Erasmus University Rotterdam
Erasmus School of Economics

Bachelor's Thesis in Quantitative Logistics and Operations Research

BSc2 Econometrics/Economics

An emission and cost analysis on stochastic arrival times and multiple quay utilization in the berth allocation and crane assignment problem

Author: D.W.F. Boissevain (424964)
Supervisor: N. Milovanovic
Second assessor: Dr. T.A.B. Dollevoet

July, 2019

Abstract

In this thesis the results from the work of Meisel and Bierwirth (2009) around the berth allocation and crane assignment problem are replicated. In summary, a construction heuristic, local refinements, squeaky wheel optimization and tabu search are re-evaluated. Although using different data sets, the main findings of Meisel and Bierwirth are confirmed. Next to this, the framework is extended to include vessel emissions, which is used to test a new strategy handling stochastic arrival times of vessels and to perform an analysis on using multiple quays. First, the cost of saving emissions is evaluated, which is found to increase exponentially by restricting vessels from speeding up. Then, a buffer zone recourse strategy is developed to deal with uncertainty of arrival times, which runs faster but performs worse than the SWO. Lastly, it is found that using one extra quay significantly reduces vessel emissions and increasingly reduces solution costs as instance sizes grow.

Contents

1 Introduction

Most of the essential products we need such as as electronics, clothing and food arrive to us over seas. More specifically, around 80% of the total volume of international trade is transported by vessel (United Nations Conference on Trade and Development, 2015). In this considerable volume, a significant number of containers are used, as their shape makes their transportation very efficient (Port of Antewerp, 2019). These containers account for a significant percentage of port throughput. In Rotterdam for instance, almost a third of the total throughput by gross weight in 2018 was accounted for by containers (Port of Rotterdam, 2019).

To handle these great volumes, ports need to handle the incoming container ships as efficiently as possible. That is, the requested number of containers should be loaded and unloaded as fast as possible. Fast service times ask for an efficient division of the available space at the quay and an effective allocation of quay cranes (QC) to unload the containers. Allocating berths and assigning cranes to the arriving container ships can be seen as two separate problems, namely the Berth Allocation Problem (BAP) and the Quay Crane Assignment Problem (QCAP).

The BAP has the goal to assign a berthing location and a berthing time to every incoming ship within a given planning horizon. The decision to plan a certain incoming vessel depends on its size and arrival time. When these are known, a solution for this problem can be portrayed by a space-time diagram as shown in Figure 1, where the vertical axis represents the quay and the horizontal axis stands for the time. Moreover, the lengths of the squares represent the length of the ships and the width of the squares its service time. For a berthing assignment solution to be feasible, the squares may not overlap to avoid vessels being at the same quay position at the same time. The berthing solution in Figure 1 is thus feasible.

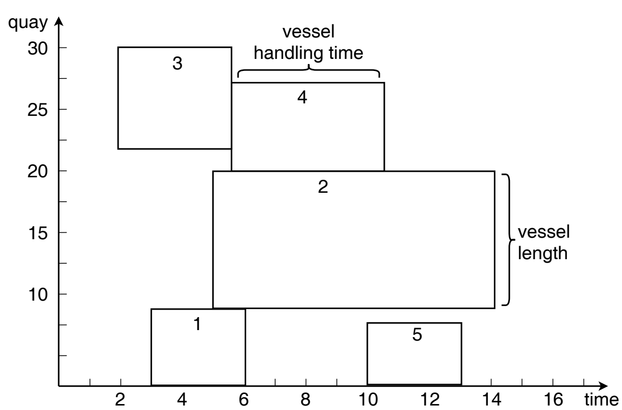

Figure 1: A feasible berth allocation

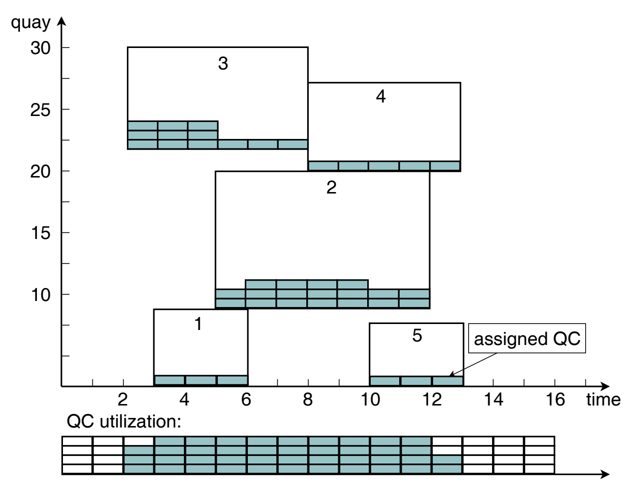

Figure 2: A feasible crane assignment

Next to the assignment of berthing positions and times, the QCAP allocates the cranes to the moored vessels. In this problem, there is a scarce amount of cranes that can be exhausted. Here, there are four cranes as seen from the QC utilization in Figure 2. In the same diagram, the grey squares stand for an hour of crane unloading allocated to a ship. Moored vessels have a minimum number of servicing cranes due to contracts and a maximum number of cranes allowed to unload the ship due to technicalities and the size of the ship. The number of cranes can change during an unloading process and assigning more cranes to a ship naturally speeds up its service time. This is illustrated in Figure 2 by assigning more cranes to vessel 2 at the expense of vessel 3. By using a different crane assignment, the berth time of vessel 4 must also be moved.

What is evident by the shifted berth time of vessel 4 is that another allocation of the QCAP could result in an unfeasible solution of the BAP, meaning the two problems are intertwined. Therefore, it is reasonable to solve the problem by using an integrated approach to the BAP and the QCAP, named the Berth Allocation and Crane Assignment Problem (BACAP). More specifically, the integrated approach of Meisel and Bierwirth (2009) will be replicated in this thesis. This framework can then be used to extend the problem domain. The goal of this thesis is therefore to enlarge the BACAP setting by adding three features: vessel emissions, arrival uncertainty and multiple quays.

First, vessel emissions as a result of the found BACAP solutions are not always measured and considered within the planning process. Moreover, it has only been treated in a handful of papers within the literature. Nonetheless, residents living around the area of a port are affected by vessel air pollution. As Bailey and Solomon (2004) note, increasing global trade leads to increasing levels of sulphur, nitrogen oxides and particle dust within ports, which has a significant effect on the health of the residents living around such areas. The pollution could cause asthma, heart problems and different forms of cancer. By extending the BACAP with emission measures, total emissions of all vessels entering the port can be measured. This metric could help to limit unnecessary diesel exhausts and thus reduce severe health effects for residents living around high transport areas. Therefore, the framework of the BACAP in this thesis will be extended to include vessel emissions.

Secondly, container terminal operators might face uncertainty due to changing arrival times, as estimated arrival times may not always be met due to weather, financial or technical conditions. Just as vessel emissions, stochastic arrivals are treated within the literature only scarcely. Uncertainty is mainly a problem for container terminal operators and vessel operators themselves. A ship being too late will namely cause a shift in the whole schedule of the terminal operators, while the vessel operators planned after a vessel with an arrival change will have to postpone their service. More specifically, large ships with variable arrival time could cause considerable bottlenecks as these vessels arriving late lead to more ships to postpone their arrival within the schedule. When the problem of arrival uncertainty can be solved, the terminal can handle more large container ships. In this way, vessel operators could benefit from economies of scale by enlarging the sizes of their vessels. Sources of economies of scale in large container ships include greater flexibility in the stowage of containers, increased weight per container slot due to increased stability and faster loading times due to decreased motion while at the quay (Cullinane & Khanna, 2000). To provide quick and effective solutions to stochastic arrival times, a new strategy will therefore be proposed.

Lastly, the incorporation of multiple quays in the problem is a feature not looked into in the literature. When solving the BACAP, all papers focus on a single quay, while in reality the problem should be solved for the port as a whole. Next to this, the solution cost analysis that can be performed after extending the BACAP by multiple quays can be used as an argument for investment in future quays. This can help ports make calculated decisions when determining whether to enlarge their container handling capacity.

When including these three main extensions, the research question for this thesis can thus be formulated by:

What are the costs of saving vessel emissions and the consequences of incorporating random arrival time and the addition of multiple quays in terms of solution cost and emissions in the Berth Allocation and Crane Assignment Problem?

In Section 2, classifications of the berth and crane problem will be given by literature performed before the work of Meisel and Bierwirth (2009). Subsequent research on vessel emissions, uncertainty and multiple quays will also be discussed here. Then, Section 3 discusses the methodology of the work done by Meisel and Bierwirth (2009) and explains the three main extensions of this thesis. After this, the data will be treated in Section 4. Finally, the results will be treated in Section 5 and a conclusion will be given in Section 6.

2 Literature

2.1 Previous research

Berth allocation models have been studied extensively before the research of Meisel and Bierwirth (2009). For these studies it is worthy to note that the BAP can be classified into a number of classes as described by Imai et al. (2001). The most important distinction is that the BAP can be discrete or continuous. In the discrete problem, the quay exists of a discrete number of berthing locations, while in the continuous case an incoming vessel can moor at any available spot at the quay, given there is space. Moreover, the BAP can have static or dynamic arrival times. Static problems ignore the arrival times of incoming ships. These problems are solved when making a schedule for ships that are currently all ready to enter the port. Dynamic problems are characterised by a constraint of the earliest arrival times of vessels, which means it is also possible to include vessels that are further away from their destination.

Next to these two characteristics a number of other BAP properties are described in the survey paper of Bierwirth and Meisel (2010). Here, the handling times of vessels are described to either be fixed (Guan & Cheung, 2004), dependent on the berthing position of the ship (Chang et al., 2008) or dependent on the assignment of the available cranes (Meisel & Bierwirth, 2009). Lastly, to evaluate whether a good berthing solution has been found, the allocation has been evaluated on the waiting, handling and completion time of vessels (Meier & Schumann, 2007).

The QCAP also has different properties, which have been described in the literature. For instance, safety margins between the cranes (Jung et al., 2006), non-crossing constraints (Meisel, 2009) and productivity loss due to potential crossing when unloading the same ship (Meisel & Bierwirth, 2009) are often used in problem formulations. Lastly, performance of the model has mostly been evaluated by the finishing time (Kim & Park, 2004), throughput (Lim et al., 2004), and movement per crane (Ak & Erera, 2006).

2.2 Subsequent research

Subsequent research performed after the paper of Meisel and Bierwirth (2009) includes many new extensions which include existing and new developments in the container shipping field. For the BAP a range of problems has been taken up, such as mobile and indented quays, tidal accessibility and handling priority which can all be found in the survey of Bierwirth and Meisel (2015).

2.2.1 Vessel emissions

An existing development is the rising pollution and its health effects in areas around ports (Bailey & Solomon, 2004). Interestingly though, this topic has been extended in only a handful of papers. One of the first of these papers is from Du et al. (2011), where an emission estimation model is developed which converts the vessel arrival time into fuel consumption and gas emissions. In this model, both the emissions obtained from mooring and sailing to the berth are taken into account. Together with a computation to obtain the optimal sailing speeds, a new strategy using a mixed integer second order cone programming model is proposed where the fuel and gas emissions are substantially decreased while retaining the same service level. Nevertheless, this paper only concerns the BAP problem.

