EELCO VAN ASPEREN

Essays on Port, Container, and Bulk Chemical Logistics Optimization

Essays on Port, Container, and Bulk Chemical Logistics Optimization

Essays on Port, Container, and Bulk Chemical Logistics Optimization

Essays over Optimalisatie van Haven, Container
en Bulk-Chemicaliën Logistiek.

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Erasmus Universiteit Rotterdam
op gezag van de rector magnificus

Prof.dr. H.G. Schmidt

en volgens het besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op

woensdag 18 november 2009 om 11.30 uur

door

EELCO VAN ASPEREN
geboren te Rotterdam.

ERASMUS UNIVERSITEIT ROTTERDAM

Promotiecommissie

Promotor: Prof.dr.ir. R. Dekker

Overige leden: Prof.dr.ir. U. Kaymak
Prof.dr. H. Haralambides
Prof.dr. S. Voß

Erasmus Research Institute of Management - ERIM
Rotterdam School of Management (RSM)
Erasmus School of Economics (ESE)
Erasmus University Rotterdam
Internet: http://www.erim.eur.nl

ERIM Electronic Series Portal: http://hdl.handle.net/1765/1

ERIM PhD Series in Research in Management, 181
Reference number ERIM: EPS-2009-181-LIS
ISBN 978-90-5892-222-9

DAI: 256946884

Design: B&T Ontwerp en advies www.b-en-t.nl
Print: Haveka www.haveka.nl

This book was typeset using the LaTeX document creation environment. The typeface for the running text is Palatino; Helvetica is used for headings. The figures and diagrams have been created using GNUplot, Metapost, and Microsoft Visio.

Acknowledgements

This thesis is the result of a prolonged effort during which many people have been involved. With this result in mind, I want to take this opportunity to acknowledge their contributions and thank them.

First and foremost, I want to thank my supervisor Rommert Dekker. Your dedication and inspiration have been essential to get this thesis finished. I have enjoyed our many, long discussions and hope that this thesis will not be the end of those.

Second, I extend a warm word of thanks to Arie de Bruin and Henk de Swaan Arons. Arie, you took me onboard in an academic role at the Computer Science department and the first couple of years were very inspirational. I have thoroughly enjoyed them. Henk, your enthusiasm for simulation and its applications was my main motivator to join the department. It has been a pleasure to work together and I particularly appreciate your efforts to coach me in doing research and finding a good focus for this thesis.

I have enjoyed being a member of the academic departments of Computer Science and, since 2005, Econometrics. In particular, I have enjoyed working with Jan van den Berg, Cor Bioch, Csaba Boer, Emiel Caron, Richard Paap, Wim Pijls, Mark Polman, Rob Potharst, Harry Trienekens, and Michiel van Wezel. Cia Scholte, as the secretary of the Computer Science department, provided the glue to keep us together as a department and as a group of individuals. In recent years I have appreciated the collaboration with Uzay Kaymak, in teaching, management, and research.

Several of the essays in this thesis are the result of collaboration with students, in particular students of the Informatics & Economics programme. As a group, they have been an inspiration and a challenge. Individually, I want to acknowledge Joost Hengstmengel, Ludo Waltman, Nees Jan van Eck, Bram Borgman (E&I), and Geerten Ochtman (Econometrics). Without your contributions, these projects might not have happened. Students of Informatics & Economics, and in particular those that were active in the study association Vrisbi, have contributed indirectly by creating an academic environment that was both challenging and fun.

Working on a PhD thesis inevitably takes a toll on your personal life. I apologize to friends and family for not being there at times and I hope to make up for it. Reino and Jolande de Boer, thank you for being such kind friends.

I thank my parents, Adrie and Hans, and my brothers Jeroen and Martijn, for their support and love.

Eelco van Asperen
Barendrecht, August 2009

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The essays in this thesis are concerned with two main themes. The first theme is the coordination of transport arrivals with the usage of the goods, the distribution processes, and the use of storage facilities. We will look at this for both containerized and bulk chemical transport, in chapters 2, 3, 4, 5 (for containerized transport), and chapter 7 (for bulk chemicals). The second theme is the uncertainty associated with the arrival time of ships with bulk chemicals and the impact on (port) logistics. This issue will be discussed in chapters 6 and 7. Both themes deal with difficult problems: a lack of coordination and information exchange between the parties involved leads to an increase in the amount of inventory required to meet demand.

In this chapter we will outline the context of the research in this thesis. We will first discuss maritime transport for containers and bulk chemicals. Next, the main actors involved in maritime transport are presented. As much of the research in this thesis is based on containerized transport, we will take a closer look at that topic and review the research issues. This is followed by a review of the research methods that are commonly used in this area. Decision making depends on the availability of accurate and timely information; we will discuss the role of information technology as a key enabler for making information available to parties in a supply chain. Finally, we will provide an outline for the remainder of the thesis.

1.1 Maritime Transport

Maritime transport plays a key role in global supply chains. Raw materials such as iron ore and crude oil are transported in bulk from their place of origin to processing facilities such as blast-furnaces or refineries. Processed goods are used by manufacturers to produce finished goods, with production facilities often located away from the main markets for the end product. The end products have to be transported from the factory to the markets. Maritime transport provides both the capacity (in case of iron ore and crude oil) and a low cost-per-unit for these transports. Air transport is limited by space and weight constraints and has a much higher cost, albeit with a higher speed. Land transport by truck or train may not be available (in case of intercontinental transports) and may also have a higher cost (both in terms of economic costs and, in case of truck transports, in terms of environmental footprint).

Many supply chains use standardized sea-containers. (The following description is based in part on Steenken et al. (2004).) Products that are shipped using these containers range from consumer electronics (TV's, DVD-players) to household appliances, clothing, and sports shoes. The processes involved consist of preparing the container (loading) before transport and then transporting the container to a port. The container terminal is a dedicated facility within a port for handling containers to and from sea-going vessels. Here, a yard or stacking area is used to decouple the land-side transport from the sea-side transport; the incoming container is temporarily stored in the container stack until its destination vessel has arrived. Once this vessel has arrived and is ready to load, the containers destined for this vessel are removed from the stack, transported to the quay, and loaded into the vessel. Once loaded, the ship can sail to its next destination port. Ships typically have a regular schedule with a number of ports to call. (For example, a very large container ship such as the Emma Maersk followed the following schedule in mid 2008: Yantian, Shanghai, Ningbo, Hong Kong, Yantian, Suez Canal, Algeciras, Rotterdam, Bremerhaven, Algeciras, Suez Canal, Yantian.) Upon arrival at the next port, the vessel berths at a container terminal and the unloading/loading operations can commence. If the container has reached its destination port, it will be offloaded and stored in the stack of the terminal. Using the stack for temporary storage separates the process of onward transport from the unloading operations of the vessel and ensures that any disruptions in the onward transport do not negatively effect the (time-critical) unloading operations. It also allows for administrative handling and customs processing. Onward transportation from the container terminal to the hinterland can involve multimodal transport such as train or inland barge. The advantages of using trains or barges rather than trucks are economies of scale (reduced cost per container) and a lower impact on the environment, albeit at the cost of speed. The final trip, either from the container terminal at the port, an inland barge container terminal or from a train terminal, is by truck to deliver the container to the customer. (Some customers may be located directly adjacent to a multimodal terminal or even have their own terminal but this is infrequent.)

Supply chains for bulk chemicals start with the mining of the raw materials. In the case of oil-based products, this involves pumping the oil from deep below the earth's surface, moving the oil to a port using pipelines, loading the crude oil into tankers, and transporting it to refineries. These refineries are usually located close to the demand regions for the finished and semifinished products. Refineries need access to large amounts of water (for cooling and generating steam); locations at a river's edge, on a sea shore, and in or near ports are common. The output of an oil refinery can consist of many products, such as kerosene and gasoline that power internal-combustion engines, and propylene that is used as a basic material in the petrochemical industry. These products are temporarily stored in tank farms. Onward transportation can be by pipeline, by vessel (sea-going, short-sea, or barge), by rail car, or by truck. The customers for semifinished chemicals such as propylene are themselves also typically located in or near ports (for access to water and for transportation of raw materials and finished products). These chemical plants use bulk chemicals such as propylene to produce materials like polyurethane that are supplied to industrial companies (polyurethane is used in seals, adhesives, and foams).

The goods that are transported in containers can be characterized as semifinished and finished goods that are transported from the manufacturing location to regions with customer demand. In bulk chemicals, the products are mainly raw materials and semifinished goods that are transported from the point of origin to a manufacturing site. The output of the manufacturing process is usually transported as a bulk good.

Containerized supply chains involve different means of transport that have to be coordinated. Each move from one means of transport to another requires equipment to perform the move. This equipment is expensive, heavy, and tied to a physical location; the transportation process must be arranged to make optimum use of the scarce equipment. This involves the use of storage areas to decouple the various stages of the transportation process; the stacking area in a container terminal is a good example of such a storage area.

There are many threats to transport planning, ranging from the weather to strikes and accidents. Thus, each individual part of the transport chain will be faced with uncertainty regarding for example the time of arrival. This uncertain time of arrival can disrupt the carefully planned operations of the next leg. We can distinguish two main approaches to handle this uncertainty: keeping additional inventory as a buffer for uncertain (late) arrival times or improved coordination. In the latter case, the coordination involves the timely exchange of information regarding the expected time of arrival and an ability to update the planning of operations accordingly.

If we compare maritime transport of containers with pure land transport, then we see that the maritime transport is much more structured around the fixed elements of the infrastructure (ports, container terminals) whereas land transport by truck is much more flexible. An advantage of liner shipping over truck transport is that the schedule of a liner is well known beforehand; this facilitates planning to make the best use of the available capacity, whereas truck transports are often arranged on short notice, making it more difficult to find a non-empty return trip. The maritime transport of containers is arranged by a relatively small number of large carriers such as AP-Moller Maersk that are either involved in container terminal operations or deal with international terminal operating companies such as DP World and PSA International (Mangan et al., 2008). Increasingly, these carriers have their own terminals. This arrangement promotes the sharing of best practices and improves terminal efficiency. By contrast, land transport by train in Europe is highly fragmented, mostly by national borders. When compared to the maritime transport sector, the train sector in Europe also suffers from a lack of standardization in terms of equipment (cars, engines) and communication and safety systems.

Air transport does use a form of standardized container (the air container or 'unit load device', see DeLorme et al. (1992)) but not for all cargo and with much less use of automated equipment to perform the loading and unloading process. Also, the air container is used only for one leg of the transport; as soon as the air transport part of the trip is over, the goods have to be unloaded from the container. This eliminates some of the main benefits of containerization; the ability to leave the goods as-is in the container for the entire trip and the ability to use standardized equipment for the handling. This results in reduced handling of the goods, and the costs and risks associated with that (such as damage, spillage, and pilferage).

1.2 Actors in Maritime Transport

The actors involved in maritime transport can be roughly divided into two categories: sea-side and land-side. On the sea-side, shipping lines or *carriers* operate a fleet of ships to transport goods between ports. Carriers in container transport mostly use fixed sailing schedules that define which ships will visit what ports and at what time. These schedules are determined at a tactical level well before the actual sailing (months to years in advance). Carriers are focused on acquiring the best ships for their operations and on designing and operating the most efficient schedules.

The process of moving cargo from the point of origin to the destination is arranged by a *shipping company* or shipper. Shippers do not necessarily own the ships; they typically purchase transportation services from carriers and forwarders (v.i.).

On the land-side, *port authorities* are tasked with the design and operation of a port; they deal with the development and use of the (typically scarce) land, and accommodate the processes of the parties involved (for example by establishing common facilities such as port community systems (v.i.)). Port authorities aim to make their port an attractive choice for the other parties involved, both at the strategic level (the physical infrastructure of the port), and at a tactical and operational level (for example by streamlining administrative procedures).

Terminal operators operate one or more (container) terminals. Their focus is on efficient operations. They have a mutual dependency with the carriers: they have to coordinate their activities to get the best overall result. (If a ship arrives late, it will disrupt the operational planning of the terminal. On the other hand, if the terminal operations are not efficient, the ship may not make its planned sailing time.) The terminal operators have to deal with many different parties, ranging from large carriers to customs, port authorities, port security, shippers, trucking companies, and barge operators.

The terms *consignee* and *merchant* are used to designate the party for which the cargo is destined. For container transport, the consignees are not usually located in the vicinity of the container terminal; onward transport is required to move the containers from the terminal to the consignee. In bulk chemical transport, the consignees are often located in or close to ports.

Forwarders take care of the transport to and from a port. As storage at the terminals is limited to short term storage for a couple of days, the forwarder can also act as a consignee for storage in a warehouse. This can happen both at import (as a buffer before onward transport) and at export (as a buffer to facilitate consolidation). Forwarders may be independent companies or divisions of carriers that specialize in the end-to-end transport. These companies arrange the end-to-end transport using the facilities offered by carriers, terminal operators, and forwarders.

The *government* is involved at various levels. At the strategic level, governments shape the overall business climate of a country or region. As ports use significant amounts of scarce land, establishing a new port or extending an existing port requires government involvement. At the tactical level, establishing a favorable fiscal or legal environment (for example using free trade zones) can improve the competitive position of a port. Finally, at the operational level, governments are involved with taxation and security.

Cargo has to be transported from the point of origin to a port and, after the maritime transport, from to the final destination. This transport is arranged by trucking, barge, and railway *transport companies*. This onward transport can be characterized as *carrier* *haulauge* or as *merchant haulage*. The former is arranged by the carrier and the latter by the merchant or consignee. This is similar to the distinction we see in distribution between factory and retailer based haulage (the latter is also known as factory-gate pricing). The party that arranges transportation will focus on aligning transportation with its own business processes and will be less focused on the needs of the receiver. In the case of carrier haulage, the influence of the receiver will typically be limited to indicating the time windows for delivery.

Overall, a large number of parties is involved with maritime transport and each party want to optimize its part of the transport chain. Many local decisions that are in themselves optimal for each of the parties do not however guarantee a transport process that is globally optimal. The party that controls (part of) the chain focuses on scheduling its activities. Arranging and coordinating the activities of subsequence links in the transport chain is challenging and requires a higher-level perspective.

A container terminal operator for example depends on transport companies to pick up import containers that have arrived by ship. If the terminal operator can coordinate this onward transportation process with those companies, leading to more precise prior information on the departure time of the container from the terminal, then the terminal operator could take this information into consideration when stacking the container, storing it in a location that would allow for fast retrieval when the truck arrives for the pick-up. This coordination clearly requires additional effort of the parties involved. It would therefore be beneficial to get a clear insight into the value of this effort, for example, in terms of improvements to terminal operations and reduced wait time for the trucks. In chapter 2 we will study (amongst other topics) the value of knowing some departure information for the stacking operations of a container terminal operator, and in chapter 4 we will propose a distribution concept that takes a supply chain perspective to take advantage of the storage options at intermodal container terminals.

Similarly, if we consider bulk chemical transport, the pressure to reduce purchasing costs tends to force customers to focus on price. Especially for commodity products this may yield significant savings. The downside however is that the customers will be forced to accept less influence on the logistics of the fulfillment process as the seller is forced to streamline his processes to meet the price demands. For bulk chemicals this means that the delivery can be part of a milk run, whereby a single vessel visits a number of ports and customers. If the delivery is controlled by the seller or its carrier, then the customer may be forced to accept a delivery time that does not align with its processes. As a consequence, the customer may be forced to hold additional inventory, either on-site or at a third-party tank farm. Again, coordination could be beneficial (reducing the cost of additional inventory itself and of the facilities required to hold the inventory) but the value of this coordination is not easily quantified and with increased coupling of the supply process to the customer's operations comes an increased risk of disruptions of those operations. A quantitative approach to asses the potential value of the coordination and the risk to operations would be helpful; in chapter 6 we will formulate such as approach.

1.3 Containerized Transport

When in April 1956, a refitted oil tanker sailed from Newark to Houston carrying fifty-eight 35-foot containers, few observers could have predicted that this would be the start of a revolution, not just in transportation but in global trade. The converted oil tanker, the *SS Ideal-X*, was the brain child of entrepreneur Malcolm McLean, who has since been named the 'father of containerization', and who went on to found Sea-Land Service, Inc. (For a detailed history of containerization we refer to Levinson (2006).)

The introduction of the container and the process of standardization to a limited number of differently sized containers has led to a shift from manual labor to mechanization for loading and unloading ships. The maritime transport industry, at least for the part that is involved with container transport, has transformed from a labor intensive to a capital and knowledge intensive industry. While containerization has reduced the number of workers required, it has also increased the speed of (un)loading dramatically and reduced shrinkage and damage to cargo. The overall effect is that the cost of maritime transport has decreased significantly. Thus, containerization is one of the enablers of globalization.

The main issues in containerized transport are the planning of the physical infrastructure (location and design of ports and container terminals), the design and operation of the equipment (ships and for example quay cranes), the design of the sailing schedules (which ports to visit and what order), the operations of the container terminals (terminal layout, effective use of equipment), security (the use of a closed box for transportation also enables all sorts of illegal uses), and the coordination of all parties involved. In this thesis, we focus on the operations of container terminals and the coordination among the parties.

There are five main components in maritime container transport: the containers themselves, the ships that carry them, the ports at which these ships call, the terminals that handle the loading and unloading of containers onto the ships, and the processes that move containers to and from these terminals. The container provides a standardized load unit in which goods can be transported. Once loaded, only the container is handled and not the goods inside. The world-wide standardization of the container is both advantageous (reduced cost and increased compatibility) and disadvantageous (the installed bases of containers and handling equipment limit the potential for innovation). An interesting innovation is the design of the foldable container (Konings, 2005), aimed at reducing the space required for empty containers. (Empty containers are not used productively and are a direct consequence of current trade imbalances.)

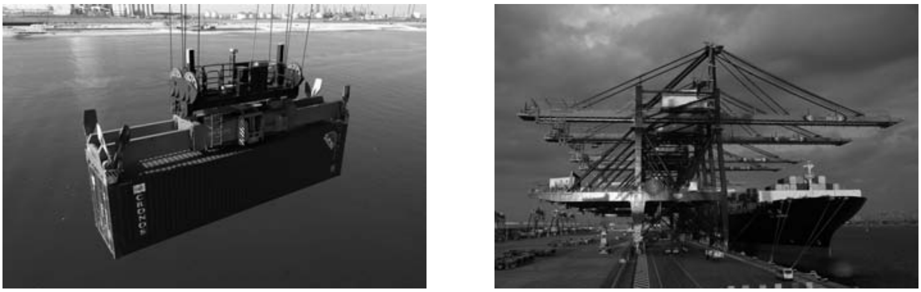

(a) A Sea container being hoisted

(b) Ship at container terminal

Figure 1.1: Images of container and Ship at berth (source: ECT website, used with permission)

1.3.1 Ships

Since the introduction of the first converted ships that could carry containers, the ships have become more specialized and have steadily increased in size. (See figure 1.1 for pictures of a container and a large container ship at berth.) As the ships get bigger, the number of containers per ship increases and the number of containers to be unloaded and loaded on a single visit to a container terminal increases. The time a ship spends in port waiting for the (un)loading operations can be considered nonproductive time and the ship's operator (the carrier) would like this time to be as short as possible. The pressure is therefore on the container terminal operator to provide the shortest possible turn-around time.

The economies of scale are clearly evident in maritime container transport but do depend on a high degree of utilization. If the utilization drops, then these large ships become less economical to run than smaller ships with a high degree of utilization. The first container ships featured on-board cranes to load and unload the containers; the port infrastructure was not yet in place to handle this type of cargo. Once the use of containers took off, dedicated facilities known as container terminals were created. In these terminals, large quay cranes can lift the containers from the ship to the quay and vice versa. Moving the cranes from the ship to the shore meant that more space became available for containers on the ship and it also eliminated restrictions on the size of the ship imposed by the reach of the on-board cranes.

1.3.2 Container terminal

A container terminal provides a coupling and transfer point between maritime transport and land transport. This is a loose coupling as a buffer storage area, the stack, is used to temporarily store containers before moving on to their destination. The quality of a container terminal can be measured both in terms of its operational efficiency (in particular fast turn-around times for sea-going vessels) and in terms of its connections. These connections can be by rail (for transport by train), by road (transport by truck), and by water (for transport by barge (inland waterways) and short-sea vessels (feeders)). See figure 1.2 for an aerial view of the ECT Delta terminal at the Port of Rotterdam. A lot of research has been devoted to the efficient operation of container terminals; for overviews, we refer to Günther and Kim (2006), Steenken et al. (2004) and Stahlbock and Voß (2008).

In many ports, getting the containers to and from the terminals becomes increasingly difficult. As many moves are done by truck, road congestion and operational bottlenecks may disrupt terminal operations. Within the terminal, the main issues are space (especially during times of extraordinary demand growth), the stacking strategy (which we will discuss in chapter 2), berth and quay crane allocation, and coordination of operations. Operations at a container terminal are tested most when a large vessel calls. The size of the vessel means that a large number of containers have to be offloaded and a comparably large number of other containers have to be loaded. As these vessels are very expensive, both to construct and to operate, there is considerable pressure to perform these operations as quickly as possible. This pressure also explains why carriers like APM Maersk have gotten involved in the container terminal business: they want to secure access to terminal capacity as well as fast and efficient handling of their ships.

The growth of containerized transport has stressed the entire transport system and in particular the container terminals as they provide the link between the maritime and hinterland transport. Ports used to be considered important in terms of employment, both direct port workers, customs, shipping agencies, tug operators) and indirect (banks, knowledge industries, consultancy, IT services). Ports have also become attractive locations for warehousing activities and value-added logistics (labeling, assembly, repackaging), creating yet more employment. Finally, there is the employment of the transport companies that provide access to the hinterland (by truck, train, or barge). In political terms, the land used by container terminals and the environmental impact of the logistics associated with global trade have caused a shift from container terminals as focus points for logistical and value-added processes to an emphasis on the disadvantages. The congestion and emissions caused by many container movements by truck are shaping future policies. The port extension plan 'Maasvlakte 2' in the port of Rotterdam for example mandates a significant shift from truck to rail and barge transport, as part of a plan to mitigate these disadvantages. In the USA, the state of California has passed a bill to regulate the queueing outside terminal gates to reduce emissions and highway congestion. While the bill did not have the desired effect (it regulates queueing outside the terminal; there are no implications for queueing inside the terminal), it has led terminal operators to move towards more controlled arrivals of trucks with the implementation of truck appointment systems (Giuliano and O'Brien, 2007).

Containerization has reduced the cost of transportation significantly (from 10% to 3%). As transport costs can be a significant portion of the total product cost, this has yielded cheaper products. The global reach of the container transport network means that there is a greater variety of products and that the availability of products has improved. The main disadvantages of modern containerized ports are pollution (from the vessels themselves but in particular from the large volume of truck trips for hinterland transport), congestion, and space occupation (ever larger areas are needed to store the containers).

Decisions regarding the terminal layout and the equipment to be used are long-term decisions at the strategic level as they involve construction and significant capital investment, especially for automated terminals; decisions regarding the stacking strategy are taken at the tactical and operational level and are mainly implemented in software that can be changed much easier. In this thesis, we will limit ourselves to the stacking strategy when we investigate container terminal operations in chapter 2. To reduce space occupation for container terminals, we could try to stack the containers higher. This however has a major drawback; as we stack higher, the probability that the container we need is not on top increases. If containers have been stacked on top of the container that is needed, then we first have to remove all these containers before we can access the container we need. These unproductive moves, called remarshalls or reshuffles, are a threat to efficient operations as the cranes that have to perform these moves are usually performance bottlenecks. If we want to avoid these unproductive moves, we should use the stack in such a way that containers can be stacked and unstacked quickly. This raises many questions: what are the trade-offs between stacking and unstacking efficiency, which components are crucial to overall performance, and what is the value of detailed information on individual containers. This calls for a smart approach to the stacking operations; we will discuss several of these approaches in chapter 2.

The environmental impact of container transport could be reduced by a shift away from road transport by truck to transport by rail (train) or water (barge). As the final leg of a trip will still have to be done by truck (few customers are located directly adjacent to a rail terminal or inland port), we would thus be dealing with multiple modes of transportation. This type of intermodal transport provides both challenges (in terms of coordination) and opportunities. We will return to this topic in chapters 4 and 5.

Figure 1.2: Aerial View of ECT's Delta Terminal (source: ECT website, used with permission)

1.3.3 Ports

The trend to ever larger container vessels puts ports that have a deep depth and easy access to its quaysides at an advantage. The port of Rotterdam, for example, is one of the few ports in Europe where the largest ships can still enter the port fully loaded at all times (i.e., not restricted by the tide). The terminals must keep up with the increases in vessel size, both in terms of quayside depth and in terms of the quay crane size. Thus, for carriers the port selection problem centers on access to quays and terminals. Shippers that use the services offered by the carriers are faced with a different port selection problem.

In the port selection literature, it is argued that the traditional criteria such as port infrastructure appear to have relatively little influence on the port selection process for shippers (Slack, 1985). In this literature, a lot of effort has been put into discovering the factors that influence the selection of a port by carriers, shippers, and customers. This research is usually qualitative and explorative in nature; quantitative and formal research in this area is scant. In this thesis (chapter 5) we will approach port selection from a quantitative angle; if we can provide more quantitative measures for the trade-offs between ports, this could be the basis for a more structured port selection decision process. Instead of a list of factors to be considered, such models could evaluate each port and provide quantitative indicators for the decision maker.

1.3.4 Knowledge infrastructure

In addition to the physical and services infrastructure that are directly related to containerized transport, a knowledge infrastructure has developed over the last forty years. This ranges from companies providing general-purpose and custom information technology (such as terminal operating systems (v.i.)), to consultancy companies such as TBA, and academic institutions. In terms of academic research, the field shows a significant development over the last five years, creating a substantial body of knowledge related to the use of containers (see Günther and Kim (2006) and Stahlbock and Voß (2008)). From these papers, we can also observe that the EU has a leading role in academic research when compared to the US and China.

Specific projects that are relevant to the knowledge development in the Netherlands are the FAMAS (1997–2002) and INCOMAAS projects that were initiated by Connekt, a Dutch knowledge center for Transport Technology with participation from government, private sector, and knowledge institutions. FAMAS (an acronym for "First All Modes All Sizes"), was aimed at developing the know-how for a new generation of container terminals that would be capable of handling all transport modes and be scalable to very large streams of containers. The INCOMAAS (Infrastructure Containers Maasvlakte, concluded in 1995) project preceded FAMAS: the aim of INCOMAAS was to define a masterplan for the infrastructure and container handling at the Maasvlakte up to 2020.

The development of container-related research is fueled by the realization of the special characteristics of containerized transport. If we compare it to warehousing, we see that it operates on a much larger scale, has to deal with additional issues such as stacking that are not typical in warehousing, that the larger number of actors involved creates additional challenges such as the coordination among transport modes, and that a very high level of investment (and a correspondingly long planning horizon) is required for automated equipment. The academic literature on warehousing tends to be focused on internal operations within the warehouse. Research on container terminals also has to deal with the interface to external operations, such as onward transportation, which complicates matters considerably. The need to include these additional elements enforces either simplifying assumptions (to use available models) or to use of research methods such as simulation that can model these elements.

1.4 Bulk Chemical Transport

In containerized maritime transport, a single ship can transport many different types of cargo, provided the cargo is stored in standardized sea containers. This standardization of the transport unit allows carriers to use the same transport device (the container ship) for many different customers. In bulk chemicals, the cargo is loaded in tanks; within a single tank, the cargo has to be homogeneous. While a ship for bulk chemicals has multiple tanks, this does not offer the same flexibility as a container ship. The maritime transport of bulk chemicals is therefore more focused on the customers and less on the infrastructure. As the bulk chemicals are the raw materials that are processed by the customers, they are critical to the operation of the customer's plant. To provide a buffer against supply disruptions, the customer will have some storage tanks on site but these have a limited size. Tight coordination of the transport process and the plants' production process is therefore essential. The cost of a plant shutdown due to lack of raw materials, or lack of storage space for the finished products, is very high as it can take a long time to restart the production process once it has been stopped.

The main components of maritime bulk chemical transport are similar to the main components of maritime container transport. The main difference is that there are far fewer different units shipped. The number of different chemicals that can be shipped with a single tanker is determined by the number of tanks onboard and is thus far less than the number of different containers that can be shipped by a large container vessel (in excess of 10,000 twenty-foot containers for the largest ships at this time). The terminals for bulk chemical transport are equipped with tank farms to allow for temporary storage.

As we have seen in our discussion of containerized transport, the structure of the infrastructure is designed to avoid a tight coupling of the maritime and the hinterland transport. In bulk chemical maritime transport, the emphasis is on linking the supply of raw materials to the plants' production process. This requires more coordination and information exchange between the customer, the shipper(s), and the supplier. Although it is possible to use decoupling points in this supply chain in the form of tank farms, this does incur additional cost and handling. Reducing the reliance on these decoupling points could therefore reduce costs and speed up the transport.

As we reduce the use of these decoupling points, we increase the level of coordination required. The alignment of the supply of raw materials with the plant production process becomes more challenging. One of the key questions then becomes how this alignment can be modeled: in chapter 7, we will use a case to discuss several ways to model this alignment.

The transport of bulk chemicals may consist of multiple stages; for example, from the supplier to a transshipment point in a port by short-sea vessel, followed by barge transport to the plant site. In such a case, it would be attractive to couple these two transport processes and eliminate the use of the decoupling point. However, if we want to directly transfer the cargo from a short-sea vessel to a barge, we are faced with the task of coordinating the movements of these two vessels. In chapter 6, we will study such a setup, create a schedule for this coordination, and evaluate it using both simulation and analytical calculations.

1.5 Methodology

A review of the literature has shown us that a lot of different research methods are employed in maritime, port, and transportation research. A common approach found in transport economics, for example regarding port selection, is exploratory research; this is typically a high-level approach that focuses on determining the main factors that influence decision makers. The research methods employed are questionnaires and interviews. This approach can yield insight into the decision making problem but one of the drawbacks is that the validity of the outcomes is unknown.

A second common method used is descriptive or case study research: a phenomenon is described in detail within its context. While such a rich description can form the basis for more formal research methods, the results are by definition limited; without additional methods, it is difficult to generalize the findings.

The results of exploratory research can provide the starting point for more formal research methods. For example, an investigation into port selection factors that involves interviewing a small number of shippers may yield some hypotheses regarding the role of the hinterland transport network. In a more formal study, these hypotheses could be tested with surveys (of a larger number of shippers) or experiments.

A large body of research output is dedicated to building and evaluating quantitative models. We can distinguish between normative and evaluative models. Normative models yield a norm for (the performance of) the system; evaluative models provide quantitative statements regarding system performance. Normative models typically involve optimization, i.e., to seek the minimum or maximum of a function by systematically choosing the values of variables from an allowed set (Wikipedia, 2009). The field of operations research has generated many different optimization models (such as linear programming) and algorithms to find optimal solutions (such as branch-and-bound) that can be used to create normative models for system performance. As a method, optimization can be used when the problem is well defined and clearly demarcated. This demarcation implies models that are limited in scope (for example to a single company or agent within a supply chain), having to meet the preconditions for the method: the result is a set of simplifying assumptions that have to be met. These assumptions and the usually limited scope of the optimization models mean that creating normative models will be of limited value for many scenarios. In situations involving multiple agents, evaluative models such as simulation models offer a means to determine the quantitative performance of a system, albeit without the benefit of having a norm for that performance. Optimization can be directed towards design (before the construction of the infrastructure) and towards operations (once the infrastructure has been built).

The main focus in this thesis is on the quantitative performance of systems used in port, container, and bulk chemical logistics. The systems we study involve multiple agents. Simulation is a commonly used method for the evaluation of system performance; it will be the core method employed in this thesis. It is a method that involves modeling the system in an executable form using a general purpose programming language (such as C++ or Java, often supplemented by a simulation library or framework), or a dedicated simulation language (such as Arena or Enterprise Dynamics). These models can be highly detailed (to evaluate operational performance) or less detailed (to evaluate global design choices). A benefit of using simulation is that it enforces fewer restrictions in the modeling phase than other methods, such as mathematics. Especially when operational decisions have a major impact on the performance of the system, simulation can capture this impact. The downside is that the models are often very involved, that they can take a lot of effort and time to create, and that they are often one-off projects. The latter means that whereas other methods can benefit from validation based on the class of problem, simulation models have to be validated individually, which is both challenging and time consuming.

Once a simulation model has been constructed, it can be used to evaluate the performance of the system. This typically involves running the simulation model for a number of different scenarios and comparing the results. Optimization can be either embedded within the simulation model (for example, when deciding where to stack a container in a yard) or be implemented outside the model. In the latter case, the optimization will use the simulation model to evaluate the performance of the system for a particular combination of values for the parameters (the simulation model performs the role of the 'function' to be minimized or maximized); the optimization process then consists of a systematic approach to choosing the particular values to evaluate and comparing the results. Both approaches will be used in this thesis.

1.6 Information Technology

Information is at the core of decision making processes. Whether it is information with the details of a container that will arrive shortly at a terminal, or inventory data of distribution centers, accurate and timely information is essential. Information technology provides the tools to exchange and process this information, often in real-time, both within and between organizations.

Just as maritime transport has been reshaped by the introduction of the container, the administrative and planning processes that are required to facilitate transport have been transformed by the introduction of information technology (IT). The use of IT has reduced manual administrative work, eliminated mistakes when transferring data, and speeded up the processing of information. The speed and low cost of information transfer has enabled new approaches. For example, the availability of detailed information on arriving containers means that it is possible to include that information when planning the unloading and stacking processes at a container terminal. Shippers that arrange end-to-end transport can use this information to plan onward transportation.

Whereas the initial use of IT was isolated (stand-alone systems or systems that were linked within a company only), the availability of low-cost, high-quality communication networks has enabled easier and faster exchange of data between companies. Similar to the standardization process of the physical containers, the IT industry has evolved to adopt a number of standards that allow the exchange of information across system and company borders. Edifact and XML are examples of these types of standards. With these standard formats for information exchange, it becomes easier to link information systems.

Information technology works as an enabler for improved cooperation and information exchange by parties in a supply chain. The ability to provide more detailed and more recent data than is possible using other means of communication can be used as a foundation for improved decision making. Using IT to share information across a supply chain can help prevent suboptimal local decisions based on partial or outdated information. The well-known bullwhip effect (Lee et al., 1997) provides a strong illustration of what can happen in a supply chain if information is not shared. Variations in demand may then get amplified along the supply chain. IT facilitates information sharing and is therefore a key component in reducing uncertainty in the supply chain. It can thus help improve the coordination among the parties involved, provided that they are willing to share such information. While IT is a critical component, it is not sufficient: the information that is exchanged using IT has to be integrated into the business processes of the partners, in order to improve decision making.

Within a port, there are many parties, each with their own information systems, ranging from simple spreadsheets to more complicated enterprise resource planning systems. To connect all these agents to a common infrastructure and promote information sharing has proven a challenge for many ports worldwide. The most successful answer to this challenge has been a class of information systems that have been labeled "Port Community Systems" or PCS's (Rodon and Ramis-Pujol, 2006). The development of these systems shows a large diversity (private sector versus government control, voluntary versus mandatory participation, ownership of the system and the data, etc.). The core function of all port community systems is to share and exchange data that facilitate port processes.

Various vertical applications have been created to support planning and operations of various types of parties. These range from Warehouse and Supply Chain Management modules that are part of larger Enterprise Resource Planning (ERP) systems from the larger software vendors such as SAP, Oracle, and Microsoft to specialized applications that are supplied by smaller vendors such as Ortec. These specialized applications can often communicate with existing ERP packages if these are used within an organization. While the generic ERP modules offer a range of standard policies that can be configured for an individual organization, the specialized applications are more tailored to specific needs and often offer advanced optimization facilities.

As the processes of container terminals are very similar, a dedicated type of vertical information system has emerged to support those processes: the so-called Terminal Operating System (Grifo, 2008). These information systems provide support for berth allocation (where should a vessel be berthed on the quay and which quay cranes are assigned), yard management (this includes the stacking strategy and optionally generating housekeeping moves), vessel planning (in which order should containers be (un)loaded)), gate control (registration of trucks arriving at the gate, truck appointment system), rail planning, equipment planning, reporting, invoicing, and external communications. As reported by Grifo (2008), almost half (43%) of the top 100 container ports of 2007 use a custom, in-house TOS. Of the packaged TOS systems, Navis has the largest market share (30%), followed by Cosmos (11%), TSB (8%), Tideworks (5%), and Cyber Logitech (3%). The large marketshare of in-house systems suggests a mismatch between the capabilities offered by the packaged systems and the requirements of the terminal operators. An alternative explanation could be the integration of existing in-house systems with other information systems that makes it difficult to replace the current TOS with a commercial off-the-shelf system.

A common problem with all packaged software is that the algorithms that are used are not always known and can not easily be replaced or augmented. While it is understandable that vendors of these packages may not want to disclose all the details of the algorithms used in their packages, this does make it difficult for potential customers to evaluate whether a package suits their needs. It also makes it very difficult to compare alternatives. Given the limited number of vendors and the infrequent nature of this type of decision problem, it may however be acceptable to terminal operators to spend a significant time and effort to evaluate the alternatives. A terminal operator who wants to compare packages in terms of their yard management capabilities will have to consult with the vendors and potentially develop and run a number of scenarios with these packages to evaluate their performance. For academic research, the limited information that is available on the strategies used for yard management make it difficult to compare research results with practice.

For a terminal operator, the terminal operating system is both a registration system (it keeps for example track of the current location of all containers) and a decision support system (it generates operational stacking decisions according to the criteria specified). In their roles as decision support systems, the TOS vendors systems offer a number of their own components (such as Expert Decking for yard management in the Navis package) and they typically provide an interface that allows external software to integrate with the system. Such an interface can facilitate the use of more specialized software that extends or augments the TOS functionality. There is no universal "connector"; the link between two packages will have to be created for each combination of packages. There are standards such as web services that standardize the mechanics of the connection (the network protocols to use and the way the messages are encoded) but the semantics and flow of the messages depend on the packages that have to be linked. The results of the work in this thesis could be used in the core of the enterprise software (such as ERP or TOS packages) or as add-on modules that augment the existing functionality.

1.7 Structure of this Thesis

This thesis consists of a number of essays on port, container, and bulk chemical logistics. These essays concern projects that presented themselves as research opportunities to the author in the 2002–2008 time frame. Thus, we do not claim to address all the issues mentioned in this introduction. Overall, the focus is on operational and real-time decisions. We will first give a high-level overview of the essays; this is followed by a more detailed description in which we also provide some background information on the projects and collaborations involved.

1.7.1 Overview

The first issue we will address is the value of detailed information on individual containers for a container terminal operator. It is clear that this operator can improve his operational decision making when more information is available: without such information, the operator can only make an uninformed or random choice. This may however mean that the decision is a poor one that will result in problems later on: if a container that has be retrieved from the stack is stored below other containers, then those other containers have to be removed first, increasing the overall time required to perform the move. Given that automated) equipment is both expensive and scarce, this is likely to be inefficient. Thus, the main issue we research in this chapter is how additional information that may be available to the terminal operator can be used to make better, more informed decisions during the stacking process.

For our second issue, we turn our attention towards the containerized transport of relatively cheap mass-market goods known as fast-moving consumer goods. Clothing and sports shoes are examples of these types of goods. These goods have to be transported to their markets quickly and at low cost. However, market demand may change: the variation among various demand regions can be high. We investigate if we can use the advantages of containerized transport (such as standard load units and some free storage time at terminals) to mitigate the effects of these shifts in demand by rerouting cargo to different destinations.

The standardization of containerized transport has lead to more choice for shippers of intercontinental cargo. A shipper can choose from many ports within Western Europe that have container terminals if he has to transport containers from Asia to Western Europe. The decision process for the selection process of a port has received some attention in the academic literature (starting with the work of Slack (1985)), but most of the existing approaches are qualitative in nature. If we can design a quantitative framework to evaluate the port options that are available, we can provide practitioners with more rigorous tools for decision making. This framework could for example be used to support a strategic (long term) decision such as the selection of a location for a European Distribution Center.

Price competition for commodity bulk chemicals can create a lack of coordination between supplier and customer. The result can be that the customer has to hold more inventory to accommodate non-optimal (from his perspective) delivery times. If we accept this effect of cost-orientation for raw materials, we can investigate if it is possible to align the transport and production processes of the customer with the fulfillment process of the supplier.

Finally, we discuss how the arrivals of ships that transport bulk chemicals should be modeled within a supply chain for a chemical plant. From studying cases described in academic and professional literature, it would appear to be obvious: arrivals of customers, trucks, and ships are modeled as Poisson processes. This however assumes that the arrivals are independent and that the interarrival times has a stationary distribution, or in other words, that the arrivals are uncontrolled. This is an odd choice for ships arriving in port as we can safely assume a fair amount of planning that has gone into determining a planned arrival time: the arrivals are the result of coordination between supplier and customer. We examine how this process of coordination can be translated into modeling constructs that control the arrival of ships in a simulation model for a chemical plant. We create some alternative constructs for different types of coordination and evaluate their relative performance.

1.7.2 Chapters

In chapter 2, we will look at a number of alternative ways of stacking containers in a yard. To evaluate the stacking strategies, a detailed simulation model is used. We will see that stacking strategies that use more detailed information, such as category stacking, have a better performance than less sophisticated strategies. Chapter 2 is based on work that was done by Rommert Dekker (Erasmus University Rotterdam (EUR)) and Patrick Voogd (at that time drs-student at Erasmus University Rotterdam) as part of the FAMAS project that aimed to design a better and more efficient container terminal, in anticipation of larger container vessel sizes (see Celen et al. (1999) and Dekker et al. (2000) for more information on the FAMAS project).

The research into stacking rules for container terminals is continued in chapter 3. Here, we have reduced the complexity of the algorithms involved and focused on a small number of basic rules. This should give us more insight into the basic trade-offs involved. For this project, we started from the specifications of the generator and simulator used in the chapter 2. As the tools for creating this type of program had improved, Bram Borgmans reimplemented both the generator and the simulator program using current (2008) technology, as part of his MSc thesis project (supervised by the author and Rommert Dekker). This reimplementation should also provide a sound basis for future experiments.

Chapter 4 looks at an intermodal alternative to truck transport in a supply chain for car tires. Using the slower and more expensive intermodal transport, we show that it is possible to reduce costs overall through smart use of the available intermodal terminals as temporary storage locations. Deploying some of the stock in advance on the basis of a demand prediction allows us to combine the lower overall cost of the intermodal transport with a lower average lead time for the customer; because some of the stock is deployed in advance of demand, we refer to this concept as “floating stock”. This work was done in cooperation with Geerten Ochtman, as part of his MSc-thesis project at EUR, Rommert Dekker, and Walter Kusters (at that time at Vos Logistics), who provided detailed case data.

If we take the floating stock approach and apply it to the Asia-Europe trade, we have to deal with much longer cycle times and we are faced with an additional decision problem: which port or ports in Europe should we use? This port selection problem is at the heart of chapter 5. In this chapter, we apply the floating stock principle to a supply chain of consumer electronics from Asia to Western Europe. We focus on the role of the ports within the road transport network and we quantify the flexibility of rerouting stock that can be offered by ports depending on their position within this road transport network. This essay is an extension of a project that was done in cooperation with Albert Veenstra (Rotterdam School of Management, Erasmus University Rotterdam), Joost Hengstmengel (internship and bachelor's thesis project (Hengstmengel, 2006)), and APL Logistics (detailed case data); the project was sponsored by the AC Transport consortium.

Next, we turn to bulk chemicals; in chapter 6 we discuss the supply of bulk chemicals to a chemical plant and evaluate the option of replacing truck transport by barge transport. Barge transport is an option because this plant is located on a waterway and the connections to the nearest seaport are available. The focus is on the coordination between the barge and seagoing vessels. This chapter is based on a consultancy project by the Erasmus School of Economics, involving Patrick Meersmans, Rommert Dekker, and Vopak.

In chapter 7, we study the design of a jetty for a chemical plant. The inspiration for this essay was provided by a consultancy project of the Erasmus School of Economics, involving Stef Kurstjens and Rommert Dekker, and Vopak. While the original research study focused on the jetty design, this chapter looks at the way in which the arrival processes of ships should be modeled. We will see that this arrival modeling has a significant impact on the outcomes and that the commonly used axiom of using a Poisson process to model arrivals in a system should be carefully considered for each model.

We summarize the results in chapter 8.

References

Branch (2007) provides an overview of the global shipping scene. The special 40th anniversary issue of the "Containerisation International" journal has several articles on the history and development of the container transport industry (Samwel (2007)). The role of ports in logistics is discussed in Mangan et al. (2008). The Transportation Institute has an online glossary of maritime transportation terms Transportation Institute (2009).

Chapter 2

Advanced Methods for Container Stacking*

In this chapter, we study stacking policies for containers at an automated container terminal. It is motivated by the increasing pressure on terminal performance put forward by the increase in the size of container ships. We consider several variants of category stacking, where containers can be exchanged during the loading process. The categories facilitate both stacking and online optimization of stowage. We also consider workload variations for the stacking cranes.

2.1 Introduction

World trade, especially the Asia-US and Asia-Europe trade, has developed rapidly over the last decades. As a result, container traffic has increased at a high rate as well. Ocean carriers have responded by ordering more and much larger ships. For example, the PONL Mondriaan can carry up to 8,450 TEU, whereas Maersk's largest ships are considered to carry up to 13,000 TEU. The consequence of having larger container ships is that terminal activities become more a bottleneck and its productivity has to go up. This was already acknowledged in the FAMAS research project started in The Netherlands in 1999 (Celen et al., 1999). In this chapter, we will report on explorative research concerning container-stacking policies at an automated terminal. Before a detailed discussion, we will first give an overview on container activities.

---

* A version of this chapter has been published as Dekker et al. (2006) and Dekker et al. (2007).

2.2 Container Operations and Trends

Several reviews on container handling have been published (Meersmans and Dekker (2001); Steenken et al. (2004); Vis and de Koster (2003)). The overview below is based on them as well as on own experience with terminal studies. Although marine container terminals vary all over the world, they have a number of similarities. Ocean-going ships moor at a berth where quay cranes unload and load containers from the ship. Containers that have been unloaded are then transported to the main stack where they are positioned by gantry cranes or straddle carriers. Containers can again be loaded in sea ships. Alternatively, they can be further transported on land by truck, train, or barge. In those cases, the container is moved from the stack to a rail or barge terminal or it is directly positioned on a truck, which has entered the terminal. Most terminals are manually operated; a few terminals use semi-automated equipment such as automatic guided vehicles (AGV), to transport containers, and automatic stacking cranes (ASC), to stack containers. These are ECT in Rotterdam, CTA in Hamburg, and Thamesport in London. In this chapter, we will focus on these automatic terminals such as the Delta Dedicated Terminals at ECT's Maasvlakte complex in Rotterdam.

2.2.1 Implications of larger ships

Large ships are more expensive to buy and to operate than small ships. As a ship's port time can be considered as nonproductive, a large ship's port time is more costly per hour than a small ship's time. Larger ships, however, take more time to unload and load due to the larger amount of cargo. This is a kind of paradox, which puts a limit to the size of ships, as pointed out in Cullinane and Khanna (2000). While the terminal operator can use more quay cranes for a larger ship to limit the impact of the number of containers to be moved, the handling time per container increases with the size of the ship as each quay crane has to reach further, both in the horizontal and in the vertical direction. The time in port consists of port entry and departure time, (un) mooring time, preparation time and the actual loading/unloading time. Larger ships are therefore likely to make fewer and larger calls than small ships to reduce unproductive time. For example, the *PONL Mondriaan* loaded and unloaded some 4,000 TEU in one port. This will put much more stress on the terminal logistics and stack.

2.2.2 Structure of stacking strategies

Several decision horizons can be identified in stacking, viz. strategic/design for the long-term, tactical for the medium-term, operational for the short-term and real-time for the direct operations. Strategic decisions concern the choice of equipment, the size of the terminal in general and the stack in particular. Automated stacks have less flexibility and apply more costly equipment than manually operated stacks; hence, the design is very important. Tactical stacking decisions concern capacity decisions on months to year. In manually operated stacks there are more tactical decision freedoms than in automated stacks, viz. layout of the stack, number of cranes employed. Decisions on a tactical level include the use of operation strategies, such as using a pre-stack or the application of stack reorganizations (also called remarshalling) at those moments where no ships need to be served. Operational decision making concerns the reservation of space for ships, the decision to store a container at a particular location, the allocation of equipment to jobs, etc. Finally, the real-time phase is mainly relevant for automated equipment, as it concerns speed control and collision avoidance of equipment. These are mostly technical decisions taken by control systems. In this chapter, we mainly address strategic and operational decisions. The way the latter are carried out is captured in a stacking strategy. The main objectives of a stacking strategy are

- efficient use of storage space,
- efficient and timely transportation from quay to stack and further destination and vice versa,
- avoidance of unproductive moves.

The second objective implies, e.g., that an export container should be stacked close to the ship with which it will sail and that its retrieval time should be short. A stack with a maximum height of one container would be optimal for the third objective. This would however lead to an inefficient ground use and long travel times, so it is rarely applied in practice (apart from some stacks on wheels in the US). Accordingly, one has to decide whether a container should be stacked on another one or not.

A main input for a stacking strategy is the information available on a container. This is usually its type (size, reefer, dangerous goods), modality and date/time of departure. Unfortunately, this information may change or not be completely known upon arrival.

There are several types of stacking strategies. In category stacking, one defines categories and stacks containers of the same category on top of each other. In the residence time strategy, one stacks a container on others if its departure time is earlier than that of all containers which will be below it.

Steenken et al. (2004) distinguish storage planning and scattered stacking. In storage planning, space in specific areas of the stack is reserved before the ship's arrival. In scattered stacking, yard areas are not assigned to a ship's arrival but to a berthing place. The stacking position is then determined in real-time and containers are stochastically distributed over the area. According to Steenken et al. (2004), scattered stacking results in higher yard utilization and a significant reduction in the number of reshuffles. The category stacking employed in this chapter is a form of scattered stacking.

Some containers (e.g., reefers) require special locations because they need to be supplied with electricity. The determination of the stack capacity is a major design problem of a terminal, as the physical space required for the stack is often restricted and expensive. Stacking high may be advocated, but the expected number of reshuffles increases sharply with the stacking height. We define a reshuffle as an unproductive move of a container, which is required to access another container that is stored beneath it (this implies that reshuffles occur only when removing containers from the stack).

Quite often stacks are separated into import and export parts. Import containers are those containers that arrive in large container ships from overseas and continue their destination through inland transport. These arrivals are somewhat predictable. The departure of import containers, however, is likely to be in an unpredictable order, so they cannot be stacked that high. Export containers that arrive via land transport may arrive somewhat randomly, but their departure is usually connected to a ship; hence, they can be stacked in a much better way.

2.2.3 Loading or stowage plan

Every ship which is loaded at a terminal has a stowage plan. According to Steenken et al. (2004), it is made in two steps. First the shipping line makes a rough plan based on categories, which is sent to the terminal. Later, somewhat before the arrival of the ship, a more detailed plan is made by the terminal planner who fills the categories in with detailed containers. The stowage plan specifies which container will be loaded at which location in the ship. As containers vary in size and weight, the load distribution is essential for the ship's stability. Heavy containers should be stored as low as possible. The stowage plan, however, also directly influences the ease of unloading the containers and, hence, containers of the same destination should be loaded on top of each other or on top of containers destined for ports further away. Apart from these restrictions, there are also containers with dangerous goods, which should be stored preferably below decks, reefers that have special positions, etc. Advanced software is used to perform offline optimization of the stowage plan also to avoid reshuffles as much as possible. Although the stowage plan fixes the load order per quay crane, it does not fix the exact order in which the containers leave the stack as the crane loading cycles are quite stochastic and a difference in progress between cranes may occur. Therefore, this software does not take the actual operations of the loading into account (Steenken et al., 2004). Online stowage planning does take the details of yard operations into account and will be employed in this chapter; it is not yet in use at container terminals.

The stacking problem can be considered to be more difficult than the stowage planning as there can be uncertainty about which container will be needed before another. For import containers, this uncertainty exists because trucks arrive more or less randomly to pick up a specific container.

2.3 Stacking Research

2.3.1 Literature overview

Little has been published in scientific literature on stacking problems. A main reason may be that the practical problems are quite complex and do not easily allow for analytical results which are relevant for practice. Steenken et al. (2004) give a high-level overview of stacking both in theory and in practice.

Stacking problems can be dealt with in two ways: simplified analytical calculations or detailed simulation studies. The first gives insight into the relationships between the various parameters on a more abstract level. The second can go in much more detail, with the negative side effect that it is time-consuming and only few people really understand its ins and outs. No comprehensive stacking theory exists today, and a good stack design not only depends on local space conditions but also on the information characteristics of the ingoing and outgoing flow of containers which may vary from place to place. Examples of both approaches are given below.

