



Discrete Optimization

The lockmaster's problem[☆]Ward Passchyn^a, Sofie Coene^a, Dirk Briskorn^c, Johann L. Hurink^d, Frits C. R. Spieksma^{a,*}, Greet Vanden Berghe^b^a KU Leuven, Faculty of Economics and Business, ORSTAT, Naamsestraat 69, 3000 Leuven, Belgium^b KU Leuven, Department of Computer Science, CODES & iMinds-ITEC, Gebroeders De Smetstraat 1, 9000 Gent, Belgium^c Chair of Production and Logistics, Bergische Universität Wuppertal, Rainer-Gruenter-Straße 21, 42119 Wuppertal, Germany^d University of Twente, Department of Applied Mathematics, P.O. Box 217, 7500 AE Enschede, The Netherlands

ARTICLE INFO

Article history:

Received 29 September 2014

Accepted 1 December 2015

Available online 11 December 2015

Keywords:

Transportation

Lock scheduling

Batch scheduling

Dynamic programming

Complexity

ABSTRACT

Inland waterways form a natural network infrastructure with capacity for more traffic. Transportation by ship is widely promoted as it is a reliable, efficient and environmental friendly way of transport. Nevertheless, locks managing the water level on waterways and within harbors sometimes constitute bottlenecks for transportation over water. The lockmaster's problem concerns the optimal strategy for operating such a lock. In the lockmaster's problem we are given a lock, a set of upstream-bound ships and another set of ships traveling in the opposite direction. We are given the arrival times of the ships and a constant lockage time; the goal is to minimize total waiting time of the ships. In this paper, a dynamic programming algorithm is proposed that solves the lockmaster's problem in polynomial time. This algorithm can also be used to solve a single batching machine scheduling problem more efficiently than the current algorithms from the literature do. We extend the algorithm such that it can be applied in realistic settings, taking into account capacity, ship-dependent handling times, weights and water usage. In addition, we compare the performance of this new exact algorithm with the performance of some (straightforward) heuristics in a computational study.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Transportation of goods by ship, over sea as well as over waterways, is a promising alternative for transport over land. Reasons are its reliability, its efficiency (a barge of 3000 tons can transport as much as 100 train wagons or 150 trucks), its relatively low operating cost and its environmental friendliness. Hence, the relative importance of this mode of transport is rising. Maritime transport for intercontinental trade receives increasing attention due to the high cargo volumes that can be shipped at a much lower energy cost compared to air cargo traffic. Inland waterways transport is therefore seen as a mode of transport that can make a significant contribution to sustainable mobility. For instance, in a recent report, the [European Commission \(2015\)](#) promotes a better use of inland waterways in order to relieve heavily congested transport corridors. Not only is the energy consumption of transport over water approximately 17 percent of that of road transport and 50 percent of rail transport, it also has

a high degree of safety and its noise and gas emissions are modest. This natural network is the only existing infrastructure with excess capacity and where congestion is limited ([Inland Navigation Europe, 2014](#)). In the US, total waterborne commerce has risen from about 1500 million tons of goods in 1970 up to 2600 million tons in 2006; due to the economic crisis it has lowered since then to a level of about 2200 million tons in 2009 ([U.S. Army Corps of Engineers, 2009](#)). In a report prepared for the State of New York ([Goodban Belt LLC, 2010](#)), the potential of the New York Canal System for container-on-barge logistics is extensively described and an increased flow is expected after the Panama Canal expansion in 2015. The cargo volume transiting the Panama Canal is expected to grow on average 3 percent per year between 2005 and 2025, mainly due to an increase in container transport ([Panama Canal Authority, 2006](#)). In China, 88 million tons of freight passed the Three Gorges Dam in 2010; this is nearly 5 times the maximal annual volume reported before 2003 ([ChinaDaily, 2011](#)).

Many of these waterways (the Panama Canal, the Three Gorges Dam, inland waterways in Europe such as the Kiel canal ([Luy, 2010](#)), and many others) are accessible through sea locks and are often interrupted by river locks. Locks are needed to control the water level so that large and heavy ships can continue to access the corresponding waterways. At several waterway networks, congestion is expected to increase, yielding extra pressure on the locks. Examples are the New York State Canal System ([Goodban Belt LLC, 2010](#)) and

[☆] An extended version of this paper is available as a research report, [Passchyn et al. \(2015\)](#)

* Corresponding author. Tel.: +3216326976; fax: +3216326732.

E-mail addresses: ward.passchyn@kuleuven.be (W. Passchyn), sofie.coene@kuleuven.be (S. Coene), briskorn@uni-wuppertal.de (D. Briskorn), j.l.hurink@utwente.nl (J.L. Hurink), frits.spieksma@kuleuven.be (F.C.R. Spieksma), greet.vandenbergh@cs.kuleuven.be (G. Vanden Berghe).

the North Sea Canal region (van Haastert, 2003). Some locks are being expanded, such as locks on the Twente Canal in the Netherlands (Rijkswaterstaat, 2010). Also, new locks are being built, for example, the Deurganckdock lock at the harbor of Antwerp (Antwerp Port Authority, 2011), which will become the world's largest lock by volume.

These locks are bottlenecks for transportation over water, and hence, operating locks wisely contributes to the attractiveness of transportation over water and can help to avoid expensive infrastructural interventions. However, the algorithmic problem of how to operate a lock has not been studied broadly in the scientific literature. The purpose of this paper is to fill this gap. Our point of view here is to see the lock as an entity providing a service to the ships. Then, it makes sense to identify a strategy for the lock that optimizes some criterion related to service (as we do in this paper). Another point of view would be that the lock announces times when ships can enter the lock in order to be transferred, and that the ships simply need to respect these times. Clearly, even in the latter model, there is still a decision to be made concerning these times.

1.1. Problem definition

We now give a formal description of a very basic situation that will act as our core problem: the lockmaster's problem. We first study this simplified problem; the results obtained for this problem will serve as a basis for dealing with more realistic settings later on.

Consider a lock consisting of a single chamber. Ships that are traveling downstream arrive at the lock at given times. Other ships traveling upstream arrive at the lock, also at given times. Let $A = \{1, \dots, n\}$ represent the set of all ships. Let $t(a)$ be the given arrival time of ship $a \in A$, and $p(a)$ the arrival position of ship $a \in A$, with $p(a) = 0$ for a downstream arrival, and $p(a) = 1$ for an upstream arrival. For convenience, we assume that a total order is imposed on A so that ships are ordered by non-decreasing arrival time, i.e. $t(i) \leq t(i+1)$ for $i = 1, \dots, n-1$. Note that multiple ships may have the same arrival time; the total order thus breaks these 'ties' arbitrarily. It is important to emphasize that we do not require ships to enter the lock in the order imposed on A , except in the cases (see Sections 5.1 and 5.3) where we explicitly state this as a requirement.

Let T denote the *lockage duration*: this is the time needed for the water level to rise from the downstream level to the upstream level (or vice versa), plus the time needed to load and unload the lock, which is assumed to be constant in the basic problem. In other words, T measures the time that elapses between opening the lock such that ships can enter, and closing the lock after ships have left. Throughout the paper, we assume $T > 0$ as the problem becomes trivial for $T = 0$. We further assume that all data are integral. Our goal is to find a feasible lock-strategy that minimizes total waiting time of all ships. The waiting time of a ship is the length of the period that elapses between the ship's arrival time and the moment in time when the ship enters the lock. Thus, we need to determine at which moments in time the water level in the lock should start to go up (meaning at which moments in time downstream ships enter the lock and are lifted), and at which moments in time the water level in the lock should start to go down. For such a strategy to be feasible, (i) going-up moments and going-down moments should alternate, and (ii) consecutive moments should be at least T time-units apart.