A second paper which concerns the emissions of incoming container ships is from Venturini et al. (2017). Here, the idea of Du et al. (2011) is extended to incorporate all available routes over the network. This enables reduction of fuel consumption and emissions during the whole trip. Still, this model does not allocate quay cranes as it concerns the BAP.

To the best of our abilities, the only paper that could be found where fuel consumption is incorporated in a complete BACAP is Hu et al. (2014). In this work, a second-order mixed-integer cone model is used to more quickly solve the BACAP, where a vessel's fuel consumption and emissions are considered. The proposed strategy improves fuel consumption and utilization of berths and cranes without reducing vessel service quality. This being the most complete model there is, it misses the operational costs of crane utilization and the crane interference as described by Meisel and Bierwirth (2009). Furthermore, the methods only concern an exact model without applications for meta-heuristics.

Therefore, we extend the BACAP of Meisel and Bierwirth (2009) with vessel fuel consumption and emissions cost factors, while keeping the crane operation cost and the crane interference effect within the objective. Furthermore, the problem is not solved through exact cone models, but through the meta-heuristics from Meisel and Bierwirth (2009). This enables analysis of the costs of emission reduction as the heuristics can easily be altered. The procedures for this analysis are explained in Section 3.3.

2.2.2 Uncertainty of arrival times

Next to the problem of vessel emission, the most important direction for future research which Meisel and Bierwirth (2009) mention is the inclusion of stochastic elements to assess the robustness of schedules. Here they hint on the variability of ship arrival times and handling time. This direction is also highlighted in their follow-up survey, where it is observed that most work on uncertain data only concerns discrete BAP's. The reason for this is that a discrete number of berths handles uncertain arrivals easier. Furthermore, papers often use methods to introduce randomness in their model that require assumptions on data availability.

The first work where randomness of arrivals and handling time is introduced is Lu et al. (2010). A genetic algorithm is used in combination with Monte Carlo simulation to find robust schedules. The problem, however with simulation is that it may take a long time and the probability distributions have to be known. Moreover, the problem is also a discrete case instead of a continuous quay.

Zhen et al. (2011) studies the BAP under stochastic arrivals and handling time. A scenario based approach is used to generate the random arrivals and handling times. When a scenario is known, a two-stage model is used together with heuristics for large problem instances. The problem with a scenario based approach however, is that a large set of parameters and historical data is needed to identify the scenario. Therefore, obtaining results from these methods require observing scenario processes over a long time.

Lastly, the literature also uses sets of lower and upper bounds to model stochastic arrivals and handling times. Two papers where this is done are Golias et al. (2014) and Xu et al. (2012). Nevertheless, these papers only concern BAP's and exact solution techniques. Therefore, in this research, each vessel will have a buffer time behind its ETA, reserving time for when the ship would arrive later than expected. To the best of our abilities, this method of including uncertainty in arrivals has not been performed in a BACAP setting yet. Therefore, the introduction of stochastic arrivals via buffer zones will be incorporated within the BACAP framework.

2.2.3 Multiple quays

The last problem feature that is not observed in the literature is the inclusion of multiple quays. Namely, in all research found quays are looked at in isolation. Imai et al. (2008) perform a genetic algorithm on a BAP with a main terminal and an external terminal where a container can resort to. This external terminal acts as a back-up quay whenever a vessel's maximum waiting time is exceeded. The algorithm is able to reduce waiting time, but the paper does not concern the assignment of quay cranes and the inclusion of more than two quays is not inspected. Therefore, it would also be of interest to extend the BACAP setting to include multiple quays.

3 Methodology

This section provides the problem description of the BACAP and discusses the methods of Meisel and Bierwirth (2009) used to solve the problem. Meisel and Bierwirth use an exact model, which they solve with CPLEX and compare to a construction heuristic, local refinements, squeaky wheel optimization and tabu search. For this thesis, the last four heuristics are replicated. Finally, the methods used for extending their work are also treated.

3.1 Problem description

3.1.1 Notation

To perform the heuristics, the following input data and decision variables are of importance. Each vessel has a length, a desired berthing position, a crane demand, a minimum and maximum number of serving cranes and cost factors associated with the vessel type. Moreover, each vessel has an estimated arrival time and an arrival time if the vessel would be sped up. The two deadlines a vessel has are the expected finishing time and its latest finishing time. The interference and berth deviation factor will be explained in Section 3.1.3. To create a schedule, a vessel's berthing position, starting and ending time have to be known. Vessels are allowed to speed up to reach its berthing time and it must be known whether a ship is too late and how many cranes are assigned to it per hour. With the last four variables the objective value of the problem instance can be determined.

Table 1: Notation for the berth allocation and crane assignment problem

Input data		*Decision variables*	
V	set of vessels to be loaded and unloaded, $V = \{1, 2, \ldots, n\}$	b_i	berthing position of vessel $i \in V$
Q	number of available cranes	s_i	starting time of vessel $i \in V$
L	number of 10 meter berth segments	e_i	ending time of vessel $i \in V$
T	set of periods, $T = \{0, 1, \ldots, H-1\}$, with time horizon H	Δb_i	berthing position deviation of vessel $i \in V$, $\Delta b_i = \lvert b_i^0 - b_i \rvert$
l_i	length of vessel $i \in V$ in 10 meter segments	ΔETA_i	speed up of vessel $i \in V$, $\Delta\text{ETA}_i = (\text{ETA}_i - s_i)^+$
b_i^0	desired berthing position of vessel $i \in V$	ΔEFT_i	hours too late for vessel $i \in V$, $\Delta\text{EFT}_i = (e_i - \text{EFT}_i)^+$
m_i	crane demand of vessel $i \in V$	u_i	equals 1 if the ending time exceeds LFT_i, 0 otherwise
$r_i^{\min}$	minimum number of cranes to serve vessel $i \in V$	r_{itq}	equals 1 if q cranes are assigned at time t, 0 otherwise
$r_i^{\max}$	maximum number of cranes to serve vessel $i \in V$		
ETA_i	expected time of arrival of vessel $i \in V$		
EST_i	earliest starting time when vessel $i \in V$ is sped up		
EFT_i	expected finishing time of vessel $i \in V$		
LFT_i	latest finishing time of vessel $i \in V$		
c_i^1, c_i^2, c_i^3	service cost factors for vessel $i \in V$ in 1000 USD per hour		
c^4	operation cost rate in 1000 USD per crane hour		
α	interference exponent		
β	berth deviation factor		

3.1.2 Objective function

The objective value Z for the problem formulation of Meisel and Bierwirth (2009) can be given by

$$Z = \sum_{i \in V} \left(c_i^1 \Delta \mathrm{ETA}_i + c_i^2 \Delta \mathrm{EFT}_i + c_i^3 u_i + c^4 \sum_{t \in T} \sum_{q \in R_i} q \cdot r_{itq} \right), \tag{1}$$

which includes service quality and operator costs. For the service costs, three main cost factors are distinguished. The cost c_i^1 when having to speed up a vessel to reach the port earlier than its earliest arrival time ETA_i, the delay cost c_i^2 for exceeding the expected finishing time EFT_i, and the penalty cost c_i^3 for exceeding the latest finishing time LFT_i of vessel i. The speedup cost c_i^1 and the delay cost c_i^2 increase over time, while the penalty cost c_i^3 is penalized only once. The development of the service costs can bee seen in Figure 3. When the vessel does not have to speed up, or is finished not later than its expected finishing time EFT_i and latest finishing time LFT_i, the vessel incurs no service quality cost.

The final cost in the objective is the cost c_i^4 for using a crane hour. The sum of all crane hours multiplied by this cost then makes up the operator cost. Summing the operator cost with the service quality cost over all vessels makes up the objective function Z.

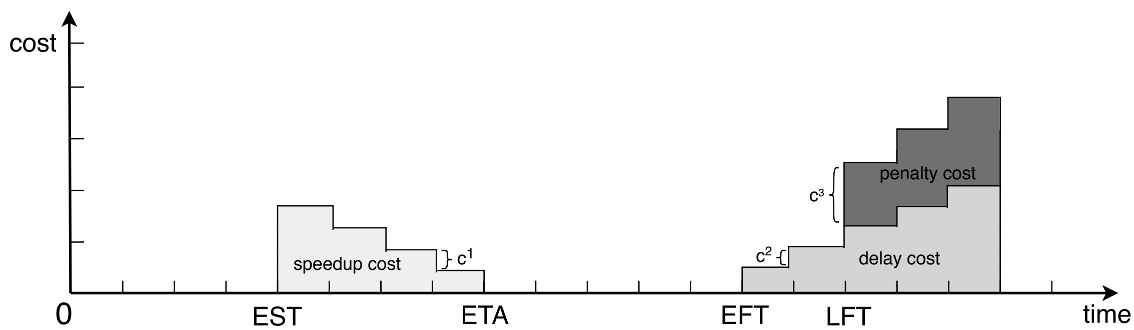

Figure 3: Service quality cost structure

3.1.3 Productivity effects

Nevertheless, in the formulation of Meisel and Bierwirth there exists a catch when assigning as much cranes as possible to avoid the delay and penalty costs. Namely, the used formulation introduces an interference effect on the cranes working on the same vessel i. The productivity obtained from assigning q crane hours to vessel i is expressed as q^α crane hours, where α $(0 < \alpha \leq 1)$ is the interference exponent. The productivity thus does not increase linearly by assigning more cranes, which makes a quick service time more expensive.

Next to the productivity hit from assigning multiple cranes to a vessel i, the productivity of a vessels service time is also affected by how far away the vessel is moored from its ideal berthing position b_i^0. When the vessel is moored further away and Δb_i increases, the crane demand increases for vessel i. The reason for this is that more transport vehicles are needed to move the containers to its preferred destination in the yard, as this destination is further away. This increase in transport vehicles at the quay reduces the average vehicle speed and so a larger Δb_i can thus be modeled by an increased crane demand. In their paper, Meisel and Bierwirth use a berth deviation factor $\beta \geq 0$, which denotes the increase in crane demand per berthing deviation in 10 meters. With this factor, a vessel which is moored Δb_i segments further from the desired berthing position needs $(1 + \beta \cdot \Delta b_i) \cdot m_i$ crane hours.

Together with the interference effect, a minimum duration needed to serve vessel i can be defined as

$$d_i^{\min} = \left\lceil \frac{(1 + \beta \cdot \Delta b_i) \cdot m_i}{(r_i^{\max})^{\alpha}} \right\rceil . \tag{2}$$

To illustrate, let the crane demand of vessel i be 15 crane hours, where the quay has 5 cranes and the vessel is scheduled 100 meters away from its desired position ($\Delta b_i = 10$). Without productivity effects, the minimum duration would be $\lceil 15/5 \rceil = 3$ crane hours. With productivity effects where $\alpha = 0.85$ and $\beta = 0.02$, the minimum handling time increases to $d_i^{\min} = \left\lceil \frac{(1+\beta \cdot \Delta b_i) \cdot m_i}{(r_i^{\max})^{\alpha}} \right\rceil = \left\lceil \frac{(1+0.02 \cdot 10) \cdot 15}{5^{0.85}} \right\rceil = 5$ crane hours.

