Modelling Crew Scheduling in Off-Shore Supply Vessels

Alex Leggate, Seda Sucu, Kerem Akartunalı, Robert van der Meer

Dept. of Management Science, University of Strathclyde, Glasgow, G4 0GE, UK,

{alexander.leggate}, {seda.sucu}, {kerem.akartunali}, {robert.van-der-meer}@strath.ac.uk

Crew scheduling problems have been widely studied in various transportation sectors, such as airlines, railways and urban buses. However to date it appears that application of these problems in sea transport has been very limited. In this paper, we explore key differences between various transport settings, and propose mixed-integer programming formulations for both the crew scheduling and re-scheduling problems for a company operating a fleet of Off-shore Supply Vessels (OSVs) on a global scale. Computational results on an extensive set of problems show that our proposed models are practically applicable to generate real-time solutions. We also present a thorough statistical analysis of key problem parameters, and share insights regarding their impacts.

Key words: Sea transport; Crew scheduling; Integer programming; Off-shore supply vessels.

Mathematics Subject Classification (2000): 90B06; 90C10; 90C59.

---

1. Introduction

Crew scheduling aims to obtain work timetables (schedules) for an organisation to meet the demands of required number of staff with eligibility while incorporating sector-specific regulations and organizational rules. In simple terms, the crew scheduling problem consists of a set of employees (the crew) and a set of tasks, which often have a defined time window to be carried out and possibly a location or route, and it allocates the crew to these tasks in the best possible way subject to various legal and contractual requirements, as well as physical constraints imposed by the geographical and temporal aspects of the tasks. Crew scheduling problems have been studied in a variety of contexts within the field of transportation and logistics, since crew costs are a large (if not the largest) single cost element in their operations and even a small percentage saving in crew cost translates to a very significant amount in real terms. Moreover, due to high number of variables that are naturally binary and the number and complexity of the constraints involved, crew scheduling problems are inherently challenging, where it is very hard to obtain optimal results minimizing crew costs and finding

even a feasible schedule is often a non-trivial task. Therefore, an extensive literature covers various solution methods ranging from exact solution methods to approximation algorithms and heuristics proposed for a variety of crew scheduling problems, see, e.g., Ernst, Jiang, Krishnamoorthy, Owens and Sier (2004); Ernst, Jiang, Krishnamoorthy and Sier (2004). It is also noteworthy to remark that Pinedo (2009) provides a good overall view of all planning, scheduling and timetabling activities with regards to the general transportation setting.

The complexity of the crew scheduling problems often causes computational times required to find good solutions to be prohibitive, in particular when schedules are needed in real-time. Therefore, a great deal of literature is devoted to developing fast algorithms, mainly in the context of airlines but also for buses and trains (Van den Bergh et al.; 2013). This is in particular crucial for crew recovery problems (also referred to as disruption management), a natural extension of the crew scheduling problem, where the active schedule needs to be changed due to incidents such as break-downs, traffic- and weather-related disruptions. Novel techniques such as the size reduction of the rescheduling problem (Rezanova and Ryan (2010)) are critical for effective planning in real-time.

With all the importance attached to crew scheduling, literature for companies operating fleets of ships appears to be extremely limited. This is surprising, since crew costs in the sea transport industry are not less significant than in other transportation areas, e.g., (Stopford; 2009, pp. 226-229) notes that crew costs, from basic wages through to pension payments and travel costs, can account for up to half of a vessel's operating costs excluding fuel. This relative lack of literature has been noted by Christiansen et al. (2007) due to four potential reasons: i) Low visibility of shipping industry compared to road, rail or air; ii) Variability of maritime problems and high degree of customization required for decision support tools; iii) Strong tradition leading to limited openness to new ideas; iv) Level of uncertainty in sea operations and thus a need for frequent replanning. While the first three points are consistent with our experiences in practice, we observed that frequent replanning can be difficult and time-pressured, and hence the last point rather offers opportunities, which we discuss in Section 3.1.

There are several modes of operations in sea transportation, which can be grouped into three categories of Liner, Tramp and Industrial shipping (Christiansen et al.; 2013). We note another category not suitable for this classification is the Offshore Supply (or Service) Vessel (OSV), which provide services such as construction support, deep-water lifting and installation to industries including offshore oil and gas, and offshore wind farms (Barrett

(2008)). In this paper, we consider the case of a large company operating many OSVs to provide services to the oil industry world-wide.

Finally, we discuss key similarities and differences between sea and other transportation settings: 1) Planning horizon: While a typical airline schedules their crew on a monthly basis and bus companies often use daily or weekly schedules, a globally-operating shipping company draws schedules for several months at a time and at least a few months in advance, mainly due to long duty periods often covering weeks and operational matters such as arrangement of transportation and visas for crew. 2) Transport arrangements: While airline scheduling operates with pairings, where a crew member follows a series of flights starting and finishing at their home base, the maritime context requires crew arrangements that are practically more complex due to the possible variations of employee's home country and the port of assigned vessels, leading to extra arrangements such as for transportation and visas. 3) Expertise and eligibility: While many settings disregard crew's expertise, e.g., identical bus drivers (Huisman and Wagelmans; 2006), some arise in the airlines in a limited fashion, such as pilot's knowledge of specific aircraft or cabin crew's expertise for long-haul flights. In comparison, maritime setting is significantly more complicated due to bigger number of tasks on board, some of which require specialist skills. Moreover, crew nationality increases complexity, e.g., ships with certain flags can employ only certain nationalities, certain ports require work permits, and some nationalities might not work along well.

The paper is organized as follows: In Section 2, we present relevant literature including other related maritime optimization problems. Then, in Section 3, we describe the crew scheduling problem as it appears at our industrial partner, with the company's current approach presented before a mathematical formulation. Computational results and discussion follow in section 4, while some conclusions drawn from this research are given in section 5.

2. Literature Review

There is a large body of crew scheduling research in other transportation settings, as highlighted in Van den Bergh et al. (2013). This is in particular true for airlines, with solution methods ranging from heuristics to column generation, see Gopalakrishnan and Johnson (2005) for an excellent overview and Clausen et al. (2010) for an overview of the recent methods of airline disruption management. The crew scheduling literature is very scarce in the area of sea transportation. However, OR techniques, in particular optimization, have

3

been used extensively in this sector, with a focus on routing and scheduling, see Christiansen et al. (2013). An overview to sea transportation is given thoroughly by Christiansen et al. (2007), which also notes that "in many aspects aircraft are similar to ships". Although this paper notes "crew scheduling for deep-sea vessels is not a major issue", our experience is on the contrary with our industrial partner, which operates OSV's providing construction support, remote operated vehicle operations, surveying, diving support, and deep-water lifting and installation. The literature related to OSV's also reveals a focus on routing or scheduling rather than crew, see, e.g., Fagerholt and Lindstad (2000), Gribkovskaia et al. (2008) and Halvorsen-Weare et al. (2012). Next, we discuss maritime problems that have some similarities to our setting, i.e., involving skill levels for different roles and long planning horizons.

