

On Improving the Maritime Transshipment

Operations of the Noble Group

The Noble Group is a market-leading global supply chain manager of agricultural products, metals and minerals, operating in more than 140 locations. This paper focuses on Noble's maritime operations in Indonesia, where coal is transported from mines to ocean-bound vessels via roads and rivers. Currently, transportation delays are causing Noble to lose tens of millions of dollars per year in demurrage and detention penalties. Additional resources such as barges and floating cranes can be hired in advance to minimize the impact of delays, but their economic benefit is often unclear. In an attempt to reduce or eliminate these delays, we develop a modeling framework and decision support system to facilitate the planning and management of Noble's transshipment operations. The system utilizes fast search algorithms that deliver efficient schedules, minimizing the cost of delays and additional resources required, resulting in monthly savings exceeding \$1 million.

Keywords: scheduling; maritime; transshipment operations; decision support system; heuristics; decomposition.

History: submitted June 2014; revised versions submitted January 2015, May 2015, August 2015, October 2015; November 2015

The Noble Group (Noble) is a market-leading global supply chain manager of agricultural products, metals, minerals and ores, which focuses on transportation links between low-cost producing countries and high-demand growth markets. With a gross revenue of \$98 billion in 2013, and the energy sector accounting for more than 60% of that revenue, the coal maritime logistics operations lie at the core of Noble's business activities.

Energy coal, which is used for power generation, is one of Noble's most important traded commodities. Noble acts as an intermediary, managing transportation of coal from diverse coal-mining supply sources to ocean vessels. Indonesia, the top exporter of energy coal globally (UNCTAD 2013), serves as an upstream supplier to high-growth and developed markets, such as Korea, Japan, India and China. The majority of the mines in Indonesia are located in Borneo, which has two major trading ports: Taboneo, located on the south, and Muara Kaman, located on the east of the island, respectively. Noble manages the first part of the supply chain, transporting the coal from the mines to large ocean-bound vessels, operating on behalf of its customers. Because the mines are located in areas not easily accessible by truck or rail, and because transferring millions of tonnes of coal using trucks is neither cost efficient nor environmentally friendly, Noble uses barges to transport the coal from jetties that are located close to the mines, to ocean vessels in the two major ports (Figure 1).

Using river transport links and barges is common practice in coal logistics. To achieve economies of scale, barges carry large quantities of coal, usually between 3,000 to 10,000 tonnes. Figure 2 shows the image of a barge and tugboat.




Figure 1: Noble transports coal from mines in Borneo to river jetties, where it is loaded on barges, which move it to ports, where it is transferred onto ocean-bound vessels.



Figure 2: Coal transport operations utilize barges that carry up to 10,000 tonnes of coal.

Own, Lease or Hire Barges?

Noble owns a large fleet of barges of varying sizes. It also has long-term contracts with barge owners, who lease barges at a pre-agreed price. Typically, such contracts specify a maximum number of barges of each size guaranteed per shipment, and allow a period of seven days for the entire barge voyage. If Noble requires more barges than what the contract allows, it can hire barges on a spot basis, at a market price that is typically higher than that of leased barges. For leased and spot barges, an additional daily fee, called a *detention penalty*, applies when they are

used beyond a time window of 7 days. Table 1 shows the cost details, the potential penalties and the operational restrictions of each barge type.

Barge type	Cost structure	Potential penalties	Operational restrictions
<i>Owned</i>	Fuel and other variable costs	None	None
<i>Leased</i>	Price per tonne, that depends on barge size and starting location	After seven days, an additional daily fee is charged	A maximum number of barges of each size are available per shipment
<i>Spot market</i>	Price per tonne, that depends on barge size and starting location	After seven days, an additional daily fee is charged	Available for one trip only

Table 1: Each barge type is characterized by a specific cost structure, potential penalties and operational constraints.

Barge Logistics

As soon as they are directed to start a voyage, barges sail from a Noble-owned hub, located centrally on several of the major river paths, to the river jetty locations of coal suppliers. Each supplier gives Noble a monthly schedule that indicates which time windows can be used for loading. However, even when Noble barges arrive within the specified window, delays may occur, either because the jetties are busy serving other customers, or because of coal shortages. Loading times typically vary between half a day and two days, depending on the quantity to be loaded, the type of coal and the quality of the loading infrastructure. After loading is complete, barges wait for the clearing of transfer documents, which typically require one working day, and can therefore span four days because of weekends. After the documents are cleared, barges sail to one of two ports (south or east), where the coal is discharged onto large ocean vessels. Vessels come in various sizes, carrying between 20,000 and 120,000 tonnes, and may require between 3 and 16 barges. Smaller vessels are typically equipped with on-board cranes that allow simultaneous

loading from both sides. For loading larger vessels, an additional vessel mounted with floating cranes is necessary. Although floating cranes can only process one barge at a time, they are faster than on-board cranes and can discharge up to three large barges per day compared to only one for on-board cranes. Before discharging begins, however, barges may have to wait because (i) the vessel may not have arrived yet; (ii) other barges are being discharged onto the vessel; or (iii) a floating crane, sometimes required to discharge the coal, may not be available. When discharging is complete, barges return to the hub, refuel and wait for their next voyage. Figure 3 shows the various stages of a barge voyage, and Figure 4 shows discharging operations with on-board and floating cranes.




Figure 3: Barges rotate from Noble’s hub to a supplier jetty, where they are loaded, to the port, where they are discharged onto a client vessel, and back to the hub. Waiting for available loading and discharging resources can happen at every stage.