This particular problem (to which we refer as the lockmaster's problem) is a simplified version of reality. However, we see this problem as a basic problem underlying any practical lock scheduling problem; and we show how to solve this basic problem by dynamic programming (DP) in Section 3. Practical problems obviously feature many properties that are absent in this basic problem. In Section 5 we give an overview of many such features and investigate the complexity of these problem extensions.

The contributions of the present paper can be summarized as follows. We show that (1) there exists an $O(n^2)$ algorithm (see Section 4)

for the lockmaster's problem; (2) this algorithm can be extended to solve variants with capacities, ship-dependent handling times, ship priorities, non-uniform lockage times, and settings with a limited number of lockages (Section 5).

In addition, we investigate the performance of several heuristics by running them on randomly generated instances that possess real-life characteristics (Section 6).

2. Literature

2.1. Lock scheduling

Lock scheduling has not been studied very thoroughly in the academic literature, although it has recently started to attract more attention. We mainly focus on the literature that considers scheduling a single lock. A relatively early work by Petersen and Taylor (1988) considers the Welland Canal in Canada. The authors present a dynamic programming algorithm for scheduling a single lock with unit capacity and extend this to obtain a heuristic result for the entire canal. A more recent paper (Nauss, 2008) deals with optimal sequencing in the presence of setup times and non-uniform processing times for the case where all arrival times are equal to zero. Smith, Sweeney, and Campbell (2009) simulate the impact of decision rules and infrastructure improvements on traffic congestion along the Upper Mississippi River. Ting and Schonfeld (2001) use heuristic methods to study several control alternatives in order to improve lock service quality. Verstichel and Vanden Berghe (2009) develop (meta)heuristics for a lock scheduling problem where a lock may consist of multiple parallel, capacitated chambers of different dimensions and lockage times, making this problem at least as hard as a bin packing problem. The question of filling a lock with ships, i.e. the packing problem, is discussed by Verstichel, De Causmaecker, Spieksma, and Vanden Berghe (2014a). Verstichel, De Causmaecker, Spieksma, and Vanden Berghe (2014b) propose a mixed integer programming model to solve a generalized lock scheduling problem for instances with a limited number of ships. Recent work by Hermans (2014) considers the optimization problem of scheduling a single lock with a capacity that allows a single ship. A polynomial time procedure asserts feasibility with respect to given deadlines in $O(n^4 \log n)$ time; it can further be used to minimize the maximum lateness.

In a related problem, ships need to pass a narrow canal, and only a restricted number of wider areas is available where ships can pass each other. Ships need to wait in these areas and are arranged in convoys that transit in a one-way direction, see e.g. (Griffiths, 1995; Günther, Lübbecke, & Möhring, 2011; Luy, 2010; Panama Canal Authority, 2006). For instances where all ships travel in the same direction, total waiting time is equal to zero, which is not necessarily the case in the lockmaster's problem.

2.2. Machine scheduling

Smith et al. (2011) relate traffic operations at a river lock with a variant of the job shop scheduling problem with sequence dependent setup times and a two-stage queuing process. The lockmaster's problem is more general since ships (i.e. jobs) have release dates and multiple ships can be locked together. The relation between the lock scheduling problem and classical machine scheduling is also observed in the literature. In fact, the lockmaster's problem introduced in this work is closely related to a batch scheduling problem. Batch scheduling involves a machine that can process multiple jobs simultaneously. Suppose that the basic lock scheduling problem only has downstream-bound ships (we will refer to this special case of the lockmaster's problem as the *uni-directional case*). The lock can be seen as a batching machine and the jobs are the arriving ships with release dates and equal processing times, and the flow time of a job is the waiting time of a ship. Following the notation of Baptiste (2000) this




Fig. 1. Lockmaster's problem: concepts.

is problem $1|p - \text{batch}, b = n, r_i, p_i = p|\sum w_i F_i$. In words: the problem has a single parallel batching machine with unrestricted capacity ($b = n$), release dates on the jobs, and uniform processing times. The objective is to minimize the sum of weighted flow times ($\sum w_i F_i$), although we first discuss the problem where only unit weights are considered. Baptiste (2000) shows that this problem is polynomially solvable for a variety of objective functions. Cheng, Yuan, and Yang (2005) developed an $O(n^3)$ algorithm for $1|p - \text{batch}, b = n, r_i, p_i = p|f$ where f can be any regular objective function. Condotta, Knust, and Shakhlevich (2010) show that feasibility of the same problem with bounded capacity and deadlines can be checked in $O(n^2)$, even for a setting with parallel batching machines. Ng, Cheng, Yuan, and Liu (2003) study a single machine serial batching scheduling problem with release dates and identical processing times. Machine setup only happens after arrival of the final job in a batch and there is a fixed setup time equal to s . The completion time of a batch is equal to the sum of the processing times of the jobs in the batch. This problem is equivalent to the uni-directional lockmaster's problem with $s := T$ and $p_i = p = 0$, and can be solved in $O(n^5)$ by a dynamic programming algorithm described in Ng et al. (2003). Clearly, the lockmaster's problem is more general. Indeed, in case of both upstream and downstream-bound ships, we are dealing with two families of jobs, and only jobs of the same family can be together in a batch. Further, in our case, processing a batch of one family needs to be alternated by processing a (possibly empty) batch containing jobs of the other family; i.e. it is not possible to process two batches of the same family consecutively.

The concept of a "family" of jobs is also described by Webster and Baker (1995), be it without a batch processing machine. They deal with a scheduling problem in which setup times can be reduced by consecutively scheduling jobs of the same family. This type of problem is also known as batch scheduling with job compatibilities. Jobs within a batch need to be pairwise compatible, and these compatibilities can be expressed using a compatibility graph. Boudhar (2003) and Finke, Jost, Queyranne, and Sebó (2008) study different variants of these batch scheduling problems when the compatibility graph is bipartite or an interval graph. The compatibility graph of the lockmaster's problem is the union of two cliques.

The lockmaster's problem can be summarized as $1|p - \text{batch}, b = n, r_i, \Phi = 2, s_{fg}, p_i = 0|\sum F_i$, with $s_{fg} = 2T$ if $f = g$ and $s_{fg} = T$ if $f \neq g$, where Φ refers to the number of families and s_{fg} to the setup times between batches. For a review on scheduling a batching machine we refer the reader to Potts and Kovalyov (2000) and Brucker et al. (1998). A related problem is studied by Lee, Uzsoy, and Martin-Vega (1992) who develop dynamic programming algorithms for scheduling a batching machine with release dates, deadlines and constant processing times when the goal is to minimize makespan or minimize the number of tardy jobs. In conclusion, the complexity of the lockmaster's problem does not follow from results in literature.

3. A polynomial time algorithm for the lockmaster's problem

In this section, we construct a graph such that the shortest path in the graph corresponds to a solution for the lockmaster's problem. We show that the corresponding shortest path algorithm runs in $O(n^3)$ time. In addition, we propose a more sophisticated implementation in Section 4, which results in an $O(n^2)$ algorithm.