Sculli and Hui (1988) were among the first to develop yardsticks for the relation between stacking height, utilization (or storage space needed) and reshuffles by applying a comprehensive simulation study. Taleb-Ibrahimi et al. (1993) discussed this relation for export containers both at a long-term scale as well as operationally. They discussed dynamic strategies that store early-arriving containers in a rough pile until a certain date, after which all containers for a ship are put in a dedicated storage area (usually close to the berthing place of the ship). The procedures developed calculate the storage space needed as function of the stacking height. Castilho and Daganzo (1993) continue these studies with the stacking of import containers. They consider two strategies; one that keeps stacks of the same size versus one that segregates the containers on arrival time. A slightly more detailed discussion resulting in tables and yardsticks (looking at stacking blocks with bays of similarly sized containers served by gantry cranes), both analytically and by simulation, was given by Kim (1997). Kim and Kim (1998) extended these studies by also taking the number of stacking cranes into account. They developed a simple cost model for optimizing this number using analytical approximations for the various performance measures.

In case the stowage plan is available some time before the sailing, the containers in the export stack may be remarshalled. This results in an "ideal" stack and, thus, less handling work during the loading operation of the vessel. Kim and Bae (1998) describe a two-stage approach to minimize the number of containers to be moved and to do so in the shortest possible traveling distance. Segregating space allocation strategies of import containers was studied by Kim and Kim (1999). In segregation strategies, stacking newly arrived containers on top of containers that arrived earlier is not allowed. Spaces are thus allocated for each arriving vessel. They study cases with constant, cyclic, and varying arrivals of vessels.

An empirical statistical analysis of the actual performance at a Taiwanese container terminal was provided by Chen et al. (2000). The number of reshuffles (Chen et al. use the term "shift moves") was related to the storage density, the volume of containers loaded and the volume of containers discharged both for stacking crane blocks and straddle carrier blocks.

Decision rules using weight groups for locating export containers were derived and validated through dynamic programming by Kim et al. (2000). Weight is a useful criterion as heavy containers are usually stored deep in a ship.

Stacking policies for automated container terminals are investigated by Duinkerken et al. (2001), who use a detailed simulation model that not only models the stack, but also the quay transport in an automated container terminal. They also apply categories, but in a much more simplified way than we do in this chapter. All in all a comprehensive analysis of how stacking should be done at an operational level is still lacking; hence, this chapter deals with it.

2.3.2 Selection of research object

In this chapter we investigate a container terminal with an automated stack as it is envisaged that future developments will move into that direction. A picture and log-

Figure 2.1: Overview of ECT's DDE Terminal (source: ECT website, used with permission)

ical layout of such a stack are given in figure 2.1 and 2.2. (Notice a slight difference in figures 2.1 and 2.2 with regard to the reefer platforms; in this chapter we follow figure 2.2.)

Figure 2.2 gives the general layout of the stack. On each lane we assume one Automatic Stacking Crane (ASC). Transfer points are located on both the sea- and land side. The lanes are perpendicular to the seaside, where jumbo (very large) and deepsea ships as well as shortsea/feeders are loaded and unloaded. The containers are placed with their long side parallel to the direction of the lane. The transfer point on the landside is used for rail, truck and barge. All lanes have the same length, width, and height, expressed in terms of locations. These locations are slightly larger than one TEU to allow for some space between the containers, used to pick up and put down containers. We define a pile as zero or more containers stacked on top of each other. We will refer to a pile with zero containers (the ground location is empty) as an empty pile; a pile that is stacked to the maximum stacking height is called a full pile.

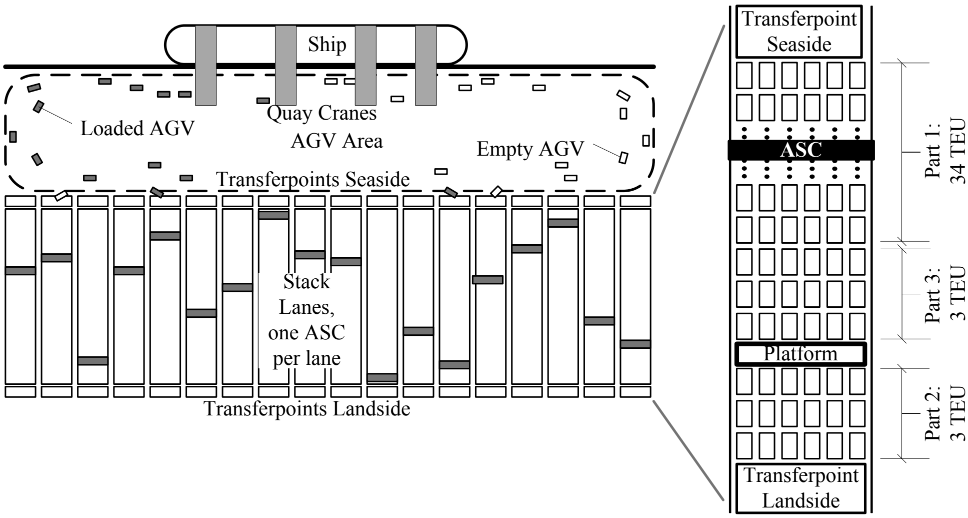

Figure 2.2: Schematic overview of the stack.

Every lane is partitioned in three parts. Part 1 starts at the seaside and is used for non-reefer containers. There are two parts adjacent to a special platform for reefer containers on the landside. Part 2 is closest to the landside, has a length of three locations, and is used for reefer containers only. Part 3 is located between part 1 and the platform; part 3 has a length of three locations, and is used for reefers and reshuffles of other containers. In part 3, the reefers are stored only directly adjacent to the platform between parts 2 and 3. The platform is 10 feet deep and supplies the reefers with electricity.

Automatic guided vehicles (AGV) transport the containers from the transfer points on the seaside to the quay cranes, and vice versa. In the stacking algorithm it is assumed that the number of AGVs is sufficient to handle all transport to and from these transfer points in time.

The base case is a stack with 27 lanes, where each lane has a length of 40 TEU, a width of 6 TEU and a height of 3 TEU. This implies that the theoretical stack capacity is 19,440 TEU. We do not make a distinction between an export and import stack. This is partly because separate space is already reserved for reefers and partly because it would cause inflexibility. The difference between export and import is implicitly incorporated in the analysis, because we will introduce different categories of containers which are stacked together and the import/export property is part of these categories.

Table 2.1: Modal Split Matrix.

To / *From*	Jumbo	Deepsea	Shortsea/ feeder	Truck	Rail	Barge	Total
Jumbo	0	2,332	3,630	407	965	1,405	8,739
Deepsea	2,691	1,344	1,466	270	389	568	6,728
Shortsea/feeder	3,870	2,000	0	967	1,876	2,735	11,448
Truck	438	368	967	0	0	0	1,773
Rail	1,047	540	1,877	0	0	0	3,464
Barge	1,524	788	2,736	0	0	0	5,048
Total	9,570	7,372	10,676	1,644	3,230	4,708	37,200

2.3.3 Simulation program

A large simulation program was developed in the MUST language Upward Systems (1994). This is a Turbo Pascal add-on which allows easy programming in Turbo Pascal, while using a number of modules from the package. It was also extensively used within ECT. It is fast, memory can be managed well and complex algorithms can easily be written and incorporated. Two separate programs have been developed: a generator program and an evaluator program.

The generator program creates entry and departure times of some 175,000 containers covering a period of 15 weeks of operation. The first three weeks are used as a warm-up period to fill up the stack. The output was written to a file which was used as input for the evaluator program where different stacking procedures could be tested. The generation of the containers was tied to the modalities with which they would arrive or depart. Several types of ships were considered, viz. deepsea ships and 8,000 TEU large jumbo ships, the latter arriving once a week with a call size of about 3,000 containers. From a high-level modal split matrix we developed cyclic ship schedules as well as detailed arrivals of all other modalities. The matrix in table 2.1 illustrates the flow between the different modes for a three week period.

We also developed detailed ship loading plans that specify the locations of individual containers and detailed crane sequences for the loading and unloading. The call size of the jumbo ship was set at some 3,000 containers. We assumed a 50%:40%:10% ratio between 20, 40 and 45 feet containers, which gives a TEU container ratio of about 1.5. This means that the jumbo ship loads some 4,500 TEU. We have also modeled other transport modalities: shortsea / feeder, rail, truck, and barge, with the daily fluctuations in truck arrivals and a stationary pattern with fluctuations for all the other modes.

An average container residence time of 3.7 days was used, in line with information available at ECT. This implies an average utilization of 50% of the base stack configuration. Detailed information about the generator program and its output is available in Voogd et al. (1999).

The evaluator program performs a deterministic simulation of an experimental setting, based on the stochastic output of the generator program. The output of the generator program contains exact departure times for all containers. The evaluator program uses these times to trigger events and adds a small perturbation for use in the stacking algorithms. These perturbations are used to model the information uncertainty that occurs in practice. AGV routing was not modeled in the simulation program. We took a constant time depending on the quay crane and ASC lane where the container came from or had to go to.

This experimental setup enables accurate evaluation of various stacking algorithms; the generator program provides the same scenario for each experiment. Any change in the results is due to the stacking algorithm selected for the experiment and to the minor perturbations. This way of experimental set-up however, does not facilitate different demand scenarios, as for each scenario a quite detailed arrival modeling needs to be constructed which is a difficult scientific problem on its own.

2.3.4 Stacking algorithms

A stacking algorithm describes the way in which containers are handled both in case of moves into and out of the stack as well as in case of reshuffles. For containers leaving the stack we have no options unless the containers are exchangeable with others. In that case (see also below) there might be other containers of the same category for the same ship (or other modality) which can be retrieved in a better way.

The main part of a stacking algorithm decides where to put a new container or a reshuffled container. In this chapter we investigate two main concepts, viz. random stacking and category stacking. In random stacking there is no preference for particular places and it is used to evenly spread containers over the stack. In category stacking, we define categories of containers on the basis of the loading plan. Containers in the same category may be exchanged freely. In category stacking one tries to exploit this property as much as possible. We supplement these concepts with decision rules for specific cases.

A stacking algorithm is also influenced by the information available at the moment of stacking. If the departure time of a container is known at stacking time, then we can stack the container on top of a pile of containers with a later departure time. This does however require a sufficiently large stack to allow the creation and maintenance of these "ordered" piles.

2.3.5 Common rules

There are some basic rules for all stacking algorithms in this chapter:

- 20 ft. containers occupy one TEU location in the stack, 40 ft. containers occupy two locations and 45 ft. containers occupy three locations.
- Containers of different sizes can not be stacked on top of each other.
- Containers have to be stacked precisely on top of each other (no overhang and a container can be on top of just one container).
- Containers can only be stacked in the direction along the lane, not transverse.
- Reefer containers are not placed on top of normal containers, or vice versa.

Reefer containers have a special requirement: the need for a power connection. This limits the locations available for stacking these containers. Thus, we have implemented the same stacking algorithm for reefers in the first five experiments. The only locations with power connections are directly adjacent to the platform. Thus, the number of locations available to reefers is limited to twice the lane width (once for each side of the platform). The stacking algorithm for reefers selects a random non-full pile within the special reefer section of the stack. If the pile is empty, the container is only stacked there when no more than three of these six reefer positions are occupied. This helps to make sure that all reefer reshuffles can be carried out. Otherwise this could cause a problem, because there are very few possibilities for the container to be reshuffled to. If the pile is not full, the reefer can be stacked if they are containers (reefers) of the same size. Whenever no suitable location is found in 5,000 random choices of a lane, the aim is changed to the reefer locations on the other side of the platform. This way of stacking probably causes low occupancy in parts two and three of the stack.

2.3.6 Random stacking

This algorithm is used as a benchmark. Suppose a 20 ft. container has to be stacked. The program uses random search to find a pile that is not full. If the pile is empty or if the containers in this pile are also 20 ft., an acceptable position has been found and the container can be stacked in this pile. If the pile consists of containers of a different size, then the container can not be stacked here. The program then determines a new random position by choosing at random a new lane, row, and position until a location is found where the container can be stacked. 40 ft. and 45 ft. containers are handled in the same way but in those cases the algorithm searches for either an existing non-empty pile of the same size or for a sufficient number of adjacent empty piles (two for 40 ft. containers; three for 45 ft. containers).

For reshuffles the program searches all piles in the lane except for the reefer positions on the landside of the platform. The container is reshuffled to one of the possible piles closest to the original pile.

2.3.7 Category stacking

This algorithm is based on defining categories of containers. These are defined through the export modality and in case of a ship, the place of a container in it. We assume that for certain categories (especially those defined for jumbo and deepsea, but not for trucks) containers are exchangeable in the loading plan or in the actual loading, if they are either in the same or different piles. The algorithm keeps track of a variable for every combination of lane, ship and category. This variable indicates how many piles of containers exist, within that lane, with only containers of that specific ship/category combination and an empty top position. The variable is used to facilitate the search for a good location (note that searching over 19,000 locations upon each of the 175,000 container entries is very time consuming in the simulation).

Now, suppose a new container has to be stacked. The first step is to determine if there is a pile that is not full and with only containers of that same category and for the same ship. All lanes are checked for such a pile; to spread the load evenly across the lanes, we start the search at a random lane. Using the variable described above, a zero indicates no such pile exists, whereas a positive value means one or more of those piles exist.

When the variable indicates that one or more of those piles exist, the program starts searching, randomly within that lane, for one of those piles. When found the container is stacked on top of that pile. If this creates a full pile, the variable associated with the current ship/category combination is decreased with one for that lane.

When no such pile can be found in the current lane, i.e. the variable has value zero for that ship/category combination in that lane, the aim shifts to the next lane. If value of the variable equals zero for that ship/category combination for all lanes, the container is stacked using random stacking (see description above).

2.3.8 Performance measures

Below we discuss appropriate performance measures of stacking policies.

Reshuffles and reshuffle occasions

There are two performance measures concerning the reshuffles. First of all, we define a reshuffle occasion as one or more reshuffle operations required to retrieve a container from the stack. We measure the reshuffle occasions as a percentage of containers that leave the stack. The total number of reshuffles is also counted (again as a percentage of the total number of containers leaving the stack). These measures are calculated separately for import as well as export containers. An export reshuffle is a container (export or import) that is reshuffled because the export container needed is under that container (so it is not necessarily an export container that is reshuffled). It seems obvious that a situation with many reshuffles or reshuffle occasions is undesirable, for reshuffling takes a valuable amount of time.

No positions available

We may not always find an empty location in the stack, especially considering the randomness in positioning a container when it enters the stack. This will most likely concern the 45 ft. containers, because they require three adjacent empty locations and they form a minority in comparison to the 20 ft. and 40 ft. containers. Therefore, there will be few piles with 45 ft. containers. Although the maximum utilization is always less than 100%, we may not find an empty location for a 40 or 45 ft. container. We assume that there is an emergency stack for these containers and leave them out of consideration, as they would otherwise cause a deadlock in the program. The aspect does imply that the real capacity is much lower than the physical capacity, which is also a known practical fact. We may also encounter this problem when reshuffling a container; if we can not find an empty location in the same lane, we move these containers to the emergency stack. A small number of reshuffles and reshuffle occasions indicate a better performance. Larger numbers indicate that the current stack size might be too small to be used with the current algorithm.

Workload of the Automatic Stacking Cranes (ASCs)

A third group of performance measures deals with the workloads of the ASCs. These workloads are determined every quarter of an hour as the proportion of time the ASCs are busy. The design of the simulation program allows ASC workloads to exceed the

Table 2.2: Typical ASC Travel Times.

Transfer point	First position	Twentieth position	Last position
Seaside	9.2 s	45.3 s	79.8 s
Landside	9.2 s	47.8 s	79.8 s

capacity, i.e. workloads of more than 900 seconds per quarter. Since the focus of this research is on the stacking algorithms and not on ASC scheduling, we have chosen to allow these overloads and consider the frequency and gravity of these occasions as one of the criteria for the performance of an algorithm. Details about ASC technical performance can be found in Voogd et al. (1999).

A move is handled at the same moment in time as specified in the container files, even when the ASC is not ready at that moment. Every move starts, when not already in the right position, with shifting the ASC from the previous position to the position for picking up the container (transfer point for containers that enter the stack) and ends at the position where the container is put down (transfer point for containers that leave the stack).

To give an indication of traveling times for ASCs, the times are calculated for going from one of the transfer points to the first container position, to the twentieth position, and to the last (fortieth) position (all positions relative to the transfer point; see table 2.2). The implementation code contains a precise model of the ASC movements, including maximum speed and acceleration along the three axes (longitudinal, lateral and vertical).

The difference between the traveling times to the twentieth position is incurred by the reefer platform.

The workloads for all ASCs are written to a file at the end of each quarter. The maximum and average workloads are determined, given as percentages of one quarter. An average workload of 50 percent means, therefore, that on average an ASC is busy half of the time, which is 450 seconds per quarter. Concerning the actual scheduling of an ASC, a workload of 80 percent is already pretty high. That's why the proportions of ASC quarters with the ASC working more than 80, 90, 100, 110 and 120 percent of the time are measured.

Occupation

The degree of occupation is measured for the ground locations. For this purpose, at the end of each quarter, the number of ground locations in use is recorded. The maximum and average numbers are calculated separately for the three parts of the stack. The overall occupation of the stack depends only on the size of the stack as the number of containers that will be handled during the simulation is constant for all experiments. The occupation is 51% for the first three experiments and 47% for the other experiments; this is low but a consequence of the large call sizes of the jumbo ships (sufficient space in the stack is required for the large number of containers to be unloaded before loading can begin).

For the ground locations, we expect a larger number of reshuffles when few ground locations are occupied. The average height of the non-empty piles is higher, which increases the possibility of reshuffles. If, on the other hand, almost all ground locations are covered, then we expect a negative influence on the number of reshuffles and new containers that can not be stacked in the regular stack.

2.4 Features of the Stacking Algorithms

In this chapter we explore the use of categories for the stacking of containers. For each experiment we will indicate for which categories containers are considered to be exchangeable. Here, we define exchangeable to mean that a different container from the same category may be substituted when a container is requested for loading. The categories defined for large containerships are typically exchangeable. All containers to be picked up by trucks also form a category, but these containers are not exchangeable. To facilitate the exchange operationally, we stack containers of the same category in the same pile as much as possible, but exchange is also possible for containers of different lanes.

The definition of the categories is based on the weight class, destination and type of container (the same criteria are mentioned in Steenken et al. (2004)). Thus, only the export modality is a feature in the definition of the categories: the import modality is not taken into consideration. Using the data from ECT, we defined some 45 different categories for jumbo ships and 90 categories for deepsea ships. Containers destined for shortsea/feeder, truck, rail and barge transport will be allocated to a single category for each mode, even though they can not be exchanged in operation. As we will see in the experiments it is not wise to stack them in the same pile.

In addition to categorization, we have implemented several other features for the stacking algorithm.

2.4.1 Preference for ground locations

We use a preference for ground locations to decrease the possibility of spoiling a uniform pile, i.e., a pile with containers that all belong to the same category. The implementation of this feature tries to avoid stacking a container of a different category onto an existing uniform pile. This causes a preference for stacking on empty piles and for stacking on multiform piles. It will reduce the number of empty piles and may cause problems for stacking or reshuffling (45 ft.) containers.

2.4.2 Workload control

The workload control feature associates a workload variable with each lane. We defined the workload variable as the percentage of time of the current quarter that the ASC for the lane was busy. When the workload variable exceeds a specified threshold, the lane is skipped in the search for a stacking position.

2.4.3 Alternative algorithm for reefers

Reefer containers can be stacked in just a small part of the stack. Therefore, our initial experiments exhibited some problems with reefer reshuffles. For every reefer reshuffle there are only up to five possible new positions (within the same lane). When stacking these containers at random, a lot of containers could not be reshuffled (within the same lane). The number of reefer reshuffles however was substantial. We therefore introduced category stacking for reefers with a modification to avoid the creation of full piles. In this way, we aim to leave a sufficient number of feasible empty positions for reefers.

2.4.4 Use empty pile closest to departure transfer point

When an empty pile has been selected for a container and multiple empty piles are available in the same lane, the algorithm will select the pile that is closest to the point where the container will leave the stack. The aim is to lower the ASC workloads during ship loading. The ASC will have to travel a shorter distance to get to the container, which decreases the time needed to unstack a container. Furthermore, it is expected that this feature will also increase the number of non-empty piles. This is due to the fact that we will now use the empty pile directly adjacent to an existing non-empty pile, leaving no space (TEU position) open. The proportion of 'unusable' empty piles will then be lower. Using more ground locations is also thought to decrease the number of reshuffles. We will explain this feature with the following example.

For example, consider the case where we have to stack four 45 ft. containers with a maximum stack height of three containers. Furthermore, suppose that all ground locations are occupied except for the last six TEU positions in front of a transfer point and that all piles of 45 ft. containers are full. In this case random stacking might put the first container on the second, third, and fourth TEU location instead of the first, second, and third TEU location. The second and third containers will then be stacked on top of the first container. Even with three TEU ground locations available, the fourth container can not be stacked in this lane: the locations that are available are not adjacent. With the new rule, the first container will be stacked upon the first three empty TEU locations, leaving the other locations open for one (or more) of the other three containers.

At first, we will use this feature for all modalities. A variation of this feature is designed to reduce the ASC travel time (and thus the workload) when unloading jumbo or deepsea ships: containers destined for the landside are not subjected to this rule. When there are no jumbo or deepsea ships present at the quays, the ASC workloads are lower and the additional travel time does not pose a problem. This feature will probably have a negative effect on the average distance to travel for export containers (because the import containers can use positions close to the sea-side). It will also decrease the effect described above concerning the use of ground locations.

2.4.5 Combine parts one and three of the stack

Initial experiments showed a low use of the locations in the third part of the stack. We therefore decided to use this part of the stack for both regular containers and reefers (for reshuffles and new containers). The reefers can still be stacked onto the last (one, two or three) piles of the second part of the stack. We expect this feature to generate a better use of ground locations and thus reduce the number of reshuffles. An obvious disadvantage of this feature is that the number of available positions for reefers is reduced.

2.4.6 Exchanging containers from different lanes

Categories can also be used to select a container from a different lane in the loading operation. We can use this to avoid overloading an ASC. This feature is therefore triggered if the ASC in a selected lane is too busy. The algorithm scans all lanes of the stack for a lane that contains a container of the required category and an ASC that has a workload below the predefined limit.

2.4.7 Using the expected departure (residence) time of the containers

The expected departure time can be used to store containers that will leave shortly, on top of containers that will stay in the stack for a longer period. This feature is used whenever a container has to be stacked and there is no non-full pile of that category. The container will then be stacked on top of a container for which the expected departure time is later than the expected departure time of the incoming container.

The expected departure time for jumbo and deepsea containers is approximated by the middle of the time interval during which the ship lies alongside the quay. For the other modalities, the average dwell time of a container is approximately half a week: the expected departure time is therefore approximated by adding 3.5 days to the time of arrival. Note that this option does not use detailed information. It can also be applied if no information on the departure time is available.

2.4.8 Choosing the ASC that has the lowest workload

The lowest ASC workload feature can be used for both incoming and outgoing containers. For incoming containers, creating uniform piles takes precedence over the lowest workload. Thus, a container will be stacked on top of a uniform pile of the same category even if the ASC for that lane is very busy. If there are uniform piles in multiple lanes, then the lane with the lowest ASC workload will be selected. For outgoing containers, we select the lane with the lowest ASC workload from lanes in which containers from the target category are stored.

2.5 Experiments

The following data applies to all experiments. The stack has 27 lanes for experiments A0 to C and 29 lanes for all other experiments; a lane is 40 TEU long, 6 TEU wide, and the maximum stacking height is 3 containers. Categories and exchanges are possible for jumbo, deepsea as well as for rail and barge; temporary substacks are used for rail and barge to loosen the loading order restrictions when leaving the main stack. Category stacking is applied for all modalities, except where stated differently. The experiments are listed in appendix A and the numerical results from the experiments are in appendix B. We will now describe the experiments and analyze the results.

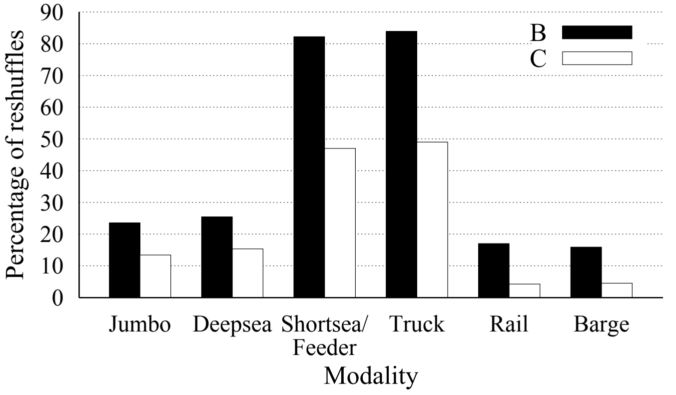

Figure 2.3: The effect of preference for ground positions on reshuffles.

2.5.1 Base Case

A null experiment (A0) uses random stacking without the possibility to exchange containers of the same category for the same (jumbo or deepsea) ship. The number of reshuffles in case of random stacking is high (89%). Although it is hard to validate such stacking programs, the number is not considered unrealistic by terminal operations people.

Experiment A considers category stacking for all modalities without any of the additional features. This yields much better results than random stacking: the percentage of reshuffles drops from 89% to 46%.

In experiment B, shortsea/feeder and truck containers are not stacked as categories, because these containers are not exchangeable. The percentage of reshuffles for these containers is reduced significantly (shortsea/feeder from 112% to 82%; truck from 104% to 84%): the percentage of reshuffles for all other modalities has increased. The average use of ground locations rises from 65% to 70%.

2.5.2 Preference for ground positions

Experiment C extends experiment B with a preference for ground locations as discussed in section 4. This has a pretty large effect, mainly on the number of reshuffles and reshuffle occasions. On aggregate, those percentages are approximately half of the percentages when using no preference. The percentages of reshuffles are shown in figure 2.3.

The workloads of the ASCs are also influenced by this preference, although the effects are moderate.

As expected, the number of empty piles drops, especially in part 1 of the stack. This causes an increase in the number of containers that can not be stacked. The probability that containers can not be stacked or can not be reshuffled is higher when there are fewer empty piles; on the other hand, the percentages of reshuffles and reshuffle occasions are lower.

In this case almost one out of every 1,000 containers can not be stacked, which is a very high proportion. One way to reduce this number is to increase the size of the stack. Therefore we added two lanes (29 instead of 27) to create experiment D (this configuration of the stack will be used for all other experiments). As can be expected, this decreases the number of reshuffles as well as the average workloads and the proportion of containers that can not be stacked. Finally, it also reduces the use of ground locations a little.

2.5.3 Workload control

In experiment E we add a workload control variable for each lane. A container is not stacked into a certain lane when the workload of the ASC in the current quarter exceeds 80%. This workload control variable is only used when (un)stacking regular (non-reefer) containers.

The workload control variables do not affect the reshuffles. The aim of this feature is to reduce the number of busy or very busy ASC quarters. The most significant impact can be observed in the percentage of ASC quarters with a workload over 100% during jumbo operations: this percentage drops from 11.8% to 8.3%.

2.5.4 Improved reefer stacking

The next experiment (F) adds the modified category stacking policy for reefers to the setup of experiment E. Experiment G adds workload control for reefer containers: the limit is set to 80%.

This seems to have a few positive effects. The overall number of reshuffles and reshuffle occasions are reduced (resp. from 19.0% to 16.0% and from 13.4% to 11.4%). There are no more reshuffles for reefer containers (this can not be deduced directly from the table). In addition, it is now possible to find a position for all new containers and reshuffles. Finally, the use of ground locations in the third part is much lower when using category stacking for reefer containers.

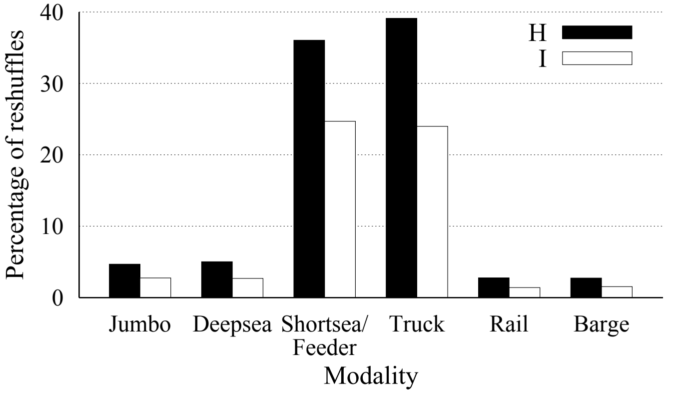

Figure 2.4: The effect of combining parts 1 and 3 of the stack on reshuffles.

Adding a workload control variable for reefers (experiment G) has a (small) positive effect (it reduces the proportions of busy ASC quarters a little).

2.5.5 Use ground position closest to transfer point for unstacking

In experiment H, whenever a container is to be stacked on an empty pile in a lane, we select the pile that is closest to the transfer point where the container will leave the stack. The result is an increase in the use of ground locations in part 1 of the stack (both average and maximum) with approximately two percent. The overall percentage of reshuffle occasions decreases from 11.4% to 10.4%; the percentage of reshuffles drops from 16.2% to 14.8%. The percentage of quarters with a high workload is lower during jumbo handling (7.8% versus 6.3% quarters with a workload over 100%). This is also true during deepsea handling and overall.

2.5.6 Combine parts one and three of the stack

Experiment I was motivated by an observed low use of ground locations in the third part of the stack. Thus, experiment I extends experiment H with the option to stack regular containers in the third part of the stack. The average and maximum use of ground locations increase and lead to a clear reduction in reshuffles and reshuffle occasions (figure 2.4). Compared to experiment H, there is a significant number of reshuffles for which no position could be found (15 per 100,000 containers). The maximum workload for jumbo containers rises from 220% to 279%.

The total percentage of reshuffles decreased from 14.8 percent to 9.7 percent. The total percentage of reshuffle occasions dropped from 10.4 to 6.9. The percentage of busy ASC quarters has decreased from 6.3 to 5.6 for jumbo containers.

2.5.7 Exchanging containers from different lanes

To study the influence of this feature, we have defined experiments J and K.

J: Experiment I modified to exclude import containers from the closest-transfer point rule.

K: Experiment J, with the added possibility of exchanges between different lanes. The exchange candidate has to be on top of its pile. The algorithm looks for exchange candidates whenever the workload of the ASC for the original container exceeds 80%.

Experiment J does not yield favorable results in comparison to experiment I: the percentages of reshuffles and reshuffle occasions are higher. Adding the exchange from different lanes feature in experiment K causes the percentage of reshuffles to drop from 9.9 to 9.5. The primary purpose of adding this feature was to obtain lower proportions of ASC quarters with high workloads. Figure 2.5 below illustrates the overall percentages of high ASC workloads: the percentage of busy quarters is reduced significantly. We have explored several additional ways to implement this feature but the results are similar. From these experiments, we conclude that adding the possibility of exchanging containers from the same category within different lanes has a positive effect. It reduces the number of reshuffles and reshuffle occasions as well as the proportion of high ASC workloads.

2.5.8 Using the expected departure time of the containers

In practice, it is often difficult to obtain a reliable indication of the departure time. We therefore use of the expected departure time to order the containers when stacking, not as a reliable indicator of the actual departure time. For this feature, we have to define a boundary value that controls whether a container can be stacked on top of another one. When we make this restriction too loose, a lot of containers will be stacked on a container that will leave earlier, which causes a reshuffle. If, on the other hand, the restriction is too tight, we will make less use of the opportunity to use the expected departure times of the containers.

To get some insight into the effects of adding a rule based on the expected departure times of the containers, we can compare the results of the experiments K, L, and M. Experiment K makes no use of this rule; experiments L and M extend experiment K with the expected departure time feature. For experiment L, the value of the boundary is three hours after the expected departure time of the container already in the stack. Experiment M sets the boundary to the expected departure time of the container that is currently on top of the selected pile.

This feature was designed to lower the number of reshuffles. The percentages of reshuffles and reshuffle occasions are lowest for experiment L (8.8 and 6.2). For experiment M these percentages (9.6 and 7.0) are even higher than for experiment L (9.5 and 6.8). The restriction on the expected departure times may be too tight for experiment M. The differences between these experiments concerning the high ASC workloads are small. Furthermore, using the departure times of the containers leads to a somewhat higher use of ground locations.

2.5.9 Choosing the ASC that has the lowest workload

We have designed two experiments to determine the effects of starting in the lane for which the ASC has the lowest workload when stacking or unstacking. Experiment K is used for comparison.

N: Algorithm K with the ASC workload feature implemented for incoming containers for which multiple uniform piles in different lanes have been found.

O: Same as experiment N, with lowest ASC workload feature implemented for outgoing, regular (non-reefer) containers.

The percentages of reshuffles and reshuffle occasions increase when adding this feature. However, the feature was designed to improve the workloads, so figure 2.5 shows the percentages of high workloads for these experiments.

As we can see, the percentage of high workloads has indeed decreased by starting in the lane where the ASC has the lowest workload. We have also experimented with the lowest ASC workload rule for reefer containers and a lower maximum stacking height (two) for truck containers as an extension of experiment O: these experiments yielded no additional benefits.

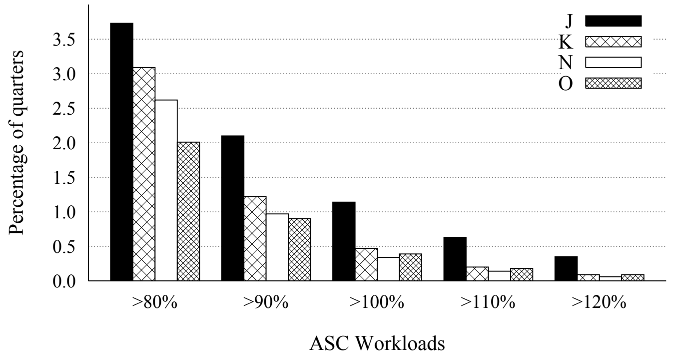

Figure 2.5: The effect of using ASC workload on the percentage of busy quarters.

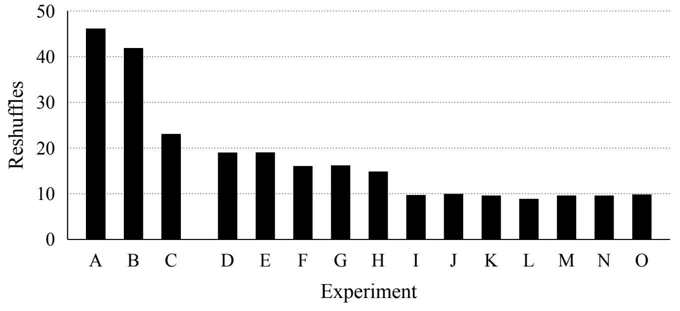

Figure 2.6: Overall percentage of reshuffles.

2.6 Comparison of All Scenarios

In this section, we will focus on the overall results rather than compare individual experiments. Again, we will visualize some of the results in graphs.

First of all, figure 2.6 indicates that the percentage of reshuffles can be significantly reduced. For our benchmark, this percentage was 46.1 percent; for experiment L, the percentage is just 8.8 percent. That is less than 20 percent of the initial percentage. The graph also shows that a number of other experiments have a similar percentage of reshuffles.

Maybe the most important performance measure is the proportion of busy ASC quarters. Figure 2.7 shows for all experiments the percentage of ASC quarters with the ASC working more than possible. We have decreased this value a lot. In the benchmark

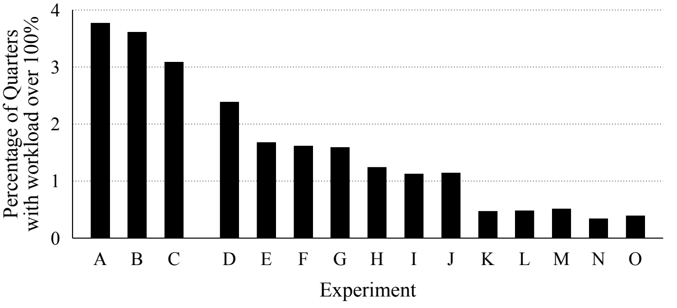

Figure 2.7: Percentage of ASC quarters with a workload over 100 percent.

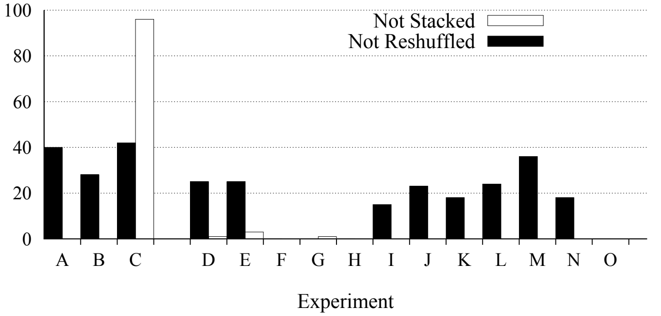

Figure 2.8: Containers that can not be stacked or shuffled.

case, this percentage is equal to 3.8 percent. The best result is obtained using experiment N (0.3%), but there are several experiments with similar performance (in terms of this percentage).

For some experiments there are (relative to the numbers for other experiments) a lot of containers that can not be stacked (either new containers or reshuffles; see figure 2.8). This is a highly undesirable effect. Note that, because of the fact that we just took these containers out of the stack, or did not stack them at all, this could also biases the results somewhat. The impact should be limited as the number of containers that can not be stacked is relatively small (the highest value is 96 per 100,000 containers).

2.7 Conclusions

In this chapter, we have investigated a number of policies for stacking containers in a yard by means of simulation. The following conclusions can be derived from the experiments.

Loading and unloading operations for jumbo containers ships creates workloads that exceed the capacity of the set of ASCs (27 to 29 in total). The average workload over time is well below a hundred percent but the workload during the handling of a jumbo ship is very high with many short-term bottlenecks. This means that the stack configuration is not able to follow the quay crane production.

Category stacking yields much better results than random stacking. Allowing exchanges for containers for the same category jumbo or deepsea ship further improves the results. The number of actual reshuffles and the number of reshuffle occasions can be reduced by adding a preference for ground locations. This also reduces the ASC workloads. There is however the possibility of creating a higher proportion of non-stackable containers due to the reduced number of empty piles; this feature requires careful implementation.

Treating containers for shortsea/feeder, rail, truck and barge as categories to be stacked together seems to have no large effect on the whole. Although it reduces the number of reshuffle occasions for jumbo and deepsea, this number increases for the other modalities. Using fewer piles, on average, for the same containers leads to a higher number of reshuffles. The effect of stacking piles with only truck containers up to height two is negligible.

The peaks in ASC workloads can be reduced by adding a workload control variable as well as stacking on piles close to the transfer point where the containers are to leave the stack. Finally, the possibility of exchanging containers of the same category within different lanes decreases the proportion of high workloads as well.

The definition of the categories is based on parameters used in stowage planning. This allows online optimization in which we can avoid sub-optimal yard operations that might be caused by a predefined (offline) stowage plan.

Overall, we conclude that detailed simulation experiments of the stacking operations can drastically improve the stacking performance and is thus essential for constructing automated container terminals.

Appendix A: Experiments

A0. This is a reference experiment that uses random stacking without exchanges.

A. This experiment considers category stacking for all modalities.

B. Category stacking without exchangeability for shortsea/feeder and truck containers.

C. Same as the previous experiment, with an added preference for ground locations in the random part of the algorithm.

D. Experiment C with 29 lanes instead of 27.

E. Experiment D, with the workload control variable set to 80% rather than 88.9%.

F. Experiment E with alternative reefer stacking policy (no workload control variable for reefer containers though).

G. Experiment F, with reefer containers also subject to the workload control feature with a limit of 80%.

H. Experiment G with the closest transfer point feature. This feature selects an empty pile closest to the transfer point at which the container will leave the stack.

I. This setup is based on experiment H: We allow the stacking of regular containers in the third part of the stack (this part is usually reserved for reefer containers).

J. In this modification of experiment I, we exclude import containers from the closest-transfer point rule.

K. Experiment I, with the option of exchanges between different lanes. Exchanges are considered whenever the ASC workload of the selected lane exceeds 80%. Feasible exchange locations are limited to the top containers of each pile and are located using a random search approach.

L. Same as experiment K, with the expected departure time rule: a container can only be stacked on top of other containers if the new container has an expected departure time less than three hours after the expected departure time of the current topmost container in the pile.

M. Experiment L, but the expected departure time of the new container must be before or equal to the expected departure time of the topmost container of the pile.

N. Experiment K with the ASC workload feature for incoming containers for which multiple uniform piles in different lanes have been located.

O. Experiment N, with the addition of the ASC workload feature for outgoing, regular (i.e., non-reefer) containers.

Appendix B: Numerical results of the experiments.

The table below contains the numerical results of the experiments.

Experiment	A0	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
Reshuffle Occasions																
Total	60.9	31.0	31.0	16.1	13.3	13.4	11.4	11.4	10.4	6.9	7.1	6.8	6.2	7.0	6.8	6.9
Jumbo	67.8	12.4	18.5	10.2	8.5	8.3	4.3	4.1	3.6	2.1	2.3	2.1	1.8	2.4	2.0	1.4
Deepsea	55.9	12.3	19.5	11.4	9.4	9.5	5.0	4.8	4.1	2.2	2.5	2.2	2.4	2.8	2.3	2.3
Shortsea/feeder	62.8	71.5	58.6	31.5	26.1	26.6	26.3	26.6	24.5	17.3	17.5	17.2	15.2	16.9	17.1	17.9
Export	55.9	35.1	34.2	18.7	15.5	15.6	12.9	12.9	11.8	8.0	8.2	7.9	7.1	8.1	7.9	8.0
Truck	–	68.0	58.4	32.3	26.7	26.4	26.5	26.6	26.0	16.2	16.3	14.9	12.5	14.9	16.1	16.1
Rail	–	9.5	14.7	3.8	3.0	3.2	2.7	3.1	2.5	1.3	1.2	1.1	1.5	1.5	1.3	1.0
Barge	–	8.8	13.8	3.9	2.8	2.9	2.8	2.9	2.4	1.4	1.4	1.1	1.5	1.6	1.2	1.3
Import	–	19.2	21.8	8.7	7.0	7.1	6.9	7.0	6.5	3.9	3.9	3.5	3.4	3.8	3.8	3.7
Reshuffle Performed																
Total	89.3	46.1	41.8	23.0	19.0	19.0	16.0	16.2	14.8	9.7	9.9	9.5	8.8	9.6	9.5	9.8
Jumbo	99.7	15.7	23.5	13.4	11.3	10.8	5.5	5.3	4.7	2.8	3.0	2.8	2.0	2.8	2.6	1.8
Deepsea	81.8	16.4	25.4	15.3	13.0	13.1	6.1	5.8	5.0	2.7	3.0	2.6	2.8	3.1	2.7	2.7
Shortsea/feeder	92.1	112.2	82.2	47.0	38.5	39.0	38.3	39.3	36.0	24.7	25.1	24.7	23.0	24.1	24.5	26.1
Export	81.9	52.9	46.6	26.8	22.2	22.3	18.3	18.5	16.8	11.2	11.5	11.2	10.3	11.1	11.0	11.4
Truck	–	103.5	83.9	49.0	40.2	39.2	39.4	39.5	39.1	24.0	24.3	22.6	20.1	22.7	24.2	24.8
Rail	–	10.8	17.0	4.3	3.4	3.7	3.0	3.5	2.8	1.4	1.4	1.2	1.6	1.6	1.4	1.1
Barge	–	9.8	15.9	4.5	3.2	3.4	3.2	3.3	2.7	1.5	5.4	1.2	1.6	1.6	1.4	1.3
Import	–	26.2	27.9	12.1	9.6	9.6	9.4	9.6	9.0	5.3	9.9	4.9	4.8	5.2	5.3	5.3
No Position (per 100,000)																
For new container	0	0	0	96	1	3	0	1	0	0	0	0	0	0	0	–
For reshuffle	74	40	28	42	25	25	0	0	0	15	23	18	24	36	18	–
Ground Locations: Maximum																
Overall	84.5	75.3	79.9	91.1	89.2	89.1	88.7	87.8	89.4	95.2	94.7	95.7	96.4	96.5	95.1	95.2
Part 1 of the stack	89.6	80.8	86.3	98.7	97.7	97.5	97.5	97.4	99.4	98.4	97.8	98.9	99.6	99.8	98.4	98.3
Part 2 of the stack	50.6	50.2	51.4	48.6	48.1	46.7	58.2	58.2	58.2	58.2	58.2	58.2	57.9	57.9	58.2	58.2
Part 3 of the stack	68.9	50.4	50.6	53.5	42.2	40.2	22.4	22.2	14.4	–	–	–	–	–	–	–

Experiment	A0	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
Ground Locations: Average																
Overall	77.4	64.7	70.3	81.3	79.0	79.0	78.3	78.2	79.4	83.9	83.5	84.0	85.1	85.3	83.9	84.3
Part 1 of the stack	83.5	69.6	76.5	89.0	87.5	87.4	87.4	87.4	89.3	87.6	87.1	87.7	88.9	89.0	87.5	88.0
Part 2 of the stack	38.7	38.9	38.6	39.2	37.5	38.1	38.7	38.7	38.7	38.7	38.7	38.6	38.7	38.6	38.6	38.7
Part 3 of the stack	47.1	36.0	31.4	36.5	24.5	24.8	13.6	13.3	8.8	–	–	–	–	–	–	–
Workload ASCs: Overall																
Maximum (%)	301.9	312.6	289.4	270.4	302.5	260.8	246.5	287.8	220.2	278.9	258.9	259.1	193.6	238.1	214.0	367.3
Average (%)	31.2	27.0	26.8	25.4	23.3	23.4	23.3	23.3	22.5	22.3	22.3	22.3	22.3	22.4	22.1	22.0
Percentage > 80 %	10.3	7.3	7.1	6.1	5.1	4.7	4.6	4.6	3.9	3.7	3.7	3.1	3.1	3.2	2.6	2.0
Percentage > 90 %	7.7	5.3	5.1	4.4	3.5	2.8	2.7	2.7	2.2	2.1	2.1	1.2	1.2	1.3	1.0	0.9
Percentage > 100 %	5.7	3.8	3.6	3.1	2.4	1.7	1.6	1.6	1.2	1.1	1.1	0.5	0.5	0.5	0.3	0.4
Percentage > 110 %	4.1	2.7	2.5	2.2	1.6	1.0	1.0	1.0	0.7	0.6	0.6	0.2	0.2	0.2	0.1	0.2
Percentage > 120 %	3.0	1.9	1.7	1.5	1.1	0.6	0.6	0.6	0.4	0.3	0.4	0.1	0.1	0.1	0.1	0.1
Workload ASCs: Jumbo																
Maximum (%)	–	312.6	289.4	270.4	302.5	260.8	246.5	225.3	220.2	278.9	258.9	259.1	193.6	218.0	214.0	367.3
Average (%)	–	59.5	60.4	57.6	53.2	53.2	53.0	52.9	50.0	49.6	49.8	49.5	49.4	49.8	49.4	49.0
Percentage > 80 %	–	28.4	28.6	25.9	22.2	21.3	21.1	20.7	17.7	16.8	17.3	14.5	14.6	15.3	12.7	9.5
Percentage > 90 %	–	22.3	22.0	19.8	16.5	13.3	13.1	12.8	10.5	9.8	10.4	5.9	6.0	6.3	4.8	4.4
Percentage > 100 %	–	17.1	16.5	14.7	11.8	8.3	8.0	7.8	6.3	5.6	5.9	2.3	2.3	2.6	1.7	1.9
Percentage > 110 %	–	13.0	12.3	10.8	8.4	5.2	5.0	5.0	3.7	3.2	3.3	1.0	1.0	1.1	0.7	0.9
Percentage > 120 %	–	9.6	8.8	7.8	5.9	3.1	3.1	3.0	2.1	1.8	1.9	0.5	0.4	0.5	0.3	0.5
Workload ASCs: Deepsea																
Maximum (%)	–	266.3	242.7	239.8	226.5	190.9	199.0	287.8	206.5	204.3	197.9	176.6	159.6	238.1	213.4	205.8
Average (%)	–	39.7	40.1	38.3	35.2	35.3	35.1	35.0	33.2	33.0	33.1	33.0	33.1	33.2	32.9	32.5
Percentage > 80 %	–	11.8	11.8	10.1	8.1	7.3	7.0	7.0	5.7	5.4	5.4	4.2	4.4	4.3	3.3	2.8
Percentage > 90 %	–	8.0	8.2	6.8	5.1	4.0	4.0	4.0	2.9	2.9	2.8	1.7	1.8	1.6	1.2	1.2
Percentage > 100 %	–	5.4	5.4	4.4	3.2	2.2	2.2	2.2	1.6	1.5	1.4	0.6	0.7	0.6	0.4	0.5
Percentage > 110 %	–	3.5	3.5	2.9	2.0	1.3	1.2	1.2	0.8	0.7	0.7	0.2	0.3	0.3	0.2	0.2
Percentage > 120 %	–	2.3	2.2	1.8	1.1	0.7	0.7	0.7	0.4	0.4	0.4	0.1	0.1	0.1	0.1	0.1

Chapter 3

Online Rules for Container Stacking*

In this chapter we continue our research into container stacking rules. In order to further investigate the performance of these stacking rules, we have taken the research design from the previous chapter and reduced its complexity in order to facilitate the analysis. In particular, we have reduced the complexity of the stacking rules and of the streams of containers that enter and depart the stack. This chapter was written as a separate article that has been submitted for publication to an academic journal.

We investigate two concepts to increase efficiency and compare them to several benchmark algorithms, using a discrete-event simulation tool. The first concept is to use knowledge about container departure times, in order to limit the number of reshuffles. We stack containers leaving shortly before each other on top of each other. The second concept is the tradeoff between stacking further away in the terminal vs. stacking close to the exit points and accepting more reshuffles. It is concluded that even the use of imperfect or imprecise departure time information leads to significant improvements in efficiency. Minimizing the difference in departure times proved to be important. It was also found that the tradeoff between stacking further away in the terminal vs. stacking close by the exit points and accepting more reshuffles leads to improvements over the benchmark.

3.1 Introduction

One of the main problems in container terminals concerns the stacking of containers. Although it is also one of the main advantages of containers, viz. that they can be stacked on top of each other, additional work is required if the bottom container is

* This chapter is under review for publication in OR Spectrum

needed. In that case the top containers have to be moved to another place, which is called a reshuffle or unproductive move.

Accordingly, every terminal needs a stacking strategy. The main objectives of such a strategy are 1) the efficient use of storage space, 2) limiting transportation time from quay to stack and beyond, (and vice versa), and 3) the avoidance of reshuffles. Of course, the importance of each criterion depends from terminal to terminal. Ports like Singapore and Hong Kong have limited land space, so they need efficiently used storage spaces. Note also that these objectives are conflicting: you cannot maximize them all. For example, the third objective would be optimized by having stacks of only one container high; however, this would lead to very inefficient use of storage space.

In stacking containers, various decision horizons can be identified. An often used classification has four temporal categories: long term (years, decades), medium-term (months), short-term (days) and real-time (minutes, seconds).

Long term decisions are strategic decisions, e.g. decisions concerning the type of equipment (automated/manual), the stacking height and the location and size of the stacking area.

Medium-term (or tactical) decisions concern capacity decisions, such as stack layout, number of vehicles used to move containers about and whether or not (and how) to do remarshalling (i.e. performing reshuffles) in the yard when no ships are being served.

Short term (or operational) decisions concern finding the storage location for a particular container and the allocation of the equipment to the various jobs scheduled in the coming hours.

Real-time decisions are decisions made when actually executing whatever part of the stacking process. It includes the speed and direction control of all vehicles, as well as the cranes, and is hence mostly of importance to automated equipment.

This paper focuses primarily on the short term decision to allocate an incoming container to a stacking position. We try to mimic the most common situation where imperfect or imprecise information about the departure time of a container is available. Moreover, we consider online stacking rules, which do not require extensive computations and can be used in many types of stacks and large number of containers. We concentrate on the trade-off between traveling and finding a position which limits the likelihood of reshuffles. We use a quite realistic simulation program to test our ideas. As benchmark we take both random stacking policies as well as policies which use precise information on the container departure times.

The paper is structured as follows. We start with giving an overview of existing literature on stacking in section 3.2. Next we explain the set-up of our simulation model in section 3.3. The basic stacking concepts are explained in section 3.4. The experimental set-up is presented in section 3.5, while section 3.6 presents our benchmark algorithms. The results from the experiments are given in section 3.7 and we finish the paper with conclusions in section 3.8.

3.2 Literature Review

Academic literature on stacking problems is not very common yet, perhaps because the problem does not easily lend itself for analytical solutions Dekker et al. (2006). However, in recent years, the subject seems to get more attention, because its importance is recognized Steenken et al. (2004). In a recent overview paper on operations research at container terminals, Stahlbock and Voß Stahlbock and Voß (2008) looked at a number of aspects of container terminal operations. Among the topics surveyed were stowage planning, berth allocation, crane optimization, terminal transport optimization, and storage and stacking logistics. Their work is an extension of an earlier overview Steenken et al. (2004), which also contains a paragraph on how stacking is done in practice.