Figure 4: Barges are discharged onto ocean vessels using on-board (left) or floating cranes (right).

Demurrage and Despatch

Overall, a voyage (from hub to hub) can require anywhere between 3 and 10 days, and therefore a barge voyage needs to start well in advance of the anticipated arrival of the customer vessel in order to be able to start discharging on time. Each vessel has an *estimated time of arrival* (ETA, see Figure 5), mutually agreed upon by Noble and the customer at least two weeks in advance. Also mutually agreed upon is the *laytime*, a time window that starts with the vessel's arrival, after which the vessel has to be fully loaded. When the vessel loading time exceeds the laytime, Noble pays a daily penalty, called *demurrage*, which can be as high as \$50,000 per vessel per day. Delays of five or more days per vessel are not uncommon, resulting in demurrage penalties that have reached over \$10 million per year. On rare occasions, loading finishes before the end of the laytime, resulting in *despatch*, a bonus for Noble, typically at half the rate of demurrage.




Figure 5: The estimated time of arrival (ETA) of a vessel denotes the start of the laytime, i.e., the maximum time window in which the vessel has to be loaded. Overruns result in demurrage penalties (left), and early completion in despatch bonuses (right).

Maritime and Barging Problems

Maritime problems similar to the one we consider appear sparingly in the operations research literature, but at a steadily increasing rate during the last few years. Christiansen et al. (2007) give a comprehensive review of advances in maritime transportation modeling problems, while Vacca et al. (2010) give an overview of the berth allocation and crane assignment problems. To the best of our knowledge, the first paper relevant to modeling barge transportation is that of O'Brien and Crane (1959), who use a simulation model to determine the allocation of tug boats and determine the optimal number of barge loads on the Ohio and Mississippi rivers. Schwartz (1968) describes a transshipment-scheduling model that minimizes barge fleet costs, but he notes that his model is so complex that it could not be used given the current technology, despite some significant simplifications, such as infinite fleet capacity and a homogenous fleet. Most practical applications, such as the studies of Richetta and Larson (1997) and Taylor et al. (2005) use simulation to capture intricacies of each problem at hand. The bulk of the maritime literature has focused on long-haul transportation problems, such as maritime inventory routing (Persson and

Gothe-Lundgren 2005, Al-Khayal and Hwang 2007, Furman et al. 2011), with the exception of Agra et al. (2013) that study a short-sea transshipment problem.

This paper describes Noble's scheduling problem and the implementation of a decision support system adopted by Noble to schedule their barge operations. Although we also discuss a mathematical programming formulation and solution algorithms that form the basis of the decision support system, an in-depth investigation of exact solution techniques and analytical properties of the optimal solution is the subject of ongoing research. It is noteworthy that our approach has similarities with algorithms found in process scheduling (Floudas and Lin, 2005), as the problem of sequencing barge voyages within a single vessel so that the loading completion time is minimized is a generalization of a two-stage hybrid flow shop problem (Johnson 1954, Ruiz and Vasquez Rodriguez, 2010).

Framing the Problem

From an economic perspective, Noble's objective is to minimize the joint cost of barges and demurrage penalties, thus striking a balance between hiring leased and spot barges and avoiding demurrage penalties that result from late cargo deliveries. Determining the optimal trade-off requires three interconnected decisions, namely (i) how many owned, leased and spot barges to allocate to each customer vessel; (ii) when to dispatch each barge; and (iii) whether or not to hire a floating crane. Clearly, leasing additional barges can reduce waiting times and demurrage, but this comes at a price. Dispatching barges early can also reduce demurrage, but will result in the barge being tied up for longer, resulting in additional daily fees and in a need for more barges (instead of re-using barges on shorter voyages). Finally, hiring a floating crane can reduce demurrage by speeding up the discharging of coal, but this also comes at a price.

The Legacy Decision Making Process

Making barge hiring and scheduling decisions is a complex process because of the interactions between the barge voyages, the scheduled vessel arrivals, the availability of resources and the propagation of delays throughout the schedule. Nevertheless, Noble used a cumbersome manual scheduling procedure, which the logistics managers had to perform multiple times per day. The manual procedure also made it impossible to take into account the complex trade-offs between delays and the cost of additional resources, due to the complex interactions between the operations and the uncertainty affecting the arrival of client vessels and the availability of jetties and floating cranes.

Often, additional resources were quickly put in place whenever an unexpected event threatened to upset a schedule, with only rough back-of-the-envelope calculations of anticipated benefits versus cost. Frequent use was made of rules of thumb, such as “allocate spot barges to suppliers in locations with low spot unit cost” or “always use leased barges for a predefined set of suppliers, namely for which the lease barge provider’s location is nearby.” Although such rules make sense intuitively, they do not take into account the complex interactions and propagation of delays. Sometimes more complex rules of thumb were used, such as allocating a spot barge to a shipment only if the estimated marginal reduction of demurrage penalties outweighs the marginal spot barge cost. Although this back-of-the-envelope calculation is optimal for a stylized situation with only one shipment, it is myopic in nature, and tends to underestimate the benefit of hiring additional barges, which can prevent propagation of delays, especially when several vessels arrive in close time intervals. Sometimes, however, the rules-of-thumb were not only myopic, but plainly incorrect, for instance, by including fixed overheads and sunk costs. Noble often chose to hire a barge rather than use its own, as the hiring cost was deemed lower

than the cost of operating its own barge. The latter, however, often included fixed costs that would be incurred regardless of whether or not the barge was used. When managers noticed that, as a result, their own barges were often idle, they enforced guidelines in the form of a minimum number (four) of voyages that each Noble-owned barge should make per month. The logistics managers viewed these rules and guidelines as confusing and often contradictory, resulting in inefficiencies.