3.1. Definitions and terminology

For clearness of presentation, let us describe some terminology. We define a *lock movement* as the act of bringing the water level down from its high point to its low point, or vice versa, plus loading and unloading the lock. Such a lock movement takes T time units, with $T > 0$; thus, when a lock movement starts at time t , it finishes at $t + T$. A pair of lock movements is called *consecutive* when the starting time of the second lock movement equals the starting time of the first lock movement plus T . Clearly, after two consecutive lock movements starting at time t , the lock is back at the same position at time $t + 2T$ as it was at time t . A set of ℓ lock movements is called *consecutive* when the pairs $(i, i + 1)$ are consecutive, $i = 1, \dots, \ell - 1$. Clearly, when a set of ℓ consecutive lock movements start at time t , it finishes at $t + \ell T$. Recall that each $a \in A$ has an associated *arrival time* $t(a) \in \mathbb{N}$ and an *arrival position* $p(a) \in \{0, 1\}$ (1 for upstream and 0 for downstream).

Definition 1. For each pair $a, b \in A$ with $p(a) = p(b)$, and $t(a) + 2T < t(b)$, we define a *block* $B = (a, b)$ as the set of $\ell(B)$ consecutive movements that starts at $t(a)$ and ends at $t(a) + \ell(B)T$, where $\ell(B)$ is the largest even integer such that $t(a) + \ell(B)T < t(b)$.

For each pair $a, b \in A$ with $p(a) \neq p(b)$, and $t(a) + T < t(b)$, we define a *block* $B = (a, b)$ as the set of $\ell(B)$ consecutive movements that starts at $t(a)$ and ends at $t(a) + \ell(B)T$, where $\ell(B)$ is the largest odd integer such that $t(a) + \ell(B)T < t(b)$.

A geometric visualization of an instance and a corresponding solution is given in Fig. 1, where upstream arrivals appear on the upper horizontal line and downstream arrivals on the lower horizontal line, and the schedule is represented by the lock movements between both lines.

We see a solution to the lockmaster's problem as a set of lock movements, each lock movement starting at some given time. Such a solution will be called a *schedule*. However, some schedules are more interesting than others. We first identify schedules that possess a certain structure.

Definition 2. A schedule is called a *block-schedule* if it consists of a sequence of blocks B_1, \dots, B_s where $B_q = (a_q, b_q)$, $q = 1, \dots, s$, and $b_q = a_{q+1}$ for each $q = 1, \dots, s - 1$; and with a final set of consecutive lock movements starting at $t(b_s)$, and ending not later than $t(n) + 3T$.

Notice that the distinguishing feature of a block-schedule is that each movement starts at an arrival time or it directly follows the previous movement. Now, we are ready to state the following characteristic.

Lemma 1. *There is an optimal lock schedule that is a block-schedule.*

Proof. We first observe that the following properties characterize block-schedules:

Property 1. Each lock movement either directly follows upon a previous lock movement, or starts upon some $t(a)$ while containing ship a .

Property 2. The length of a period in which there is no lock movement (a so-called *empty period*) is at most $2T$.

Property 3. The final lock movement does not end later than $t(n) + 3T$.

It is easily verified that any schedule satisfying **Properties 1–3** is a block-schedule, and any block-schedule satisfies these properties.

We now prove **Lemma 1** by showing that any schedule not satisfying the above properties can be transformed into a schedule satisfying them without increasing the objective value. Consider some schedule for the lockmaster’s problem not satisfying **Property 1**. Let t denote the earliest time where a lock movement starts, such that t is neither the ending time of a previous lock movement, nor the arrival time $t(a)$ of some ship $a \in A$ contained in this lock movement. Since at $t - \epsilon$, with $\epsilon > 0$ and small, the lock is waiting to go up, and since no ships contained in the lock movement arrive between $t - \epsilon$ and t , we could have started this lock movement at $t - \epsilon$ without increasing the objective function value. Thus, we start the lock movement earlier in time, at either (1) the latest arrival time of a ship contained in the lock movement, or (2) the ending time of the preceding lock movement.

Next, consider a solution where no lock movement occurs during a period $S > 2T$. At the beginning of this period, two additional movements can be scheduled, reducing S by $2T$ time units. By repeating this step we end up only with empty periods shorter than $2T$ (**Property 2**).

Further, assume a solution exists with a final lock movement that does not end before $t(n) + 3T$. Then, the last movement contains no ships and, hence, can be dropped. By repeating, we obtain a solution which ends not later than $t(n) + 3T$ (**Property 3**). □

Furthermore, since all ships are identical, ships can be interchanged in any solution to the lockmaster’s problem, such that the following property holds:

Property 4 [FCFS]. For each pair of ships $a, b \in A$ with $p(a) = p(b)$: if $a < b$, then ship a will leave the lock not later than ship b .

Now, we describe an algorithm solving the lockmaster’s problem in polynomial time. The basic idea is to build a directed acyclic graph $G = (V, E)$ with a given cost c_e for each $e \in E$. The arcs in the graph represent the blocks of a schedule, the situation before the first block or the final set of movements. After describing this graph, we argue that a path in this graph with a certain cost, corresponds to a block-schedule with a total waiting time equal to this cost, and vice versa. Thus, a shortest path corresponds to a solution to the lockmaster’s problem.

3.2. Constructing the graph

There is a node s , a node t , and for each $a \in A$, there is a set of $n + 2$ nodes that we denote by a *layer* of nodes $L(a)$. One node from this layer is called the *top node*, indicated by a_{top} . All other nodes of the layer are indexed by $k = 0, \dots, n$, and are denoted by a_k . Hence V has $O(n^2)$ nodes in total.

There are five types of arcs, namely arcs leaving s , arcs entering t , so-called block1 arcs, block2 arcs, and 0-cost arcs; we now describe these arcs and the corresponding costs. See **Fig. 3** for a graphical representation, corresponding to the instance shown in **Fig. 2**, assuming $T = 30$.




Fig. 2. Lockmaster’s problem: example instance.

There is an arc from s to a single node from each layer $L(a)$, $a \in A$, say node a_k , with $k \in \{0, \dots, n\}$. The value of k is determined by the number of ships arriving in $(-\infty, t(a) + T]$ at position $1 - p(a)$. This arc represents the situation where the first block of a schedule starts at time $t(a)$ and position $p(a)$. Accordingly, the cost of this arc equals the waiting time accumulated at time $t(a)$ of all ships arriving in $(-\infty, t(a)]$ at position $p(a)$, plus the waiting time accumulated at time $t(a) + T$ of all ships arriving in $(-\infty, t(a) + T]$ at position $1 - p(a)$.

There is an arc from a top node a_{top} from each layer $a \in A$ to node t . The cost of this arc equals the waiting time of all ships arriving in $[t(a), t(n)]$ at position $p(a)$, and in $[t(a) + T, t(n)]$ at position $1 - p(a)$, in a solution where at time $t(a)$ a set of consecutive movements starts, serving all these ships.

We now describe the block2 arcs. For each block $B = (a, b) \in A \times A$ with $\ell(B) \geq 2$, there is an arc from node a_{top} to node $b_l \in L(b)$, where l equals the number of ships that arrive at position $1 - p(b)$ in the interval $(t(a) + (\ell(B) - 1)T, t(b) + T]$. The cost of such an arc equals the waiting time accumulated in block B at time $t(b) + T$, of, in case $p(a) = p(b)$, all ships arriving in $(t(a), t(b)]$ at position $p(a)$, and in $(t(a) + T, t(b) + T]$ at position $1 - p(a)$; in case $p(a) \neq p(b)$, all ships arriving in $(t(a), t(b) + T]$ at position $p(a)$, and in $(t(a) + T, t(b)]$ at position $1 - p(a)$. These are the block2 arcs.

Fig. 4 illustrates a block (a, b) of length at least 2. It can be seen easily that all ships arriving in $p = 1$ in $(t(a) + T, t(a) + 3T]$ are moved at time $t(a) + 3T$. Similarly, all ships arriving in $p = 1$ in $(t(a) + (\ell(a, b) - 2)T, t(b)]$ are moved at time $t(b)$. Finally, notice that ships arriving in $p = 0$ in $(t(a) + (\ell(a, b) - 1)T, t(b) + T]$ are moved at time $t(b) + T$.