Various methods are used to tackle the stacking problem, but two main approaches can be distinguished like in job scheduling Dekker et al. (2006). *Analytical calculations* with full information on the moment a container will be retrieved from the stack. They are often based on integer programming and take relatively much computation time. Next there are *detailed simulation studies* which evaluate various stacking strategies. These strategies can be *online*, in which they determine for each container separately where to place it independently of other incoming containers and *offline*, where locations are found simultaneously for all containers to be offloaded from a ship. So far only online rule-based strategies have been studied in simulation studies. These rules can handle imperfect or imprecise information on departure times of containers. Their study takes a lot of time and the results may be dependent on the simulation set-up. Dekker et al. Dekker et al. (2006) distinguish two types of stacking strategies: category stacking and residence time stacking. The former strategy assumes that containers of the same category (e.g. having the same size, destination, weight, etc.) are interchangeable, and can thus be stacked on top of each other without the risk a lower container in a stack is needed before the ones on top of it have been removed. The latter strategy does not use categories, but instead looks at the departure times of the containers: a container can only be stacked on top of containers that all have a (planned) departure time that is later than the departure time of the new container.

Recent examples of the analytical approach include Kim and Hong Kim and Hong (2006a), who use branch-and-bound to find an optimal solution to a stacking problem and then propose several heuristics to try to come close to the optimum, and Kang et al. Kang et al. (2006), who use simulated annealing to find good solutions reasonably fast. Caserta et al. Caserta et al. (2010) combine metaheurstics and dynamic programming to improve upon the known results of Kim and Hong Kim and Hong (2006b). The problem most of these optimization approaches have, is that they assume perfect prior knowledge on the order in which the containers will be picked up. However, this information is usually not known in advance. Nevertheless, finding a theoretical optimum can be very useful as a benchmark (although the metaheuristic approaches used to make this approach computationally feasible are not guaranteed to find the global optimum). Other methods that have been used for this include Q-Learning Hirashima et al. (2006) and critical-shaking neighborhood search Lim and Xu (2006). Han et al. Han et al. (2008) use integer programming with tabu search to generate an entire yard template, which should minimise reshuffling moves. Froyland et al. Froyland et al. (2008) also use integer programming, optimizing the entire terminal in an effort to maximize quay crane performance.

Detailed simulations were performed by several authors. Dekker et al. Dekker et al. (2006) simulated different stacking policies for containers in automated terminals. In particular, several variants of category stacking (with up to 90 different categories) were examined and compared with a base case in which containers are stacked randomly. The simulations demonstrated very high peak workloads during the handling of very large container ships. Category stacking was found to significantly outperform random stacking. Considering the workload of each automated stacking crane (when selecting the lane for an incoming container) and stacking close to the export transfer point were found to provide additional performance benefits. There was no significant benefit to using category stacking for containers with onward transport by short-sea/feeder, rail, truck, or barge. As the category definitions are based on information used in stowage planning, they advocate an integration of terminal operations and stowage planning.

Duinkerken et al. Duinkerken et al. (2001) also used simulation and category stacking, albeit with only a limited number of categories. Several (reactive) reshuffling rules were tested. Reactive reshuffling means the reshuffling is done when a container on which other containers are stacked is demanded for retrieval, leading to a number of reshuffling operations. This in contrast to “proactive” reshuffling, which is done when stacking cranes are idle. They evaluated several reshuffling rules (random, leveling, closest position, and minimizing remaining stack capacity reduction). The use of categories was compared with a model that required specific containers and found that the categories lead to much better performance. Also, it was shown that the remaining stack capacity strategy lead to big improvements when compared to the other three. Duinkerken et al. Duinkerken et al. (2001) also tested two "normal" stacking strategies (i.e. for when a container has just arrived), namely random and with dedicated lanes for a quay crane. However, using dedicated lanes is hard to do in practice, as load plans are not known in advance. Also, this strategy did not yield much improvement over random stacking.

Saanen and Dekker Saanen and Dekker (2006a,b) went into great detail in simulating a (transshipment) container terminal with rubber-tired gantry cranes (RTG), carefully simulating all movements of the trucks and cranes. They made a "comparison between a refined, but still traditional, strategy for operating a transshipment RTG terminal with a simple random stacking strategy for this type of terminal", and measured the differences in quay crane productivity (in lifts per hour), as this is considered the most important indicator for terminal efficiency. They found that the differences between the strategies was very small (about 0.7 lifts per hour). However, it was found that the number of gantry movements is a major factor in limiting the quay crane productivity.

Finally, Park et al. Park et al. (2006) used simulation to determine the best combination of an AGV dispatching rule with a reactive reshuffling rule, for various amounts of containers and AGV's in an automated container terminal. Their goal is to minimize the number of reshuffling operations. Park et al. used the random, closest position (but with residence time taken into account) and minimal RSC reduction reshuffling rules from Duinkerken et al. (2001). In most cases, the minimal RSC reduction rule, combined with the Container Crane Balancing (CCB) dispatching rule (the AGV is sent to the crane which has the most containers waiting for it), led to the least reshuffling operations.

This paper elaborates on residence time stacking. In particular we consider several residence time classes and use that information to limit the number of reshuffles. We compare a number of stacking rules where we consider trade-offs between further traveling and the possibility of reshuffles. We also consider the case of full information, on one hand as a benchmark, but also to get some insight into the structure of good policies. The research approach and experimental setup of this paper build on prior work Dekker et al. (2006) but here we have limited ourselves to relatively simple stacking rules that use less information in order to get more insight into the basic performance of these rules.

3.3 Simulation Model

The simulation model that was developed for the experiments in this paper consists of two major components: a generator and a simulator. Although based on the same specifications as the simulator model described in Dekker et al. (2006), the code for both programs was rewritten from scratch. The existing code could not easily facilitate some of the new experiments. As the tools for developing discrete-event simulation models, especially the language and library that were used in the original implementation (Pascal and MUST) are less prevalent today, and programming languages in general have improved since the original implementation, we have chosen to rebuild the system in a modern programming language (Java) using a solid discrete-event simulation library (SSJ, L'Ecuyer and Buist (2005)).

The generator program creates arrival and departure times of some 76,300 containers covering a period of 15 weeks of operation, including a three week warm-up period to initialize the stack. The generator is based on the same data as the generator in Dekker et al. (2006), including sailing schedules and a modal-split matrix. The output of the generator is a file that contains the ship arrivals, details of the containers to be unloaded and loaded, and the specification of the destination of each container. The departure time is specified as the planned (a.k.a. expected) departure time and the actual (a.k.a. real) departure time. In the implementation of the generator program the actual departure time is generated first on the basis of the sailing schedules; the expected departure time is created from this actual departure by applying a perturbation function (the parameters of this function depend on the departure mode (ship, short-sea vessel, train, or truck)). The destination can be another deep-sea vessel or (for import containers) a short-sea vessel, barge, train or truck. For each container the location of the individual container within a ship is specified. The generator takes the detailed quay crane sequences for loading and unloading into account.

The average residence time of a container is 3.8 days; the 90%-percentile of the dwell time is 5.3 days, and the maximum dwell time is 8 days. The specifications of the input for the generator are detailed in Voogd et al. (1999); the current implementation is documented in Borgman (2009). The experiments in this paper are done with a stack that has a total capacity of 3,672 TEU; the average utilization of the yard is therefore $(76,300 \times 3.8)/(3,672 \times 12 \times 7) \approx 75\%$.

The simulator program reads the output of the generator and performs the stacking algorithms. The core of the simulator itself is deterministic: the stochastic components are in the generator and, optionally, in the stacking algorithm. This setup facilitates a comparison of stacking algorithms as any changes in the statistical output of the simulator must be caused by the stacking algorithm.

Within the simulation program, the containers are loaded and unloaded from ships and other transport modes (trucks and trains). The transport of containers from the quay to the stack is performed by Automated Guided Vehicles (AGV's); the simulation does not contain a detailed model of the AGV's (issues such as routing and traffic have not been modeled). Once an AGV with a container arrives at the stack, the Automated Stacking Crane (ASC's) for the lane is tasked with lifting the container from the AGV and storing it in the stacking lane. As previous research into this container terminal has shown the ASCs to be performance bottlenecks, they have been modeled in detail: the simulator calculates the time it takes for all the motion components (hoisting, lengthwise and widthwise movement) where hoisting and movement are sequential and the length- and widthwise movement are done simultaneously. There is a single ASC per lane (based on the ECT terminal, v.i.) and the simulation program maintains a job queue for each ASC. Containers that have to be reshuffled are always stored within the same lane. On the land side, the containers are moved to and from the stack using straddle carriers.

To verify that the generator works correctly, we have performed a number of tests on the output (such as testing the distribution of the container dwell times). The simulator was verified using a number of test scenarios for which the values of the statistical indicators could be determined analytically. Once the simulator had passed these tests, the simulator was benchmarked against the model described in Dekker et al. (2006); the performance of random stacking and category stacking (experiments A0 and A from Dekker et al. (2006)) was similar but not identical. The main difference in the two models is that the current simulation model has a very detailed simulation of the ASC's whereas the original model had a very simplistic model for the ASC's. As other authors (such as Axelrod (1997)) have noted, achieving numerical identity for simulation models is hard; the detailed descriptions required to achieve this can rarely be published in papers and are too much work (with little reward) to describe in internal reports.

3.4 Basic Concepts

In this section we present some generic concepts that form the basis of the stacking rules we will evaluate in this paper. The precise formulations will depend on the layout of the stacking area. While we only present these formulations for one particular layout (in order to clarify the presentation and analysis), the formulations can be adapted for other layouts.

We want to investigate the basic concepts of stacking containers in a yard. The core dilemma is that we would like to stack a container that arrives and departs at the quay side as close to the transfer point quay-side as possible because this will minimize the total travel time of the stacking crane when the container enters and exits the stack. As there are many of these sea-sea containers, this would require us to stack high. Unfortunately, when we start to stack containers on top of each other, we face the risk of stacking a container on top of a container that will depart before the incoming container. This will lead to a reshuffle, which takes time. We will thus have to balance the travel and hoisting time of the stacking crane and the time taken by reshuffle moves.

If we consider a single lane of the type of container terminal under investigation, we see that there is a single rail-mounted stacking crane that has to perform all the stacking moves for that lane. We can distinguish between containers that are moving into the lane (i.e., that are being stacked) and containers that are moving out of the lane (they are being "unstacked"). Containers can enter and leave the lane at two sides: at the quay side (for containers that are coming from or going to deep-sea ships) and at the land side (for all other modes of transport). Figure 3.1 provides a schematic overview of the terminal layout. The layout of this terminal has the stacking lanes perpendicular to the quay. Each lane has a length, a width, and a height. We will refer to a single line along the length of the lane as a lane segment. A layout within this terminal configuration will be denoted as 'number of lanes × length × width × height'; the basic configuration for our experiments will be '6 × 34 × 6 × 3'.

The first trade-off that is worthy of investigation is the trade-off between the time it takes the ASC to travel to a certain location and the amount of time required to (un)stack a container. For a container that has arrived on a deep-sea vessel and that will also depart on another deep-sea vessel, it is attractive to stack it as close to the transfer point at the quay side as possible. If we can stack the container close to the transfer point we save travel time of the ASC both when the container is stacked and when it is unstacked. Clearly the same applies for containers that arrive and depart at the land-side of the stack. There is no obvious best location for containers that arrive at the quay-side and will depart at the land side and for containers that arrive at the land side and will leave the stack at the quay side. In both cases the ASC will have to (in two stages) move the container along the entire length of the lane. At first glance stacking the container close to the planned exit transfer point seems beneficial; however, this would imply a longer travel time of the crane. Since sea-to-land moves will occur most when large ships are being unloaded, it would seem more beneficial to stack these container as fast as possible in order to release the crane more quickly for other moves. This would however conflict with the desire to stack sea-sea containers

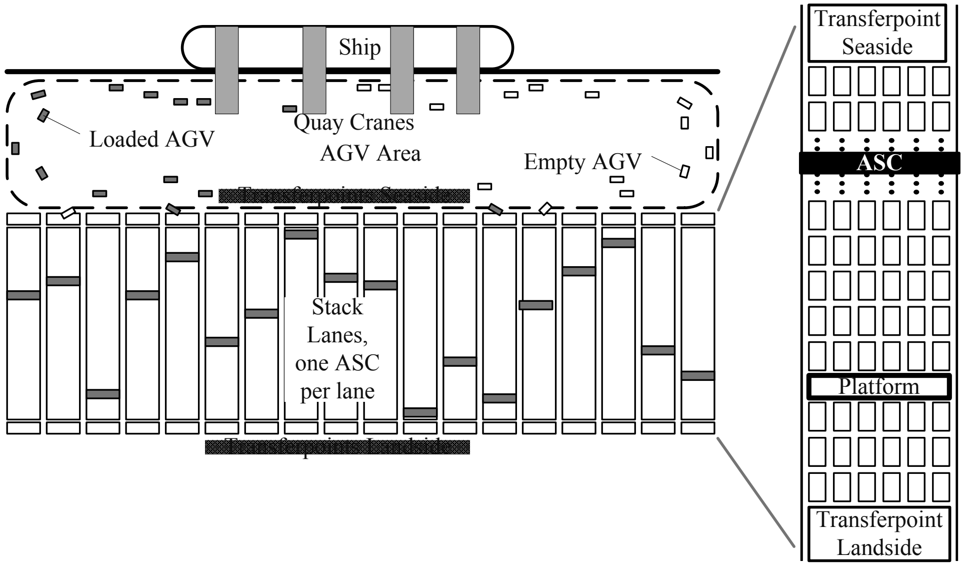

Figure 3.1: Terminal layout with details of a single lane

as close to the quay-side transfer point as possible.

The second trade-off we want to research is between the time required to stack a container and the number of reshuffles. Although reshuffles as such should be avoided, it is interesting to test a strategy that favors a fast stacking time during peak times with a resulting reshuffle that may occur at an off-peak time.

The overall approach of the experiments in this paper is focused on the operational decisions that have to be made by terminal operators. Specifically, we take the arrivals and departures that are specified as part of the generator output and perform these operations. There is no global optimization or explicit planning; the operations are performed one at a time, i.e. in a greedy fashion, whenever a container arrives. We do not consider future events such as other incoming containers.

3.5 Experimental Setup

The experiments in this paper all use the following configuration. The experiments are run for a 15 week period, of which three weeks are used for warm-up (to initialize the stack). Each experiment consists of ten replications in order to get statistically robust results and compute the 95%-confidence intervals of the means.

There are sufficient AGVs and straddle carriers to ensure that these resources do not act as bottlenecks. The basic configuration for the stacking area is modeled on part of the automated ECT Delta Terminal at the Port of Rotterdam, The Netherlands. We have chosen to use only a part of the actual stack area in these experiments to clarify the discussion and to facilitate the analysis. Thus, our stacking area has far fewer lanes than the actual terminal; the length of the lanes, the maximum stacking height, and the number of ASC's per lane are based on the configuration at the ECT Delta terminal.

There are thus just six lanes, each equipped with a single ASC. Each line is 34 TEU long, for a total of $6 \times 34 = 204$ ground positions (measured in TEU). (The Delta terminal has some additional room for reefer containers in each lane but we have not taken these containers into account for our experiments so we present the layout without this reefer area.) The maximum stacking height is three containers. Most experiments in this paper are done with six lane segments per lane; we denote the configuration as $6 \times 34 \times 6 \times 3$ (six lanes, each lane being 34 TEU long, 6 segments per lane, and a maximum stacking height of 3 containers). The number of lanes and the maximum stacking height will be changed for some experiments to evaluate the performance of the stacking rules under investigation. For these experiments we use a single size of container, the 20 ft. container. All other types were removed from the generator's output, including reefers. The number of 20 ft. containers is a good fit with the base layout.

We will use Random Stacking (RS) and an implementation of the Leveling algorithm (LEV) described in Duinkerken et al. (2001) as benchmarks for the experiments.

3.6 Benchmark algorithms

In this section we introduce the basic algorithms used for comparison in the experiments.

3.6.1 Random stacking

Random stacking is a straightforward way of determining a stacking position for a new container. Basically, the new container is placed at a randomly chosen allowed location, with every allowed location having an equal probability of being chosen. We have implemented this as follows:

1. Select a random lane
2. Select a random position in the lane

3. Check whether we could stack at this position
4. If so: stack here
5. If not: start again in the next lane

Given enough tries this algorithm is guaranteed to find an available location (in our implementation, we have set the limit at 5,000 tries, which has proven sufficient). This algorithm is applied for reshuffling, with the difference that we then only want to search the lane the container is in.

3.6.2 Leveling

The idea is to fill lanes in layers, so that all empty ground positions are filled with containers first, before containers are stacked upon others. The stacking lane is filled from the transfer point quayside on. This strategy is taken from the earlier work of Duinkerken et al. Duinkerken et al. (2001). It is an intuitive strategy, but it does not use most of the available information. We thus get the following steps:

1. Choose a random lane with at least one available position.
2. Search for the first empty location, from the transfer point quayside towards the transfer point landside, row for row (i.e. widthwise).
3. If found: stack there
4. If not found: search all existing piles (of the same size and type), from the transfer point landside towards the transfer point quayside, row for row, for the lowest (i.e. search lowest piles first) stack location and stack on the location found first.

This algorithm is also applied for reshuffling, with the difference that we only search the lane the container is in.

3.7 Experiments

In this section we will present our experiments with a number of stacking rules. We follow the same structure for each experiment. We first present the design of the stacking rule. Next, we formulate a number of hypotheses regarding the performance of the stacking rule. The results are presented in tabular form and we discuss the results in terms of our hypotheses. The hypotheses are tested using the 95% confidence intervals of the mean; we accept that there is a significant difference if these intervals do not overlap. In the interest of clarity and brevity we only present a subset of the total experimental results; a comprehensive list of results is listed in Borgman (2009).

The performance of a stacking algorithm is measured with the following statistics:

Exit Time (ETQ and ETL). The exit time is the time (in hours) it takes to remove a container from the stack and have it ready for onward transport (to the quay or to a truck/train/barge). This time is measured for each side (quay-side and land-side) of the stack; they will be abbreviated as ETQ and ETL respectively. Exit time is the main performance indicator for a stacking algorithm. When a container enters the stack, the time is takes to perform this operation is determined by the workload of the ASC (how many jobs are in the current job queue) and the time it takes the ASC to move the container to its position. There are no reshuffles when containers are stored in the stack; reshuffles only occur when a container has to leave the stack.

ASC Workload (ASC). The automated stacking cranes are critical components for the overall performance so we measure the percentage of time that the ASC's are busy (an ASC is busy if is moving to a new location and while actually moving a container; it is idle otherwise). (The ASC workload will be denoted as ASC in the results.)

Reshuffles (RDC and ROC). For the unproductive reshuffle moves, we measure the number of reshuffles (denoted as RDC) as a percentage of the total number of container movements. To get an indication of the number of reshuffles that happen per move, we also measure the reshuffle occasions (as a percentage of the total number of container movements, denoted as ROC); a single reshuffle occasion implies one or more reshuffles. These numbers are not absolute indicators of performance as the time of the reshuffle is not taken into consideration. A reshuffle that occurs when the workload is low has less impact on the overall performance than a reshuffle during a peak workload, for example when (un)loading a very large vessel.

Ground Position Usage (GPU). We report only the average percentage of ground positions that are in use (denoted by GPU) as the various stacking strategies have differing preferences for stacking on the ground.

3.7.1 Experiment 1: Leveling with Departure Times (LDT)

The first experiment is on the influence of knowing the exact departure times of all containers. This can be exploited by only stacking containers on top of other containers, if the new container departs at an earlier time than the one below it, i.e. residence time stacking.

In practice the actual departure time is not known. However, it is valuable as a reference case to determine the best possible performance of such a stacking rule. For a more realistic scenario, we use the expected departure time that is also part of the generator output file.

We differentiate between containers both arriving and leaving at the quay and other containers (i.e. containers arriving or departing at the truck loading point). The containers both arriving and leaving at the quay (the sea-sea containers) should be stacked close to the transfer point quayside, because every meter they move towards the transfer point landside is a waste.

The dwell time, or better the departure time of a container, can be used when determining where to stack it. Containers departing before the containers below them will never lead to a reshuffle. We would like to exploit this to the maximum and stack as high as possible, because this means other positions remain free for other containers which depart later. The first priority when searching for a place to stack is thus to find the piles which have such a container on top.

Second, and again in the interests of keeping options open, we want to find the position where the difference between departure times is as small as possible, since this means that there are more possibilities for stacking other containers on top of them, using as little space as possible. We thus select from the piles the one where the difference in departure times would be smallest.

If we can find no pile with a container on top that will depart after the new container, we have to stack it elsewhere. Preferably, we do this on the ground, so no reshuffles will occur. From the available positions on the ground, we want to stack it as close to the transfer point as possible, so that travel times are minimized.

Should we still not be able to find a position, we have to stack the container on an existing pile, rendering a reshuffle inevitable. To minimize the number of reshuffles, we place the container on the highest pile available, so that no or few containers can be stacked on top of them, each of which would lead to another reshuffle. In case several of these piles are available, we select the pile closest to the transfer point, in order to minimize travel time for the ASC.

In summary, we use the following algorithm, which we shall call "LDT" *(Leveling* *with Departure Times)* (in each step we look at all lane segments):

1. Stack the new container (departing at $T = T_n$) at that pile, where the top container departs at time $T = T_o$, $T_o > T_n$, and $T_o - T_n$ is minimal and on which the container may technically be stacked (i.e. pile is not full).

2. If no position was found yet: stack the container on an empty ground location. Sea-sea containers are to be stacked as close to the transfer point quayside as possible.

3. Stack at a pile of the highest height available, as close to the transfer point as possible.

For sea-land and land-sea containers we do not have a preference for a particular part of the lane, since they have to traverse it in full anyway. However, because sea-sea containers prefer the sea (quay) side of the lane, sea-land containers should be stacked away from them, as close to the transfer point landside as possible, when the two types are both included in an experiment. Since we select the stacking location based on departure times, the two types may become mixed. If only sea-sea containers are included, the algorithm automatically stacks the containers near the transfer point quayside, but with land containers included, it may not do this. Since this removes the advantage of stacking near the quayside, we can choose to separate the piles, so that sea-sea containers may not be stacked on land-sea containers and vice versa. However, as this means there are less options for optimizing residence times, it remains to be seen which is best.

We compare the LDT algorithm with random stacking, leveling, and a modified version of random stacking (RS-DT), in which the algorithm searches for a random pile with the top container's departure time being after the new container's departure time. If no such pile is found, the container is stacked randomly, according to the random stacking algorithm. We use this to see which part of the differences between ordinary random stacking and the LDT algorithm is caused by the "random" part and which part is caused by the lack of perfect information.

Hypotheses

Because of the great advantage of perfect information regarding the departure times, we expect to see a very big improvement for relevant statistics, compared to random stacking and leveling. In particular, we look at the reshuffle percentages (which we expect will be lower for this algorithm), time to exit (will also be lower), ASC workload (will also be lower, as it is related to the previous ones), stack usage (will be slightly lower due to containers exiting quicker), and ground position usage (will be much lower than with leveling, which maximizes ground usage. We cannot say in advance how it compares with random stacking).

The effects of the RS-DT algorithm should be similar to those of LDT with regard to random stacking and leveling, because of the extra information. However, because the pile selection process is still very basic, it probably won't perform as good as LDT.

The experiments with expected, rather than actual, departure times should still be better than random stacking and leveling, i.e. the same effects should occur as with perfect information. We do expect these effects to be somewhat weaker, since the information is less reliable and hence some poor decisions are likely to be made.

On the basis of these considerations we formulate the following hypotheses:

Hypothesis 1.1 The LDT stacking algorithm will have a lower number of reshuffles, a lower exit time and a lower ASC workload than the benchmarks RS and LEV.

Hypothesis 1.2 The RS-DT stacking algorithm will have a better performance than RS, but worse than LDT.

Hypothesis 1.3 Mixing piles in the LDT stacking algorithm will lead to less reshuffles when compared to not mixing piles.

Hypothesis 1.4 Mixing piles in the LDT stacking algorithm will lead to higher exit times compared to not mixing piles.

Experimental setup

We have varied two parameters for this experiment; the first is the departure time (actual or real departure time versus expected departure time) and the second parameter controls whether mixed piles are allowed (mixed versus unmixed).

In this experiment we are particularly interested in the value of the perfect information regarding departure times. We therefore compare the results to random stacking. We would also like to know whether it is better to allow mixed piles or not. The difference is measured by looking at the times it takes for a container to enter and leave the stack.

Results

We have listed the results of all experiments in a single table to facilitate comparison (see table 3.4). The first column of table 3.4 lists the number of the experiment. The results for the benchmark algorithms are included as experiment "0". As expected, the LDT algorithm in its various forms outperforms the random stacking and leveling benchmarks. Less reshuffles occur and exit times and ASC workloads are lower. Ground position usage is also lower.

The relative performance of RS-DT is as predicted: better than the benchmark but worse than LDT. It is also of note that the modified random stacking algorithm (using real departure times), outperforms the LDT when the latter is using expected departure times. Using expected data, rather than actual, leads to a big performance drop for LDT.

The effects of mixed piles are inconclusive. When using actual data, mixed piles lead to less reshuffles but an increased exit time. However, using expected data, they lead to more reshuffles (and also an increased exit time). These effects are also found in the other stack layouts, albeit somewhat weaker.

We have tested this strategy for some larger stacks as well. In larger stacks there is little improvement in the results for expected times, unlike those for actual times. Moreover, the results for expected times actually deteriorate when going from 6 lanes, 3 high to 6 lanes, 4 high.

Discussion

From the results, it becomes clear that LDT's departure times have a big impact, even when they are not exactly known. Even partially random stacking, using departure times only to a limited extent, leads to big improvements across the board. Still, there remains a big gap between the results of expected and actual departure times, especially in bigger stacks, where little improvement, if any, is seen in comparison to smaller stacks. This is most dramatically the case when going from height 3 to 4 with 6 lanes, where performance actually drops, despite an increase of stacking options. This is probably due to any mistakes made being punished more heavily with extra reshuffles at higher stack levels.

Using mixed piles seems not to be such a good idea. In some cases, there is an improvement in the number of reshuffles (as was expected), but in others there is none. In all cases, mixed piles lead to longer exit times.

Our conclusion with respect to the hypotheses is:

Hypothesis 1.1: Confirmed

Hypothesis 1.2: Confirmed

Hypothesis 1.3: Rejected

Table 3.1: The five classes of container departure times used in experiment 2. The maximum time any container in the used arrivals file will stay is 192.2 hours.

	Actual Times		Expected Times	
Class	From (hrs.)	To (hrs.)	From (hrs.)	To (hrs.)
1	0	68.2	0	68.0
2	68.2	79.4	68.0	79.4
3	79.4	92.4	79.4	92.3
4	92.4	111.6	92.3	112.5
5	111.6	∞	112.5	∞

Hypothesis 1.4: Confirmed

3.7.2 Experiment 2: LDT with Departure Time Classification (DTC)

In the first experiment we have used the information from the generator program on the actual and expected departure times. We now take a different approach to model uncertain departure time information.

We use the data from the arrivals file to define a limited number of classes. The boundaries of these classes are calculated from the arrivals file by taking the quintiles or the 20th, 40th, 60th, 80th and 100th percentiles of the residence time. This gives us five classes of almost equal size, for the initial residence times at least. The classes used can be seen in table 3.1.

We use the algorithm from experiment 1 for this experiment too, only when the time difference is calculated we do not use the actual time or the expected time from the file, but instead use the class value from table 3.1 (based on either the actual or expected departure time). This means that lower classes will be stacked on top of higher classes, thereby ensuring no reshuffles occur, unless no suitable pile or ground position could be found.

Note that a container's class will change over time as its departure time comes nearer. This means that the lower classes will be more prevalent in the stack, because every high class will at one time become a low class.

Experimental setup

To test this algorithm, we used the settings of experiment 1 (see section 3.5), because we want to compare the different method of estimating departure times with the original and with perfect knowledge. We also experiment with classes based upon expected departure times, to see whether the double uncertainty gives any different results. We compare this algorithm to the entire experiment 1 (LDT), including random stacking (RS), leveling (LEV), and RS-DT. We also modified the RS-DT algorithm to work with departure time classes (this version being referred to as RS-DTC).

Hypotheses

We expect this algorithm to have a better performance than LDT-exp because it uses real departure times to define the classes. The algorithm may have a problem with small stacks, where space is in short supply and some suboptimal decisions may have to be made. In larger stacks, we expect that LDT-DTC will perform almost as good as normal LDT with real departure times. We expect similar effects with the use of expected departure times. Because the classes are then based on imperfect information, we do expect a drop in performance, when compared to LDT-DTC.

Hypothesis 2.1 The LDT-DTC algorithm with real departure times will have a lower number of reshuffles, a lower exit time, and a lower ASC workload than LDT-exp, but higher than LDT.

Hypothesis 2.2 The LDT-DTC algorithm with expected departure times will have a higher number of reshuffles, a higher exit time, and a higher ASC workload than LDT-exp and LDT.

Hypothesis 2.3 The LDT-DTC-exp will have a higher number of reshuffles, a higher exit time, and a higher ASC workload than LDT-DTC-real.

Results

The results for a $6 \times 34 \times 6 \times 3$-stack are in table 3.4. For brevity, we have not included the results for mixed piles, since their performance is similar to the previous experiment and they only complicate the presentation.

For the $6 \times 34 \times 6 \times 3$-stack, the LDT-DTC performs much worse than normal LDT with expected departure times. The random stacking version of the departure time classes algorithm (RS-DTC) also performs significantly weaker, compared to RS-DT, and even compared to LEV; in fact, the performance of RS-DTC is similar to RS. Interestingly, the results for LDT-DTC-exp are slightly (yet significantly) better than LDT-DTC-real. For the RS-DTC algorithm, the same applies. This is probably due to the following effect: when no "nice" stacking position can be found (this happens in as much as 40% of the cases, we found), a container is put on another pile, certainly causing a reshuffle when real departure times are used. However, when expected departure times are used, there is a small probability that, because of the error in departure time estimation, no reshuffle is caused. This leads to a slight advantage for expected times, and hence to this counterintuitive insight.

However, the results differ significantly for larger stacks, as can be seen in table 3.5 (a $8 \times 34 \times 6 \times 4$-stack). In this larger stack, there are no problems finding a suitable spot for the LDT-DTC-real algorithm, and thus its advantage of certainly knowing whether a reshuffle will occur is enough to yield better results. This leads to results almost as good as when using normal LDT. There is also a big improvement for the LDT-DTC-exp algorithm, which now actually performs better than LDT-exp. This is probably due to the used classes providing a bigger "margin of error", leading to less mistakes in determining which container departs first.

Interestingly, RS-DTC-real algorithm does not seem to benefit much from a larger stack. Apparently, its method of trying to find any suitable pile, without regard for the smallest class difference, leads to very inefficient stacking.

Discussion

Using the suggested five classes gives a very good result and a very good approximation of the results of actual departure times, provided there is enough space in the stack. For that case, we can confirm the hypotheses but as the results for the smaller stack differ, we can not confirm them.

Hypothesis 2.1: Rejected

Hypothesis 2.2: Rejected

Hypothesis 2.3: Rejected

3.7.3 Experiment 3: Traveling Distance vs. Reshuffling (TVR)

In this experiment, we do not use residence time knowledge, but rather try to optimise the selection of a location in some other way.

Suppose we only have sea-sea containers. As stated previously, these containers should be stacked as close to the transfer point quayside as possible, because they both arrive and leave the lane there. Leveling all the way down the lane would thus not be a good idea; we want to stack as high as possible at the beginning of the lane, and leave the end of the lane empty, if possible. Of course, we could just stack new containers as close to the transfer point as possible, by stacking them on top of each other, but

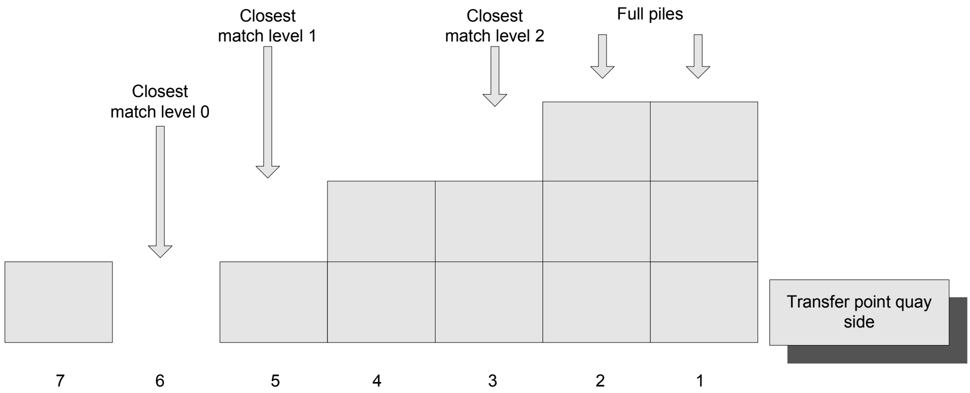

Figure 3.2: Example case for experiment 3. Maximum stacking height is 3 containers. The two piles to the right are closest to the transfer point, but are full. Three other options are available though.

this has the unwanted side effect of leading to reshuffles, because the order in which containers arrive and leave is mostly random.

However, since we also don't want to level too much, a compromise solution should be better. Whenever a container arrives at a lane (maximum stacking height n), we can say there are n possibilities to stack it: at every possible height layer, as close to the transfer point quayside as possible. For example, if $n = 3$, we may get the scenario from figure 3.2. The ground positions closest to the transfer point are full, but the next one is only stacked 2 out of 3 high. This is the best position from the distance viewpoint, but it has a high risk of leading to reshuffles, with two containers under it. The next position has the same risk, but it is further away, so therefore not as good, and we henceforth discard it for this decision.

We then encounter a pile with only one container. This means that the risk of reshuffling is lower, but unfortunately the pile is also further away from the transfer point. Still, it may be a good candidate to investigate. The sixth ground position is empty, and while it is still further away from the transfer point, there is no risk of causing reshuffles. This is the final candidate for stacking, because although there are more ground positions ahead, none of them is better in terms of reshuffling risk than the ones we already selected, and they are all worse in terms of distance to travel.

Now that we have found the "best" (note that we do not use residence times or categories here, these factors seriously complicate matters) three piles, we need to choose which one is best, i.e. which one has the lowest associated cost in time. Time costs consist of two parts: extra ASC driving time and time due to reshuffling. The latter will often be greater, but it is not certain it will occur.

The ASC driving time can easily be calculated using the distance from the transfer point to the selected ground position and the ASC's lengthwise and widthwise speeds. Lifting times can also be calculated, since we know exactly at which level the container is and will be stacked. The cost of reshuffling is more difficult to determine, because it is not known in advance whereto a container would be reshuffled. We assume this to be a (configurable) time period of driving away. The lifting times are also not known, so we estimate the lifting time to be that of a container in the middle of a pile (i.e. at the level of half the maximum height). The cost of a reshuffle is then multiplied by the expected number of reshuffles that putting the container at a particular place will generate. This gives the following cost function to minimize:

$$\min[2 \cdot TT_{qp} + L_{qp} + (TT_r + L_r) \cdot ER(n)] \tag{3.1}$$

Where TT_{qp} is the travel time from the transfer point quayside to the target pile, TT_r is the (fixed) travel time for reshuffles, L_{qp} and L_r are the lifting times per container (including pickup, lift up, lift down and set down) for the container itself and reshuffles, respectively, and $ER(n)$ is the expected number of reshuffles. The travel time to the pile has to be doubled because containers need to go back to the transfer point at one time; reshuffling can lead to a position closer to the transfer point, so we do not double that time. We will use TT_r as a penalty factor to vary the relative cost of reshuffles in our experiments.

A new container can, by itself, only generate one *extra* reshuffle, at most. Any other reshuffles were already in the pile when the container was stacked, or are added later. This means that we have to calculate the probability $P(nr)$ of the new container leading to an extra reshuffle. Since we assume every container in a pile has an equal chance of being chosen, this gives the following formula:

$$P(nr) = \frac{n-1}{n}, \tag{3.2}$$

where n is the height of the pile. We use this probability as $ER(n)$ in equation 3.1.

Summarizing, we use the following algorithm:

1. Select for every possible stacking level the (available) position closest to the transfer point quayside, if any position is available.

2. Calculate the costs of every position, according to equation 3.1 (using equation 3.2).

3. Select the position with the lowest cost and stack there.

While this algorithm can stack sea-sea containers of a single type, it is easy to extend the algorithm for other containers. As discussed in the previous experiment, sea-land and land-sea containers should be stacked as close to the transfer point landside as possible. We can determine costs in the same way.

Experimental setup

In this experiment we are particularly interested in the value of putting sea-sea containers close to the transfer point quayside. We therefore compare the results to random stacking. We would also like to know what the effects of different penalties (i.e., TT_r in equation 3.1) are. (Here, we report a subset of the experiments described in Borgman (2009) with values of the reshuffling movement penalty ranging from -0.03 to 0.04.) The difference is measured by looking at the times it takes for a container to enter and leave the stack.

We compare this algorithm with random stacking, leveling and a modified version of random stacking, which we shall refer to as TPRL (Transfer Point Random Level), in which the algorithm chooses one of the possibilities offered (i.e. it randomly selects the level where the container is to be stacked). This means containers will be near the quay, and since we know one of the positions is the "best" one, we can see the influence of the complicated calculations, when we compare it with chance.

Hypotheses

Again, we predict this algorithm will always outperform the basic random stacking and leveling algorithms. We again look at the reshuffle percentage (which we expect will be lower for this algorithm), time to exit (will also be lower), ASC workload (will also be lower, as it is related to the previous ones), stack usage (will be slightly lower due to containers exiting quicker) and ground position usage (will be somewhat lower than with leveling, which maximizes ground usage).

The ground usage depends on the reshuffling movement penalty applied. A higher penalty will lead to more containers being stacked on the ground and this leads to a higher ground position usage. A high penalty would lead to the algorithm behaving as normal leveling. This would also mean that the effect of allowing to stack 4 containers on top of each other, rather than 3, would be almost completely gone (there is still a minor effect on crane lifting times). Low penalties, on the other hand, would lead to a very low ground position usage, with high piles near the transfer points.

In any case, when only sea-sea containers are included, the algorithm will stack containers mostly next to the transfer point quayside and will stack less and less containers further away. This would mean that the average pile height would be decreasing in a monotone way, when going away from the transfer point.

We expect the modified random stacking algorithm TPRL to perform worse than the original random stacking (and, for that matter, all other algorithms tested in this section), especially when there is a lot of space in the stack. This is because modified random stacking will too often build high piles, while the other algorithms place more containers on the ground, which leads to less reshuffles.

On the basis of these considerations we define these hypotheses:

Hypothesis 3.1 The best TVR stacking algorithm will have a lower number of reshuffles, a lower exit time and a lower ASC workload than the benchmarks RS and LEV.

Hypothesis 3.2 The TPRL stacking algorithm will have a worse performance than TVR.

Hypothesis 3.3 The TPRL stacking algorithm will have a worse performance than RS and LEV.

Hypothesis 3.4 The TVR stacking algorithm is equal to LEV with a high penalty and only sea-sea containers.

Hypothesis 3.5 The best TVR stacking algorithm will lead to a monotone decreasing average pile height away from the transfer point quayside, when using only sea-sea containers. and TPRL.

Hypothesis 3.6 The TVR stacking algorithm, with very low penalties, will have a worse performance than RS, LEV, and TPRL

Results

See table 3.4 for the results of experiment 3. We can see that the TVR algorithm performed better than random stacking on all statistics, but compared to basic leveling the differences are minute. Even the best TVR version, in this experiment with a moving penalty for reshuffles of 0.0 hours, scores only marginally less reshuffles and slightly lower exit times. It should be noted, however, that these results are still statistically significant and outside the 95% confidence interval.

The very low penalty value of -0.03 hours results in very low ground position usage, meaning high piles. This in turn leads to more reshuffles and higher exit times.

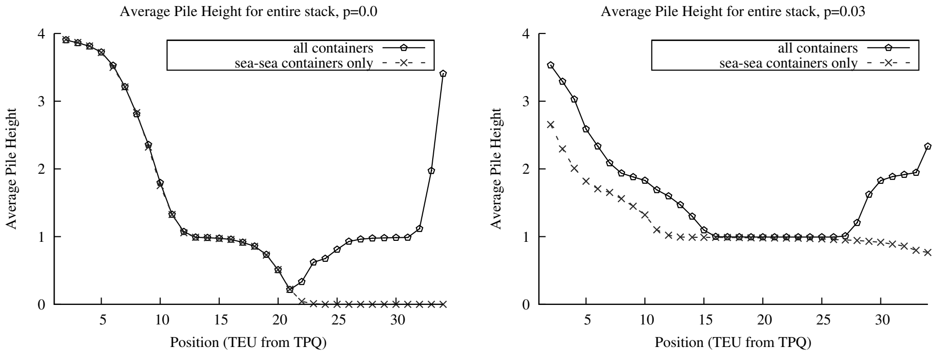

Figure 3.3: Time-average pile height for the entire stack for penalty value 0 (left) and 0.03 (right). Terminal layout: $8 \times 34 \times 6 \times 4$.

What further becomes clear is that penalties from -0.01 and higher yield almost the same results. Lower penalties give different (worse) scores (with the low ground position usage as predicted). With higher penalties, the results approach the benchmark result of leveling, but a penalty of 0.03 does not mean that the algorithm is the same yet. For larger stacks, the same behavior is observed, although the value of -0.01, from whereon the results are very similar, appears to be somewhat higher for larger stacks. (The full set of results are in Borgman (2009).)

The TPRL performs slightly worse than leveling, especially on larger stacks, in terms of exit times and reshuffle occasions, but not by the amount we expected. Moreover, it performs better than normal random stacking (RS). The number of actual reshuffles is higher though, because the ground position usage is lower and piles are higher.

To illustrate the effect of the penalty value on the pile heights along the length of the lane, we have graphed the time-averaged pile height in figure 3.3. We have selected a slightly larger and higher stack configuration ($8 \times 34 \times 6 \times 4$ layout) as this configuration provided the clearest graphical illustration. If the penalty is zero (left), then the average pile height is the same for both types of containers up to position 20. Beyond that, there are no sea-sea containers which means that the penalty value causes the two types of containers (sea-sea versus land-sea/sea-land) to be spatially separated. If we increase the penalty value, to e.g. 0.03, we see more sea-sea containers being moved along the entire length of the lane. The line for all containers shows that the sea-sea containers can no longer be stacked at the landside: more are now stacked towards the quayside and the average pile height increases.

Discussion

The TPRL algorithm performs not as bad as expected. This is likely partially due to the forced stacking near the transfer points, which leads to lower exit times, even compared to leveling. However, reshuffles are also consistently lower than with random stacking. A possible explanation lies in the fact that, on average, there is about a 33% probability (25% for stacking of height 4) of TPRL stacking in any of the possible stack layers. With random stacking, however, the probability for stacking on top of a high pile increases with an increased stack usage. In the biggest stack tested (8 lanes, 4 high, which naturally had the lowest stack usage), random stacking had a ground position usage of 76.1% (1632 piles on average) and a stack usage of 42.1% (average 2748.3 TEU in the stack at any time). This means piles were, on average, 1.68 (out of 4) containers high. Random stacking thus had a $1 - 0.761 = 23.9\%$ probability of selecting an empty ground position and 76.1% probability of selecting one of another height (which was, on average, 1.68). This gives the expected pile height of random stacking as $0.239 \times 0 + 0.761 \times 1.68 \approx 1.28$). TPRL, on the other hand, had 25% probability of choosing an empty ground position, and 75% of choosing an existing pile, of which the average height was $1 \times 0.25 + 2 \times 0.25 + 3 \times 0.25 = 1.5$. This gives the expected pile height stacked upon as $0.25 \times 0 + 0.75 \times 1.5 = 1.125$. Since this is a lower number, the expected number of reshuffles caused is also lower for TPRL than for random stacking. For smaller stacks, the probability for random stacking to stack on an existing pile only becomes greater, so this explanation applies there as well.

A reshuffle movement penalty of -0.01 hours appears to not lead to very bad results. There are only slight drops in performance compared to the higher penalties. We expect this is due to the feature of the algorithm which estimates the lifting time for reshuffles (L_r in equation 3.1). This variable is 0.019 hours for a 4-high stack. When the penalty of -0.01 is added, this still leaves a sizable reshuffling penalty of 0.009 hours, which is usually more than the cost of driving a little further, which the algorithm chooses to do. Hence, the value of -0.01 hours, which is not possible to have as a travel time in reality, still gives acceptable results.

The results show that different reshuffle movement penalties lead to different outcomes. Also, the best value in terms of the primary performance measures (ETQ and ETL) is not the same for every stack configuration. See table 3.2 for an overview of the best penalties per setup we found in the experiments. More specifically, a maximum height of 3 requires a lower penalty than one of height 4.

TVR performs better than the benchmark tests if we focus on our primary performance indicator, the exit time; when no residence time information is available, this is

Table 3.2: Best penalties found in experiment 3. Penalties are reshuffle movement penalties in hours.

Max Height	Lanes	Best penalty All	Best Penalty Sea-Sea
3	6	0.01	0.01
3	8	0.01	0.01
4	6	0.03	0.03
4	8	0.03	0.03
5	5	0.04	0.04

a good strategy to use. TPRL is also better than normal random stacking, but TVR yields much greater benefits and, since it requires no extra information, is the preferred option.

Thus, we accept the hypotheses H3.1, H3.2, H3.4, H3.5, and H3.6; we reject hypothesis H3.1.

Hypothesis 3.1: Confirmed

Hypothesis 3.2: Confirmed

Hypothesis 3.3: Rejected

Hypothesis 3.4: Confirmed

Hypothesis 3.5: Confirmed

Hypothesis 3.6: Confirmed

3.7.4 Experiment 4: Peak-Adjusted TVR

In experiment 3, we argued that sea-sea containers should be stacked as close to the transfer point quayside as possible, because every move further on is a waste of time. Likewise, we stated that sea-land and land-sea containers should be stacked near the transfer point landside, because for them distance does not matter (as they need to traverse the entire lane anyway), and they are out of the way for sea-sea containers there.

There is a slight problem with this reasoning, though. The above is true only if the time spent driving across the lane by ASCs for sea-land containers is valued the same at every point in time. However, since sea containers often arrive and depart many at a time (i.e. in a jumbo or deep sea ship), there are big peaks in crane workload. At these times, it would be not such a great idea to move sea-land containers all the way across the lane.

In this experiment, we try to counter this problem and extend the algorithm of experiment 3, by not putting the sea-land containers next to the transfer point landside, but somewhat further away. This is achieved by dividing every lane segment into two parts; one for sea-sea containers (near the transfer point quayside) and one for other containers (near the transfer point landside). We then stack all containers as close to the transfer point quayside, but in their own part of the segment. The size of the two parts is a configurable parameter. In the experiments it was set to 74% for sea-sea containers (which is roughly the fraction of that type in 20 ft. containers). There is also an option to allow stacking of sea-sea containers in the land part.

Experimental setup

In this experiment we are particularly interested in the value of putting sea-sea containers close to the transfer point quayside. We therefore compare the results to random stacking. We would also like to know what the effects of different penalties are. The difference is measured by looking at the times it takes for a container to enter and leave the stack.

It would not make much sense to test this algorithm with sea-sea containers only, because it is aimed only at improving the combination of both types.

Hypotheses

In this experiment, we expect roughly the same results as in experiment 3. The question is which algorithm will perform better.

We also expect a slightly lower exit time in experiment 4 than in experiment 3, when using a low penalty in both. This is because there should be slightly less pressure on the crane at peak times, while there is no other change (sea-sea containers should not be interfering with other containers). With a high penalty and a small stack, the leveling process is less efficient because of the two parts, which probably increases the exit time and the number of reshuffles, when compared to experiment 3.

Regarding the mixed or unmixed version of the algorithm, we expect there to be very little difference (if any at all) between both results when using a big stack. This is because there will likely not be a need for any sea-sea containers to be put in the "land" part, especially with low penalties, and even if there was a need, there is plenty of space. In small stacks, on the other hand, there will probably be larger differences, since a lack of space is far more an issue. We cannot predict in advance which version is best, because they both have their advantages and drawbacks. Mixed segments leave more room for the sea-sea containers, but this goes at the expense of land containers.

Thus, allowing sea-sea containers in the land part offers some extra possibilities, which may be needed in a small stack. However, it could also limit the options to stack sea-land containers, which could undo this. Generally speaking, a high reshuffling penalty will level out the stack, and thus also increase the number of sea-sea containers in the land part. Conversely, there should be very little difference in the results of allowing and not allowing the mix, when a low penalty (which encourages high piles) is used.

Our hypotheses for TVR-PA are:

Hypothesis 4.1 The best TVR-PA stacking algorithm will have a lower exit time than the best TVR.

Hypothesis 4.2 Allowing sea-sea containers in the land part in the TVR-PA stacking algorithm will lead to more reshuffles for smaller stacks (compared to not mixing).

Hypothesis 4.3 Allowing sea-sea containers in the land part in the TVR-PA stacking algorithm will lead to longer exit times (compared to not mixing).

Results

The results for experiment 4 with the stack configuration that is used throughout this paper are in table 3.4. Furthermore, we have also done some experiments with a stack that has a higher capacity and thus a lower utilization; see table 3.3 for the results of this experiment for a relatively big 8 lane, 4 high stack.

In the case of a small stack, the TVR-PA algorithm seems to perform slightly worse compared to normal TVR. Both exit times and reshuffles are up, as well as ASC workload. In the table two penalty values are shown, but these results also appear with other penalties.

With a larger stack, however, the results are less clear. With a low penalty, reshuffles are up, but with a high penalty they are down, all compared to the TVR equivalent. Normal TVR's exit times appear to be slightly lower than those of TVR-PA, but the differences are minute and well inside the 95% confidence intervals. ASC workloads also differ slightly, with a slightly lower value for normal TVR.

In both setups, there are differences between the results for mixed (i.e. allowing sea-sea containers to be stacked in the "land" part) and unmixed TVR-PA. Unmixed exit times are generally lower than mixed. These differences are small, but (for the small stack) outside the 95% confidence interval limits, so they are significant.

Table 3.3: Results of experiment 3 vs 4. Terminal layout: $8 \times 34 \times 6 \times 4$. The "90%" values are the average 90% percentile values.

Experiment	ROC %	RDC %	GPU %	ASC %	ETQ hrs.	ETL hrs.	90% ETQ	90% ETL
LEV	62.20	80.60	99.35	63.17	0.84	0.53	2.52	1.44
RS	69.32	134.11	84.68	69.41	1.83	1.15	5.48	3.64
TPRL	58.91	116.67	90.21	61.89	0.77	0.49	2.26	1.33
TVR (penalty 0.0)	45.64	92.03	89.30	45.66	0.21	0.16	0.38	0.28
TVR-PA (mixed, 0.0)	46.07	85.99	86.98	46.89	0.21	0.18	0.40	0.29
TVR-PA (unmixed, 0.0)	46.07	85.99	86.98	46.89	0.21	0.18	0.40	0.29
TVR (penalty 0.03)	45.12	64.38	99.52	47.20	0.19	0.17	0.34	0.26
TVR-PA (mixed, 0.03)	44.42	64.36	98.72	47.43	0.20	0.17	0.37	0.28
TVR-PA (unmixed, 0.03)	43.89	64.91	98.74	47.14	0.20	0.17	0.35	0.27

Discussion

When using TVR-PA, it appears that not mixing the two parts of the stack is best. The sea-sea containers in the land section take up much valuable space and also have to move further to get there. Whether to mix or not to mix only matters when space is tight, with large stacks, the algorithm almost never puts sea-sea containers in the land section.

TVR-PA is, compared to TVR, almost the same in terms of results. As predicted, TVR performs relatively best in a small stack, because the leveling part of the algorithm has to go up another stacking level a bit sooner. For larger stacks, it appears normal TVR still holds a tiny advantage. We can therefore not say that TVR-PA is better than TVR.

TVR-PA does not lead to improvements, compared to TVR. Mixing TVR-PA is not a good idea, because it increases exit times.

We infer the following from these experiments:

Hypothesis 4.1: Rejected

Hypothesis 4.2: Rejected

Hypothesis 4.3: Confirmed

3.7.5 Experiment 5: TVR with departure time classes (TVR-DTC)

In the past two experiments, any knowledge about departure times was ignored. In this experiment, we put it back into the equation, to combine the two ideas of using residence time knowledge and calculating the costs and reshuffle probabilities of possible locations. We use the departure time classes idea from section 3.7.2.

Departure time classes are pretty easy to use with the existing TVR algorithm. With the classes, we can more accurately estimate the number of expected reshuffles, which leads to a better calculation of costs for each pile, which, in turn, leads to better stacking decisions.

One major difference with the original TVR algorithm that we will have to make, is that it is now no longer necessarily optimal to stack on the closest positions available near the transfer points. This is because piles further away may have more favorable departure times and lead to less reshuffles. We thus have to check every position from the transfer point on further down the lane. We can only stop once we have found a pile where stacking would lead to no extra reshuffles, for every possible stacking level. This will make the algorithm much slower.