Also, the existing manual system lacked crucial information, such as the current state of operations, e.g., the location of barges, and cost information, which were only recorded ex-post, due to the fact that no proper cost estimation system was in place. Additionally, much of the required information, such as updates on supplier availability, prices of spot barges, the cost of fuel, and the availability of floating cranes, was often held and maintained by other Noble divisions. This data was not always accessible by the logistics managers in a timely manner. The lack of up-to-date information on the state of the system required managers to frequently call operators, asking for the location and state of each barge, the loading progress of each vessel and supplier cargo and jetty slot availabilities. Then they had to estimate, given the current start time of each barge, when it would be next available, and to allocate it to a new voyage, while taking into account each supplier's availability and the interactions with other barges. This procedure requires considerable cognitive effort, and, without a proper decision support system, can generate wildly optimistic estimates, due to underestimating the time of each operation, ignoring the cumulative effects of delays, or failing to incorporate the availability of resources.

The Barge Rotation System

Data Input

Our system, hereafter called the *barge rotation system*, integrates large amounts of information in order to optimize the barge allocation and scheduling process, and provides the logistics managers with an Excel-based graphical user interface. Hard-to-find and incorrect data, combined with frequent and time-consuming updates, rendered the existing decision making process cumbersome and ineffective. Therefore, we integrated all the required information into one spreadsheet model, with data located in different sheets depending on the frequency with which they are updated. For example, supplier locations, which are not updated very often, and available jetty slots and vessel ETAs, which are updated on a daily basis, are located in different sheets. This makes the system more ergonomic and greatly facilitates the data entry task. Figures 6 through 8 show data entry tables containing information on supplier locations (updated monthly), shipments (daily), and barges (daily).

Logistics managers typically complete the data entry process in a few minutes, and notification messages, in the form of pop-up boxes, are used to cross-validate data consistency. This is an important feature that was missing from the existing manual system, in which data entry and scheduling could take up to half a day, without any data validation.

Location	COA (\$/Ton) 300ft	COA Detention (\$/Day) 300ft	COA (\$/Ton) 270ft	COA Detention (\$/Day) 270ft	COA (\$/Ton) 230ft	COA Detention (\$/Day) 230ft
DONDANG		1750	4.18	1600		1250
SANGA SANGA	3.4	1750	4.18	1600		1250
BALIKPAPAN	6	1750	6	1600		1250
LOA KULU	3.57	1750	4.3	1600		1250
SEPARI	3.82	1750	4.3	1600		1250
SEBULU	3.92	1750	6	1600		1250
M KAMAN	4.2	1750	6	1600		1250
MANAU		1750		1600	8.22	1250
ULAK	5.8	1750	6	1600	8.22	1250
TANJUNG REDEB		1750		1600		1250

Figure 6: The location specifies the cost per tonne and detention penalty (for utilizing a barge for more than 7 days) for leased barges of various sizes for each supplier (location).

Customer Names	Vessel Names	Vessel Status	Estimated Time of Arrival (ETA)	Contracted Loading Rate	Expected Loading Rate	Demurrage	Vessel_Type
Customer 1	MV Vessel 1	Vessel in Service	14/02/2014	12,000	12,000	9,000	GG 2 Barges
Customer 2	MV Vessel 2	Vessel in Service	20/02/2014	15,000	20,000	14,000	FC/FT 1 Barge
Customer 3	MV Vessel 3	Vessel in Service	19/02/2014	15,000	20,000	20,000	Spot FC / GG 1
Customer 4	MV Vessel 4	Vessel in Service	26/02/2014	15,000	20,000	12,500	FC/FT 1 Barge
Customer 5	MV Vessel 5	Laycan Agreed	02/03/2014	15,000	20,000	13,000	FC/FT 1 Barge
Customer 6	MV Vessel 6	Vessel in Service	25/02/2014	15,000	20,000	15,000	FC/FT 1 Barge

Customer Names	Stowage Plan (MT)	Supplier 1	Supplier 2	Supplier 3	Supplier 4	Supplier 5	Supplier 6	Supplier 7	Supplier 8
Customer 1	110,000	55,000							55,000
Customer 2	66,600		10,400		15,400				
Customer 3	70,400				22,500		5,200		
Customer 4	135,000	45,000				22,500			45,000
Customer 5	150,000	52,500				22,500			52,500
Customer 6	101,400		15,000			23,000	5,200		

Figure 7: The shipment list includes customer data (top) and quantities sourced from suppliers (bottom). A pre-agreed contracted loading rate (tonnes/day) and the stowage plan determine the lay-time, whereas the expected loading rate depends on the loading infrastructure.

Barges	Barge Type	Size	Current Vessel	Current Jetty	Current Operation
Barge 1	BB	270 ft	MV Vessel 1	KBM	Discharging at MV
Barge 2	BB	270 ft	MV Vessel 1	BMSA	Waits to discharge at Muara Berau
Barge 3	Spot	300 ft	MV Vessel 1	ORBIS	Waits to discharge at Muara Berau
Barge 4	Spot	300 ft	MV Vessel 1	ORBIS	Awaiting Loading
Barge 5	BB	270 ft	MV Vessel 1	KBM	Loading
Barge 6	BB	270 ft	MV Vessel 1	PEI - ALH	Waiting for Documents
Barge 7	BB	270 ft	MV Vessel 1	PEI - ALH	Processing Documents
Barge 8	CoA	300 ft	MV Vessel 2	BJA	Sailing from Jetty to Muara Berau
Barge 9	BB	300 ft	MV Vessel 2	BJA	Waits to discharge at Muara Berau

Figure 8: The *barge list* describes the current state of each barge. Some data has been anonymized for confidentiality purposes.