We now describe the block1 arcs. For each block $B = (a, b) \in A \times A$ with $\ell(B) = 1$, there is an arc from each node a_k ($k = 0, \dots, n$) to some node b_l from layer $L(b)$ where l equals the number of ships that arrive at position $1 - p(b)$ in the interval $(t(a), t(b) + T]$. The cost of such an arc consists of two parts, first, the waiting time accumulated at time $t(b) + T$ of all ships arriving in $(t(a), t(b) + T]$ at position $p(a)$, and in $(t(a) + T, t(b)]$ at position $p(b)$; and second, $k \times (t(b) - t(a) - T)$. **Fig. 5** illustrates a block (a, b) of length 1. Ships arriving in the time interval represented by the dashed section at $p = 0$, are transported by a movement starting at time $t(b)$. Since $t(b) > t(a) + T$, we need to take this additional waiting time into account, which is achieved by the second term described above. Note that, by construction of the block2 arcs, the index k of a_k corresponds precisely to the number of these waiting ships for block2 arcs ending at a_k .

Finally, within each layer $L(a)$, $a \in A$, there is an arc with cost 0 that goes from each node a_k , $k = 0, \dots, n$, to node a_{top} ; these are the 0-cost arcs. This completes the description of graph G . We now prove a lemma that establishes the correspondence between block-schedules and paths in G .

Lemma 2. *An s, t path in G corresponds to a block schedule and vice versa.*

Proof. Consider some path from s to t in G . Clearly, the path will visit some nodes of some layers. More precisely, the path can only visit a node in some layer by entering a node a_k in $L(a)$ for some $k \in \{0, \dots, n\}$. The cost defined above, of any arc entering node a_k assumes that the lock will move at $t(a)$. Next, there are two ways of leaving node a_k : either, the path proceeds with a 0-cost arc to the top




Fig. 3. Lockmaster's problem: the graph.




Fig. 4. Waiting times covered by block (a, b) of length at least 2.




Fig. 5. Waiting times covered by block (a, b) of length 1.

node, or it proceeds with a block1 arc. Proceeding with a 0-cost arc means that the lock will move at $t(a) + T$, since all arcs leaving the top nodes correspond to blocks of length at least 2. Proceeding with a block1 arc to some node b_l means that there is no lock movement starting at $t(a) + T$. In this situation, by construction, k ships are waiting at position $1 - p(a)$ at time $t(a) + T$. The waiting time of these ships after $t(a) + T$ is included in the cost of the block1 arc leaving node a_k . Hence, the cost on the arc entering node a_k is defined appropriately. Finally, note that the arcs leaving s represent the waiting time before the first block and the arcs entering t represent the final set of consecutive movements. It follows that the cost of each path

from s to t corresponds to the total waiting time of a block schedule. The reverse is easy to see: any block schedule can be mimicked by choosing the appropriate arcs. \square

Theorem 1. *The lockmaster's problem can be solved in $O(n^3)$ by a straightforward implementation of a shortest path algorithm on the acyclic graph G .*

Proof. The previous Lemma, combined with Lemma 1 and the observation that the graph G is acyclic, and contains $O(n^3)$ arcs, imply this result. \square

Notice that the problem definition does not impose a starting position for the lock. In case the starting position of the lock is pre-specified, it is easy to modify the algorithm to deal with this feature.

4. A faster algorithm for the lockmaster's problem

We describe here a more efficient implementation that uses the structure in the graph G to solve the lockmaster's problem. The resulting complexity is $O(n^2)$. The method used to obtain this speed-up shows similarities to the approaches proposed by Wagelmans, Van Hoesel, and Kolen (1992), Federgruen and Tzur (1991), and Aggarwal and Park (1993), where an improvement procedure is described to solve the well-known lot sizing problem. We use a specific forward labeling algorithm, where we label the nodes layer by layer, and within a layer, we start with a_0 , then a_1 , up to a_n , and finally a_{top} .

Define $sp(a_k)$ as the shortest path length from s to node a_k , with $k = 0, \dots, n$. Further, define $sp^{a_k}(b_i)$ as the shortest path length to node b_i , with $i = 0, \dots, n$; $b = a + 1, \dots, n$, when the node visited just before b_i is a_k . Note that the arc (a_k, b_i) is a block1 arc in G , which is defined only when $t(a) + T < t(b)$ and $p(a) \neq p(b)$.

There are $O(n^2)$ arcs leaving nodes from layer $L(a)$, $O(n)$ block2 arcs and $O(n^2)$ block1 arcs. A shortest path to some node b_i has as last arc either a block1 arc, or some arc that is not a block1 arc. It follows that, for each $b_i \in L(b)$,

$$sp(b_i) = \min \left\{ \min_{\substack{a < b \\ k=0, \dots, n}} \{sp^{a_k}(b_i)\}, \min_{a < b} \{sp^{a_{top}}(b_i)\}, sp^s(b_i) \right\}.$$

In the following, we argue that once $sp(a_k)$ is determined for arbitrary a and each $k = 0, \dots, n$, all $sp^{a_k}(b_i)$, i.e. shortest paths determined by block1 arcs leaving $L(a)$, can be obtained in $O(n)$ time.

For a given b_i , it holds that $sp^{a_k}(b_i) = sp(a_k) + c(a_k, b_i)$, whereby $c(a_k, b_i) = w(a, b) + k(t(b) - t(a) - T)$ is the length of the block1 arc (a_k, b_i) in G . The term $w(a, b)$ represents the waiting time accumulated at time $t(b) + T$ of all ships arriving in $(t(a), t(b) + T]$ at position $1 - p(b)$ and the waiting time of all ships arriving in $(t(a) + T, t(b)]$ at position $p(b)$; this is a constant over all a_k, b_i for which a block1 arc (a_k, b_i) exists. Note that for each a_k , there is at most one b_i from layer $L(b)$ for which an arc (a_k, b_i) exists. We argue that finding all $sp^{a_k}(b_i)$ for $k = 0, \dots, n$ and for $b = a + 1, \dots, n$ (i is fixed, given a_k and b), corresponds to finding the minimum value of at most n linear functions ($\forall k = 1, \dots, n$) for no more than n inputs ($\forall b = a + 1, \dots, n$).

Lemma 3. Given m linear functions $f_i(x) = \alpha_i \cdot x + \beta_i$ with $\alpha_i, \beta_i \in \mathbb{R}$ for each $i = 1, \dots, m$, $\alpha_i \leq \alpha_{i+1}$ for each $i = 1, \dots, m - 1$, and an ordered finite set Q . Then the indices $\arg \min\{f_i(q) \mid i = 1, \dots, m\}$ for each $q \in Q$ can be found in $O(m + |Q|)$.

Proof. It is clear that we may assume w.l.o.g. that no two of the given functions are identical. We start by eliminating dominated functions. We say that f_i is dominated by $f_j, j < i$, if $\beta_j \leq \beta_i$ because then $f_j(q) \leq f_i(q)$ for each $q \geq 0$. Henceforth, we assume that $\beta_i > \beta_{i+1}$ for each $i = 1, \dots, m - 1$ in the following.