In addition, equation 3.2 is not valid for this algorithm, because we no longer assume that every pile has the same probability of being chosen. Rather, the probability of an extra reshuffle for every pile is determined using the following algorithm:

1. Set c_{min} as the earliest departure class in the current pile and c_{new} as the departure class of the container to be stacked.
2. if $c_{min} > c_{new}$, there is no reshuffle ($P(nr) = 0$).
3. else if $c_{min} < c_{new}$, there is definitely a new reshuffle ($P(nr) = 1$).
4. If the classes are equal, count the number of times the class occurs in the current pile as n. Then $P(nr) = \frac{n}{n+1}$.

We have added a variation of this TVR-DTC algorithm (TVR-DTC-MD) in which we also incorporate the minimization of the difference in departure time class between levels of the pile.

Experimental setup

We use the same setup as experiment 3 with the classes of departure times from experiment 2; see table 3.1.

Hypotheses

Hypothesis 5.1 The best TVR-DTC stacking algorithm will have lower exit times, reshuffles and ASC workloads than the best TVR algorithm.

Hypothesis 5.2 The best TVR-DTC stacking algorithm will have lower exit times, reshuffles and ASC workloads than the best LDT-DTC algorithm.

Hypothesis 5.3 The TVR-DTC-MD stacking algorithm will have lower exit times, reshuffles and ASC workloads than the TVR-DTC algorithm.

Since this algorithm combines "the best of both worlds", in this case of the two ideas we use to improve stacking efficiency (LDT and TVR), we expect it to perform better than the two ideas individually, when comparing the "best" penalties for both algorithms. For other penalties, this may not be the case, especially for negative penalties. This is because the residence time knowledge allows us to make better estimates of reshuffle probabilities, which will lead to a very high number of reshuffles, since these penalties favor reshuffles, rather than penalize them.

Results

The results in table 3.4 show that TVR-DTC outperforms normal TVR and LDT-DTC. The TVR-DTC-MD variation further improves the exit times; this combination of features provides the best performance for this stack configuration. However, we have also performed this experiment for the larger $8 \times 34 \times 6 \times 4$ stack (table 3.5) and in that case neither TVR-DTC nor TVR-DTC-MD improve upon LDT-DTC.

Discussion

Apparently, the combination of TVR and LDT-DTC is a good idea for a relatively full stack. The TVR-DTC-MD algorithm displays better performance still, which indicates that minimizing the difference in classes between levels of the piles is worthwhile. For a larger stack, the performance of LDT-DTC and TVR-DTC is very similar. From this, we conjecture that there is little room for further improvement of these algorithms and that the larger stack that was evaluated in these experiments does not highlight the differences. From experiment 5 we conclude:

Hypothesis 5.1: Confirmed

Hypothesis 5.2: Rejected

Hypothesis 5.3: Rejected

3.8 Conclusion

In this paper we have evaluated the performance of a number of online stacking strategies. We have used data from practice to generate scenarios of container movements for an automated container terminal. These scenarios were then processed by a simulation model in which the various stacking strategies were implemented. The main results are listed in tables 3.4 and 3.5.

The relatively simple, greedy strategies that were evaluated in this paper provide more insight into the basic trade-offs for stacking at an automated container terminal. The strategies operate in an online mode: the stacking location is selected on the basis of the current state of the stack and the parameters of the incoming container. We do not consider the stream of containers that will follow it.

Using detailed simulation experiments we have evaluated stacking rules from two perspectives. On the one hand we have investigated stacking rules that are based on departure time information. As a reference case, precise information on the actual time of departure was used. For a more realistic rule, we have used the actual departure times to formulate a number of classes and then used this class information to evaluate potential stacking locations. The aim is to create the piles in such a way that the higher containers have a lower class (i.e., will depart sooner) than the containers below them. The best performance is achieved if every container has a class that directly precedes the class of the container below it.

On the other hand we have looked into the trade-off between the travel time of the stacking crane versus the probability of reshuffles. We aim to stack an incoming container as fast as possible but we are willing to accept a longer use of the stacking crane if we can reduce the probability of future reshuffles.

We have found that rules with a limited number of classes for the remaining residence time work very well. The experiments in this paper show that, for the larger stack, the algorithms that use these classes perform similar to algorithms that use the exact departure times. From this we conclude that even imprecise information on this departure time is very valuable.

The performance of the travel time versus reshuffling stacking rules shows a clear advantage with respect to the reference stacking rules. We have formulated an extension of this stacking rule to attempt to reduce the exit time during future peak workload periods but we found no significant improvement. A combination of the travel

Table 3.4: Results of experiments 1–5. Terminal layout: 6 × 34 × 6 × 3. The "90%" values are the average 90% percentile values. All normal 20 ft. containers were included in the simulation. "p" is used to indicate the values for the reshuffle penalty TT_r.

Exp.	Description	ROC %	RDC %	GPU %	ASC %	ETQ hrs.	ETL hrs.	90% ETQ	90% ETL
0	LEV	62.92	81.10	99.60	59.44	0.48	0.33	1.30	0.76
0	RS	69.44	107.69	89.56	62.70	0.73	0.47	2.09	1.23
1	RS-DT (real)	35.27	52.10	82.94	54.44	0.37	0.26	0.94	0.56
1	RS-DT (exp)	38.11	55.53	82.81	54.80	0.40	0.28	1.04	0.61
1	LDT (unmixed, real)	10.20	15.21	87.39	49.93	0.21	0.17	0.42	0.30
1	LDT (unmixed, exp)	17.49	23.19	87.12	50.61	0.24	0.19	0.51	0.35
1	LDT (mixed, real)	9.09	13.62	86.64	49.62	0.22	0.18	0.44	0.32
1	LDT (mixed, exp)	17.12	22.52	86.37	50.52	0.25	0.20	0.54	0.38
2	RS-DTC (real)	68.28	101.06	87.54	61.26	0.73	0.47	2.20	1.21
2	RS-DTC (exp)	67.59	99.87	87.44	61.17	0.72	0.46	2.18	1.18
2	LDT-DTC (unmixed, real)	39.83	62.13	96.72	56.53	0.53	0.34	1.59	0.73
2	LDT-DTC (unmixed, exp)	38.15	58.97	96.45	56.24	0.50	0.32	1.46	0.69
3	TPRL	62.82	95.35	96.41	57.78	0.43	0.29	1.13	0.66
3	TVR (p=-0.03)	80.01	131.68	75.85	57.90	0.66	0.42	1.87	1.09
3	TVR (p=-0.01)	64.85	102.93	96.68	55.87	0.43	0.29	1.11	0.66
3	TVR (p=0)	62.24	94.95	99.35	57.01	0.40	0.28	1.02	0.62
3	TVR (p=0.03)	63.71	86.50	99.65	57.57	0.40	0.28	1.00	0.61
3	TVR (p=0.04)	63.52	84.76	99.65	57.61	0.40	0.27	1.00	0.60
4	TVR-PA (unmixed, p=-0.03)	77.73	126.89	76.00	56.86	0.56	0.38	1.54	0.92
4	TVR-PA (unmixed, p=0)	63.41	95.58	98.70	57.11	0.41	0.28	1.07	0.62
4	TVR-PA (unmixed, p=0.03)	64.80	90.13	99.48	57.43	0.40	0.28	1.03	0.62
4	TVR-PA (mixed, p=-0.03)	77.34	126.12	75.94	56.97	0.57	0.38	1.57	0.94
4	TVR-PA (mixed, p=0.0)	63.23	95.49	98.66	57.13	0.41	0.28	1.07	0.63
4	TVR-PA (mixed, p=0.03)	63.52	90.60	99.50	57.69	0.42	0.28	1.08	0.63
5	TVR-DTC (p=0)	32.62	47.44	95.20	52.28	0.28	0.21	0.65	0.39
5	TVR-DTC (p=0.01)	30.19	43.26	97.43	52.37	0.28	0.21	0.67	0.39
5	TVR-DTC (p=0.03)	30.82	43.68	97.91	52.63	0.30	0.22	0.74	0.41
5	TVR-DTC-MD (p=0.01)	22.78	29.88	91.60	50.65	0.19	0.17	0.38	0.30

Table 3.5: Results of the experiment for terminal layout: $8 \times 34 \times 6 \times 4$. The "90%" values are the average 90% percentile values. All normal 20 ft. containers were included in the simulation.

Exp	Description	ROC %	RDC %	GPU %	ASC %	ETQ hrs.	ETL hrs.	90% ETQ	90% ETL
0	RS	65.77	118.35	76.14	54.38	0.41	0.30	0.99	0.61
0	LEV	52.63	52.64	99.54	48.53	0.21	0.19	0.38	0.31
1	RS-DT (real)	14.16	27.44	53.92	43.15	0.17	0.16	0.31	0.26
1	RS-DT (exp)	19.43	33.69	53.79	43.67	0.19	0.17	0.34	0.28
1	LDT (unmixed, real)	0.04	0.07	54.88	39.68	0.13	0.12	0.20	0.19
1	LDT (unmixed, exp)	10.55	11.97	54.62	40.42	0.14	0.13	0.22	0.20
2	RS-DTC (real)	51.82	97.88	68.32	51.29	0.35	0.26	0.82	0.50
2	RS-DTC (exp)	51.02	96.70	68.03	51.19	0.34	0.25	0.77	0.48
2	LDT-DTC (real)	1.94	3.97	75.64	42.06	0.13	0.13	0.21	0.20
2	LDT-DTC (exp)	3.64	5.49	75.54	42.16	0.13	0.13	0.21	0.20
3	TVR (p=0)	47.68	96.34	85.22	45.85	0.21	0.18	0.40	0.30
5	TVR-DTC (p=0)	11.74	20.35	74.96	41.51	0.14	0.13	0.22	0.20
5	TVR-DTC (p=0.01)	1.16	1.25	81.61	41.22	0.13	0.13	0.20	0.19
5	TVR-DTC (p=0.03)	0.02	0.02	82.22	41.24	0.13	0.13	0.20	0.19
5	TVR-DTC-MD (p=0.01)	8.21	10.49	67.59	41.09	0.13	0.13	0.20	0.20

times versus reshuffling stacking rule with the departure time classes was also evaluated. Experiments with this rule confirm that it is beneficial to create piles of high quality, i.e. piles where the difference in departure time class between levels is exactly one.

We have tested the strategies using a small number of variations of the basic layout (in terms of the number of lanes, the lane width and length, and the maximum stacking height). The overall layout of the yard has been the same for all experiments and it would be interesting to evaluate these simple rules for other basic layouts such as lanes that are parallel rather than perpendicular to the quay or with multiple ASC's per lane. The experiments in this paper have been limited to the standard twenty foot container; further research is needed to investigate the performance of the online rules for heterogeneous container types (forty foot, reefers).

Chapter 4

Floating Stocks: Using intermodal transport to facilitate advance deployment*

In this chapter we present a distribution concept called "floating stocks", which uses intermodal transport to deploy inventories in a supply chain in advance of retailer demand. A significant drawback of intermodal transport is the longer transit time. In the floating stock strategy, we mitigate this disadvantage by supplying part of the demand by road. First an analytical comparison is made which shows that this concept has advantages in inventories over pure road and intermodal transport. Next a simulation study of a real case is made which quantifies the cost-differences in detail.

4.1 Introduction

Intermodal transport can be defined (cf. ECMT, 1993) as the movement of goods in one and the same load unit or vehicle by successive modes of transport without handling the goods themselves during transfers between the modes, e.g. container transport via rail and road. It has been advocated for a long-time by the European Union and national governments because of its environmental friendliness and reduction of traffic congestion. It is often used when a modality change has to be made anyhow, e.g. in case of importing goods from China to Germany and using rail transport from the sea port to an inland terminal. Within a continent, intermodal transport is competitive if large amounts have to be transported over long distances. In today's supply chains

* A version of this chapter has been published in the International Journal of Physical Distribution and Logistics Management (Dekker et al., 2009)

where responsiveness is important transportation has to be done often within a short time frame. Hence road transport is dominant transport mode, although load factors are low, because it is difficult to find return trips. It is therefore important if intermodal transport could be made more competitive on intra continental flows. A major problem in this respect, is that rail transport takes a long time, especially if rail cars have to be shunted. Therefore it does not seem to be suited for fast-moving consumer goods, where just-in-time (JIT) principles demand responsive supply chains and fast transport is applied.

In this chapter we investigate a new distribution concept, called floating stocks, to overcome some of these problems. It is intended for fast-moving consumer goods, where products are produced in large quantities in a batch mode. If there is not a large product variety, or if a standard product mix can be defined, then products can be shipped in containers in the direction of customers, before they have really called the demand off. Next we use the intermodal rail terminals as temporary storage places and await the final order call-off at these locations. This looks like the concept of virtual warehouses, yet it is different from what is normally understood by that concept in literature (see e.g. Landers et al., 2000 as they stress real-time global visibility of assets).

The idea has been applied in some cases, although not with one clear name. In the North American lumber industry, western lumber producers would ship loads to north central and eastern customers before demand had finalized (Sampson et al., 1985). The flatcars or boxcars were held at transit yards in the Midwest until a customer order was received. This practice enabled western producers to compete in the eastern market with their southern competitors in terms of lead times. The floating stock idea has also been mentioned in publications on the Dutch Distrivaart project (Boerema et al., 2003 and Teulings and van der Vlist, 2001), neither of which does a detailed supply chain analysis. The floating stocks concept draws on the areas of transportation and inventory control: we will therefore briefly recapitulate relevant transportation and inventory control literature.

Road transport and intermodal transport are compared in many papers, e.g. Bookbinder and Fox (1998), Rutten (1995), and Konings (1996); they typically focus on a comparison of transportation characteristics, such as costs, transit time, reliability, etc. Intermodal transport is said to be slower and have a lower reliability. It is competitive for longer distances (more than 750 km, according to Konings (1996)). In this chapter we do not consider transport on its own, but take a supply chain perspective, as advocated by Chopra and Meindl (2004). Supply chain responsiveness can be obtained by fast transport and by keeping inventories close to customers. Normally, inventories are stored in distribution centres (DC's) which have fixed cost elements and increase safety stocks. In this chapter we will therefore compare a number of distribution concepts in terms of transportation options and inventories location. This has been done in few papers only. Evers (1996, 1997, 1999) has studied risk pooling of demand and lead times in relation to transshipments, but he does not consider transport costs explicitly. Other papers on inventory control typically consider only two transportation options: a regular one and a faster one for emergency shipments (cf. Moinzadeh and Nahmias, 1988). Some studies also consider lateral transshipment in multi-echelon chains, but mostly again only in the case of stockouts (cf. Minner, 2003 and Diks et al., 1996). Other streams investigate the relation between transport frequency and inventory control (cf. Tyworth and Zeng (1998)), but they focus on the relation between either transport frequency or transit time reliability and inventory control. No studies seem to exist which integrate intermodal transport and inventory control, according to reviews on intermodal research such as Bontekoning et al. (2004) and Macharis and Bontekoning (2004).

The structure of this chapter is as follows. In section 4.2 we state our methodology, while in section 4.3 we give a conceptual model. We demonstrate the concept in a case with real cost data provided by Vos Logistics. This case is presented in section 4.4 and results in section 4.5. To ensure generality we have applied a sensitivity analysis in section 4.6 and give a cost analysis in section 4.7. In a follow-up paper, an optimisation model for timing shipments in the floating-stock concept was presented (Pourakbar et al., 2009).

4.2 Methodology

We use a conceptual model to allow a qualitative comparison between four distribution concepts that differ in the use of intermodal transport and inventory deployment. To avoid many complicating and potentially conflicting aspects, we confine ourselves to a part of a Fast Moving Consumer Goods (FMCG) (such as cosmetics and batteries) supply chain: from the manufacturer to the retailer's distribution center (DC). Moreover, we aggregate all products to one standard mix. For this case we also make an analytical comparison. Next we numerically evaluate our concepts in a case study in Europe taking data from Vos Logistics, a logistic service provider. We use simulation as the main method and check its outcomes with the analytical calculations. The advantage of this approach is that we can get an estimate of the real savings, yet the disadvantage is that the calculations are only done for one specific case. To get some idea of generality we also perform a sensitivity analysis.

4.3 Conceptual Model

The conceptual model consists of a general network representation of the distribution process, together with assumptions. First, we explain the assumptions and the construction of the network model with the possible choices in this model. We then formulate four distribution strategies based on the general choices on the position of inventories in the chain. Next we define the performance criteria and evaluate the different strategies.

4.3.1 Model

We consider fast moving consumer products that are made in batches. A production cycle starts with the production of a new batch and ends when the next batch is produced. We assume this production cycle length to be fixed. The size of a production batch is based on the remaining number of products from the last production cycle and a demand forecast for the new cycle. The demand forecast relies on information provided by the retailers. The production time is neglected.

The distribution process starts right after the production of a new batch. The output of a production batch can be stored in a storage location near the factory (which we call the factory storage) or can be transported to a regional stocking point (or a terminal used as such). All costs caused by these products from this moment are taken into account, whether they are for the manufacturer or retailer in reality. In our model the distribution process ends when a product arrives at the retailer's distribution center. (We will refer to this distribution center as "DC" in the remainder of this chapter.) In the supply chain between the factory and the DC, there can be one or two transshipment or stocking points. These points are used if the transportation is intermodal or if the storage is decentralized. In this chapter we will refer to these points as terminals, but they could be regional distribution centers as well.

In order to make a good comparison between the distribution strategies, we assume that all orders and deliveries consist of full-truck-loads (FTL's). If, for example, a retailer is supplied using 40 ft. containers, then the order size of its DC must be exactly the number of products that fills a 40 ft. container (or a multiple of this number). The demand for fast moving consumer goods is high enough to make it possible to transport only FTL's of a single product. If the demand is too low for a single product, a standard product mix can be used to create FTL-transports (Teulings and van der Vlist, 2001). The composition of this mix has to be fixed, because the products must remain in the load unit during the distribution process. In our model, every transport is a direct run from departure to destination. Vehicle route planning is not taken into account and we assume that a transportation vehicle is always available when needed.

All DC's can be reached both by a direct (road a.k.a. unimodal) connection and an indirect (intermodal) connection with one or two transshipment points (regional terminals). In these regional terminals the products can be stored for a short period (this free time is typically a couple of days). When a new production batch is ready, the manufacturer has to choose where to store the products. The products can either be stored on-site in the factory storage or transported to a regional terminal immediately after the production. For each order, the manufacturer has to choose from which stocking point the order will be fulfilled and which transportation mode will be used. We assume that the transit time of a direct transport from both the factory and the regional terminals is short enough to be acceptable for the retailer as order lead time.

4.3.2 Distribution strategies

In this chapter we examine four distribution strategies. For every full truck load unit, we have to decide whether it will be stored in a centralized or a decentralized location, and whether to use road or intermodal transport.

The first strategy is based on the just-in-time concept and applies direct road transport only. This is frequently used in FMCG-supply chains. The second strategy is completely based on floating stock: all transports are intermodal. This strategy is especially popular in supply chains where an intermodal connection has lower transport costs than a road connection. The third and fourth strategies aim to take as much advantage of floating stock as possible. Figure 4.1 illustrates the four strategies. In this chapter we limit ourselves to using readily available storage locations, either at the factory or at the intermodal terminals. A strategy with delivery by truck to a regional terminal is not considered as this approach would require either a facility with container-handling equipment or storage on trailer.

Strategy CS: Centralized storage and unimodal transport

Using this just-in-time based strategy means that the whole production batch and the safety stock are stored on-site at the factory storage. When an order arrives, it is always fulfilled using road transport from the on-site inventory. In this strategy the emphasis is on fast transportation and easy coordination.

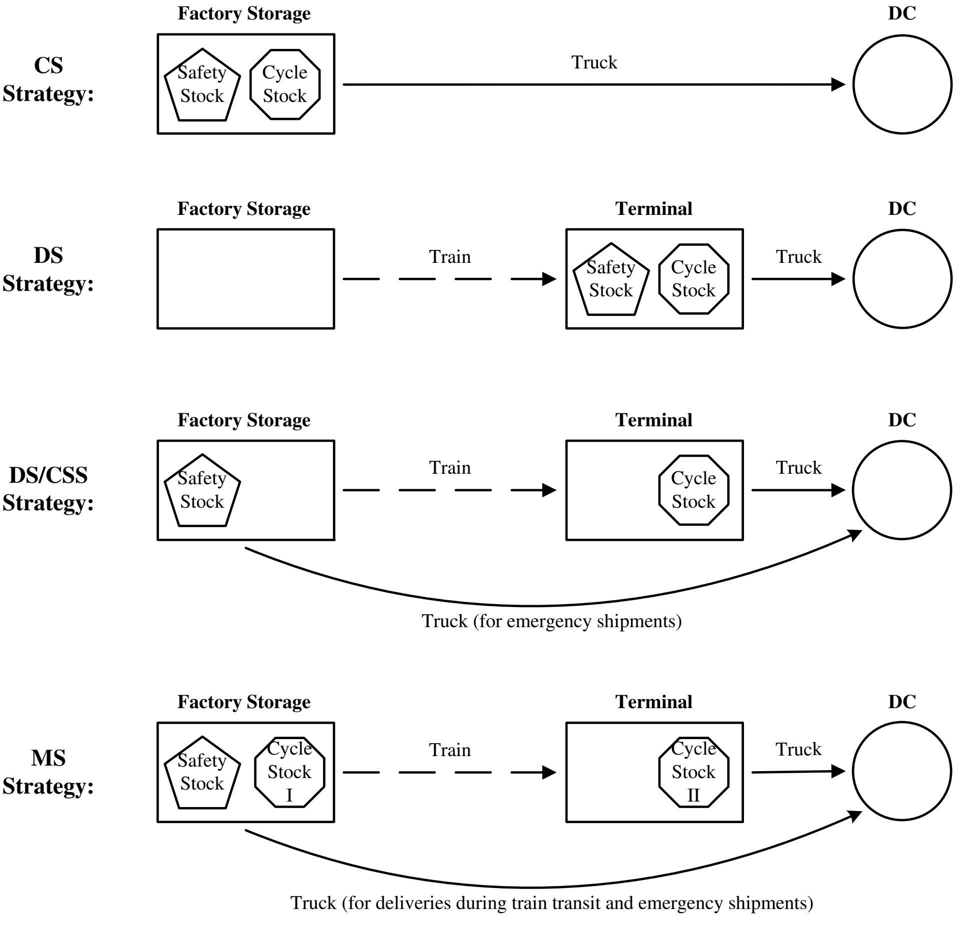

Figure 4.1: The four strategies and the stocking points

Strategy DS: Decentralized storage and intermodal transport

The complete production batch is shipped to regional terminals using intermodal transport. A demand prediction is used to determine the split of the production batch over the regional terminals. Orders are delivered by truck from these terminals to the DCs. The safety stock is also stored in these regional terminals. The emphasis is on using intermodal transportation and short order lead times (because the order lead time from the terminal will be shorter than from the factory). If the safety stocks are depleted at a terminal, lateral transshipments from other terminals are made.

Strategy DS/CSS: Decentralized storage, intermodal transport, and centralized safety stock

In this case the safety stock is stored at the factory storage, whereas the production batch is shipped to the terminals using intermodal transport and stored there. The regular deliveries to the retailers are fulfilled from the terminals. In a period of excess demand, we first consider lateral transshipments from other terminals: if all terminals are without stock, emergency deliveries are done from the factory storage. These emergency deliveries are transported by road, because the intermodal transit time is much longer.

The safety stock storage costs will probably be lower in the DS/CSS strategy when compared to the DS strategy. This is because long storage on-site is in general cheaper than long storage in a terminal. Furthermore, reliability increases if the safety stock is stored in a central location.

Strategy MS: Mixed storage

The mixed storage strategy stores part of the production batch in the factory storage (centralized) and part of the production batch is stored in decentralized terminals. The latter part is shipped using intermodal transport. All orders that are placed while the intermodal transport is in transit, are fulfilled from the on-site inventory at the factory using road transport. Once the products have arrived at the terminal, the orders are delivered from the terminal (with a shorter order lead time). Emergency orders in a period of excess demand are delivered using road transport from the safety stock stored at the factory. If the safety stock at the factory is depleted, lateral transshipments from other regional terminals are considered.

This strategy is designed to benefit from cost advantages of floating stock storage without having to increase the total inventory level in the supply chain. The DS strategy ships the complete production batch using intermodal transport. This batch cannot be used to fulfill orders until it has arrived at the regional terminal*. Any orders coming in during this transit time can only be fulfilled using products from a previous production cycle. This increases storage time and costs. If we split the batch into a part that is stored in the central factory storage and a part that will be stored in decentralized locations, then it is possible to benefit from the cost advantages of floating stock storage without suffering additional inventory costs. Orders received during the transit time of the intermodal transport can now be fulfilled using the on-site inventory from the current production cycle. In this way, the total stored inventory is low during the intermodal transport transit time and the order fill rate is high. Centralized storage of the safety stock and the inventory required for the expected orders during the intermodal transit time maximizes the savings. If more products were to be stored on-site, then the floating stock part (which generates the storage costs savings) would decrease. A lower level of centralized inventory will either lower the order fill rate or increase the storage costs (for products stored centrally from previous production cycles).

4.3.3 Performance criteria

The following criteria are relevant for evaluation of the strategies: expected costs, average order lead time, and order fill rate.

The expected costs are divided into transportation and handling costs, storage costs and holding costs. Transportation and handling costs differ per transportation route. They contain all costs that result from using the specific transportation route: these costs depend on the number of transported load units (FTL's in our model). Therefore, transportation costs can cause differences in the total costs of each strategy, but these are independent of the inventory levels during a production cycle. The storage costs are the direct costs for storing a certain number of products for a certain period. These costs depend on the storage tariff at the specific point, the storage time, and volume of the products (or load units) stored. The holding costs are the indirect costs for keeping inventory in the supply chain. Examples of holding costs are cost of capital and obsolescence cost. Storage costs are usually considered part of the total holding costs, but in this chapter we list them separately to support our analysis.

The average order lead time for a DC (customer) is the average time between placement of an order by a customer and the supply moment of this DC. If intermodal transport is combined with decentralized storage, the order lead time is shorter than

* this is a limitation of our current simulation model

a strategy with centralized storage and road transport, although intermodal is slower than road transport in general. Figure 4.2 shows an example of this with an order lead time of two days using centralized storage and one day using decentralized storage.

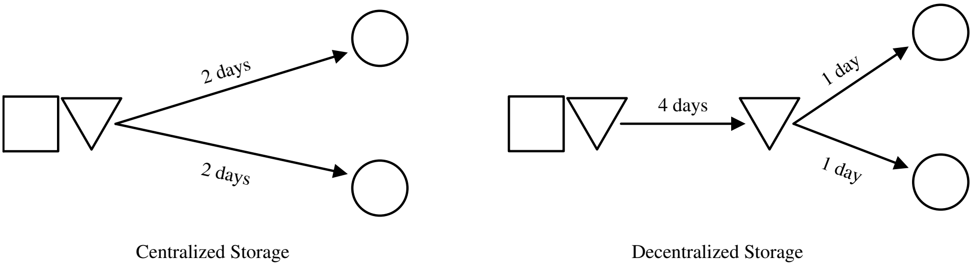

Figure 4.2: Centralized storage leads to longer order lead times

Orders can only be supplied from the static (non-moving) inventory, so inventory in transit (pipeline inventory) is not considered when an order arrives. If the available static stock is too low to fulfill the order, the order is rejected. There is no back-ordering. The order fill rate is the percentage of the orders that can be fulfilled. If a strategy's order fill rate is less than the required fill rate, the safety stock must be increased. This causes additional holding and storage costs, so the increase should be the smallest possible increase that will lead to the required order fill rate.

4.3.4 Example

Consider a simple supply chain that consists of a factory site, a regional stocking point close to the customer, and a single customer. Intermodal transport from the factory to the regional stocking point takes four days; truck transport from the regional stocking point to the customer takes a single day. In a situation with a constant demand of one FTL per day and a cycle length of fourteen days, the DS and MS strategies would yield the results shown in table 4.1. The intermodal strategies reduce static inventories in favor of pipeline inventory; these strategies use the slower intermodal transport to reduce storage costs and can further reduce these storage costs by exploiting the free time offered by intermodal terminals. These strategies also benefit from the reduced lead time when compared to the CS strategy; the reduction in average order lead time causes the inventory to leave the supply chain sooner and this reduces holding hosts.

Table 4.1: Results for simple supply chain

	Unit	CS	DS	MS
Demand per cycle	FTL	14	14	14
Inventory	FTL	8.5	11.5	7.8
Pipeline inventory	FTL	2.0	5.0	4.1
Static inventory	FTL	6.5	6.5	3.6
static inventory at factory	FTL	6.5	0.0	0.4
static inventory at regional terminal	FTL	0.0	6.5	3.2
Delivered from factory	FTL/cycle	14	0	4
Delivered from regional terminal	FTL/cycle	0	14	10
Order fill rate	%	100	100	100
Average order lead time	days	2	1	1.3

4.3.5 Impact of distribution strategy on inventory

In this section we compare the four distribution strategies on their average storage levels. This gives insight into the storage and holding costs per strategy.

Consider the supply chain from factory to a single retailer's DC. The demand of the DC is assumed constant at rate r (shipping units per day). The production cycle has length t_c, so on day t_c, $2t_c$, $3t_c$ etc. a new batch is produced of size $Q = t_c r$. Furthermore, the manufacturer uses a safety stock of size s. The intermodal transport from the factory to the terminal has transit time $t_i < t_c$.

Using the CS-strategy, the manufacturer has $t_c r + s$ in storage at the start of a production cycle, because in this strategy the whole new batch is stored at the factory storage immediately after the production. During the production cycle, this decreases linearly to the safety stock level s at the end of the production cycle. A new batch is then produced and the process is repeated. The average storage level for static inventory is $\frac{t_c r}{2} + s$. Figure 4.3 shows the inventory profile of this process.

Using the DS strategy, the new production batch is shipped to the terminal at the start of a production cycle using intermodal transport with transit time t_i. Therefore, a storage level of $t_i r$ from the previous production cycle is necessary at the terminal to be able to deliver the orders during t_i. The safety stock is stored at the terminal as well. The inventory level $t_c r + s$ is reached at time t_i. The inventory profile (see figure 4.3) is identical to the profile of the CS strategy with a delay of t_i days. Thus the average storage level of the DS-strategy is $\frac{t_c r}{2} + s$.

The DS/CSS-strategy differs only from the DS strategy in the location of the safety stock. This location makes no difference for the total average storage level, so the storage levels of these three strategies are all equal if they use the same level of safety stock. However, as the amount of pooling is different for the strategies, the safety stock level could differ.

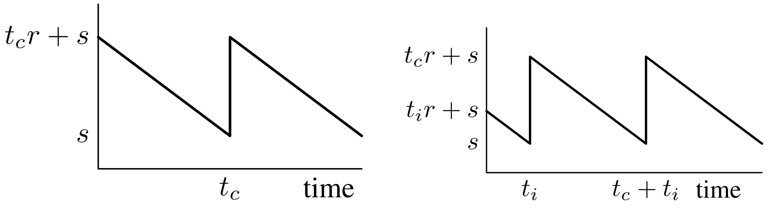

Figure 4.3: Storage at factory for CS-strategy (left) and at terminal for DS-strategy (right)

Using the MS-strategy, the new production batch is split into two parts. The first part is required to deliver the orders in the first t_i days of the production cycle: this part and the safety stock are stored at the factory. In total this amounts to $t_i r + s$. The second part is used to deliver the orders in the last $t_c - t_i$ days of the production cycle: $(t_c - t_i)r$ units are transported to the terminal using intermodal transport. In this strategy, the average storage level at the factory is $\frac{t_i}{t_c} t_i \frac{r}{2} + s$. The average storage level at the terminal is $\frac{t_c - t_i}{t_c} \frac{(t_c - t_i)r}{2}$. The total average storage level is the sum of the average storage level at the factory and the average storage level at the terminal: $(\frac{2t_i^2}{t_c} - 2t_i + t_c)r/2 + s$. So by this advanced positioning the MS strategy has a lower average storage level than the other three if $2t_i^2/t_c - 2t_i < 0$ and because $t_i < t_c$, this is always true. This storage level difference is optimal in the case that $t_i = t_c/2$. Note that delivering $t_c r$ directly is optimal, as a higher or lower amount does not reduce inventories.

The storage levels at the factory storage and the terminal in this strategy are shown in figure 4.4.

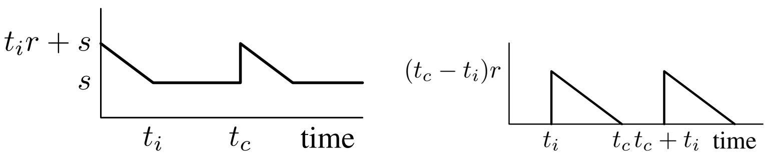

Figure 4.4: Storage at factory (left) and terminal (right) for the MS-strategy

The average pipeline inventory level, i.e. the average number of shipping units in transit, depends on the transportation mode. A strategy has a higher average pipeline inventory level if more intermodal transports are used because of the longer transit time of intermodal transport. Therefore, the DS and DS/CSS strategies always have a higher average pipeline inventory level than the other two. More pipeline inventory does not lead to higher storage costs, but it does lead to higher holding costs so this effect should be taken into account when the strategies are compared.

The amount of safety stock needed to reach a certain order fill rate can also differ between the strategies. The CS, MS and DS/CSS all apply centralized safety stocks, which can therefore be lower than the total decentralized safety stock for the DS strategy. Moreover, the CS and, to a lesser extent, the MS strategy can also benefit from the safety aspect of a pooled cycle stock which may also lead to a lower safety stock (if demand at one location is low, cycle stock can be used for another location). In the MS strategy the cycle stock is partially pooled during the free time at the terminals; if actual demand shifts within an area serviced by a terminal and if the total demand can be met from this stock, we can adjust the cycle stock split accordingly. This effect and the degree of pooling clearly depend on the layout of the actual distribution network. In our model the safety stock is held in an integer number of full truck units; this, a small effect may not be detected.

In table 4.2 we summarize the performance differences between the various distribution strategies (IM indicates intermodal transport.)

Table 4.2: Centralized storage leads to longer order lead times

	Strategy			
Aspect	CS	DS	DS/CSS	MS
Transportation	Road	IM	mainly IM	Road & IM
Centralized safety stock	Yes	No	Yes	Yes
Advanced deployment	No	Yes	Yes	Yes
Pooling effect of cycle stock	Yes	No	No	Partial
Pipeline stocks	Low	High	High	Moderate
Average order lead time	Long	Short	Short	Varying

To test how large these differences are and whether the storage advantage has any negative effect on the order fill rate of the MS-strategy, we performed a case simulation.

4.4 Case Description

Below we present a real case and match it to the conceptual model that was presented in section 4.3.5. The case uses realistic data from logistic service provider Vos Logistics.

In the next section we will describe a simulation model that was developed for the case study.

An FMCG-manufacturer runs a factory in Poznan (Poland) and distributes its products to four retail DCs (customers) in Germany, viz. in Dortmund, Köln, Rüsselsheim (near Frankfurt), and Appenweier (near Strasbourg). At this moment all orders are transported FTL by truck. The load unit is 40 ft. container. An alternative intermodal route is a rail connection from a station in Gadki (15 km from Poznan) to two train terminals in Duisburg and Mannheim; trucks are used for the transport from the train terminals to the customer locations (DCs). The customers keep a limited supply of the goods in their DCs and rely on replenishment orders with a short lead time (two days). The conceptual network representation for this case is displayed in figure 4.5.

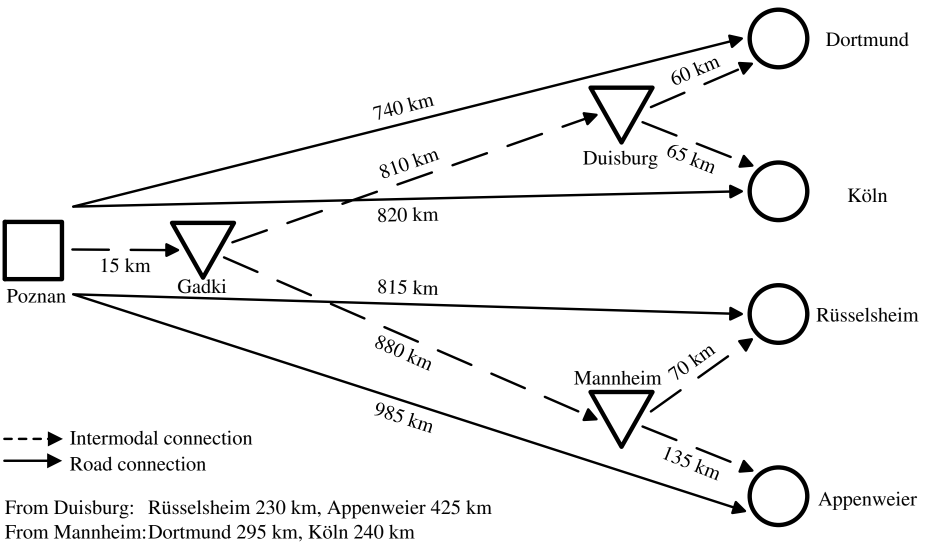

Figure 4.5: Network representation of the case

The transit time for all four direct truck routes is two days including handling time for in- and outbound in the on-site DC. The intermodal connection makes use of the rail connection. Due to the long time needed for shunting, the transit time of the train transport to both terminals is 2.5 days. The total transit time of the intermodal transport, including handling and waiting times, is five days (the individual steps are shown in table 4.3). If a stock-out occurs at the regional terminal, the DC is supplied by the other terminal: in this case, the final truck transport takes a full day and the total transit time will be 5.5 days.

Table 4.3: Steps in Intermodal Transport

Step	Duration
Transport Poznan – Gadki and inbound Gadki	0.25 days
Expected waiting time Gadki	0.75 days
Loading time train	0.25 days
Transit time rail transport Gadki - Duisburg/Mannheim	2.50 days
Inbound regional terminal	0.25 days
Outbound regional terminal	0.25 days
Transit time final truck transport	0.50 days
Inbound retailer's DC	0.25 days

The cost components which are used to estimate the costs are linear per FTL container delivery and are detailed in table 4.4 (these numbers were established in consultation with Vos Logistics). The intermodal transportation costs are higher than the road transportation costs because the distance is relatively short for a train connection and the overhead of getting the container on the train is relatively high.

4.5 Experiments

In this section we introduce the simulation program. Next we present the results of simulation of the case and explain them. Additionally a sensitivity analysis is performed to investigate the influence of a number of factors on the results. Finally, a cost analysis is done and the results of the experiments will be discussed.

4.5.1 The simulation program

The simulation program has been implemented in Arena (Kelton et al., 2004). The core of the simulation program consists of three processes: the (stochastic) order generation process, the production process, and the distribution process.

Orders are generated by a random number generator using a probability distribution. The generated number is the interarrival time between two orders from one DC. Every DC uses its own random number generator so every time a retailer places an order, the time until the next order of that DC is drawn. This makes an order by a DC independent of the other DC's orders and of the orders from that DC in the past. The probability distribution used for the case is a negative-exponential distribution with a

Table 4.4: Cost Components

Component	Costs
Transport and handling:	
For the direct road connection from factory to DC	€ 880 per container
For the intermodal connection from factory to DC	€ 900 per container
Extra costs for transport from terminal outside region of DC	€ 100 per container
Storage:	
Centralized at factory storage	€ 8 for volume of one container per day
Decentralized in terminal	€ 16 per container per day
No charge period at terminals (free time)	3 days
Holding:	
15% interest over € 41,370 (value of products in 40 ft. container FTL)	€ 17 per container per day

mean of 2 days.

The production process takes place every time a new production cycle starts. The production batch size depends on the demand forecast in the new cycle and the remaining inventory from the last period. The exact algorithms used to determine the batch size differ per distribution strategy, but they always target to keep the cycle stock equal to the expected demand in one production cycle, taking into account the average order volume per day. (The algorithms are described in (Ochtman et al., 2004).) Production time is neglected.

The distribution process models the distribution of the new production batch to the appropriate stocking points and the selection of the stocking point for order fulfillment. In our simulation model, only stored inventory can be used to deliver orders; pipeline inventory cannot be used for this. The CS strategy delivers all orders from the factory storage. The DS strategy generally delivers an order from the terminal in the same region as the DC that placed the order. If this terminal does not have sufficient inventory, the order is delivered from the terminal in the other region (which causes higher transportation costs and a longer transit time). The DS/CSS and MS strategies use the same sequence as the DS strategy, but now the safety stock at the factory storage might deliver the order if both terminals are out of inventory. If the stocking points do not have enough inventory when an order arrives, the order is rejected. At the end of the simulation the order fill rate of the distribution strategy is determined by dividing the total number of supplied orders (= total orders – rejected orders) by the total number of orders. If this order fill rate is less than the required level, then the simulation is restarted with a higher safety stock level, using a step of one full truck load.

4.5.2 Simulation Results

Table 4.5 lists the parameters used for the case simulation.

Table 4.5: Parameters for the case

Parameter	Value
Transit time intermodal transport from factory to terminal (t_i)	4 days
Production cycle length (t_c)	14 days
Variation coefficient of the order interarrival times*	0.5
Demand forecast per DC per production cycle ($t_c r$)	7 FTL's
Minimum order fill rate	99 %
Train departure frequency	daily
Demand ratio region 1 vs. region 2	50-50

* Variation coefficient is the standard deviation divided by the mean.

A simulation run consists of 35 independent replications. Every replication consists of a warm-up cycle and 500 production cycles (7,000 days). During the warm-up period, safety stock is produced and stored at the appropriate locations. The simulation output are the total average amount of inventory and transportation per production cycle. These are the averages of 35 replications. The 95%-confidence intervals of these averages are very small so the results are very reliable. Table 4.6 shows the results of the simulation program for the four distribution strategies.

Inventory

The average inventory levels of the four strategies are quite different. The DS and DS/CSS strategy need a lot more inventory than the other two strategies. This is because these strategies lead to a high pipeline inventory due to the use of the (slow) intermodal transport, whereas no savings on storage are obtained. Compared to the CS strategy, the MS strategy has a high pipeline inventory as well, but in this strategy the average storage level is low as explained in the conceptual model.

Table 4.6: The results of the case simulation

	Unit	CS	DS	DS/CSS	MS
Demand per cycle	FTL	28.0	28.0	28.0	28.0
Inventory	FTL	25.0	32.1	29.7	23.7
Pipeline inventory	FTL	4.0	9.9	9.6	8.3
Static inventory	FTL	21.0	22.1	20.1	15.4
static inventory in Poznan	FTL	21.0	0.0	5.5	7.9
static inventory in Duisburg	FTL	0.0	11.1	7.3	3.7
static inventory in Mannheim	FTL	0.0	11.1	7.3	3.7
Delivered from factory	FTL/cycle	27.7	0.0	1.9	7.9
Delivered from regional terminal	FTL/cycle	0.0	27.1	24.5	18.1
Delivered from other terminal	FTL/cycle	0.0	0.6	1.4	1.7
Rejected orders	FTL/cycle	0.2	0.3	0.2	0.2
Required safety stock	FTL	7	8	7	7
Order fill rate	%	99.1	99.1	99.3	99.1
Average order lead time	days	2.0	1.0	1.1	1.3

In the analysis in section 4.3.5, the average storage formula derived for the first three strategies was $t_c r/2 + s$. For the case this is equal to $\frac{14*2}{2} + 7 = 21$ FTL's for the CS and DS/CSS strategies and 22 FTL's for the DS strategy (because of the higher safety stock required). The average storage level for the MS strategy is $(2{t_i}^2/t_c - 2t_i + t_c)r/2 + s$. In this case, this is 15.3 FTL's. The simulation results agree with this with a little aberration because of the stochastic order process.

This analysis shows that the DS and DS/CSS strategy are inefficient. Although the total storage of these strategies is equal to the total storage of the CS strategy, they need much more pipeline inventory. This will cause more holding costs. On the contrary, the MS strategy has a slightly lower total inventory level than the CS strategy. Moreover this strategy makes efficient use of the floating-stock advantages, which leads to less storage and more pipeline inventory. In this way, the MS strategy could save on storage costs.

Other simulation results

The other simulation results are simply explained by the definition of the four strategies. The number of FTL's delivered from the terminal outside the region appears to be very small for every strategy. The extra transport costs, caused by this inefficient way of delivering is therefore marginal. The order lead time depends on whether the orders are delivered only from the factory (in two days for CS), mostly from the terminals (in a single day for DS and DS/CSS) or both (MS).

4.6 Sensitivity Analysis

In the sensitivity analysis all seven parameters listed in table 4.5 were varied individually to measure their influence on the simulation results of the four strategies. (The details of this analysis are described in Ochtman (2003).) Only the intermodal transit time caused the differences between the strategies to change significantly. The simulation results with varying transit times are shown in table 4.7 (the unit of measurement is FTL 40 ft. container).

These results show that an intermodal transit time from the factory to the terminal that is closer to half of the production cycle length makes the MS advantage in storage bigger with respect to the other three as described by the analysis in the conceptual model section.

Table 4.7: Results Sensitivity Analysis with Varying Intermodal Transit Time

	CS					DS				
	2	4	6	7	10	2	4	6	7	10
Avg. inventory	31.9	31.9	31.9	31.9	31.9	35.0	39.0	43.0	46.0	53.0
Avg. pipeline inventory	4.0	4.0	4.0	4.0	4.0	6.0	9.9	13.9	15.9	21.9
Avg. storage	28.0	28.0	28.0	28.0	28.0	29.0	29.1	29.1	30.2	31.1
Safety Stock	14.0	14.0	14.0	14.0	14.0	15.0	15.0	15.0	16.0	17.0
Order Fill Rate (%)	99.2	99.2	99.2	99.2	99.2	99.3	99.1	99.0	99.0	99.4

	DS/CSS					MS				
	2	4	6	7	10	2	4	6	7	10
Avg. inventory	25.8	29.6	33.7	35.8	42.0	23.5	23.7	23.9	24.0	24.4
Avg. pipeline inventory	5.9	9.6	13.3	15.2	20.7	5.7	8.3	9.7	10.0	9.1
Avg. storage	20.0	20.1	20.4	20.6	21.3	17.8	15.4	14.2	14.0	15.3
Safety Stock	15.0	15.0	15.0	16.0	17.0	15.0	15.0	15.0	16.0	17.0
Order Fill Rate (%)	99.6	99.2	99.1	99.1	99.1	99.2	99.2	99.1	99.1	99.1

4.7 Cost Analysis

In this paragraph we make a cost comparison between the strategies. The costs not only depend on used transport mode and average inventory levels, but also on the considered cost tariffs for calculating the transport, holding and storage costs. In practice cost calculations are rather intricate as they depend on many details and vary over time. That's why this comparison can only give an impression on the possible differences of the strategies without guaranteeing that these differences will hold in another situation as well. The estimated costs by simulation of the case are shown in table 4.8.

Table 4.8: Cost Comparison for the Case Simulation

	Unit	CS	DS	DS/CSS	MS
Transport costs	€ per FTL	880	902	904	901
Storage costs	€ per FTL	81	125	92	58
Holding costs	€ per FTL	214	275	254	203
Total costs	€ per FTL	1175	1303	1250	1162
Order fill rate	%	99.1	99.1	99.3	99.1
Average order lead time	days	2.0	1.0	1.1	1.3

These results show that in the simulated case the MS strategy is cheaper than the other three strategies. Although intermodal transport is more expensive than road transport, the MS strategy has lower total costs than the CS strategy. By making efficient use of floating stock, the storage and holding costs advantages are big enough to compensate the higher transport costs. Furthermore, the average order lead time of MS is shorter than the lead time of CS, so in this case the MS strategy should be preferred over the CS strategy anyway. The DS and DS/CSS strategies always need more inventory than the other two strategies as shown in the conceptual model analysis. This is why the holding costs of these strategies are always higher than those of the other two. Because in this case the transportation costs and storage costs are higher as well, these two strategies are inefficient with respect to the other two.

4.8 Discussion

The results in the previous section show that under the assumed conditions, the MS strategy is the most efficient of the four strategies in the area of inventory management. Using this strategy leads to the lowest storage level without significantly effecting the order fill rate. Given the sensitivity analysis results it has been shown that the efficiency of this strategy does not depend on the used data. In every simulation experiment this storage level advantage existed and it even increased with a longer intermodal transit time.

The average order lead time when using the MS strategy is always shorter than when using the CS strategy so on this performance criterion, the MS strategy beats the CS strategy in any case. However, whether the storage level advantage actually leads to storage costs savings depends partly on the storage tariffs as well, so we cannot draw a general conclusion about this. In the case situation the MS strategy is slightly cheaper than the CS strategy, despite the higher transport costs for using the intermodal transport connection; however, this does not need to hold in general. It proves that it is possible to obtain cost advantages by switching partially from using road transport to intermodal transport even on a route where intermodal transport is more expensive.

The condition that the production takes place in batches is essential for these results to hold in general. This is because the costs advantages of the MS strategy are obtained by keeping the part of the inventory moving (without causing storage costs) that is not expected to be ordered on the short run. If on the contrary the production is continuous or order-based, this part does not exist and these advantages cannot be obtained. We would like to note that in the intermodal distribution strategies it is not essential to send the whole batch directly by intermodal transport. One may send the first containers and some days later the others. This prevents a long residence time at the terminal. These calculations are much more complex and have been studied in a follow-up paper (Pourakbar et al., 2009).

The case explanation proves that the MS strategy can be profitable on both one single transportation lane and a whole distribution network. So the presence of a network with a couple of terminals and DCs as in the case is not necessary. However, a greater number of terminals and DCs cause the demand forecast to be more accurate, because an aggregate demand distribution has less uncertainty.

Finally we would like to remark that in reality one can make use of Megatrailers for truck transport, which carry 100 m^3. Although this changes most of the cost calculations, our conclusion that the use of the MS strategy has advantages over the other two intermodal strategies and that it improves the cost efficiency of intermodal transport compared to direct transport remains valid. Some calculations on this case were done in Ochtman et al. (2004).

4.9 Conclusions

Floating stock is a concept where a new production batch is (partly) pushed into the supply chain, without determining the exact destination for each product beforehand. Using this concept may lead to lower storage costs and a shorter order lead time, without a decrease in the order fill rate. This is possible if immediately after the production a part of the batch is centrally stored at the factory to deliver the orders in the first part of the production cycle, while the other part of the batch is transported intermodal to a regional stocking point. Orders in the last part of the production cycle are then fulfilled from these regional stocking points. This strategy offers the best opportunities to benefit from low storage levels, which is the goal of the floating stock concept.

The case study shows that the floating stock strategy can reduce costs and lead times, and improve the order fill rate, in spite of the possible higher transportation costs of an intermodal connection. The much longer transit time of intermodal transport is nullified by the floating-stock concept. Additionally, the strategy can help green the supply chain as the environmental impact of intermodal transport per ton per kilometer is lower for intermodal transport than it is for road transport. Application of the concept does require batch production, a somewhat predictable demand (in terms of volume), containerized transport, and a standardized product mix. Finally, we would like to advocate that inventory costs as well as transportation costs are taken into account when a move from road transport to intermodal transport is considered.

An additional benefit that has not been monetized in this thesis is the potential for reduced CO_2 emissions for the transport by rail or inland waterways instead of road. As the cost of transportation is likely to include a factor for environment impact (for example as part of road pricing schemes) within the foreseeable future, this will become a significant factor and provide an interesting avenue for future research.

Chapter 5

Flexibility in Port Selection*

Ports provide a number of logistical choices concerning storage, onward transport, and postponement. We investigate the routing flexibility offered by ports with a central location with respect to the hinterland. This flexibility is investigated using an illustrative case in which a number of alternative strategies are evaluated by means of simulation. Detailed cost data was used for the illustrative case. The combination of a simulation model and detailed cost data allows us to quantify the value of the rerouting flexibility. A combination of using regional distribution centers and a European Distribution Center results in the lowest cost per container.

5.1 Introduction

Academic research on ports has evolved from a focus on the physical infrastructure to the supply chain perspective, where a port is seen as a node in the supply chain network. Robinson (2002) promotes the paradigm of ports as elements in value-driven chains and contrasts this paradigm with previous paradigms such as the morphological framework (ports as places), the operational efficiency framework (ports as operating systems), the economic principles framework (ports as economic units), and the governance and policy framework (ports as administrative units). One aspect that has received little explicit attention is the role of ports as natural locations where multiple logistical choices are available regarding storage, onward transport, and postponement. Containerization has lowered the cost of transport and increased the speed of cargo handling through the use of a single, standardized type of load unit. It also allows rerouting standard cargo.

* This chapter has been submitted for publication to the "Maritime Economics & Logistics" journal.

The configuration of the logistics for a supply chain of fast-moving consumer goods such as consumer electronics or sports shoes involves many choices. These products are relatively expensive and have to reach the markets quickly but at a limited cost. The demand is that large that standardized containers can be used for the sea transport. Many manufacturers are located in Asia, particularly China, and the main markets for these goods are in Western Europe and the USA. Transportation from China to the markets is mostly done by sea, using standardized containers, as the costs for this mode of transport are low and the volumes large. If we look at the situation in Western Europe, we see several ports that are both close to the demand regions and able to handle the larger container ships. The focus of this paper is on the factors that influence the selection of these ports and in particular on the value provided by the flexibility of ports with a central location with regard to the hinterland.