Modeling and Algorithms

After the data entry phase is completed, the barge rotation algorithm is invoked. Ideally, the algorithm should incorporate uncertainties that affect the schedule, such as vessel arrival dates, loading times and supplier availability. However, data about uncertainties was not readily available, and managers were not comfortable with assigning probabilities to uncertain events. Also, a stochastic version of our system would be computationally intractable. Therefore, we decided to build a reactive deterministic model, in the sense that frequent rescheduling takes place, often multiple times per day, in order to incorporate unforeseen changes and new information. A reactive system works well in practice, as uncertainty in the short term is not high, and therefore current decisions are not significantly affected by uncertainty, and longer-term decisions need not be made until most uncertainty is resolved. Nevertheless, to create some protection against longer-term uncertainty, we have also inflated some nominal operation times, such as loading and sailing times, based on feedback from the operators, in order to create buffers. A detailed mathematical programming formulation of the voyage allocation problem is available in the appendix. Next we provide a general overview of the algorithm.

The Barge Rotation algorithm: voyage allocation, scheduling and improvement

The barge rotation algorithm decomposes the problem into a voyage *allocation* part and a voyage *scheduling* part, which are invoked initially to generate a feasible schedule, and then are called iteratively in a *voyage improvement* heuristic, i.e., a local search procedure that modifies the initial allocation decisions. Figure 9 shows the main blocks of the algorithm.

The initial voyage allocation algorithm (block *I*) determines the number of voyages of each barge type for each supplier and vessel that minimizes the transportation costs, ensuring that (i) the quantity that must be sourced from each supplier is covered, and (ii) the number of voyages allocated to each vessel does not exceed the maximum number of barges available of each type and size, including leased and spot barges. Note that using leased or spot barges may be cheaper than using owned barges, as leased or spot barges are often larger and can combine shipments that otherwise would require several owned barges.

Next, the voyage scheduling algorithm (block *II*) creates a feasible schedule for each vessel, while adhering to the voyage allocation decisions made in block *I*. Vessels are scheduled in order of non-increasing demurrage penalties, and the schedule for each vessel takes into consideration restrictions on the availability of barges, floating cranes and jetties imposed by vessels already scheduled.




Figure 9: The barge rotation algorithm is comprised of a voyage allocation (block V) and a voyage scheduling part (block VI). Blocks III-IX iterate between the two parts, in order to generate improved feasible solutions.

Once a feasible schedule is determined for each vessel, the voyage improvement algorithm initiates (blocks *III* to *IX*). This procedure is necessary because the initial allocation of voyages tends to over-utilize owned barges, which are typically less expensive, but which might create excessive delays resulting in high demurrage. The key idea of the voyage improvement algorithm is to identify the vessel with the highest demurrage (block *III*) and check whether substituting owned barges with leased or spot barges would lead to a lower total cost (blocks *IV* – *VIII*). We check for substituting sets of barges in addition to one-for-one swaps. In particular, block (*IV*) records the number of owned barges of each size that are to be substituted in each iteration, which is then implemented in block (*V*) that re-optimizes the voyage allocation in a way similar to block (*I*). Every time an improved schedule is found, it is stored (block *VII*), and the number of owned barges that are to be substituted is increased by one, until all owned barges are substituted (block *IX*). When all owned barges of a vessel have been considered for substitution, then that vessel is not considered again, even if it still has the highest demurrage (blocks *VIII* and *III*). The algorithm terminates when all owned barge substitutions in all vessels have been considered.

Voyage Scheduling

The voyage scheduling section is the backbone of the barge rotation algorithm, and the quality of the generated schedules relies heavily on its efficacy. Since, for a given allocation, the barge transportation costs are fixed, the voyage scheduling algorithm aims to minimize any penalties due to delays, i.e., demurrage and detention. In terms of the mathematical programming formulation in the appendix, the voyage scheduling algorithm minimizes (heuristically) the objective components (I) and (III), subject to constraints (2)-(10) and (13)-(31), where the barge-type allocation variables in (13) are fixed based on the result of the voyage allocation algorithm.

The voyage-scheduling algorithm is invoked for each vessel in turn, starting with the one with the highest demurrage rate. For any given vessel, the sequencing of the barge voyages can be seen as a variant of the two-stage multi-machine hybrid flow shop problem (Ruiz and Vázquez Rodríguez 2010), with (i) loading as the first-stage operation and discharging the second stage; (ii) the suppliers as the machines; and (iii) the barge voyages as individual jobs, with the objective to minimize the discharge time of the last barge. Therefore, we use a list scheduling algorithm that determines the sequence of barge voyages by taking into account (i) the earliest availability of each barge and (ii) the loading restrictions of each supplier. Transforming the sequence of voyages into a schedule for each voyage is done by starting from the discharging operations and propagating backwards, using shifting operations to accommodate any loading, discharging and barge availability restrictions. Finally, a swapping operation checks if an alternative allocation of loading and discharging slots among the voyages would result in a lower maximum waiting time, thereby reducing detention penalties.