One relevant study investigates the on-shore shipping operations, in particular the scheduling of employees at a marine container terminal (Legato and Monaco; 2004). However, in their case, the workforce demand is not known, and based on their skill levels, workers can be assigned one main task and additional secondary tasks. Another relevant study is the work of Li and Womer (2009) on the shipboard manpower scheduling, which considers crew to be multi-skilled as we do, but with the limitation that for any single duty period, they can be allocated only to a single role. Moreover, no consideration is given by Li and Womer (2009) to the issue of transportation, whereas our problem maintains a significant consideration for the cost of, and indeed time required to arrange, the transportation of a crew member to and from the vessel to which they are assigned. The short-sea crew scheduling problem of Wermus and Pope (1994) considers only small problems with regards to number of crew, and the lack of a mathematical formulation limits any further comparisons. Another relevant problem is the sailor assignment problem (Blanco and Hillery; 1994; Holder; 2005). One similarity to our problem is the long time scale for planning, however, each sailor goes through a rotation every three years, and will carry out their new assignment until the time of their next rotation. There is also no mention of tasks being changed once they are assigned. The integrated vehicle and crew scheduling problem from the Australian Navy (Horn et al.; 2007) has a fairly limited crew perspective, and the problem was still found to be too challenging to solve.

The work of Ammar et al. (2013) in crew scheduling for ferries relates closely to our problem. Similar to Horn et al. (2007), crew are allocated in teams rather than as individuals, and timescale considered is very short, with tasks in the order of hours, hence presenting

more similarities to airline or urban mass transit system settings. To the best of our knowledge, the most similar study to our problem in terms of problem description is the crew scheduling problem of Giachetti et al. (2013) in the cruise industry. The global characteristics of the company, consideration of transportation costs and international crew members have similarities with the features of our setting. However, cruise schedules are drawn up well in advance and not changed, and they consider only crew members with identical skills and experience. They propose a goal programming approach, and argue that different crew categories are independent and therefore can be solved separately. Their proposed solution approach involves a demand planning module, which takes past data and uses it to predict no shows and thereby estimate the overbooking requirement for each vessel. This is in contrast to our problem, where vessel requirements are considered to be known (but changeable), and hence a demand model is not used.

3.   Vessel Crew Scheduling Models for OSVs

In this section, we present a full description of the problem as observed by our industrial partner, which follows predictive-reactive scheduling with revisions made periodically on a weekly basis (and occasionally with event-driven changes). Because of the lengthy duty periods, and the need to arrange travel in advance for the crew, the company has a much longer planning horizon than other transportation settings, planning at least 13 weeks ahead, and arranging the logistics for the crew up to four weeks before the task. The length of the duty periods as well as the planning horizon lead to an accumulating uncertainty in planning, which can be as a result of changes in crew availability or in the vessels' assignments. As a result, the scheduling process follows a rolling horizon approach, i.e., using the current schedule for the coming thirteen weeks as a starting point and making any necessary changes. The planners of our industrial partner often attempt to make back-to-back scheduling, i.e., assigning crew on the same vessel and sharing their role with another employee in turn. However, deviations from back-to-backs are often required making these problems practically complex, where the planners need to identify a new possible schedule, contact the affected crew and confirm the changes, which might be refused by crew entitled to do so, and hence returning the planner to the first stage of identifying a schedule.

We will develop mathematical formulations in two steps: 1) While describing the problem in detail and defining notation, a first model constructing all schedules from scratch is built

considering cost of crew as objective, allowing us to set out the basics of the model. 2) We then extend this model to take into account the recovery-type nature of the real problem to minimize the cost of possible changes, and also discuss an alternative approach by minimizing number of changes instead.

3.1 Problem Description and Basic Scheduling Formulation

We have a wide set of roles which must be covered by crew members, and which can be broadly divided into two categories: 1) "Marine crew" include roles such as captain, bridge and engineering crew, and must be covered at all times, even when a vessel is unassigned or undergoing maintenance; 2) "Project crew" vary depending on a vessel's assignment, e.g. diving crew when a vessel is engaged in a diving project, and also include crew working on deck such as riggers and deck foremen, who will not be required when a vessel is unassigned or undergoing maintenance. Vessels require only one of certain role types, such as a captain, but may require multiple diving or engineering roles to be covered.

We consider two types of crew: regular crew, denoted by the set $\mathcal{F} = \{1, ..., m\}$, and agency crew, represented by $\{m+1\}$. We also define the set of all employees, denoted by $\mathcal{E} = \{1, ..., m + 1\}$. Regular crew are permanent employees of the company, and include a subset of fixed contract crew, denoted by $\mathcal{G}$. Fixed contract crew are paid a salary to work a certain number of days at sea per year, and using this contracted number of days (as well as the number of days worked so far in the year), the company estimates for each employee $i \in \mathcal{G}$ the number of days this employee is expected to work in a given planning period, which is denoted by G_i. This quantity is used in order to estimate the costs associated with overtime and undertime, as follows: if employee $i \in \mathcal{G}$ works more than G_i days in the planning period, each additional day will be charged by the overtime rate of Φ_i; on the other hand, if this employee works less than G_i days in the planning period, then this wastage of salary will be charged by the effective rate per day denoted by Υ_i. The remaining regular crew are the so-called day rate crew, who are paid per day at sea with a rate up to 50% higher than the rate of the fixed contract crew. Agency crew $\{m+1\}$ are outsourced by external agencies and are available at short notice, albeit more expensive as up to twice of day rate crew. In general, the company relies on regular employees, and uses agency crew only when absolutely necessary.

Crew availability for a specific role at a given time depends on a number of factors, including other commitments (e.g., training), unexpected unavailability (e.g., illness), their

training and experience level, their nationality, and legal and contractual requirements. For example, ships registered under certain flags can employ only crew from certain nationalities, costly visas may be required in certain regions, and there might be undesirable combinations of certain nationalities. On the other hand, each employee $i \in \mathcal{E}$ has a maximum duration they can be assigned to work, denoted by Ω_i, and a minimum duration of rest, denoted by P_i. Notwithstanding individual differences stemming from nationalities or contractual specifics, regular working patterns for the vast majority of crew is 4-on-4-off (4 weeks at sea followed by a 4 week rest), and the rest 5-on-5-off (often in remote regions) (European Union; 2003). We assume that the regular assignment pattern on each vessel is pre-determined, thus dividing each role into four- or five-week duty periods, or tasks. We assume there are n tasks in total, which is composed of the set of tasks which are to be carried out, denoted by $\mathcal{J}$, and the set of rest tasks, denoted by $\mathcal{N}$. It is possible that a task $j \in \mathcal{J} \cup \mathcal{N}$ can cover a number of consecutive roles, and each task is designed for a single employee, with a starting time S_j and duration D_j. Without loss of generality, we assume that the indices of all tasks in the set $\mathcal{J} \cup \mathcal{N}$ are ordered in the nondecreasing order of starting times, i.e., $S_{j-1} \leq S_j$. Crew changes take place each week on each vessel, affecting only a subset of the crew, and the cost of assigning employee $i \in \mathcal{E}$ to task $j \in \mathcal{J}$, denoted by C_{ij}, considers all penalties and financial costs (including day rates of day rate and agency crew).