If demand is less than a full container load, then it is not possible to store the goods in the transport unit (container); storage and handling would be required at a (European) distribution center. The sea transport from Asia to Europe is followed by inland transport by barge, train or truck. The container terminals provide some short-term storage capacity; while this terminal storage is intended to decouple the stages of the intermodal transport chain, it can also be used to postpone the routing decision (i.e., to which demand region the container will be shipped). This flexibility can be used to accommodate demand variations between regions. To ensure fast delivery to customers in the face of long supply chains, safety stock has to be held close to the demand region. This safety stock has to be stored at a physical location such a container terminal or a distribution center, incurring storage costs. Throughout the supply chain, one also incurs holding costs (the products tie up capital and depreciate in value over time). Both the location for and the amount of safety stock should therefore be carefully selected to minimize these costs.

The contribution of this paper is to quantify the routing flexibility that can be provided by ports with a central location. This makes the paper unique in the literature (see section 5.3). We use a case to illustrate our approach. The rerouting flexibility is used to some degree in practice at an operational level, where it is known as container rerouting.

This paper is structured along the following lines. We first formulate some hypotheses. Next, we review the academic literature, focusing on port selection and inventory control. We formulate an illustrative case (section 5.4) which provides the basis for numerical experiments using simulation (sections 5.5 and 5.6). With these experiments we evaluate the performance of the strategies and present the outcomes (section 5.7). We then present the results of a sensitivity analysis to assess the robustness of the outcomes (section 5.8). We close with a discussion of the results and the conclusion.

5.2 Hypotheses

We formulate the following hypotheses regarding the application of intermodal transport and distribution chains for a long sea-transport to a continent (e.g. the China–Europe route).

1. Ports with a central location with regard to the hinterland have a competitive edge due to greater flexibility in (re)routing traffic.

2. The value of this flexibility depends on the value of the products and the demand uncertainty.

We investigate these hypotheses using an illustrative case. The analysis is generic but the illustrative case provides more insight into the consequences. For this case we have gathered realistic data; we use real ports but aggregate the customers to demand regions that are represented as a point. The cost structure for the road haulage is based on the road distances covered; the cost data are based on consultation with practitioners as well as academic and professional literature. For the shipping network, we used expert judgements.

We use a simulation model to compare the effects of various strategies for the illustrative case. Simulation was selected as an evaluation method because some of the strategies we want to evaluate can not be analyzed analytically. This case study focuses on Western-Europe but the approach can also be applied to other regions. The simulation model was verified using a set of test cases.

5.3 Literature Review

In the academic literature the topics of shipping networks, logistics planning and inventory policies are most commonly studied separately. In this paper, we apply the concept of floating stocks to the transport of fast-moving consumer goods (FMCG) from Asia to Western Europe. The floating stocks concept draws on the areas of transportation and inventory control: in addition to the literature on port selection, we will therefore briefly recapitulate relevant transportation and inventory control literature.

5.3.1 Transportation

We will first look at the topic of port competition, followed by the topic of port selection, and the role of container terminals.

The geographical location of ports and demand regions on a continent are important factors in port selection. For the Asia-Western Europe route, one could argue that offloading cargo in Southern Europe would be beneficial for the carriers as this would shorten the trip time. The main disadvantages of this approach are that there are few ports that can handle the larger container ships and that most of the cargo would be far from the demand region. Therefore, most cargo from Asia to Western Europe is discharged in the Hamburg-Le Havre range in North-Western Europe; from the ports in this range, the main demand regions in Western Europe can be reached within days. Thus, one should also look at the connectivity of a port; when there is choice amongst multiple ports in a region, the position of that port in the transport network becomes an important characteristic. When routing most cargo through a single or a small number of ports, there is also the possibility of postponement; cargo can be rerouted at a fairly late moment in the transportation process. Offloaded cargo that is stored in a container terminal can be rerouted as demand shifts across regions. A port with good connectivity and support for multiple modes of transportation will be a more attractive choice. Port selection is an issue for both carriers (as part of the network design) and shippers (as part of the transport choice decision process). In this paper we take a supply-chain perspective and will thus focus on the transport choice rather than the network design.

Port competition

Chang and Lee (2007) provide an overview of the literature on port competition. They note that port competition has risen in prominence as a result of containerization and identify five main topics: governance, performance, cooperation, competitive policy, and port selection factors. In their review of the literature on port selection, they conclude that most studies have focused on the shippers rather than on other stakeholders. The methodologies that were applied to performance evaluation tended to be quantitative; cooperation was researched using conceptual, descriptive and case studies research; qualitative surveys were used for studies on governance and port selection. For the latter, surveying shippers and port authorities are popular methods. Several papers use this method to determine the factors that influence port selection and port performance (see for example Tongzon (2002)); Yeo et al. (2008) used literature review and a regional survey to find determinants of port competitiveness in Korea and China.

They group the determinants into seven main categories (port service, hinterland condition, availability, convenience, logistics cost, regional center, and connectivity). The last category, connectivity, partially matches our flexibility (in terms of land distance, connectivity to major shippers, and efficient inland transport network). They also mention terminal free dwell time (as part of the logistics cost) as a significant factor for port selection.

Notteboom and Rodrigue (2005) argue that inland distribution is an important factor in port competition. They propose that regionalization expands the hinterland reach by linking the port more closely with inland freight distribution centers. Notteboom and Rodrigue (2008) indicate that terminal managements skills (software and know how) and hinterland size are key to productivity gains for container terminal operators. They also signal the development of multi-port gateway regions such as those in the Hamburg and Rotterdam–Antwerp regions and the integration of ports, liner shipping networks and hinterland transport.

Veldman and Bückmann (2003) use a logit model to forecast a port's market share in terms of container throughput, based on demand choice models. The logit model was used for an economic analysis of the port-extension project "Maasvlakte-2" in Rotterdam. Their approach includes model and route choices but at an aggregate level only. In line with the regionalization trend described by Notteboom and Rodrigue (2005), they focus on the European end; the Hamburg-Le Havre range. The analysis is limited to the factors of shipping costs for a route, the transit time, and the frequency of service. In their paper, the authors note that carriers use "the same tariff to each of the continental seaports".

Port selection

In his seminal paper on port selection, Slack (1985) argues that the traditional criteria such as port equipment appear to have relatively little influence on the port selection process. The flexibility of hinterland transportation is discussed only in an indirect fashion (for example, as 'number of sailings' and 'possibility of intermodal links', (Slack, 1985)). Some papers do touch on this topic; consider de Langen (2007), in which the port selection process for cargo destined for Austria is considered. In that paper, the flexibility of onward transport for imports gave the port of Rotterdam an edge over competing ports due to the possibility of barge traffic; there is a choice between fast and expensive transport (by road) and slower but less expensive transport (by inland waterways). In Wiegmans et al. (2008), the focus is on the selection of ports and terminals in the Hamburg–Le Havre range; for both decisions, the availability of hinterland connections is a key determinant (the methodologies employed were literature review, interviews with industry practitioners and application of decision-making theory). The immediacy of consumers (large hinterland) is found to be a determinant for the port selection.

Transportation choice

Meixell and Norbis (2008) provides a review of the transportation mode choice and carrier selection literature. The academic literature is categorized by topic, methodology, and challenges. They note the low use of simulation and interviews as methodologies. McGinnis (1989) classifies the models of freight transportation choice as classical economic, inventory-theoretic, trade-off, and constrained optimization, and identifies the variables involved in transportation and non-transportation costs through a review of empirical literature. (Flexibility is not mentioned.) The approach in our paper has elements of the inventory-theoretic and constrained optimization models. Naim et al. (2006) is one of the few papers that discuss flexibility in transportation. Starting from the use of flexibility in the field of manufacturing, they perform a synthesis of the literature to identify the key components of transport flexibility. The flexibility as used in our paper would classify as external (i.e., provided by transportation providers to customers), volume (range of and ability to accommodate changes in transport demand), delivery (range of and ability to change delivery dates) and access flexibility (ability to provide extensive distribution coverage).

Container terminals

The role of container terminals in supply chain logistics from the perspective of the terminal operator is discussed by Panayides and Song (2008). They identify four key variables for the integration of terminal operators in the supply chain: information and communication systems, value-added services, multimodal systems and operations, and supply chain integration practices. van der Horst and de Langen (2008) acknowledge the important role of hinterland transport (with costs often exceeding that of the maritime transport) and focus on the coordination between seaport actors involved in hinterland chains. Mangan et al. (2008) provide an overview of port-centric logistics. The introduction of ever larger vessels causes a concentration of traffic to larger ports. This in turn creates hub and spoke networks with feeder ports. 'Lean' supply chains with relatively long lead times and predictable demand cause a focus on cost-effective storage capabilities; 'leagile' supply chains with long lead times and unpredictable demand lead to postponement of manufacturing/assembly. Olivier et al. (2007) signal the development of container terminal from cost centers to profit centers: increasingly, container terminals are operated by transnational companies, and while Asian companies have come to dominate the global terminal business, the European carriers have taken the lead in delivering total logistics packages. The shift from cost centers to profit centers may negatively affect the free dwell time.

Christiansen et al. (2004) present an overview of ship routing and scheduling literature. They distinguish routing (sequences of ports to be visited) and scheduling (timing the sequences) within network design. They do not mention port selection issues.

5.3.2 Inventory

In this paper we do not consider transport on its own, but take a supply chain perspective; we look at the total system of inventory and transportation facilitites to satisfy customer demand. Supply chain responsiveness can be obtained by fast transport and by keeping inventories close to customers. Normally, inventories are stored in distribution centers (DC's) which have fixed cost elements: setting up a DC requires an investment up-front and annual costs for operation that are incurred irrespective of the volume of inventory kept.

We consider a form of inventory speculation, which is appropriate for low customer order-to-delivery time and high-delivery frequency (Wallin et al. (2006), Pagh and Cooper (1998)). Inventory speculation is also identified as the method of choice by Baker (2007), if supplier lead times far exceed customer lead times.

Many papers in the area of supply chain strategies are qualitative in nature and do not provide quantitative results (such as Pagh and Cooper (1998)). The floating stock concept (Ochtman et al., 2004), in which intermodal transfer points are used as short-term storage locations for advance deployment of stock, attempts to find a middle ground between speculative and postponed logistics in the terminology of Pagh and Cooper (1998) (see their figure 6). In this concept, the geographical layout of the demand region and the available transport infrastructure are exploited to delay the final choices until retailer demand materializes. Ochtman et al. (2004) apply this to a case within Europe; they use simulation to numerically evaluate the performance of several stock deployment and transport strategies. Pourakbar et al. (2009) provide a mathematical analysis of the floating stock distribution concept; they present two mathematical models to analyze the floating stock policy with backlogging allowed and determine the optimal shipping time of containers through intermodal routes.

Baker (2007) discusses the role of inventory and warehousing in international supply chains, the role of decoupling points, and distribution centers; he maps 13 different supply chains from six companies, including some FMCG supply chains, on the basis of a survey.

Huggins and Olsen (2003) present an analytical model that takes the expediting and holding costs into account but their expediting cost model is fairly simple (fixed cost plus a linear function of volume) and they don't include lead times.

5.3.3 Summary of literature review

Most papers on port competition and selection are qualitative in nature and focus on identifying the key variables. Few papers approach this from a quantitative angle. The flexibility aspect has received some attention in the transportation literature, but again mostly from a qualitative perspective.

The main trends we synthesize from the literature are the shift towards terminal management and hinterland size as determinants for productivity gains for container terminal operators, the regionalization of ports, the advent of "port-centered logistics", and the development of hub-and-spoke networks. When combined with the industry trends towards ever larger container vessels and a reduction in the number of ports called, we conclude that an approach that can quantify the routing flexibility of a port could be a valuable tool in the port selection decision process.

5.4 Illustrative Case

In this case we consider a supplier of fast-moving consumer goods (FMCG) with variable demand, such as DVD-players, LCD televisions, sports shoes, or clothing. The supplier is typically located in China (for the case study, we will use a supplier close to the port of Shanghai). The goods are shipped in containers via Shanghai to Western Europe. (See figure 5.1.) We distinguish five demand regions: Austria (AT), Belgium (BE), Germany (DE), The Netherlands (NL), and the United Kingdom (UK). Here we limit ourselves to these five countries to limit the amount of data to analyze. (A larger case has also been evaluated and we will return to this topic in the discussion section.)

For the distribution in Western Europe, we are faced with a choice of ports and storage locations. We consider five ports (La Spezia in Italy, Antwerp in Belgium, Hamburg in Germany, Rotterdam in The Netherlands, and Southampton in the United Kingdom) that are able to meet the following requirements: they are physically close to the demand regions, they feature in Asia to Europe container shipping schedules, and they are visited by large to very large container ships. The frequency for the

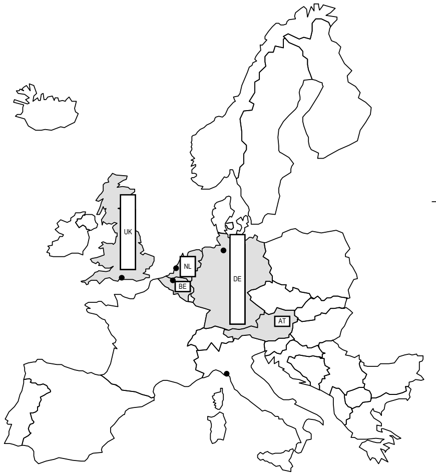

Region	Ratio	Interarrival time in days ($1/\lambda$)
NL	2.1	4.3
BE	1.2	7.3
DE	8.9	1.0
UK	7.4	1.2
AT	1.0	8.9

Figure 5.1: Map of Western Europe with ports and Relative demand per region

Shanghai–Rotterdam and Shanghai–Hamburg connections is high and costs are (comparatively) low. For the Shanghai–Antwerp and Shanghai–Southampton connections, the frequency is lower but the costs are identical to the Rotterdam and Hamburg connections (Veldman and Bückmann (2003) confirm that the tariffs for the continental ports are identical). The frequency for the Shanghai–La Spezia connection is low and the costs are higher.

The demand ratios for the regions are based on the 2007 Gross Domestic Product (GDP) per country, based on the April 2008 World Economic Outlook Database (IMF (2008), see figure 5.1). So, for example, the demand for Germany is approximately four times the demand for the Netherlands and eight times the demand for Austria and Belgium. The demand is modeled by a Poisson process at each location; the interarrival time between demands (customer orders in units of a full container) thus follows a negative-exponential distribution function. The parameters for the distribution function for each demand region are based on the demand ratios and are listed in the table next to figure 5.1. The size of an order is always one full container or TEU (twenty-foot equivalent unit).

Given this geographical layout of the demand regions and ports, we can determine which port has the most 'central' location. The parameters are the distance from each port to the demand regions, the cost per trip, and the volume for these links. Using these parameters we can calculate the average cost to deliver a container from a port to the demand regions. Using the road distances from each port to each demand region

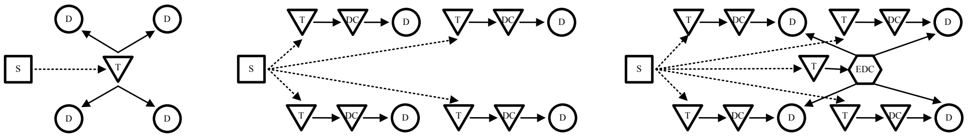

Figure 5.2: Diagram of centralized (left), Decentralized (middle), and EDC (right) Strategies. (S indicates Supplier; T–container terminal; DC–distribution center; D–demand region. Solid lines represent truck transport and dashed lines represent ship transport.)

and the truck transport cost model that we used for this case (v.i.), we can compute this average cost per trip. Antwerp has the lowest average cost, at $764, closely followed by Rotterdam at $839; Hamburg, Southampton, and La Spezia trail at $994, $1,049, and $1,396 respectively. On the basis of this calculation, Antwerp and Rotterdam are the best candidates to locate a central distribution center.

Taking the location of the ports and the demand regions into consideration, we now formulate a number of alternative layouts or strategies. In the first strategy the containers are shipped from China to the port closest to demand regions. From these ports the containers are then transported by road to distribution centers located in the demand regions. Some safety stock is required in these DC's to secure fast delivery. We will refer to this strategy as the decentralized strategy ("DEC"). (Figure 5.2 provides a schematic representation of the strategies; "S" represents the supplier in China, "T" the container terminals, "DC" the distribution centers, and "D" the demand regions. In this diagram, dashed lines represent sea transport and solid lines are use for land transport by truck.)

Alternatively, we can pool all demand and fulfill it from a central location. In this centralized strategy ("CEN") the containers are shipped to a container terminal at a port with a central location (for this case we selected Rotterdam). Customer demand is fulfilled by road transport from the central location; we keep some safety stock at this location to ensure fast delivery.

The centralized strategy uses the container terminal for temporary storage. This is fine if the dwell time is short; if the dwell time is longer, then this may present problems (especially higher costs). Both the centralized and the decentralized strategies depend on moving the load unit (container) as-is; there is no opportunity to rebatch the goods into smaller units for regions with lower demand. We therefore define a third strategy that includes a European Distribution Center. In this strategy (labeled "EDC") we will ship most containers to the ports closest to the demand region; onward transport to the distribution centers in the demand regions is done by truck. Safety stock is kept at the European Distribution Center, which we located close to the port of Rotterdam, and at the distribution centers in the demand regions.

Most stock is held at distribution centers, except for the centralized strategy. In the centralized strategy we use the container terminal as a storage location. In this strategy the demand from all regions is pooled. We expect the average dwell time of a container to be short and can thus take advantage of terminal free time (which is also mentioned by Slack (1985) as a relevant port selection criterium). In the DEC and EDC strategies the dwell times will be longer; stock is therefore stored at distribution centers located close to the demand regions.

The main performance indicators are the fraction of orders fulfilled within a preset time limit (three days) as a measure of responsiveness, and the total cost per container. The total cost per container includes the costs of shipping, inland transportation, handling, storage (at terminals and distribution centers), and holding (depreciation). We hold sufficient safety stock to ensure that 95% of orders are fulfilled within the preset limit. To minimize the cost associated with this safety stock (both storage and holding costs), we need to determine the inventory levels that will satisfy the order fulfillment requirement while minimizing the overall cost per container. Secondary indicators are the lead time, as an indicator of supply chain responsiveness, the residence times, and the inventory levels.

In this model we use a continuous-review base-stock policy with parameters (S-1,S) as we assume the absence of economies of scale in container transport for ordering more than one container. These parameters are defined for each storage location (terminal or DC). When an order for a container is placed, a replenishment order is initiated. Orders are fulfilled from either the on-hand stock or from the virtual (in-transit) stock; the former is shipped immediately, the latter is shipped when the stock arrives at the storage location.

5.4.1 Costs

The transport network consists of a number of storage locations such as warehouses, distribution centers, and container terminals. The storage locations are connected via transport links. The transport links are either shipping links (sea-going vessels) or trucking links (land transport using trucks). We do not consider barge or rail transport to limit the complexity of the model.

To determine the total cost of transporting a container to the customer, we introduce the following variables:

L the set of transport links

S the set of storage locations

t_i^t the time (in days) is takes to transport an individual container on transport link i.

c_j^1 cost of handling incoming goods at storage location j.

c_j^2 cost of handling outgoing goods at storage location j.

c_j^s cost of storing a container at location j for one day.

t_j^s the time (in days) a container is stored at location j (this includes the time for handling at arrival and departure).

t_j^f free dwell time for a container at location j.

c_i^t the cost of transporting a container along transport link i.

c^h the holding cost per day for a single container.

The total cost per trip for a container can then be formulated as:

$$\sum_{i \in L} \left(c_i^t + (t_i^t \times c^h) \right) + \sum_{j \in S} \left(c_j^1 + (t_j^s \times c^h) + (max(t_j^s - t_j^f, 0) \times c_j^s) + c_j^2 \right)$$

The total cost can vary for each individual container because the times a container is stored or transported (t_j^s and t_i^t) can vary, depending on the strategy selected. In this formulation, some of the t_i^t and t_j^s will be zero (for transport links and for storage locations that were not used). The time a container is stored at a location will vary, depending on the time required to arrange for pickup by truck or on the sailing schedule.

5.4.2 Inventory model

The base scenario of our illustrative case can also be analyzed with the METRIC model (Sherbooke, 2004). In the METRIC model a multi-echelon inventory chain is considered with a central depot and several bases where demand occurs. Demand at each base is modeled by a Poisson process. The bases apply a (S-1,S) model for stock replenishment. Every base has a lead time for replenishment from the central depot, which can reorder at a supplier. The output of the METRIC model comprises the service level (fraction of customers supplied in time) as well as the distribution of the inventory level as function of the base stock levels. In our case the bases are the demand region DC's and the central depot is the EDC or central port. We applied the METRIC model to verify our simulation model and the results were within the confidence bounds. As it is very difficult to accommodate other demand distributions in the METRIC model and to accommodate for the free time at terminals, we have chosen for simulation as evaluation method.

5.5 Simulation Model

The simulation model was implemented in the Java programming language, using the open-source SSJ discrete-event framework (L'Ecuyer et al., 2002). The configuration for a particular experiment is specified in a spreadsheet file. This includes the transport network, inventory locations, the costing parameters, and the experimental setup (warm-up time, length of a run, and the number of replications). This implementation facilitates easy experimentation without having to modify the Java source code. (Tables 5.1–5.3 are near verbatim copies from these spreadsheets.)

The simulation model tracks each container individually from creation (at the factory), through transport and storage to the final delivery at the customer. This tracking allows for detailed calculation of the costs (transport, storage, holding).

The source for this simulation model totals around 6,500 lines of code. It takes ten seconds on a 2.4 Ghz Intel Core 2 Duo processor to run 30 replications of six years each. The number of replications and the length per run were set to these values to get small confidence intervals on the statistical outputs. (For example, the 95% confidence interval on the 'average total cost' statistic is approximately 0.1 percent.)

5.6 Experimental Setup

5.6.1 Decentralized strategy (DEC)

In the first scenario we ship the goods directly from China to the port closest to each demand region. The safety stock is held at distribution centers close to the local ports; the level of safety stock required is determined by repeatedly running the simulation model and increasing S until the order fulfillment requirement (95% of orders is delivered within three days) is met for each location. (For this scenario, it is possible to determine the parameters for the inventory policy analytically (f.e., using the method discussed by Chopra and Meindl (2004) on p.326), but in more complicated scenarios

Table 5.1: Network Links for Decentralized Strategy (DEC)

From	To	Modality	Duration (days) t_i^t	Cost ($/TEU) c_i^t	Interval (days)	Distance (km)	City
China vendor	Shanghai port	truck	0.5	150			
Shanghai port	La Spezia port	ship	19.0	1580	14		
Shanghai port	Antwerp port	ship	21.0	1340	7		
Shanghai port	Hamburg port	ship	21.0	1340	3.5		
Shanghai port	Rotterdam port	ship	21.0	1340	3.5		
Shanghai port	Southampton port	ship	21.0	1340	7		
La Spezia port	DC in AT	truck	1.1	920		920	Vienna
Antwerp port	DC in BE	truck	0.1	79		20	
Hamburg port	DC in DE	truck	0.1	79		20	
Rotterdam port	DC in NL	truck	0.1	79		20	
Southampton port	DC in UK	truck	0.1	79		20	
DC in AT	Customer AT	truck	0.1	79		50	Vienna
DC in BE	Customer BE	truck	0.1	130		130	Liege
DC in DE	Customer DE	truck	0.4	580		580	Nuremberg
DC in NL	Customer NL	truck	0.1	100		100	Amersfoort
DC in UK	Customer UK	truck	0.2	330		330	Manchester

this becomes more difficult; thus, we have chosen the same approach throughout.)

We assume there are sailings from Shanghai to Rotterdam and from Shanghai to Hamburg twice a week; Antwerp and Southampton are visited once a week, and La Spezia once every fortnight. While the frequencies of all carriers on these routes combined may be higher, an individual client of a carrier will usually be limited to the sailings offered by that carrier. For each demand region, we have selected a location for the average customer (listed in the 'City' columns below) to determine transport distances, and thus times and costs. (See table 5.1 for the details of the links in the network.) It takes two days to unload a container in a port, stack it in the yard, get cleared through customs, and arrange onward truck transport to a DC. The time to arrange transport from a DC to the customer is included in the order generation process.

Table 5.2: Network Links for Centralized Strategy (CEN)

From	To	Modality	Duration (days) t_i^t	Cost ($ per TEU) c_i^t	Interval (days)	Distance (km)	City
Supplier China	Shanghai port	truck	0.5	150			
Shanghai port	Rotterdam port	ship	21.0	1340	3.5		
Rotterdam port	Customer AT	truck	1.2	1200		1200	Vienna
Rotterdam port	Customer BE	truck	0.2	250		250	Liege
Rotterdam port	Customer DE	truck	0.9	710		710	Nuremberg
Rotterdam port	Customer NL	truck	0.1	150		150	Amersfoort
Rotterdam port	Customer UK	truck	1.3	1212		780	Manchester

5.6.2 Centralized strategy (CEN)

The second scenario routes all transports through a centralized port, in this case a container terminal in Rotterdam. Using the container terminal eliminates extra handling times and costs when compared to storing in a European Distribution Center, even though the storage costs can be high. Onward transport from Rotterdam to the demand regions is by truck. The details of the transport network links are listed in table 5.2.

5.6.3 EDC strategy (EDC)

In this scenario, the goods are shipped to the regional ports close to the demand regions. Some stock is kept at distribution centers close to these regional ports and an additional safety stock is kept at a European Distribution Center (EDC) in Rotterdam. Fulfillment from the EDC to the customers is via truck. (See table 5.3 for details.) Storing at an EDC implies extra handling costs; for long dwell times, storing at an EDC is cheaper than at a terminal. We do not explicitly consider the fixed costs of an EDC or regional warehouse.

5.6.4 Cost parameters

For the cost associated with stocks, we specify the holding costs and the storage costs separately (Chopra and Meindl (2004) refer to these as the ‘cost of capital’ and the ‘occupancy costs’). Here, the holding cost is the money spent to maintain a stock of

Table 5.3: Network for European Distribution Center Strategy (EDC)

From	To	Modality	Duration (days) t_i^t	Cost ($ per TEU) c_i^t	Interval (days)	Distance (km)	City
The first six rows are identical to the CEN strategy (see table 5.1)							
La Spezia port	DC in AT	truck	1.1	920		920	
Antwerp port	DC in BE	truck	0.1	79		20	
Hamburg port	DC in DE	truck	0.1	79		20	
Rotterdam port	DC in NL	truck	0.1	79		20	
Southampton port	DC in UK	truck	0.1	79		20	
DC in AT	Customer AT	truck	0.1	79		50	Vienna
DC in BE	Customer BE	truck	0.1	130		130	Liege
DC in DE	Customer DE	truck	0.4	580		580	Nuremberg
DC in NL	Customer NL	truck	0.1	100		100	Amersfoort
DC in UK	Customer UK	truck	0.2	330		330	Manchester
Rotterdam port	EDC	truck	0.1	79		20	Rotterdam
EDC	Customer AT	truck	1.2	1200		1200	Vienna
EDC	Customer BE	truck	0.2	250		250	Liege
EDC	Customer DE	truck	0.9	710		710	Nuremberg
EDC	Customer NL	truck	0.1	150		150	Amersfoort
EDC	Customer UK	truck	1.0	1212		780	Manchester

goods, excluding the cost of storing those goods. For this case, the holding costs are based on a cargo of approximately 2,000 DVD players per container valued at $45 each. This means that the total value of a single container is $91,250. At an interest level of 8%, the holding cost per container per day c^h is then $20.

The storage costs c_j^s are 10 dollar per TEU per day, 5 free days for the container terminals (t_j^f); at 5 dollar per TEU per day, no free days for the European Distribution Center; and at 6 dollar per TEU per day, no free days for the regional DC's. The regional DC's are more expensive due to fewer economies of scale than the EDC.

The container terminal handling charges are based on expert opinion at $120 for Antwerp, $140 for Rotterdam, $160 for Hamburg, and $180 for La Spezia and Southampton*.

5.6.5 Transport parameters

The costs for truck transport are based on a simple model that is linear in the distance covered. The cost is $1 per TEU per kilometer, with a minimum of $79 per trip. (These parameters were based on expert opinion; Notteboom (2004) cites a range of $0.8 to $2 per TEU-kilometer for inland haulage per truck.) These parameters are used to calculate the c_i^t for the truck transports.

The distances from the terminals to the customers were estimated using the 'Driving Directions' feature of Google maps (Google, 2008); a sample of these distances was verified using the Microsoft AutoRoute 2007 software package. The travel times for trucks (t_i^t) have been calculated using the Dutch regulations for driving/rest-times, an average speed of 80 km/hr and a one-hour overhead per trip.

The tariff for shipping one TEU from Shanghai to ports in the Hamburg–Le Havre range is 1340 dollar per TEU (c_i^t for shipping routes); there is a 20% premium for the Shanghai–La Spezia route. For trips between the European continent and the UK, we include in the inland transport costs the cost of a channel tunnel crossing, which is $432 one way (based on tariff from the Eurotunnel website).

* As we did not have data for La Spezia and Southampton, we have selected the highest charge (Hamburg) and added a small premium to model the lack of economies of scale. No terminal handling charges were defined for the Shanghai terminal because they would not cause any difference in the results.

Table 5.4: Inventory Policy Parameter (S)

Location	DEC	CEN	EDC
DC in AT	7	-	4
DC in BE	8	-	4
DC in DE	32	-	21
DC in NL	10	-	7
DC in UK	30	-	21
Rotterdam terminal	-	69	-
EDC (Rotterdam)	-	-	18
	87	69	75

5.7 Experimental Results

The experiments consist of 30 replications of six years; before the start of each replication the system is warmed-up by ordering and delivering the base-stock level for each location. Here, we report the means over the number of replications; more detailed statistics, including the 95% confidence intervals and 90% quantiles are available in a separate technical report.

We have run the simulation model for the three strategies described above (DEC, CEN, and EDC). The order-up-to levels (parameter S of the inventory policy) are listed in table 5.4. The order-up-to levels were determined by repeatedly running the model, increasing S if necessary, until the fraction of orders that could be fulfilled within three days exceeded 0.95 for all regions. For the EDC strategy, we determined the order-up-to levels by starting with no EDC stock and determined the required stock levels for the other DC's. We then increased the EDC level by one unit (container) at a time; for each EDC level we decreased the local levels until we found the minimum level necessary to meet the order lead time requirement. We selected the setting with the lowest total cost.

The inventory required to meet this requirement is smallest with the centralized strategy (CEN); as expected, pooling demand clearly has a significant impact on the level of stock required. The results of the EDC strategy exhibit a similar effect with regard to the pooling of the safety stock; this strategy has the additional benefit of a lower average lead time because most of the stock is held closer to the demand regions.

For easy comparison, we have calculated the average lead time and costs; these reflect the differing demand volumes per region. These indicators are listed in table 5.5. The average lead-time is the number of days it takes to fulfill an order from the demand

Table 5.5: Overview Results (averages per container)

Totals	Unit	DEC	CEN	EDC
Avg. Lead Time	days	0.5	1.0	0.7
Avg. Handling Cost	$/TEU	164	140	159
Avg. Holding Cost	$/TEU	759	618	663
Avg. Storage Cost	$/TEU	71	29	40
Avg. Shipping Cost	$/TEU	1352	1340	1349
Avg. Transport Cost	$/TEU	660	980	740
Avg. Total Cost	$/TEU	3006	3106	2951

region. The cost parameters are the average costs per container. The EDC strategy has the lowest average total cost; the centralized strategy has the highest total cost and the longest lead-time; the savings in holding costs and storage costs due to pooling are offset by higher inland transport costs. In this case, the trips to the UK demand region are relatively expensive due the additional charges for the channel tunnel. The decentralized strategy has the best performance in terms of the lead time; this is expected as the inventory is held close to the demand regions at the regional DC's.

If we look at the overall cost per container delivered to the customers (table 5.6), we can see that centralized strategy lowers the costs most for regions with relatively low demand (Austria and Belgium); the cost savings for regions with high demand are more modest in comparison to the decentralized strategy. The EDC strategy leads to an increase in storage costs (as this strategy does not take full advantage of the free dwell time at the container terminals) when compared to the centralized strategy but this is balanced by a reduction in the inland transport costs.

The lead times per demand region (i.e., the time between the moment of ordering by the customer and the actual delivery to the customer) are in line with expectations (table 5.7). As the inventory levels were set on the basis of the fulfillment requirement (95% fulfillment within three days), we can expect the best performance from the strategy that places most of the stock closest to the demand region. The centralized strategy has a higher average order lead time: the stock is now further from the demand regions and final delivery from the central stock to the customer by truck takes longer than delivery from the regional port.

The EDC strategy fulfills most orders from the regional distribution centers and some orders from the European Distribution Center. The inventory policy settings for this strategy place a modest amount of stock at the EDC and significant amounts at the regional DC's (table 5.4); the lead time is thus longer than the lead time of the DEC

Table 5.6: Cost per Customer

Customer	Total Costs			Transport Costs		
	DEC	CEN	EDC	DEC	CEN	EDC
Customer AT	4371	3483	3897	1149	1350	1209
Customer BE	3118	2516	2687	359	400	391
Customer DE	3006	2989	2977	809	860	841
Customer NL	2779	2414	2602	329	300	337
Customer UK	2870	3491	2935	559	1362	727
	Storage Costs			Holding Costs		
	DEC	CEN	EDC	DEC	CEN	EDC
Customer AT	223	28	110	1239	624	877
Customer BE	177	30	68	1122	607	763
Customer DE	42	29	27	655	620	614
Customer NL	107	29	63	863	604	722
Customer UK	59	29	36	732	620	660

Table 5.7: Lead Times (in days)

Customer	DEC	CEN	EDC
Customer AT	0.4	1.4	0.8
Customer BE	0.3	0.4	0.6
Customer DE	0.6	1.1	0.8
Customer NL	0.4	0.3	0.3
Customer UK	0.5	1.1	0.6

strategy but shorter than the lead times of the CEN strategy. The lead time increases most for the demand region that is furthest from the EDC (Austria) as more orders are fulfilled by a long truck trip from the EDC.

For the CEN strategy, the average dwell time at the container terminal is within the terminal free time at 4.8 days (see table 5.8; the 95% quantile just exceeds the free dwell time at 5.2 days). This is reflected in lower storage costs (table 5.6).

5.8 Sensitivity Analysis

We have performed a sensitivity analysis to investigate the robustness of the outcomes. We have focused on the holding costs, the inland transport costs, the demand functions, and the free time for container terminals. Table 5.9 contains the total cost for

Table 5.8: Residence Times (in days)

Location	DEC	CEN	EDC
DC in AT	32.5	-	16.8
DC in BE	29	-	12.3
DC in DE	7.1	-	3.0
DC in NL	17.8	-	10.8
DC in UK	9.3	-	4.7
Antwerp terminal	2	-	2
Hamburg terminal	2	-	2
La Spezia terminal	2	-	2
Rotterdam terminal	-	4.8	2
Shanghai terminal	2.7	1.8	2.5
Southampton terminal	2	-	2
EDC (Rotterdam)	-	-	10.8

each scenario of the sensitivity analysis results. In the 'absolute' column, the best performing strategy (indicated with the value '0') is used as a benchmark for the other strategies. The 'relative' column shows the relative differences of each scenario when compared to the base case. The EDC strategy has the lowest total cost for most scenarios. The exceptions are the extreme holding cost scenarios and the less variable demand (Erlang(9)). The decentralized strategy is most sensitive to the demand distribution function; as the demand becomes less variable, it becomes easier to meet demand from the decentralized stocks. Overall, the performance of the simulation model appears to be sensitive to the holding costs and the demand functions.

5.8.1 Holding costs

In our base case, the holding cost is $20 per TEU per day; for the sensitivity analysis, we have also run the model with values of $5, $10, $40, and $100 per TEU per day to reflect two lower and two higher valued scenarios. Changing the holding costs will only affect the 'Holding Cost' and 'Total Costs' outputs. The strategies that include pooling (CEN and EDC) benefit from increases in the holding costs. For the lowest holding costs of $5 per TEU per day, the disadvantage of higher overall inventory for the decentralized strategy is offset by lower inland transport costs. As the holding costs increase, the EDC strategy offers a nice balance between inventory pooling and lower inland transport costs caused by keeping some inventory closer to the demand region. Finally, for the

Table 5.9: Overview of Sensitivity Analysis (total cost per TEU)

	Relative to base case			Absolute		
	DEC	CEN	EDC	DEC	CEN	EDC
Base Case				+55	+155	0
Holding Cost 5	-19%	-15%	-17%	0	+206	+16
Holding Cost 10	-13%	-10%	-11%	+8	+179	0
Holding Cost 40	+25%	+20%	+22%	+151	+110	0
Holding Cost 100	+101%	+80%	+90%	+465	0	+25
Demand Erlang(2)	-4%	-1%	-2%	+18	+179	0
Demand Erlang(9)	-7%	-3%	-5%	0	+213	+9
Inland Transport 2$/km	+14%	+22%	+16%	+6	+349	0
Free dwell time 0 days	+1%	+1%	+1%	+57	+149	0
Free dwell time 2.5 days	-	-	-	+57	+157	0
Free dwell time 7.5 days	-	-	-	+55	+154	0
Free dwell time 10 days	-	-	-	+55	+151	0
40 ft container	-8%	-16%	-15%	+222	+87	0

highest holding costs of $100 per TEU per day, the centralized strategy provides the lowest total costs; the higher inland transport costs from the central location to the demand regions are offset by savings in the holding costs due to pooling. Conversely, the average cost per container increases significantly for the decentralized strategy as this strategy does not feature any pooling.

5.8.2 Inland transport costs

For the base case, we use a tariff of $1 per TEU-km for inland (road) transport. As Notteboom (2004) mentions a range of $0.8 to $2 per TEU-km, we have also done an experiment using the upper limit of this range, $2 per TEU-km. (In line with Notteboom (2004) we assume that there are no economies of distance.) The tariff per TEU-km has the biggest impact on the centralized strategy as it uses the most and the longest truck transport trips. The difference between the EDC and the DEC strategy is now very small. The EDC strategy could additionally benefit from a location adjacent to the terminal. If the EDC could be reached by the terminal transporters, the transport from the terminal to the EDC could be performed at the discretion of the terminal operator. This could benefit both the costs of the move (even a short move by truck costs $79) and the operation of the terminal itself as it would allow the terminal operator to schedule these moves away from peak times. (Consider, for example, the Distripark concept used at the Maasvlakte in Rotterdam (United Nations (2002), p.44); a site directly adjacent to the ECT container terminal with a dedicated internal transport track.)

5.8.3 Demand function

In the base case, we modeled demand using the familiar negative-exponential distribution function for the order interarrival times per region. This distribution function generates a large proportion of very short interarrival times. The negative-exponential function is a specific case (shape parameter $k = 1$) of the more general Erlang distribution function. To examine the sensitivity of the simulation results for the demand function, we have done two additional experiments with the Erlang distribution function with the shape parameter value set to $k = 2$ and $k = 9$. For higher values of the shape parameter, the proportion of very short interarrival times will diminish; in essence, the order arrivals will be more evenly distributed over time, modeling more predictive demand.

The Erlang distribution function can model the sum of a number of exponential distributions; thus, the $k = 2$ and $k = 9$ cases are a model for a number (two, nine) of customers within a demand region. For this analysis, the second parameter of the Erlang distribution (the scale parameter θ) was set to have the same mean for all three functions ($k\theta$ is constant).

The order-up-to parameter S was determined separately for each distribution function. The stock required to meet the lead time constraint (95% of orders delivered within three days) is lower for higher values of the shape-parameter k. As k increases, the scale parameter θ decreases ($k\theta$ is constant). This implies that the variance ($k\theta^2$) decreases. All strategies benefit in a similar way; holding and storage costs are reduced. With more predictable demand (higher values of k), the DEC strategy benefits most: the safety stock required drops from 87 ($k = 1$) to 77 ($k = 2$) and 68 ($k = 9$) units (for CEN and EDC this numbers are 69-65-61 and 75-72-67, respectively). Less variable demand reduces the advantage of the strategies that involve pooling, making the decentralized strategy that places the inventory close to the demand regions the most attractive in term of total cost per TEU. (for $k = 9$, the total costs per TEU are \$2,802 (DEC), \$3,015 (CEN), and \$2,811 (EDC)).

5.8.4 Free dwell time on terminal

The influence of the free dwell time at the container terminals was tested for 0, 2.5, 5, 7.5, and 10 days. The analysis showed that the free dwell time on container terminals has little influence on the overall cost level unless the free dwell time is less than the time required for handling and arranging onward transportation (in our case, less than two days). This is, however, unlikely to happen in practice. The differences in storage costs between 2.5 and 10 days of free dwell time are $9, $13, and $7 per TEU for the DEC, CEN, and EDC strategies; the differences between no free dwell time and 2.5 days of free dwell time are $38, $30, and $38 respectively.

5.8.5 40ft container

If the inland transport costs do not depend on the size of the container, then it would be attractive to use 40ft rather than 20ft containers. To analyze the impact of this change, we have run the base case configuration with all the settings adjusted for the use of 40ft containers (assuming that the shipping tariff for a 40ft container is twice that of a 20ft container). The data for this experiment in table 5.9 have been scaled back to TEU. As expected, the strategies with longer truck transport trips benefit most from this change. Looking at the detailed data (which is not listed in the table), we see that the handling and inland transport costs per TEU decrease whereas the holding and storage costs increase.

5.9 Central Location

Our initial calculation in section 5.4 indicated that Rotterdam and Antwerp have the most central location. To further investigate this, we have performed some additional experiments with Antwerp, Hamburg, Southampton, and La Spezia as the ports for the centralized strategy. Table 5.10 displays the outcomes. In line with our initial calculation, the total costs for Antwerp and Rotterdam are lowest; Antwerp has the lowest total costs due its more central location, in spite of the lower sailing frequency (once a week rather than twice a week for Rotterdam and Hamburg). The low frequency of sailings to La Spezia means that a higher level of safety stock has to be kept. Combined with a higher shipping rate, the total costs are therefore higher. The geographical position of La Spezia means that the order lead time is also significantly higher. Hamburg benefits from the higher frequency of sailings (the same as Rotterdam); the less central geographical location means that the inland transport costs are higher. The transport

Table 5.10: Comparison of Different Ports for Centralized Strategy

	Unit	Rotterdam	Antwerp	Hamburg	Southampton	La Spezia
S for central terminal	TEU	69	74	70	76	90
Avg. Lead Time	days	1.0	1.0	1.0	0.9	1.6
Avg. Handling Cost	$ per TEU	140	120	160	180	180
Avg. Holding Cost	$ per TEU	618	661	627	677	810
Avg. Shipping Cost	$ per TEU	1340	1340	1340	1340	1580
Avg. Storage Cost	$ per TEU	29	29	32	35	78
Avg. Inland Transport Cost	$ per TEU	980	912	1142	1199	1544
Avg. Total Cost	$ per TEU	3106	3062	3302	3430	4192

costs from Southampton are impacted by the cost of using the Euro Tunnel for transport to the European mainland.

5.10 Discussion and Conclusions

From our study, ports with a central location with respect to the hinterland in a region or on a continent enjoy a competitive advantage; when cargo is shipped to such a port, it can be redirected before arrival, when unloaded and stacked in the container terminal or when stored in a (European) distribution center. The value of this flexibility depends on the hinterland (where does the demand originate) and the value of the products. In this chapter we have looked at a sample supply chain in which Rotterdam is used as an example of a port that can offer this type of flexibility. We used a simulation model to quantify the value of the flexibility.

The port selection criteria that are discussed in the literature are rather abstract. We provide a more precise, quantified interpretation of criteria such as flexibility, location, shipping frequency, and charges. The results of our simulation model for the case show that the free dwell time of the container terminal does not have a large impact on the total cost provided the free time does allow sufficient time for the onward transport to be arranged. For the centralized strategy, the average dwell time is just below the terminal free dwell time. Although terminal operating companies might want to reduce the free dwell time in order to reduce yard congestion, they would thereby also en-danger the potential for the terminal to be used as a temporary storage location. If the yard is very congested, then a setup such as the Distripark Maasvlakte in Rotterdam (United Nations, 2002) could provide a solution: an off-terminal location that is linked to the terminal via a dedicated internal track. The transport costs to the Distripark can be significantly lower than truck transport to an external distribution center.

The flexibility that is offered by ports with a central location with regard to the hinterland enables pooling of safety stock. This flexibility is useful when there is variation in demand across the regions. With highly predictable demand, it would be more beneficial to keep stocks close to demand regions. However, with less predictable demand or with a high variation in demand across regions, pooling stocks at ports with a central location and a good transport network becomes more attractive. This pooling opportunity provides ports with a central location and a good hinterland transport network with a competitive edge. The regions with relatively low demand can then benefit from the safety stock that is also used for the regions with relatively high demand. The pooled demand reduces the average residence time of the stock; this in turn reduce storage cost and especially holding costs.

In the illustrative case study, we have looked at two variations of centralization. The centralized strategy uses a container terminal for temporary storage. Within the case setup, the dwell times are such that the storage costs remain low because we can take advantage of the free dwell time. This strategy has a slightly higher average cost per container than the decentralized strategy; however, as the value of the goods increases (and thus the holdings costs), the pooling advantage of the centralized strategy enables it to outperform the other two strategies. The EDC strategy has the lowest total cost per container and a shorter lead time than the centralized strategy. As demand variance was reduced, it became more attractive (cheaper) to hold more stock in the regional DC's and the role of the EDC was reduced. The storage costs for the EDC are such that it is a more attractive choice for stock with longer dwell times. The sensitivity analysis indicates that this strategy is attractive for moderately high holding costs; for very low holding costs, the decentralized strategy is preferred, and for very high holding costs, the centralized strategy performs best. Additionally, the EDC enables value-added logistics and less than full container shipments to regions with lower demand. These options are not available for the centralized strategy as the stock remains in the load unit (container) and they are less efficient if implemented in all the regional distribution centers (for the decentralized strategy).

In the illustrative case, we have looked at a limited number of ports and demand regions to enable a clear presentation of the results. Using the same methodology, we have also evaluated a larger case that includes 15 demand regions (the regions from the base case plus Denmark, the Czech Republic, France, Spain, Portugal, Switzerland, Poland, and Hungary) and nine ports (the ports from the base case plus Barcelona, Le Havre, Marseille, and Trieste). Initial analysis has shown that the results match the results of the base case.

The geographical layout of Western Europe provides a number of ports in the Hamburg–Le Havre range with a beneficial, central location that facilitates the centralization approaches included in our model. The East coast of the USA has somewhat similar characteristics; ports such as Savannah, Norfolk, Baltimore, and New Jersey serve an overlapping hinterland and most industrial areas in the Eastern USA can be reached by truck within three days. The addition of the new set of locks for the Panama canal (planned for 2014) which can handle larger and longer ships may cause a shift from using the West Coast ports with onwards transport by rail to the East Coast to using the East Coast ports. The carriers and their customers will then face a new port selection problem. Once the shipping tariffs for the new routes to the East Coast are known, customers could employ the model presented in this paper to evaluate their options.

An obvious extension would be to include barge and train transport. This would require a fairly detailed model of the hinterland transport network for these modes as well as accurate costing data. The location of a port in relation to these transport networks could, however, be an important factor in the overall flexibility of (re)routing traffic and could thus be worthwhile.

Acknowledgements

The authors want to thank prof.dr. Peter de Langen (Port Authority of Rotterdam) for his valuable suggestions and cost estimates. We thank Joost Hengstmengel (MSc student at Erasmus University Rotterdam) for his implementation of a simulation model that provided the basis for the simulation model used in this paper.

Chapter 6

Coordination In A Supply Chain For Bulk Chemicals*

A chemical plant in The Netherlands uses large annual supplies of a bulk chemical. A number of suppliers deliver their parcels from overseas by short sea vessel to a transshipment point where they are stored using a tank farm. Transportation from the transshipment point to the plant takes place by barge. Coordination of the schedules of vessels and barge provides the opportunity for board to board loading. Board to board loading provides clear benefits for the plants' operator, as it requires less handling and intermediate storage at the transshipment point. We demonstrate this by experiments conducted with a simulation model. The results are confirmed by analytical means.

6.1 Introduction

A supply chain can be characterized by four drivers, viz. facilities (such as plants, warehouses), transportation (different modes), (product) inventories and information (Chopra and Meindl, 2004). Bulk chemicals are characterized by high volumes of liquids, which are transported by pipelines, tank trucks, rail wagons, and ships: the inventories are stored in tanks. Several manufacturing stages exist, e.g., one producer makes bulk raw materials which are further processed by other producers whose products are in turn used by other manufacturers.

Proximity to users of bulk chemicals, such as car manufacturers, and cheap transport of supply are essential aspects in location decisions of chemical plants. Usually the bulk of supply and sales for a chemical plant occurs within a single continent, with occasional imports or exports to other continents.

* This chapter was previously published as de Swaan Arons et al. (2004).

As many chemical products have commodity characteristics, there is severe price competition. However, the number of suppliers and users is limited, and prices are not always listed on markets. Hence, from a strategic/cost perspective, it seems wise to use several suppliers instead of one. However, scheduling deliveries among multiple suppliers is hard, which negatively impacts the logistical process.

In this chapter, we consider the case of an inland bulk chemical plant in the Netherlands trying to improve its supply chain. The plants' domestic supply can be arranged by pipeline, train, truck or barge, depending on the proximity of the supplier. Alternatively, several suppliers may transport by ship to a port, where they transship their products into tanks using a third party tank farm. Inland transport from this transshipment point occurs either by pipeline, truck, train or barge. Tankage is available in flexible amounts at a third party tank farm in a port. This provides a buffer for any uncertainty in the supply coming from sea transport. Further inland transport is more or less controlled by the plant, which allows it to limit the amount of dedicated tankage at its site.

The plant decided on a transition of its inland transport mode from truck to barge. Barge transport has several advantages: it is cheaper, safer and does not create dangerous transports through populated areas. Barge transport, however, causes larger lots to be transported, requires investment in jetties and pipelines, and is restricted to waterways only.

Barges sail back and forth from the plant to a transshipment point where short sea vessels deliver their parcels from overseas. When the company decided to make the move to supply by barge it also wanted to assess the performance of its supply chain, limit the number of suppliers somewhat to reduce the effort of managing the supply chain, and increase the probability of direct-on transport by so-called board to board loading from sea ship to inland barge. This is cheaper as it bypasses the tank farm. For reasons of presentation, some of the details of the case presented in this chapter have been omitted. The actual plant uses several bulk chemicals. Because the products are used in fixed proportions and are delivered simultaneously in separate tanks on the same vessels and barges, we will not distinguish the individual chemicals but rather refer to this product mix as a 'bulk chemical'.

In the next section relevant literature on this subject is reviewed. In section 6.3 a conceptual model of a supply chain for bulk chemicals is given. This model is applied to a case study in section 6.4. A simulation model based on this case study is outlined in section 6.5. For this chapter, the simulation model is then used in section 6.6 to assess the probability of board to board loading and the effect of tighter delivery scheduling on this probability. The results of the simulation experiments are compared with analytical calculations. In section 6.7 we draw conclusions.

6.2 Literature Review

The topic of logistics planning in bulk chemicals has received little attention in scientific literature. Literature dealing with inventory and supply chain management tends to focus on a supplier's perspective, where one supplier has to move products to many customers. In the current situation, however, one customer (the chemical plant) receives products from several suppliers. In order to minimize logistic costs and increase reliability, it strives for increased control over its suppliers. Therefore, this supply chain has to be looked at from a customer's point of view.

Nieboer and Dekker (1995) discuss a model for tankage assessment (i.e., determining how much storage capacity is needed) and stock control in refineries. They consider a tank that is continuously fed by a production unit and two types of demand. First, there is a demand for large parcels, which is planned in advance. Second, there is demand for small parcels, modeled by a Poisson process. The difference between the inflow from production and the outflow from small demands is modeled as a Brownian motion. Silver and Peterson (1985) present a decision rule that can be applied to calculate the required safety stock, given a certain probability of stock out during a replenishment cycle. However, this rule is based on variability in demand, not in supply. Newhart et al. (1993) discuss a simple method for incorporating the variability in supply within the variability of demand. However, this method assumes normality of demand and lead times, which is not realistic. Haehling von Lanzenauer et al. (1992) use a stochastic process approach to calculate the probability of insufficient supply (stock out) of natural gas, which is essentially risk analysis. In this case, the authors use the concept of a Design Day; a date for which extreme demand is assumed. van Asperen et al. (2003) describe the role of arrival processes in a port simulation and demonstrate the impact of increased coordination in terms of ship waiting times and required storage capacity.

All in all, these references give us little help in tackling the problem. Other relevant results may be found in the supply chain literature. Sometimes it is advocated to reduce the number of suppliers in order to improve coordination in the supply chain (e.g. shorter response time, collaborative planning and forecasting) or at least to carefully consider the number of suppliers needed in the presence of risks (Berger et al., 2004). In the present case, using a supplier that is located in close proximity to the plant would greatly simplify the logistics. Yet, the product under consideration is a commodity for which price competition exist. Multiple suppliers are used in order to get the lowest price and to keep several supply chain options open. The advantages of having multiple suppliers have to outweigh the associated logistical problems. How to achieve this has received little attention in literature.

