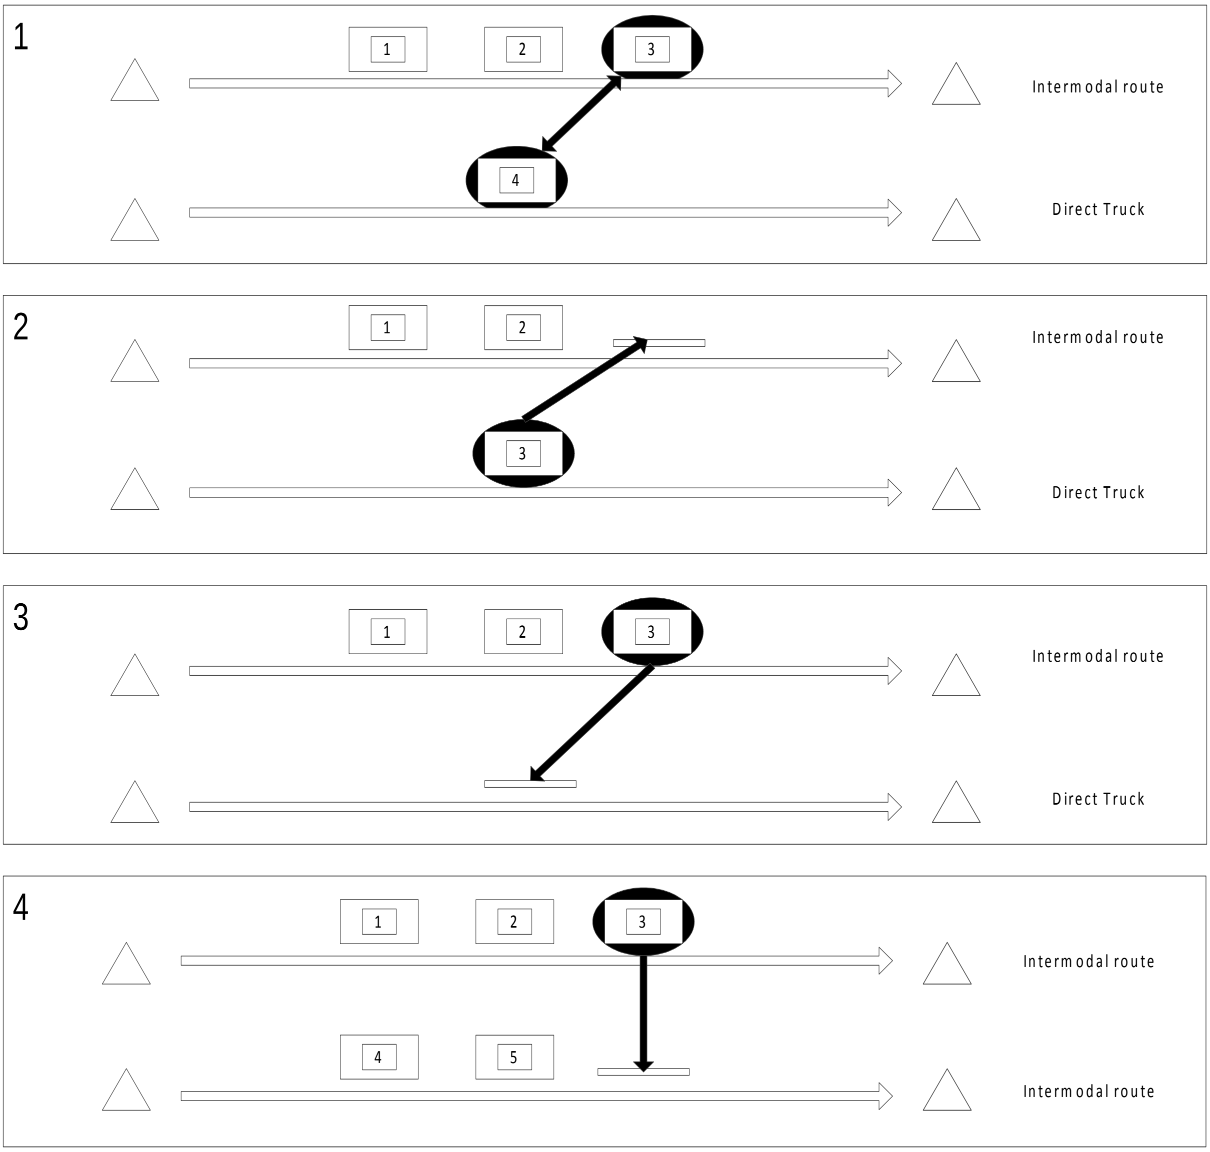

Figure 4.6: Swaps applied in the simulated annealing procedure.

Empties

After performing simulated annealing, there may be capacity remaining on the transport routes that are used. This capacity can be used for the movement of empty containers to solve the imbalance issues. For the determination of empty moves we use the procedure as introduced in our benchmark with empty moves, see Section 4.3.1.

4.4 Conclusions

In this chapter we gave a detailed description of the problem that we deal with. To address the problem we presented our planning heuristic together with several benchmarks to assess the performance of the planning heuristic. Two benchmarks are proposed:

1) A direct truck benchmark where every order is trucked (containerised or in a trailer) between origin and destination.
2) A “Best” route with empty container movements. The same as 2, but with repositioning of empty containers.

Further we presented our planning heuristic that is able to replan orders through a simulated annealing procedure. After optimisation of the planning, a decision policy determines the maximum number of empty containers to be placed on transport routes already in use.

5. Numerical experiments

This chapter focusses on setting up, testing, and analysing our planning methods. Section 5.1 addresses our simulation model. Next, in Section 5.2 we address our choices with regard to generic network settings. Section 5.3 covers our experimental design. Finally, Section 5.4 covers the analysis of the output from the simulation model. This chapter is concluded in Section 5.5.

5.1 Simulation model

Orders arrive at a terminal according to a Poisson process determined in Section 3.3.2. At the point of creation, the order gets an origin and destination terminal assigned. Finally it gets a due date, which indicates a latest delivery time at the destination terminal. The due date is determined with a triangular distribution based on findings from Section 3.3.3. Because the due date gained through the triangular distribution is the door to door lead time, we deduct two random values determined by a Weibull distribution based on Section 3.3.3. These two random variables represent the required pre- and end-haul time to truck orders from and to the terminal. Finally, orders can be of two types: either an order is standard or it is a temperature controlled order. A standard order can be placed in both a dry box and a reefer container. The temperature controlled orders can only be placed in reefer containers. The exact parameter values are given in Section 5.3.

After determination of the parameters, the order gets placed at its origin terminal. The planning heuristic plans the order on an intermodal route or assigns it to a direct trucking route depending on the outcome of availability of containers and capacity on the legs within a route. This assignment is preliminary as we need to optimise the plan. Optimisation can be done after each order arrival but this would drastically increase the computation time. Therefore optimisation happens once per day at the same point in time. Depending on the heuristic that is running, optimisation of the planning and assignment of empty containers to transport routes happens. In short: throughout the day orders are planned as they arrive. At a given point in time the plan is optimised. Next, the algorithm determines assigns empty containers. Finally, the orders and empty containers are actually transported.

With regard to costs, two different structures are applied. A fixed tariff and a variable tariff. In the case of fixed tariff, costs are paid for all available capacity on a on a leg if at least one container is assigned to this leg. The variable tariff only allocates costs to containers that are assigned to a leg. To give a simple example, consider a train with a capacity of 5 containers, typically the utilisation rate for this train is 50% and the fixed price for this train is 50. We have 3 containers loaded onto the train. Because the utilisation rate is on average 50% the variable costs per container are 20. In the variable case our costs are 60 for transport, whereas in the fixed case the transport costs are 50. Adding a container to the train in the variable case increases the costs by 20, whilst in the fixed case the costs remain 50.

5.2 Generic network generation

In this section we consider the generation of a generic network. As our case network with 6 terminals and 9 legs is quite limited, we want to test the effect of our heuristic on somewhat larger networks containing more legs. For the generation of a generic network we take an area of 1500 by 250km. This is done as this is approximately equal to the area considered in the case. For this area we consider two different randomly generated networks. We consider a 10 terminal network with 20 and 30 connections. Although it is favourable to generate new networks for every instance of a network, it would require to run a k-shortest path algorithm to create new routes over and over again, since we are working with pre-determined input to reduce the computation time. Therefore we decide to use the two predetermined networks (20/30 connections).

First, for each terminal, a random X and Y coordinate is generated. Next, based on the amount of connections required, a random number is drawn and compared to a threshold. This threshold is nothing more than a percentage of the maximum connections possible, e.g. if there are a maximum of 20 connections possible and we want 10 connections the threshold is 50%. Based on this a connection is either established or not. Networks are generated until the number of connections required is established. Logically, there are other methods to create a random network, e.g., picking a random origin and destination terminal 20 times. We decide to use our method because we think it is an interesting way although it is cumbersome.

Each terminal has access to truck transport. This is to ensure that each order can be delivered directly. For each connection in the network there is an 80% probability that it is a train leg and 20% probability that it is a barge leg. In the generic network each mode has one departure per day. All transport is executed at the same point in time on a given day, for example, all departures are at 12 AM.

Finally, we want insight in the effect of imbalance in the network. To this end, the network area is cut in half, effectively dividing the area in east and west. Each area gets a percentage of the orders as origin based on an imbalance percentage. The origin and destination area of the orders are always different. The orders are then uniformly distributed over the terminals in the corresponding area.

Figure 5.1 shows the 10 terminal network with 20 and 30 connections. Solid lines indicate a train connection and dashed lines indicate a barge connection. These two networks are used for testing our heuristics. In Appendix E we give details on the locations of terminals and the type of connections between each terminal pair.

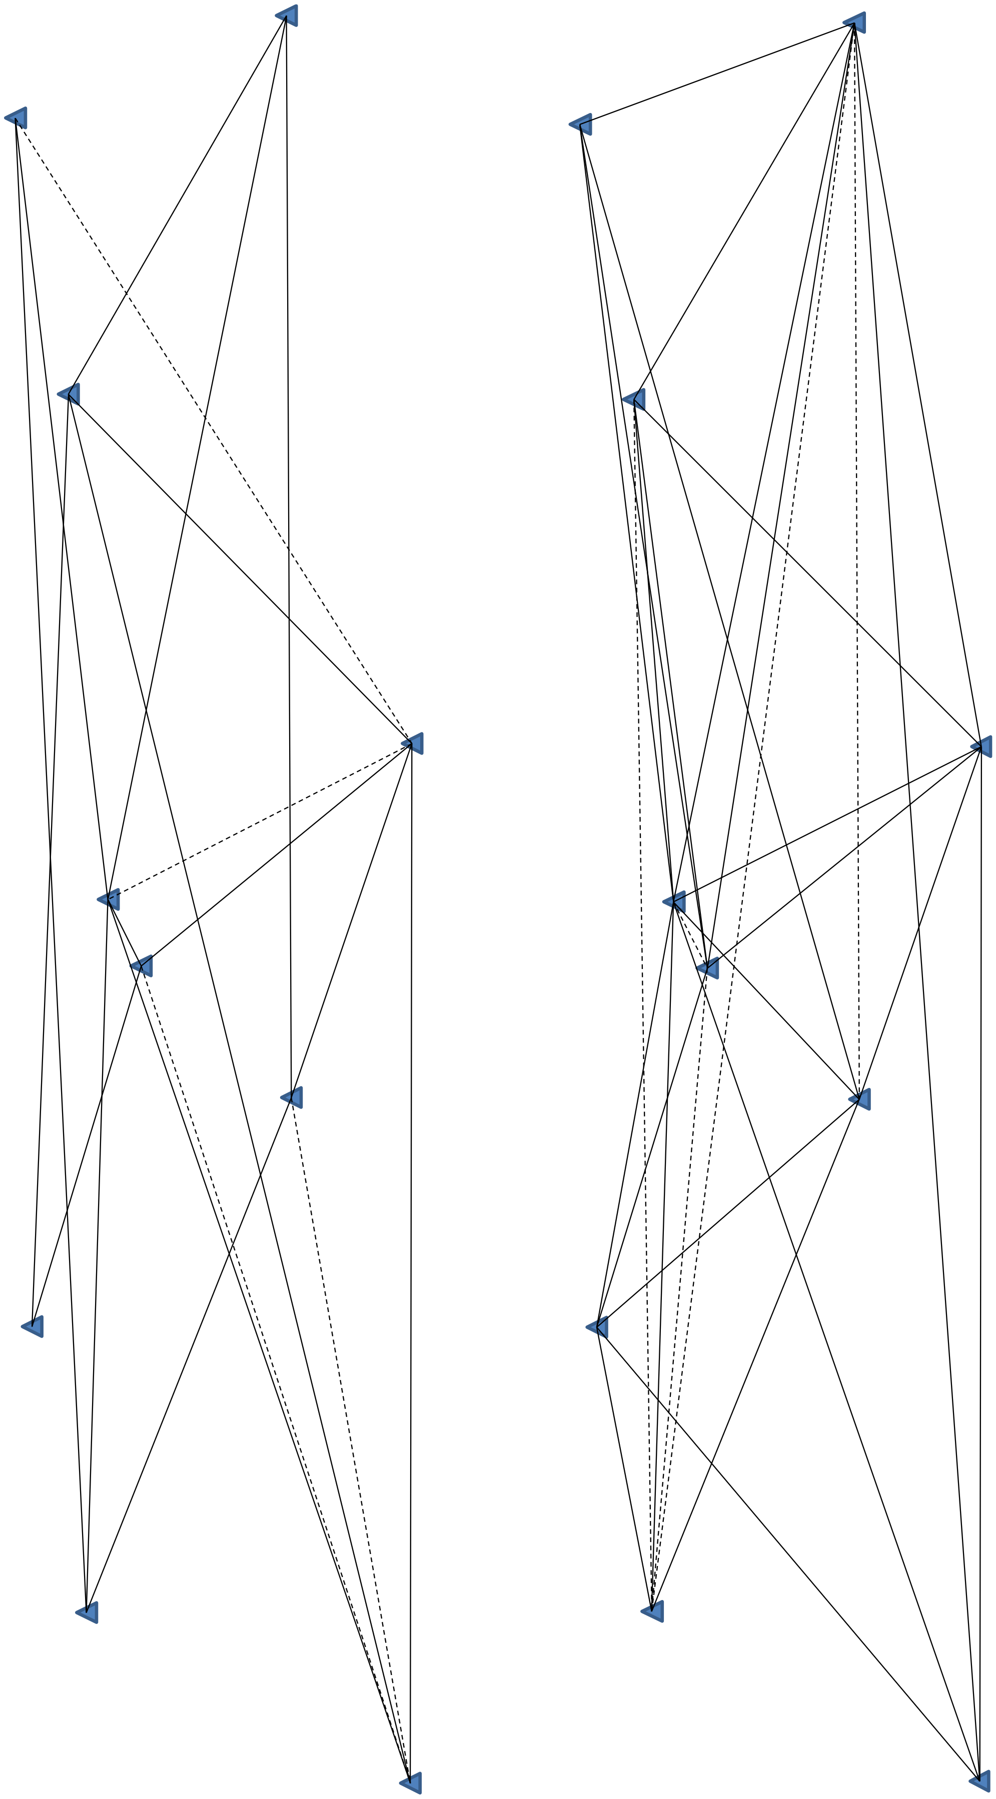

Figure 5.1: 10 terminal network with 20 connections (left) and 30 connections (right)

5.3 Experimental design

This section covers our experimental design decisions. Section 5.3.1 covers the parameter settings. Next, Section 5.3.2 covers the experimental settings.