Some of the projects run by our industrial partner require specialist knowledge, training or experience. We denote by $\mathcal{K}$ the set of these projects, and for any such project $k \in \mathcal{K}$, the company determines the minimum total experience required across the tasks in the project, denoted by H_k. In order to be able to calculate how much actual total experience is available in a project $k \in \mathcal{K}$ for a given allocation of crew, the company identifies the so-called experience score E_{ij} of employee $i \in \mathcal{E}$ for each task $j \in P_k$.

Before formally stating the mathematical formulation of the problem, we define next our decision variables. For each employee $i \in \mathcal{E}$ and task $j \in \mathcal{J} \cup \mathcal{N}$, the binary variable x_{ij} takes a value of 1 if employee i is allocated to task j, and 0 otherwise. For each fixed contract employee $i \in \mathcal{G}$, we define the continuous variables u_i and o_i to indicate the number of days under and over the guaranteed days expected in the planning period, respectively. Finally, for each employee $i \in \mathcal{E}$ and task $j \in \mathcal{J} \cup \mathcal{N}$, we define the variables w_{ij} and r_{ij} to represent the accumulated work (and rest, respectively) resource value for employee i once all tasks up to and including the task j have been considered, where w_{ij} is a continuous variable counting the number of consecutive working days whereas r_{ij} is a binary variable indicating

whether the employee needs a rest at this point or not. In line with these variables and the ordered set of $\mathcal{J} \bigcup \mathcal{N}$, we also define the parameters w_{i0} and r_{i0}, where the former indicates the number of consecutive working days employee i has accumulated immediately prior to the start of the planning period, and the latter indicates whether employee i requires rest at the start of the planning period or not. Finally, in order to facilitate the calculations of the accumulated work and rest resource value variables, we define the parameters W_j and R_j for each task $j \in \mathcal{J} \bigcup \mathcal{N}$, where the work resource value W_j is either set to the duration of the task j, i.e., $W_j = D_j$, if $j \in \mathcal{J}$, or to a sufficiently large negative number if j is a rest task, while the rest resource value R_j is either set to 1 if $j \in \mathcal{J}$, or to -1 otherwise.

We can now state the task-based formulation (denoted as TBF) as follows:

$$\min \sum_{i \in \mathcal{E}} \sum_{j \in \mathcal{J} \cup \mathcal{N}} C_{ij} x_{ij} + \sum_{i \in \mathcal{G}} (\Upsilon_i u_i + \Phi_i o_i) \tag{1}$$

$$\text{subject to:} \quad \sum_{i \in \mathcal{E}} x_{ij} = 1 \qquad \forall j \in \mathcal{J} \tag{2}$$

$$\sum_{j \in \mathcal{J} \cup \mathcal{N}} x_{ij} \leq 1 \qquad \forall i \in \mathcal{F} \tag{3}$$

$$\sum_{i \in \mathcal{E}} \sum_{j \in \mathcal{P}_k} E_{ij} x_{ij} \geq H_k \qquad \forall k \in \mathcal{K} \tag{4}$$

$$w_{i,j-1} + W_j x_{ij} \leq w_{ij} \qquad \forall j \in \mathcal{J} \cup \mathcal{N}, i \in \mathcal{F} \tag{5}$$

$$w_{ij} \leq \Omega_i \qquad \forall j \in \mathcal{J} \cup \mathcal{N}, i \in \mathcal{F} \tag{6}$$

$$r_{i,j-1} + R_j x_{ij} \leq r_{ij} \qquad \forall j \in \mathcal{J} \cup \mathcal{N}, i \in \mathcal{F} \tag{7}$$

$$u_i \geq G_i - \sum_{j \in \mathcal{J}} D_j x_{ij} \qquad \forall i \in \mathcal{G} \tag{8}$$

$$o_i \geq \sum_{j \in \mathcal{J}} D_j x_{ij} - G_i \qquad \forall i \in \mathcal{G} \tag{9}$$

$$x_{ij} \in \{0, 1\} \quad \forall i, j \text{ s.t. } x_{ij} \text{ is defined} \tag{10}$$

$$r_{ij} \in \{0, 1\}, \; w_{ij} \geq 0 \qquad \forall j \in \mathcal{J} \cup \mathcal{N}, i \in \mathcal{F} \tag{11}$$

$$u_i, o_i \geq 0 \qquad \forall i \in \mathcal{G} \tag{12}$$

The objective function (1) minimises the sum of the direct costs of assigning employees to tasks and the costs incurred due to guaranteed days of fixed contract crew. Constraints (2) ensure each task is covered, and constraints (3) prevent an employee to be assigned two overlapping tasks, where the agency crew (indexed $m + 1$) is not included since as many agency crew as needed are assumed to be available. The minimum experience required for

certain projects is ensured by constraints (4), and employee's consecutive working period lengths are calculated using constraints (5), with constraints (6) enabling maximum permitted working durations. Similarly, constraints (7) cover the minimum between-task rest period duration for employees, and finally, constraints (8) and (9) enable the calculation of under- and over-time for each fixed contract employee. We also note that in particular the constraints (4), (8) and (9) are specific to this maritime setting, whereas the remaining constraints can be found in various other transportation crew scheduling problems, expressed either implicitly or explicitly.

3.2 Extending to a Recovery-type Formulation

The nature of our particular maritime problem requires frequent changes on the existing schedule in light of new information about crew availabilities or vessel requirements. To accommodate this, we first define a binary $(m + 1) \times |\mathcal{J}|$ matrix $\mathbf{X}^*$ indicating assignments of the current schedule. Then, we define new binary decision variables y_{ij} to keep track of changes from the current schedule, which is 1 if there is a change to employee i's schedule with respect to task j, and 0 otherwise. For notational simplicity and preserving our previously defined formulation, we note that x_{ij} is now a dependent variable and can be written in terms y_{ij} and x_{ij}^*, where $x_{ij} = x_{ij}^* - y_{ij}$ if $x_{ij}^* = 1$ or $x_{ij} = x_{ij}^* + y_{ij}$ if $x_{ij}^* = 0$. This is equivalent to:

$$x_{ij} = x_{ij}^* + (1 - 2x_{ij}^*)y_{ij}; \forall i \in \mathcal{E}, \forall j \in \mathcal{J} \cup \mathcal{N} \tag{13}$$

As we focus now on changes made to an existing schedule, the cost of changing an employee's assignment to a task (including e.g. costs/savings with respect to wage and transportation) becomes relevant. Hence, we define two additional parameters, C'_{ij}, the cost (or saving if $C'_{ij} < 0$) of changing the assignment of employee i with respect to task j, and Ξ_i, the additional overtime/undertime costs for employee $i \in \mathcal{G}$ with respect to $\mathbf{X}^*$. Then, the objective function is as follows:

$$\min \sum_{i \in \mathcal{E}} \sum_{j \in \mathcal{J} \cup \mathcal{N}} C'_{ij} y_{ij} + \sum_{i \in \mathcal{G}} (\Upsilon_i u_i + \Phi_i o_i - \Xi_i) \tag{14}$$

Finally, we formally state our formulation, denoted as $RF1$, as follows: $\min_{y,r,w,u,o}$ $((14)|(y, r, w, u, o) \in X^{RF1})$ and $X^{RF1} = \{(y, r, w, u, o)| (2)-(9), (11)-(13), y \in \{0, 1\}^{|\mathcal{F}| \times |\mathcal{J}|}\}$.