3.2 Heuristic methods

3.2.1 Construction heuristic

As the exact CPLEX model in the work of Meisel and Bierwirth takes a long time to solve for large instances, a construction heuristic is used as a basis for all their methods. In this heuristic, all vessels get inserted into the space time diagram in order of early to late arrival times. When vessel i is inserted, it gets assigned a berthing time s_i, an ending time e_i, a berthing position b_i, the number of cranes assigned to the vessel at period t and a cost. Figure 4 shows the construction heuristic, which performs some basic steps for each inserted vessel i.

In the initialization step (a), the cost of current vessel i is initially set to infinity. When this is done, the starting time s_i of vessel i is immediately set to its ETA in step (b) and the berthing position b_i to its preferred berthing position b_i^0 in step (c), as this assignment is the most cost efficient.

Then, a crane assignment will be made in step (d) for the current berthing time s_i and position b_i. By using the expression for $d_i^{\min}$ in equation 2, the fastest service time is computed which also enables the computation of the ending time by $e_i = b_i + d_i^{\min}$. In the service interval from s_i to e_i, the maximum possible number of cranes is assigned to the vessel for each hour. If this maximum is not available for one of the service hours, the number of assigned cranes is decreased until it reaches a number of cranes which is available. When the number of available cranes is below the minimum allowable cranes $r_i^{\min}$, the assignment is infeasible and a new berthing time is selected in step (g). Moreover, it could also be that the number of assigned crane hours does not meet the vessel's crane demand m_i within the given service time from s_i to e_i. In this case, the ending time e_i is incremented such that more cranes can be assigned for an extra hour. This is done until an assignment can be made which meets the crane demand. When $e_i > H$, the crane assignment is also unfeasible and a new berthing time is selected in step (g).

When cranes have successfully been assigned to the vessel, the ending time e_i is fixed. The service time from s_i to e_i together with the ships length l_i enables us to check whether the current vessel overlaps with other already inserted vessels in the space time diagram. If it does overlap, a new berthing position b_i is selected in step (f), where the ship is moved upward until quay position $L - l_i$ and downward until position 0. The new berthing position b_i is selected which is closest to b_i^0 which had not been checked yet and the crane assignment in step (d) is repeated. If no b_i can be found, the heuristic continues with a new berthing time s_i in step (g). On the other hand, when the insertion does not overlap, the cost of the ship Z_i is computed using the procedures in section

3.1.2 in step (e). If the cost is better than its current cost, the berthing time s_i, position b_i, ending time e_i, crane assignment and its best cost Z_i^* get updated. After this is done, the heuristic also continues with step (g).

In step (g) a new starting time s_i is selected to find a possible better or feasible insertion. These starting times are taken from a list $[\mathrm{ETA}_i - 1, \mathrm{ETA}_i + 1, \mathrm{ETA}_i - 2, \dots]$ where the s_i's should always bigger than the vessel's earliest starting time EST_i and may only be less than the time horizon H. The new starting time is evaluated by beginning in step (c) and continuing from there as described in this section. If no s_i is left anymore in the list, the insertion procedure ends and returns the best found solution for vessel i. A procedural example can be found in the paper of Meisel and Bierwirth (2009).

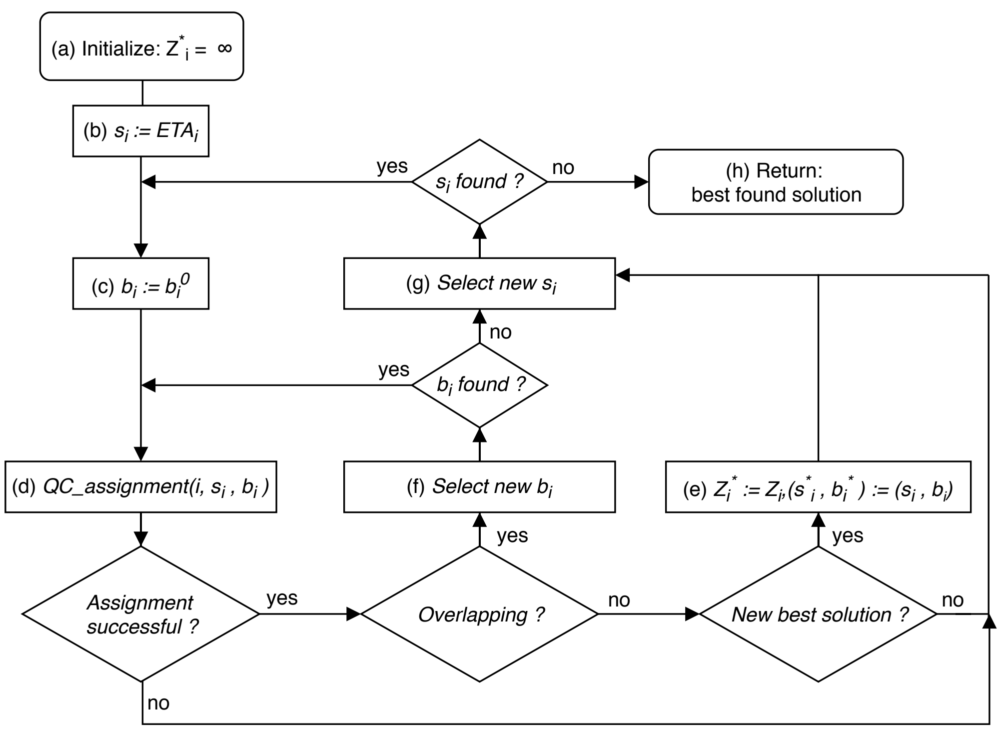

Figure 4: The construction heuristic

3.2.2 Local refinements

To improve the solution of the construction heuristic, two refinements are made after inserting all the vessels into the space time diagram. First, crane resources are leveled over all vessels. The idea behind crane resource leveling is that vessels that are first inserted have a good chance to get a crane assignment with a maximum number of cranes assigned to them. As this favorable crane assignments puts the other vessels in a disadvantage, it is also possible to restrict the maximum allowable cranes to vessel i to $r_i^{|v|}$ below $r_i^{\max}$. This restriction enables other vessels which are inserted later on to also have some cranes available to them.

Therefore, the first refinement procedure considers the construction heuristic which inserts vessels p_i with $i \in V$ in a priority list $P = (p_1, p_2, \dots, p_n)$. First, vessel p_1 is inserted for all crane resource levels $r_{p_1}^{|v|}$ where v runs from $r_1^{\min}$ to $r_1^{\max}$ in different berth plans. For each of these plans, the remaining vessels $p_2, \dots, p_n$ are inserted without a restriction. It could be that there still are left over cranes for vessel 1 after inserting all the remaining vessels, so therefore vessel 1 is removed from the plans and inserted again without a resource restriction. After computing all the costs of the berthing plans with different resource restrictions for vessel 1, the best plan is chosen and vessel 1 is fixed to a partial berthing plan. This partial berthing plan is extended by doing the same procedure but for vessel 2. We insert vessel 2 to the partial berthing plan containing vessel 1 for every resource level v and insert the remaining vessels. Then, vessel 2 is removed again and inserted without restriction and the best plan is chosen. This procedure is performed for all vessels in priority list P until the partial berthing plan gets completed.

The second local refinement which is performed is the spatial and temporal cluster shifting. A spatial cluster is a cluster of vessels that are connected on the space time diagram on the top or the bottom, meaning the vessels are scheduled next to each other at the quay during the same time. A temporal cluster is a cluster of vessels that are connected on a vessels left or right border in the diagram, meaning the vessels are scheduled next to each over time at the same quay position. The idea behind cluster shifting is that one vessel in a cluster could have been scheduled in its ideal way, while the second connected vessel had a similar ETA and desired berthing position b_i^0, meaning its cost is increased as the other vessel is in the way. If we could shift both vessels in a direction such that the first vessel's costs increases by a small amount, but the second vessel's cost decreases substantially, we can reduce the costs of the overall solution.

The cluster shifting performs shifts of one hour and one quay segment towards both sides of the planning interval and the quay respectively. Every time the cluster is moved, the crane assignment is performed again for all ships as the berth deviation could change and another number of cranes could be available during the new service time. When a crane assignment is unfeasible for a vessel, it is reinserted with the resource level $r_i^{|v|}$ from the resource leveling refinement. The procedure ends when all clusters have been shifted in the time or space axis, depending on which type of cluster it is. In the end, the cluster position with the lowest cost is chosen for each cluster. To speed up the procedure though, a cluster is not shifted anymore after 20 shifts without finding a better solution.

3.2.3 Squeaky wheel optimization

To improve the solutions, Meisel and Bierwirth also use two meta-heuristics. The first one being a squeaky wheel optimization (SWO). In a SWO, vessels contributing a large part of the costs are prioritized over vessels with low costs. First, a priority list P is used with a vessel order based on ETA. With this list, first a construction heuristic is performed together with the local refinements, which leads to costs for every individual vessel.

By looking at the costs of these vessels, changes can be made in the priority list. Namely, two vessels are swapped in the list when the first vessel has a smaller service cost than the second vessel. Only service costs are used to avoid any bias from a large crane demand m_i. These swaps are performed $n-1$ times starting from the beginning of the list P. After these swaps, the construction and local refinement heuristics are performed for the changed priority list P. Then, changes are made again, but now based upon the average costs over all iterations. In this way, the vessels are given priority based on the performance in all solutions so far. Furthermore, whenever a new best solution has been found, the best performing priority list P gets updated.

It could be that the priority list does not change after one iteration of construction and local refinements. In this case, the SWO will only perform the construction heuristic without local refinements in the next iteration to allow for slightly worse solutions, which enable changes in the priority list P. The local refinements are activated again when a priority list P has been made which has not been seen before. The SWO ends after a given number of iterations without seeing a new best solution.

3.2.4 Tabu search

The last meta-heuristic Meisel and Bierwirth propose to solve the BACAP is a tabu search (TS). It also performs the construction and local refinements given a priority list P, but here TS makes pairwise exchanges over the whole list instead of swaps. This procedure enables the TS to explore the whole neighborhood of the current solution. The exploration of the whole neighborhood was left somewhat ambiguous by Meisel and Bierwirth, but here every vessel is exchanged with all vessel in front of it as can be seen in Figure 5 for exchanges of vessel 1 and 2. In this way, every exchange is evaluated by the construction heuristic and local refinement is only done for the best performing neighboring solution of all the construction heuristic solutions. Moreover, the best found priority list will be saved and all priority list evaluated are put in a tabu set such that they are not evaluated anymore. When the local refinement is finished for the best performing neighbor list, the TS performs a neighborhood search again on that priority list with the construction heuristic. This search continues until no better solution can be found for a number of iterations.

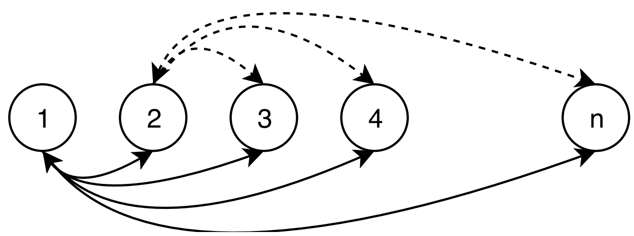

Figure 5: Pair-wise exchange procedure of the tabu search for vessels 1 and 2

3.3 Vessel emissions

Then, the framework will be extended to perform a vessel emissions cost analysis based on the computations and data from Du et al. (2011). When a ship is sped up ($\Delta ETA > 0$), extra fuel has to be burned to accelerate the ship. To compute the extra fuel burned, first the fuel consumption in gallons per hour (r_F) can be described by equation (3). This means that the fuel consumption is mostly dependent on the speed (v) of the vessel (Schrady & Wadsworth, 1991):

$$r_F = c^0 + c^1 v^\mu. \tag{3}$$

Here, c^0 and c^1 are estimation parameters and μ equals 3.5 for Feeder container ships, 4 for medium sized ships and 4.5 for Jumbo container ships as suggested by MAN Diesel & Turbo (2008). The parameters c^0 and c^1 can be estimated by ordinary regression and are given by Du et al. (2011) in Table 2 below.