5.3.1 Parameter settings

In a simulation study several aspects have influence on the output data. In this section we cover transient/steady state system behaviour, terminating/non-terminating simulations, warm up periods, and number of replications. According to Law (2007) theses settings need to be determined to increase the reliability of our output data.

Transient/steady state system behaviour

Transient system behaviour indicates that system performance is dependent on the initial state of the system whereas steady state behaviour means that the performance is not dependant on the initial state (anymore) (Law, 2007). In reality there are usually orders already in the pipeline, however our model starts with an empty system. Therefore at the start of a simulation run we do not have a steady state yet. To deal with this, a warm up period is required.

Warm up period

A warm up period is used to remove data from the transient phase within our simulation model. The warm up period for our model is determined by using the procedure of Welch (Law, 2007). The number of orders delivered each day is used for determining after which point a steady state behaviour is reached. We use the costs per day as indicator to determine the warmup period, because costs is one of the important performance indicators. We use a warm up period of 30 days, see Appendix F for details with regard to this decision.

Terminating/non terminating simulation

We are dealing with a nonterminating simulation. This is due to the fact that we cannot identify a specific moment in time where the system is empty, e.g., all orders have been delivered (Law, 2007). After a run is finished we always have a situation where orders are planned or being transported. A run represents a period of simulated time. The state of the simulation model at the end of a simulation run in turn is an initial situation for the next run effectively ensuring that no new warm up period is needed. Using one warmup period and then performing multiple runs is known as the batch means method. A batch means method is a valid approach for running a simulation model if batch observations are uncorrelated. This is the case for batches that span 30 days. Therefore we use a batch means method where each run spans a period of 30 days, see Appendix F for details.

Number of runs

To increase the reliability of the outcomes, multiple simulation runs per experiment are performed. We determine the required number of batches we use the total costs as indicator. After the warm up period of 30 days, each run spans a 30 day period. The average outcome values of every run are indicators for the performance of our model. Again, costs are used as indicator. To determine the required number of runs the following equation should hold:

$$\frac{\delta(n,\alpha)}{\bar{X}_n} \leq \frac{\gamma}{1+\gamma}$$

The left hand side of the equation represents the confidence interval of the average performance in a batch (run) that should be below a certain threshold (right hand side). The equation holds for n = 18 runs and therefore we use 18 runs for the experiments, see Appendix F for the details.

5.3.2 Experimental settings

In Chapter 3, the imbalance of orders between east and west is addressed. We want to know the effect of imbalance on the solution of our optimisation heuristic. To this end, we consider:

- Two settings: no imbalance where both “sides” have a 1:1 order ratio and the case with imbalance using a ratio of 3:2.
- A division of 1000 containers in a 3:1, 1:1, and 1:3 Dry box to Reefer ratio.
- The order ratio (cooled and dry goods) remains equal throughout all experiments and is set on 1:1.
- A maximum of 10 routes between each origin destination pair, in practice having more than 10 different routes between a given origin and destination rarely happens. A more common approach is to have one main route and a few alternative routes in case of disruptions on the main route. However we want to give our heuristic more freedom than the limited set of routes usually found in practice.
- The order arrival rate that we decide to use for our generic networks is higher than we use in our case study, see Chapter 3. Due to having more connections we also have more capacity available. To reach a realistic utilisation rate we increase the number of order that arrive per day. After some preliminary testing we find that intermodal routes are not used in the case of 28.6 orders per day, unless the container capacity on most legs for barges and trains are set to 1. A capacity of 1 in turn does not allow for consolidation of orders. Preliminary testing with 50 orders per day does show effects and therefore we set the order arrival rate to 50 orders per day.

The travel speeds for truck and train are based on realistic values that are used at Seacon. The travel speed for barges that we use is derived from Spikker (2014). In realistic networks, capacity for trains tends to vary between 36 and 42 containers. The capacity for barges is usually higher. If we are to use these capacity values in our model we would have to increase the order arrival rate accordingly to it as well because of the desired utilisation rate, realistically we are looking at arrival rates of 1000 orders per day and higher. The result of increasing the order arrival rate is that the computation time would most likely exceeds reasonable boundaries. If for example our heuristic has to be executed once per day this is not a major issue. However during the testing we have to execute our heuristic several hundreds of times. Therefore we lower the capacity of trains and barges. For each leg in the network we uniformly draw a number. Train capacity is drawn uniformly between 3 and 4 and barge capacity is uniformly drawn between 6 and 8. These values are then fixed for the network. As the capacity on a leg in reality is far larger, the capacity we use can be seen as a reservation on that leg. Therefore the capacity in one direction does not have to be equal to capacity in the other direction on that same leg, e.g., going from A to B by barge has 4 capacity but going from B to A by barge has 5 capacity. Table 5-1 gives an overview of the fixed and experimental settings.

Variable	Fixed settings	Experimental settings
Transhipment costs	Values according to Table G-4 in Appendix G	
Applied heuristics		Direct trucking, best route, and the optimisation heuristic.
Transportation costs (per km per container)	Truck 1.00 Train 0.51 Barge 0.33	Tariff per container and fixed price per leg
Order arrival rate (per day)	50	
Order lead time	Triangular distributed (7, 7.5, 8)	
Order imbalance rate (east/west)		1:1, 3:2 (Table 5-2)
# of available containers (Reefer/Dry box)	1000	1:3, 1:1, 3:1 (Table 5-2)
Order types	2 (cooled and dry) with a 1:1 ratio	
Max number of routes per origin-destination terminal pair	10	
Run length	30 days	
Warm up period	30 days	
Number of runs	18	
Travel speeds	Truck: 68 km/h Train: 65 km/h Barge: 15 km/h	
Number of terminals	10	
Number of legs		20, 30 (Table 5-2)
Vehicle capacity[2]	Train: between 2 and 3 Barge: between 4 and 6	

Table 5-1: fixed and experimental settings for generic networks

Network breakdown

Each network experiment consists of running a direct truck and a best route with empty repositioning benchmark together with our optimisation heuristic. All of these benchmarks and the heuristics are performed for both variable (per container) and fixed (per transport mode) transport costs, except for the direct truck benchmark, which is performed only once. Table 5-2 provides an overview of our network settings.

[2] Leg specific capacity is given in Appendix E.

	Connections	containers (reefer : dry box)	Imbalance
Generic 1 (G1)	30	1:1	1:1
Generic 2 (G2)	30	1:1	3:2
Generic 3 (G3)	30	3:1	1:1
Generic 4 (G4)	30	3:1	3:2
Generic 5 (G5)	30	1:3	1:1
Generic 6 (G6)	30	1:3	3:2
Generic 7 (G7)	20	1:1	1:1
Generic 8 (G8)	20	1:1	3:2
Generic 9 (G9)	20	3:1	1:1
Generic 10 (G10)	20	3:1	3:2
Generic 11 (G11)	20	1:3	1:1
Generic 12 (G12)	20	1:3	3:2

Table 5-2: Network settings.

The settings result in a total of 72 experiments. However, as the cost structure (fixed or variable costs) does not have influence on the direct trucking benchmark we can leave out 12 experiments. Leaving out 12 experiments leaves a total of 60 experiments. Further, in the case network we look at the effect on costs when we consider no stops at terminal 3, one stop per week at terminal 3, and two stops per week at terminal 3. As the other experimental settings are fixed for the case network, we are left with 9 experiments for the case network (3 types of heuristics and 3 different train schedules).

Simulated annealing parameters

One of the important aspects of simulated annealing (SA) is ensuring a good cool down schedule. The cool down schedule has influence on the performance of simulated annealing.

In Figure 5.2 we see that the initial solution is approximately 118,000. From there the SA selects neighbour solutions moving in either direction (increase or decrease in costs). In the early stages this should give a broad range of accepted solutions, gradually proceeding to the point where it is unlikely that a worse solution is accepted. In Figure 5.2 we identify the behaviour as described as well. Ensuring the perfect settings for SA is a difficult task on its own, we tested with several settings for our cool down schedule resulting in all kind of curves. The most favourable curve that we find is as depicted in Figure 5.2. The parameters used to create the curve are:

- Start temperature: 500
- End temperature: 0.1
- Cooling parameter: 0.9
- Iterations 10

Each iteration in the SA procedure is equal to one swap or move of order(s). When a set number of iterations from the SA procedure are done, the process is cooled down by the cooling parameter. At this point we increase the number of iterations by 5. We found that when allowing a lot of iterations in the early stages would result in not finding any improvements to our schedule, because we accepted too many “bad” solutions. Therefore we increase the number of iterations during the execution of SA. The process is repeated until the end temperature is reached.

As we find the most favourable curve with the parameters described before, we use these settings in our SA procedure as well.

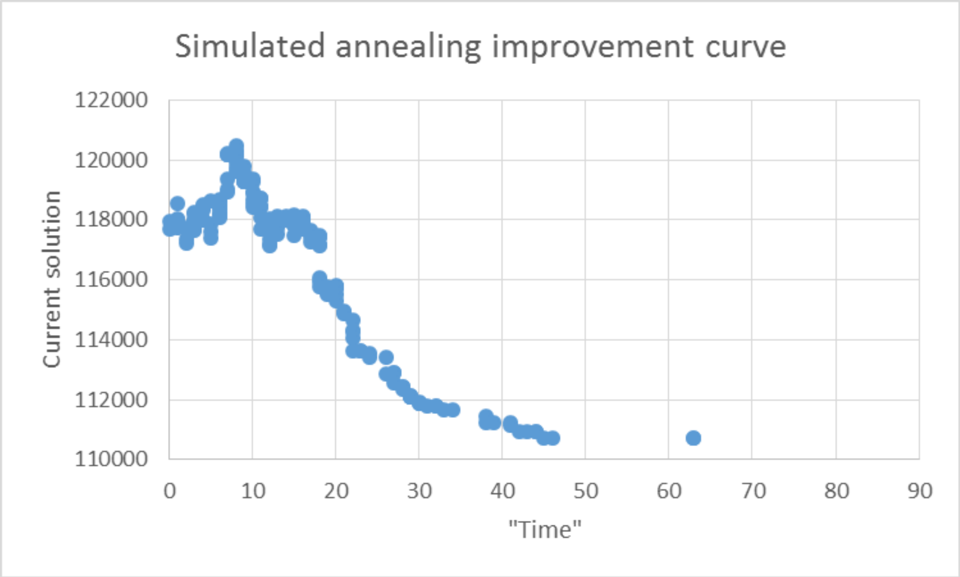

Figure 5.2: Simulated annealing improvement curve.

In Figure 5.3 we show the acceptance ratio curve belonging to the cool down schedule. The acceptance ratio curve gives information about the percentage of feasible neighbour solutions that get accepted. Normally, the acceptance curve starts at approximately 100% acceptance. However as we see in Figure 5.3 the curve belonging to our cool down schedule starts at approximately 50%. In essence we can adjust this increasing the start temperature. Preliminary testing shows us that the acceptance curve starts at 100% in that case. However, the results with regard to the optimal solution show the same behaviour as when we start with a high amount of iterations. In the early stages too many “bad” solutions get accepted and in the end the SA procedure finds no improvement. Therefore, we decide to lower our start temperature that results in an acceptance curve that not completely matches the standard procedure for SA.