Next, we find the lower envelope of the functions f_1, \dots, f_m . We refer to Sack and Urrutia (2000) for an overview of methods devoted to computing the lower envelope of a set of functions. We maintain a list \mathcal{L} of active functions which potentially contribute to the lower envelope. Note that functions may be dropped from this list later on. Initially, this list contains f_m and f_{m-1} . We consider the functions $f_{m-2}, f_{m-3}, \dots, f_1$ in this order. Let f_i be the function to be considered, and f_v and f_u the last and next-to-last function contained in \mathcal{L} , respectively. We determine the q -coordinate $q_{i,v} \equiv \frac{\beta_i - \beta_v}{\alpha_v - \alpha_i}$ of the intersection of f_i and f_v and the q -coordinate $q_{i,u} \equiv \frac{\beta_i - \beta_u}{\alpha_u - \alpha_i}$ of the intersection of f_i and f_u . If $q_{i,u} \geq q_{i,v}$, then we remove f_v from \mathcal{L} since $f_v(q) \geq \min\{f_i(q), f_u(q)\}$ for each $q \geq 0$, and we repeat this step with




Fig. 6. Lower envelope.

$\mathcal{L} := \mathcal{L} \setminus \{f_v\}$ until $q_{i,u} < q_{i,v}$. We then add f_i to \mathcal{L} , i.e. let $\mathcal{L} := \mathcal{L} \cup \{f_i\}$, and consider the next function. Note that at the end of each iteration, \mathcal{L} must contain at least two functions since f_m is not dominated due to assumption.

Finally, by scanning through the sorted set of intersections of linear functions and through values in Q , we find $\arg \min\{f_i(x) \mid i = 1, \dots, m\}$ for each $q \in Q$. For a graphical representation of the lower envelope, see Fig. 6.

It remains to show that this procedure runs in linear time. Eliminating dominated functions takes $O(m)$. During the construction of \mathcal{L} , each function is added at most once and each function is deleted at most once. The decision whether to add the next function at the end of \mathcal{L} or remove the currently last function from \mathcal{L} can be taken in constant time. Thus, constructing \mathcal{L} takes $O(m)$ and, consequently, finding $\arg \min\{f_i(q) \mid i = 1, \dots, m\}$ for each $q \in Q$ takes $O(m + |Q|)$. \square

We now use Lemma 3 and consider all block1 arcs leaving $L(a)$, for a given $a \in A$. Consider the linear functions $sp^{a_k}(b_i) = sp(a_k) + w(a, b) + k(t(b) - t(a) - T)$, for $k = 0, \dots, n$ and for $b = a + 1, \dots, n$. As mentioned before, $w(a, b)$ is a constant for any given a and b . We thus exclude it from the linear functions to be considered for the lemma, and add this constant later. Let $Q = \{t(b) - t(a) - T \mid b = a + 1, \dots, n\}$ and $f_k(q) = sp(a_k) + q \cdot k$ for each $0 \leq k \leq n$. Thus, referring to Lemma 3, $\alpha_k = k$, $\beta_k = sp(a_k)$, and $q \in Q$. In Lemma 3 it is assumed that the α_i and the set Q are ordered, which is obviously the case. Now, $\arg \min\{sp^{a_k}(b_i) \mid 0 \leq k \leq n\}$ for each b_i determines the shortest path to b_i having a block (a_k, b_i) of length one as last part. Note that its length is $\min\{f_k(q) \mid 0 \leq k \leq n\} + w(a, b)$. Hence, we can update all $sp(b_i)$ taking into account all block1 arcs leaving layer $L(a)$ in $O(n)$ time.

Applying this procedure to all layers $L(a)$, with $a \in A$, all shortest paths with correct cost values are obtained in $O(n^2)$. The overall procedure to find the shortest path in $O(n^2)$ is shown in Algorithm 1. We can summarize the above in the following theorem.

Theorem 2. The lockmaster's problem is solvable in $O(n^2)$.

5. Extensions

It can be argued that the basic problem defined in Section 1, due to different assumptions regarding the input, ignores a number of issues regarding practical lock operation. We now proceed by showing how the procedure described in Section 3 can be extended towards a closer approximation of reality.

5.1. Capacity

Section 3 did not take any capacity restrictions into account. We now discuss two different settings with capacity restrictions. In one setting, the lock can accommodate at most a given number of ships; in another setting each ship has a size, the total size of ships within a lock should then not exceed the given lock's size. A fundamental difference with the basic lockmaster's problem is that ships may be

Algorithm 1 $O(n^2)$ algorithm for finding the shortest path in G .

```

construct graph  $G$  as described above
 $sp(a_k) \leftarrow \infty$  for all  $a \in A$  and  $0 \leq k \leq n$ 
for  $a \in A$  do
  consider block  $(s, a)$  and update  $sp(a_k)$  for  $k$  corresponding
  to  $(s, a)$ 
end for
for  $a = 1, \dots, n$  do
  set  $sp(a_{top}) = \min \{sp(a_k) \mid 0 \leq k \leq n\}$ 
  determine the lower envelope for all block1 arcs  $(a_k, b_i)$ 
  with  $k = 0, \dots, n$  and  $b = a + 1, \dots, n$  ( $i$  follows from  $a_k$ 
  and  $b$ )
  for  $b = a + 1, \dots, n$  do
    if  $l(a, b) == 1$  then
      update  $sp^{a_k}(b_i)$  using the lower envelope
    else
      update  $sp^{a_k}(b_i)$  using  $sp(a_{top})$ 
    end if
  end for
end for

```

“left behind”. Thus, ships that lay waiting may not all be transferred in the same lock movement.

When the number of ships that may be contained within the lock at any given time is bounded by a constant, a similar procedure to Section 3 may be used.

Theorem 3. *The lockmaster’s problem with a bound on the number of ships in the lock is solvable in $O(n^4)$.*

Proof. Suppose that the lock can accommodate at most c ships at once. We can reformulate Property 3 as follows: the final lock movement does not end later than $(1 + \lceil \frac{n}{c} \rceil)2T$. For this generalization, Lemma 1 remains true. Indeed, Properties 1 and 2 remain trivially true; moreover, the ships being identical (except for their arrival time) implies that there exists an optimal schedule that satisfies Property 4. \square

Definition 3. For each $a \in A$, let $U_a = \{b \in A \mid p(b) = 1, b < a, t(b) \leq t(a) - |p(b) - p(a)|T\}$. For each $a \in A$, let $D_a = \{b \in A \mid p(b) = 0, b < a, t(b) \leq t(a) - |p(b) - p(a)|T\}$.

Thus, the set $U_a (D_a)$ contains all ships b for which $b < a$, that arrive upstream (downstream), not later than $t(a)$ if ship a arrives upstream (downstream), and not later than $t(a) - T$ if ship a arrives downstream (upstream).

We now design the following directed acyclic graph $G = (V, E)$ in order to represent the problem. In contrast to the graph used for the lockmaster’s problem, here, each arc directly corresponds to a set of ships transferred, and there is no need to distinguish block1 and block2 arcs. We have $V = \{s, t\} \cup V^A$ with $V^A = \{(a, u, d) \mid a \in A, u \in U_a, d \in D_a\}$. A node (a, u, d) corresponds to a block starting at $t(a)$ and position $p(a)$ while u and d represent the latest (with respect to the ordering in A) served downstream- and upstream-bound ship, respectively. It is easily seen that whenever a ship a does not enter a lockage starting not earlier than $t(a)$ from position $p(a)$ while sufficient lock capacity is available to serve this ship, the solution is sub-optimal. In combination with Property 4 it follows that, within each block, each ship is served by the next appropriate lockage with sufficient capacity, and it is thus known when each ship is handled. We thus know which ships u' and d' become the latest ships served upstream and downstream, respectively. We record this observation by writing $u' = f_{up}((a, b), u)$ and $d' = f_{down}((a, b), d)$ for each block $(a, b), u \in U_a, d \in D_a$.