When the parameters are known, the number of gallons of fuel burned (F_i) by a vessel i can be computed by the formula provided by Du et al. (2011):

$$F_i = r_F s_i = \left[c_i^0 + c_i^1 (v_i)^{\mu_i}\right] s_i = \left[c_i^0 + c_i^1 \left(\frac{k_i}{s_i}\right)^{\mu_i}\right] s_i, \tag{4}$$

where k_i is the distance that has to be traveled to the port. As the distance is not available in the given data, it will be approximated by $k_i = ETA_i \cdot v_i^0$, where v_i^0 is a vessel's design speed given in Table 2. This equation is not an ideal metric, as we neglect the distance traveled before the time of the schedule creation at time zero. In this way, vessel i with a larger ETA$_i$ emits more although it could have travelled less distance than other vessels. Nevertheless, ETA$_i$ is believed to be the best available measure for approximating the distance yet to be covered by the vessel as it stands for the time it takes to sail directly to the quay at design speed. Lastly, s_i is the chosen arrival time for vessel i as described in Section 3.1. This means that if s_i is chosen to be lower, a vessel is sped up as $v_i = \frac{k_i}{s_i}$ increases, which increases the fuel consumption F_i.

With the number of gallons of fuel burned by a vessel, the total emissions can now be computed by using emission factors. This is 3.154 kg fuel per fuel gallon burned and 3.179 kg/kg-fuel for CO_2, in the case of 'in port' container ship operations as estimated by Entec UK in request of the European Commission (Entec UK Ltd, 2002). The CO_2 emissions (E_i) of vessel i will thus be computed by

$$E_i = 3.154 \cdot 3.179 F_i. \tag{5}$$

Now the CO_2 emissions for each sped up vessel can be computed, an emission cost analysis can be performed by using the SWO. First, a SWO solution where all ships are allowed to arrive between their EST and ETA is performed. Then, a solution is obtained where all ships except one are allowed to arrive between their EST and ETA. Then a solution with all ships except two, until we have V solutions where in each solution one extra ship is restricted from speeding up. By decreasing the number of vessels that can speed up, we force ships to arrive later. With these solutions we are then able to quantify and compare how much emissions can be saved in comparison to more ships being late. A ship is too late when it is finished after its expected finishing time ($\Delta \text{EFT}_i > 0$), so the number of total hours too late will be computed over all vessels as $\sum_{i \in V} \Delta \text{EFT}_i$. We will average the CO_2 savings and hours late of all vessels over all the given instances. The relation over the V solutions between how much average CO_2 can be saved and how much total average hours all vessels are too late will be analyzed. With this relation we are then able to compute how much saving CO_2 emissions cost, as the costs of a late hour is given by c_i^2.

Table 2: Parameters used for computing the vessel emissions

Vessel class	$\mathbf{c_0}$	$\mathbf{c_1}$	$\boldsymbol{\mu}$	$\mathbf{v_0}$ (knots)
Feeder	598.65	0.0198	3.5	14.67
Medium	649.65	0.004004	4	15.25
Jumbo	600.45	0.000918	4.5	14.84

3.4 Uncertainty of arrival times

Secondly, uncertainty of arrival times will be incorporated into the heuristics and vessel emissions framework. To account for the uncertainty of a ship being late, buffer zones will be made within the space-time diagram, as can be seen in Figure 6. When a vessel is added in the construction heuristic, the new vessel cannot overlap with an existing buffer zone of another vessel. These zones could help to solve two main problems encountered when dealing with uncertainty, as will be explained below.

Nevertheless, first the impact of the buffer zones on the quality of the solution and on the vessel emissions will be analysed. Adding buffers to the berth schedule will most likely increase the costs of the solution. On the other side, it is expected that the buffer zones will cause the next ship to be unable to move to the front by speeding up ($\Delta ETA > 0$) as the ships may not overlap with buffer zones. This causes less ships to speed up, meaning less fuel is consumed. To study these two effects, the SWO will be used on all instances and for different buffer sizes. These buffer sizes are chosen to incorporate the modeled uncertainty of arrival times, to be discussed next.

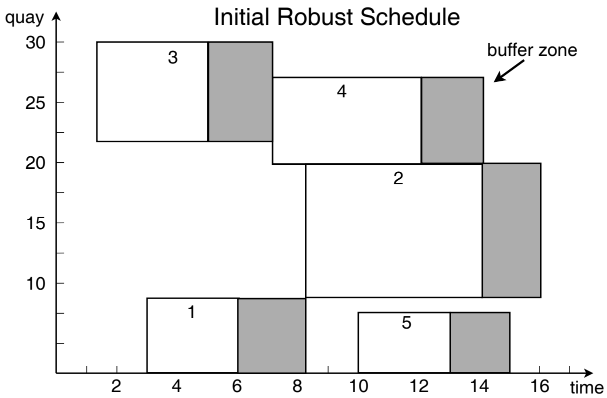

Figure 6: The initial robust schedule with buffer zones

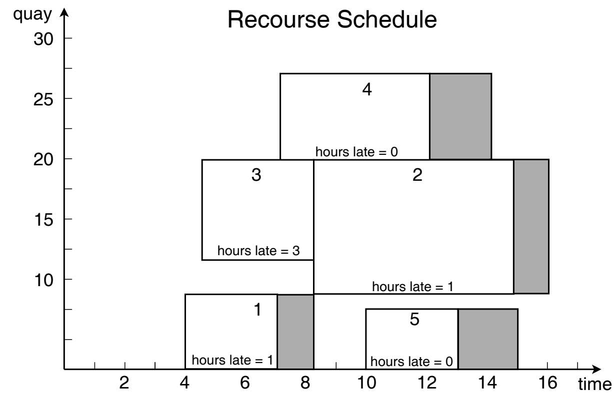

Figure 7: An example of the buffer zone recourse strategy

When the buffer zones are incorporated into the model, we obtain a schedule which we call the initial robust schedule. We can use this schedule to handle the uncertainty of arrival times, as a vessel can always be shifted up in time, should it be late. How much as vessel is late will be modeled by an exponential distribution as data studies on container ships usually encounter exponential arrival times (Imai et al., 2008; Kuo et al., 2006). This means that if we let the hours late for vessel $i \in V$ be denoted by w_i, then $w_i \sim EXP(\lambda)$. Moreover, the exponential distribution makes sure that only positive values are sampled. We will make robust schedules for each instance with the SWO with different buffer zone sizes. Then, we will sample w_i for all vessels $i \in V$, after which we have to make a new schedule.

For this new schedule, one could simply run the SWO again with a buffer size of zero. This procedure, however, takes a long time and can thus not be run for every arrival time update of an incoming vessel. Moreover, before running the SWO again, the terminal manager does not have a good indication whether the new solution brings higher or lower costs and vessel emissions. Therefore, a buffer zone recourse (BZR) strategy will be used, which uses the initial robust schedule and the priority list order of the optimal solution given by the SWO. Every time an update comes in of an arrival time, the BZR gets the order of the performed SWO and tries to shift all vessels on the same quay position within its buffer zone. When the vessel overshoots its buffer zone or when a crane assignment is not feasible, the vessel will simply be inserted by using the construction heuristic. Using the priority list found by the prior work of the SWO and using the same berthing positions is believed to considerably speed up the procedure of making a recourse schedule while still keeping a low objective value Z.

For an example, see Figure 7. First an initial robust schedule is created with buffer zones of size 2. Then, new arrival times of all ships come in and thus some ships are shifted forward based on the priority list found by the SWO performed for the initial robust schedule. For example, vessel 1 and 2 are moved one hour to the front into its buffer zone, after having checked the crane assignment feasibility. Vessel 4 and 5 will simply be added on the same position after checking the crane assignment feasibility, as they do not arrive late. As vessel 3 overshoots its buffer of two hours with three hours too late, it cannot be placed on the same quay position. Therefore it gets inserted by the construction heuristic, giving it its place in front of vessel 2.

3.5 Multiple quays

The last extension is the inclusion of multiple quays. To include the possibility of residing to another quay, the construction heuristic will be adjusted such that all quays will be considered in the insertion of a vessel as can be seen in Figure 8. In the end when all quays are checked, the berthing quay, place and time will be chosen for which the vessel attains the least cost. With this adjustment a sensitivity analysis will be performed. When increasing the number of quays, it is expected that both the solution costs decrease as well as the emissions by speeding up. Moreover, ten instances of 60 vessels will be introduced using combined data from instances of 30 vessels to more accurately analyze the cost decrease when adding more quays to the problem. It is expected that these instances need at least two quays and that the solution costs decrease more than the smaller instances when adding quays. Here we assume that all vessels can go to any quay without competitive restrictions.