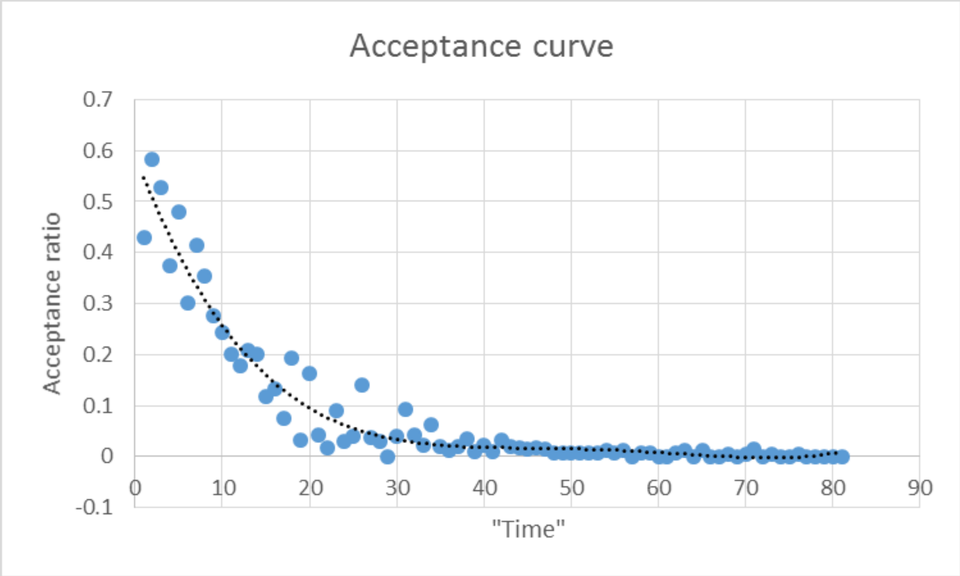

Figure 5.3: Simulated annealing acceptance curve.

5.4 Analysis of results

In this section we analyse the results from our simulation experiments. Section 5.4.1 focusses on the output from our generic networks with fixed costs, which we consider as the most interesting to analyse. In Section 5.4.2, we briefly address the generic networks with variable costs. As the case uses a fixed cost structure for the long-haul, the variable cost structure is less interesting for this research. In Section 5.4.3, we focus on the results from the case network. Finally, in Section 5.4.4, we give insight into the performance with different order arrival rates.

The notations G1-G12 used in tables or figures refer to the network settings as introduced in Table 5-2.

5.4.1 Generic network with fixed costs

In this section we analyse the results from our experiments on the generic networks with fixed costs. We subdivide the analysis in network utilisation, costs, empties, and lead times.

In the comparison between the direct trucking benchmark with our optimisation heuristic we only consider the cost aspect. Because network utilisation and empties do not play a role in direct trucking and lead time is not our main optimisation criterion.

Network utilisation

Table 5-3 covers the average utilisation rate and utilisation rate improvement of the transport legs that are in use. In all instances, our optimisation heuristic shows an increase in utilisation rate compared to the best route benchmark. On average we find 12.1% and 9.8% increase in utilisation rate for the 30 and 20 connection networks respectively. Further, we find that none of the network settings reaches the mentioned utilisation rate of 90% that we found in the literature. We suspect that the main reason for not achieving a utilisation rate of 90% is the fine tuning between capacity and demand.

Network	Average utilisation without optimisation	Average utilisation with optimisation	Improvement
G1	72.1%	84.3%	12.2%
G2	67.8%	79.6%	11.8%
G3	71.8%	84.6%	12.8%
G4	67.7%	79.9%	12.2%
G5	71.3%	83.2%	11.9%
G6	67.8%	79.5%	11.7%
G7	78.4%	87.6%	9.1%
G8	73.9%	84.4%	10.5%
G9	78.0%	87.1%	9.0%
G10	74.1%	84.2%	10.2%
G11	77.6%	87.6%	10.0%
G12	73.8%	83.7%	9.9%

Table 5-3: Utilisation rate improvement.

We further find that the number of orders getting trucked increases heavily in imbalanced networks, compared to their respective balanced counterpart. Table 5-4 shows the difference in number of orders trucked per setting. As we see, the 30 connection network (G1-G6) has a major increase in trucked orders. The increase ranges from 70 to 95%. For the 20 connection network we see a slight decrease in number of orders trucked for the balanced network (G7, G9, and G11), whilst we see an increase in the imbalanced situation (G8, G10, and G12). Partially this increase can be attributed to the fact that in most of the networks the average number of legs that are used per day decreases. This decrease in number of legs used results in less capacity available for transportation. On average, the balanced networks use 3.7% more legs and imbalanced networks use 11.4% less legs in a 30 connection network. For the 20 connection network we are looking at an average of 3.8% less and a 14.7% less legs for balanced and imbalanced networks respectively.

Network	Number of orders trucked without optimisation	Number of orders trucked with optimisation	Difference
G1	5.9	10.1	4.1
G2	10.8	21.0	10.2
G3	5.9	9.9	4.1
G4	10.8	20.4	9.5
G5	6.9	11.7	4.7
G6	11.0	22.1	11.2
G7	21.9	21.2	-0.7
G8	24.9	27.7	2.8
G9	21.8	21.1	-0.7
G10	24.8	27.8	2.9
G11	22.4	21.6	-0.8
G12	25.1	29.3	4.2

Table 5-4: Number of orders trucked before and after optimisation.

Costs

Figures 5.4 and 5.5 show the average cost reduction for our optimisation heuristics compared to the best route heuristic. As we see, in all cases the best route benchmark gets outperformed by our optimisation heuristic. The actual cost savings appear to be dependent on the imbalance ratio. When an imbalance ratio is present, the cost savings are lower. The main source for imbalanced networks having higher costs stems from the fact that within imbalanced networks more orders get trucked directly to their destination.

Further, we can see in Figures 5.5 and 5.6 no major influence on the average cost changes when we consider the division of reefers and dry boxes. Finally, Table G-5 in Appendix G summarises the absolute costs and cost reduction per setting. From a relative point of view our heuristic yields better results when we apply it to a smaller network.

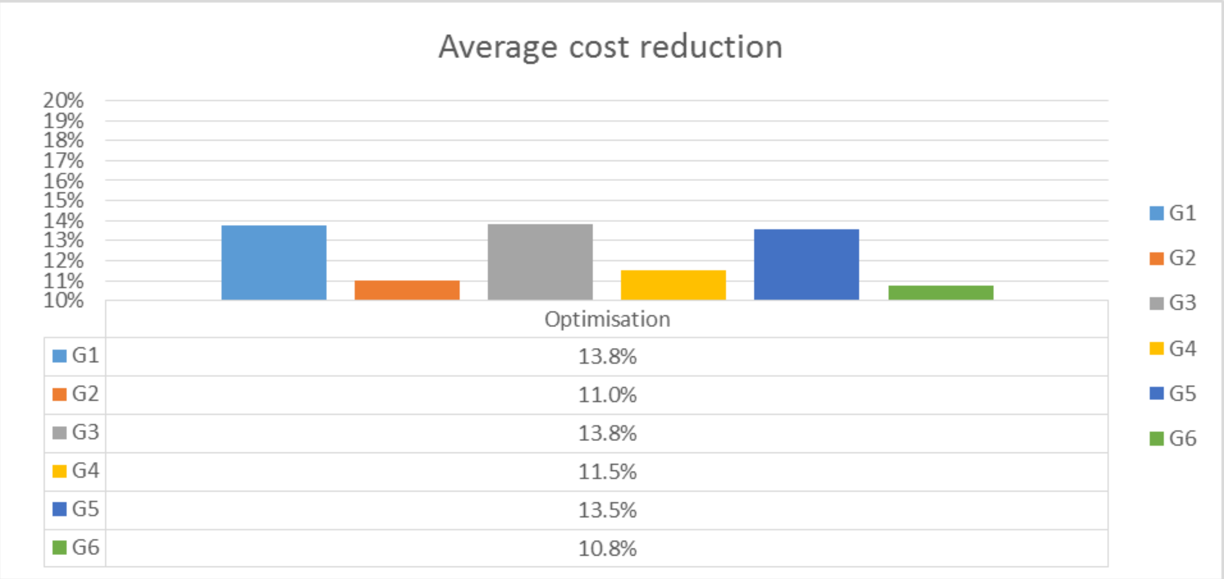

	Optimisation
G1	13.8%
G2	11.0%
G3	13.8%
G4	11.5%
G5	13.5%
G6	10.8%

Figure 5.4: Average cost reduction in a network with 30 connections.

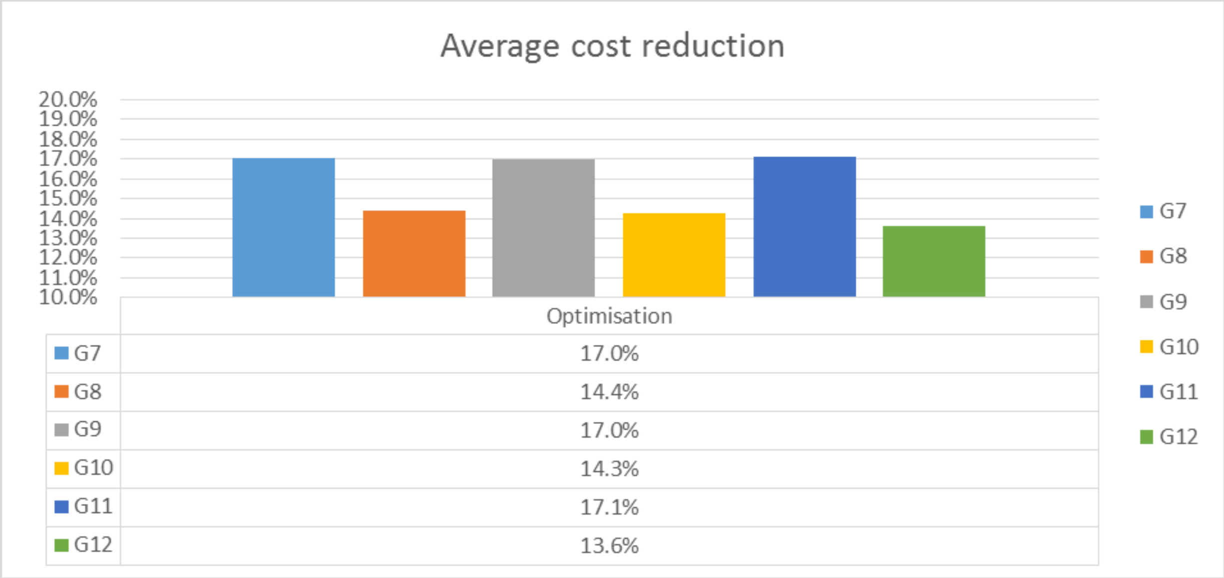

Figure 5.5: Average cost reduction in a network with 20 connections.

Figures 5.6 and 5.7 show how the costs are divided over intermodal costs and direct truck costs. We already gave an overview of the change in number of orders trucked in Table 5-4. The increase and decrease in number of orders trucked is closely reflected in the cost division as well. The difference percentage wise, if any, is smaller than 2.5%.

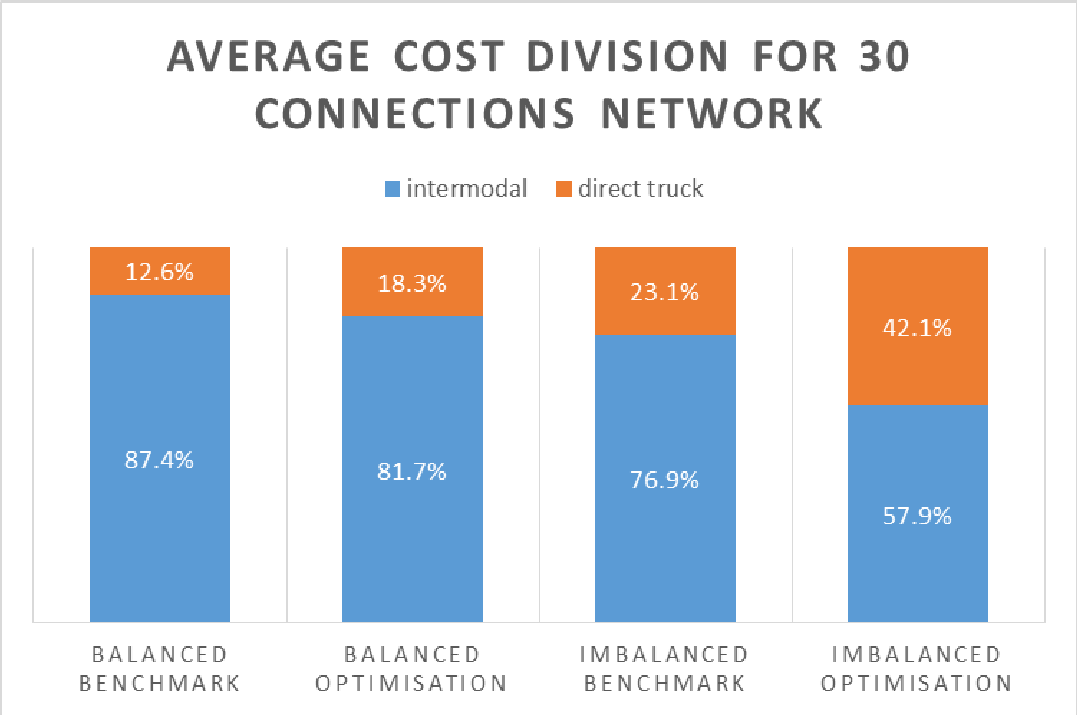

Figure 5.6: Average cost division in a network with 30 connections.