6.3 A Conceptual Model

In this section a model is presented in order to help understand the dynamics of a simple supply chain for bulk chemicals.

A chemical plant (from now on denoted as the plant) uses large annual supplies of a bulk chemical. A number of suppliers deliver their parcels by short sea vessel to a transshipment point (TSP). Here, the chemical is stored in tank that are part of a third-party tank farm. The plant rents a prespecified number of tanks from the tank farm operator. Transportation from the TSP to the plant can be performed using either trucks or barges.

Since disruptions in the plants' production process are very expensive, buffer tank capacity at the plant site is required for sustained production and tolerance towards variations in supply.

Further transport to the plant used to be carried out by trucks but for a variety of reasons such as cost and safety aspects, as well as uncertainty caused by traffic congestion, the plant has moved to using a single dedicated river barge instead. The supply chain is illustrated in figure 6.1

At scheduled times short sea vessels deliver parcels at the TSP which will be unloaded in the shore tanks. The tanks must have an appropriate capacity in order to deal with uncertainty regarding the arrivals of the vessels. Weather influences, but also a lack of coordination between the supply chain partners make the actual time of arrival (ATA) of vessels differ from their expected time of arrival (ETA). An example distribution function of the deviation in hours from the ETA is the following. If the deviation in hours is denoted as x, then:

$$x = U(-120, -48) \quad \text{with } p = 0.1 \tag{6.1}$$

$$U(-48, 48) \quad \text{with } p = 0.8 \tag{6.2}$$

$$U(48, 120) \quad \text{with } p = 0.1 \tag{6.3}$$

where U is the uniform distribution function. This means that 80% of the vessels arrive within plus or minus two days of their ETA; within this interval the probability density is constant. 10% of the vessels arrive between two or five days later and another 10%

arrive between two or five days earlier than the ETA, in both cases also with constant probability density.

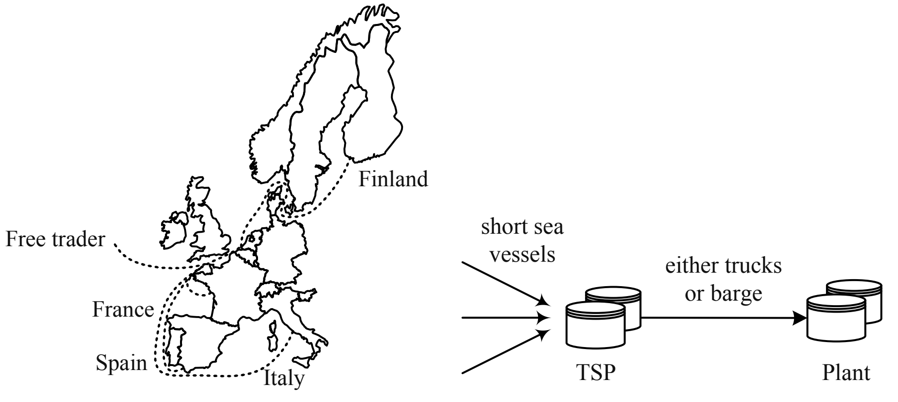

Figure 6.1: Supply Chain Overview.

Figure 6.2 depicts the graph of this discrete distribution function. The deviations to the ETA cause the level in the shore tanks at the TSP to vary stochastically. For this reason, a safety stock level is maintained ensuring continued supply to the plant.

Below, the two modes for the onward transport from the TSP to the plant are discussed: road transport using trucks, and water transport using a barge.

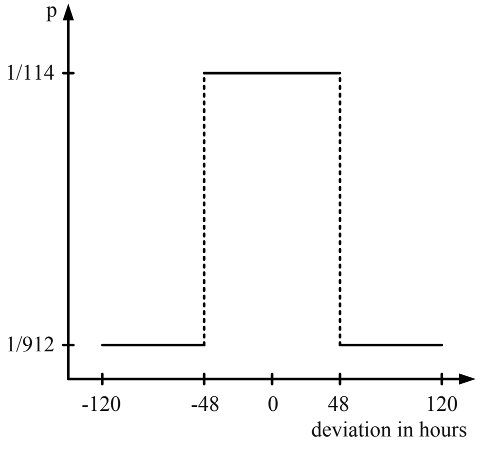

Figure 6.2: Distribution Function of the Deviation from the ETA of Short Sea Vessels.

6.3.1 Transport by Trucks

This supply process is characterized by a large number of small deliveries, as the constant daily intake of the plant is tens of times larger than the capacity of a single truck. The trucks pick up their cargo from the shore tanks at the TSP and shuttle between the plant and the TSP. This provides a steady stream of equidistantly spaced deliveries.

In this scenario, the tank farm at the TSP provides the primary buffer against disruptions in the supply by sea-going vessels. The buffer capacity at the plant is designed to accommodate small disruptions to the supply process from the TSP to the plant. The required tankage at the plant is therefore quite small.

6.3.2 Transport by Barge

A dedicated river barge carrying the same parcel size as the sea-going vessels provides a shuttle service from the TSP to the plant. This supply process features a smaller number of larger deliveries.

Using a river barge with the same parcel size as the sea-going vessels offers the opportunity to load the cargo directly from a vessel into the barge. This board to board loading is discussed in section 6.3.3. If board to board loading is not possible, the vessel will unload into, and the barge will load from the shore tanks.

The deliveries at the plant are stored in tanks. The cycles are stock-controlled (van Asperen et al., 2003): the barge is scheduled to arrive at the plant when the stock in the tanks has reached the safety stock level.

The barge sails back and forth from the plant to the TSP. A cycle starts and ends at the plant when the barge arrives at the plant just before unloading. We have defined the cycle in such a way that the opportunity for board to board loading is maximized (the details of the cycle are described in the next section). Thus, both vessel and barge are scheduled to meet at the TSP somewhere in the middle of the cycle (the exact time depends on case-specific parameters). By maximizing the opportunity for board to board loading, the plant can reduce the tank capacity it rents at the tank farm and thus reduce costs.

Note that the return trip from plant to TSP is an empty run. In reality, the barge will sometimes use the idle time to transport cargo for another company.

6.3.3 Board to Board Loading

Board to board loading takes place if the short sea vessel and the river barge meet at the transshipment point. From the plant operator's point of view, this board to board loading is faster (no intermediate unloading and loading) and cheaper (less handling and storage capacity at the TSP is required). The supplier does not have a direct interest in board to board loading and he will therefore be willing to wait for a barge to arrive for a very limited time. Board to board loading requires close coordination among the partners in the supply chain to make sure that the ships meet at the appropriate time.

The probability of board to board loading can be defined as the fraction of the arrivals of the river barge at the transshipment point that result in board to board loading. It depends on the length of the time window in which the barge is at the transshipment point. Figure 6.3 illustrates this time window as a part of the cycle of the river barge.

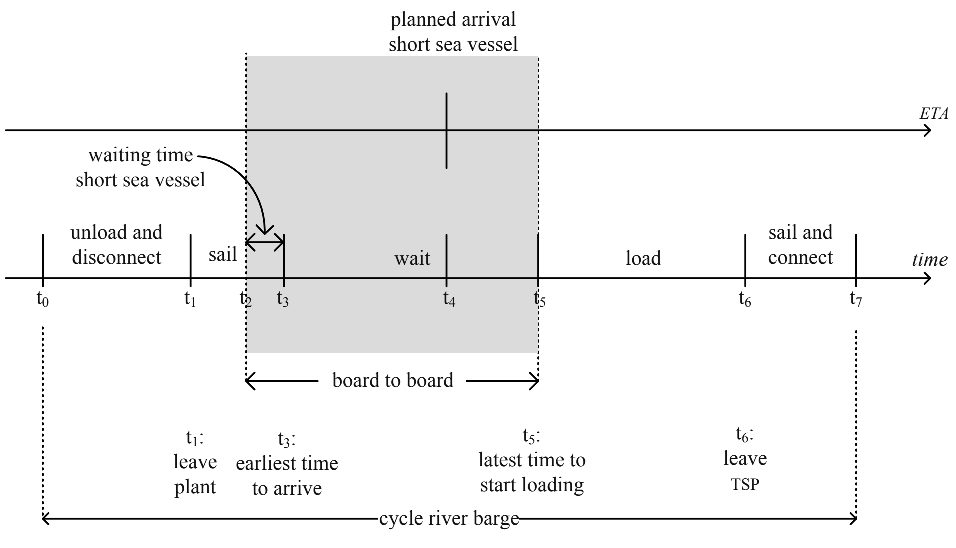

Figure 6.3: Time Window for Board to Board Loading.

The arrivals of the short sea vessels are scheduled somewhere in the middle of the cycle of the river barge. At that moment the barge is waiting at the transshipment point. The time window during which board to board loading is possible, can be determined as follows (see also figure 6.3):

cycle = unloading, disconnecting at plant,
+ sailing to TSP + waiting time at TSP
+ connecting, loading, disconnecting at TSP
+ sailing to and connecting at plant

The planned arrival time for the short sea vessel is t_4. If the vessel arrives between t_2 and t_5, then board to board loading is possible. If the vessel arrives before t_2, the maximum waiting time for the vessel ($t_3 - t_2$) will make board to board loading impossible; if the vessel arrives after t_5, then the loading ($t_6 - t_5$) and sailing time ($t_7 - t_6$) for the barge will force it to load from the shore tanks in order to reach the plant in time (i.e., before the stock at the plant drops below the safety stock level).

6.4 A Case Study

The chemical plant in this case study is located in the south-western region of The Netherlands. It is located about 50 km from Antwerp. Existing waterways connect the plant to the port of Antwerp. Roads to the plant site run through a densely populated area. The plant uses 219 kiloton of a bulk chemical annually. There are five suppliers for reasons of price sensitivity and competition: table 6.1 specifies the annual volumes of the contracts per supplier. Suppliers S_1 to S_4 are located at various remote sites in Europe and deliver their parcels by short sea vessel to the TSP at Antwerp. A dedicated river barge transports the bulk chemical to the plant. Short sea vessels and the barge carry 1,450 tons of bulk chemicals. A local supplier S_L is located near the TSP at Antwerp and loads its parcels directly into the barge.

Table 6.1: Annual Volume of Contracts per Supplier.

Supplier	Bulk Chemical (10^3 metric tons)	Rounded % of total
S_1	14	6
S_2	38	17
S_3	22	10
S_4	50	23
S_L	95	43
Total	219	100

The (un)loading rate at the TSP is 150 ton/hr. The barge shuttles between the plant and either the TSP or the supplier S_L. This takes four hours in each direction. The geographical layout is illustrated in figure 6.4.

If the barge can not load at the TSP (board to board is not possible and the shore tanks have insufficient stock), then the barge can pick up an emergency delivery at S_L (this is part of the contract with this supplier). Sufficient stock is always available at S_L. Sometimes, in case of a possible overflow of the tanks at the TSP, the barge will skip a scheduled visit to S_L.

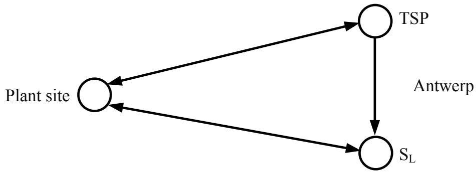

Figure 6.4: River Barge Sails between the Plant, and the TSP and Supplier S_L.

The deliveries at the plant are stored in tanks with a maximum capacity of 2,700 ton; the unloading rate is 200 ton/hr. The plant uses these tanks for a constant daily intake of 620 ton. A safety stock level of 1,240 ton must be maintained.

The barge sails from the plant to Antwerp (either to TSP or S_L) and back. A cycle starts and ends when the barge arrives at the plant just before unloading. Given the plants' annual need for the bulk chemical as outlined in table 6.1 and the transport capacity of the barge, the plant needs to be supplied three times a week. Consequently, the barge cycle has a length of 56 hours. A possible schedule is depicted in figure 6.5. The ratio of visits to the TSP and S_L (5:4) corresponds to the percentages in table 6.1 (56% and 44%, respectively).

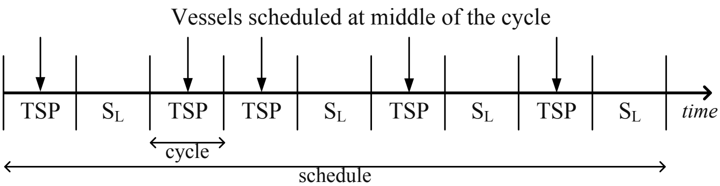

Figure 6.5: A Schedule of River Barge and Short Sea Vessels.

If an incoming vessel does not meet a waiting barge, then the vessel will wait up to two hours for the barge to arrive. To maximize the opportunities for board to board loading, the barge will wait for a vessel until the latest possible moment. The barge does not have to wait for the actual arrival of the vessel: the ATA of vessels is known six hours beforehand.

Figure 6.6 provides an overview of vessel and barge operations. It shows the three locations (plant, TSP and S_L), and can be seen as divided into two parts. The left hand side describes the arrival of the vessel at the TSP whereas the right hand side shows how the barge sails either to the TSP or to S_L. In this flowchart, the vessel and the barge meet at the box Connect to vessel if the vessel is ready or will arrive soon. If possible, the vessel and barge will perform the board to board loading and disconnect: upon completion, the vessel leaves the system and the barge sails to the plant.

6.5 A Simulation Model

In section 6.3 we used a simplified model describing a base scenario of a supply chain for bulk chemicals. When applied to the case described in the previous section, it allows us to calculate the probability of board to board loading. This probability can only be determined analytically if a number of simplifying assumptions hold (see below). In reality, the supply chain is much more complicated and answers are hard to get analytically if these assumptions are dropped. For this reason a simulation model was developed which is outlined in section 6.5.1.

One simplifying assumption concerns the sailing times of the barge between the various locations. They were previously assumed to be constant (four hours) but this is not plausible. All kinds of delays can occur (e.g., caused by locks and weather) which affect the behavior of the supply chain. Similar arguments hold with respect to a constant intake by the plant. Plant data show that quite regularly the intake is much less, sometimes up to 20%. A varying intake can be caused by rejected deliveries or failures. The impact of these kinds of events cannot easily be calculated analytically although it certainly affects the performance of the supply chain.

6.5.1 Implementation Model

Based on what is outlined in the previous section, a simulation model has been implemented in Enterprise Dynamics (InControl Enterprise Dynamics, 2003), a simulation package for discrete-event simulation. Simulation environments such as Enterprise Dynamics and Arena are generally easy to use, and allow for quick model construction. They provide built-in animation, generate statistics, and form well-tested simulation environments. The implementation model comprises various types of atoms, the Enterprise Dynamics equivalents of objects. Some of the atoms implement the simulation's logic, others hold the simulation data (tables), define the types of experiments or provide the desired output (e.g., graphs).

The scripting language of Enterprise Dynamics and the ability to open and close connections (known as channels in ED) between elements of the model were used in this implementation. The implementation of the rendezvous of vessels and barge with separate time-windows for both agents posed a particular challenge.

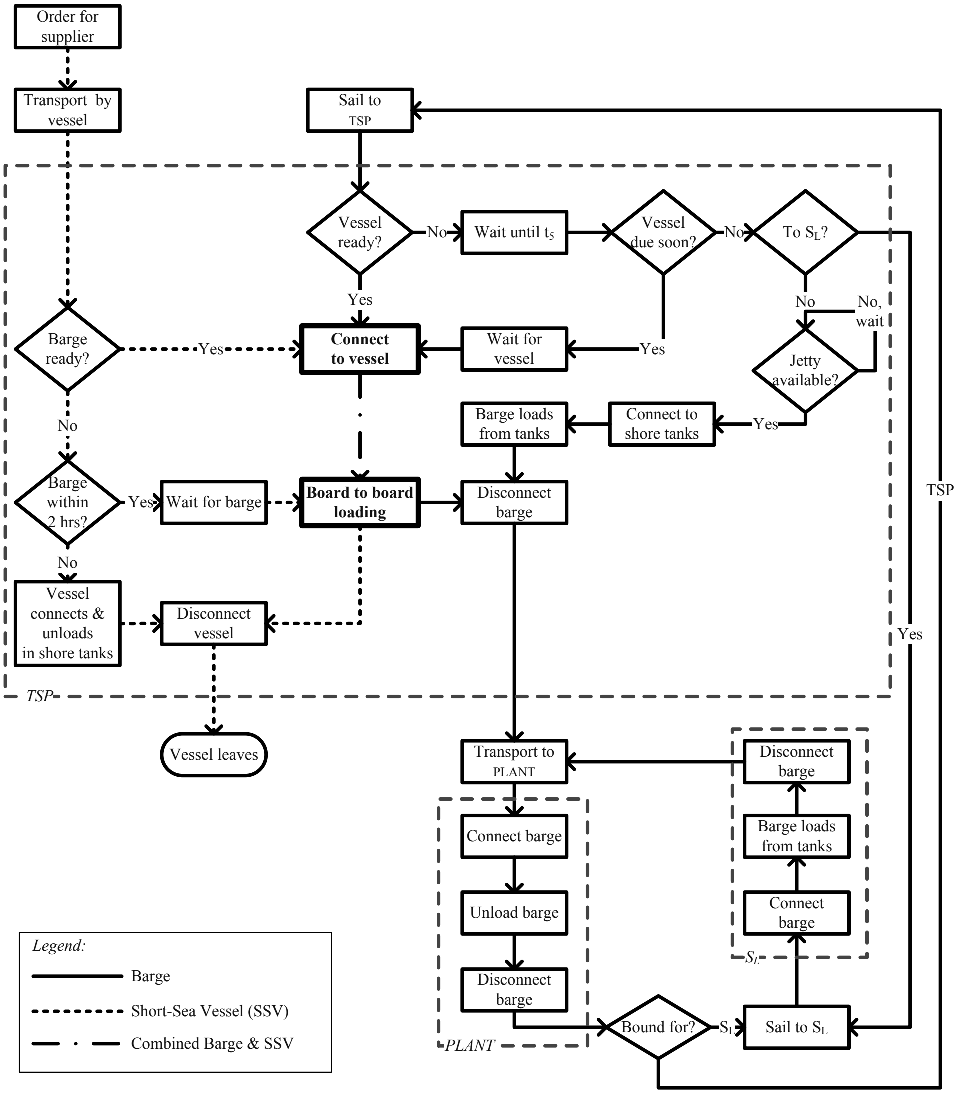

Figure 6.6: Flow Chart of Barge Operations.

6.6 Experiments and Results

We expect the efficiency of the logistics process to be sensitive to more intensive coordination by the plant. Improved coordination could for example reduce the center section of the discrete distribution described in section 6.3, i.e. the (-48,+48) hours interval. As a consequence, one would expect the number of stock outs or overflows to decrease and the probability of board to board loading to increase.

Better flow management could also increase the performance of the barge. The barge as discussed in section 6.3.2 faces considerable idle time. At the TSP it may be waiting for many hours for a vessel that may not even show up in time. Improved coordination between partners in the supply chain could prevent such waiting times, enabling the barge to be deployed for other tasks.

In this chapter, we focus on the effect of the improved coordination on the probability of board to board loading. To this end, we have performed a number of experiments with different time windows for the center section of the discrete distribution, using both an analytical and a simulation-based approach. For the analytical approach, we maintain the simplifying assumptions mentioned in section 6.5. The simulation-based approach uses stochastic sailing times from plant to TSP and S_L, but maintains the assumption of a constant intake at the plant. All experiments were conducted using a simulation run of twenty years to obtain statistically reliable results. Table 6.2 displays the results of both approaches. Here, the number N represents the maximum deviation to the ETA of arrivals in the 80% section of the discrete distribution function. As can be observed, the results of the analytical approach closely match the outcomes of the simulation experiments. Figure 6.7 shows the board to board probability results of the simulation experiments and the results of the analytical calculations in a graph.

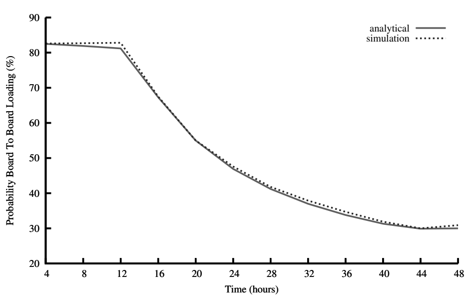

Figure 6.7: Board to Board Probability from Simulation Experiments and Analytical Calculations.

Table 6.2: Results of Analytical and Simulation Approach.

	Analytical Results	Simulation Results	
N	Board to board %	Board to board %	Std.dev.
4 hr	82.5	82.6	0.8
8 hr	81.9	82.7	0.8
12 hr	81.2	82.8	0.8
16 hr	67.3	67.5	1.1
20 hr	55	55	0.9
24 hr	46.9	47.6	0.9
28 hr	41.2	41.8	1
32 hr	37	37.9	0.9
36 hr	33.8	34.7	0.9
40 hr	31.3	31.9	1.2
44 hr	29.9	30	0.9
48 hr	30	30.9	1.2

Clearly, coordination efforts by the plant reducing the deviations to the ETA of the bulk (80 percent) of the vessels, positively impact the board to board percentage. However, below 12 hours, further reduction is pointless: the board to board percentage remains the same. This is due to the fact that if a vessel arrives within 12 hours of its ETA, it automatically (i.e., given the barge's schedule) arrives within the interval that the barge is ready and waiting for board to board loading. Since, according to the discrete distribution, this goes for at least 80 percent of the vessels (the other 20 percent is still somewhere between 120 hours early or 120 hours late), the board to board percentage should be at least 80 percent as well.

Another observation that can be made from the graph, is that when N increases beyond the 40-hour limit, the board to board probability does not further deteriorate. This is due to the fact that at this point, many vessels start to arrive very early or late, to the extent that they will arrive within the board to board window of the previous or next barge cycle, thus still enabling board to board loading.

6.7 Conclusions

The results of the experiments conducted in this chapter clearly demonstrate the beneficial influence of improved coordination on the logistics of a supply chain in the bulk chemical sector. The simulation outcomes were confirmed by analytical calculations based on a number of simplifying assumptions.

Chapter 7

Arrival Processes in Port Modeling*

This chapter investigates the impact of arrival processes on the ship handling process. Two types of arrival processes are considered: controlled and uncontrolled. Simulation results show that uncontrolled arrivals of ships perform worst in terms of both ship delays and required storage capacity. Stock-controlled arrivals perform best with regard to large vessel delays and storage capacity. The combination of stock-controlled arrivals for large vessels and equidistant arrivals for barges also performs better than the uncontrolled process. Careful allocation of ships to the mooring points of a jetty further improves the efficiency.

7.1 Introduction

In this chapter we investigate the impact of ship arrival processes and jetty allocation schemes on the efficiency of the loading and unloading process in a port simulation. An arrival process is a formal specification of how entity arrivals in a system are scheduled. In our case, it determines, among others, the likelihood of several ships arriving simultaneously, which is an important aspect in, for example, determining the required jetty capacity. Our research was triggered by work done on a confidential case study with the objective to help determine the optimal layout of the jetty owned by a new chemical plant in the port of Rotterdam. The original tender of that case study provided detailed data on the types and numbers of ships to be handled annually, but failed to specify their arrival process. However, an initial simulation model described by van Asperen et al. (2003) demonstrated a considerable impact of the type of arrival process on system performance, in terms of both waiting times of ships waiting to load or unload at the jetty, and stock fluctuations in the tanks on the chemical

* This chapter was first published as van Asperen et al. (2005).

plants' facilities. In this chapter, we further develop and analyze the arrival processes themselves, evaluate their impact on system performance, and evaluate several jetty allocation schemes for additional performance enhancement.

A basic distinction can be made between uncontrolled and controlled arrival processes. Uncontrolled arrivals are typically modeled by a Poisson process, a common assumption, for example, in modeling incoming telephone calls in call center simulations. Controlled arrivals concern scheduled arrivals, such as scheduled airline flight arrivals to an airport (Banks et al., 2000). For our port system we distinguish two types of controlled arrivals. The first type are the so-called stock-controlled arrivals, i.e., ship arrivals are scheduled in such a way, that a base stock level is maintained in the plants' tanks. The second type is based on equidistant arrivals per ship type and relates to contracts prescribing product supply and pick-up at regular time intervals, e.g., once a month. We compare model outcomes based on four different arrival processes: uncontrolled, stock-controlled, equidistant and a blend of stock controlled arrivals for the larger ships and equidistant arrivals for the smaller ones. Furthermore, for all four types of arrival process, it is investigated to what extent careful allocation of ships to the jetty's mooring points enhances system performance.

Apart from some very scattered material, little practice with the simulation of port facilities can be drawn from existing literature. van Nunen and Verspui (1999) provide insight in simulation and logistics in ports, but it is in Dutch only. Here, we briefly recapitulate the literature review on jetty design from Dekker (1999) in that volume. Well-known to insiders are the reports from UNCTAD (1978) on the design of jetties. They report results from both queuing theory and simulation applied in studies on jetty capacity. However, the reports are difficult to obtain and they give yardsticks for simple cases only. Other papers more or less describe particular simulation studies, without trying to generalize their results: Philips (1976) and Andrews et al. (1996) describe the planning of a crude-oil terminal; Baunach et al. (1985) deal with a coal terminal; Heyden and Ottjes (1985), Ottjes (1992) and Ottjes et al. (1994) deal with the set-up of the simulation programs for terminals. None of these papers however, deals explicitly with arrival processes. Kia et al. (2002) do mention the arrival process in the context of a port simulation: they assume a Poisson process.

In section 7.2 we provide a detailed description of the conceptual model of the system. In section 7.3 the various types of arrival processes are discussed in detail. Three schemes for the allocation of ships to the jetty's mooring points are given in section 7.4. Section 7.5 provides a brief discussion of how the simulation models have been implemented. The experiments conducted with the model and their results are discussed in section 7.6, and the conclusions are presented in section 7.7.

7.2 The Conceptual Model

The system considered in this chapter involves a chemical plant with a continuous production process. Both the supply of raw materials and the export of finished products occur through ships loading and unloading at a plant-owned jetty. Since disruptions in the plants' production process are very expensive, buffer tank capacity is required for sustained production and tolerance towards variations in ship arrivals and overseas exports through large ships. With respect to the original case study, some simplifications apply. For reasons of confidentiality, the diversity of ships has been skewed down, and their numbers modified. Also, details concerning tank operation, tank farm layout, and inland transport have been abstracted from. Still, the resulting model is general enough to draw conclusions applicable to many jetty simulation studies.

Operational costs of such a facility increase when ships have to wait to (un)load, or amplitudes in stock level fluctuations widen (tankage is costly as well). Causing factors of such events include the shape of the jetty, the number of mooring points it has and their restrictions with respect to the types of ships and cargo they can handle, and whether the port is an open port or has locks. Other possibly relevant factors - and that is the key subject of this chapter - are the arrival processes of the various types of incoming ships and the allocation of ships to the jetty's mooring points. Figure 7.1 provides a schematic outline of the model as a whole. Apart from the arrivals of ships, it comprises a jetty with a number of mooring points, several storage tanks and a chemical plant, which are described in sequence below.

7.2.1 The Jetty

The jetty provides four mooring points (numbered 1 to 4) in a T-shaped layout (figure 7.2). Ships arriving at the jetty to load or unload cargo dock at one of these. Mooring points 1 and 2 are suited to handle ships of all sizes; mooring points 3 and 4 can handle only short ships (see also Table 7.1).

Incoming ships unload raw materials (A or B), or load finished products (C or D). Pipes facilitate the transport of all chemicals to and from the ships. Since cost considerations are a limiting factor on their construction, not every type of raw material and finished product can be (un-)loaded at every mooring point. For example, mooring point 1 can handle A, B, and C, whereas mooring point 2 can only handle products C and D.

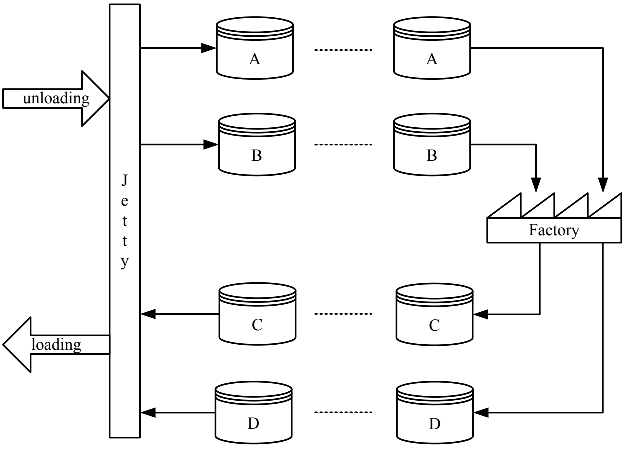

Figure 7.1: A schematic outline of the loading and unloading process, including jetty, tanks and plant.

7.2.2 Tanks and Stocks

After unloading, raw materials are stored in tanks A and B, for later extraction and processing by the plant. Finished products are transferred to tanks C and D, to be loaded into ships. Tanks can be used for just one type of raw material or finished product. The transfer of products from ships into tanks, from tanks to the plant, and from the plant into the tanks are continuous processes, which, in reality, are subject to several restrictions. One restriction prescribes that there shall be no simultaneous pumping and running into and out of a tank. Another restriction is that stocks are limited due to finite storage capacity available. However, for simplicity we allow them

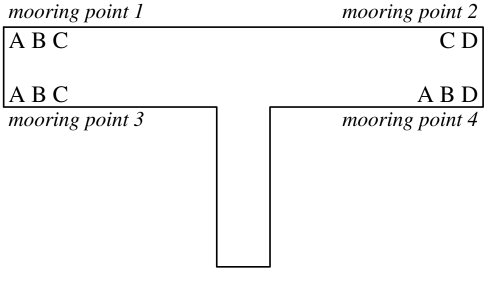

Figure 7.2:]
The jetty layout.

to take on any value, and neglect ship delays because of stock outs or lack of ullage (available tank space). We ignore all these restrictions, because they do not affect the comparison between the arrival processes.

7.2.3 Ships

Ships (ocean-going vessels, short-sea vessels, and inland barges) unload raw materials or load finished products. Each ship has five defining properties relevant to our model:

- size (tonnage);
- length (a distinction between long and short suffices);
- product (each ship handles just one specific type of cargo);
- (un)loading time (in hours);
- priority (a distinction between high and low suffices).

When a ship has arrived in the port, a suitable mooring point is selected according to a set of rules, which are discussed below. Table 7.1 shows all types of ships loading and unloading at the jetty along with their values for the aforementioned properties. For example, every year, a total of fourteen short vessels arrive carrying 4,000 tons of product B, with a loading time of 26 hours. Columns "Ships per year", "Priority", and "Tons per year" are discussed in more detail later.

7.3 The Arrival Process

In many simulation studies it is assumed that arrivals in client-oriented processes cannot be controlled. Simulation languages and environments acknowledge this and tend to offer Poisson as a first-choice option for the specification of arrival processes. However, in some port situations, a definite measure of control over the arrival process can be observed. This suggests that care should be taken in settling on a process to feed the simulation with ship arrivals.

In order to understand how such considerations affect our simulation study, one should first analyze the plant's planning process and organizational structure. Usually, every month or few months, depending on the company, the sales/marketing department sets up tactical sales plans, including contract sales over a long period, new contract sales and spot sales. In order to see whether the production required

Table 7.1: Ship types, properties, and arrival rates

Number	Type	Size (metrics tons)	Length	Product	Loading time (hrs)	Ships per year	Priority	Tons per year
1	barge	1,500	short	A	8	196	low	294,000
2	vessel	2,000	short	A	8	48	low	96,000
3	vessel	4,000	short	A	20	80	low	320,000
4	vessel	6,000	long	A	26	60	high	360,000
								1,070,000
5	barge	1,000	short	B	10	38	low	38,000
6	vessel	2,000	short	B	11	161	low	322,000
7	vessel	4,000	short	B	26	14	low	56,000
8	vessel	6,000	short	B	26	12	low	72,000
								488,000
9	barge	1,000	short	C	10	180	low	180,000
10	vessel	2,000	long	C	14	126	high	252,000
								432,000
11	barge	1,500	short	D	8	134	low	201,000
12	vessel	2,000	short	D	8	300	low	600,000
13	vessel	10,000	long	D	44	14	high	140,000
14	vessel	20,000	long	D	56	8	high	160,000
								1,101,000

for sales fulfillment can be achieved, possible bottlenecks in the production process need to be identified. Sales plans and bottleneck analysis together constitute the primary building blocks for a tactical production/sales plan. For our chemical plant this plan ultimately determines the required production level for the coming period, and provides direction for the logistics department to plan order pickups and deliveries.

However, many long-term contracts in the bulk oil and chemical sector, while including detailed price specifications (to avoid uncertainties as a consequence of market fluctuations), are considerably less rigid about the exact delivery dates. It is up to the waterfront part of the logistics department to agree with clients and suppliers on pickup and delivery schedules. Furthermore, short term sales and purchases require additional planning effort, since these often depend on ad hoc opportunities as short term traders tend to focus on prices, disregarding logistical feasibility. The logistics department is now faced with the challenge of accommodating this type of deals as well.

Finally, it should be noted, that during the design of a new plant, it is often unclear to both the logistics department and the construction engineers what purchasing/sales contracts will be used by the marketing/sales department in the future, and in what the ratio of short term deals and long term contracts will be.

In designing a simulation model for such a logistical process, one cannot but make some assumptions about the level of control that the logistics department maintains over ship arrivals. Several possibilities for modeling such control (or lack thereof) are described below. Their impact on simulation outcomes is this chapter's main subject.

7.3.1 Expected Times of Arrival (ETA)

Modeling control over ship arrivals involves the notions of Expected Times of Arrival (ETAs) and Actual Times of Arrival (ATAs). Here, the time of arrival is the time at which a ship arrives before the jetty. Let us start with the ETA. We consider two major types of controlled arrival processes yielding ETAs: stock-controlled and equidistant arrivals, and a third type, hybrid, which is a blend of of these two types.

7.3.2 Stock-controlled Arrivals

The plant management's aim is to achieve efficient production, avoiding costly interruptions such as those caused by stock-outs in the raw materials tanks. Further efficiency can be attained through prevention of stock-outs in the finished products tanks. These would cause ships to have to wait around for cargo, which is also costly. In case ship arrivals can be planned by plant management, stock-controlled arrivals can be used to maintain a target base stock level in the tanks as a buffer for production (raw materials) and transport (finished products). In our model, this is implemented as follows. For the loading process, it implies that the arrival time of the next ship is planned to coincide with the moment that, through production, there is sufficient stock in the tank to load the ship without dropping below base stock level. In this calculation, the parameters are the loading time of the present ship, the cargo capacity and loading time of the next ship, and the production capacity of the plant. Setting the appropriate base stock level for a tank involves an estimation of the tendency of ships to arrive ahead of schedule (see below), this being the only threat to maintaining base stock level.

For the unloading process, maintaining base stock levels in the raw materials tanks is achieved by planning the next ship's arrival to coincide with the moment that, through extraction of raw material during production, base stock level will be reached.

In this calculation, the parameters are the cargo capacity of the present ship, and the rate at which the plant extracts material from the tank. Here, the danger of stock dropping below base stock level comes from ships arriving late (or from ships unable to instantly find an unoccupied mooring point).

To illustrate the above, Figure 7.3 shows stock level fluctuations in raw material tank A over time with stock-controlled arrivals. At time t_1, when the tank contents is at base stock level, a 1,000 ton barge arrives, unloading its cargo into the tank over an 8 hour period. This implies that 8 hours later, the tank will contain an extra 1,000 tons of raw material, minus the volume of raw material pumped out of the tank by the plant. After this point, the tank's contents will steadily decrease back to base stock level. The next ship's arrival is planned to coincide with this moment t_{2p} ('p' for 'planned'). However, this ship could arrive ahead of time (see section 7.3.5), for example at t_{2a} ('a' for 'actual'), causing stock to start rising again before reaching base level. The dashed line shows how stock level would develop if all ships arrived exactly as planned. The solid line shows actual stock level development. After the last ship's early arrival, the next ship is again scheduled to arrive when stock reaches base level (t_{3p}). However, it arrives late at time t_{3a}, causing stock to drop below base level.

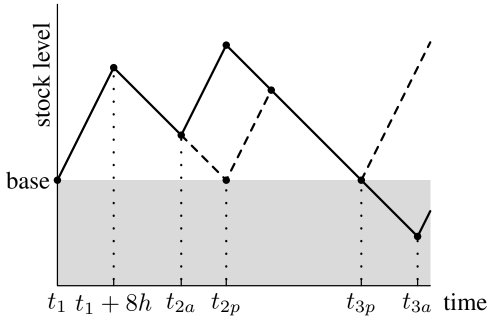

Figure 7.3: Stock level fluctuations in raw material tank with stock-controlled arrivals

7.3.3 Equidistant Arrivals

Equidistant arrivals model situations in which loading and unloading ships arrive at regular intervals. This regularity could, for example, be the consequence of year-based contracts specifying annual amounts of raw product to be delivered in equal batches every n weeks.

In our model, equidistant arrivals imply that arrivals of ships within the same ship type are assumed to be evenly spread over the year. For example, per year, twelve vessels carrying 6,000 ton of product B arrive (see Table 7.1). With equidistant arrivals, this means a 1-month inter-arrival period between such ships. Note that ships from different ship types may still arrive simultaneously.

7.3.4 Hybrid Arrivals

In a hybrid arrival process, the total population of ships is partitioned along some criterion, after which each type is assigned an arrival process for scheduling the arrivals of its members. In this chapter, we consider one hybrid process, in which the smaller ships (below 6,000 tons) arrive equidistantly, whereas the larger ones are subject to stock-controlled scheduling. The arrivals of all larger ships are scheduled on a per-product basis whereas the smaller ships are scheduled per ship type.

The underlying assumption is that contracts with clients and suppliers are such that alignment of the corresponding shipments with the production process is, in principle, hard. Hence, the majority of deals results in equidistant pick-ups and deliveries, partly due to transportation-related clauses in the contracts, and partly due to a client/supplier (especially those transporting many smaller shipments) preference for regularity in their logistical processes. Under such circumstances, the logistics department's focus will be on aligning the larger shipments with the production process. This is feasible for two reasons. First of all the number of large shipments is limited. Second, the plant operator and clients and suppliers requiring large shipments have a shared interest in coordinating ship arrivals and thus reducing waiting times. From the plants' point of view, large shipments are most likely to cause stock-outs or lack of available tankage, and from the client/supplier's point of view, avoiding delays for their large ships pays off (waiting by large ships is relatively costly).

Obviously, when simulating with this hybrid arrival process one implicitly assumes that stock-controlling large ship arrivals is feasible.

7.3.5 Actual Times of Arrival (ATA)

In reality ships will seldom exactly meet the schedule as defined by the ETAs. Most ships arrive within a relatively short interval around their expected time of arrival, while some arrive significantly earlier or later. Such deviations are modeled by a disturbance to the ETA. An ETA together with a disturbance yields the actual time of arrival (ATA) of a ship. The parameters of the disturbances were set together with shipping experts, taking into account the fact that the Port of Rotterdam is an open port, with relatively stable weather conditions.

The distribution function of the deviation in hours from the ETA can be described as follows. If the deviation in hours is denoted as x, then:

$$x = U(-12, -2) \quad \text{with } p = 0.1 \tag{7.1}$$

$$= U(-2, 2) \quad \text{with } p = 0.8 \tag{7.2}$$

$$= U(2, 12) \quad \text{with } p = 0.1 \tag{7.3}$$

where U is the uniform distribution function. This means that all ATAs are within a margin of twelve hours before and twelve hours after the corresponding ETA. Eighty percent of these are within a margin of two hours before and two hours after the corresponding ETA, in all cases with constant probability density (see figure 7.4).

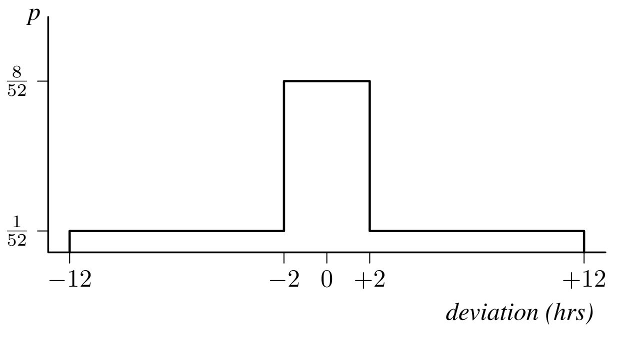

Figure 7.4: Distribution of disturbances to expected times of arrival

7.3.6 Uncontrolled Arrivals

The assumption underlying uncontrolled arrivals is that — in contrast to both stock-controlled and equidistant arrivals — there is no control by plant management over the intervals at which ships arrive. In that case, opting for a Poisson process is the logical choice. This does imply that the number of arrivals per year can vary. In the process industry, however, annual throughput is more or less fixed. As a consequence, in our model, the total number of arrivals per year within each ship type is fixed across all arrival processes. This implies that, if the distribution function of interarrival times is exponential, the arrival times are uniformly distributed (Banks et al., 2000). When simulating uncontrolled arrivals, we therefore draw the arrival times per ship type from a uniform distribution over the year.

7.3.7 Ship Arrival Rates

Table 7.1 shows how many ships of each type arrive per year. For each product/cargo type, the number of ships carrying it is chosen such that the total amount of cargo transported matches the plants' capacity. For instance, per year, the plant processes 1,070,000 tons of raw material A. Therefore, the total cargo capacity of ships carrying product A into the port needs to be 1,070,000 tons, which can be verified from the table.

This implies that among simulation runs, only the mutual order of arriving ships and their interarrival times are variable. Thus comparisons regarding port efficiency among arrival processes are kept clean (i.e., devoid of other circumstantial factors such as random fluctuations in production).

With constant loading and unloading times per ship type, fixing the number of ships implies that the utilization rate of the jetty will be the same for all arrival processes. In our case, the utilization rate is 61%. According to industry norms, this is considered to be busy but not overloaded.

7.3.8 Input Analysis

As was mentioned in the introduction, the case study's original tender did not specify the ships' arrival process, providing only the estimated numbers of ships arriving annually per ship type. This is a quite common phenomenon in simulation studies: the distribution functions of the various stochastic processes governing a system, such as interarrival times, service times etc., are often unavailable. In the case of arrivals, a Poisson process has proven to be a reliable choice when arrivals appear to be random. As a consequence, many simulation development environments present the Poisson arrival process as a first option for configuring simulation entity sources, see e.g., (InControl Enterprise Dynamics, 2003), (Kelton et al., 2004), and (Rockwell Software, 2003). If historical arrival data is available, one may attempt to fit a distribution function onto the dataset, and use it in the simulation model to generate arrivals. However, this strategy can easily lead to serious errors.

To illustrate, suppose that the actual system is fed by a hybrid arrival process as outlined before (ships of 6,000 tons and up arrive stock-controlled and ships of less than 6,000 ton arrive equidistantly), but the modeler is not aware of this. He may then use a data fitting program on the collected historical arrival data to help select the distribution function for his model. The results are displayed in figure 7.5 and table 7.2.

The figure was conceived as follows. Arrival data from a hybrid arrival process as generated by our own simulation model were fed to a data fitting program (Arena's

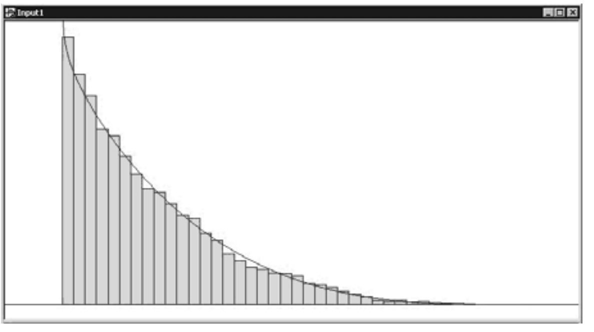

Figure 7.5: Result of Arena's Input Analyzer fit of a hybrid arrival process

Table 7.2: Result of Arena's Input Analyzer fit of a hybrid arrival process

Function	Sq. Error
Beta	0.000147
Weibull	0.000386
Gamma	0.000449
Erlang	0.000508
Exponential	0.000508
Lognormal	0.004450
Normal	0.014700
Triangular	0.015400
Uniform	0.036300

Input Analyzer (Rockwell Software, 2003)). According to this data 13,708 ship arrivals occurred over a ten year period. The figure displays their interarrival times divided over 40 intervals. The distribution function that fits best based on criteria such as the square error is the beta distribution function -0.001 + 2010 * beta(0.957, 4.04), but the exponential distribution function is still quite close (see Law and Kelton (2000) for both distribution functions). Based on these results the beta or even the exponential function appear suitable candidates for modeling the arrival process. However, as can be learned from Table 7.5 (shown and discussed in section 7.6.3) experiments show dramatic differences in simulation outcomes among the various arrival processes considered. This suggests that a data fitting strategy should be preceded by a thorough arrival process analysis to eliminate the possibility of the process being controlled instead of truly random.

7.4 Jetty Scheduling

The arrival process determines when a ship arrives at the port. Next, a scheduling algorithm can be used to control how the ship will be handled in the port. A ship entering the port will eventually be assigned a free mooring point which suits the ship's cargo type and length. The simplest mooring point allocation scheme we consider is one in which the ship is assigned the shortest suitable and available mooring point. If all suitable mooring points are occupied, the ship is placed in a queue before the mooring point with the smallest workload (The workload of a mooring point at instant t is defined as the total time from t that the mooring point will be occupied by the ship currently using it, and the ships currently in the queue before it.) , or, in case of equal workloads, the shortest queue so far. Such a scheme disregards any information on future ship arrivals that might be available.

However, in reality, the ATA of a ship is known to plant management, sometimes days beforehand, by a so-called pre-arrival notice, which can be used in more advanced mooring point allocation algorithms. The general idea is to incorporate all ships within an n-hour horizon into the choice of a mooring point for an incoming ship. Given the fact that for some ship types waiting is more expensive than for others (e.g., dependent on the type of cargo, the ship's capacity or crew size), adequate priority rules might reduce total costs induced by waiting for available mooring points. Also an enumeration algorithm may be applied to select the optimal allocation schedule of all possible schedules within the look-ahead time window. In general, this is a time-consuming approach.

In this chapter we use the ATA information gained from the pre-arrival notices to implement a simple priority scheme with two priority classes (high and low), in which long ships get high priority, and short ones get low priority. The time horizon is 36 hours, i.e., the pre-arrival notice is received 36 hours before the ship's ATA. The priority scheme makes reservations for the high-priority ships based on their ATA. The assignment of a ship to a mooring point can be done as follows. A high-priority ship entering the port is in principle assigned to a free mooring point that suits its cargo type and length. If all suitable mooring points are occupied, the ship is placed in a queue before the mooring point with the smallest workload.

For low-priority ships, the situation is similar, apart from an additional condition. To explain this, let s be a low-priority ship, let t be the current time, let Wi(t) be the workload of mooring point i at time t, and let Di(s) be the time that ship s needs if serviced at mooring point i. Then mooring point i is considered reserved if a high-priority ship arriving within a 36-hour horizon will need mooring point i between t and t + Wi(t) + Di(s). If this is the case, s is not assigned to i, or enqueued before i. Note, that the shorter mooring points at the jetty are never reserved by high-priority ships, since all high-priority ships are too long for these mooring points. Hence, a low-priority ship will always either be assigned to a mooring point directly or placed in a queue before one.

In the presentation of the results in section 7.6, we will make a distinction between model outcomes with and without priority-based mooring point allocation, so that the impact of incorporating such allocation is clearly visible. We will also consider an enumeration algorithm to find the optimal allocation schedule within a 36 hour window.

7.5 The Implementation Model

The model outlined in section 7.2 has been implemented in Enterprise Dynamics (InControl Enterprise Dynamics, 2003), a simulation environment for discrete-event simulation. With this implementation, the experiments in van Asperen et al. (2003) were carried out. Later the model has been implemented in Java using a simulation library. The results presented in this chapter are based on both implementations.

Simulation environments are generally easy to use, and allow for quick model construction. Also they provide built-in animation, generate statistics, and form well-tested simulation environments. Unfortunately, they also have their weak points. Relevant in this context is that, generally speaking, their programming facilities are poor and communication with other programming languages such as Java usually is laborious. General purpose programming languages such as Java or C++ lack the inherent advantages of the simulation environments. On the other hand, they provide a powerful, flexible and fast programming environment. This quality may be indispensable for solving some specific modeling problems, such as complex jetty allocation algorithms.

The initial simulation model was constructed fairly quickly using the Enterprise Dynamics (ED) environment. This implementation provides animation, which facilitates debugging and communication about the simulation model. However, ED's scripting language proved to be too limited for the implementation of complex issues, most notably stock-controlled arrivals. Hence, we implemented the arrival processes in an external (Java) program. The resulting list of interarrival times was used by a custom-built ED object to generate ship arrival events.

Due to more implementation problems concerning the mooring point allocation (e.g. using priorities) and the need for increased runtime speed, the second simulation model was developed in the Java programming language, using the DESMO-J library (University of Hamburg, Department of Computer Science, 2003). This discrete-event simulation framework has been a sound platform for our work.

7.6 Experiments and Results

The Java implementation of the model outlined in the previous section has been used to carry out experiments. While it is capable of generating results on a variety of topics, and on many levels of detail, we focus on the ones relevant to our objective: assessing the impact of using different arrival processes on stock levels and ships' waiting times.

We consider four arrival processes: a Poisson process as described in section 7.3.6, equidistant arrivals per ship type, stock-controlled arrivals per product type, and arrivals modeled using the hybrid process described in section 7.3.4. Each run starts in a steady-state situation, with the tanks partly filled.

Table 7.3 through table 7.5 show the relevant simulation outcomes. Table 7.3 contains waiting statistics for ships with the simplest mooring point allocation scheme as outlined in section 7.4, divided into separate columns for high and low-priority ships (The distinction between high and low-priority ships is made here to facilitate a comparison with the results of simulation runs that do include a priority scheme.) . Table 7.4 reports on the maximum and minimum stock levels reached for each of the arrival processes, both in raw material and finished product tanks. Table 7.5 adds the results of using the simple priority scheme outlined in section 7.4 and an enumeration algorithm to determine the mooring point allocation that yields the least waiting by ships within a 36 hour planning horizon. This is further discussed in section 7.6.4.

7.6.1 Waiting Times

From Table 7.3 it can be observed that the choice for an equidistant, stock-controlled or hybrid arrival process shows a significant difference in terms of the number of waiting ships and the number of hours spent waiting by these ships when compared to the uncontrolled arrival process. This holds for both high and low-priority ships.

Clearly, a mechanism to keep ships apart, such as equidistant or stock-controlled arrival planning, prevents clusters of ships arriving within a small time frame, causing queues. For both low and high-priority ships, the stock-controlled arrival process "outperforms" the equidistant arrival process. The results of the hybrid arrival process are in between those of the equidistant arrival process and those of the stock-controlled process.

Table 7.3: Waiting times per arrival process
Means over a 10-year period; standard deviation is based on ten runs of one year.

	Ship Priority			
	High		Low	
	mean	st.dev.	mean	st.dev.
Percentage of ships that had to wait (%)				
Uncontrolled	45.7	2.1	35.2	2.0
Equidistant	34.7	1.8	23.5	0.8
Stock-controlled	21.1	3.7	12.0	1.0
Hybrid	31.4	3.6	20.6	1.1
Average waiting time of ships that had to wait (hours)				
Uncontrolled	12.3	1.8	7.5	0.9
Equidistant	9.5	0.6	6.2	0.2
Stock-controlled	7.9	1.1	3.5	0.2
Hybrid	8.3	0.7	5.6	0.3

The explanation for this is manifold. For one, stock-controlled arrivals are more efficient overall since they tend to keep ships of identical cargo types apart, whereas equidistant arrivals keep ships of identical types apart. With multiple ship types per cargo type this is an advantage. Furthermore, simulation-specific factors have to be taken into account. Consider the arrival rates of the individual ship types. Here, care has been taken to avoid introducing unrealistic queuing situations. With equidistant arrivals, for example, spreading the arrivals of the first ship of each type, seeks to prevent the scheduling for multiple ship types in such a way, that they all coincide several times a year. Not all such mechanisms are that obvious though, especially when related to another simulation-specific aspect: the jetty layout.

However, the observed differences in waiting time statistics among the arrival processes, whatever their causing factors, clearly demonstrate the need for careful arrival process modeling, which is this chapter's primary objective. Obviously, arrival process modeling requires a careful look at the real situation, involving expert input on many subjects. Only then are simulation results valid, and can they be used in corporate decision-making. Alternatively stated, providing only the numerical data from Table 1, and simply assuming an uncontrolled process, is not sufficient, rendering any subsequent decision (for example on an expensive alternative jetty layout to reduce waiting times) ill founded.

7.6.2 Stock Levels

Table 7.4 shows 10-year stock level statistics in terms of the difference between minimum and maximum levels reached. As could be expected, stock fluctuations are smallest with stock-controlled arrivals, whereas uncontrolled arrivals allow for the largest. The results of the hybrid arrival process are again a blend of the equidistant and stock-controlled results.