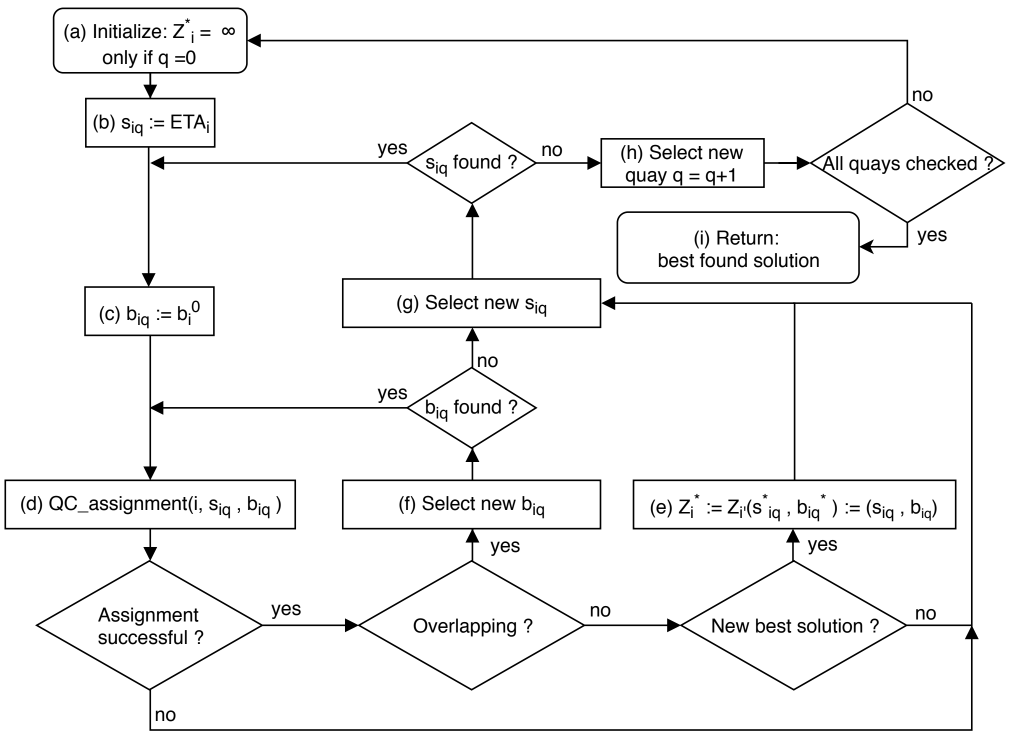

Figure 8: The insertion heuristic diagram with multiple quays

4 Data

The test instances that will be used are not from Park and Kim (2005) nor from Meisel and Bierwirth (2009). Instead, three sets of new instances are newly generated containing 10, 20 and 30 vessels, with the sets having 20 instances each. Similarly as in Meisel & Bierwirth (2009), three vessel classes are used (Feeder, Medium and Jumbo), where the technical specifications are listed in Table 3. Within each instance roughly 60% of the vessels belongs to the Feeder class, 30% belongs to the Medium class and 10% belongs to the Jumbo class. As the table shows, each vessel has a specific length and crane demand given by a uniform distribution U and each vessel in a class has the same minimum and maximum allowed number of unloading cranes. Moreover, each vessel i in a class has the same cost for speeding up c_i^1, the same delay cost c_i^2 and the same penalty cost when arriving late c_i^3. The ETA_i's for each vessel in each instance is uniformly distributed over the time horizon H of one week (168 hours). For each instance, the earliest starting time equals $\lceil EST = 0.9 \cdot ETA \rceil$, which means a vessel can be sped up at most 10%. The latest finishing time LFT_i is computed by $\text{EFT}_i = \text{ETA}_i + 1.5\lceil m_i / r_i^{max} \rceil$ and the expected finishing time EFT_i by $\text{EFT}_i = \text{ETA}_i + \lceil m_i / r_i^{max} \rceil$.

A vessel i's desired berthing position is uniformly distributed as $U[0, L-l_i]$. Moreover, the container terminal has the following data: L = 100 (1000m), Q = 10 cranes and $c^4 = 0.1$ thousand USD per crane-hour. The interference exponent equals to $\alpha = 0.9$ and the berth deviation factor equals $\beta = 0.01$.

What should be noted, is that some differences could exist in the data generation between this thesis and the work by Meisel and Bierwirth. First, in the data for this thesis, vessels i that were generated with a latest finishing time LFT_i greater than time horizon H were removed from the vessel list. This means the division of vessel classes may be favored to Feeders as Mediums and Jumbo's more often have a greater crane demand and thus have a greater chance to fall off the time horizon H with regards to LFT_i. Furthermore, the instances of Meisel and Bierwirth were generated and only used if the construction heuristic was able to make a feasible solution. This check could result in a bias toward easier solutions that are quicker to solve by the SWO and TS. The instances in this thesis were not checked for feasibility before hand. Nevertheless, it could be that the local refinements and the priority changing procedures of the SWO and TS change an unfeasible solution into a feasible one.

Table 3: Technical specifications and costs for the different vessel classes

Class	l_i	m_i	r_i^{min}	r_i^{max}	c_i^1	c_i^2	c_i^3
Feeder	U[8, 21]	U[5,15]	1	2	1	1	3
Medium	U[21, 30]	U[15,50]	2	4	2	2	6
Jumbo	U[30, 40]	U[50,65]	4	6	3	3	9

5 Results

In this section, the four heuristics are compared, the impact of crane productivity parameters on the solutions are analyzed and the three main extensions are performed. The SWO ends after 200 iterations without finding a better solution and the TS ends after 50 iterations without finding a better solution. The algorithms were run in Java JDK 12 on an i7-7700HQ at 3.20 GHz.

5.1 Results validation

As the construction heuristic serves as a base for all heuristics, the algorithm was thoroughly checked for correctness. The data from Meisel and Bierwirth (2009) is not used, so it cannot be said for certain that the construction heuristic from Meisel and Bierwirth and from this thesis are exactly the same. Nevertheless, an example positioning of a vessel is given in Figure 4 of Meisel and Bierwirth (2009), which shows and explains the steps for inserting one vessel next to two other vessels. The exact same example was set up to test the replicated construction heuristic, where it was concluded that the algorithms behave in exactly the same way for this small example. More specifically, the algorithm gives the same costs, same vessel positioning and same crane assignment for each vessel.

The behavior for the construction heuristic and the local refinements were validated by visualized schedules as can be seen in Figure 9 for which instance `n30i11` is optimized by SWO. As in Figure 1, the vertical and horizontal axis in Figure 9 stands for the quay and the time respectively. Moreover, the vessels are portrayed as boxes with their vessel number, vessel cost, and crane demand below the vessel number and cost. Lastly, the crane assignments are visualized within the vessels and the total crane assignment can be seen, where a gray square stands for one hour of crane service.

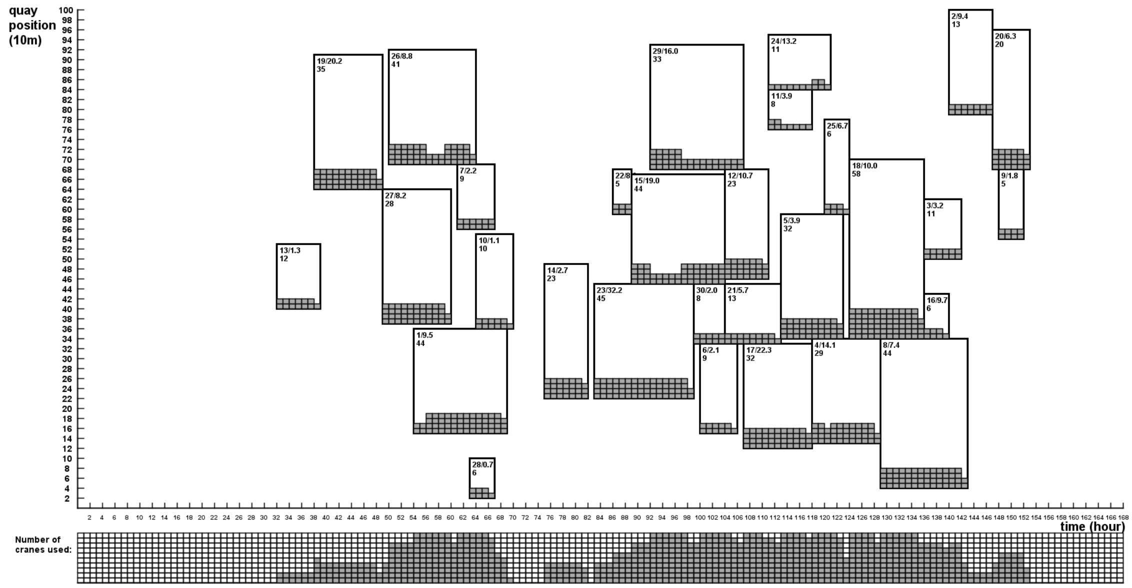

Figure 9: A visualized schedule automatically generated by the Java code for instance `n30i11`

5.2 Comparison of heuristics

The construction heuristic, local refinements, SWO and TS were performed for all instances, as can be seen in Table 4. First, the construction heuristic is listed as FCFS, because it inserts the vessels sorted on ETA's, which acts as a First-Come-First-Serve rule. Then, the FCFS results are improved with the local refinements, referred to as FCFS_{LR}. These refinements include the resource leveling and the spatial and temporal cluster shifting. To improve these solutions further, the SWO and TS are used. The Z value denotes the total sum of all vessel costs. Lastly, the time is only reported for the SWO and TS as the construction and local refinements are executed in a matter of seconds.

First, two instances were not feasible when using the construction heuristic. These were instances `n30i8` and `n30i13`, which is reasonable as most 30 vessels construction solutions looked fairly dense in terms of schedule. Such a dense solution is illustrated in Figure 9. What can be seen from the Z values when comparing the FCFS results with the local refinements is that in only some cases are the local refinements able to improve the solution of the construction heuristic for ten vessels. This improves when increasing the number of vessels to 20 where more solutions can be improved, while all solutions can be improved significantly for instances of 30 vessels. This could be likely due to the fact that there is not much room for improvement when using a small amount of vessels, as most of the ten vessels can have their preferred position and time. Moreover, we see that the local refinements were able to generate a feasible solution for instance `n30i8`.

Furthermore, we see that the SWO improves the solutions over the local refinements for almost all instances with 20 vessels, and improves all solutions with 30 vessels significantly. Moreover, `n30i13` is solved, such that all generated instances are feasible with SWO. Comparing the SWO to the TS, it can be concluded that the TS sometimes performs a little better and for some instances a little worse, as Meisel and Bierwirth also found. For most instances, the running time of the TS was higher than that of the SWO, which also confirms the findings of Meisel and Bierwirth. Lastly, we also validate that the runtime demand significantly rises when using larger instances. Still, the running time for all instances in this thesis were reasonable, all solving within six minutes.

An adjusted tabu search can be seen in the Appendix. Here, the idea was to combine the adjacent swaps of the SWO with the search procedures of the tabu search. In this way, running times should be reduced by a smaller search neighborhood while retaining the strengths of the TS.

Table 4: Performance results of the four heuristics on the given instances

		FCFS	$FCFS_{LR}$	SWO		TS	
n	#	Z	Z	Z	Time	Z	Time
10	0	40.8	38.7	38.7	2	38.7	1
	1	11	11	11	0	11	0
	2	38.9	36.6	27.1	2	27.1	15
	3	30.8	29.8	29.8	2	29.8	0
	4	28	24.8	21.5	4	21.5	10
	5	52.5	45.4	45	3	45.4	23
	6	25.9	25.9	25.9	0	25.9	0
	7	44.9	44.9	44.9	6	44.9	21
	8	41.5	40.5	30.9	1	30.9	10
	9	20.3	20.3	20.3	0	20.3	0
	10	12.4	12.4	12.4	0	12.4	0
	11	47.4	40.7	37.4	2	37.4	9
	12	33.3	30.3	30.3	9	30.3	15
	13	41.2	38.9	38.9	3	38.9	14
	14	16.6	14.4	14.4	0	14.4	4
	15	27.6	27.6	27.6	1	27.6	0
	16	22.8	22.8	22.8	0	22.8	0
	17	36.4	31.4	30.5	2	30.5	15
	18	32.6	32.6	32.6	2	32.6	0
	19	30.8	30.8	30.8	1	30.8	1
20	0	49.2	47	47	7	47.7	75
	1	54.1	54.1	53.1	46	53.1	77
	2	54.2	48.9	48.8	51	48.8	77
	3	38.5	33.3	33.1	6	33.1	48
	4	93.9	76.2	68.2	89	68.2	104
	5	91.2	76.9	68.9	26	68.8	85
	6	54.6	54.6	53.5	56	54.6	4
	7	177.8	148	142.1	106	150.5	123
	8	132.6	107.4	101	62	100.1	141
	9	77.9	67.3	64.2	36	64.2	114
	10	44	40.8	35.8	9	35.8	54
	11	54.9	52.6	51.3	23	50.3	67
	12	63.4	59.3	58.3	39	58.3	88
	13	100.6	86.3	73.6	16	73.6	76
	14	134.8	117.7	101.2	37	108.6	92
	15	126.4	120.2	109.6	27	117.9	100
	16	61.6	56.2	56.2	52	57.3	87
	17	47.1	45.8	43.6	8	44.6	65
	18	64.6	59.4	57.6	22	57.6	78
	19	97	81.4	78.2	35	79	87
30	0	64.7	63.1	60.9	72	60.8	188
	1	249.1	241.8	148.5	87	135.9	320
	2	77.6	67.5	65.3	61	65.5	216
	3	204.6	166.8	142.7	157	143.1	326
	4	181.3	140.8	110.8	95	102.1	248
	5	216.8	182.2	158	167	157.6	360
	6	211.6	166.1	159.2	102	153.5	305
	7	247	158	135.8	211	135.1	279
	8	-	231.1	229.1	160	219.4	313
	9	52.2	52.1	48.8	42	48.8	170
	10	249.7	183.5	177.2	238	174.7	350
	11	378.4	267.7	223.7	237	220.4	324
	12	186.5	153.2	116.3	93	117.9	277
	13	-	-	270.9	253	275.8	414
	14	95.6	91.5	87.7	131	89.4	268
	15	147.3	121.8	104.5	135	106.8	207
	16	189	159.2	134.4	153	134.4	233
	17	180.8	165.2	144.5	254	137.1	235
	18	226.9	186.7	167.5	108	158.9	282
	19	78.1	71.3	65.4	68	65.4	200
Average					60.3		121.6