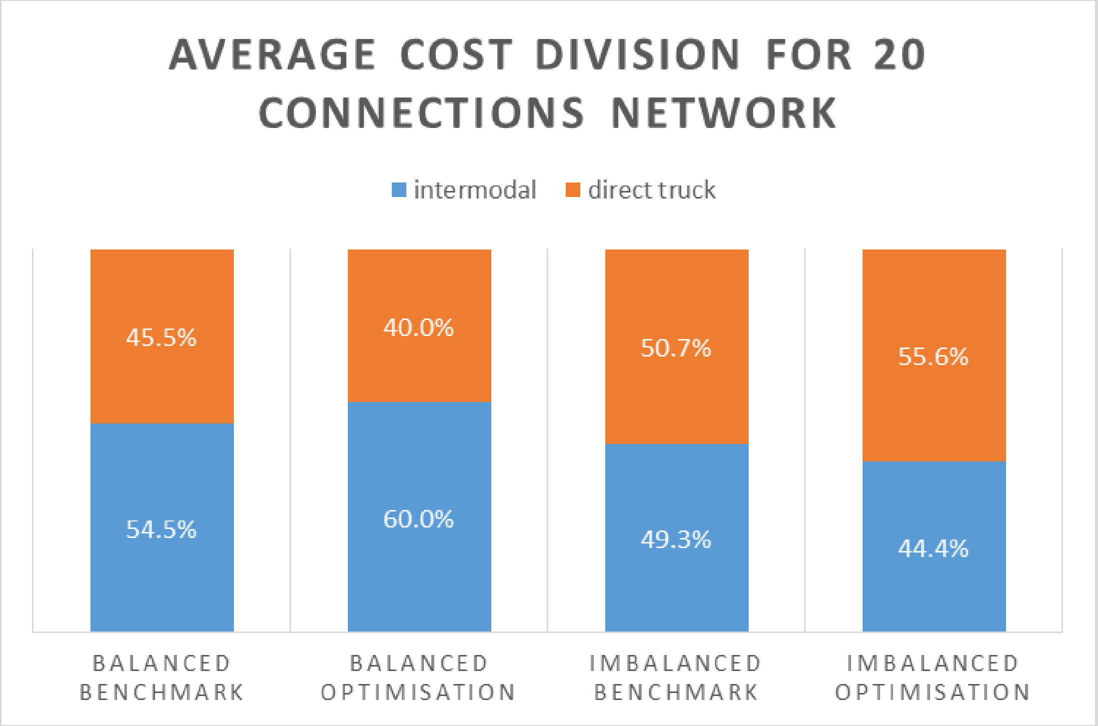

Figure 5.7: Average cost division in a network with 20 connections

Finally, we look at the performance of our heuristic compared to direct trucking. In all cases the optimisation heuristic outperforms direct trucking. We summarise the relative and absolute difference between the two in Table G-6 in Appendix G. When our heuristic is applied to the network with 30 connections we see better results than the 20 connections counterpart. The reason that better results are found in the 30 connection network is found in the fact that the "shorter" intermodal routes are possible in a 30 connection networks.

Empties

Figures 5.8 and 5.9 show the change in number of empty moves. As we see, the empty containers transported are one average 47.2% and 37.8% in the 30 and 20 connection network respectively. The behaviour of fewer empty moves goes against our expectations as the imbalance has not disappeared and we would expect to see an equal amount of empty moves before and after optimisation. We expect that this behaviour stems from the fact that the number of intermodal legs that are in use are lower in the optimised case when compared to the "best" route heuristic. Less legs in use in turn means that there is fewer capacity available to transport containers on intermodal legs. As we see in Table 5-5 the difference between the different network settings (G1-G6 and G7-G12) is quite small.

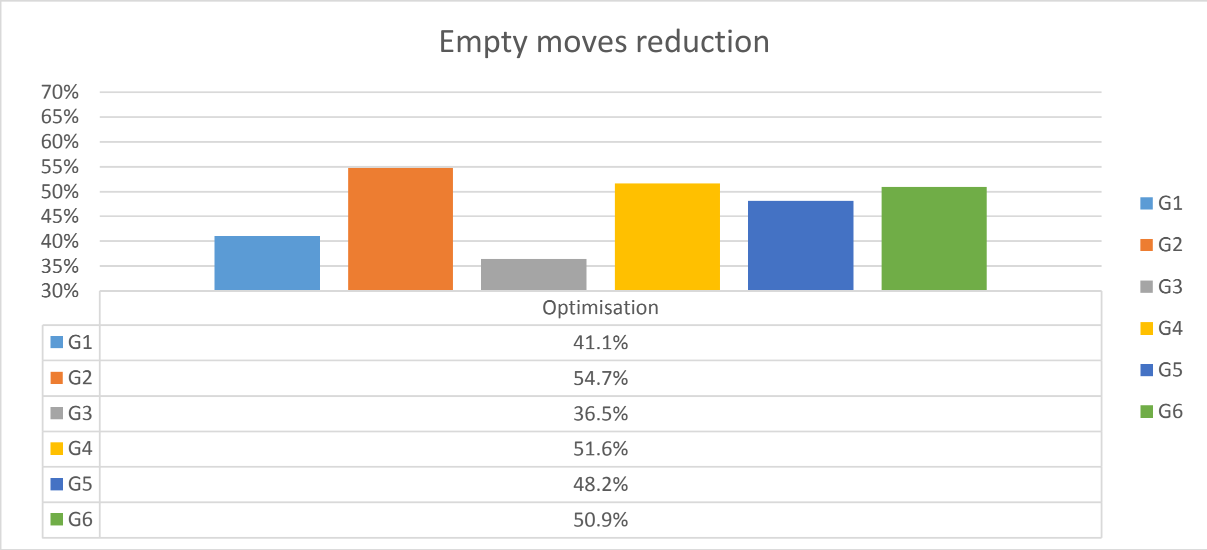

Figure 5.8: Average changes for empty moves in a 30 connection network.

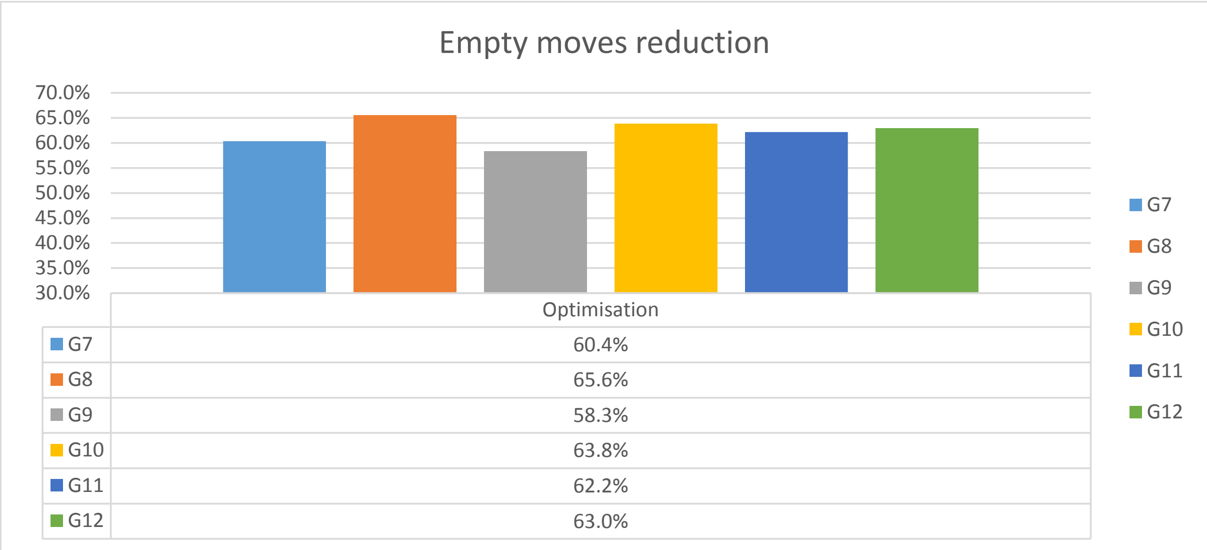

Figure 5.9: Average changes for empty moves in a 20 connection network.

Network	Average reduction (per day)	Network	Average reduction (per day)
G1	2.2	**G7**	1.1
G2	2.5	**G8**	0.8
G3	1.9	**G9**	1.1
G4	2.3	**G10**	0.8
G5	2.5	**G11**	1.0
G6	2.3	**G12**	0.8

Table 5-5: Absolute reduction in empty movements per day.

Lead and dwell time

The lead time is not part of our objective function but it is an interesting indicator to look into. The lead time here is the moment an order arrives at its origin terminal until it is delivered at its destination terminal. Because we do not optimise on lead time we do not expect to find major changes with regard to the average lead time. On average we find an improvement of 0.4% in the 30 connection networks and 0.9% in the 20 connection networks is seen. To put these percentages in perspective, for an average lead time of 5 days we have a reduction of approximately 1 hour. Given the minor impact, we decide not to look into further detail on lead times.

Another interesting aspect that we consider is the dwell time. Dwell time is defined as the period between an order arriving at its origin terminal and departure on its first transport leg. Figures 5.10 and 5.11 show the average dwell improvements for the different settings when compared to the “best” route benchmark. For the 30 connection network we see an average reduction of 14.5% whereas the 20 connection network has an average reduction of 10.4%. We consider a reduction of dwell time as improvement. Table 5-6 shows the absolute values for dwell time reduction. In all cases we find a reduction in dwell time. The larger network has an average reduction of 19.2 hours and the smaller network 13.6 hours.

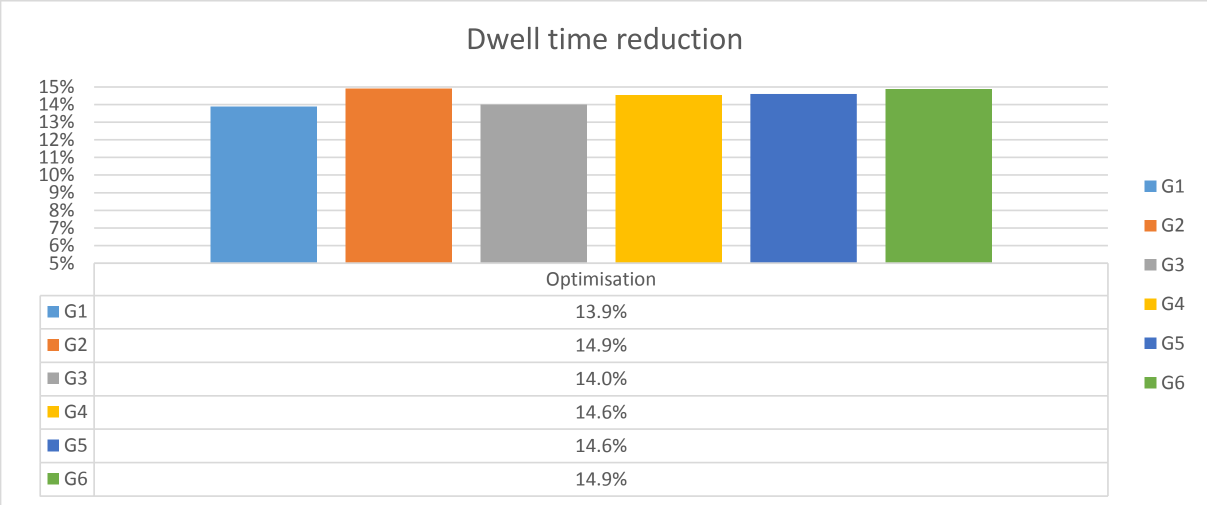

Figure 5.10: Average dwell time improvement in a 30 connection network.

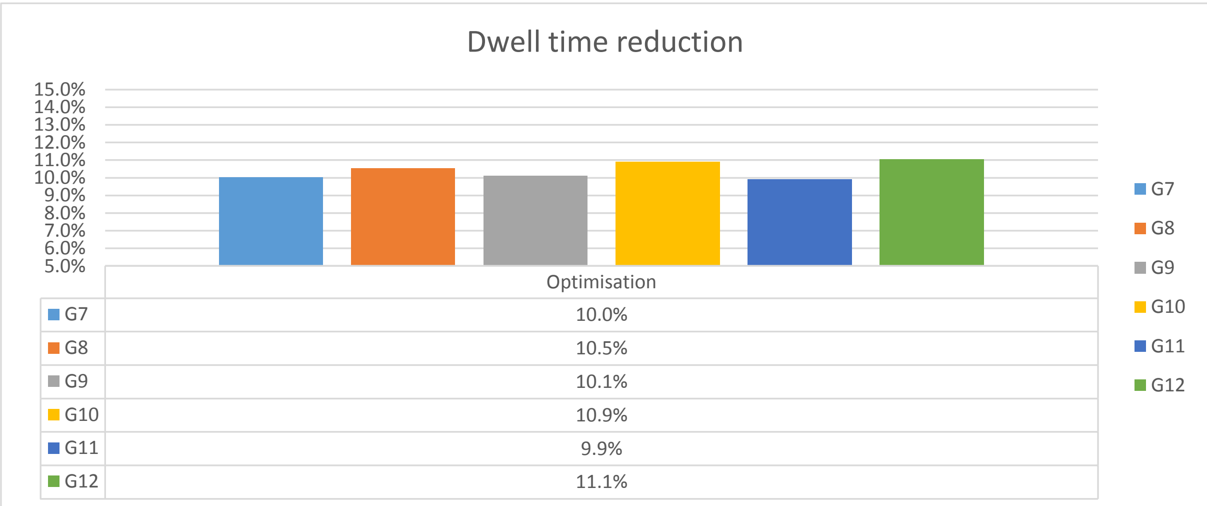

Figure 5.11: Average dwell time improvement in a 20 connection network.

Network	Average improvement (hours)	Network	Average improvement (hours)
G1	18.4	G7	13.1
G2	19.8	G8	13.7
G3	18.5	G9	13.2
G4	19.3	G10	14.1
G5	19.3	G11	13.0
G6	19.7	G12	14.3

Table 5-6: Dwell time improvement in hours.

Considering that the lead time shows no major improvements and dwell time decreases, it appears that orders are assigned to routes with a longer duration when using simulated annealing. The assignment to routes with longer durations could potentially have an effect on the total costs on the long term because containers are unavailable for a longer period of time. The longer unavailability of containers could then result in the need for more containers.

5.4.2 Generic networks with variable costs

As we mentioned in the introduction of Section 5.4, we address the variable cost networks only briefly. Because the results for the variable costs networks are less relevant with regard to the case.