Table 7.4: Stock level ranges in tons per arrival process
Results based on ten runs of one year.

	Product							
	A		B		C		D	
	mean	st.dev.	mean	st.dev.	mean	st.dev.	mean	st.dev.
Uncontrolled	74,396	18,333	48,058	11,789	32,045	9,112	89,177	15,112
Equidistant	10,756	273	11,245	312	3,381	283	27,474	574
Stock-controlled	6,970	468	5,890	294	3,012	320	15,982	578
Hybrid	8,212	508	8,032	274	3,369	274	20,932	623

Figure 7.6 shows example stock behavior over time for product D over a one-year period. (As stated before, arrivals are aligned with production in such a way, that

stock does not structurally grow or shrink over a one-year period. Any difference between stock levels at the start or the end of a year are due to ships still being loaded and unloaded at the end.) The initial stock level for each arrival process was set to a value that would prevent stock-outs. Figure 6a shows the results of an uncontrolled arrival process. Note that the scale of figure 6a differs from the scales of the other three graphs: the uncontrolled nature of this arrival process causes large fluctuations in the stock level.

The largest available vessel (see Table 7.1) comes in to load product D eight times a year. This is clearly visible in the graph for the equidistant arrival process (Figure 7.1b). Figure 7.1c shows the typical stock fluctuation pattern for stock-controlled arrivals. Peak levels are reached whenever large ships are scheduled to arrive for loading. A late arrival around day 220 causes a larger peak due to continued production whereas the early arrival of the next ship makes the stock level drop below the base stock level. Figure 7.1d shows the stock level fluctuations for the hybrid arrival process, with stock-controlled arrivals for the larger ocean-going vessels and equidistant arrivals for all other vessels. Notice that in case of product D, stock fluctuation is almost completely determined by the size of the largest vessel, which makes it easy to determine the required tank capacity. The fluctuation patterns observed with the other products are similar in shape. However, their amplitude is considerably smaller, as product D is the only product transported by ships carrying as much as 10,000 and 20,000 tons of chemicals. So, again, the choice of an arrival process is an important factor in simulation outcomes. For example, should the simulation be part of a cost-benefit analysis to the acquisition of additional tankage, then its results are of no value without realistic arrival process modeling.

7.6.3 The Effect of Using a Priority Scheme

In section 7.4 it was explained that a priority scheme is expected to reduce the waiting costs of high-priority ships. A simple priority scheme was considered with two priority classes (high and low), where long ships get high priority, and short ones low priority.

Table 7.5 shows ship waiting statistics over a ten-year simulation period for the same types of arrival process, both with and without a priority scheme. Standard deviations have been omitted for brevity.

In all cases, applying priorities indeed reduces the percentage of high-priority ships, while increasing the percentage of low-priority ships waiting. All waiting time means go up, for which there are, again, multiple causing factors. One seemingly obvious mechanism is that high-priority ships are now very rarely blocked from suitable moor-

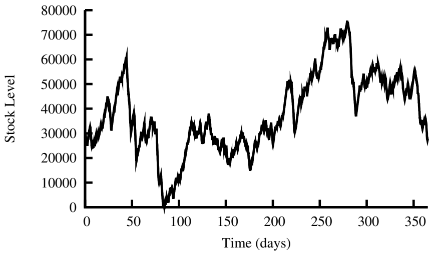

a. Uncontrolled arrival process

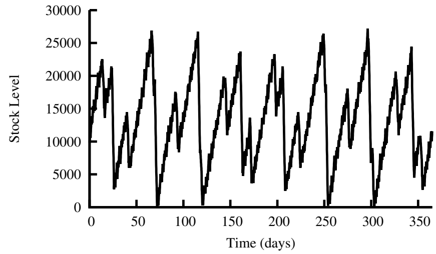

b. Equidistant arrival process

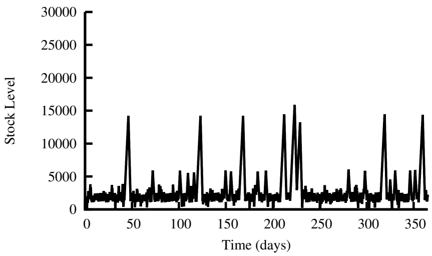

c. Stock-controlled arrival process

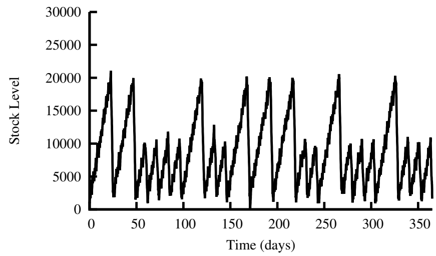

d. Hybrid arrival process

Figure 7.6: Level of tank D over a one year period

ing points by low-priority ships. Hence, if a high-priority ship has to wait, it is probably for another high-priority ship, which takes longer to (un)load, causing longer delays.

The question as to whether total waiting costs are reduced by incorporating priorities, or to what extent, depends on how much more expensive an idle high-priority ship is over a low-priority ship.

7.6.4 Exhaustive Search for the Best Berthing Sequence

In addition to the simple priority scheme for mooring point allocation, we have implemented an enumeration algorithm. This algorithm uses the same information as the simple priority scheme: the pre-arrival notices that are available a number of hours before the actual time of arrival. Rather than looking at just the high-priority ships, the enumeration algorithm evaluates the waiting time for all ships. Every time a ship arrives or a mooring point becomes available, this algorithm determines the best berthing sequence by evaluating all possible berthing sequences. The sequence with the least amount of waiting is then selected as the best possible sequence (in this implementa-

Table 7.5: The effect of using a priority scheme and the optimal berthing sequence with a 36-hour horizon. (Means over a 10-year period.)

	Ship Priority					
	High			Low		
	def.	pri.	enum.	def.	pri.	enum.
Percentage of ships that had to wait (%)						
Uncontrolled	45.7	18.5	15.1	35.2	40.1	34.9
Equidistant	34.7	9.2	3.5	23.5	28.7	23.0
Stock-controlled	21.1	8.5	3.6	12.0	14.2	7.3
Hybrid	31.4	10.1	3.7	20.6	24.6	19.4
Average waiting time of ships that had to wait (hours)						
Uncontrolled	12.3	14.5	13.7	7.5	9.3	7.7
Equidistant	9.5	9.5	9.3	6.2	7.2	5.9
Stock-controlled	7.9	10.0	8.7	3.5	3.8	2.6
Hybrid	8.3	9.7	11.0	5.6	6.3	5.0

tion, we do not distinguish among ship types).

As Table 7.5 shows, the application of the enumeration algorithm provides a clear improvement over the simple priority scheme, both in the percentage of ships that had to wait and in the average number of hours that were spent waiting. The process that best aligns the arrivals with production (the stock-controlled arrival process) achieves the best results. The percentage of larger vessels that have to wait can be reduced to around 3.5% by the application of this enumeration algorithm.

7.7 Conclusions

The importance of careful arrival process modeling is clearly demonstrated in this chapter. Model outcomes over various arrival processes vary significantly, e.g. the uncontrolled process has by far the worst performance of the three processes discussed, both in terms of waiting times and in terms of the required storage capacity, whereas the stock-controlled process performs best overall. An optimization procedure for jetty allocation yields a substantial performance improvement over a first-come-first served allocation, especially in combination with the stock-controlled or hybrid arrival process. Although these results were obtained in a specific case with a relatively high jetty utilization, they are general enough to be appropriate for many port and jetty simulation studies, when the logistical process is directly linked to the production process. The stock-controlled arrival process works well in case of a limited number of products and a large variety in ship sizes. It does however, not coordinate arrivals of ships for different products. The hybrid process provides an alternative in situations where only limited control over arrivals can be implemented. In any case, as soon as there is some sort of control over arrivals, it should be explicitly incorporated in the model.

Obviously, the challenge in shaping and managing these logistical processes is to realize the importance of arrival processes and to assess which one can be actually realized. This requires close collaboration between production, logistics and the sales or marketing functions within a company. If such cooperation is lacking, a marketing department might buy or sell large quantities to meet sales targets, causing serious disruptions in planned arrivals, yielding costly delays. In this case, brute overcapacity in terms of available jetty facilities, piping and tankage is the only alternative.

Acknowledgment

The authors would like to thank Ludo Waltman for his solid implementations of the simulation model in both Enterprise Dynamics and DESMO-J and for carrying out the experiments. They also appreciate his many valuable suggestions.

Chapter 8

Summary

In this chapter we will summarize the main findings from this thesis.

The *first theme* of this thesis is the coordination of transport arrivals with the usage of the goods, the distribution processes, and the use of storage facilities. We have looked at this theme for both containerized and for bulk chemical transport, in chapters 2, 4, 5 (for containerized transport), and 7 (for bulk chemicals). Overall, we conclude that information sharing amongst partners in transport or supply chains is highly beneficial. Even inaccurate information on for example the departure time of a container in a container terminal can help to improve terminal operations.

In chapter 2, we have looked at various approaches for stacking containers that can be employed in a container terminal. On the basis of detailed data from practice, a generator for ship arrivals and a simulator for evaluating the stacking strategies were developed. The peak loads caused by handling very large container ships stress the terminal's operational performance. If the container terminal operator does not have any information on the departure of a container that has just arrived and is about to be stored in the stack area, then the operator can only make an uninformed, random decision. We have however found that using detailed data on the containers and their (planned) movement was very beneficial. We use this detailed data to define groups of similar containers. The similarity of containers is based on characteristics such as the destination (for example, to which port will this container be shipped), the size (20 ft., 40 ft., or 45 ft. container), and the weight (which is important to determining a location within a ship). When a container is to be removed from the stack, we can then use these groups to have more choices. If the container we initially want to remove is stored below some other containers, we can use the groups to locate similar containers in the same area of the stack and see if they can be retrieved faster. If such a container is available, we substitute it for the container we initially wanted to retrieve from the stack. We have labeled this approach “category stacking” and have found it to yield much better results than random stacking. We present several refinements of the basic stacking approach; some of these provide significant benefits whereas others were found to provide little gains. The basic premise of category stacking is that detailed information on the departure of each container is available to the terminal operator. This means that a high level of coordination and information exchange is needed with the parties that arrange onward transportation, whether that transportation is by sea or land. It is clear from practice that achieving this level of information exchange is not easy. The parties are each focused on optimizing their internal operations. Yet, we also see that this information is very valuable and would benefit many parties.

The follow-up project on container stacking described in chapter 3 looked at the influence of using departure time information and the trade-off of stacking crane travel time versus the reshuffling. We have found that rules with a limited number of classes for the remaining residence time work very well. From this we conclude that even imprecise information on this departure time is very valuable and that creating piles of high quality, i.e. piles where the difference in departure time class between levels is exactly one, is beneficial.

The performance of the travel time versus reshuffling stacking rules shows a clear advantage with respect to the reference stacking rules. We have formulated an extension of this stacking rule to attempt to reduce the exit time during future peak workload periods but we found no significant improvement. A combination of the travel times versus reshuffling stacking rule with the departure time classes was also evaluated. Experiments with this rule confirm that it is worthwhile to aim for adjacent containers to have adjacent departure time classes.

The floating stock concept introduced in chapter 4 uses the transfer points between transport modes (such as train and truck) as temporary locations to hold inventory. It also exploits the relatively slow transport modes to reduce the cost of keeping stock at fixed locations such as warehouses because the stock is in-transit, and thus not in a warehouse for a longer time. The longer transport times means that we have to look ahead to make decisions on where to send the inventory. We have used a prediction of future demand to make these decisions to deploy some of the inventory before demand materializes. Using detailed data from a third party logistics service provides (Vos Logistics), we have illustrated this concept with a case in which car tires are shipped from Poland to Germany. We have found that the higher transport costs and longer transit times of intermodal transport can be offset by lower inventory costs. An additional benefit of moving some of the inventory towards the areas of customer demand based on the prediction is that the average time it takes to deliver an order to the customer is reduced. Here, we conclude that it is advantageous to take a more global perspective that takes the entire supply chain into account. By taking this perspective, we can incorporate elements such inventory costs, transport costs, and logistics choices into the model. This means that we can make decisions that are based on considerations of the overall supply chain rather than focusing on local issues such as minimizing transport costs or the customer lead time. The logistics service provider Vos Logistics has applied the concept in practice and uses it for a supply chain of granulate from Rotterdam to Italy (Smit and van Nederpelt, 2009).

Shippers and their customers who want to access the markets in Western Europe have a lot of choices with regard to maritime transport of containers. There are many ports and container terminals that can provide that access. The standardization of load units that has been a benefit of the widespread adoption of sea containers makes it easier to compare the terminals and the facilities they offer. We argue that ports that have a central location in the transport network with respect to the hinterland are at an advantage. The central location allows for rerouting of containers. This is especially beneficial for intercontinental container transport as the long transit time of this type of transport make it more likely that demand may shift among the regions. This flexibility has been investigated using an illustrative case in which a number of alternative scenarios have been evaluated by means of simulation. Detailed cost data was used for the illustrative case. The combination of a simulation model and detailed cost data allows us to quantify the value of the rerouting flexibility. In the case, a combination of using regional distribution centers and a European Distribution Center results in the lowest cost per container. While the specific outcome of this illustrative case can not directly be generalized as it depends on case-specific parameters, the overall approach to quantify the routing flexibility that was outlined in chapter 5 has a more general application. The model that was built in chapter 5 could be used to create a component for a decision support system that could help shippers with this type of problem. For such a component, the current model should be extended to include barge and rail transport.

The *second theme* is the uncertainty associated with the arrival time of ships with bulk chemicals and the impact on (port) logistics.

In chapter 6 we have investigated the coordination between bulk chemical transport by short sea vessel and onward transportation by barge. A chemical plant in The Netherlands uses large annual supplies of a bulk chemical. A number of suppliers deliver their parcels from overseas by short sea vessel to a transshipment point where they are stored using a tank farm. Transportation from the transshipment point to the plant takes place by barge. In this case the tank farm acts as a decoupling point: the logistics process of the supplier is not directly linked to the production process of the customer. Disruptions to the supply of raw materials can be accommodated by keeping sufficient inventory at the tank farm. The main question here is how beneficial improved coordination among suppliers and customer can be to the level of inventory and the associated costs. We have constructed a schedule for the barge that is aimed at meeting the plants' deadlines for delivery and maximizing the opportunities to transfer the bulk chemicals directly from the short-sea vessels to the barge. This is attractive as it reduces the use of the tank farm (potentially reducing handling at the transshipment point and annual hiring costs of tank capacity). We have evaluated the robustness of the barge schedule using simulation. The probability of board to board loading was determined using this simulation model and confirmed using analytical means; with this probability, the benefit of coordination can be quantified.

Chapter 7 deals with the impact of arrival processes modeling on the ship handling process. Whereas the literature on simulation clearly states that input modeling should always be focused at modeling the relevant real-world aspects (see for example Leemis (2004)), many cases that are described in the academic and professional literature assume that arrival processes should be modeled as Poisson processes. This assumption implies that there is no control over the arrivals. This may be an appropriate assumption for customers entering a bank or telephone calls at a call center, but this is not appropriate for ships arriving in port. A ship does not arrive at a completely random point in time: as we have seen, a large number of parties is involved in arranging maritime transport and while these parties may not have perfect control over the time of arrival (controlling the weather and accidents are still beyond their grasp), there is sufficient control to warrant a different way of modeling these arrivals. Although the simulation literature describes approaches to find structure in arrival data in order to create credible input models, this message seems to have a hard time reaching the practitioners and researchers in transportation. We have investigated the potential impact of this modeling choice in the context of a consultancy project on the design of a jetty for a bulk chemical plant. In this project, the arrivals had initially been specified as a number of ships arriving per year, with detailed specifications for each type of ship. The timing was however not explicitly defined. While those involved in the (confidential) consultancy project wisely chose to use arrival processes linked to the plants' inventory levels (the so-called stock-controlled arrivals) and equidistant arrivals (i.e., evenly spread the arrivals of the ships over the time period), they did not investigate the impact of this choice on the outcomes. In chapter 7 we have performed a sensitivity analysis of the model outcomes with regard to the modeling of the ship arrivals. The main result is that this modeling choice has a very significant impact on the outcomes. The arrival modeling approach that requires the highest level of coordination, i.e. the approach in which the planned arrival of a ship is aligned with the plants' inventory, has the best performance. As coordination can require a significant effort, we have introduced a fourth model in which the arrival a small number of larger vessels is aligned with the plants' inventory and a larger number of smaller vessels is planned using equidistant arrivals. The combination of the two arrival coordination models provides an attractive mix: the effort to manage the larger number of small vessels is small, and any variations caused by deviations from the schedule can be absorbed into the stock-controlled arrivals of a small number of large vessels. As expected, this approach has a better performance than the uncontrolled arrival model but less than the stock-controlled model.

In both bulk chemical cases, plant management was not initially aware of the importance of modeling arrivals and coordination. Here, the use of simulation as a method is not just important as a quantitative evaluation method but also as a modeling technique. The conceptual phase of building a simulation model forces the modeler to unearth detailed information on the processes involved.

Samenvatting

Dit proefschrift doet verslag van een aantal onderzoeken naar het optimaliseren van de logistiek in havens, containers en bulk chemicaliën. Het eerste hoofdthema betreft de coördinatie van transport aankomsten met het gebruik van de goederen, de distributie processen en het gebruik van opslagfaciliteiten.

In het eerste onderzoek (hoofdstuk 2) hebben we gekeken naar verschillende manieren om zee-containers (voortaan kortweg containers genoemd) te stapelen in een container terminal. Het gaat daarbij niet om de fysieke manier van stapelen maar om de keuzes welke containers boven op elkaar worden geplaatst. Het is duidelijk dat het niet verstandig is om een container die later zal worden opgehaald te stapelen op een container die eerder weg gaat aangezien we in dat geval eerst de bovenste container moeten weghalen om bij de gewenste container te kunnen komen. In dit hoofstuk hebben we op basis van data uit de praktijk twee programma's gemaakt: een generator, die de aankomsten van schepen en de daarbij behorende lijsten met containers die uit- en ingeladen moeten worden, genereert en een simulator, die op basis van zo'n lijst en een gekozen stapel-strategie alle bewerkingen uitvoert. Met die simulator kunnen we dus de prestaties van een stapel-strategie bepalen en vergelijken met andere strategiën. Deze vergelijkingen zijn interessant omdat de komst van steeds grotere containerschepen, met de daaraan verbonden piek-belastingen voor de container terminals, hoge eisen stellen aan de manier waarop de terminal wordt bestuurd.

In de kern gaat het om informatie. Indien niet bekend is wanneer en hoe containers zullen vertrekken, dan kan de beheerder van de terminal slechts gokken bij de keuze voor een plek om een binnenkomende container te plaatsen. In dit onderzoek hebben we echter gezien dat de prestaties van stapel-strategiën die wel beschikken over gedetailleerde informatie over de containers en hun (geplande) vertrektijd veel beter zijn dan ongeinformeerde strategiën. We hebben deze gedetailleerde informatie gebruikt om categoriën van vergelijkbare containers te formuleren. Een voorbeeld van zo'n categorie kan zijn containers van een bepaald formaat (bijvoorbeeld een van de industriestandaardmaten van 40 ft.), die ongeveer hetzelfde wegen en die met hetzelfde schip zullen worden verscheept. Op het moment dat nu een container uit de stapel moet worden gehaald, kijken we bij deze strategie naar alle containers uit die categorie om meer keuzemogelijkheden te hebben. Indien de gevraagde container niet bovenop staat en er dus eerst een (niet-productieve) verkassingsoperatie zou moeten worden uitgevoerd, kijken we of andere containers uit dezelfde categorie als de gevraagde container wel bovenop staan; als er zo'n container is, dan halen we die in plaats van de gevraagde container uit de stapel en besparen zo op dit moment een verkassingsoperatie. Deze strategie, die we categorie-stapelen hebben genoemd, presteert veel beter dan de referentie strategie, willekeurig stapelen. Naast dit basisconcept van categorie-stapelen hebben we diverse verfijningen en uitbreidingen onderzocht, met wisselend succes.

In hoofdstuk 3 zijn deze stapel strategiën verder onderzocht. Door terug te gaan relatief eenvoudige regels voor het stapelen hebben we geprobeerd meer inzicht in de onderliggende wisselwerkingen te krijgen. De strategiën die gebruik maken van een beperkt aantal klassen van de resterende statijd van een container werken goed. Hierin maken we op dat zelf het gebruik van relatief grove informatie over de vertrektijd de moeite waard is. Ook het maken van stapels met een hoge 'kwaliteit', dat wil zeggen stapels waarbij de klassen van op elkaar gestapelde containers ook precies opvolgend zijn, komt de prestaties ten goede. De wisselwerking tussen rijtijd van de stapelkraan en de kans op toekomstige verkassers leidt tot minder duidelijke resultaten. Deze strategiën doen het duidelijk beter dan de referentie-strategiën, maar minder goed dan de strategiën die gebruik maken van de informatie over de vertrektijd. Een combinatie van de drie ideeën (klassen van resterende statijd, de afweging van rijtijd versus kans op verkassers en het maken van de stapels van hoge kwaliteit) geeft goede resultaten bij een relatief volle stapel.

De essentiële aanname bij categorie-stapelen is dat gedetailleerde informatie over de containers en met name over het geplande vertrek van de terminal beschikbaar is voor de terminal beheerder. Dit impliceert in de praktijk een hoge mate van coördinatie tussen de betrokken partijen, zowel voor het zee-transport als voor het land-transport. De transporten worden veelal door verschillende partijen uitgevoerd. Deze partijen richten zich in eerste instantie op het optimaliseren van hun eigen deel van het transport. Het (tijdig) delen van informatie kan echter zoals we hebben gezien in het geval van categorie-stapelen leiden tot operationele voordelen voor de terminal en daarmee ook voor de transportbedrijven. Als een transportbedrijf nauwkeurige informatie heeft verstrekt over het moment dat een container per truck zal worden opgehaald bij de terminal, dan kan de terminal beheerder deze informatie gebruiken bij het stapelen en kan de tijd die nodig is om zo'n container uit de stapel te halen worden teruggebracht.

In hoofdstuk 4 hebben we het ‘floating stock’ (FS) concept geïntroduceerd. Hierbij wordt gebruik gemaakt van de opslagcapaciteit van de overslagpunten die worden gebruikt bij intermodaal vervoer. Op deze overslagpunten kunnen containers vaak enkele dagen blijven staan zonder dat daar directe opslagkosten aan verbonden zijn. Deze speling tussen het moment waarop een container op zo’n overslagpunt aankomt en het moment waarop de container wordt opgehaald voor het volgende deel van het transport is noodzakelijk om de twee transportbewegingen niet te stringent aan elkaar te koppelen; een dergelijk harde koppeling is erg kwetsbaar voor verstoringen zoals vertragingen. We maken bij dit FS concept ook gebruik van de relatief langzamere transport mogelijkheden zoals transport per trein; zolang de containers met daarin de voorraden onderweg zijn, hoeven we geen kosten voor opslag te betalen. De langere transporttijden betekenen wel dat we verder vooruit moeten kijken; op basis van een voorspelling van de toekomstige vraag versturen we een deel van de voorraad al via die langzamere transportmodi. We hebben dit concept concreet toegepast in een studie met het bedrijf Vos Logistics. In deze casus worden autobanden verscheept van een fabriek in Polen naar klanten in het westen van Duitsland. Met de gedetailleerde kosteninformatie van Vos Logistics hebben we met behulp van simulatie gezien dat de hogere kosten en de langere doorlooptijd van het transport per trein (ten opzichte van transport per vrachtwagen) kunnen worden gecompenseerd door besparingen op de voorraad kosten. Een bijkomend voordeel is dat de reactietijd op een bestelling met dit concept gemiddeld korter wordt omdat een deel van de voorraad zich al dicht bij de klanten bevindt. Het is derhalve de moeite waard om niet alleen naar de losse onderdelen van het transport te kijken, maar juist naar de hele keten. Vanuit dat perspectief kunnen dan keuzes ten aanzien van voorraad- en transport-kosten, en de te kiezen logistieke aanpak worden gemaakt. (Vos Logistics gebruikt het concept inmiddels voor het transport van granulaat van Rotterdam naar Italië (Smit and van Nederpelt, 2009).)

Er zijn veel mogelijkheden om containers over zee van Azië naar (west) Europa te vervoeren; er zijn verschillende havens en daarbinnen container terminals die dergelijke faciliteiten bieden. Door de standaardisatie van zee-containers is het gemakkelijker geworden om terminals en de geboden faciliteiten te vergelijken. Havens en derhalve terminals binnen die havens met een centrale locatie ten opzichte van het achterland hebben naar onze mening een voordeel. Door de centrale locatie wordt het gemakkelijker om de bestemming van een container te veranderen. Vooral bij intercontinentaal transport met de daaraan verbonden lange reistijden is dit aantrekkelijk; als de containers uitwisselbaar zijn, dan kan een verschuiving in de vraag (bijvoorbeeld tussen verschillende landen in West Europa) worden opgevangen door een deel van de containers een nieuwe bestemming te geven. Om dit mogelijke voordeel van een centraal gelegen haven te onderzoeken, hebben we in hoofdstuk 5 een illustratieve casus geformuleerd, waarbinnen enkele verschillende scenarios met behulp van simulatie zijn onderzocht. Voor deze casus was gedetailleerde informatie over kosten beschikbaar dankzij input vanuit de praktijk. Hiermee hebben we de flexibiliteit van havens kunnen kwantificeren. In de casus bleek een combinatie van regionale distributie-centra en een centraal gelegen Europees distributie-centrum de laagste kosten per container te geven.

Hoewel de specifieke uitkomsten niet direct kunnen worden gegeneraliseerd, biedt de aanpak om de flexibiliteit op deze manier te kwantificeren wel een goed uitgangspunt voor een beslissingsondersteunend systeem voor scheepvaartbedrijven. In dat geval zou het huidige model wel nog moeten uitgebreid met mogelijkheden om ook transport per trein en binnenvaartschip te modelleren.

Het tweede thema van dit proefschrift gaat over de aankomsten van schepen met bulk chemicaliën en de impact daarvan op (haven)logistiek.

De coördinatie van het transport van bulk chemicaliën met behulp van een short sea vessel en een binnenvaartschip is het onderwerp van hoofdstuk 6. Een chemische fabriek in Nederland verbruikt jaarlijks grote hoeveelheden chemische grondstoffen. Deze grondstoffen worden betrokken van verschillende leveranciers en per short sea vessel getransporteerd naar een overslagpunt in een haven. De grondstoffen worden aldaar opgeslagen in een verzameling opslagtanks. Het transport van het opslagpunt naar de fabriek wordt gedaan per binnenvaartschip. Door het gebruik van de tanks bij het opslagpunt en enkele (veel kleinere) tanks bij de fabriek zelf wordt het productieproces losgekoppeld van de bevoorrading. Verstoringen van de aanvoer kunnen worden opgevangen indien er voldoende voorraad in de tanks wordt aangehouden. De kernvraag is in hoeverre door een verbeterde afstemming tussen de klant (fabriek) en de leverancier van grondstoffen de kosten voor het aanhouden van dergelijke voorraden kan reduceren. We hebben een schema opgesteld voor het binnenvaartschip waarbij de mogelijkheden voor het direct overslaan van het short sea vessel naar het binnenvaartschip (dus zonder tussenkomst van de tanks) wordt gemaximaliseerd, zonder dat daardoor de deadlines voor de levering (en daarmee het productieproces) in gevaar komen. Directe overslag is aantrekkelijk omdat daarmee het beroep op de (gehuurde) capaciteit van de opslagtanks kan worden teruggebracht. Met behulp van simulatie hebben we onderzocht in hoeverre het opgestelde schema robuust is. Tevens kan hiermee de kans op directe overslag worden bepaald. Voor een eenvoudig scenario hebben we deze kans tevens analytisch bepaald en dit bevestigde de uitkomsten van het simulatiemodel.

De wetenschappelijke en professionele literatuur over simulatie geeft duidelijk aan dat het modelleren van de invoer van simulatiemodellen gebaseerd moet zijn op de relevante aspecten van het te modelleren systeem (zie bijvoorbeeld Leemis (2004)). In de praktijk en ook in veel wetenschappelijke publicaties wordt veelal verondersteld dat aankomsten het beste kunnen worden gemodelleerd als een Poisson proces. Deze aanname impliceert een gebrek aan controle over de aankomsten. Hoewel dit in een simulatiemodel van klanten bij een bankfiliaal of binnenkomende gesprekken in een call center een redelijke aanname is, lijkt dit niet logisch voor aankomsten van schepen in een haven. Het is immers onwaarschijnlijk dat dit een willekeurige gebeurtenis is. Zoals ook uit de andere hoofdstukken in dit proefschrift is gebleken, zijn er vele partijen betrokken bij zee-transport en hoewel er wellicht geen perfecte controle over het preciese aankomsttijdstip is (men denke aan de invloed van het weer), is er toch voldoende sturing van de aankomsttijdstippen dat dit aspect in de modellering van het aankomstproces dient te worden opgenomen. Ondanks de beschikbare literatuur blijkt dit in diverse gepubliceerde projecten niet door de onderzoekers te worden gedaan.

Om de mogelijke impact van een dergelijke modelleerkeuze te onderzoeken, hebben we een bestand simulatiemodel voor het ontwerp van een aanlegsteiger voor een nieuw te bouwen chemische fabriek op de Maasvlakte opnieuw onderzocht. Bij dit (vertrouwelijke) consultancy project waren er gedetailleerde specificaties van de type schepen en het aantal schepen van ieder type dat per jaar zou aankomen. De timing van deze aankomsten binnen het jaar was echter niet gespecificeerd. In het consultancy project was er gekozen voor twee modelleringen: enerzijds het modelleren van de aankomstprocessen van schepen op basis van de voorraadniveaus van de fabriek en anderzijds op basis van een gelijkmatige spreiding van het (bekende) aantal aankomsten over een jaar. Naast deze twee modelleringen hebben wij ook een Poisson aankomstproces en een hybride vorm onderzocht. Bij de hybride vorm wordt de aankomst van een relatief klein aantal grotere schepen gepland op basis van het voorraadniveau van de fabriek en de aankomsten van een groter aantal kleinere schepen worden gelijkmatig in de tijd verdeeld. De modellering die de meeste coördinatie vergt (dus op basis van het voorraadniveau) heeft de beste prestaties. De ongecontrolleerde Poisson aankomstmodellering heeft de slechtste resultaten. De hybride vorm zit daar tussen en kan een aantrekkelijk compromis bieden waarbij de benodige inspanning voor de coördinatie kan worden beperkt.

In deze laatste twee projecten was het management van de fabriek zich in eerste instantie niet bewust van het belang van de aankomstmodellering en coördinatie. Hier blijkt simulatie niet alleen handig als een manier om scenarios op een kwantitatieve manier te evalueren; het gebruik van simulatie als methode dwingt de modelleerder om gedetailleerde informatie over de relevante processen boven tafel te krijgen.

Bibliography

Andrews, S., F. Murphy, X. Wang, and S. Welch (1996). Modeling crude oil lightering in delaware bay. *Interfaces* 26(6), 68–78.

Axelrod, R. (1997). *Simulating Social Phenomena*, Chapter Advancing the Art of Simulation in the Social Sciences, pp. 21–40. Springer, Berlin.

Baker, P. (2007). An exploratory framework of the role of inventory and warehousing in international supply chains. *The International Journal of Logistics Management* 18(1), 64–80. doi:10.1108/09574090710748171.

Banks, J., J. C. II, B. Nelson, and D. Nico (2000). *Discrete-Event System Simulation*. Prentice-Hall.

Baunach, G., E. Wibberley, and B. Wood (1985). Simulation of a coal transshipment terminal: Batam island,indonesia. *Mathematics and Computers in Simulation* 27(2-3), 115–120. doi:10.1016/0378-4754(85)90029-1.

Berger, P. D., A. Gerstenfeld, and A. Zeng (2004). How many suppliers are best? a decision-analysis approach. *Omega* 32(1), 9–15. doi:10.1016/j.omega.2003.09.001.

Boerema, R., B. Groothedde, E. Guis, M. Leijnse, K. Rancuret, C. Ruijgrok, A. Vermunt, M. Verwijmeren, and F. Waals (2003). Distrivaart 2: Aanzet tot een business plan. project report. Technical report, Connekt, Delft. In Dutch.

Bontekoning, Y., C. Macharis, and J. Trip (2004). Is a new applied transportation research field emerging?–a review of intermodal rail-truck freight transport literature. *Transportation Research Part A: Policy and Practice* 38(1), 1–34. ISSN 0965-8564. doi:10.1016/j.tra.2003.06.001.

Bookbinder, J. and N. Fox (1998). Intermodal routing of Canada–Mexico shipments under NAFTA. *Transportation Research Part E: Logistics and Transportation Review* 34(4), 289–303. doi:10.1016/S1366-5545(98)00017-9.

Borgman, B. (2009). Improving container stacking efficiency using simulation. Master's thesis, Erasmus University Rotterdam, School of Economics. Economics & Informatics programme, forthcoming.

Branch, A. E. (2007). *Elements of Shipping* (8th ed.). Routledge.

Caserta, M., S. Voß, and M. Sniedovich (2010). Applying the corridor method to a blocks relocation problem. *OR Spectrum*. doi:10.1007/s00291-009-0176-5. Forthcoming, available online since 18 June 2009.

Castilho, B. d. and C. Daganzo (1993). Handling strategies for import containers at marine terminals. *Transportation Research Part B: Methodological* 27(2), 151–166. ISSN 0191-2615. doi:10.1016/0191-2615(93)90005-U.

Celen, H., R. Slegtenhorst, R. van der Ham, and others (1999). *FAMAS – NewCon: Phase 1: Starting points, Phase 2: Architecture integrating information system*, Volume 32. CTT. ISBN 90-76091-36-6. In Dutch.

Chang, Y.-T. and P. T. Lee (2007). Overview of interport competition: Issues and methods. *Journal of International Logistics And Trade* 5(1), 99–121. ISSN 1738-2122.

Chen, T., K. Lin, and Y.-C. Yuang (2000). Empirical studies on yard operations part 2: Quantifying unproductive moves undertaken in quay transfer operations. *Maritime Policy & Management* 27(2), 191–207. doi:10.1080/030888300286608.

Chopra, S. and P. Meindl (2004). *Supply Chain Management* (2 ed.). Prentice-Hall, New-Jersey.

Christiansen, M., K. Fagerholt, and D. Ronen (2004, February). Ship routing and scheduling: Status and perspectives. *Transportation Science* 38(1), 1–18. doi:10.1287/trsc.1030.0036.

Cullinane, K. and M. Khanna (2000). Economies of scale in large containerships: optimal size and geographical implications. *Journal of Transport Geography* 8(3), 181–195. doi:10.1016/S0966-6923(00)00010-7.

de Langen, P. (2007). Port competition and selection in contestable hinterlands; the case of Austria. *European Journal of Transport and Infrastructure Research* 7(1), 1–14. ISSN 1567-7141.

de Swaan Arons, H., E. van Asperen, R. Dekker, and M. Polman (2004). Coordination in a supply chain for bulk chemicals. In *Proceedings of the 2004 Winter Simulation Conference*. December 5-8, 2004, Washington, DC.

Dekker, R. (1999). *SimLog, Simulatie en logistiek rond de haven*, Chapter Simulation of jetty and storage activities for oil and chemicals, pp. 105–116. Eburon Delft, Netherlands.

Dekker, R., E. van Asperen, G. Ochtman, and W. Kusters (2009). *International Journal of Physical Distribution & Logistics Management* 39(8), 632–648. doi:10.1108/09600030910996297.

Dekker, R., P. Voogd, L. Nagy, and P. Meersmans (2000). Famas - newcon: Long-term stacking experiments for the reference case. Technical report, TRAIL.

Dekker, R., P. Voogd, and E. van Asperen (2006). Advanced methods for container stacking. *OR Spectrum* 28(4), 563–586. doi:10.1007/s00291-006-0038-3.

Dekker, R., P. Voogd, and E. van Asperen (2007). Advanced methods for container stacking. In K. Kim and H.-O. Günther (Eds.), *Container terminals and cargo systems Design, operations management, and logistics control issues*. Springer-Verlag Berlin Heidelberg.

DeLorme, P., J. Procter, S. Swaminathan, and T. Tillinghast (1992). Simulation of a combination carrier air cargo hub. In J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson (Eds.), *Proceedings of the 1992 Winter Simulation Conference*.

Diks, E., A. de Kok, and A. Lagodimos (1996). Multi-echelon systems: A service measure perspective. *European Journal of Operational Research* 95(2), 241–263. doi:10.1016/S0377-2217(96)00120-8.

Duinkerken, M., J. Evers, and J. Ottjes (2001). A simulation model for integrating quay transport and stacking policies in automated terminals. In *Proceedings of the 15th European Simulation Multiconference (ESM2001)*, Prague. ISBN 1-5655-225-3.

ECMT (1993). *Terminology on combined transport*. Publications Services, Paris, France: OECD. URL `http://www1.oecd.org/cem/online/glossaries/`.

Evers, P. T. (1996). The impact of transshipments on safety stock requirements. *Journal of Business Logistics* 17(1), 109–133.

Evers, P. T. (1997). Hidden benefits of emergency transshipments. *Journal of Business Logistics* 18(2), 55–76.

Evers, P. T. (1999). Filling customer order from multiple locations: A comparison of pooling methods. *Journal of Business Logistics* 20(1), 121–139.

Froyland, G., T. Koch, N. Megow, E. Duane, and H. Wren (2008). Optimizing the landside operation of a container terminal. *OR Spectrum* 30, 53–75.

Giuliano, G. and T. O'Brien (2007). Reducing port-related truck emissions: The terminal gate appointment system at the ports of los angeles and long beach. *Transportation Research Part D: Transport and Environment* 12(7), 460 – 473. ISSN 1361-9209. doi:10.1016/j.trd.2007.06.004.

Google (2008). Google maps. URL `http://maps.google.nl`. accessed in november 2008.

Grifo, C. (2008, december). The economic value of terminal operating systems. Master's thesis, Maritime Economics & Logistics, Erasmus University Rotterdam, The Netherlands. In Dutch.

Günther, H.-O. and K.-H. Kim (2006). Container terminals and terminal operations. *OR Spectrum* 28(4), 437–445. doi:10.1007/s00291-006-0059-y.

Haehling von Lanzenauer, C., W. James, and D. Wright (1992). Insufficient supply in a natural gas distribution system: a risk analysis. *European Journal of Operational Research* 56(1), 41–53. doi:10.1016/0377-2217(92)90291-G.

Han, Y., L. H. Lee, E. P. Chew, and K. C. Tan (2008). A yard storage strategy for minimizing traffic congestion in a marine container transshipment hub. *OR Spectrum* 30, 697–720. doi:10.1007/s00291-008-0127-6.

Hengstmengel, J. (2006). Kostenbesparing door toepassing van floating stock in de maritieme logistiek. Bachelor's Thesis Informatics & Economics, Erasmus University Rotterdam, The Netherlands. In Dutch.

Heyden, W. v. d. and J. Ottjes (1985, December). A decision support system for the planning of the workload on a grain terminal. *Decision Support Systems* 1(4), 293–297. doi:10.1016/0167-9236(85)90169-1.

Hirashima, Y., K. Takeda, S. Harada, M. Deng, and A. Inoue (2006). A Q-learning for group-based plan of container transfer scheduling. *JSME International Journal Series C* 49(2), 473–479.

Huggins, E. L. and T. L. Olsen (2003, SEP). Supply chain management with guaranteed delivery. *Management Science* 49(9), 1154–1167. ISSN 0025-1909. doi:10.1287/mnsc.49.9.1154.16573.

IMF (April 2008). World economic outlook database. Technical report, International Monetary Fund. URL `http://www.imf.org`.

InControl Enterprise Dynamics (2003). *Enterprise Dynamics User Guide*. InControl Enterprise Dynamics.

Kang, J., M.-S. Oh, E. Y. Ahn, K. R. Ryu, and K. H. Kim (2006). Planning for intra-block remarshalling in a container terminal. In M. Ali and R. Dapoigny (Eds.), *Advances in applied artificial intelligence*, pp. 1211–1220. Springer.

Kelton, W., R. Sadowski, and D. Sturrock (2004). *Simulation with Arena* (3 ed.). McGraw–Hill.

Kia, M., E. Shayan, and F. Ghotb (2002). Investigation of port capacity under a new approach by computer simulation. *Computers & Industrial Engineering* 42(2-4), 533–540. doi:10.1016/S0360-8352(02)00051-7.

Kim, K. H. (1997). Evaluation of the number of rehandles in container yards. *Computers & Industrial Engineering* 32(4), 701 – 711. ISSN 0360-8352. doi:10.1016/S0360-8352(97)00024-7. New Advances in Analysis of Manufacturing Systems.

Kim, K. H. and J. W. Bae (1998). Re-marshalling export containers in port container terminals. *Computers & Industrial Engineering* 35(3-4), 655–658. ISSN 0360-8352. doi:10.1016/S0360-8352(98)00182-X. Selected Papers from the 22nd ICC and IE Conference.

Kim, K. H. and G.-P. Hong (2006a). A heuristic rule for relocating blocks. *Computers & Operations Research* 33, 940–954.

Kim, K. H. and G.-P. Hong (2006b). A heuristic rule for relocating blocks. *Computers & Operations Research* 33(4), 940–954. doi:10.1016/j.cor.2004.08.005.

Kim, K. H. and H. B. Kim (1998). The optimal determination of the space requirement and the number of transfer cranes for import containers. *Computers & Industrial*

Engineering 35(3-4), 427 – 430. ISSN 0360-8352. doi:10.1016/S0360-8352(98)00125-9. Selected Papers from the 22nd ICC and IE Conference.

Kim, K. H. and H. B. Kim (1999). Segregating space allocation models for container inventories in port container terminals. *International Journal of Production Economics* 59(1-3), 415–423. ISSN 0925-5273. doi:10.1016/S0925-5273(98)00028-0.

Kim, K. H., Y. M. Park, and K.-R. Ryu (2000). Deriving decision rules to locate export containers in container yards. *European Journal of Operational Research* 124(1), 89–101. doi:10.1016/S0377-2217(99)00116-2.

Konings, J. (1996). Integrated centres for the transshipment, storage, collection and distribution of goods: A survey of the possibilities for a high-quality intermodal transport concept. *Transport Policy* 3(1–2), 3–11. doi:10.1016/0967-070X(96)00007-8.

Konings, R. (2005). Foldable containers to reduce the costs of empty transport? a cost-benefit analysis from a chain and multi-actor perspective. *Maritime Economics & Logistics* 7(3), 223–249. doi:10.1057/palgrave.mel.9100139.

Krimpen, H. v. (1986). *Boek over het maken van boeken* (2 ed.). Veenendaal, NL: Gaade. ISBN 906175212.

Landers, T., M. Cole, B. Walker, and R. Kirk (2000). The virtual warehousing concept. *Transportation Research Part E* 36(2), 115–125. ISSN 1366-5545. doi:10.1016/S1366-5545(99)00024-1.

Law, A. M. and W. D. Kelton (2000). *Simulation, Modeling and Analysis* (3rd ed.). McGraw-Hill.

L'Ecuyer, P. and E. Buist (2005). Simulation in java with SSJ. In M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines (Eds.), *Proceedings of the 2005 Winter Simulation Conference*, pp. 611–620.

L'Ecuyer, P., L. Meliani, and J. Vaucher (2002). SSJ: A framework for stochastic simulation in Java. In E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes (Eds.), *Proceedings of the 2002 Winter Simulation Conference*, pp. 234–242. IEEE Press. ISBN 0-7803-7615-3.

Lee, H. L., V. Padmanabhan, and S. Whang (1997). The bullwhip effect in supply chains. *Sloan Management Review* 38(3), 93–102. ISSN 0019-848X.

Leemis, L. M. (2004). Building credible input models. In R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters (Eds.), *Proceedings of the 2004 Winter Simulation Conference*.

Levinson, M. (2006). *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger*. Princeton University Press. ISBN 0691123241.

Lim, A. and Z. Xu (2006). A critical-shaking neighborhood search for the yard allocation problem. *European Journal of Operational Research* 174, 1247–1259.

Macharis, C. and Y. Bontekoning (2004). Opportunities for OR in intermodal freight transport research: A review. *European Journal of Operational Research* 153(2), 400–416. doi:10.1016/S0377-2217(03)00161-9.

Mangan, J., C. Lalwani, and B. Fynes (2008). Port-centric logistics. *International Journal of Logistics Management* 19, 29–41. doi:10.1108/09574090810872587.

McGinnis, M. A. (1989). A comparative evaluation of freight transportation choice models. *Transportation Journal* 29(2), 36–46. ISSN 00411612.

Meersmans, P. and R. Dekker (2001). Operations research supports container handling. Technical Report EI/2001-22, Report Econometric Institute, Erasmus University Rotterdam.

Meixell, M. J. and M. Norbis (2008). A review of the transportation mode choice and carrier selection literature. *The International Journal of Logistics Management* 19(2), 183–211. ISSN 0957-4093. doi:10.1108/09574090810895951.

Minner, S. (2003). Multiple-supplier inventory models in supply chain management: a review. *International Journal of Production Economics* 81–82, 265–279. doi:10.1016/S0925-5273(02)00288-8.

Moinzadeh, K. and S. Nahmias (1988). A continuous review model for an inventory system with two supply modes. *Management Science* 34(8), 761–773.

Naim, M. M., A. T. Potter, R. J. Mason, and N. Bateman (2006). The role of transport flexibility in logistics provision. *International Journal of Logistics Management* 17(3), 297 – 311. ISSN 09574093. doi:10.1108/09574090610717491.

Newhart, D., K. Stott, and F. Vasko (1993). Consolidating product sizes to minimize inventory levels for a multi-stage production and distribution system. *Journal of the Operational Research Society* 44(7), 637–644. ISSN 01605682. doi:10.1057/jors.1993.113.

Nieboer, R. and R. Dekker (1995). Brownian motion approximations for tankage assessment and stock control. *European Journal of Operational Research* 85(1), 192–204. doi:10.1016/0377-2217(93)E0356-3.

Notteboom, T. (2004). Container shipping and ports: An overview. *Review of Network Economics* 3(2), 86–106.

Notteboom, T. and J.-P. Rodrigue (2005). Port regionalization: towards a new phase in port development. *Maritime Policy & Management* 32(3), 297–313. doi:10.1080/03088830500139885.

Notteboom, T. and J.-P. Rodrigue (2008). Containerisation, box logistics and global supply chains: The integration of ports and liner shipping networks. *Maritime Economics & Logistics* 10(1-2), 152–174. doi:10.1057/palgrave.mel.9100196.

Ochtman, G. (2003, december). De mogelijkheden van floating stock in een FMCG suppy chain. Master's thesis, Erasmus School of Economics, Erasmus University Rotterdam, The Netherlands. In Dutch.

Ochtman, G., R. Dekker, E. van Asperen, and W. Kusters (2004). Floating stocks in a fast-moving consumer goods supply chain: Insights from a case study. Technical Report ERS-2004-010-LIS, ERIM Report Series Research in Management, Erasmus University Rotterdam, The Netherlands.

Olivier, D., F. Parola, B. Slack, and J. J. Wang (2007). The time scale of internationalisation: The case of the container port industry. *Maritime Economics & Logistics* 9, 1–34. ISSN 1479-2931. doi:10.1057/palgrave.mel.9100169.

Ottjes, J. (1992). *Modeling and Simulation of logistic systems*. Delft University. in Dutch.

Ottjes, J., S. Hengst, and W. Tutuarima (1994). A simulation model of a sailing container terminal service in the port of rotterdam. In Guasch and Huber (Eds.), *Proceedings of the 1994 Conference on Modelling and Simulation*, pp. 876–880.

Pagh, J. D. and M. C. Cooper (1998). Supply chain postponement and speculation strategies: How to choose the right strategy. *Journal of Business Logistics* 19(2), 13–33. ISSN 0735-3766.

Panayides, P. M. and D.-W. Song (2008). Evaluating the integration of seaport container terminals in supply chains. *International Journal of Physical Distribution & Logistics Management* 38(7), 562–584. doi:10.1108/09600030810900969.

Park, B. J., H. R. Choi, H. K. Kwon, and M. H. Kang (2006). *AI 2006: Advances in Artificial Intelligence*, Volume 4304 of *Lecture Notes in Computer Science*, Chapter Simulation Analysis on Effective Operation of Handling Equipments in Automated Container Terminal, pp. 1231–1238. Berlin/Heidelberg: Springer Verlag.

Philips, O. (1976). *Optimization models for a crude oil storage export system*. Ph. D. thesis, Pennsylvania State University.

Pourakbar, M., A. Sleptchenko, and R. Dekker (2009). The floating stock policy in fast moving consumer goods supply chains. *Transportation Research Part E: Logistics and Transportation Review* 45(1), 39 – 49. ISSN 1366-5545. doi:10.1016/j.tre.2008.06.001.

Robinson, R. (2002). Ports as elements in value-driven chain systems: the new paradigm. *Maritime Policy & Management* 29(3), 241–255. ISSN 03088839. doi:10.1080/03088830210132623.

Rockwell Software (2003). *Arena User's Guide*. Rockwell Software. URL `http://www.software.rockwell.com/arenasimulation/`.

Rodon, J. and J. Ramis-Pujol (2006). Exploring the intricacies of integrating with a port community system. In *Bled eConference, Bled, Slovenia*, pp. 1–15.

Rutten, B. (1995). *On Medium Distance Intermodal Rail Transport*. Ph. D. thesis, Faculty of Mechanical Engineering, TU Delft, Delft University Press, Delft, The Netherlands.

Saanen, Y. A. and R. Dekker (2006a). Intelligent stacking as way out of congested yards? part 1. *Port Technology International* 31, 87–92.

Saanen, Y. A. and R. Dekker (2006b). Intelligent stacking as way out of congested yards? part 2. *Port Technology International* 32, 80–86.

Sampson, R. J., M. T. Farris, and D. L. Shrock (1985). *Domestic transportation: practice, theory, and policy* (5 ed.). Boston/London: Houghton Mifflin. ISBN 0-395-35713-6.

Samwel, E. (2007). IT: past, present and future. *Containerisation International* 40, 51–57. special issue.

Sculli, D. and C. Hui (1988). Three dimensional stacking of containers. *Omega* 16(6), 585–594. ISSN 0305-0483. doi:10.1016/0305-0483(88)90032-1.

Sherbooke, C. C. (2004). *Optimal Inventory Modeling of Systems: Multi-echelon techniques* (2nd ed.). Kluwer Academic Publishers, Boston/Dordrecht/London. ISBN 1-4020-7849-8.

Silver, E. and R. Peterson (1985). *Decision systems for inventory management and production planning*. New York: John Wiley & Sons.

Slack, B. (1985). Containerization, inter-port competition, and port selection. *Maritime Policy & Management* 12(4), 293–303. doi:10.1080/03088838500000043.

Smit, S. and M. van Nederpelt (2009). Onthaasten levert co2- en kostenreductie op. URL http://www.logistiek.nl/experts/id12243-Onthaasten_levert_CO_en_kostenreductie_op.html. accessed february 2009.

Stahlbock, R. and S. Voß(2008). Operations research at container terminals: a literature update. *OR Spectrum* 30(1), 1–52. doi:10.1007/s00291-007-0100-9.

Steenken, D., S. Voß, and R. Stahlbock (2004). Container terminal operation and operations research – a classification and literature review. *OR Spectrum* 26(1), 3–49. doi:10.1007/s00291-003-0157-z.

Taleb-Ibrahimi, M., B. De Castilho, and C. F. Daganzo (1993). Storage space vs handling work in container terminals. *Transportation Research Part B: Methodological* 27(1), 13–32. ISSN 0191-2615. doi:10.1016/0191-2615(93)90009-Y.

Teulings, M. and P. van der Vlist (2001). Managing the supply chain with standard mixed loads. *International Journal of Physical Distribution & Logistics Management* 31(3), 169–186. doi:10.1108/09600030110389442.

Tongzon, J. (2002). Port choice determinants in a competitive environment. In *International Association of Maritime Economists Annual Conference*, Panama, pp. 1–22. URL http://www.eclac.cl/Transporte/perfil/iamepapers/papers.asp.

Transportation Institute (2009). Maritime glossary. URL http://www.trans-inst.org/seawords.htm. accessed April 2009.

Tyworth, J. and A. Zeng (1998). Estimating the effects of carrier transit-time performance on logistics cost and service. *Transportation Research Part A* 32(2), 89–97. doi:10.1016/S0965-8564(97)00020-7.

UNCTAD (1978). *chapter II: the break-bulk berth group*, pp. 108–128. United Nations, New York. TD/B/C.U/175.

United Nations (2002). *Commercial Development of Regional Ports as Logistics Centres*. United Nations Publications. ISBN 9211201551. Economic and Social Commission for Asia and the Pacific.

University of Hamburg, Department of Computer Science (2003). *DESMO-J*. University of Hamburg, Department of Computer Science. URL `http://www.desmoj.de/`.