System Output

Figure 10 shows the result of the scheduling algorithm for one particular customer vessel, which visualizes the schedules for each barge allocated to that vessel. The system also provides a breakdown of the total cost into vessel- and barge-related costs, and calculates a cost per tonne, an important performance indicator of the economic efficiency of a schedule.




Figure 10: Timing and duration of the barge operations for each vessel. The top bar is a timeline, with below the vessel laytime and an estimate of demurrage. The voyages of each barge are indicated by their states (SJ=sailing to jetty, WL=waiting for loading, L=loading, D=waiting documents, S=sailing to anchorage, WA=waiting at anchorage, DA=discharging at anchorage, R=returning to hub). Slack time shows when a barge idles before the start of a voyage.

Implementation and Adoption

The barge rotation system is currently in operation at both the Taboneo port, in South Kalimantan, and the Muara Kaman port, in East Kalimantan. Originally the authors were instructed to design and implement the system in the south port only, but because an early prototype (in June 2012) showed great promise, this was later extended to the east port as well (in February 2013). Overall, the cycle of designing, developing, testing and refining the system lasted approximately a year.

A design constraint posed by Noble's senior management was that the barge rotation system should run in a spreadsheet environment, such as Microsoft Excel, and that users should be able to run it without installing any additional software. This was a rigid requirement because the spreadsheet is circulated via internal email so that managers from other divisions can review, modify and invoke the scheduling process. In order to ensure maximum compatibility with the spreadsheet environment, we developed a custom algorithm in Visual Basic. The integer programs in the initial allocation phase are well within the variable and constraint limits posed by the standard Excel Solver.

We also developed a procedure for monitoring the quality of the solutions generated by the system, by comparing them to a lower bound for the total cost computed using a column-generation approach based on a Dantzig-Wolfe decomposition of the mathematical formulation (where the vessel-specific constraints, such as the covering of demand from each supplier, are included at the subproblem level, and the vessel-crossing constraints, such as the allocation of loading time slots at jetties, are included at the master level). The column generation process is invoked from Excel but solved with an advanced solver. However, this means that the column generation part is not portable, and is currently only used by one dedicated logistics manager,

responsible for ensuring the quality of the schedules. We expect that over time, confidence will grow in the capabilities of the system, so that checking the solution quality will no longer be required. We continue to support Noble with maintaining and refining the barge rotation system at the time of this writing.

Realized Benefits

The adoption of the barge rotation system brought multiple benefits to the Noble Group, both quantitative and qualitative. Results collected during a period before and after the system implementation indicate that the realized benefits are on the order of \$1 million per month. Figure 11 shows the evolution of the monthly cost per carried tonne, which includes demurrage costs, dispatch bonuses, barge hiring and detention costs, and transport cost of owned barges, before and after the implementation of the barge rotation model in the two ports, as well as a six-month moving average. In particular, it shows that the six-month average cost per tonne in the east port was reduced from \$3.7 to \$2.2, a reduction of \$1.5, while in the south port it was reduced by \$1.8. Given the number of tonnes transported during the observation period, this represents a total savings of approximately \$1.3 million per month or \$15 million per year. Despite the rather high volatility in the monthly average cost, a statistical analysis confirms that the reduction in average cost per tonne is significant at a 5% level.




Figure 11: The system implementation reduced the average cost per tonne (dotted lines) in both ports significantly. The bold lines denote six-month moving averages.

Naturally, other factors may also have had an influence on the transportation costs, such as supplier availability, the amount of vessel traffic intensity, the availability of floating cranes and the oil price. In an attempt to isolate the effect of the model implementation and control for the impact of supplier availability, traffic intensity and the other factors, we carried out a regression analysis focusing on the South Kalimantan port, for which we had impact data over a longer period. We used an aggregate measure of supplier availability to control for supply disruptions, and the carried tonnage per month as a proxy for vessel traffic intensity. As delays can propagate to following months, we also tested a version of carried tonnage lagged by one month. Further, we combined supplier availability with carried tonnage and lagged carried tonnage to test if heavy vessel traffic has an impact only when it is combined with bad supplier performance. Our dependent variable was the monthly cost per tonne. Finally, we controlled for the cost of oil and

for floating crane availability. We ran four variable selection methods, namely backward, forward, stepwise and best-subsets regression, in order to see which combination of explanatory variables yielded the best outcome, as measured by the adjusted R^2 of each model. The dataset we used has a balanced number of observations before and after the barge rotation system implementation.

<i>Explanatory variable</i>	<i>Best Subset</i>	<i>Forward Selection</i>
Supplier performance	2.34 (1.34)	—
System implementation	-2.47 (0.60)**	-2.04 (0.45)**
Tonnes carried	—	—
Tonnes carried lagged	-4.61 (2.45)	-1.78 (1.63)
(Tonnes carried) * (supplier performance)	2.95 (2.60)	—
(Tonnes carried lagged) * (supplier performance)	4.72 (3.36)	—
Oil price	0.06 (0.02)*	0.05 (0.02)
Adjusted R^2	0.62	0.61

Notes. Standard errors of regression coefficients appear in parenthesis. Stepwise regression and backward regression selected the same model as best subset selection. * $p < 0.05$; ** $p < 0.01$.

Table 2: A regression analysis for the South Kalimantan port reveals a significant decline of cost per tonne due to the system implementation.