Utilisation

Opposed to the fixed costs network, we find an average decrease in utilisation rate in the variable cost networks. We give the results Table 5-7.

	30 connections		20 connections
G1	2.8%	G7	18.9%
G2	14.0%	G8	20.1%
G3	2.9%	G9	18.6%
G4	13.0%	G10	19.8%
G5	4.0%	G11	18.0%
G6	16.0%	G12	21.1%

Table 5-7: Average utilisation rate reductions.

Costs

Contrary to what we expect, the cost results for the variable cost network are worse after optimisation for the 30 connection network. The main reason that we suspect for the behaviour is that a positive change (reduction in costs) for the planning during an optimisation has negative impact on the costs for the next few days. For example, it could be that an order is assigned to an intermodal route with a longer travel duration because this assignment yields better results cost wise. But due to this increase in travel duration the container that is used for this transport is unavailable for a longer period of time as well. Which in turn could force orders to be assigned to direct trucking due to unavailability of a container. We have already addressed the fact that orders get assigned to routes with a longer duration in Section 5.4.3. Therefore, the scenario described above is not unthinkable. We address the other reason in Section 5.4.1 (orders getting assigned to routes with a longer duration). We are looking at an increase in costs varying from 1.8% and 7.4% for the balanced and imbalanced networks respectively. The 20 connection network on the other hand shows a slight decrease in costs, 3.3% for balanced and 1.7% for imbalanced networks. The results are summarised in Table 5-8.

	30 connections		20 connections
G1	-1.8%	**G7**	3.2%
G2	-7.6%	**G8**	1.8%
G3	-1.5%	**G9**	3.4%
G4	-6.3%	**G10**	1.8%
G5	-2.1%	**G11**	3.3%
G6	-8.4%	**G12**	1.5%

Table 5-8: Average cost decrease.

5.4.3 Case network analysis

With regard to the case network, we consider the same aspects as in the fixed cost network, namely, network utilisation, costs, empty moves, and lead time. We distinguish between three different train schedules: one where the train does not stop at terminal 3 (Case 0), one where the train stops once per week at terminal 3 (Case 1), and one where the train stops twice per week at terminal 3 (Case 2).

Network utilisation

As we can see in Table 5-9, the utilisation rate of the long haul transport between the United Kingdom and Poland varies dependent on the applied train schedule. When 0 or 2 stops are made at terminal 3, we find a decrease in the utilisation rate compared to the best route without optimisation. In case of 1 stop at terminal 3, an increase in the utilisation rate is expected. In the case of 0 stops we have a decline in arrival orders that are sent through intermodal transport because there is no available route to terminal 3 (the train does not stop there). Further, in case of 2 stops at terminal 3 the capacity for containers is effectively doubled which can results in having too few orders to make train transport viable. As we see, none of the methods reaches an average utilisation rate of 90%. The 90% utilisation rate is the boundary that we found in literature, see Chapter 2, to make intermodal freight trains viable. With 84.8% utilisation, the network with 1 stop at terminal 3 performs best with regard to utilisation rate. Although the literature gave 90% utilisation rate is a boundary, in the costs section we see that with the utilisation rates presented here, intermodal transport is still viable.

Netwerk	Average utilisation without optimisation	Average utilisation with optimisation	Improvement
Case 0	87.9%	80.4%	-7.6%
Case 1	82.5%	84.8%	2.3%
Case 2	88.4%	83.5%	-4.9%

Table 5-9: Utilisation rate improvements.

In Table 5-10 we summarise the change in average number of orders that are sent by direct trucking. The case 1 network has no shocking changes, the case 0 and case 2 network see a large increase in the amount of orders that are being trucked directly.

This is most likely a result from the fact that in the case 0 case, the amount of orders available is too little to make the train transport viable every time. For the case 2 case the available capacity is most likely too high.

Netwerk	Before optimisation	After optimisation	Absolute increase
Case 0	9.0	10.6	1.6
Case 1	8.0	8.2	0.1
Case 2	7.1	8.7	1.6

Table 5-10: Changes in trucked orders per day.

Costs

The effect that optimisation has on the total costs for our case is less than we saw when considering the generic networks. One of the main reasons for this is the fact that the case network is a very limited network since we have 6 terminals with 9 connections in total. Further only a limited amount of long haul departures occur per week. The results are summarised in Figure 5.12, here we see that in case of 0 or 1 stop in terminal 3, a decrease of 1% and 4.2% respectively are expected. For the schedule with 2 stops in terminal 3 we find a slight increase in costs, namely 1.7%. Table G-7 in Appendix G summarises the absolute values corresponding to the cost changes. For the 0 stops and 1 stop in terminal 3 we find a decrease in costs. In case of 2 stops at terminal 3 an increase in costs is expected. Further, we see that 0 stops at terminal 3 has the best performance cost wise.

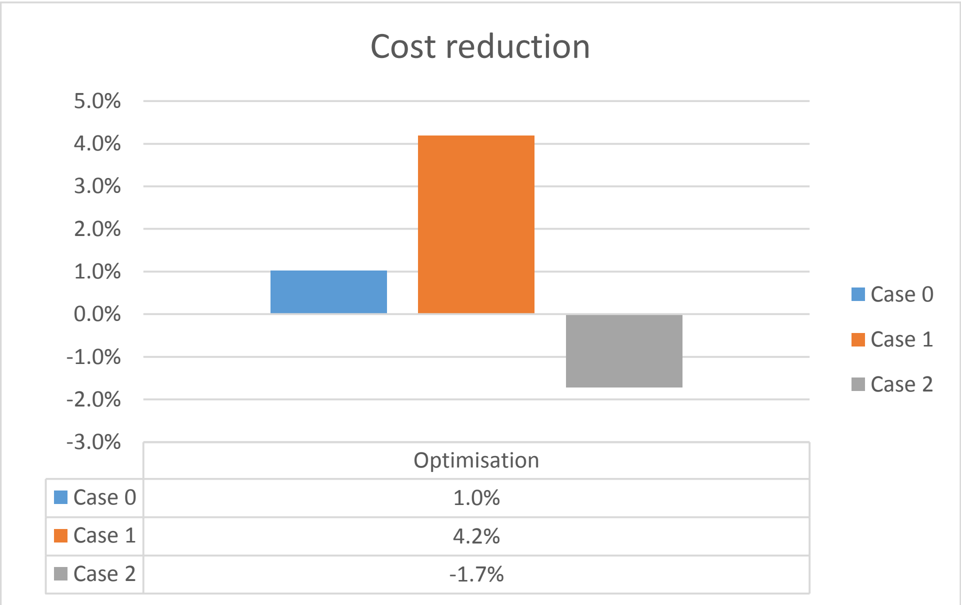

Figure 5.12: Average costs changes for the case networks.

Figure 5.13 gives insight into the division of costs. In line with the utilisation rate and average cost changes we see that for both the 0 stops and the 2 stops in terminal 3 an increase of truck costs is expected.

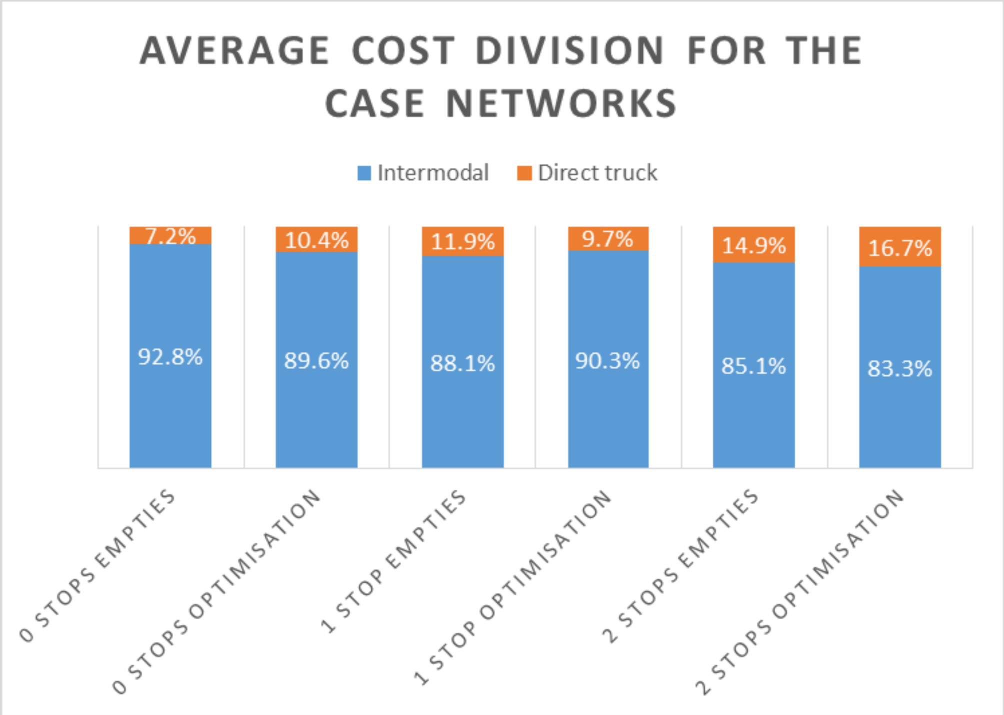

Figure 5.13: Average cost division.

Finally, comparing the optimisation heuristic to direct trucking, we find the results as shown in Table G-8 in Appendix G. In all cases our optimisation heuristic performs better than direct trucking.

Empties

As we can see in Figure 5.14, when we leave out the stop in terminal 3 we get a relative heavy decline in empty moves, on average we look at 9.7%. For one stop in terminal 3 we also see a decrease in empty moves (5%). Finally, two stops in terminal 3 result in an increase in empty moves by 2.4%. Table 5-11 shows the absolute values corresponding to the relative changes presented in Figure 5.14. When we consider the absolute values for empty moves, we look at a range of -0.3 to 1.6 containers per day, which we consider quite small.

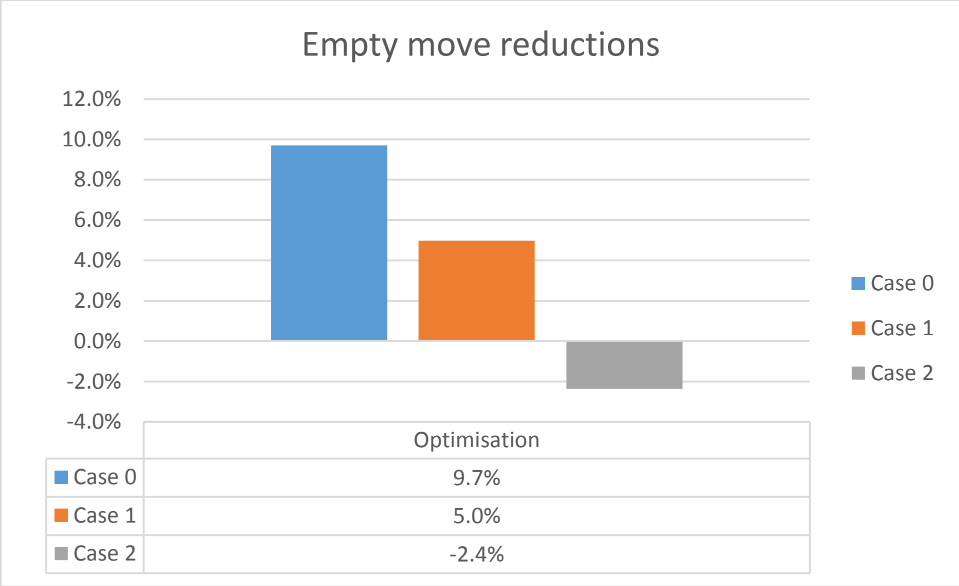

Figure 5.14: Average change in empty moves.

Network	average improvement (per day)
Case 0	1.6
Case 1	0.8
Case 2	-0.3

Table 5-11: Average number of containers moved per day.

Lead time

Just as with the generic networks we find little change in the total lead time of orders. On average we are looking at an increase of lead time by 1 to 2.2%. In the generic networks we found a decrease in dwell time, in the case networks we find an increase in dwell time. We are looking at an increase varying between 6.4 to 9.4%, see Figure 5.15. As we see in Figure 5.15 all values reductions are negative thus indicating an increase in dwell time. Reasons for the increase in the dwell time might be attributed to the point that we are very limited in the transport possibilities in our case network.

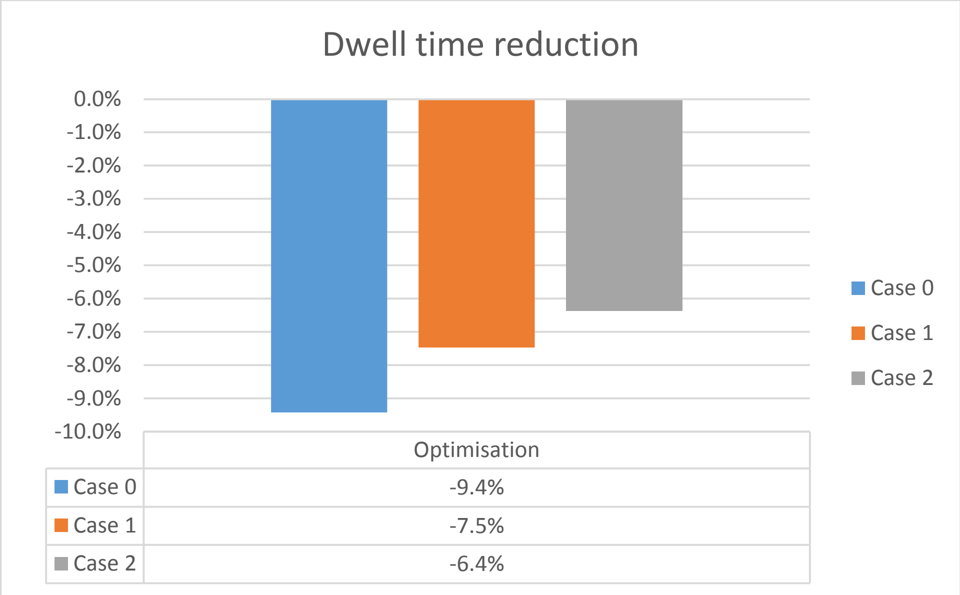

Figure 5.15: Average reduction in dwell time compared to the best route benchmark.