3.3 An Alternative Recovery-Type Formulation

An alternative approach for recovery is to seek to minimise the number, rather than the cost, of the changes. As highlighted by our industrial partner, this is particularly relevant when a planner seeks a solution which is easy to implement and also requires minimal discussions with employees, as planning is often an iterative process between plan generation and human feedback. Therefore, we define the new objective function as follows:

$$\min \sum_{i \in \mathcal{E}} \sum_{j \in \mathcal{J} \cup \mathcal{N}} y_{ij} \tag{15}$$

In line with the practice, an upper bound Λ on the cost of making changes is also required:

$$\sum_{i \in \mathcal{E}} \sum_{j \in \mathcal{J} \cup \mathcal{N}} C'_{ij} y_{ij} + \sum_{i \in \mathcal{G}} \left(\Upsilon_i u_i + \Phi_i o_i - \Xi_i \right) \leq \Lambda \tag{16}$$

If a sensible upper limit is not known, one can set Λ to its maximum value, i.e., $\Lambda = \sum_{i \in \mathcal{E}} \sum_{j \in \mathcal{J} \cup \mathcal{N}} \left| C'_{ij} \right| + \sum_{i \in \mathcal{G}} \left(366 \left(|\Upsilon_i| + |\Phi_i| \right) + |\Xi_i| \right)$, and if necessary, iteratively reduce it.

Next, we note that inequalities (8) and (9) are now not sufficient to ensure that u_i and o_i take the correct values due to the new objective function. Therefore, we introduce two additional binary variables, θ_i^u, which is 1 if the undertime value for employee i is non-negative and 0 otherwise, and θ_i^o, which is 1 if the overtime value for employee i is non-negative and 0 otherwise. Then, we replace (8) with the following constraints:

$$u_i \leq G_i - \sum_{j \in \mathcal{J}} D_j x_{ij} + M \theta_i^u \quad \forall i \in \mathcal{G} \tag{17}$$

$$u_i \leq M \left(1 - \theta_i^u \right) \quad \forall i \in \mathcal{G} \tag{18}$$

and similarly, use the following constraints to replace (9):

$$o_i \geq \sum_{j \in \mathcal{J}} D_j x_{ij} - G_i + M \theta_i^o \quad \forall i \in \mathcal{G} \tag{19}$$

$$o_i \leq M \left(1 - \theta_i^o \right) \quad \forall i \in \mathcal{G} \tag{20}$$

Here, M is a suitably large number greater than the maximum possible value for the number of days of under- or over-time, and can be set to $M = 366$. Finally, to ensure u_i and o_i are not both set to non-zero for the same i, we add the constraint:

$$\theta_i^u + \theta_i^o \leq 1 \quad \forall i \in \mathcal{G} \tag{21}$$

Then, our alternative formulation, denoted as $RF2$, can be stated as $\min_{y,r,w,u,o,\theta^u,\theta^o} \left((15) | (y,r,w,u,o,\theta^u,\theta^o) \in X^{RF2} \right)$ where $X^{RF2} = \{(y,r,w,u,o,\theta^u,\theta^o) | \ (2) - (7), (11) - (13), (16) - (21), y \in \{0,1\}^{|\mathcal{F}| \times |\mathcal{J}|}, \theta^u \in \{0,1\}^{|\mathcal{G}|}, \theta^o \in \{0,1\}^{|\mathcal{G}|} \}$.

4. Computational Results

Although access to real problem instances was not possible due to data confidentiality, our industrial partner was able to provide us several specially identified key parameters and possible values or value ranges for these parameters. The company was also interested to investigate the effect of such parameters on solving these problems, and this process enabled us to generate randomized but realistic data sets covering all possible combinations. We used a full-factorial design with different number of levels involving the following four experimental factors identified by the company, which were varied across all instances: i) The probability that an employee is available on a given day, which is dependent on their availability the previous day (for which we define the parameters p (and q), i.e., the probability of an employee being unavailable on a certain day given they were unavailable (available) on the previous day); ii) Use of a probability reduction factor, $r(d)$, that increases uncertainty about the availability of an employee; iii) A disruption penalty K (which is split into K_N for **N**ear-term (i.e. 4 weeks), and K_L for **L**ong-term); iv) An agency penalty factor K_{AG}. The levels for each of these factors is listed in Table 1, and the interested reader can refer to VCS-data (2015) for the data of the resulting 240 instances. We note that these instances concern problems with ship captains only in consistency with the approach used by the company solving separate problems for each employee category, where the number of employees was set to 48 in line with the biggest problems of the company with 40 vessels and 13 weeks planning horizon. Moreover, in order to provide extensive computational results in Section 4.1, we generated 240 very large scale instances with the number of captains doubled to 96 by replicating the data of the original 240 instances.

Table 1: Levels used for the full factorial design

Factor	Set of levels
p	{0.2, 0.5, 0.8}
$r(d)$	{Use, Not Use}
(K_N, K_L)	{(1,1), (2,1) ,(2,2), (5,1), (5,2), (5,5), (10,1), (10,2), (10,5), (10,10)}
K_{AG}	{1,2,5,10}

As discussed in Section 3.1, the primary concern of the company is to find feasible solutions quickly. Hence, rather than investigating how long the optimal solution would take to find, the aim of the computational tests is to discover the number and quality of solutions

found within a practically acceptable time limit. Since the planners may run a model multiple times altering settings, it was agreed with the company to test two time limit settings, namely 2 and 10 minutes. In order to provide comprehensive computational results, we also solve the test instances that could not be solved in 10 minutes with a time limit of 1 hour. All test runs were carried out on a Dell Optiplex 790 PC (Windows 7 32-bit, Intel Core i5 3.10 Ghz, 4GB RAM), and the formulations $RF1$ and $RF2$ were implemented and tested using FICO® Xpress-MP (Mosel v3.6.0, Xpress-MP v7.7).

For the cost-minimization approach, our preliminary tests suggested changing the default settings to a maximum of 30 rounds of cover and 10 rounds of Gomory cuts, with an acceptable gap of 5% set as a cutoff. For change-minimization, preliminary tests indicated that when the upper bound on cost was relatively high, the problem could be solved very quickly, albeit often with a very high cost. Hence, a simple iterative algorithm that progressively lowers the upper bound on cost and re-solves the problem within total available time of 2 minutes is implemented to obtain a set of pareto-optimal solutions. We refer the interested reader to Algorithm 5.2 on p.75 of Leggate (2016) for the details of this algorithm.