All regression models show a significant (<1%) reduction in the average cost per tonne as a result of the system implementation (see Table 2). The forward-selection model is most conservative, and shows the lowest impact at \$2.04/tonne, which corresponds to estimated savings of \$9 million per year for South Kalimantan only.

Alongside the improved efficiency of operations and associated cost reductions, the barge hiring recommendation made by the system also yielded important qualitative insights. For example, for cases in which many vessels arrive within a small time interval, the barge rotation system tends to recommend either hiring a large number of leased and spot barges, or hiring no leased or spot barges at all, depending on the corresponding demurrage penalties and hiring

costs. Solutions that utilize both barge hiring while also incurring demurrage are typically not optimal. This was an unexpected result, and scheduling managers did not anticipate the fact that batch hiring can be optimal in busy periods, but when they realized that it could be beneficial, they started adopting this practice. We were able to verify the optimality of this extreme-hiring structure in small examples, with up to four vessels, using mixed-integer programming.

Finally, it is interesting to note that the portability of the system enabled its circulation to other departments that cooperate with the logistics department. In particular, it is now also used by the marketing team when they negotiate the arrival dates of new shipments and by the floating crane management team that needs to know when the floating cranes can be made available to external customers in order to generate additional income. An important factor that contributed to this wider adoption is the user-friendly graphical interface, which makes it easy to observe and amend the barge, cargo and floating crane availability.

Challenges and Opportunities in the Maritime Industry

Although a significant number of operations research applications in the maritime world have been reported in recent years (e.g., Furman et al. 2011, Wagner and Radovilsky 2012, Agra et al. 2013, Varelas et al. 2013), many large maritime businesses continue to make complex operational decisions based primarily on manual interventions, using intuition and limited data. Significant advances that have been made in optimization mean that a broader class of problems can now be tackled successfully, although customization is often still required. With the maritime environment being a fruitful area for operations research applications, and operations research being able to bring tangible benefits to the maritime businesses, we hope that our application inspires a closer collaboration between the two communities.

Appendix: The Barge Rotation Model

In this appendix, we outline a mixed-integer linear programming (MILP) formulation for the barge rotation problem. The formulation subsumes the voyage allocation model (Figure 9, blocks I and V), which minimizes the barge transportation cost (component II of the objective function, see below) subject to Constraints 11 and 12, barge availability and vessel demand, respectively. With given voyage allocation decisions, the voyage scheduling algorithm (Figure 9, blocks II and IV) generates a solution that is feasible for the remaining constraints, (2)-(10) and (13)-(31). Realistic instances include as many as 15 vessels with up to 16 barges per vessel, which cannot currently be solved with commercial MILP solvers; their limit is around 4 vessels and 7 barges per vessel. We note that the model is amenable to a Dantzig-Wolfe decomposition, with each subproblem corresponding to a single vessel, and Constraints (18)-(19) and Constraints (27), the loading and voyage-sequencing constraints, respectively, as linking constraints. The exact solution of real instances is the subject of current and future research.

The barge rotation model uses the following notation.

Sets and Indexes

$b \in \mathcal{B}$:	Regular (Noble-owned) barges.
$s, \bar{s} \in \mathcal{S}$:	Suppliers.
$\tau \in \mathcal{T}$:	Barge types.
$t \in T$:	Periods in the horizon.
$v, \bar{v} \in \mathcal{V}$:	Vessels.
$o \in \mathcal{O} := \{load, dis\}$:	Set of transshipment operations.

Subsets and Indexed Sets

$\mathcal{R} \subseteq \mathcal{T}$:	Regular barge types.
$\mathcal{R}_s \subseteq \mathcal{T}_s$:	Regular barge types allowed at supplier s .
$\mathcal{S}_t \subseteq \mathcal{S}$:	Suppliers whose jetty is blocked at time t .
$\mathcal{S}_v \subseteq \mathcal{S}$:	Suppliers who serve vessel v .
$\mathcal{T}_s \subseteq \mathcal{T}$:	Barge types that can be sent to supplier s .
$\mathcal{F} \subseteq \mathcal{V}$:	Vessels that need a floating crane.
$\mathcal{V}_t \subseteq \mathcal{F}$:	Set of vessels for which a floating crane is not available at time t .

Parameters

$n_{\tau v}^{max}$:	Maximum number of barges of type τ that can be allocated to vessel v .
$c_{\tau s}^{ton}$:	Tonnage cost of barge type τ when directed to supplier s [$\$/tonne$].
$c_{\tau s}^{det}$:	Detention cost of barge type τ when directed to supplier s [$\$(tonne \cdot day)$].
$t_{\tau s}^{det}$:	Detention time window of barge type τ when directed to supplier s [days].
r_v^{dem} :	Agreed demurrage rate of vessel v [$\$/day$].
r_v^{des} :	Agreed despatch rate of vessel v [$\$/day$].
t^{doc} :	Documents processing time [days].
t^{ret} :	Time to return to hub from the port [days].
t_s^{sail} :	Sailing time to supplier s [days].
$t_{\tau s}^{load}$:	Duration of the loading operation of barge type τ at supplier s [days].
$t_{\tau v}^{dis}$:	Duration of discharge operation of barge type τ at vessel v [days].
l_v :	Agreed laytime of vessel v [days].
eta_v :	Estimated time of arrival of vessel v .
q_{sv} :	Quantity to be carried from supplier s to vessel v [tonnes].
cap_{τ} :	Capacity of barge type τ [tonnes].