5.3 Impact of the productivity parameters

Lastly, we look at what effect the crane productivity has on the solution quality. In a similar fashion, we also adjust α and β and run the SWO on all test instances. For the adjustments, α is varied from 1.0 to 0.8 and β from 0 to 0.02, each time keeping the other parameter constant at the default levels $\alpha = 0.9$ and $\beta = 0.01$. For these values, the average cost of all instances are computed. The axis for the α is decreasing as a lower α results in a higher productivity loss when assigning more crane hours to a vessel and thus higher average cost. The parameter β is increasing as a larger β causes larger productivity losses by berth deviation and thus higher costs.

From Figure 10 and Figure 11 we see that the average instance costs for all adjustments are almost twice as low as the average costs in Meisel and Bierwirth (2009). This is mainly due to the fact that we use instances of ten vessels instead of 40, which results in lower costs. For the trend of the average cost when decreasing α, we see a somewhat similar but different sloping pattern. Here, the slope of the average cost steepens and then flattens after $\beta = 0.85$, while in the paper the slope of the curve keeps increasing. A possible explanation for the fact that SWO did not perform well only for $\beta = 0.85$ could be due to the implementation of the heuristics. Moreover, the average cost does not increase as much in the end as in the paper. Here, the average cost increased around 1.8 times, while in the paper it increases over 2.5 times. This could be due to the fact that in this paper smaller problem instances are used, which means less crane hours are assigned, leading to less productivity loss.

For the trend of the average cost when increasing β, we see a very similar pattern with decreasing slope. Moreover, the cost increases in the end around the same order of 1.4 times. Therefore, we confirm the findings of increasing β, while we see a slightly different result for the impact of α.

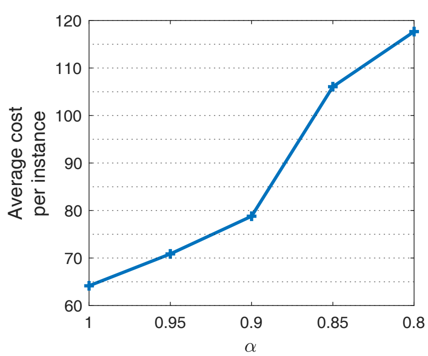

Figure 10: The impact of altering α on the average instance cost

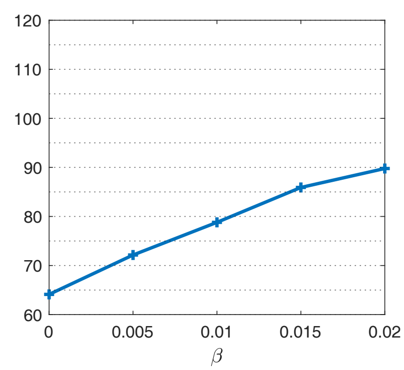

Figure 11: The impact of altering β on the average instance cost

5.4 Costs of saving vessel CO_2 emissions

For the first extension, instances with an increasing number of vessels restricted from speeding up were solved by the SWO to save CO_2 emissions. To make the comparison more direct and as the largest impact of these restrictions are found in large instances, only instances of 30 vessels are used. Moreover, the SWO was performed without any local refinements to speed up the running time. This was done as all 20 instances had to be run 30 times, amounting to an approximate running time of sixteen hours. For all 30 restrictions, the total number of kilograms CO_2 saved by restricting more vessels from speeding up are averaged over all instances.

This resulted in Figure 12 where restricting more vessels clearly leads to increasing average CO_2 savings, as expected. Nevertheless, due to the nature of the scheduling, for some increasing number of restrictions, the savings went down. This can be explained by the fact that some schedules with a given number of restrictions save less CO_2 as their objective value in this way was lower. For the whole trend, we clearly see a decreasing marginal returns of savings when increasing the restrictions. Restricting all vessels thus does not give as much savings benefit as restricting a few. This is likely due to the fact that it gets harder to create a good schedule when increasing the number of restrictions, thus increasing the number of vessel speed-ups and decreasing the slope of emissions saving.

When we also compute how many hours all vessels are too late by $\sum_{i \in V} \Delta \text{EFT}_i$ due the restrictions and average this over all instances, we are able to quantify how expensive it would be to save CO_2 emissions. The data points of how many hours vessels were too late which was the result from saving CO_2 emissions is shown in Figure 13. As the nature of the scheduling procedures resulted in a non-monotonic increasing curve in Figure 12, we regard points that are 1.5 times the standard deviation of the residuals out as outliers in Figure 13. Without these outliers, a non-linear least squares line is fitted of the form $ax^2 + b$ as it clearly shows a non-linear trend, due to the reasons discussed above. The results from this fit can be found in Table 5 below.

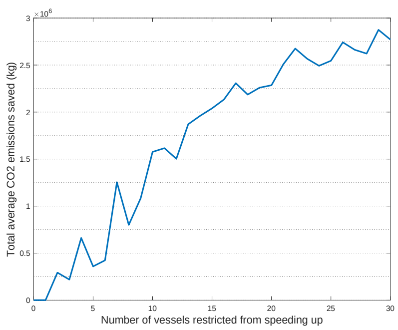

Figure 12: Total average of CO_2 emissions saved by increasing the number of restricted vessels

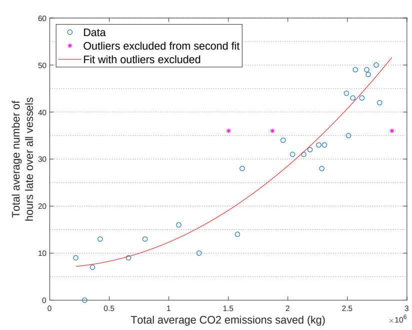

Figure 13: Relation between average total saved CO_2 and the total number of late hours

With these results, the cost of saving a kilogram of CO_2 can be calculated. Using the least squares results, when we would want to save a 1000 tonnes of CO_2 emissions using the cost of being an hour late of Medium container ships, this results in an average cost increase of $5.412 \cdot 10^{-12} \cdot 1000,000^2 \cdot 2000 =$ €$10,824$. Nevertheless, this cost already amounts to €97,416 when we would want to save 3000 tonnes of CO_2 emissions, due to the quadratic nature of the relationship. Lastly, there is a limit of how much CO_2 can be saved on average due the number of total possible restrictions equaling the number of vessels. This limits the least squares line and thus the maximum number of kilograms that can be saved is just below 3000 tonnes CO_2 for these instances.

Table 5: Results for the least squares estimation of the model $y = ax^2 + b$

Variable	**Model coefficient**
a	5.412e-12***
b	6.92*
N	26
R^2	0.9186

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

5.5 Impact of incorporating buffer zones

For the second extension, we incorporate uncertainty of arrival times by first introducing buffer zones into the schedule. These solutions will be used as a basis to incorporate late vessels and thus do not reflect the final solution quality of the BZR. Nevertheless, it is interesting to analyze the effects of incorporating buffers into the schedule, as these schedules will form the basis of our BZR schedules. These schedules thus do not contain any uncertainty of arrival times yet. The SWO is run with buffer zones ranging from zero to four hours to compute the average cost and average CO_2 vessel emission over all instances.

When looking at the average costs in Figure 14 we see that incorporating buffer zones has a significant effect on the solution quality. When adding an extra hour to the buffer, the average costs increase by around 7 in a linear trend. As adding time spaces between the vessels increases the difficulty to make a schedule, this increasing trend was expected. These average costs, however, do not reflect the final costs by the buffer zone recourse strategy yet. So with these results we cannot conclude what buffer size to use.

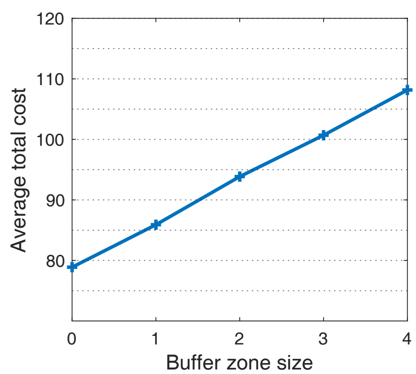

Figure 14: The impact of altering the size of the buffer zones on the average solution costs

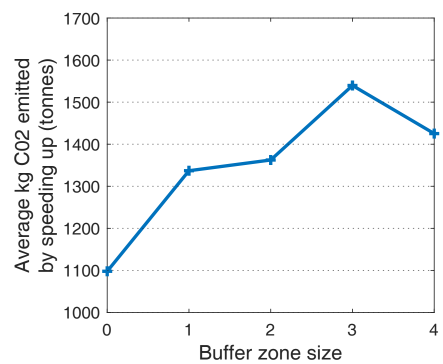

Figure 15: The impact of altering the size of the buffer zones on the total average vessel emissions

In terms of CO_2 emissions in Figure 15, it was expected that the average emissions would go down when using larger buffer zones. Instead, the average emissions significantly rose for buffer zones up to three hours. An explanation for this increase is that due to the buffers, all ships are likely to incur delay cost c_i^2 and penalty cost c_i^3. To avoid these relatively large costs, all ships are sped up contributing to increased emissions. This is also in line with the discussed increasing average costs. The average CO_2 decreased from a buffer zone of 3 hour to 4 hours, which could be the cause of the nature of scheduling. For example, a better solution could be found with these vessel emissions. Still, these average vessel emissions do not tell anything about the final emissions of using buffer zones in the recourse strategy. That is something we will turn to next.

5.6 Buffer zone recourse strategy

After having obtained the initial robust schedule including buffer zones by SWO, we can incorporate the uncertainty of arrival times by using the buffer zone recourse strategy. The number of hours late w_i for vessel $i \in V$ will be exponentially distributed with a mean of two hours. Based on $w_i \sim EXP(\lambda = 2)$ we both sample new ETA's for all vessels and then perform the buffer zone recourse strategy for a 100 iterations. To test the performance of different recourse strategies, the solution costs and CO_2 emissions are measured for buffers ranging from zero to four hours. For each buffer size the number of hours late of all vessels are pseudo-randomly sampled with the same seed.