4.1 Cost-minimization results

In order to test the cost minimization model, we solved 240 problem instances of the data set VCS-data (2015) with 48 captains as well as the extended 240 instances with 96 captains using 2-min, 10-min and 1 hour time limits. We note that a typical instance contains around 23 thousand rows and 17 thousand columns for problems with 48 captains, and 89 thousand rows and 67 thousand columns for problems with 96 captains. We present overall results in Table 2, where the number of test instances that were completed in 2-min/10-min/1-hour runs are noted in columns "Opt." (for optimal solutions) and columns "< 5%" (for instances that were stopped due to gap in a run being less than the acceptable gap of 5%). In addition, the columns "Ave." indicate the average gaps and the columns "Med." indicate the median gaps for all 240 instances in each set, where we calculated the gaps using the overall best bound for each instance, i.e., the bound obtained from the 1-hour run, rather than the bound obtained in each run with different time limits, in order to provide a clear comparison of solution qualities between different lengths of runs.

As Table 2 clearly indicates, the biggest instances the company deals with in practice, i.e., those with 48 captains, can be in general solved very effectively. Even using the 2-min time limit, only 13 out of 240 instances cannot be solved within 5% of optimality, and 2 of

Table 2: Cost minimization results for instances with 48 and 96 captains.

Time	48 captains				96 captains			
	Opt.	< 5%	Ave.	Med.	Opt.	< 5%	Ave.	Med.
2-min	106	121	2.47%	0.27%	20	107	33.04%	4.76%
10-min	106	123	2.42%	0.27%	21	128	9.81%	4.22%
1-hr	106	123	2.41%	0.27%	21	133	9.70%	4.22%

these instances can be brought within 5% of optimality using the 10-min time limit. This is certainly encouraging for the practical use of the approach, as only 11 out of 240 instances ($\approx 4.6\%$ of the total) do not achieve the desired optimality gap of 5%. On the other hand, extending the time limit to 1-hour does not provide any benefit for these instances, only with marginal improvement of gaps for the remaining 11 instances (and no further instance brought within 5% of optimality).

Next, we discuss the instances with 96 captains that provide further insights with regards to the computational capability of the models proposed for large-scale problems. Although these significantly bigger problems are naturally more challenging, resulting in average percentage gaps of 33.04%, 9.81% and 9.70% for 2-min, 10-min and 1-hour runs, respectively, the proposed models still remain practically effective, indicating further possibility for use even in case of significant growth of the company, such as in case of a merger. As results indicate, 149 out of 240 instances ($\approx 60.1\%$ of the total) can achieve the desired optimality gap of 5% in a 10-min run. Although extending the time limit to 1-hour helps to achieve the desired optimality gap for 5 more instances, the improvement of gaps remain marginal, similar to our previous experience with the 48 captain instances. In addition, the histograms in Figure 1 provide a breakdown of solutions with respect to gaps for all these tests. These figures present a better understanding of the improvement of gaps with longer computational times, though the improvements remain marginal, in particular for the case of 48 captains.

The results obtained so far also motivated us to further investigate challenging test instances with extensive computational runs. Therefore, we have applied a 6 hour time limit and executed the approach on the 5 instances with 48 captains that could not be solved to optimality in 1 hour and had a gap of 6% or higher, in order to ensure not to be too close to the 5% gap cutoff. As the results in Table 3 indicate, these most challenging 5 instances improved only marginally between 1-hour to 6-hour runs, the best performing instance (R205) having still a gap of 6.1%. Our preliminary tests with some of the unsolved 96 captain in-
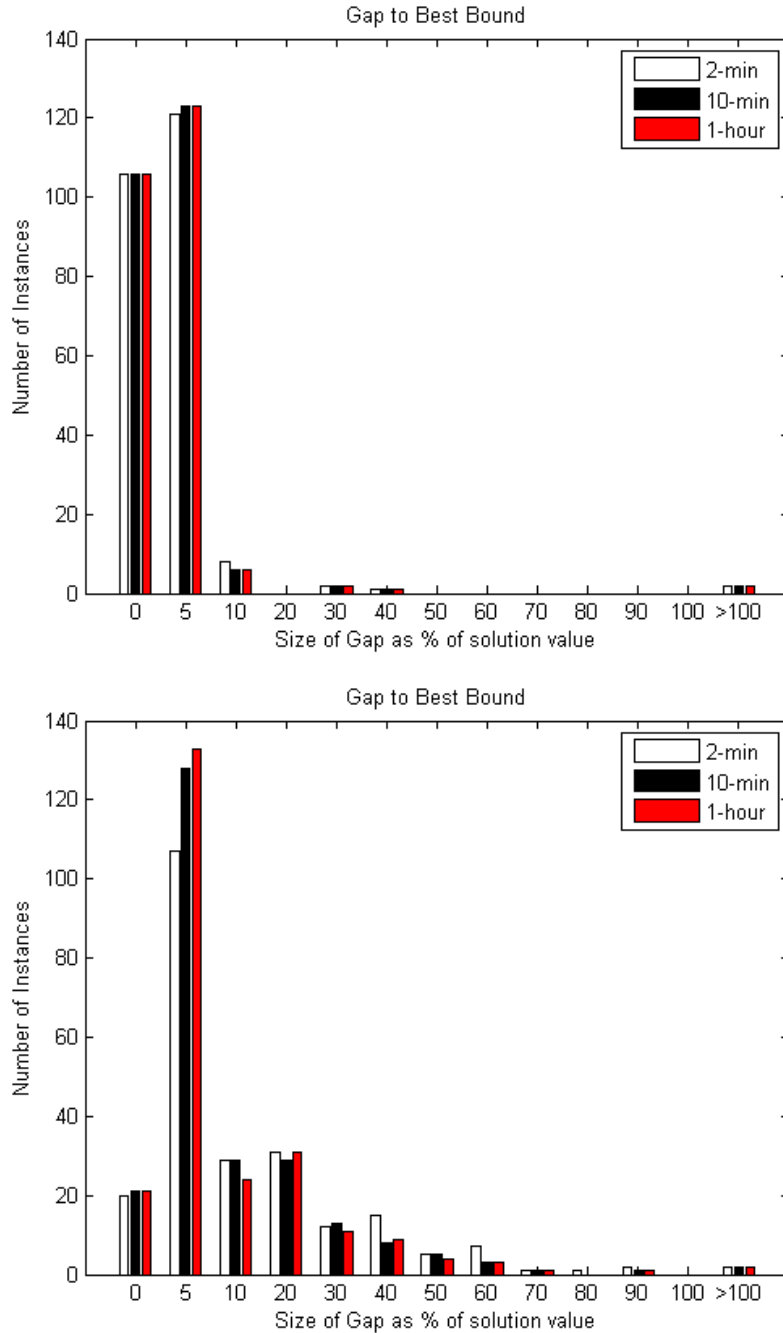
Figure 1: Gaps found in 2-min, 10-min and 1-hour cost-min runs for 48 captains (top) and 96 captains (bottom), based on best known bounds. x-axis indicates the integrality gaps, and y-axis indicates number of instances in each bracket.

stances indicated a similar pattern of very limited improvement over extended runs, though it is also worth to note that some of these large scale instances experienced memory issues.