We define a set of integers, which denotes voyages associated with each supplier $s \in \mathcal{S}_v$ and vessel $v \in \mathcal{V}$:

$$w, \bar{w} \in \mathcal{W}_{sv} := \left\{ 1, \dots, \left\lceil \frac{q_{sv}}{\min_{\tau} cap_{\tau}} \right\rceil \right\}, \forall s \in \mathcal{S}_v, v \in \mathcal{V},$$

where $\left\lceil \frac{q_{sv}}{\min_{\tau} cap_{\tau}} \right\rceil$ indicates the maximum number of barges needed to carry q_{sv} tonnes to vessel v . Note that the actual number of voyages depends on the size of the allocated barges, and can be less than the maximum. In particular, if barges larger than the minimum size are allocated, fewer voyages might be needed. We call each chosen voyage *active*, and assign a binary variable showing when a voyage is active, as explained below. In addition, we denote a voyage w to supplier s of vessel v as (w, s, v) , and define the following sets of *pairs of voyages* to facilitate the notation:

$$\begin{aligned} \mathcal{P}^{load} &:= (\mathcal{W}_{sv} \times \mathcal{S}_v \cap \mathcal{S}_{\bar{v}} \times \mathcal{V}) \times (\mathcal{W}_{\bar{s}\bar{v}} \times \mathcal{S}_v \cap \mathcal{S}_{\bar{v}} \times \mathcal{V}): && \text{Pairs } (w, s, v), (\bar{w}, s, \bar{v}) \text{ that load from the same supplier.} \\ \mathcal{P}^{dis} &:= (\mathcal{W}_{sv} \times \mathcal{S}_v \times \mathcal{V}) \times (\mathcal{W}_{\bar{s}\bar{v}} \times \mathcal{S}_v \times \mathcal{V}): && \text{Pairs } (w, s, v), (\bar{w}, \bar{s}, v) \text{ that discharge on the same vessel.} \\ \mathcal{P}^{voy} &:= (\mathcal{W}_{sv} \times \mathcal{S}_v \times \mathcal{V}) \times (\mathcal{W}_{\bar{s}\bar{v}} \times \mathcal{S}_{\bar{v}} \times \mathcal{V}): && \text{All pairs } (w, s, v), (\bar{w}, \bar{s}, \bar{v}). \end{aligned}$$

Decision Variables

$z_{\tau sv} \in \mathbb{N}$:	Number of barges of type τ allocated to supplier s for vessel v .
$x_v^{dem} \geq 0$:	Amount of demurrage for vessel v , [\$].
$x_v^{des} \geq 0$:	Amount of despatch for vessel v , [\$].
$x_{wstv}^{det} \geq 0$:	Amount of detention of voyage (w, s, v) that uses a type τ barge, [days].
$x_{wsv}^{load} \geq 0$:	Start of loading for voyage (w, s, v) , [time].
$x_{wsv}^{dis} \geq eta_v$:	Start of discharge for voyage (w, s, v) , [time].
$x_v^{com} \geq eta_v$:	Loading completion time for vessel v .
$y_{wstv}^{voy} \in \{0, 1\}$:	=1 if voyage (w, s, v) uses a type τ barge, 0 otherwise.
$y_v^{dem} \in \{0, 1\}$:	= 1 if the vessel is in demurrage ($x_v^{dem} \geq 0$), 0 otherwise.
$y_{wsv}^{bar} \in \{0, 1\}$:	=1 when voyage (w, s, v) is allocated to barge b , 0 otherwise.
$y_{wstv}^{load} \in \{0, 1\}$:	=1 if loading for voyage (w, s, v) starts on day t , 0 otherwise.
$y_{wstv}^{dis} \in \{0, 1\}$:	=1 if discharge for voyage (w, s, v) starts on day t , 0 otherwise.
$y_{wv\bar{w}\bar{v}s}^{seql} \in \{0, 1\}$:	=1 if voyage (w, s, v) loads cargo before voyage (\bar{w}, s, \bar{v}) , 0 otherwise.
$y_{wsv\bar{w}\bar{v}}^{seqd} \in \{0, 1\}$:	=1 if voyage (w, s, v) discharges cargo before voyage $(\bar{w}, \bar{s}, \bar{v})$, 0 otherwise.
$y_{wsv\bar{w}\bar{v}}^{seqv} \in \{0, 1\}$:	=1 if voyage (w, s, v) is completed before the start of voyage $(\bar{w}, \bar{s}, \bar{v})$, 0 otherwise.

Objective Function

The objective function takes into consideration three cost components: (1) the joint demurrage cost or despatch bonus for all vessels, (2) the total transportation cost, which depends on the barge type (i.e., its size and its contract structure), and (3) the penalty detention, which occurs if the voyage transshipment operations exceed a predefined time window.

$$\min \sum_{v \in \mathcal{V}} \underbrace{(x_v^{dem} - x_v^{des})}_I + \sum_{v \in \mathcal{V}} \sum_{s \in \mathcal{S}_v} \sum_{\tau \in \mathcal{T}_s} \underbrace{c_{\tau s}^{ton} cap_{\tau} z_{\tau sv}}_{II} + \sum_{v \in \mathcal{V}} \sum_{s \in \mathcal{S}_v} \sum_{\tau \in \mathcal{T}_s} \sum_{w \in \mathcal{W}_{sv}} \underbrace{c_{\tau s}^{det} cap_{\tau} x_{wstv}^{det}}_{III}. \quad (1)$$

Demurrage and Despatch: Definitions and Penalties

Constraints (2)–(10) model the penalties incurred when the completion of loading exceeds the ETA by more than the laytime, the bonuses received for early completions, and the detention

amount of each voyage.