Upward Systems (1994). *MUST simulation software: User and reference manual*. Delft, The Netherlands: Upward Systems.

van Asperen, E., R. Dekker, M. Polman, and H. de Swaan Arons (2003). Modeling ship arrivals in ports. In S. Chick, P. Sánchez, D. Ferrin, and D. J. Morrice (Eds.), *Proceedings of the 2003 Winter Simulation Conference*, New Orleans.

van Asperen, E., R. Dekker, M. Polman, H. de Swaan Arons, and L. Waltman (2005). Arrival processes in port modelling: Insights from a case study. In *Proceedings of the International Workshop on Harbour, Maritime & Multimodal Logistics Modelling and Simulation (HMS2005), Marseille, France*.

van der Horst, M. R. and P. W. de Langen (2008). Coordination in hinterland transport chains: A major challenge for the seaport community. *Maritime Economics & Logistics* 10(1/2), 108–129. doi:10.1057/palgrave.mel.9100194.

van Nunen, J. and L. Verspui (1999). *SimLog, Simulatie en logistiek rond de haven*. Eburon Delft, Netherlands.

Veldman, S. J. and E. H. Bückmann (2003). A model on container port competition: An application for the west european container hub-ports. *Maritime Economics & Logistics* 5(1), 3–22. ISSN 1479-2931. doi:10.1057/palgrave.mel.9100058.

Vis, I. F. and R. de Koster (2003). Transshipment of containers at a container terminal: An overview. *European Journal of Operations Research* 147(1), 1–16. ISSN 0377-2217. doi:10.1016/S0377-2217(02)00293-X.

Voogd, P., R. Dekker, and P. Meersmans (1999). Famas – newcon: A generator program for stacking in the reference case. Technical Report EI-9943/A, Econometric Institute, Erasmus University Rotterdam.

Wallin, C., M. J. Rungtusanatham, and E. Rabinovich (2006). What is the “right” inventory management approach for a purchased item? *International Journal of Operations & Production Management* 26(1), 50–68. doi:10.1108/01443570610637012.

Wiegmans, B. W., A. v. d. Hoest, and T. E. Notteboom (2008). Port and terminal selection by deep-sea container operators. *Maritime Policy & Management* 35(6), 517–534. ISSN 03088839. doi:10.1080/03088830802469329.

Wikipedia (2009). Optimization (mathematics). URL `http://en.wikipedia.org/wiki/Optimization_(mathematics)`. accessed March 2009.

Yeo, G.-T., M. Roe, and J. Dinwoodie (2008). Evaluating the competitiveness of container ports in Korea and China. *Transportation Research Part A: Policy and Practice* 42(6), 910 – 921. ISSN 0965-8564. doi:10.1016/j.tra.2008.01.014.

About the author

Eelco van Asperen (11 april 1965, Rotterdam) heeft na het VWO (OSG Gemini te Ridderkerk) de studie Bestuurlijke Informatica gevolgd aan de Erasmus Universiteit Rotterdam (EUR, afgestudeerd in 1993).
Van 1990 tot 2000 was hij als software expert werkzaam bij de Faculteit der Economische Wetenschappen (FEW) van de EUR.
Sinds 2000 was hij universitair docent bij de FEW. Hij heeft onderwijs verzorgd aan meerdere opleidingen (Economie & Bedrijfseconomie, Informatica & Economie, Global E-Commerce Masters, Maritime Economics & Logistics, ESAA Auditing) en geparticipeerd in diverse projecten (o.a. met AC Transport en de EUR Universiteitsbibliotheek).
Met ingang van 1 september 2009 werkt hij als lecturer bij het Center for Maritime Economics & Logistics, Erasmus Universiteit Rotterdam. Zijn onderzoek is vooral gericht op het gebruik van zee containers in de logistiek.

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT (ERIM)

ERIM PH.D. SERIES
RESEARCH IN MANAGEMENT

ERIM Electronic Series Portal: http://hdl.handle.net/1765/1

Agatz, N.A.H., *Demand Management in E-Fulfillment*, Promotor: Prof.dr.ir. J.A.E.E. van Nunen, EPS-2009-163-LIS, ISBN: 978-90-5892-200-7, http://hdl.handle.net/1765/15425

Althuizen, N.A.P., *Analogical Reasoning as a Decision Support Principle for Weakly Structured Marketing Problems*, Promotor: Prof.dr.ir. B. Wierenga, EPS-2006-095-MKT, ISBN: 90-5892-129-8, http://hdl.handle.net/1765/8190

Alvarez, H.L., *Distributed Collaborative Learning Communities Enabled by Information Communication Technology*, Promotor: Prof.dr. K. Kumar, EPS-2006-080-LIS, ISBN: 90-5892-112-3, http://hdl.handle.net/1765/7830

Appelman, J.H., *Governance of Global Interorganizational Tourism Networks: Changing Forms of Co-ordination between the Travel Agency and Aviation Sector*, Promotors: Prof.dr. F.M. Go & Prof.dr. B. Nooteboom, EPS-2004-036-MKT, ISBN: 90-5892-060-7, http://hdl.handle.net/1765/1199

Assem, M.J. van den, *Deal or No Deal? Decision Making under Risk in a Large-Stake TV Game Show and Related Experiments*, Promotor: Prof.dr. J. Spronk, EPS-2008-138-F&A, ISBN: 978-90-5892-173-4, http://hdl.handle.net/1765/13566

Baquero, G, *On Hedge Fund Performance, Capital Flows and Investor Psychology*, Promotor: Prof.dr. M.J.C.M. Verbeek, EPS-2006-094-F&A, ISBN: 90-5892-131-X, http://hdl.handle.net/1765/8192

Berens, G., *Corporate Branding: The Development of Corporate Associations and their Influence on Stakeholder Reactions*, Promotor: Prof.dr. C.B.M. van Riel, EPS-2004-039-ORG, ISBN: 90-5892-065-8, http://hdl.handle.net/1765/1273

Berghe, D.A.F. van den, *Working Across Borders: Multinational Enterprises and the Internationalization of Employment*, Promotors: Prof.dr. R.J.M. van Tulder & Prof.dr. E.J.J. Schenk, EPS-2003-029-ORG, ISBN: 90-5892-05-34, http://hdl.handle.net/1765/1041

Berghman, L.A., *Strategic Innovation Capacity: A Mixed Method Study on Deliberate Strategic Learning Mechanisms*, Promotor: Prof.dr. P. Mattyssens, EPS-2006-087-MKT, ISBN: 90-5892-120-4, http://hdl.handle.net/1765/7991

Bijman, W.J.J., *Essays on Agricultural Co-operatives: Governance Structure in Fruit and Vegetable Chains*, Promotor: Prof.dr. G.W.J. Hendrikse, EPS-2002-015-ORG, ISBN: 90-5892-024-0, http://hdl.handle.net/1765/867

Bispo, A., *Labour Market Segmentation: An investigation into the Dutch hospitality industry*, Promotors: Prof.dr. G.H.M. Evers & Prof.dr. A.R. Thurik, EPS-2007-108-ORG, ISBN: 90-5892-136-9, http://hdl.handle.net/1765/10283

Blindenbach-Driessen, F., *Innovation Management in Project-Based Firms*, Promotor: Prof.dr. S.L. van de Velde, EPS-2006-082-LIS, ISBN: 90-5892-110-7, http://hdl.handle.net/1765/7828

Boer, C.A., *Distributed Simulation in Industry,* Promotors: Prof.dr. A. de Bruin & Prof.dr.ir. A. Verbraeck, EPS-2005-065-LIS, ISBN: 90-5892-093-3, http://hdl.handle.net/1765/6925

Boer, N.I., *Knowledge Sharing within Organizations: A situated and Relational Perspective,* Promotor: Prof.dr. K. Kumar, EPS-2005-060-LIS, ISBN: 90-5892-086-0, http://hdl.handle.net/1765/6770

Boer-Sorbán, K., *Agent-Based Simulation of Financial Markets: A modular, Continuous-Time Approach,* Promotor: Prof.dr. A. de Bruin, EPS-2008-119-LIS, ISBN: 90-5892-155-0, http://hdl.handle.net/1765/10870

Boon, C.T., *HRM and Fit: Survival of the Fittest!?,* Promotors: Prof.dr. J. Paauwe & Prof.dr. D.N. den Hartog, EPS-2008-129-ORG, ISBN: 978-90-5892-162-8, http://hdl.handle.net/1765/12606

Braun, E., *City Marketing: Towards an Integrated Approach,* Promotor: Prof.dr. L. van den Berg, EPS-2008-142-MKT, ISBN: 978-90-5892-180-2, http://hdl.handle.net/1765/13694

Brito, M.P. de, *Managing Reverse Logistics or Reversing Logistics Management?* Promotors: Prof.dr.ir. R. Dekker & Prof.dr. M. B. M. de Koster, EPS-2004-035-LIS, ISBN: 90-5892-058-5, http://hdl.handle.net/1765/1132

Brohm, R., *Polycentric Order in Organizations: A Dialogue between Michael Polanyi and IT-Consultants on Knowledge, Morality, and Organization,* Promotors: Prof.dr. G. W. J. Hendrikse & Prof.dr. H. K. Letiche, EPS-2005-063-ORG, ISBN: 90-5892-095-X, http://hdl.handle.net/1765/6911

Brumme, W.-H., *Manufacturing Capability Switching in the High-Tech Electronics Technology Life Cycle,* Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr.ir. L.N. Van Wassenhove, EPS-2008-126-LIS, ISBN: 978-90-5892-150-5, http://hdl.handle.net/1765/12103

Burgers, J.H., *Managing Corporate Venturing: Multilevel Studies on Project Autonomy, Integration, Knowledge Relatedness, and Phases in the New Business Development Process,* Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2008-136-STR, ISBN: 978-90-5892-174-1, http://hdl.handle.net/1765/13484

Campbell, R.A.J., *Rethinking Risk in International Financial Markets*, Promotor: Prof.dr. C.G. Koedijk, EPS-2001-005-F&A, ISBN: 90-5892-008-9, http://hdl.handle.net/1765/306

Chen, C.-M., *Evaluation and Design of Supply Chain Operations Using DEA*, Promotor: Prof.dr. J.A.E.E. van Nunen, EPS-2009-172-LIS, ISBN: 978-90-5892-209-0, http://hdl.handle.net/1765/1

Chen, H., *Individual Mobile Communication Services and Tariffs*, Promotor: Prof.dr. L.F.J.M. Pau, EPS-2008-123-LIS, ISBN: 90-5892-158-1, http://hdl.handle.net/1765/11141

Chen, Y., *Labour Flexibility in China's Companies: An Empirical Study*, Promotors: Prof.dr. A. Buitendam & Prof.dr. B. Krug, EPS-2001-006-ORG, ISBN: 90-5892-012-7, http://hdl.handle.net/1765/307

Damen, F.J.A., *Taking the Lead: The Role of Affect in Leadership Effectiveness*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2007-107-ORG, http://hdl.handle.net/1765/10282

Daniševská, P., *Empirical Studies on Financial Intermediation and Corporate Policies*, Promotor: Prof.dr. C.G. Koedijk, EPS-2004-044-F&A, ISBN: 90-5892-070-4, http://hdl.handle.net/1765/1518

Delporte-Vermeiren, D.J.E., *Improving the Flexibility and Profitability of ICT-enabled Business Networks: An Assessment Method and Tool*, Promotors: Prof. mr. dr. P.H.M. Vervest & Prof.dr.ir. H.W.G.M. van Heck, EPS-2003-020-LIS, ISBN: 90-5892-040-2, http://hdl.handle.net/1765/359

Derwall, J.M.M., *The Economic Virtues of SRI and CSR*, Promotor: Prof.dr. C.G. Koedijk, EPS-2007-101-F&A, ISBN: 90-5892-132-8, http://hdl.handle.net/1765/8986

Diepen, M. van, *Dynamics and Competition in Charitable Giving*, Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2009-159-MKT, ISBN: 978-90-5892-188-8, http://hdl.handle.net/1765/14526

Dietz, H.M.S., *Managing (Sales)People towards Perfomance: HR Strategy, Leadership & Teamwork*, Promotor: Prof.dr. G.W.J. Hendrikse, EPS-2009-168-ORG, ISBN: 978-90-5892-210-6, http://hdl.handle.net/1765/1

Dijksterhuis, M., *Organizational Dynamics of Cognition and Action in the Changing Dutch and US Banking Industries*, Promotors: Prof.dr.ir. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2003-026-STR, ISBN: 90-5892-048-8, http://hdl.handle.net/1765/1037

Eijk, A.R. van der, *Behind Networks: Knowledge Transfer, Favor Exchange and Performance*, Promotors: Prof.dr. S.L. van de Velde & Prof.dr.drs. W.A. Dolfsma, EPS-2009-161-LIS, ISBN: 978-90-5892-190-1, http://hdl.handle.net/1765/14613

Elstak, M.N., *Flipping the Identity Coin: The Comparative Effect of Perceived, Projected and Desired Organizational Identity on Organizational Identification and Desired Behavior*, Promotor: Prof.dr. C.B.M. van Riel, EPS-2008-117-ORG, ISBN: 90-5892-148-2, http://hdl.handle.net/1765/10723

Erken, H.P.G., *Productivity, R&D and Entrepreneurship*, Promotor: Prof.dr. A.R. Thurik, EPS-2008-147-ORG, ISBN: 978-90-5892-179-6, http://hdl.handle.net/1765/14004

Fenema, P.C. van, *Coordination and Control of Globally Distributed Software Projects*, Promotor: Prof.dr. K. Kumar, EPS-2002-019-LIS, ISBN: 90-5892-030-5, http://hdl.handle.net/1765/360

Fleischmann, M., *Quantitative Models for Reverse Logistics*, Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr.ir. R. Dekker, EPS-2000-002-LIS, ISBN: 35-4041-711-7, http://hdl.handle.net/1765/1044

Flier, B., *Strategic Renewal of European Financial Incumbents: Coevolution of Environmental Selection, Institutional Effects, and Managerial Intentionality*, Promotors: Prof.dr.ir. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2003-033-STR, ISBN: 90-5892-055-0, http://hdl.handle.net/1765/1071

Fok, D., *Advanced Econometric Marketing Models*, Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2003-027-MKT, ISBN: 90-5892-049-6, http://hdl.handle.net/1765/1035

Ganzaroli, A., *Creating Trust between Local and Global Systems*, Promotors: Prof.dr. K. Kumar & Prof.dr. R.M. Lee, EPS-2002-018-LIS, ISBN: 90-5892-031-3, http://hdl.handle.net/1765/361

Gertsen, H.F.M., *Riding a Tiger without Being Eaten: How Companies and Analysts Tame Financial Restatements and Influence Corporate Reputation*, Promotor: Prof.dr. C.B.M. van Riel, EPS-2009-171-ORG, ISBN: 90-5892-214-4, http://hdl.handle.net/1765/1

Gilsing, V.A., *Exploration, Exploitation and Co-evolution in Innovation Networks*, Promotors: Prof.dr. B. Nooteboom & Prof.dr. J.P.M. Groenewegen, EPS-2003-032-ORG, ISBN: 90-5892-054-2, http://hdl.handle.net/1765/1040

Gijsbers, G.W., *Agricultural Innovation in Asia: Drivers, Paradigms and Performance,* Promotor: Prof.dr. R.J.M. van Tulder, EPS-2009-156-ORG, ISBN: 978-90-5892-191-8, http://hdl.handle.net/1765/14524

Gong, Y., *Stochastic Modelling and Analysis of Warehouse Operations*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr. S.L. van de Velde, EPS-2009-180-LIS, ISBN: 978-90-5892-219-9, http://hdl.handle.net/1765/1

Govers, R., *Virtual Tourism Destination Image: Glocal Identities Constructed, Perceived and Experienced*, Promotors: Prof.dr. F.M. Go & Prof.dr. K. Kumar, EPS-2005-069-MKT, ISBN: 90-5892-107-7, http://hdl.handle.net/1765/6981

Graaf, G. de, *Tractable Morality: Customer Discourses of Bankers, Veterinarians and Charity Workers*, Promotors: Prof.dr. F. Leijnse & Prof.dr. T. van Willigenburg, EPS-2003-031-ORG, ISBN: 90-5892-051-8, http://hdl.handle.net/1765/1038

Greeven, M.J., *Innovation in an Uncertain Institutional Environment: Private Software Entrepreneurs in Hangzhou, China*, Promotor: Prof.dr. B. Krug, EPS-2009-164-ORG, ISBN: 978-90-5892-202-1, http://hdl.handle.net/1765/15426

Groot, E.A. de, *Essays on Economic Cycles*, Promotors: Prof.dr. Ph.H.B.F. Franses & Prof.dr. H.R. Commandeur, EPS-2006-091-MKT, ISBN: 90-5892-123-9, http://hdl.handle.net/1765/8216

Guenster, N.K., *Investment Strategies Based on Social Responsibility and Bubbles*, Promotor: Prof.dr. C.G. Koedijk, EPS-2008-175-F&A, ISBN: 978-90-5892-206-9, http://hdl.handle.net/1765/1

Gutkowska, A.B., *Essays on the Dynamic Portfolio Choice*, Promotor: Prof.dr. A.C.F. Vorst, EPS-2006-085-F&A, ISBN: 90-5892-118-2, http://hdl.handle.net/1765/7994

Hagemeijer, R.E., *The Unmasking of the Other*, Promotors: Prof.dr. S.J. Magala & Prof.dr. H.K. Letiche, EPS-2005-068-ORG, ISBN: 90-5892-097-6, http://hdl.handle.net/1765/6963

Halderen, M.D. van, *Organizational Identity Expressiveness and Perception Management: Principles for Expressing the Organizational Identity in Order to Manage the Perceptions and Behavioral Reactions of External Stakeholders*, Promotor: Prof.dr. S.B.M. van Riel, EPS-2008-122-ORG, ISBN: 90-5892-153-6, http://hdl.handle.net/1765/10872

Hartigh, E. den, *Increasing Returns and Firm Performance: An Empirical Study*, Promotor: Prof.dr. H.R. Commandeur, EPS-2005-067-STR, ISBN: 90-5892-098-4, http://hdl.handle.net/1765/6939

Hermans. J.M., *ICT in Information Services; Use and Deployment of the Dutch Securities Trade, 1860-1970,* Promotor: Prof.dr. drs. F.H.A. Janszen, EPS-2004-046-ORG, ISBN 90-5892-072-0, http://hdl.handle.net/1765/1793

Hessels, S.J.A., *International Entrepreneurship: Value Creation Across National Borders*, Promotor: Prof.dr. A.R. Thurik, EPS-2008-144-ORG, ISBN: 978-90-5892-181-9, http://hdl.handle.net/1765/13942

Heugens, P.P.M.A.R., *Strategic Issues Management: Implications for Corporate Performance*, Promotors: Prof.dr.ir. F.A.J. van den Bosch & Prof.dr. C.B.M. van Riel, EPS-2001-007-STR, ISBN: 90-5892-009-9, http://hdl.handle.net/1765/358

Heuvel, W. van den, *The Economic Lot-Sizing Problem: New Results and Extensions*, Promotor: Prof.dr. A.P.L. Wagelmans, EPS-2006-093-LIS, ISBN: 90-5892-124-7, http://hdl.handle.net/1765/1805

Hoedemaekers, C.M.W., *Performance, Pinned down: A Lacanian Analysis of Subjectivity at Work*, Promotors: Prof.dr. S. Magala & Prof.dr. D.H. den Hartog, EPS-2008-121-ORG, ISBN: 90-5892-156-7, http://hdl.handle.net/1765/10871

Hooghiemstra, R., *The Construction of Reality: Cultural Differences in Self-serving Behaviour in Accounting Narratives*, Promotors: Prof.dr. L.G. van der Tas RA & Prof.dr. A.Th.H. Pruyn, EPS-2003-025-F&A, ISBN: 90-5892-047-X, http://hdl.handle.net/1765/871

Hu, Y., *Essays on the Governance of Agricultural Products: Cooperatives and Contract Farming*, Promotors: Prof.dr. G.W.J. Hendrkse & Prof.dr. B. Krug, EPS-2007-113-ORG, ISBN: 90-5892-145-1, http://hdl.handle.net/1765/10535

Huij, J.J., *New Insights into Mutual Funds: Performance and Family Strategies*, Promotor: Prof.dr. M.C.J.M. Verbeek, EPS-2007-099-F&A, ISBN: 90-5892-134-4, http://hdl.handle.net/1765/9398

Huurman, C.I., *Dealing with Electricity Prices*, Promotor: Prof.dr. C.D. Koedijk, EPS-2007-098-F&A, ISBN: 90-5892-130-1, http://hdl.handle.net/1765/9399

Iastrebova, K, *Manager's Information Overload: The Impact of Coping Strategies on Decision-Making Performance*, Promotor: Prof.dr. H.G. van Dissel, EPS-2006-077-LIS, ISBN: 90-5892-111-5, http://hdl.handle.net/1765/7329

Iwaarden, J.D. van, *Changing Quality Controls: The Effects of Increasing Product Variety and Shortening Product Life Cycles,* Promotors: Prof.dr. B.G. Dale & Prof.dr. A.R.T. Williams, EPS-2006-084-ORG, ISBN: 90-5892-117-4, http://hdl.handle.net/1765/7992

Jansen, J.J.P., *Ambidextrous Organizations*, Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2005-055-STR, ISBN: 90-5892-081-X, http://hdl.handle.net/1765/6774

Jaspers, F.P.H., *Organizing Systemic Innovation*, Promotor: Prof.dr.ir. J.C.M. van den Ende, EPS-2009-160-ORG, ISBN: 978-90-5892-197-), http://hdl.handle.net/1765/14974

Jennen, M.G.J., *Empirical Essays on Office Market Dynamics*, Promotors: Prof.dr. C.G. Koedijk & Prof.dr. D. Brounen, EPS-2008-140-F&A, ISBN: 978-90-5892-176-5, http://hdl.handle.net/1765/13692

Jiang, T., *Capital Structure Determinants and Governance Structure Variety in Franchising*, Promotors: Prof.dr. G. Hendrikse & Prof.dr. A. de Jong, EPS-2009-158-F&A, ISBN: 978-90-5892-199-4, http://hdl.handle.net/1765/14975

Jiao, T., *Essays in Financial Accounting*, Promotor: Prof.dr. G.M.H. Mertens, EPS-2009-176-F&A, ISBN: 978-90-5892-211-3, http://hdl.handle.net/1765/1

Jong, C. de, *Dealing with Derivatives: Studies on the Role, Informational Content and Pricing of Financial Derivatives*, Promotor: Prof.dr. C.G. Koedijk, EPS-2003-023-F&A, ISBN: 90-5892-043-7, http://hdl.handle.net/1765/1043

Kaa, G. van, *Standard Battles for Complex Systems: Empirical Research on the Home Network,* Promotors: Prof.dr.ir. J. van den Ende & Prof.dr.ir. H.W.G.M. van Heck, EPS-2009-166-ORG, ISBN: 978-90-5892-205-2, http://hdl.handle.net/1765/1

Keizer, A.B., *The Changing Logic of Japanese Employment Practices: A Firm-Level Analysis of Four Industries,* Promotors: Prof.dr. J.A. Stam & Prof.dr. J.P.M. Groenewegen, EPS-2005-057-ORG, ISBN: 90-5892-087-9, http://hdl.handle.net/1765/6667

Kijkuit, R.C., *Social Networks in the Front End: The Organizational Life of an Idea,* Promotor: Prof.dr. B. Nooteboom, EPS-2007-104-ORG, ISBN: 90-5892-137-6, http://hdl.handle.net/1765/10074

Kippers, J., *Empirical Studies on Cash Payments,* Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2004-043-F&A, ISBN: 90-5892-069-0, http://hdl.handle.net/1765/1520

Klein, M.H., *Poverty Alleviation through Sustainable Strategic Business Models: Essays on Poverty Alleviation as a Business Strategy,* Promotor: Prof.dr. H.R. Commandeur, EPS-2008-135-STR, ISBN: 978-90-5892-168-0, http://hdl.handle.net/1765/13482

Knapp, S., *The Econometrics of Maritime Safety: Recommendations to Enhance Safety at Sea*, Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2007-096-ORG, ISBN: 90-5892-127-1, http://hdl.handle.net/1765/7913

Kole, E., *On Crises, Crashes and Comovements*, Promotors: Prof.dr. C.G. Koedijk & Prof.dr. M.J.C.M. Verbeek, EPS-2006-083-F&A, ISBN: 90-5892-114-X, http://hdl.handle.net/1765/7829

Kooij-de Bode, J.M., *Distributed Information and Group Decision-Making: Effects of Diversity and Affect*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2007-115-ORG, http://hdl.handle.net/1765/10722

Koppius, O.R., *Information Architecture and Electronic Market Performance*, Promotors: Prof.dr. P.H.M. Vervest & Prof.dr.ir. H.W.G.M. van Heck, EPS-2002-013-LIS, ISBN: 90-5892-023-2, http://hdl.handle.net/1765/921

Kotlarsky, J., *Management of Globally Distributed Component-Based Software Development Projects,* Promotor: Prof.dr. K. Kumar, EPS-2005-059-LIS, ISBN: 90-5892-088-7, http://hdl.handle.net/1765/6772

Krauth, E.I., *Real-Time Planning Support: A Task-Technology Fit Perspective,* Promotors: Prof.dr. S.L. van de Velde & Prof.dr. J. van Hillegersberg, EPS-2008-155-LIS, ISBN: 978-90-5892-193-2, http://hdl.handle.net/1765/14521

Kuilman, J., *The Re-Emergence of Foreign Banks in Shanghai: An Ecological Analysis,* Promotor: Prof.dr. B. Krug, EPS-2005-066-ORG, ISBN: 90-5892-096-8, http://hdl.handle.net/1765/6926

Kwee, Z., *Investigating Three Key Principles of Sustained Strategic Renewal: A Longitudinal Study of Long-Lived Firms,* Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2009-174-STR, ISBN: 90-5892-212-0, http://hdl.handle.net/1765/1

Langen, P.W. de, *The Performance of Seaport Clusters: A Framework to Analyze Cluster Performance and an Application to the Seaport Clusters of Durban, Rotterdam and the Lower Mississippi*, Promotors: Prof.dr. B. Nooteboom & Prof. drs. H.W.H. Welters, EPS-2004-034-LIS, ISBN: 90-5892-056-9, http://hdl.handle.net/1765/1133

Le Anh, T., *Intelligent Control of Vehicle-Based Internal Transport Systems*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr.ir. R. Dekker, EPS-2005-051-LIS, ISBN: 90-5892-079-8, http://hdl.handle.net/1765/6554

Le-Duc, T., *Design and Control of Efficient Order Picking Processes*, Promotor: Prof.dr. M.B.M. de Koster, EPS-2005-064-LIS, ISBN: 90-5892-094-1, http://hdl.handle.net/1765/6910

Leeuwen, E.P. van, *Recovered-Resource Dependent Industries and the Strategic Renewal of Incumbent Firm: A Multi-Level Study of Recovered Resource Dependence Management and Strategic Renewal in the European Paper and Board Industry*, Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2007-109-STR, ISBN: 90-5892-140-6, http://hdl.handle.net/1765/10183

Lentink, R.M., *Algorithmic Decision Support for Shunt Planning*, Promotors: Prof.dr. L.G. Kroon & Prof.dr.ir. J.A.E.E. van Nunen, EPS-2006-073-LIS, ISBN: 90-5892-104-2, http://hdl.handle.net/1765/7328

Li, T., *Informedness and Customer-Centric Revenue Management,* Promotors: Prof.dr. P.H.M. Vervest & Prof.dr.ir. H.W.G.M. van Heck, EPS-2009-146-LIS, ISBN: 978-90-5892-195-6, http://hdl.handle.net/1765/14525

Liang, G., *New Competition: Foreign Direct Investment and Industrial Development in China,* Promotor: Prof.dr. R.J.M. van Tulder, EPS-2004-047-ORG, ISBN: 90-5892-073-9, http://hdl.handle.net/1765/1795

Liere, D.W. van, *Network Horizon and the Dynamics of Network Positions: A Multi-Method Multi-Level Longitudinal Study of Interfirm Networks,* Promotor: Prof.dr. P.H.M. Vervest, EPS-2007-105-LIS, ISBN: 90-5892-139-0, http://hdl.handle.net/1765/10181

Loef, J., *Incongruity between Ads and Consumer Expectations of Advertising*, Promotors: Prof.dr. W.F. van Raaij & Prof.dr. G. Antonides, EPS-2002-017-MKT, ISBN: 90-5892-028-3, http://hdl.handle.net/1765/869

Maeseneire, W., de, *Essays on Firm Valuation and Value Appropriation,* Promotor: Prof.dr. J.T.J. Smit, EPS-2005-053-F&A, ISBN: 90-5892-082-8, http://hdl.handle.net/1765/6768

Londoño, M. del Pilar, *Institutional Arrangements that Affect Free Trade Agreements: Economic Rationality Versus Interest Groups*, Promotors: Prof.dr. H.E. Haralambides & Prof.dr. J.F. Francois, EPS-2006-078-LIS, ISBN: 90-5892-108-5, http://hdl.handle.net/1765/7578

Maas, A.A., van der, *Strategy Implementation in a Small Island Context: An Integrative Framework,* Promotor: Prof.dr. H.G. van Dissel, EPS-2008-127-LIS, ISBN: 978-90-5892-160-4, http://hdl.handle.net/1765/12278

Maeseneire, W., de, *Essays on Firm Valuation and Value Appropriation,* Promotor: Prof.dr. J.T.J. Smit, EPS-2005-053-F&A, ISBN: 90-5892-082-8, http://hdl.handle.net/1765/6768

Mandele, L.M., van der, *Leadership and the Inflection Point: A Longitudinal Perspective*, Promotors: Prof.dr. H.W. Volberda & Prof.dr. H.R. Commandeur, EPS-2004-042-STR, ISBN: 90-5892-067-4, http://hdl.handle.net/1765/1302

Meer, J.R. van der, *Operational Control of Internal Transport*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr.ir. R. Dekker, EPS-2000-001-LIS, ISBN: 90-5892-004-6, http://hdl.handle.net/1765/859

Mentink, A., *Essays on Corporate Bonds*, Promotor: Prof.dr. A.C.F. Vorst, EPS-2005-070-F&A, ISBN: 90-5892-100-X, http://hdl.handle.net/1765/7121

Meyer, R.J.H., *Mapping the Mind of the Strategist: A Quantitative Methodology for Measuring the Strategic Beliefs of Executives*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2007-106-ORG, ISBN: 978-90-5892-141-3, http://hdl.handle.net/1765/10182

Miltenburg, P.R., *Effects of Modular Sourcing on Manufacturing Flexibility in the Automotive Industry: A Study among German OEMs*, Promotors: Prof.dr. J. Paauwe & Prof.dr. H.R. Commandeur, EPS-2003-030-ORG, ISBN: 90-5892-052-6, http://hdl.handle.net/1765/1039

Moerman, G.A., *Empirical Studies on Asset Pricing and Banking in the Euro Area,* Promotor: Prof.dr. C.G. Koedijk, EPS-2005-058-F&A, ISBN: 90-5892-090-9, http://hdl.handle.net/1765/6666

Moitra, D., *Globalization of R&D: Leveraging Offshoring for Innovative Capability and Organizational Flexibility,* Promotor: Prof.dr. K. Kumar, EPS-2008-150-LIS, ISBN: 978-90-5892-184-0, http://hdl.handle.net/1765/14081

Mol, M.M., *Outsourcing, Supplier-relations and Internationalisation: Global Source Strategy as a Chinese Puzzle*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2001-010-ORG, ISBN: 90-5892-014-3, http://hdl.handle.net/1765/355

Mom, T.J.M., *Managers' Exploration and Exploitation Activities: The Influence of Organizational Factors and Knowledge Inflows*, Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2006-079-STR, ISBN: 90-5892-116-6, http://hdl.handle.net/1765

Moonen, J.M., *Multi-Agent Systems for Transportation Planning and Coordination*, Promotors: Prof.dr. J. van Hillegersberg & Prof.dr. S.L. van de Velde, EPS-2009-177-LIS, ISBN: 978-90-5892-216-8, http://hdl.handle.net/1

Mulder, A., *Government Dilemmas in the Private Provision of Public Goods*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2004-045-ORG, ISBN: 90-5892-071-2, http://hdl.handle.net/1765/1790

Muller, A.R., *The Rise of Regionalism: Core Company Strategies Under The Second Wave of Integration*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2004-038-ORG, ISBN: 90-5892-062-3, http://hdl.handle.net/1765/1272

Nalbantov G.I., *Essays on Some Recent Penalization Methods with Applications in Finance and Marketing*, Promotor: Prof. dr P.J.F. Groenen, EPS-2008-132-F&A, ISBN: 978-90-5892-166-6, http://hdl.handle.net/1765/13319

Nederveen Pieterse, A., *Goal Orientation in Teams: The Role of Diversity*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2009-162-ORG, http://hdl.handle.net/1765/15240

Nguyen, T.T., *Capital Structure, Strategic Competition, and Governance*, Promotor: Prof.dr. A. de Jong, EPS-2008-148-F&A, ISBN: 90-5892-178-9, http://hdl.handle.net/1765/14005

Niesten, E.M.M.I., *Regulation, Governance and Adaptation: Governance Transformations in the Dutch and French Liberalizing Electricity Industries*, Promotors: Prof.dr. A. Jolink & Prof.dr. J.P.M. Groenewegen, EPS-2009-170-ORG, ISBN: 978-90-5892-208-3, http://hdl.handle.net/1765/1

Nieuwenboer, N.A. den, *Seeing the Shadow of the Self*, Promotor: Prof.dr. S.P. Kaptein, EPS-2008-151-ORG, ISBN: 978-90-5892-182-6, http://hdl.handle.net/1765/14223

Ning, H., *Hierarchical Portfolio Management: Theory and Applications*, Promotor: Prof.dr. J. Spronk, EPS-2007-118-F&A, ISBN: 90-5892-152-9, http://hdl.handle.net/1765/10868

Noeverman, J., *Management Control Systems, Evaluative Style, and Behaviour: Exploring the Concept and Behavioural Consequences of Evaluative Style*, Promotors: Prof.dr. E.G.J. Vosselman & Prof.dr. A.R.T. Williams, EPS-2007-120-F&A, ISBN: 90-5892-151-2, http://hdl.handle.net/1765/10869

Oosterhout, J., van, *The Quest for Legitimacy: On Authority and Responsibility in Governance*, Promotors: Prof.dr. T. van Willigenburg & Prof.mr. H.R. van Gunsteren, EPS-2002-012-ORG, ISBN: 90-5892-022-4, http://hdl.handle.net/1765/362

Oostrum, J.M., van, *Applying Mathematical Models to Surgical Patient Planning*, Promotor: Prof.dr. A.P.M. Wagelmans, EPS-2009-179-LIS, ISBN: 978-90-5892-217-5, http://hdl.handle.net/1765/1

Paape, L., *Corporate Governance: The Impact on the Role, Position, and Scope of Services of the Internal Audit Function*, Promotors: Prof.dr. G.J. van der Pijl & Prof.dr. H. Commandeur, EPS-2007-111-MKT, ISBN: 90-5892-143-7, http://hdl.handle.net/1765/10417

Pak, K., *Revenue Management: New Features and Models*, Promotor: Prof.dr.ir. R. Dekker, EPS-2005-061-LIS, ISBN: 90-5892-092-5, http://hdl.handle.net/1765/362/6771

Pattikawa, L.H, *Innovation in the Pharmaceutical Industry: Evidence from Drug Introduction in the U.S.*, Promotors: Prof.dr. H.R.Commandeur, EPS-2007-102-MKT, ISBN: 90-5892-135-2, http://hdl.handle.net/1765/9626

Peeters, L.W.P., *Cyclic Railway Timetable Optimization*, Promotors: Prof.dr. L.G. Kroon & Prof.dr.ir. J.A.E.E. van Nunen, EPS-2003-022-LIS, ISBN: 90-5892-042-9, http://hdl.handle.net/1765/429

Pietersz, R., *Pricing Models for Bermudan-style Interest Rate Derivatives*, Promotors: Prof.dr. A.A.J. Pelsser & Prof.dr. A.C.F. Vorst, EPS-2005-071-F&A, ISBN: 90-5892-099-2, http://hdl.handle.net/1765/7122

Poel, A.M. van der, *Empirical Essays in Corporate Finance and Financial Reporting*, Promotors: Prof.dr. A. de Jong & Prof.dr. G.M.H. Mertens, EPS-2007-133-F&A, ISBN: 978-90-5892-165-9, http://hdl.handle.net/1765/13320

Popova, V., *Knowledge Discovery and Monotonicity*, Promotor: Prof.dr. A. de Bruin, EPS-2004-037-LIS, ISBN: 90-5892-061-5, http://hdl.handle.net/1765/1201

Pouchkarev, I., *Performance Evaluation of Constrained Portfolios*, Promotors: Prof.dr. J. Spronk & Dr. W.G.P.M. Hallerbach, EPS-2005-052-F&A, ISBN: 90-5892-083-6, http://hdl.handle.net/1765/6731

Prins, R., *Modeling Consumer Adoption and Usage of Value-Added Mobile Services*, Promotors: Prof.dr. Ph.H.B.F. Franses & Prof.dr. P.C. Verhoef, EPS-2008-128-MKT, ISBN: 978/90-5892-161-1, http://hdl.handle.net/1765/12461

Puvanasvari Ratnasingam, P., *Interorganizational Trust in Business to Business E-Commerce*, Promotors: Prof.dr. K. Kumar & Prof.dr. H.G. van Dissel, EPS-2001-009-LIS, ISBN: 90-5892-017-8, http://hdl.handle.net/1765/356

Quak, H.J., *Sustainability of Urban Freight Transport: Retail Distribution and Local Regulation in Cities*, Promotor: Prof.dr. M.B.M. de Koster, EPS-2008-124-LIS, ISBN: 978-90-5892-154-3, http://hdl.handle.net/1765/11990

Quariguasi Frota Neto, J., *Eco-efficient Supply Chains for Electrical and Electronic Products*, Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr.ir. H.W.G.M. van Heck, EPS-2008-152-LIS, ISBN: 978-90-5892-192-5, http://hdl.handle.net/1765/14785

Radkevitch, U.L, *Online Reverse Auction for Procurement of Services*, Promotor: Prof.dr.ir. H.W.G.M. van Heck, EPS-2008-137-LIS, ISBN: 978-90-5892-171-0, http://hdl.handle.net/1765/13497

Rinsum, M. van, *Performance Measurement and Managerial Time Orientation*, Promotor: Prof.dr. F.G.H. Hartmann, EPS-2006-088-F&A, ISBN: 90-5892-121-2, http://hdl.handle.net/1765/7993

Romero Morales, D., *Optimization Problems in Supply Chain Management*, Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Dr. H.E. Romeijn, EPS-2000-003-LIS, ISBN: 90-9014078-6, http://hdl.handle.net/1765/865

Roodbergen, K.J., *Layout and Routing Methods for Warehouses*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr.ir. J.A.E.E. van Nunen, EPS-2001-004-LIS, ISBN: 90-5892-005-4, http://hdl.handle.net/1765/861

Rook, L., *Imitation in Creative Task Performance*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2008-125-ORG, http://hdl.handle.net/1765/11555

Rosmalen, J. van, *Segmentation and Dimension Reduction: Exploratory and Model-Based Approaches*, Promotor: Prof.dr. P.J.F. Groenen, EPS-2009-165-MKT, ISBN: 978-90-5892-201-4, http://hdl.handle.net/1765/15536

Rus, D., *The Dark Side of Leadership: Exploring the Psychology of Leader Self-serving Behavior*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2009-178-ORG, http://hdl.handle.net/1765/1

Samii, R., *Leveraging Logistics Partnerships: Lessons from Humanitarian Organizations,* Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr.ir. L.N. Van Wassenhove, EPS-2008-153-LIS, ISBN: 978-90-5892-186-4, http://hdl.handle.net/1765/14519

Schaik, D. van, *M&A in Japan: An Analysis of Merger Waves and Hostile Takeovers,* Promotors: Prof.dr. J. Spronk & Prof.dr. J.P.M. Groenewegen, EPS-2008-141-F&A, ISBN: 978-90-5892-169-7, http://hdl.handle.net/1765/13693

Schauten, M.B.J., *Valuation, Capital Structure Decisions and the Cost of Capital,* Promotors: Prof.dr. J. Spronk & Prof.dr. D. van Dijk, EPS-2008-134-F&A, ISBN: 978-90-5892-172-7, http://hdl.handle.net/1765/13480

Schramade, W.L.J., *Corporate Bonds Issuers,* Promotor: Prof.dr. A. De Jong, EPS-2006-092-F&A, ISBN: 90-5892-125-5, http://hdl.handle.net/1765/8191

Schweizer, T.S., *An Individual Psychology of Novelty-Seeking, Creativity and Innovation,* Promotor: Prof.dr. R.J.M. van Tulder, EPS-2004-048-ORG, ISBN: 90-5892-077-1, http://hdl.handle.net/1765/1818

Six, F.E., *Trust and Trouble: Building Interpersonal Trust Within Organizations*, Promotors: Prof.dr. B. Nooteboom & Prof.dr. A.M. Sorge, EPS-2004-040-ORG, ISBN: 90-5892-064-X, http://hdl.handle.net/1765/1271

Slager, A.M.H., *Banking across Borders*, Promotors: Prof.dr. R.J.M. van Tulder & Prof.dr. D.M.N. van Wensveen, EPS-2004-041-ORG, ISBN: 90-5892-066–6, http://hdl.handle.net/1765/1301

Sloot, L., *Understanding Consumer Reactions to Assortment Unavailability*, Promotors: Prof.dr. H.R. Commandeur, Prof.dr. E. Peelen & Prof.dr. P.C. Verhoef, EPS-2006-074-MKT, ISBN: 90-5892-102-6, http://hdl.handle.net/1765/7438

Smit, W., *Market Information Sharing in Channel Relationships: Its Nature, Antecedents and Consequences*, Promotors: Prof.dr.ir. G.H. van Bruggen & Prof.dr.ir. B. Wierenga, EPS-2006-076-MKT, ISBN: 90-5892-106-9, http://hdl.handle.net/1765/7327

Sonnenberg, M., *The Signalling Effect of HRM on Psychological Contracts of Employees: A Multi-level Perspective*, Promotor: Prof.dr. J. Paauwe, EPS-2006-086-ORG, ISBN: 90-5892-119-0, http://hdl.handle.net/1765/7995

Speklé, R.F., *Beyond Generics: A closer Look at Hybrid and Hierarchical Governance*, Promotor: Prof.dr. M.A. van Hoepen RA, EPS-2001-008-F&A, ISBN: 90-5892-011-9, http://hdl.handle.net/1765/357

Stam, D.A., *Managing Dreams and Ambitions: A Psychological Analysis of Vision Communication*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2008-149-ORG, http://hdl.handle.net/1765/14080

Stienstra, M., *Strategic Renewal in Regulatory Environments: How Inter- and Intra-organisational Institutional Forces Influence European Energy Incumbent Firms,* Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2008-145-STR, ISBN: 978-90-5892-184-0, http://hdl.handle.net/1765/13943

Sweldens, S.T.L.R., *Evaluative Conditioning 2.0: Direct versus Associative Transfer of Affect to Brands,* Promotor: Prof.dr. S.M.J. van Osselaer, EPS-2009-167-MKT, ISBN: 978-90-5892-204-5, http://hdl.handle.net/1765/1

Szkudlarek, B.A., *Spinning the Web of Reentry: [Re]connecting reentry training theory and practice*, Promotor: Prof.dr. S.J. Magala, EPS-2008-143-ORG, ISBN: 978-90-5892-177-2, http://hdl.handle.net/1765/13695

Teunter, L.H., *Analysis of Sales Promotion Effects on Household Purchase Behavior*, Promotors: Prof.dr.ir. B. Wierenga & Prof.dr. T. Kloek, EPS-2002-016-MKT, ISBN: 90-5892-029-1, http://hdl.handle.net/1765/868

Tims, B., *Empirical Studies on Exchange Rate Puzzles and Volatility*, Promotor: Prof.dr. C.G. Koedijk, EPS-2006-089-F&A, ISBN: 90-5892-113-1, http://hdl.handle.net/1765/8066

Tuk, M.A., *Is Friendship Silent When Money Talks? How Consumers Respond to Word-of-Mouth Marketing*, Promotors: Prof.dr.ir. A. Smidts & Prof.dr. D.H.J. Wigboldus, EPS-2008-130-MKT, ISBN: 978-90-5892-164-2, http://hdl.handle.net/1765/12702

Valck, K. de, *Virtual Communities of Consumption: Networks of Consumer Knowledge and Companionship*, Promotors: Prof.dr.ir. G.H. van Bruggen & Prof.dr.ir. B. Wierenga, EPS-2005-050-MKT, ISBN: 90-5892-078-X, http://hdl.handle.net/1765/6663

Valk, W. van der, *Buyer-Seller Interaction Patterns During Ongoing Service Exchange*, Promotors: Prof.dr. J.Y.F. Wynstra & Prof.dr.ir. B. Axelsson, EPS-2007-116-MKT, ISBN: 90-5892-146-8, http://hdl.handle.net/1765/10856

Verheul, I., *Is There a (Fe)male Approach? Understanding Gender Differences in Entrepreneurship*, Prof.dr. A.R. Thurik, EPS-2005-054-ORG, ISBN: 90-5892-080-1, http://hdl.handle.net/1765/2005

Verwijmeren, P., *Empirical Essays on Debt, Equity, and Convertible Securities*, Promotors: Prof.dr. A. de Jong & Prof.dr. M.J.C.M. Verbeek, EPS-2009-154-F&A, ISBN: 978-90-5892-187-1, http://hdl.handle.net/1765/14312

Vis, I.F.A., *Planning and Control Concepts for Material Handling Systems*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr.ir. R. Dekker, EPS-2002-014-LIS, ISBN: 90-5892-021-6, http://hdl.handle.net/1765/866

Vlaar, P.W.L., *Making Sense of Formalization in Interorganizational Relationships: Beyond Coordination and Control,* Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2006-075-STR, ISBN 90-5892-103-4, http://hdl.handle.net/1765/7326

Vliet, P. van, *Downside Risk and Empirical Asset Pricing*, Promotor: Prof.dr. G.T. Post, EPS-2004-049-F&A, ISBN: 90-5892-07-55, http://hdl.handle.net/1765/1819

Vlist, P. van der, *Synchronizing the Retail Supply Chain*, Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr. A.G. de Kok, EPS-2007-110-LIS, ISBN: 90-5892-142-0, http://hdl.handle.net/1765/10418

Vries-van Ketel E. de, *How Assortment Variety Affects Assortment Attractiveness: A Consumer Perspective*, Promotors: Prof.dr. G.H. van Bruggen & Prof.dr.ir. A. Smidts, EPS-2006-072-MKT, ISBN: 90-5892-101-8, http://hdl.handle.net/1765/7193

Vromans, M.J.C.M., *Reliability of Railway Systems*, Promotors: Prof.dr. L.G. Kroon, Prof.dr.ir. R. Dekker & Prof.dr.ir. J.A.E.E. van Nunen, EPS-2005-062-LIS, ISBN: 90-5892-089-5, http://hdl.handle.net/1765/6773

Vroomen, B.L.K., *The Effects of the Internet, Recommendation Quality and Decision Strategies on Consumer Choice*, Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2006-090-MKT, ISBN: 90-5892-122-0, http://hdl.handle.net/1765/8067

Waal, T. de, *Processing of Erroneous and Unsafe Data*, Promotor: Prof.dr.ir. R. Dekker, EPS-2003-024-LIS, ISBN: 90-5892-045-3, http://hdl.handle.net/1765/870

Wall, R.S., *Netscape: Cities and Global Corporate Networks,* Promotor: Prof.dr. G.A. van der Knaap, EPS-2009-169-ORG, ISBN: 978-90-5892-207-6, http://hdl.handle.net/1765/1

Watkins Fassler, K., *Macroeconomic Crisis and Firm Performance,* Promotors: Prof.dr. J. Spronk & Prof.dr. D.J. van Dijk, EPS-2007-103-F&A, ISBN: 90-5892-138-3, http://hdl.handle.net/1765/10065

Weerdt, N.P. van der, *Organizational Flexibility for Hypercompetitive Markets: Empirical Evidence of the Composition and Context Specificity of Dynamic Capabilities and Organization Design Parameters,* Promotor: Prof.dr. H.W. Volberda, EPS-2009-173-STR, ISBN: 978-90-5892-215-1, http://hdl.handle.net/1765/1

Wennekers, A.R.M., *Entrepreneurship at Country Level: Economic and Non-Economic Determinants,* Promotor: Prof.dr. A.R. Thurik, EPS-2006-81-ORG, ISBN: 90-5892-115-8, http://hdl.handle.net/1765/7982

Wielemaker, M.W., *Managing Initiatives: A Synthesis of the Conditioning and Knowledge-Creating View*, Promotors: Prof.dr. H.W. Volberda & Prof.dr. C.W.F. Baden-Fuller, EPS-2003-28-STR, ISBN: 90-5892-050-X, http://hdl.handle.net/1765/1042

Wijk, R.A.J.L. van, *Organizing Knowledge in Internal Networks: A Multilevel Study*, Promotor: Prof.dr.ir. F.A.J. van den Bosch, EPS-2003-021-STR, ISBN: 90-5892-039-9, http://hdl.handle.net/1765/347

Wolters, M.J.J., *The Business of Modularity and the Modularity of Business*, Promotors: Prof. mr. dr. P.H.M. Vervest & Prof. dr. ir. H.W.G.M. van Heck, EPS-2002-011-LIS, ISBN: 90-5892-020-8, http://hdl.handle.net/1765/920

Yang, J., *Towards the Restructuring and Co-ordination Mechanisms for the Architecture of Chinese Transport Logistics*, Promotor: Prof.dr. H.E. Harlambides, EPS-2009-157-LIS, ISBN: 978-90-5892-198-7, http://hdl.handle.net/1765/14527

Yu, M., *Enhancing Warehouse Perfromance by Efficient Order Picking*, Promotor: Prof.dr. M.B.M. de Koster, EPS-2008-139-LIS, ISBN: 978-90-5892-167-3, http://hdl.handle.net/1765/13691

Zhang, X., *Strategizing of Foreign Firms in China: An Institution-based Perspective*, Promotor: Prof.dr. B. Krug, EPS-2007-114-ORG, ISBN: 90-5892-147-5, http://hdl.handle.net/1765/10721

Zhu, Z., *Essays on China's Tax System*, Promotors: Prof.dr. B. Krug & Prof.dr. G.W.J. Hendrikse, EPS-2007-112-ORG, ISBN: 90-5892-144-4, http://hdl.handle.net/1765/10502

Zwart, G.J. de, *Empirical Studies on Financial Markets: Private Equity, Corporate Bonds and Emerging Markets*, Promotors: Prof.dr. M.J.C.M. Verbeek & Prof.dr. D.J.C. van Dijk, EPS-2008-131-F&A, ISBN: 978-90-5892-163-5, http://hdl.handle.net/1765/12703

Erasmus Research Institute of Management - ERIM

ESSAYS ON PORT, CONTAINER, AND BULK CHEMICAL LOGISTICS OPTIMIZATION

The essays in this thesis are concerned with two main themes in port logistics. The first theme is the coordination of transport arrivals with the distribution processes and the use of storage facilities. We study this for both containerized and bulk chemical transport. The second theme is the uncertainty associated with the arrival time of ships with bulk chemicals and the impact on port logistics. Each essay describes a case study where quantitative methods, especially simulation, are used.

The operation of container terminals and in particular the way in which containers are stacked in a yard is influenced by information about the departure of a container. We find that even inaccurate information is valuable and helps to reduce unproductive moves.

Next, we present the ``floating stocks'' distribution concept which uses intermodal transport to deploy inventories in a supply chain in advance of retailer demand. We demonstrate that a main drawback of intermodal transport, a longer transit time, can be mitigated using this concept. This concept also influences the choice of a port: we provide a quantitative interpretation of routing flexibility in port selection.

ERIM

The Erasmus Research Institute of Management (ERIM) is the Research School (Onderzoekschool) in the field of management of the Erasmus University Rotterdam. The founding participants of ERIM are Rotterdam School of Management (RSM), and the Erasmus School of Economics (ESE). ERIM was founded in 1999 and is officially accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW). The research undertaken by ERIM is focussed on the management of the firm in its environment, its intra- and interfirm relations, and its business processes in their interdependent connections.

The objective of ERIM is to carry out first rate research in management, and to offer an advanced doctoral programme in Research in Management. Within ERIM, over three hundred senior researchers and PhD candidates are active in the different research programmes. From a variety of academic backgrounds and expertises, the ERIM community is united in striving for excellence and working at the forefront of creating new business knowledge.

ERIM PhD Series
Research in Management

Erasmus Research Institute of Management - **ERIM**
Rotterdam School of Management (RSM)
Erasmus School of Economics (ESE)
P.O. Box 1738, 3000 DR Rotterdam
The Netherlands

Tel. +31 10 408 11 82
Fax +31 10 408 96 40
E-mail info@erim.eur.nl
Internet www.erim.eur.nl

Design & layout: B&T Ontwerp en advies (www.b-en-t.nl) Print: Haveka (www.haveka.nl)