Finally, we note that we will discuss in Section 4.3 the effect of parameter values, using

Table 3: 6-hour extensive runs for challenging instances with 48 captains.

Instance	Opt.(1-hour)	Opt.(6-hour)	Best Bound (6-hour)
R001	1827	1801	1029.25
R005	20.5	20.5	-50.515
R056	22620.8	22618.8	18822.11767
R081	-409	-409	-507.333
R205	-394	-396	-424.008

statistical analysis tools such as Kruskal-Wallis test.

4.2 Change-minimization results

As we approach change minimization in a different fashion, we focus on the two factors that matter: the cost of the solution found, and the number of changes it entails. Unlike cost-minimization, this approach generates multiple solutions along the pareto-optimal frontier. Figure 2 presents an example of the multiple solutions and the associated pareto-optimal frontier obtained for the specific instance R133, which in fact generated 34 solutions, highest number of solutions for all instances, within the allocated two minute time limit. It is noteworthy to remark that on average we found 13.3 solutions for each instance, and all instances found at least 5 solutions within the time limit, demonstrating that in every one of these cases a planner would have been presented with at least some degree of choice.

In terms of the lowest cost solution found for each instance, the iterative approach is very competitive: As Figure 3 (left) illustrates, for 198 instances (82.5% of the total) a solution within 5% of the best known bound was found; of these, 21 instances (8.75% of the total) were solved to optimality within the 2-min time limit. Moreover, when we compare these solutions with the solutions found by the 2-min cost-minimization (see Figure 3 (right)), we note that 208 instances ($\approx 86.67\%$ of the total) had a solution within 5% of the 2-min cost-minimization solution, and 64 ($\approx 26.67\%$ of the total) had an equal or slightly better solution (due to the cost-minimization method stopping when a near-optimal solution is found). We also note that, as expected, earlier solutions found in the change-minimization process are likely to have much greater gaps to the cost-minimization solution, e.g., the first (i.e. minimal-change) solution has a median gap of $\approx 398.4\%$ to the best known bound.

With respect to the number of changes, it is most insightful to compare this quantity between the solutions found by the cost-minimization method and the cheapest solutions
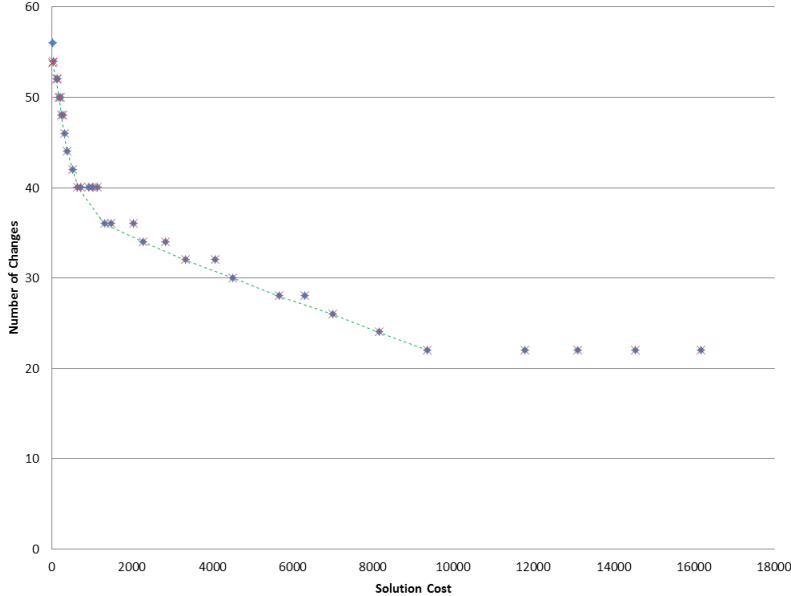
Figure 2: Solutions obtained by the change minimization approach for the test instance R133 and pareto-optimal frontier (highlighted with the dashed line).

(rather than the ones with least number of changes) found by change-minimization algorithm. As seen in Figure 4 (left), we see only 4 instances had fewer changes in the solution found by the cost-minimization and only 27 had the same number of changes, while the maximum difference appears when the change-minimization finds a solution with $\approx 59.6\%$ fewer changes. On the other hand, as Figure 4 (right) illustrates, the number of changes in the first solution of change minimization (i.e., the solution with the least number of changes) is between 36% and 89% worse than the solution found by the cost-minimization approach.

Finally, we also note that minimum number of changes required for all instances varied between 11 and 33 changes, with a mean of 20.95. The interested reader is referred for further detailed discussions and results on this to pp.82-92 of Leggate (2016).

Evaluating these results from a practical perspective, we conclude that in general it is more preferable for the planners to use the change-minimization approach. Although it will often not find the cheapest solution, it is capable of obtaining good quality solutions with regards to cost. Moreover, change minimization has the advantage of entailing less disruption (hence easier to implement in practice) and generates multiple solutions, while cost minimization can have a significantly bigger number of changes in comparison.
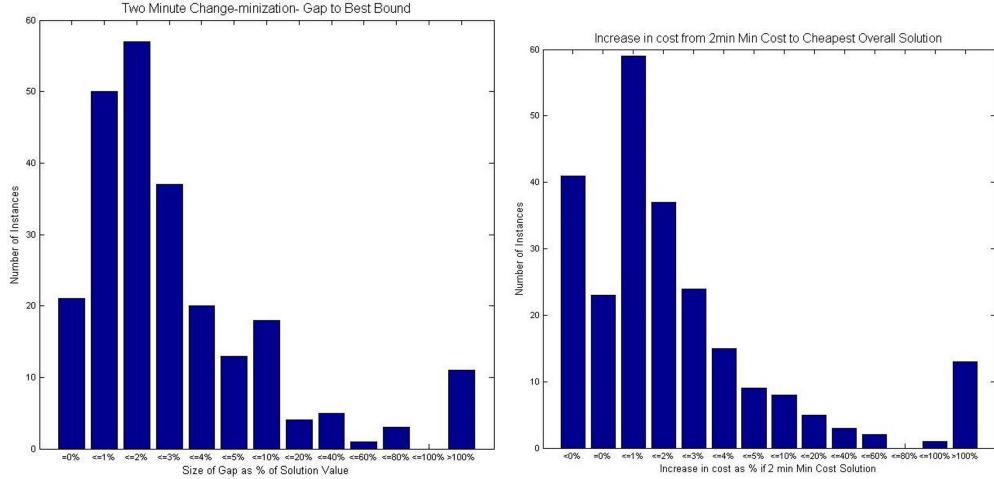
Figure 3: Gap to best bound for lowest cost solutions found in 2-min change-min run (left) and increase in cost from 2-min cost-min solution to the change-min solution (right).
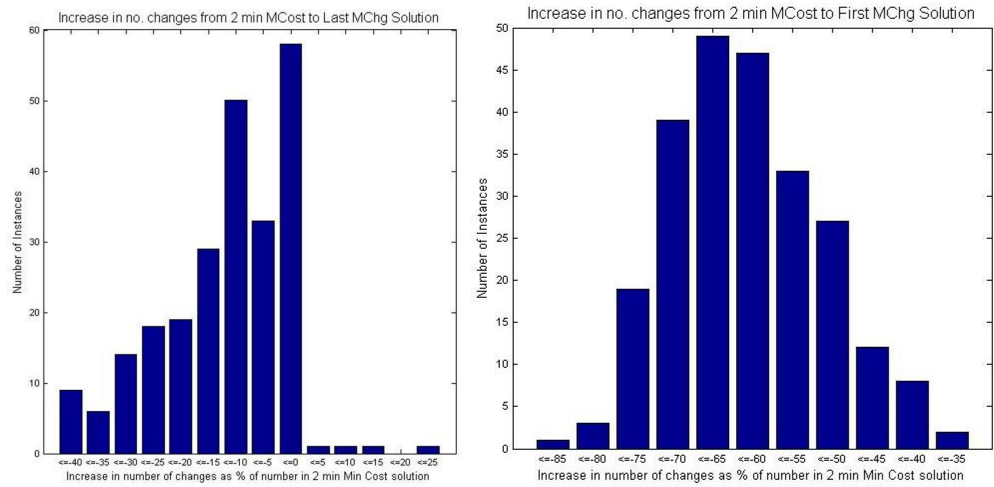
Figure 4: Comparison of: i) number of changes in cost-min solution vs. cheapest solution by change-min method (left), and ii) number of changes in cost-min solution vs. the first change-min solution (right)