Definition 4. We say that two nodes (a, u, d) and (b, u', d') are compatible if (a, b) is a block and if $u' = f_{up}((a, b), u)$ and $d' = f_{down}((a, b), d)$.

The set of edges in our graph is given as $E = E^s \cup E^B \cup E^t$, with

- $E^s = \{(s, v) \mid v = (a, \emptyset, \emptyset) \in V^A\}$,
- $E^B = \{(v, v') \mid v \text{ and } v' \text{ are compatible}\}$, and
- $E^t = \{(v, t) \mid v = (a, u, d) \in V^A\}$.

The cost $c(v, v')$ of the edges in E^B , with $v = (a, u, d)$ and $v' = (b, u', d')$, equals the sum of waiting times of all ships k with $u < k \leq u'$ arriving at $p = 1$ and the sum of waiting times of all ships ℓ with $d < \ell \leq d'$ arriving at $p = 0$, such that these ships are handled while respecting the lock capacity.

The costs $c(s, v)$ on edges leaving the origin are equal to 0, since no ship has been transferred yet. The costs $c(v, t)$, with $v = (a, u, d)$, represent the waiting times of ships k with $u < k$ arriving at $p = 0$ and the waiting times of ships ℓ with $d < \ell$ arriving at $p = 1$. These ships are served by the final sequence of lock movements.

Then, a path from s to t represents a feasible lock schedule and the shortest path represents the lock schedule having minimum total waiting time of ships. G is acyclic and contains $O(n^3)$ nodes, each node being source of at most $O(n)$ arcs. Furthermore, we can determine u', d' , and total waiting of ships according to all arcs emerging from $v \in \{s\} \cup V^A$ in $O(n)$ time. Constructing G and using a straightforward shortest path algorithm allows us to find a solution in $O(n^4)$ time.

A more general setting assigns to each ship and lock an arbitrary size. When each ship and lock has an associated length and width, the problem is easily seen to be NP-hard by reduction from rectangle packing, as mentioned by Hermans (2014). We extend this result to instances where the size of the lock is represented by a scalar value. For the details of the proof, based on reduction from 3-PARTITION, we refer to Passchyn et al. (2015).

Theorem 4. *The lockmaster’s problem with a (scalar) bound on the size of the lock is NP-hard.*

We may, however, impose the requirement that all ships traveling in the same direction must be handled in a predetermined order. More specifically, we will require that ships traveling in the same direction are handled according to the total order on A . We will say that solutions satisfying this requirement adhere to a first-come first-served (FCFS) policy. The FCFS policy is currently applied as the standard handling policy for ships at many locks, see e.g. Smith et al. (2009); Ting and Schonfeld (2001); van Haastert (2003). When this requirement is enforced, an efficient procedure exists for finding an optimal solution.

Theorem 5. *Under a FCFS policy, the lockmaster’s problem with a bound on the size of the lock is solvable in $O(n^4)$.*

Proof. In a preprocessing step, the subsets of consecutive ships that fit together in the lock, are determined. Here, it is possible to include all kinds of filling and entering rules that can be important in practice, see e.g. Verstichel and Vanden Berghe (2009). This step can be executed in $O(n^2)$ time. We use the same graph as defined for the setting with bounded capacity, described above. Note that when determining u', d' , and total waiting time of ships according to arcs, we can only fill the lock with ships that fit together in the lock, which we determined initially. Still, we can treat all arcs emerging from a node in $O(n)$ time. \square

5.2. Ship priorities

Minimizing the weighted sum of waiting times allows us to take into account ship priorities. This is often relevant, for instance, in

locks operating in ports. It is indeed quite common to distinguish between sea ships (having limited manoeuvrability) and inland ships. In addition, ships transporting dangerous goods receive priority over regular cargo ships (Du & Yu, 2003), which in turn may have priority over leisure ships, see e.g. (Smith et al., 2009; Verstichel & Vanden Berghe, 2009). All this can be dealt with by assigning a weight w_a to each ship $a \in A$, revealing their priority. We can generalize the construction of the graph G , and in particular the arc costs, to this setting by multiplying the waiting time of ship a by its weight w_a . We state without proof:

Theorem 6. *The lockmaster's problem with objective to minimize total weighted waiting time is solvable in $O(n^3)$.*

Notice that when the ships are weighted and the capacity of the lock is limited, the algorithm described in Section 5.1 might fail to find an optimal solution. Earlier, specifying a block would implicitly specify the set of ships that are transferred in this block. This, however, need then no longer be the case; although the optimal solution can still be seen as consisting of blocks, it is not clear which particular ships are transferred during which movements in a block. When considering the uni-directional case, Baptiste's algorithm (Baptiste, 2000) (see Section 2) yields a polynomial time procedure.

5.3. Handling times

In practice, placing a ship into a lock takes a certain amount of time, such that the total time a ship spends in the lock may depend on the other ships present. Smith et al. (2011) describe a lock scheduling problem where handling times are not only ship-dependent, but also sequence dependent. We here consider the setting where each ship $a \in A$ requires an integral handling time h_a . We model this by modifying the lockage time so that it is no longer constant: $T_j = T + \sum_{a \in A_j} h_a$, where A_j refers to the set of ships transferred in the j th lock movement, and T_j denotes the corresponding lockage duration, $j = 1, 2, \dots$

Observe that, with distinct h_a values, it may be optimal to let a ship with a relatively large handling time wait while handling subsequent arrivals first. In general, all optimal solutions may thus violate Property 4.

We first state that the general setting for the lockmaster's problem with handling times is NP-hard. The proof, by reduction from 3-PARTITION, is described by Passchyn et al. (2015).

Theorem 7. *The lockmaster's problem with ship-dependent handling times is NP-hard.*

Observe that, because of the handling times, the number of ships contained in a lockage influences the possible starting times of later lockages. Thus, it may be optimal for a ship not to enter the first possible lockage in the ship's direction of travel. In the absence of handling times, i.e. Section 3, not entering the first available lockage would clearly be suboptimal. Despite this, we show that, in the FCFS case, a polynomial time dynamic programming approach exists. Note that this also applies to the setting where all handling times are equal, i.e. the case where $h = h_a$ for all $a \in A$. For a detailed description of the algorithm, and thus the proof of Theorem 8, we refer to Passchyn et al. (2015).

Theorem 8. *The lockmaster's problem with arbitrary handling times under a FCFS policy can be solved in $O(n^{10})$ time.*

5.4. Non-uniform lockage duration

Lockage times for upstream-bound (T_u) and downstream-bound (T_d) may differ. Passchyn et al. (2015) argue that the procedure for the lockmaster's problem outlined in Section 3 is easily adjusted to take non-uniform lockage durations into account.

Theorem 9. *The lockmaster's problem with a lockage duration that depends on the position of the lock is solvable in $O(n^3)$.*

5.5. Water usage

Due to organizational/environmental reasons, the number of lock movements allowed in some time-interval may be limited. In particular when water is scarce (e.g. after dry seasons). In such a case, the number of allowed lockages is bounded in order to keep the water at a navigable level (Verstichel & Vanden Berghe, 2009). Again a modification to the original procedure can be found which finds an optimal solution. This procedure runs in polynomial time provided that the bound Q on the number of lockages is part of the input. We refer to Passchyn et al. (2015) for the full description.

Theorem 10. *The lockmaster's problem with a bound Q on the number of lock movements is solvable in $O(Q^2 n^4)$.*

6. Computational study

Is it worth the effort to solve instances of the lockmaster's problem to optimality? Or, are heuristics sufficient to achieve near-optimal solutions? In this section we answer this question by performing computational experiments. We consider the basic problem setting as well as a number of extensions and compare the optimal solution to a number of straightforward heuristics.