The hypothesis was that the the buffer zone strategy with buffer zones of two hours obtained the least costs, as two hours is the best approximation for $E[w_i]$ and limits the buffer sizes. A limited buffer size enables better schedules in the SWO, which causes better berthing places for the vessels, which are fixed when shifting vessels in the BZR.

When looking at the costs and emissions in Figure 16 and Figure 17, this hypothesis is not fully confirmed, as the average cost is lowest for buffers of one hour. The average cost does not differ too much from the costs with two hour buffers, so this could be due to the nature of scheduling. Nevertheless, the CO_2 emissions are at its lowest for buffer zones of two hours. For both figures, the costs and emissions show a similar trend: a buffer zone recourse strategy without any buffer zone is quite costly. Then the costs and emissions drop significantly when using a buffer zone of one hour. Therefore, shifting the vessels in the BZR is a significantly better strategy than inserting the vessels in the same order found by the SWO using the construction heuristic. Then, the costs and emissions gradually increase when increasing the size of the buffers. This increasing trend is likely be due to the fact that the SWO can find better berthing places with smaller buffer zones. The berthing places found by SWO are namely fixed when shifting up the vessels in the BZR. This indicates that a small buffer zone in the BZR is recommended for arrival times with this exponential distribution.

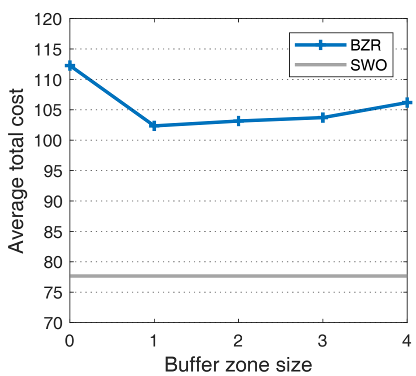

Figure 16: Average cost with using different buffer zone recourse strategies

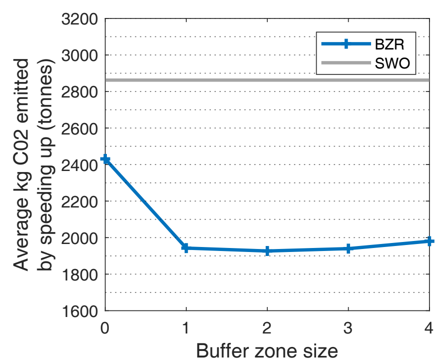

Figure 17: Average total emissions with using different buffer zone recourse strategiess

When we use the SWO to get a new schedule for the new sampled ETA's instead of the BZR, we see that it performs significantly better than the BZR with an average cost below 80. Nevertheless, this is mainly due to the fact that the SWO can be run without buffer zones. When comparing these costs to the costs of the SWO with buffer zones in the previous section, we see that the costs with two hour buffers namely rise from 95 of the SWO in Figure 14 to 105 of the BZR in Figure 16. Moreover, the running time of the SWO is significantly greater than that of the BZR: an average SWO runs for 98 seconds while the BZR can adjust an average schedule within 83 milliseconds. Therefore, the BZR can be useful for ports with a high stochastic environment where the schedule needs to be adjusted constantly. When there is enough time and computation power on hand though, the SWO is a better alternative. On the other hand, the CO_2 emissions are significantly higher with SWO. This is mainly due to the fact that vessels are only added in BZR from their ETA, so they are not allowed to speed up unless they are inserted due to an unfeasible move.

5.7 Multiple quay benefits

Lastly, the problem is extended to incorporate multiple quays. The SWO is run for all instances using a range of one until five quays. In Figure 18 and Figure 19 the solution costs and emissions by speeding up are averaged over all instances with the same number of vessels.

We see that the average solution costs decreases only slightly when adding a second quay for the instances with ten vessels. The costs do decrease significantly for instances with 20 vessels. Nevertheless, a third quay for instances with 20 vessel does not make sense as the average costs already converge at this point. Adding a third quay is only somewhat viable for instances of 30 vessels. When we also include newly created instances of 60 vessels, we are able to analyze the effects further. We see that one quay is not enough as all ten instances were not solvable using one quay. By starting at two quays and adding a third and even a fourth significantly drops the costs. After a fourth quay, again the costs start to converge. Therefore, larger instances benefit more from adding another quay.

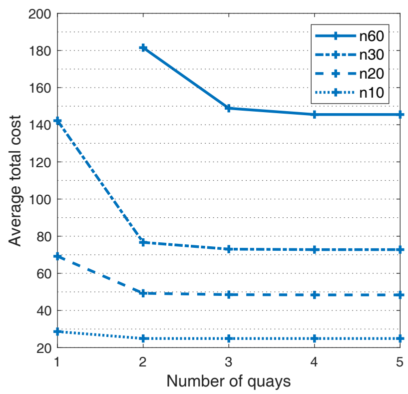

Figure 18: Average costs effects of adding quays to the problem

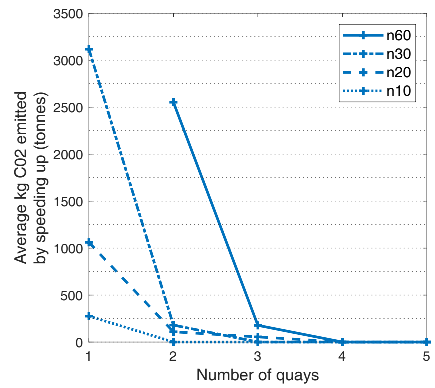

Figure 19: Total average emissions effects of adding quays to the problem

When looking at the CO_2 emissions with different level of quays, we notice that adding one quay to the minimum number of needed quays drastically reduces the emissions by speeding up to almost zero for all instance sizes. This is likely be due to the fact that all vessels have significantly more space in the schedule reducing the need for speeding up. Therefore, to save emissions it is viable to add another quay to the minimum number of quays needed to make a feasible schedule. When making a cost calculation for instances of 60 vessels, we can use the fact that the costs decrease from roughly 180 to 150 which is a saving of $30,000 * 52 = €\,1,560,000$ per year. Moreover, when incorporating the cost of saving vessel emissions by restricting vessels from speeding up in Section 5.4, we can save roughly 2250 tonnes of CO_2 per week which costs $5.412 \cdot 10^{-12} \cdot 2250^2 \cdot 2000 = €\,54,797$ and amounts to 1,170,000 tonnes of CO_2 per year. Therefore, next to extra benefits such an increased satisfaction of vessel operators, the investment in an extra quay can be justified by saving $€\,1,560,000$ in service and operator cost and 1,170,000 tonnes of extra CO_2 per year.

6 Conclusions

Replication and extensions

The need for efficiently solving the berth allocation and crane assignment problem is of considerable value especially for terminal operators, vessel operators and residents working and living around ports. Therefore, finding rich BACAP models and well performing solution techniques are an important aspect in contributing to the research around this topic. In this thesis, the goal has mostly been to enrich the BACAP model by incorporating vessel emissions, uncertainty of arrival times and utilization of multiple quays. This goal resulted in the research question asking what the costs are of saving vessel emissions. Moreover, the consequences of arrival uncertainty and of using multiple quays within the problem in terms of costs and emissions were of interest. Before these questions could be answered though, the results from the work of Meisel and Bierwirth were replicated, which focused upon incorporating handling time due to crane productivity and berthing position.

First, the construction heuristic was performed and validated by visualized schedules. This heuristic provided satisfactory results, although it was not able to solve instances `n30i8` and `n30i13`. By creating a local refinement heuristic, which leveled crane resources over all vessels and moved spatial and temporal clusters over the schedule, instance `n30i8` was able to be solved and most instances improved in solution quality. After this, the instructions from Meisel and Bierwirth were followed to create a squeaky wheel optimization and tabu search. These behaved very similar to the results in their work, with the SWO and TS obtaining comparable costs although SWO obtained better running times. Lastly, the impact of altering the two productivity parameters was checked. These had the comparable effects as found in Meisel and Bierwirth (2009).

Then, the model was first extended by incorporating vessel emissions. The emissions from speeding up a vessel were calculated and could be used to determine how expensive it would be to try to save these emissions by restricting vessels from speeding up. This was done by performing the SWO for a number of instances where for each instance, an extra vessel was restricted from speeding up. It was concluded that the costs of saving extra vessel emissions increased exponentially, thus making it increasingly expensive to save more CO_2 emissions.

The second extension used buffer zones to keep space for late-arriving vessels in the schedule. First, an initial robust schedule was generated with buffer zones for which the late-arriving vessels were shifted up in time. Vessels that were later than expected were inserted via the construction heuristic. After the impact of buffer zones had been inspected, the results of the BZR was analyzed. The BZR performed best for buffer zones of one to two hours, although the SWO performed significantly better. Nevertheless, an average SWO takes longer to solve than the BZR. Furthermore, the BZR was able to solve a significant amount of CO_2 emissions.

Lastly, utilization of multiple quays was incorporated into the model. The SWO was performed with one to five available quays, which yielded expected results for the solution costs and a surprising result for the CO_2 emissions. Using one extra quay on top of the minimum number of needed quays drastically decreased the CO_2 emissions by speeding up. Finally, it was concluded that the larger the problem instances are, the larger the cost benefits from using more quays.

Recommendations

To highlight the conclusions retrieved from the vessel emissions, arrival uncertainty and multiple quay utilization, the most important recommendations are given separately. First, the costs of saving vessel emissions by restricting vessels from speeding up seem to be fairly low when saving small amounts of emissions. Therefore, we recommend saving a third of the average emissions as this only bares roughly a tenth of the total costs, due to the found quadratic relationship of cost versus saving. Terminal operators can namely save a 1000 of the total 3000 tonnes of CO_2 by a solution cost increase of € 10,824 instead of € 97,416 when trying to save all emissions from speeding up.

Next, for ports with a highly stochastic environment where the schedule needs to be updated constantly, the BZR strategy proposed in this thesis is necessary. The BZR is namely able to create well-performing schedules in a fraction of the time the SWO can and it also gives a better approximation of the costs for terminal operators beforehand. Although the results in this thesis points at smaller buffer zones, we therefore recommend the BZR for these ports with buffer zones of two hours whenever the arrival process can be portrayed by an exponential distribution with a mean of two hours. Setting the buffer size equal to the expected number of hours late is done, because the costs between one and two hours does not differ too much and theoretically it is more underpinned.

Lastly, for ports using multiple quays at once, solving the BACAP integrated with multiple quays is essential. Not only does it enable solving the problem over the whole port, but it also enables the analysis and simulation for investment in future expansion. The most important takeaway for this feature is to invest in another quay on top of the minimum number of needed quays, as it drastically reduces the CO_2 emissions.

Limitations and further research

Nevertheless, there are some limitations to the extensions in this work. First, in the vessel emissions framework the distance k_i is not available and is approximated by the ETA and design speed of the vessels. This results in the fact that the computed vessel emissions are rough approximations. To solve this problem, future research should keep in mind the distance covered over the whole travel from port to port. Moreover, the emissions caused by burning extra fuel from slowing down could also be considered. Secondly, the BZR strategy obtains low running times for all instances, but finds considerably worse schedules than the SWO. Therefore, some work still could be done to adjust the strategy such that it keeps its low running time, but such that it approaches the average costs of the SWO. Lastly, the addition of multiple quays can be of particular interest for future research, as it enables solving the BACAP integrally over the whole port area. The quays could even be extended to only allow vessels of a certain type, such as container, dry and liquid bulk. Moreover, in the visualized schedules seen in performing the extension, some quays had been assigned very little vessels. When we would also incorporate the container yard capacity and work division of personnel on each quay, it might be wise to equally distribute the vessels over each quay to ensure that some quays are left unproductive and some overfull of containers. The work remains how to do this while still ensuring minimum solution costs.