4.3   Effect of parameter values on results

As discussed before, several parameters were varied for the generation of data instances. In this part, we investigate the effect of changing these parameter values using "Analysis of Variance". We used the ANOVA and F-test where possible; however, we note that for some of the parameters investigated, the required assumption of normal distribution did not necessarily hold and hence, the non-parametric equivalent "Kruskal-Wallis test" was carried

out instead. Table 4 shows the p-values for these tests, with respect to the varied parameters and the key output measures. We grouped the outputs in 3 main groups depending on the main objective functions that we run for experimental results. In Table 4 these 3- main groups and the depending performance measurements are given on the y-axis and the parameters are given on the x- axis. The type of applied tests is also stated in the x-axis of Table 4. The null hypothesis that we use for testing the parameter values effect is that there is no difference between different levels of a given parameter. Accordingly an asterisk next to a p-value indicates that, at the 5% significance level, there is a significant difference between the values of the given key output measure for instances generated using different values of the given parameter. For example, where the '# changes - 2mins (only)' row intersects with the 'p and q' column the value of $0.004*$ indicates that there is a significant difference between the number of changes in the min-cost solution between instances which were generated using different values of availability probabilities p and q. Those parameters which returned a significant result as shown in Table 4 required further investigation to determine in what way the change in parameter values was significant. We note that the quantity $\bar{K}$ is an average of the Near and Long disruption factors K_N and K_L, weighted according to the number of weeks to which each apply (i.e. 4 and 9 weeks in a 13-week planning period).

Further investigation allows us to make some observations about how the values of these parameters may influence the running of the models in practice. For example, it can be seen from the table that the factors relating to crew availability have a significant influence on certain outputs relating to changes. As might be expected, we can conclude that if in reality crew absences tend to be longer but less frequent, then we would expect to see fewer changes being found in the 2-min cost-minimization and in the change-minimization solutions, and a greater gap between the first change-minimization solution and the cost-minimization solution in terms of number of changes. In this case, we would also expect a smaller number of iterations to be required in the change-minimization run. A similar pattern might be observed if the time reduction of the absence probability proved to be a reasonable assumption in reality.

Meanwhile, the factors which relate to costs have a considerably wider influence. If the company wishes to apply disruption factors (i.e. penalties for making changes to the schedule), then in general we expect the effects to be: i) shorter run times in both approaches, ii) fewer iterations for change-minimization, iii) fewer changes in 2-min cost-minimization and in final change-minimization solutions, iv) a greater percentage decrease in changes

18

Table 4: Results of ANOVA and Kruskal-Wallis tests for influence of parameters used in instance generation

	Outputs of interest		K-W or ANOVA	p and q	use $r(d)$?	Parameter K_N	K_L	$\bar{K}$	K_{AG}
Cost-min settings	Run time	10mins	K-W	0.283	0.137	0.006*	0.000*	0.000*	0.201
		2mins	K-W	0.217	0.351	0.003*	0.000*	0.000*	0.008*
	# changes	2mins (only)	ANOVA	0.004*	0.000*	0.000*	0.000*	0.000*	0.960
	Gap found in each run	Actual; 10mins	K-W	0.115	0.383	0.168	0.000*	0.010*	0.560
		Actual; 2mins	K-W	0.701	0.430	0.024*	0.000*	0.000*	0.015*
		Adjusted; 10mins	K-W	0.458	0.576	0.063	0.074	0.213	0.208
		Adjusted; 2mins	K-W	0.538	0.825	0.432	0.011*	0.104	0.056
	Gap to best known bound	Actual; 10mins	K-W	0.158	0.152	0.227	0.001*	0.033*	0.767
		Actual; 2mins	K-W	0.341	0.231	0.169	0.000*	0.001*	0.106
		Adjusted; 10mins	K-W	0.434	0.636	0.049*	0.054	0.151	0.032*
		Adjusted; 2mins	K-W	0.592	0.975	0.158	0.027*	0.277	0.101
Change-min settings	Run details	Running time	K-W	0.101	0.528	0.045*	0.001*	0.023*	0.000*
		Iterations	ANOVA	0.001*	0.000*	0.003*	0.125	0.037*	0.000*
	# of changes	First sol found	ANOVA	0.000*	0.002*	0.515	0.808	0.941	0.780
		Last sol found	ANOVA	0.002*	0.002*	0.003*	0.007*	0.010*	0.561
	Cost of soln.	First min-chng sol	ANOVA	0.280	0.393	0.001*	0.012*	0.005*	0.000*
		Chpst min-chng sol.	ANOVA	0.212	0.198	0.002*	0.029*	0.011*	0.000*
	% gap to best known bound	Chpst ovrll	K-W	0.187	0.481	0.010*	0.000*	0.000*	0.000*
		Adjusted chpst ovrll	K-W	0.446	0.846	0.308	0.001*	0.018*	0.000*
Comparison with 2-min min-cost	% diff in # of changes	First sol found	ANOVA	0.000*	0.002*	0.008*	0.000*	0.000*	0.868
		Last sol found	K-W	0.576	0.830	0.000*	0.000*	0.000*	0.413
	% diff in cost	First min-chng sol	K-W	0.439	0.074	0.002*	0.093	0.024*	0.000*
		Chpst min-chng sol	K-W	0.013*	0.122	0.008*	0.064	0.032*	0.000*
		Chpst ovrll	K-W	0.413	0.676	0.181	0.014*	0.168	0.000*

for both first and final change-minimization solutions, v) smaller cost gaps both for the individual cost- and change-minimization approaches and when comparing the two. The additional penalty for using agency crew has a similar influence with respect to running time, iterations and gaps on cost, but has no significant effects on the number of changes. It should be noted that when the agency penalty is not applied, particularly long run times and large percentage gaps on cost are observed.