In Table 5-12 we show the absolute dwell time changes. As we see, for each network an increase of dwell time is expected. However, given the increase of approximately 6-8 hours on the total dwell time we see no reason to look at the dwell time in further detail.

Network	average improvement (hours)
Case 0	-6.4
Case 1	-6.2
Case 2	-5.1

Table 5-12: Absolute dwell time changes.

5.4.4 Sensitivity

We suspect that the tuning of capacity and arrival rate have an influence on the performance of our heuristic. Specifically the influence on utilisation rates and the number of orders that are assigned to direct trucking. To test this we perform additional tests with adjusted arrival rates. The tests are performed on both a balanced and an imbalanced network. We take the settings of the G1 and G2 network and adjust the arrival rate of these networks. The tests are performed with an arrival rates of -10%, -5%, +5%, and +10% of the original 50 orders per day. Table 5-12 summarises the results with regard to the utilisation rate, the average number of active legs, and the number of orders assigned to direct truck routes. When we increase the arrival rate of orders we see an increase in utilisation rate and active legs. However, we also see an increase in number of orders trucked which smaller than or equal to the increase in arrival rate.

Change in order arrival rate	-10%	-5%	0%	5%	10%
G1 utilisation without optimisation	69.6%	70.8%	72.1%	72.6%	73.5%
G1 utilisation with optimisation	83.4%	83.8%	84.3%	84.5%	85.0%
G1 average number of legs in use without optimisation	30	30.6	31.3	31.6	32.2
G1 average number of legs in use with optimisation	29.9	31.2	32.6	33.7	35
G1 orders trucked without optimisation	4.4	4.9	5.9	6.8	7.5
G1 orders trucked with optimisation	9.5	9.6	10.1	10.6	11.1
G2 utilisation without optimisation	66.1%	66.9%	67.8%	67.7%	68.8%
G2 utilisation with optimisation	78.5%	78.9%	79.6%	80.1%	80.4%
G2 average number of legs in use without optimisation	28.3	28.9	29.4	29.9	30.5
G2 average number of legs in use with optimisation	24.1	24.9	26.2	27.3	28.4
G2 orders trucked without optimisation	8.5	9.3	10.8	11.5	12.7
G2 orders trucked with optimisation	19.1	20.1	21	21.7	22.8

Table 5-13: Results with regard to the order arrival rate changes.

5.5 Conclusions

In this chapter we addressed our simulation model and methods with regard to creating generic networks. Further we tested the heuristics that we proposed in Chapter 4. The output from these tests have been analysed.

We find that our optimisation heuristic shows promising results for some fixed cost networks. Amongst others we find:

- Cost reductions varying between 11 to 17%.
- Utilisation rate increase on transport legs by 9 to 12%.
- Dwell time reduction by 10-15%
- Change in orders trucked varying between -3.5 to 95%.
- Reduction in empty containers moved varying between 38 to 47%.

Finally, when we apply our planning heuristic to the case network we find less impact on the change in costs. The cases are defined as zero stops in terminal 3, one stop in terminal 3, and two stops in terminal 3. We mainly attribute this to the fact that the case is a very limited network with 6 terminals, 9 connections, and limited long-haul departures.

6. Conclusions and discussion

In this chapter we conclude our research. In Section 6.1 we give our conclusion and answer our research question. Our research question is:

> ***In what way and up to what degree is it possible to support real-time consolidation and equipment repositioning decisions in a synchromodal environment?***

In Section 6.2 we discuss our model, heuristic, and results. Next, in Section 6.3 we give our recommendations. Finally, we address directions for further research in Section 6.4

6.1 Conclusions

The aim of this thesis is to develop a dynamic planning model for full truckload consolidation that supports consolidation, replanning, and empty container movements.

We addressed the shift from intermodal to synchromodal transport planning. As becomes clear from literature, synchromodal transport planning is the act of planning and scheduling intermodal transportation in a flexible fashion based on real-time information. Because we were not able to find a model in literature that covered consolidation decisions, replanning, and empty repositioning, we developed our own heuristic.

As the heuristic is applied on a case study, the network corresponding to the case is analysed. Further, we gathered data with regard to more generic settings to be able to assess the heurstic in a broader setting.

We developed a heuristic combined with an optimisation heuristic to improve synchromodal planning. Our heuristic is able to replan orders through a simulated annealing procedure. We defined four different methods of creating neighbour solutions.

Our optimisation heuristic covers the planning of orders with a best route policy, where best is defined as able to deliver before the due date against lowest possible costs. Next, it is able to replan orders through a SA annealing procedure where four different neighbour structures are used. Finally, a simplified method for the movement of empty containers is applied.

The heuristic is tested on both generic networks and on the case. The generic networks have been created at random in a plane covering an area that is approximately equal to the area in the case network. Care has to be taken as, although the two generic networks are created at random, no further effect of constantly altering networks is taken into account.

We see a decrease in costs when replanning of orders is considered in comparison to no replanning of orders. Key improvements on the various settings for the network that we find are:

- Relative cost reductions varying between 11 to 17%.
- Relative increase of utilisation rates on transport legs by 11 to 18%.
- Relative dwell time reductions of approximately 10%.
- Relative change in orders trucked varying between -3.5 to 95%.
- Relative reduction in empties moved varying between 38 to 47%.

Tables 6-1 and 6-2 summarise the absolute key results from our optimisation heuristic when it is applied on fixed cost generic networks.

Balanced networks	**30 connections**	**20 connections**
Costs for direct trucking (€ per day)	29,355	29,320
Costs before optimisation (€ per day)	25,274	29,459
Costs after optimisation (€ per day)	21,807	24,436
Cost reduction (€ per day)	3,467	5,022
Utilisation rate increase (%)	12.3	9.4
Dwell time reduction (hours)	18.8	13.1
Orders trucked decrease (# of orders per day)	-4.3	0.7
Empty moves decrease (# of containers per day)	2.2	1.1

Table 6-1: Key results for the balanced networks.

Imbalanced networks	**30 connections**	**20 connections**
Costs for direct trucking (€ per day)	29,364	29,339
Costs before optimisation (€ per day)	27,020	29,759
Costs after optimisation (€ per day)	24,019	25,567
Cost reduction (€ per day)	3,001	4,191
Utilisation rate increase (%)	11.9	10.2
Dwell time reduction (hours)	19.6	14.0
Orders trucked decrease (# of orders per day)	-10.3	-3.3
Empty moves decrease (# of containers per day)	2.4	0.8

Table 6-2: Key results for the imbalanced networks.

As we can see for the indicators cost reduction, utilisation rate increase, and dwell time reduction, our heuristic outperforms the benchmark. For the number of orders trucked we only find a decrease in one out of the 4 cases. Main reason for this behaviour is the fact that we have a reduction in number of legs used per day. This reduction in turn results in less capacity that is available. As we show in Section 5.4.4 the tuning of capacity and arrival rate has influence on this behaviour. Finally, in all cases we find a slight reduction in number of empty containers moved. The reduction in empty moves is questionable due to the fact that the imbalance rate does not change per comparison. However, as empty containers are assigned to routes that are in use, it is logical to see a reduction due to the fact that the utilisation rate is higher and in most cases the number of legs in use lower.

Next to the generic networks, we have the case network. We present the key results for the case network optimisation in Table 6-3 and Table G-9 in Appendix G, the latter of the two tables contains the results with regard to costs. The difference between the three networks (case 0, case 1, and case 2) is the train schedule with regard to number of stops at an intermediate hub, where case 0 corresponds to 0 stops, case 1 to 1 stop, and case 2 to 2 stops. As we see, if we consider two stops then the results are negative on all aspects. Both the 0 and 1 stop at the intermediate hub have a positive result with regard to costs. The effect of 0 stops is smaller than the 1 stop. However, the total cost before and after optimisation score in favour of the 0 stops setting. In all cases we find an increase in dwell time. The increase however is small when considering a lead time of 7 to 8 days.

	Case 0	Case 1	Case 2
Utilisation rate increase (%)	-8.6	2.8	-5.5
Dwell time reduction (hours)	-6.4	-6.2	-5.1
Orders trucked decrease (# of orders per day)	-1.6	-0.1	-1.6
Empty moves decrease (# of containers per day)	1.6	0.8	-0.3

Table 6-3: Key results for the different case network settings.

When we look at the results from the case network, we identify several aspects that could have major influence on the transport costs.

- The availability of containers
- The applied train schedule
- Capacity availability
- Order arrival

If there is a lack of containers, logically, orders get trucked. Considering that we see an increase in the number of orders trucked, it appears that either there is an actual lack of containers or the empty repositioning part of our heuristic is not as good as we initially expected. The latter might stem from the fact that we apply a decision rule rather than optimising the empty flow.

The other three points (applied train schedule, available capacity, and order arrival rate) are intertwined with each other. Performing a stop at an intermediate terminal effectively doubles the potential available capacity. Because from the start terminal we can transport only containers for the intermediate terminal and from the intermediate terminal everything for a destination terminal. If the order arrival rate cannot cope with this increase in capacity, we see an increase in orders being delivered by truck because this is potentially cheaper. Therefore, effectively deviating from the initial aim of increasing the intermodal utilisation rate through consolidating containers on long haul transport.

All in all we find that our heuristic is capable of improving a transport plan with positive results to both the costs and utilisation rate.

6.2 Discussion

In this section we discuss our model, heuristic, and results.

Because we use a model to approach reality several aspects that occur in real life have been left out due to the scope of this research. Amongst these aspects we identify disruptions during transport, travel times that are stochastic, customs, rules & regulations, and terminal operations. All of these aspects are able to influence the results presented in this thesis.

Most of the before mentioned aspects induce extra time with regard to transportation. The extra time is most likely not static and in some cases it is hard to predict the actual effect of the aspects. The unpredictable additional time results in variance in the actual door-to-door lead time. To prevent delivery after the due date, transportation at the latest departure time is not always the best option. However, if we plan orders at earlier departure times we might lose out on consolidation opportunities. Therefore, a trade-off between consolidation opportunities and risk of delivery after the due date should be considered.

Our model only focusses on long haul transport, whereas in practice pre and end haulage is involved as well. In our simulation model we addressed the pre and end haul time by drawing random numbers from a Weibull distribution that represent the pre and end haul time. These times are deducted from the allowable door-to-door lead time that are assigned to each order. Further, the containers that are used for transportation are unavailable during pre and end haul. In our simulation model we addressed this by making containers unavailable for a fixed period of time. In reality pre and end haul time and the unavailability of containers are related to each other.

Further in a generic setting we considered two fixed random networks whereas the effect of different networks with the same kind of settings (number of terminals or number of legs) are left out. Therefore the results presented in this thesis should be seen in that perspective. In theory this is easily solved by adding additional experiments. However, as we addressed in earlier chapters, computation time becomes an issue due to complexity of the simulation model. Therefore, we refrain from performing more experiments.

Finally, the procedures applied for selecting a neighbour solution in the optimisation heuristic do not cover all possibilities. We already addressed a few other options that could probably have an influence with regard to optimisation. Amongst them we mentioned:

- Moving multiple orders from one intermodal route to another intermodal route.
- Taking orders from multiple intermodal routes and place them on one different intermodal route.

If these types of swaps are considered it might be that our heuristic has a better performance.

6.3 Recommendations

As our planning heuristic is only aimed at the long haul (inter-terminal) transport, integrating the heuristic with pre and end haul planning is recommended to get insight in the interaction between long haul and pre- and end haul transport. We advise this because in our model these factors are simplified but may influence the outcome.

The case network that we consider is relative small when compared to a generic network. Further, due to the fact that the long haul between the United Kingdom and Poland only occurs twice a week, little to no freedom in alternative route choices is possible. Given the results from the case network alone it is questionable whether the heuristic should be applied to it. Mainly due to the fact that there are several assumptions that possible might not hold, e.g., 100% availability of truck transport (either containerised or non-containerised) or setting a fixed time period for the unavailability of containers after delivery at the destination terminal. However, considering the results on the larger generic networks we recommend to consider testing the heuristic as a shadow project on somewhat larger networks with more freedom in choices of routes.

Considering the three options we tested with regard to case network (0 stops, 1 stop, and 2 stops per week at an intermediate hub), we advise to consider either 0 or 1 stop per week at the hub. The 0 stops is advised due to the fact that it yields the best results cost wise, however major downside here is that it requires more truck transport. We consider the increase in truck transport unfavourable. The 1 stop yields more costs per day on average than the 0 stops variant, but it makes better use of intermodal transport. However, given the results, we do not advise to use the heuristic on the case network as it is now. If in the future the network becomes larger the use of the heuristic should be reconsidered.

Finally, given the results on variable cost networks we do not recommend to use the heuristic. The expected results are too little or even negative.

6.4 Further research

When looking at our research there are several aspects that require further research. In the total overview of network optimisation these aspects should be considered but due to the scope of our research they are left out.

In Section 1.3 we briefly addressed less than truckload (LTL) consolidation. In this thesis we considered every container to be a full truckload, this is not the case in practice. Therefore, the added effect of consolidating LTL containers before transporting them on the long haul may yield additional benefits.