6.1. Instances and problem setting

As far as we are aware, the only publicly available instances are maintained by Verstichel and Vanden Berghe (2009). Using the shortest path procedure described in Section 3, we can find a solution for these instances. We refer to Passchyn et al. (2015) for these results.

As these instances do not contain all information needed for the extensions we cover here, we also generate new instances for further testing. We simulate a realistic arrival process with ships arriving independently over a time horizon of 24 hours. The time unit is one minute and the lockage duration is assumed to be 30 minutes. The arrival rate was estimated from a dataset with arrival times for the locks of Terneuzen, The Netherlands, provided by Rijkswaterstaat, the dutch waterway management organization. The values $p = \frac{1}{30}$, $p = \frac{1}{15}$, and $p = \frac{1}{10}$ yield instances with an average of approximately 50, 100, and 150 ships respectively. Passchyn et al. (2015) provide more details on the instances and arrival distribution.

We consider the following problem settings:

1. **Basic:** The basic settings described in Section 3. This considerably simplified problem serves as the starting point for all cases below. Further, it acts as a reference to compare the relative performance of the different solution methods.
2. **Capacity:** The extension covered in Section 5.1. We assume that the number of ships in the lock is limited for any lock movement. We arbitrarily set the bound equal to 3 for all capacitated instances below. Separate simulations with capacity values of 6 and 12 respectively, resulted only in minor differences with the basic setting, indicating that the optimal solution is rarely restricted by these capacity bounds. For this reason, results for the larger capacity bounds are omitted.
3. **Weights:** The extension considered in Section 5.2. Lock capacity is not limited in this case. The generated instances each have ship priority values chosen randomly from $\{1, 2, 3\}$ with equal probability.
4. **Handling times:** The extension considered in Section 5.3. Handling time values are chosen uniformly random from $\{0, 2, 5\}$. Ship priorities and lock capacity are not considered. Note that finding the optimum solution is NP-hard as shown in Section 5.3.

Table 1

Results for different problem settings. Each value represents the average of 25 instances simulating the same arrival process. The arrival parameter equals 1/30.

$p = 0.0333$ exact (min.)	OPT	Basic 727.3 (percent)	Capacity 793.8 (percent)	Weights 1381.4 (percent)	Handling X	cap + W + HT X
Heuristic (percent)	CUD	201.6	195.8	214.0	1555.6	3296.1
	MA	192.3	187.8	204.2	1557.6	3357.5
	LA2T	107.2	109.4	104.6	X	X

Table 2

Results for different problem settings. Each value represents the average of 25 instances simulating the same arrival process. The arrival parameter equals 1/15.

$p = 0.0666$ exact (min.)	OPT	Basic 1884.9 (percent)	Capacity 3768.8 (percent)	Weights 3644.0 (percent)	Handling X	cap + W + HT X
Heuristic (percent)	CUD	152.3	130.4	158.5	3524.6	11972.6
	MA	153.1	132.3	159.2	3516.6	12049.2
	LA2T	105.8	107.6	104.9	X	X

Table 3

Results for different problem settings. Each value represents the average of 25 instances simulating the same arrival process. The arrival parameter equals 1/10.

$p = 0.1$ exact (min.)	OPT	Basic 3361.3 (percent)	Capacity 21592.2 (percent)	Weights 6474.6 (percent)	Handling X	cap + W + HT X
Heuristic (percent)	CUD	137.1	109.7	141.1	6166.0	60722.2
	MA	134.8	108.4	139.6	6199.2	60003.0
	LA2T	105.3	103.7	105.3	X	X

6.2. Heuristics

Let us also define the behaviour of each of the heuristics. Each of these will be applied to all problem settings where possible. Since the solution values for these heuristics depend on the initial position of the lock (whereas it does not for the exact solution), we obtain a heuristic solution for both starting positions and report the best of both solutions. Some additional heuristics not mentioned here, and their results, are discussed by Passchyn et al. (2015).

Continuous up/down (CUD): The lock will move whenever possible, resulting in a continuous up/down movement. This strategy uses no information whatsoever, and even ignores the number of ships present at the lock. The first lock movement is assumed to start at $t = 0$. Lock capacity and ship handling times and weights are taken into account where relevant.

Move upon arrival (MA): The lock will not move unless a ship arrives or a ship is already waiting. Whenever a ship arrives on either side, the lock will immediately move and handle this ship (potentially by first performing an empty lock movement). When given knowledge concerning the ships that have arrived so far while having no future arrival information, this may be the most straightforward way to operate the lock. Capacity, handling times and weights are straightforward to take into account.

Look-ahead 2T (LA2T): In contrast to the CUD and MA heuristics, this procedure makes use of future arrival information. A subproblem is solved with the exact solution procedure, using arrival information for the following $2T$ time units. The solution to this subproblem then determines if the lock should move or wait for the following arrival. A more detailed description of this procedure can be found in Passchyn et al. (2015).

6.3. Results

We summarize the results of all simulations in Tables 1–3. Each table shows values obtained by averaging over 25 instances for each heuristic under different problem settings. We do not separately list the computation time as minimizing the solution time was not our main priority. Exact solutions were obtained in less than a second on average, except for the bounded capacity setting, where the in-

creased graph complexity increases the computation time up to a 10 minute average on a 3.4 GigaHertz machine with 4GigaByte RAM. Exact results are reported in minutes of total waiting time. Heuristic values are shown as a percentage relative to the corresponding exact solution where possible, and as total waiting time in minutes where the exact solution is not available. All generated instances and results for each individual instance are available online at <https://perswww.kuleuven.be/~u0086328/lockmasterdata.html>.

We see from the tables that, as should be expected, the straightforward heuristics CUD and MA perform significantly worse than LA2T. Making use of future arrival information clearly pays off. In addition, the optimality gap induced by all heuristics tends to decrease as the arrival rate increases. This is easily explained since the waiting ships will more frequently outnumber the lock capacity in this case, at which time it is always optimal to move the lock immediately with full capacity.

LA2T appears the best available strategy, provided that the optimal solution to subproblems can be found. The smaller size of subproblems may also allow an exact MIP solution in reasonable time, even for settings that are known to be NP-hard. This suggests application of this heuristic as a 'rolling horizon' strategy for large instances where the exact solution is infeasible. Increasing the look-ahead horizon will further improve the performance of this heuristic, at the cost of a rapidly increasing computation time as the horizon increases. For our instances, the LA2T computation was at least an order of magnitude faster than the exact procedure, and frequently the difference was even larger. For one of the larger instances in the basic setting however, the exact solution for the entire problem was in fact found faster, albeit only slightly, than the heuristic result.

7. Conclusion

This paper introduced the lockmaster's problem, a new problem that is closely linked to batch scheduling problems. The problem can be solved in polynomial time; we were able to build a graph such that applying a basic shortest path algorithm solves the lockmaster's problem, and several extensions, to optimality. Computational experiments confirm that this exact algorithm outperforms a number of basic heuristics.

8. Further research

In this work we study the lockmaster's problem for a single lock. A relevant question is how to deal with the problem in case of multiple locks in series, either with ships sailing downstream at the first lock and upstream at the last lock, or, more complex, with ships also arriving at intermediate locks. In general, more complex waterway networks with several locks would be a nice subject for future work. In reality, lockmasters do not know all the ship arrival times in advance, except for ships that are already within a certain distance of the lock. Studying the online version of the lockmaster's problem could capture this element in a better way. A different direction for future research is to go further into batch scheduling with three or more job families, instead of the two families related to the lockmaster's problem studied in this paper.