5. Conclusions and Future Work

In this paper, we have presented a vessel crew scheduling problem pertaining to a large global company operating Offshore Supply Vessels (OSVs) providing services particularly to the oil and gas industry. Having noted that there is an extensive literature of crew scheduling in other transportation settings as well as of other maritime OR problems, very limited literature in our specific problem area appear to exist. Discussions with our industrial partner have also allowed us to present formulations that could serve as part of a decision support tool, allowing the planners to find better quality solutions in real time. The current practice for crew scheduling at our industrial partner involves manual processes carried out on various spreadsheets. Because of the difficulties involved in the process, and the lack of an automated tool to aid decision making, the primary concern is to obtain feasible solutions. Even this however can be difficult, in particular when there is time pressure. Hence, our optimization approach effectively provides feasible, low-cost crew schedules when necessary.

A further research step is to carry out the implementation of a decision support tool to be used by our industrial partner, which the models developed in this paper will be a significant component of. On the other hand, we plan to address some limitations of the model presented here in the future. One important aspect of the current model is the implicit assumption that the time period over which a role requires to be covered can be divided into four- or five-week blocks. In the practical setting, employees could potentially be asked to work more irregular patterns even if not desirable. Although our current model can be solved quickly, a more sophisticated model will increase the flexibility for the planners as well as the potential quality of schedules generated, albeit with a need for more customized solution methodologies in order to obtain solutions in real time. This is part of our ongoing work (Sucu et al.; 2017), where heuristics and Benders decomposition are investigated for further insights.

References

Ammar, M., Benaissa, M. and Chabchoub, H. (2013). Grasp for seafaring staff scheduling: Real case., *in* M. Abed and M. Benaissa (eds), *2013 International Conference on Advanced Logistics and Transport*, IEEE, New York, pp. 427–433.

Barrett, D. (2008). The offshore supply boat sector, *Marine and Commerce* pp. 36–41.

Blanco, T. and Hillery, R. (1994). A sea story - implementing the navy personnel assignment system., *Operations Research* **42**(5): 814–822.

Christiansen, M., Fagerholt, K., Nygreen, B. and Ronen, D. (2007). Maritime transportation., *in* C. Barnhart and G. Laporte (eds), *Handbook in OR & MS*, Vol. 4, Elsevier B.V., pp. 189–284.

Christiansen, M., Fagerholt, K., Nygreen, B. and Ronen, D. (2013). Ship routing and scheduling in the new millennium., *European Journal of Operational Research* **228**(3): 467–483.

Clausen, J., Larsen, A., Larsen, J. and Rezanova, N. (2010). Disruption management in the airline industry - concepts, models and methods., *Computers & Operations Research* **37**(5): 809–821.

Ernst, A., Jiang, H., Krishnamoorthy, M., Owens, B. and Sier, D. (2004). An annotated bibliography of personnel scheduling and rostering., *Annals of Operations Research* **127**(1-4): 21–144.

Ernst, A., Jiang, H., Krishnamoorthy, M. and Sier, D. (2004). Staff scheduling and rostering: A review of applications, methods and models., *European Journal of Operational Research* **153**: 3–27.

European Union (2003). Directive 2003/88/EC of the European Parliament and of the Council of 4 November 2003 concerning certain aspects of the organisation of working time, *Official Journal of the European Union* **L299**(46): 9–19.

Fagerholt, K. and Lindstad, H. (2000). Optimal policies for maintaining a supply service in the norwegian sea., *Omega* **28**(3): 269–275.

Giachetti, R., Damodaran, P., Mestry, S. and Prada, C. (2013). Optimization-based decision support system for crew scheduling in the cruise industry., *Computers & Industrial Engineering* **64**(1): 500–510.

Gopalakrishnan, B. and Johnson, E. (2005). Airline crew scheduling: State-of-the-art., *Annals of Operations Research* **140**(1): 305–337.

Gribkovskaia, I., Laporte, G. and Shlopak, A. (2008). A tabu search heuristic for a routing problem arising in servicing of offshore oil and gas platforms., *Journal of the Operational Research Society* **59**(11): 1449–1459.

Halvorsen-Weare, E., Fagerholt, K., Nonås, L. and Asbjørnslett, B. (2012). Optimal fleet composition and periodic routing of offshore supply vessels., *European Journal of Operational Research* **223**(2): 508–517.

Holder, A. (2005). Navy personnel planning and the optimal partition., *Operations Research* **53**(1): 77–89.

Horn, M., Jiang, H. and Kilby, P. (2007). Scheduling patrol boats and crews for the Royal Australian Navy., *Journal of the Operational Research Society* **58**(10): 1284–1293.

Huisman, D. and Wagelmans, A. (2006). A solution approach for dynamic vehicle and crew scheduling, *European Journal of Operational Research* **172**(2): 453–471.

Legato, P. and Monaco, M. (2004). Human resources management at a marine container terminal., *European Journal of Operational Research* **156**(3): 769–781.

Leggate, A. (2016). *A vessel crew scheduling problem: formulations and solution methods*, PhD thesis, Dept. of Management Science, University of Strathclyde. Available online via `http://ethos.bl.uk/`.

Li, H. and Womer, K. (2009). A decomposition approach for shipboard manpower scheduling., *Military Operations Research* **14**(3): 67–90.

Pinedo, M. L. (2009). *Planning, Scheduling, and Timetabling in Transportation*, Springer New York, pp. 255–290.

Rezanova, N. and Ryan, D. (2010). The train driver recovery problem - a set partitioning based model and solution method., *Computers & Operations Research* **37**: 845–856.

Stopford, M. (2009). *Maritime Economics*, 3rd edn, Routledge, New York.

Sucu, S., Leggate, A., Akartunalı, K. and van der Meer, R. (2017). Customized heuristics and Benders decomposition for the vessel crew scheduling problem. Working paper, University of Strathclyde.

Van den Bergh, J., Beliën, J., Bruecker, P. D., Demeulemeester, E. and Boeck, L. D. (2013). Personnel scheduling: A literature review., *European Journal of Operational Research* **226**(3): 367–385.

VCS-data (2015). Vessel crew scheduling for off-shore supply vessels: A library of test instances. `http://dx.doi.org/10.15129/5dddf224-deaa-4b4b-aeb2-ef60c0d910c9`.

Wermus, M. and Pope, J. (1994). Scheduling harbor pilots., *Interfaces* **24**(2): 44–52.