Although the aim of our planning heuristic is to improve consolidated planning, we could potentially see other benefits of the planning heuristic. When considering the generic setting, applicability with regards to capacity reservation on a corridor is a possibility. In Section 5.3.2 we elaborated on the capacity settings. We deliberately choose to use lower capacity than trains and barges have in practice. With our heuristic it should in theory be possible to assess the effects with regard to tweaking capacity on a leg. A simple question that could be answered is:

- What is the effect on total costs if we increase or decrease capacity (reserve spots on a train or barge) on a certain leg.

As we consider only specific networks and therefore the conclusions cannot be generalised, more research should be performed on this.

As addressed in the recommendations, the interaction between long haul and pre and end haul transport is interesting to study further. An interesting point that should be considered is the determination of start and end terminals. Due to transport schedules on the long-haul transport it might be beneficial to send an order to a different start or end terminal than the terminal that is closest to the origin or destination.

Finally, the method that we apply for the empty repositioning is a decision rule rather than an optimisation policy. Improving this method might prove beneficial with regard to transport costs. Having an effective method for empty repositioning allows for more containers being available at the right place and time. The first step is to consider the empty containers that are already on their way to a terminal. This knowledge is available but our heuristic does not work with this knowledge. Further, we can look at probability theory. Probability theory might be useful in this case, where the decision criterion could be based on for example: the probability of a having a new order arriving at the current location versus moving the container to a terminal where there is a guaranteed load available.

Bibliography

Andersen, J., Crainic, T. G., & Christiansen, M. (2009a). Service network design with asset management: formulations and comparative analyses. *Transportation research part C: emerging technologies*, 197-207.

Behdani, B. (2013, December 13). Retrieved from http://www.slideshare.net/BehzadBehdani/synchromodal-freight-transport-system

Behdani, B., Fan, Y., Wiegmans, B., & Zuidwijk, R. (2014). Multimodal design for synchromodal freight transport systems.

Blum, C., & Roli, A. (2003). Metaheuristics in combinatorial optimzation: overview and conceptual comparison. *ACM Computing Surveys*, 268-308.

Bock, S. (2010). Real-time control of freight forwarder transportation networks by integrating multimodal transport chains. *European journal of operational research*, 733-746.

Bontekoning, Y. M., Macharis, C., & Trip, J. J. (2004). Is a new applied transportation research field emerging? - A review of intermodal rail-truck freight transportation literature. *Transportation research part A: Policy and practice*, 1-34.

Braekers, K., Caris, A., & Janssens, G. K. (2013). Optimal shipping routes and vessel size for intermodal barge transport with empty container repositioning. *Computers in industry*, 155-164.

Campbell, J. (1990). Freight consolidation and routing with transportation economies of scale. *Transportation research B*, 345-361.

Caris, A., Macharis, C., & Janssens, G. K. (2008). Planning problems in intermodal freight transport: accomplishments and prospects. *Transportation planning and technology*, 277-302.

Christiansen, M., Fagerholt, K., & Ronen, D. (2004). Ship routing and scheduling: status and perspectives. *Transportation science*, 1-18.

Crainic, T. G., & Laporte, G. (1997). Planning models for freight transportation. *european journal of operational research*, 409-438.

Crainic, T. G., & Rousseau, J. M. (1986). Multicommodity, multimode freight transportation: a general modeling and algorithmic framework for the service network design problem. *transportation research part B: methodological*, 225-242.

Crainic, T. G., Gendreau, M., & Dejax, P. (1993). Dynamic and stochastic models for the allocation of empty containers. *Operations research*, 102-126.

Crainic, T. G., Hewitt, M., Toulouse, M., & Vu, D. M. (2014). Service network design with resource constraints. *Transportation science*.

Dang, Q. V., Yun, W. Y., & Kopfer, H. (2012). Positioning empty containers under dependent demand process. *Computers & industrial engineering*, 708-715.

den Boer, E., Otten, M., & van Essen, H. (2011). *Comparison of various transport modes on a EU scale with the STREAM database.* Delft: CE Delft.

DINALOG. (2015). *Dinalog*. Retrieved May 11, 2015, from http://www.dinalog.nl/en/themes/synchromodal_transport/

European container terminal. (2011, October). Retrieved from ect: http://www.ect.nl/en/content/future-freight-transport

European Parliament. (2006, 03 15). *Driving time and rest periods*. Retrieved from Mobility and transport: http://ec.europa.eu/transport/modes/road/social_provisions/driving_time/index_en.htm

Hall, R. W. (1987). Consolidation strategy: inventory, vehicles and terminals. *Journal of business logistics*.

Ham, P. (2012, November). Retrieved from European Gateway Services: http://www.managementproducties.com/files/8613/5469/9178/Paul_Ham.pdf

Higginson, J. K., & Bookbinder, J. H. (1994). Policy recommendations for a shipment-consolidation program. *Journal of Business Logistics*, 87-112.

Janic, M., Regglani, A., & Nijkamp, P. (1999). Sustainability of the European freight transport system: evaluation of innovative bundling networks. *Transportation planning and technology*, 129-156.

Krauth, E., Moonen, H., Popova, V., & Schut, M. (2005). Understanding performance measurement and control in third party logistics. *Proceedings of the 13th European Conference on Information Systems, Information systems in a rapidly changing economy.*

Kreutzberger, E. D. (1999). *Innovative networks and new-generation terminals for intermodal transport.* Delft: Delft University of Technology.

Kreutzberger, E. D. (2010). Lowest cost intermodal rail freight transport bundling networks: conceptual structuring and identification. *European journal of transport and infrastructure research*.

Lai, K. H., Ngai, E. W., & Cheng, T. C. (2004). An empirical study of supply chain performance in transport logistics. *International journal of production economics*, 321-331.

Law, A. M. (2007). *Simulation modeling & analysis.* New York: McGraw-Hill.

Lin, C. C., & Chen, S. H. (2004). The hierarchical network design problem for time-definite express common carriers. *Transportation research part B: methodological*, 271-283.

Lium, A. G., Crainic, T. G., & Wallace, S. W. (2009). A study of demand stochasticity in service network design. *Transportation science*, 144-157.

Meng, Q., & Wang, S. (2011). Liner shipping service network design with empty container repositioning. *Transportation research part E*, 695-708.

Mes, M. R., van der Heijden, M. C., & van Harten, A. (2007). Comparison of agent-based scheduling to look-ahead heuristics for real-time transportation problems. *European journal of operational research*, 59-75.

Newman, A. M., & Yano, C. A. (2000). Centralized and decentralized train scheduling for intermodal operations. *IIE Transactions*, 743-754.

Pedersen, M. B., & Crainic, T. G. (2007). Optimization of intermodal freight train service schedules on train canals. *CIRRELT*.

Pedersen, M. B., Crainic, T. G., & Madsen, O. B. (2009). Models and tabu search metaheuristics for service network design with asset-balance requirements. *Transportation science*, 158-177.

Port of Rotterdam Authority. (2015, April 16). Retrieved from Port of Rotterdam Authority: http://www.portofrotterdam.com/en/News/pressreleases-news/Pages/strong-quarter-port-rotterdam.aspx

Powell, W. B. (2007). *Approximate dyanmic programming: solving the curses of dimensionality.* John Wiley & Sons.

Rivera, A. P., & Mes, M. R. (2015). Dynamic multi-period freight consolidation.

Roth, M., Klarmann, A., & Franczyk, B. (2013). Future logistics - challenges, requirements and solutions for logistics networks. *international journal of mechanical, aerospace, industrial and mechatronics engineering*, 898-903.

Song, D. P., & Dong, J. X. (2013). Long-haul liner service route design with ship deployment and empty container repositioning. *Transportation research part B: methodological*, 188-211.

Spikker, D. (2014). *Planning synchromodal transport at a logistics service provider.* Enschede: University of Twente.

SteadieSeifi, M., Dellaert, N. P., Nuijten, W., van Woensel, T., & Raoufi, R. (2014). Multimodal freight transportation planning: A literature review. *European journal of operational research*, 1-15.

Stuurgroep synchromodaliteit. (2013, February 25). *Platform synchromodaliteit.* Retrieved from http://www.synchromodaliteit.nl/wp-content/uploads/2013/08/ECT_SynchroBoekje_1t_LR.pdf

Stuurgroep synchromodaliteit. (2015). *Platform synchromodaliteit*. Retrieved May 11, 2015, from http://www.synchromodaliteit.nl/nl/definitie/

Theofanis, S., & Boile, M. (2009). Empty marine container logistics: Facts, issues and management strategies. *GeoJournal*, 51-65.

Unit 45. (2015). *Unit 45 intermodal innovators*. Retrieved from http://www.unit45.com/

van Heeswijk, W. J., Mes, M. R., Schutten, J. M., & Zijm, W. H. (2014). Freight consolidation in networks with transshipments. *Beta working paper series 462*.

van Riessen, B., Negenborn, R., Dekker, R., & Lodewijks, G. (2013). Impact and relevance of transit distrubances on planning in intermodal container networks. *Econometric insitute research papers*, 1-23.

Verweij, C. A. (2011, December). *Buck consultants international.* Retrieved from http://www.bciglobal.com/publicaties_detail.asp?dc=26162&cat=5022&hl=synchromodaal

Verweij, C. A. (2013, January 31). *Buck consultants international*. Retrieved from http://www.bciglobal.com/news_detail.asp?cat=25002&dc=26325

Yun, W. Y., Lee, Y. M., & Choi, Y. S. (2011). Optimal inventory control of empty containers in inland transportation system. *International journal of production economics*, 451-457.

Zhu, E., Crainic, T. G., & Gendreau, M. (2014). Scheduled service network design for freight rail transportation. *Operations research*, 383-400.

List of abbreviations

ADP	Approximate dynamic programming
CT	Control tower
DP	Dynamic programming
FTL	Full truckload
ILP	Integer linear programming
KPIs	Key performance indicators
LP	Linear programming
LSP	Logistics service provider
LTL	Less than truckload
MILP	Mixed integer linear programming
NFP	Network flow problem
SCT	Synchromodal control tower
SND	Service network design
TEU	Twenty-foot equivalent unit
VRP	Vehicle routing problem
3PL	Third party logistics

Appendix A: ILP solution method for SND

For the standard SND the following parameters and decision variables are defined.

Indices

$i: set\ of\ nodes\ i\epsilon N$

$k: set\ of\ commodities\ k\epsilon K$

Further for each node we define the sets

$\delta^+(i) = \{(j, j')\epsilon A: j = i\}$

$\delta^-(i) = \{(j, j')\epsilon A: j' = i\}$

These are the sets with tail and head in node i, respectively.

Parameters

$c_{ij}^k: costs\ for\ sending\ one\ unit\ of\ commodity\ k\ along\ edge\ i, j$

$u_{ij}: capacity\ for\ edge\ i, j$

$b^k: amount\ of\ commodity\ k\ available\ for\ transport$

$o(k): origin\ node\ of\ commodity\ k$

$d(k): destination\ node\ of\ commodity\ k$

$f_{ij}: fixed\ cost\ for\ the\ use\ of\ edge\ i, j$

$$b(i, k) = \begin{cases} b^k & if\ i = o(k) \\ -b^k & if\ i = d(k) \\ 0 & otherwise \end{cases}$$

Decision variables

$x_{ij}^k: amount\ of\ commodity\ k, sent\ along\ edge\ i, j$

$y_{ij}: binary\ variable\ indicating\ whether\ or\ not\ edge\ i, j\ is\ used$

The SND is then described by the following linear model.

$$\min \sum_{(i,j)\epsilon A} \sum_{k\epsilon K} c_{ij}^k x_{ij}^k + \sum_{(i,j)\epsilon A} f_{ij} y_{ij} \quad (1)$$

$$s.t. \sum_{(j,j')\epsilon\delta^+(i)} x_{jj'}^k - \sum_{(j,j')\epsilon\delta^-(i)} x_{jj'}^k = b(i, k), i\epsilon N, k\epsilon K \quad (2)$$

$$\sum_{k\epsilon K} x_{ij}^k \leq u_{ij} y_{ij}, (i, j)\epsilon A \quad (3)$$

$$x_{ij}^k \geq 0, (i, j)\epsilon A, k\epsilon K \quad (4)$$

$$y_{ij}\epsilon\{0,1\}, (i, j)\epsilon A \quad (5)$$

The objective function (1) minimises the costs for using edges and transporting goods on an edge. Constraint (2) ensures that all goods are delivered. Constraint (3) forces capacity for an edge i,j to 0 if it is not used. Finally, constraints (4) and (5) are sign restricting constraints.

Appendix B: Statistical testing

This appendix covers the different statistical tests performed throughout this thesis.

One-sided t-test

We would like to know if the orders in the period December-June are significantly lower than the orders in the period July-November. To test this we calculate the mean and standard deviation per group. This is given in Table B-1.

	Dec-June (Group 1)	July-Nov (Group 2)
Mean	858.1	1007.6
Standard deviation	36.5	38.5
N	7	5

Table B-1: Statistics per group.

To test whether there is a significant difference, we define our null and alternative hypothesis as follows.

H_0: $\bar{X}_1 = \bar{X}_2$ the averages of both groups are equal.

H_a: $\bar{X}_1 < \bar{X}_2$ group 1 has a lower average than group 2.

Our Test-value T is given by:

$$T = (\bar{x}_1 - \bar{x}_2) / \sqrt{\left(\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}\right)}$$

This gives a T-value of -6.8, further there are 10 degrees of freedom. We reject H_0 if $T < -t_{\alpha,10}$ and thus have a significant difference.

Confidence level	0.95	0.975	0.99
T	-6.8	-6.8	-6.8
$-t_{\alpha,10}$	-1.8125	-2.2281	-2.764
Significant	Yes	Yes	Yes

Table B-2: One-sided t-test.

Table B-2 gives the results of the one-sided t-test. From this we can conclude that at all three confidence levels we reject the null-hypothesis. Therefore we conclude that there is a difference, in the average number of orders per month, between the two groups. Therefore we assume that there are seasonality effects present in the data.