Finally, after having performed extensions to enlarge the potential for heuristics to solve the BACAP in a broader setting, it remains a challenge to minimize costs within reasonable time. Not only as constant updates of the schedules are sometimes needed due to uncertainty, but also due to the fact that even better solutions can be attained within the same running time. Moreover, it was also found that enlarging the framework to multiple quays increases the computation time significantly. To improve on running times, other meta-heuristics could be evaluated in future research, which should be compared within the same problem setting to enhance comparability and with this the resulting conclusions.

References

Ak, A., & Erera, A. L. (2006). Simultaneous berth and quay crane scheduling for container ports. *Milton Stewart School of Industrial and Systems Engineering*.

Bailey, D., & Solomon, G. (2004). Pollution prevention at ports: clearing the air. *Environmental impact assessment review*, *24*(7-8), 749–774.

Bierwirth, C., & Meisel, F. (2010). A survey of berth allocation and quay crane scheduling problems in container terminals. *European Journal of Operational Research*, *202*(3), 615–627.

Bierwirth, C., & Meisel, F. (2015). A follow-up survey of berth allocation and quay crane scheduling problems in container terminals. *European Journal of Operational Research*, *244*(3), 675–689.

Chang, D., Yan, W., Chen, C. H., & Jiang, Z. (2008). A berth allocation strategy using heuristics algorithm and simulation optimisation. *International Journal of Computer Applications in Technology*, *32*(4), 272. doi: 10.1504/ijcat.2008.021382

Cullinane, K., & Khanna, M. (2000). Economies of scale in large containerships: optimal size and geographical implications. *Journal of transport geography*, *8*(3), 181–195.

Du, Y., Chen, Q., Quan, X., Long, L., & Fung, R. Y. (2011). Berth allocation considering fuel consumption and vessel emissions. *Transportation Research Part E: Logistics and Transportation Review*, *47*(6), 1021–1037.

Entec UK Ltd. (2002). *Quantification of emissions from ships.* European Comission. Retrieved 12 May, 2019, from http://ec.europa.eu/environment/air/pdf/chapter2_ship_emissions.pdf

Golias, M., Portal, I., Konur, D., Kaisar, E., & Kolomvos, G. (2014). Robust berth scheduling at marine container terminals via hierarchical optimization. *Computers & Operations Research*, *41*, 412–422.

Guan, Y., & Cheung, R. K. (2004, Jan 01). The berth allocation problem: models and solution methods. *OR Spectrum*, *26*(1), 75–92. Retrieved from https://doi.org/10.1007/s00291-003-0140-8 doi: 10.1007/s00291-003-0140-8

Hu, Q.-M., Hu, Z.-H., & Du, Y. (2014). Berth and quay-crane allocation problem considering fuel consumption and emissions from vessels. *Computers & Industrial Engineering*, *70*, 1–10.

Imai, A., Nishimura, E., & Papadimitriou, S. (2001). The dynamic berth allocation problem for a container port. *Transportation Research Part B: Methodological*, *35*(4), 401–417.

Imai, A., Nishimura, E., & Papadimitriou, S. (2008). Berthing ships at a multi-user container terminal with a limited quay capacity. *Transportation Research Part E: Logistics and Transportation Review*, *44*(1), 136–151.

Jung, D. H., Park, Y. M., Lee, B. K., Kim, K. H., & Ryu, K. R. (2006). A quay crane scheduling method considering interference of yard cranes in container terminals. *Mexican International Conference on Artificial Intelligence*, 461-471.

Kim, K. H., & Park, Y.-M. (2004). A crane scheduling method for port container terminals. *European Journal of Operational Research*, *156*(3), 752 - 768. Retrieved from http://www.sciencedirect.com/science/article/pii/S0377221703001334

Kuo, T.-C., Huang, W.-C., Wu, S.-C., & Cheng, P.-L. (2006). A case study of inter-arrival time distributions of container ships. *Journal of Marine science and technology*, *14*(3), 155–164.

Lim, A., Rodrigues, B., Xiao, F., & Zhu, Y. (2004). Crane scheduling with spatial constraints. *Naval Research Logistics (NRL)*, *51*(3), 386-406. Retrieved from https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.10123 doi: 10.1002/nav.10123

Lu, Z.-q., Xi, L.-f., et al. (2010). A proactive approach for simultaneous berth and quay crane scheduling problem with stochastic arrival and handling time. *European Journal of Operational Research*, *207*(3), 1327–1340.

MAN Diesel & Turbo. (2008). *Basic principles of ship propulsion.* Retrieved May 11, 2019, from http://www.manbw.com/files/news/filesof3859/P254-04-04.pdf

Meier, L., & Schumann, R. (2007). Coordination of interdependent planning systems, a case study. *GI Jahrestagung*, *1*, 389-396.

Meisel, F. (2009). *Seaside operations planning in container terminals.* Berlin: Phsyica-Verlag.

Meisel, F., & Bierwirth, C. (2009). Heuristics for the integration of crane productivity in the berth allocation problem. *Transportation Research Part E: Logistics and Transportation Review*, *45*(1), 196–209.

Park, Y.-M., & Kim, K. H. (2005). A scheduling method for berth and quay cranes. In *Container terminals and automated transport systems* (pp. 159–181). Springer.

Port of Antewerp. (2019). *Types of goods.* Retrieved May 7, 2019, from https://www.portofantwerp.com/en/types-goods

Port of Rotterdam. (2019, April 18). *Throughput.* Retrieved May 7, 2019, from https://www.portofrotterdam.com/en/our-port/facts-and-figures/facts-figures-about-the-port/throughput

Schrady, D., & Wadsworth, D. (1991). Naval combat logistics support system. *Journal of the Operational Research Society*, *42*(11), 941–948.

United Nations Conference on Trade and Development. (2015). *Review of maritime transport.* United Nations.

Venturini, G., Iris, Ç., Kontovas, C. A., & Larsen, A. (2017). The multi-port berth allocation problem with speed optimization and emission considerations. *Transportation Research Part D: Transport and Environment*, *54*, 142–159.

Xu, Y., Chen, Q., & Quan, X. (2012). Robust berth scheduling with uncertain vessel delay and handling time. *Annals of Operations Research*, *192*(1), 123–140.

Zhen, L., Lee, L. H., & Chew, E. P. (2011). A decision model for berth allocation under uncertainty. *European Journal of Operational Research*, *212*(1), 54–68.

A Appendix

A.1 Adjusted tabu search with adjacent swaps

Below, results for all instances of 20 and 30 vessels are reported for the SWO, TS and an adjusted tabu search (TS_{AS}). This adjusted tabu search explores the whole neighborhood not by all possible pair-wise exchanges, but through all possible adjacent swaps. The thought behind this is that this resembles the SWO more with its adjacent swaps and reduces the search neighborhood compared to the normal tabu search. Therefore, the TS_{AS} is expected to perform close to the SWO, while reducing the running time.

The SWO is run with 200 iterations and the two tabu searches with 50 iterations without finding a better solution. When looking at the results in Table 6, we see that the TS_{AS} performs slightly worse than the TS for most instances, although improving the solution over the TS in some instances such as `n20i7` and `n20i15`. For most instances, the extensive neighborhood search of the regular TS clearly dominates TS_{AS}. Nevertheless, all running times of the TS_{AS} are significantly lower, resulting in six times smaller average running times over the regular TS and three times over the SWO. Therefore, the TS_{AS} is a competitor for the TS and even for the SWO: when the TS performs better than the SWO for instances of size 30, the TS_{AS} too obtains better solutions than the SWO (instances `n30i1`, `n30i4` and `n30i18`).

Table 6: Performance results of the SWO and TS compared to the adjusted TS_{AS}

		SWO		TS		TS_{AS}	
n	#	Z	Time	Z	Time	Z	Time
20	0	47	7	47.7	75	47.7	16
	1	53.1	46	53.1	77	54.1	1
	2	48.8	51	48.8	77	48.9	19
	3	33.1	6	33.1	48	33.1	11
	4	68.2	89	68.2	104	72.8	25
	5	68.9	26	68.8	85	69.8	23
	6	53.5	56	54.6	4	54.6	1
	7	142.1	106	150.5	123	148	36
	8	101	62	100.1	141	102.4	42
	9	64.2	36	64.2	114	64.2	29
	10	35.8	9	35.8	54	35.8	12
	11	51.3	23	50.3	67	51.4	14
	12	58.3	39	58.3	88	59.3	19
	13	73.6	16	73.6	76	76.8	18
	14	101.2	37	108.6	92	113.1	22
	15	109.6	27	117.9	100	113.7	24
	16	56.2	52	57.3	87	56.2	24
	17	43.6	8	44.6	65	44.6	15
	18	57.6	22	57.6	78	57.6	16
	19	78.2	35	79	87	87.4	27
30	0	60.9	72	60.8	188	60.8	29
	1	148.5	87	135.9	320	139.6	56
	2	65.3	61	65.5	216	65.5	37
	3	142.7	157	143.1	326	141.8	67
	4	110.8	95	102.1	248	103.1	41
	5	158	167	157.6	360	164.9	57
	6	159.2	102	153.5	305	163.6	60
	7	135.8	211	135.1	279	168.8	47
	8	229.1	160	219.4	313	231.7	57
	9	48.8	42	48.8	170	51.2	27
	10	177.2	238	174.7	350	179.6	60
	11	223.7	237	220.4	324	263.6	58
	12	116.3	93	117.9	277	123.1	48
	13	270.9	253	275.8	414	295.4	61
	14	87.7	131	89.4	268	89.4	47
	15	104.5	135	106.8	207	115.8	35
	16	134.4	153	134.4	233	136.6	42
	17	144.5	254	137.1	235	146.7	46
	18	167.5	108	158.9	282	158.9	47
	19	65.4	68	65.4	200	65.4	28
Average			60.3		121.6		23.3

A.2 Program files listing

The Java project to perform the heuristics of Meisel and Bierwirth (2009), the vessel emission analysis, the buffer zone recourse strategy and the multiple quay analysis in this thesis can be found in the given Zip file. Below, the file structure of this Java project is given. First, there are two packages: one package for the multiple quay analysis and the default package for all the other heuristics and extensions. The reason for this is that the framework in the multiple quay analysis is slower than the original version. Canvas.java prints the space-time diagram solutions in a JFrame as depicted by Figure 9, both Mains perform the analyses and heuristics and Vessel.java contains all the information of a vessel within the instance. All methods contained within the .java files are documented. In the `instances` folder, all the used instances are given including the small example from Meisel and Bierwirth (2009). Lastly, the folder `lib` contains the `commons_lang_3` library to read the CSV instances.

```
Meisel&Bierwirth2009
├─ src
│  ├─ (default package)
│  │  ├─ Canvas.java
│  │  ├─ MeiselandBierwirth2009Main.java
│  │  └─ Vessel.java
│  └─ multipleQuays
│     ├─ Canvas.java
│     ├─ multipleQuaysMain.java
│     └─ Vessel.java
├─ instances
└─ lib
```

Figure 20: The file structure of the used Java project