Acknowledgments

This research has been partially funded by the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office. This paper has grown out of an extended abstract (Coene & Spieksma, 2011) presented at the Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems. We would like to thank Cor Hurkens for an interesting discussion on the subject, Rolf Möhring for pointing out references Du and Yu (2003) and Luy (2010), and Wenchao Wei for roughly translating Du and Yu (2003). We are also grateful to Rijkswaterstaat for providing real-world arrival data for the Terneuzen locks.

References

- Aggarwal, A., & Park, J. K. (1993). Improved algorithms for economic lot size problems. *Operations Research*, 41, 549–571.
- Antwerp Port Authority (2011). Port of Antwerp starts building the largest lock in the world. <http://www.deurganckdoksluis.be/en/press/port-antwerp-starts-building-largest-lock-world>. (accessed October 7, 2011).
- Baptiste, P. (2000). Batching identical jobs. *Mathematical Methods of Operations Research*, 52, 355–367.
- Boudhar, M. (2003). Scheduling a batch processing machine with bipartite compatibility graphs. *Mathematical Methods of Operations Research*, 57, 513–527.
- Brucker, P., Gladky, A., Hoogeveen, H., Kovalyov, M. Y., Potts, C. N., Tautenhahn, T., & van de Velde, S. L. (1998). Scheduling a batching machine. *Journal of Scheduling*, 1, 31–54.
- Cheng, T. C. E., Yuan, J. J., & Yang, A. F. (2005). Scheduling a batch-processing machine subject to precedence constraints, release dates and identical processing times. *Computers and Operations Research*, 32, 849–859.
- ChinaDaily (2011). Three gorges ship lock marks 8 years of operation. http://usa.chinadaily.com.cn/china/2011-06/18/content_12730491.htm. (accessed October 7, 2011).
- Coene, S., & Spieksma, F. C. R. (2011). The Lockmaster's problem. In A. Caprara, & S. Kontogiannis (Eds.), *11th workshop on algorithmic approaches for transportation modelling, optimization, and systems: vol. 20* (pp. 27–37). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/OASICS.ATMOS.2011.27.
- Condotta, A., Knust, S., & Shakhlevich, N. V. (2010). Parallel batch scheduling of equal-length jobs with release and due dates. *Journal of Scheduling*, 13, 463–477.
- Du, J. N., & Yu, S. M. (2003). Dynamic programming model and algorithm of shiplock scheduling problem. *Computer and Digital Engineering*, 31, 47–50. (in Chinese).
- European Commission (2015). Promotion of inland waterway transport. <http://ec.europa.eu/transport/inland/promotion/promotion-en.htm>. (accessed 21 November 2015).
- Federgruen, A., & Tzur, M. (1991). A simple forward algorithm to solve dynamic lot sizing models with n periods in $o(n \log n)$ or $o(n)$ time. *Management Science*, 37, 909–925.
- Finke, G., Jost, V., Queyranne, M., & Sebö, A. (2008). Batch processing with interval graph compatibilities between tasks. *Discrete Applied Mathematics*, 156, 556–568.
- Goodban Belt LLC (2010). New York State Canal System. *Technical Report*.
- Griffiths, J. D. (1995). Queueing at the suez canal. *Journal of the Operational Research Society*, 46, 1299–1309.
- Günther, E., Lübbecke, M. E., Möhring, R. H. (2011). Challenges in scheduling when planning the ship traffic on the kiel canal. Extended abstract, 10th Workshop on Models and Algorithms for Planning and Scheduling Problems.
- Hermans, J. (2014). Optimization of inland shipping – a polynomial time algorithm for the single ship single lock optimization problem. *Journal of Scheduling*, 17, 305–319.
- Inland Navigation Europe (2014). Annual Report. *Technical Report*.
- Lee, C., Uzsoy, R., & Martin-Vega, L. A. (1992). Efficient algorithms for scheduling semiconductor burn-in operations. *Operations Research*, 40, 764–775.
- Luy, M. (2010). *Algorithmen zum scheduling von schleusungsvorgängen am beispiel des nord-ostsee-kanals*. Master's thesis. TU Berlin. (in German).
- Nauss, R. M. (2008). Optimal sequencing in the presence of setup times for tow/berge traffic through a river lock. *European Journal of Operational Research*, 187, 1268–1281.
- Ng, C. T., Cheng, T. C. E., Yuan, J. J., & Liu, Z. H. (2003). On the single machine serial batching scheduling problem to minimize total completion time with precedence constraints, release dates and identical processing times. *Operations Research Letters*, 31, 323–326.
- Panama Canal Authority (2006). Proposal for the expansion of the Panama Canal. *Technical Report*.
- Passchyn, W., Coene, S., Briskorn, D., Hurink, J. L., Spieksma, F. C. R., & Vanden Berghe, G. (2015). The Lockmaster's Problem. *Research report KBI_1530*. Leuven, Belgium: KU Leuven, Department of Decision Sciences and Information Management https://lirias.kuleuven.be/bitstream/123456789/520009/1/KBI_1530.pdf.
- Petersen, E. R., & Taylor, A. J. (1988). An optimal scheduling system for the Welland Canal. *Transportation Science*, 22, 173–185.
- Potts, C. N., & Kovalyov, M. Y. (2000). Scheduling with batching: a review. *European Journal of Operational Research*, 120, 228–249.
- Rijkswaterstaat (2010). Twentekanalen: uitbreiding sluis eefde. <http://www.rijkswaterstaat.nl/water/projectenoverzicht/twentekanalen-uitbreiding-sluis-eefde/>. (in Dutch, accessed October 7, 2011).
- Sack, J. R., & Urrutia, J. (Eds.) (2000). *Handbook of Computational Geometry*. Amsterdam, The Netherlands: North-Holland - Elsevier.
- Smith, L. D., Nauss, R. M., Mattfeld, D. C., Li, J., Ehmke, J. F., & Reindl, M. (2011). Scheduling operations at system choke points with sequence-dependent delays and processing times. *Transportation Research Part E*, 47, 669–680.
- Smith, L. D., Sweeney, D. C., & Campbell, J. F. (2009). Simulation of alternative approaches to relieving congestion at locks in a river transportation system. *Journal of the Operational Research Society*, 60, 519–533.
- Ting, C., & Schonfeld, P. (2001). Control alternatives at a waterway lock. *Journal of Waterway, Port, Coastal, and Ocean Engineering*, 127, 89–96.
- U.S. Army Corps of Engineers (2009). *Waterborne commerce from the United States*. Technical Report.
- van Haastert, M. (2003). *Planningsmodel scheepvaartafhandeling bij de noordersluis in IJmuiden*. Master's thesis. TU Delft. (in Dutch).
- Verstichel, J., De Causmaecker, P., Spieksma, F. C. R., & Vanden Berghe, G. (2014a). Exact and heuristic methods for placing ships in locks. *European Journal of Operational Research*, 235, 387–398.
- Verstichel, J., De Causmaecker, P., Spieksma, F. C. R., & Vanden Berghe, G. (2014b). The generalized lock scheduling problem: an exact approach. *Transportation Research Part E*, 65, 16–34.
- Verstichel, J., & Vanden Berghe, G. (2009). A late acceptance algorithm for the lock scheduling problem. *Logistik Management*, 5, 457–478.
- Wagelmans, A., Van Hoesel, S., & Kolen, A. (1992). Economic lot sizing: an $o(n \log n)$ algorithm that runs in linear time in the wagner-whitin case. *Operations Research*, 40, S145–S156.
- Webster, S., & Baker, K. R. (1995). Scheduling groups of jobs on a single machine. *Operations Research*, 43, 692–703.