Chi-square tests

We want to know if the assumption of a Poisson arrival processes for the orders is justified. Therefore we perform a chi-square test on both the data for the low and the high season, group 1 and group 2 respectively. The observed data (O_i) per month is presented in Table B-3.

Months Group1	**Orders Group 1**	**Months Group 2**	**Orders Group 2**
Dec	883.5	**July**	996.4
Jan	843.0	**Aug**	947.9
Feb	869.7	**Sep**	1044.9
March	880.5	**Oct**	1036.4
April	781.3	**Nov**	1012.1
May	868.2		
June	880.5		

Table B-3: Observed orders per month, divided by group.

Our Test-value χ^2 is given by:

$$\chi^2 = \sum_{i=1}^{N} \frac{(E_i - O_i)^2}{E_i}$$

The expected value (E_i) for group 1 is 858.1 orders per month, for group 2 this is 1007.6. This results in a χ^2-value of 9.33 and 5.89 for groups 1 and 2 respectivly.

Confidence level	0.95	0.975	0.99
χ^2	9.33	9.33	9.33
$\chi^2_{\alpha,6}$	10.645	12.592	16.812
Significant	No	No	No

Table B-4: Chi-square for group 1.

Confidence level	0.95	0.975	0.99
χ^2	5.89	5.89	5.89
$\chi^2_{\alpha,4}$	7.779	9.488	11.345
Significant	No	No	No

Table B-5: Chi-square for group 2.

Tables B-4 and B-5 give an overview of the chi-square test. In both cases we can conclude that no evidence is found to doubt the assumption of Poisson distributions for the arrival processes. Therefore we assume that the low season has a Poisson(858.1) and the high season a Poisson(1007.6) distribution per month.

Appendix C: Distribution fitting absence of data

For the door-to-door lead time and the total weight of an order we lack data for proper distribution fitting. Therefore we turn to heuristic procedures proposed in (Law, 2007).

Triangular distribution

The easiest distribution we can use is the triangular distribution. Figure C.1 shows the general concept of the triangular distribution. As we can see, it is determined by the three input parameters a, b, and m. Here a and b represent the lower and upper bound whilst m is the mode.

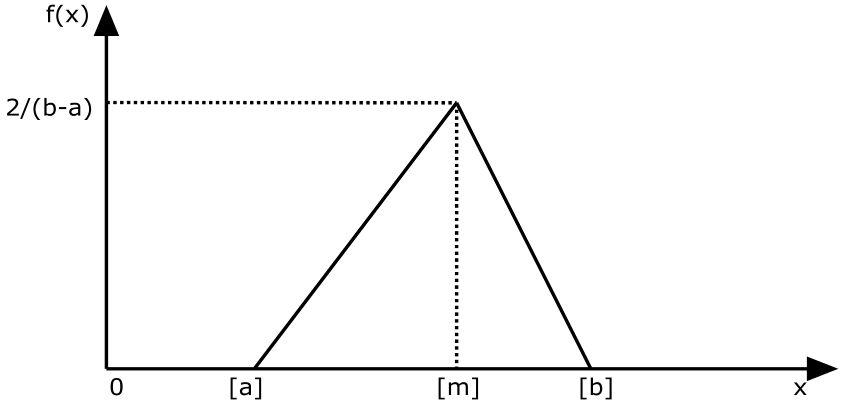

Figure C.1: Triangular distribution

Law (2007) addresses two downsides to the triangular distribution. The values for a and b are minimum and maximum estimated values. The question however is whether these minimum and maximum values are valid for a short or long time period. Another problematic issue with the triangular distribution is the fact that it cannot have a long right tail, this becomes relevant when modelling times to perform a task.

We believe that the above mentioned issues should be of no concern in our case. We are looking at weight distribution of freight and delivery time periods, which are not expected to change in the foreseeable future. We determine the actual values for a, m, and b by consulting experts.

Appendix D: Distribution fitting

To fit the data with regard to pre= and end-haul travel times, we made use of the software program EasyFit. Depending on the goodness of fit test used, we get different suitable distributions. The three applied goodness of fit tests are: Kolmogorov Smirnov (KS), Anderson Darling (AD), and Chi-Squared (CS). Further, the data for the CS fitting has been divided using both equal bin width and equal probability. Only continuous non-negative distributions are considered as we are dealing with travel times. Table D-1 gives the top 5 distributions for each of these tests. For the KS and the CS tests the p-values are listed as well.

Rank	**KS**	**p-value**	**AD**	**CS (equal bin width)**	**p-value**	**CS (equal probability)**	**p-value**
1	Log-Logistic (3P)	0.39403	Inv. Gaussian (3P)	Weibull (3P)	0.54862	Frechet (3P)	0.58404
2	Frechet (3P)	0.35845	Fatigue Life (3P)	Gen. Gamma (4P)	0.51027	Inv. Gaussian	0.51252
3	Lognormal (3P)	0.30748	Lognormal (3P)	Weibull	0.5033	Log-Logistic (3P)	0.44535
4	Inv. Gaussian (3P)	0.29741	Frechet (3P)	Gamma (3P)	0.50131	Lognormal (3P)	0.03772
5	Fatigue Life (3P)	0.26657	Log-Logistic (3P)	Gamma	0.4957	Gen. Gamma (4P)	0.02388

Table D-1: Distribution rankings

As Plant Simulation does not have all of the above mentioned probability distributions embedded, in special with regard to the 3 or 4 parameters, we have only limited options. Therefore we choose the Weibull distribution with parameters $\alpha = 1.1436$ and $\beta = 1.4873$.

Appendix E: Network details

Tables E-1 and E-2 show the connections between the terminals for both the 20 and 30 connection networks, a “1” indicates a rail connection and a “2” indicates a water connection. Table E-3 shows the X and Y coordinates of each terminal.

	T 1	T 2	T 3	T 4	T 5	T 6	T 7	T 8	T 9	T 10
T 1										
T 2										
T 3										
T 4	2	1								
T 5	2		1							
T 6	1	1			1					
T 7	1			1	1	2				
T 8	1		1				1			
T 9		1				1	2			
T 10				1		1		1		

Table E-1: Connections for the 20 connection network

	T 1	T 2	T 3	T 4	T 5	T 6	T 7	T 8	T 9	T 10
T 1										
T 2										
T 3	1	1								
T 4		1	1							
T 5		2	1							
T 6	1	1	1	1	2					
T 7	1			1	1	1				
T 8		2			1	1	1			
T 9				1	1	1				
T 10	1	2		2	1	1	1	1	1	

Table E-2: Connections for the 30 connection network

	X	Y		X	Y
terminal 1	299	5	**terminal 6**	843	188
terminal 2	404	201	**terminal 7**	939	4
terminal 3	580	234	**terminal 8**	1154	212
terminal 4	721	77	**terminal 9**	1324	244
terminal 5	802	168	**terminal 10**	1387	80

Table E-3: X and Y coordinates for the 10 terminal network.

In Tables E-4 and E-5 we give the details for the capacity on the legs in both the 20 and 30 connection network.

	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10
T1			2			3	3			3
T2			3	3	5	3		6		5
T3	2	3		2	2	3				
T4		3	3			3	3		3	5
T5		5	2			5	3	3	3	3
T6	3	3	3	3	4		3	3	2	3
T7	3			3	3	3		3		3
T8		6			3	3	2			3
T9				3	3	3				3
T10	3	5		5	3	3	3	3	3	

Table E-4: Capacity in the 30 connection network.

	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10
T1				5	6	3	3	3		
T2				2		2			2	
T3					2			3		
T4	4	3					2			2
T5	5		2			3	3			
T6	3	2			2		6		3	2
T7	2			2	3	4		2	6	
T8	2		2				2			3
T9		3				3	4			
T10				3		2		2		

Table E-5: Capacity in the 20 connection network.

Appendix F: Simulation settings

This appendix covers settings with regards to warm up period, batch correlation, and the number of runs, all of the settings were determined using the 10 terminal, 30 connection network using the "best" route benchmark policy. We decide to use this network as it has a larger variety of routes than its 20 connection counterpart. Further, as the "best" route benchmark policy does not optimise the planning we expect a larger variance in costs.

Warm up

For the three different methods, direct trucking, best route, and best route with optimisation we determine the transport costs per day. Next, we apply the graphical procedure of Welch where we take different widths for the average. This is graphically depicted in Figure F.1

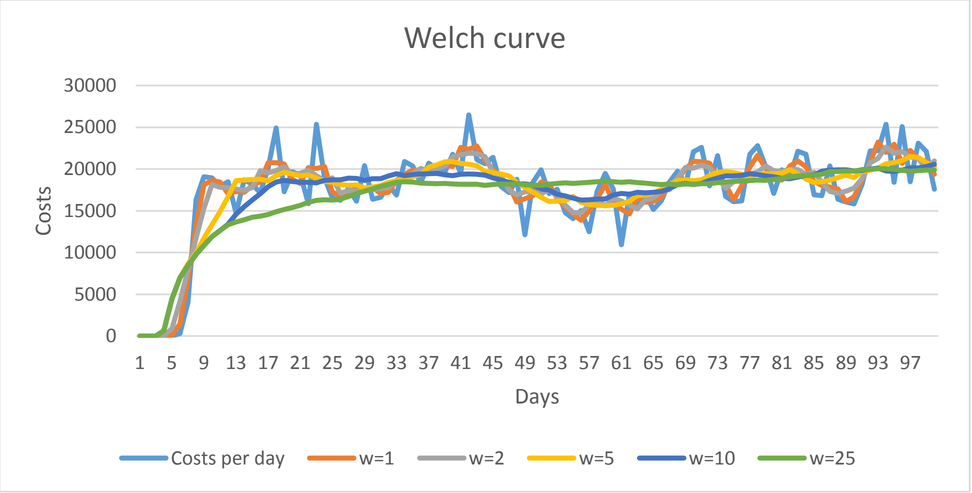

Figure F.1: Moving average for best route with empty repositioning.

As we see in Figure F.1, the first five days show some interesting behaviour. Namely, the average costs are approximately 0. These costs are 0 at for those five days due to the fact that our heuristic plans orders at the last point in time. This results in no transport costs for the first few days. Further we see that after a period of approximately 30 days, the costs per day becomes more or less stable. Therefore for each of our simulation experiments a period of 30 days is used as warmup.

Correlation batch means

To check whether our batch means are uncorrelated we apply the following formulas to estimate the correlation coefficient:

$$\hat{C}_j = \frac{\sum_{i=1}^{n-j}[X_i - \bar{X}_{(n)}][X_{i+j} - \bar{X}_{(n)}]}{n-j}$$

$$\hat{\rho}_j = \frac{\hat{C}_j}{S^2(n)}$$

Data with regard to the calculation of the estimated correlation coefficient is summarised in Table F-1.

We find an estimated correlation coefficient of $\hat{\rho}_j = -0.143$. As the correlation coefficient should lie within the range of [-1,1], where -1 is negative, 1 is positive, and 0 is no correlation. We argue that the estimated correlation coefficient is relatively low, but a relative weak negative correlation between batches is present.

avg	598194.27
var	821100210
n	33
j	1

Table F-1: summarised data for correlation estimation.

Number of runs

In table F-2 data with regards to the average number of orders delivered per run is summarised. Based on this data we can calculate the required number of runs. We apply the following procedure to determine for which value of n the equation in Section 5.3 holds.

I. Perform $n_0 > 2$ replications and set $n = n_0$

II. Compute $\bar{X}_n$, S_n^2 and the confidence interval half-width $\delta(n, \alpha) = t_{n-1,1-\alpha/2}\sqrt{\frac{S_n^2}{n}}$

III. If $\frac{\delta(n,\alpha)}{\bar{X}_n} \leq \gamma'$, then stop, confidence interval is $[\bar{X}_n - \delta(n, \alpha), \bar{X}_n + \delta(n, \alpha)]$

Else

- Set n := n + 1
- Perform another replication
- Go to step II

X	n	Average	Variance	t(n-1,1-a/2)	Confidence half width	relative error	y'
552105.5	1						
548809.5	2	550457.5	5431808	12.706	20939.488	0.038	0.024
618039.9	3	572984.9667	1525176167	4.303	97022.112	0.169	0.024
565454.89	4	571102.4475	1030959625	3.182	51084.719	0.089	0.024
551127.73	5	567107.504	853017586.7	2.776	36258.802	0.064	0.024
633252.79	6	578131.7183	1411613879	2.571	39435.233	0.068	0.024
564647.06	7	576205.3386	1202321472	2.447	32069.726	0.056	0.024
622724.86	8	582020.2788	1301069496	2.365	30160.352	0.052	0.024
597162.22	9	583702.7167	1163911185	2.306	26223.942	0.045	0.024
597239.93	10	585056.438	1052913335	2.262	23210.736	0.040	0.024
605941.69	11	586955.0973	987275978.7	2.228	21107.565	0.036	0.024
622569.45	12	589922.96	1003222127	2.201	20124.633	0.034	0.024
613477.46	13	591734.8446	962298319.3	2.179	18747.373	0.032	0.024
591605.91	14	591725.635	888276559.1	2.16	17205.363	0.029	0.024
620028.77	15	593612.5107	878232730.1	2.145	16412.944	0.028	0.024
599624.25	16	593988.2444	821942694.6	2.131	15273.699	0.026	0.024
619880.14	17	595511.2971	810005997.3	2.12	14633.746	0.025	0.024
598421.4	18	595672.9694	762829069	2.11	13735.985	0.023	0.024

Table F-2: Determination of the required number of runs.

As we can see, after 18 runs we have a relative error smaller than $\gamma' = 0.024$, resulting in a confidence interval of [581936.98, 609408.95].

Appendix G: Confidential information

In this appendix all the tables that are considered confidential are shown.