$$x_v^{dem} \leq My_v^{dem}, \quad \forall v \in \mathcal{V}. \quad (2)$$

$$x_v^{dem} \leq r_v^{dem}(x_v^{com} - eta_v - l_v) + M(1 - y_v^{dem}), \quad \forall v \in \mathcal{V}. \quad (3)$$

$$x_v^{dem} \geq r_v^{dem}(x_v^{com} - eta_v - l_v) - M(1 - y_v^{dem}), \quad \forall v \in \mathcal{V}. \quad (4)$$

$$x_v^{des} \leq M(1 - y_v^{dem}), \quad \forall v \in \mathcal{V}. \quad (5)$$

$$x_v^{des} \leq r_v^{des}(eta_v + l_v - x_v^{com}) + My_v^{dem}, \quad \forall v \in \mathcal{V}. \quad (6)$$

$$x_v^{des} \geq r_v^{des}(eta_v + l_v - x_v^{com}) - My_v^{dem}, \quad \forall v \in \mathcal{V}. \quad (7)$$

$$x_v^{com} - eta_v - l_v \leq My_v^{dem}, \quad \forall v \in \mathcal{V}. \quad (8)$$

$$x_v^{com} - eta_v - l_v \geq M(y_v^{dem} - 1), \quad \forall v \in \mathcal{V}. \quad (9)$$

$$x_{wstv}^{det} \geq x_{wsv}^{dis} + (t_{\tau s}^{dis} - t_{\tau s}^{det})y_{wstv}^{voy} - x_{wsv}^{load} - M(1 - y_{wstv}^{voy}), \quad \forall w \in \mathcal{W}_{sv}, \tau \in \mathcal{T}_s, s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (10)$$

Barge Capacities and Links of Barge Allocation, Voyage Allocation, and Operational Decisions

Constraints (11)–(16) model the allocation of barge types to vessels and voyages. Specifically, Constraints (11) impose an upper limit on the maximum number of barges for each barge type and vessel. For owned barges, this upper limit is simply the number of owned barges of each size. For other barges, it is specified by the corresponding contract. Constraints (12) impose that the barges allocated to each supplier should carry the agreed quantity for each vessel, and Constraints (13) indicate the total number of voyages taken from each barge type to the suppliers of each vessel. In addition, Constraints (14) express that at most one barge type should be used for any voyage, and Constraint (15) imposes that if a voyage is served by a regular barge type, then there must be exactly one regular barge of that type that serves it. Note that because Constraints (15) are restricted to regular barges, different barge types denote a difference in size only. The last constraints of this block, Constraints (16), hold true for both the loading and discharging operations, and employ the notation y_{wstv}^o and $o \in \mathcal{O} := \{load, dis\}$ to signify y_{wstv}^{load} and y_{wstv}^{dis} , respectively, and avoid repetition. They indicate that if a voyage is allocated to a barge (and therefore the right side is 1), then loading and discharging should each start at some period, while when a voyage is not allocated any barge (and therefore the right side is 0), the

voyage is not used; therefore, both the loading and discharging operations do not start at any period.

$$\sum_{s \in \mathcal{S}_v: \tau \in \mathcal{T}_s} z_{\tau sv} \leq n_{\tau v}^{max}, \quad \forall \tau \in \mathcal{T}, v \in \mathcal{V}. \quad (11)$$

$$\sum_{\tau \in \mathcal{T}_s} cap_{\tau} z_{\tau sv} \geq q_{sv}, \quad \forall s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (12)$$

$$\sum_{w \in \mathcal{W}_{sv}} y_{wstv}^{voy} = z_{\tau sv}, \quad \forall \tau \in \mathcal{T}_s, s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (13)$$

$$\sum_{\tau \in \mathcal{T}_s} y_{wstv}^{voy} \leq 1, \quad \forall w \in \mathcal{W}_{sv}, s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (14)$$

$$\sum_{b: \tau_b = \tau} y_{wsbv}^{bar} = y_{wstv}^{voy}, \quad \forall w \in \mathcal{W}_{sv}, \tau \in \mathcal{R}_s, s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (15)$$

$$\sum_{t \in T} y_{wstv}^o = \sum_{\tau \in \mathcal{T}_s} y_{wstv}^{voy}, \quad \forall o \in \mathcal{O}, w \in \mathcal{W}_{sv}, s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (16)$$

Linking the Timing of Operations

Constraints (17) and the continuous variables denote the start of an operation with the corresponding time-indexed binary variable. The next group of constraints, (18)–(21), express that no two barges can load simultaneously at the same jetty or discharge simultaneously at a vessel served by a floating crane. For vessels that do not have a floating crane, and therefore can simultaneously serve two barges, Constraints (22) express that any triplet of barges that discharges on the same vessel must have at least one nonoverlapping pair of barges (because at most two barges can discharge simultaneously). This constraint is necessary because if all pairs of a triplet of barges overlap, then three barges will discharge simultaneously for some period. Because at most two barges can discharge simultaneously, there must be at least one nonoverlapping pair in each triplet. The set of constraints, Constraints (23)–(26), models the specifics of the discharge operation. Specifically, Constraint (23) links the end of discharge operations with the vessel's discharge completion time, Constraint (24) imposes the end of sailing to the port as a lower bound on the start of discharging, Constraint (25) expresses that if a pair of voyages is ordered, the completion time of the first is a lower bound on the starting time of the second, and Constraint (26) imposes that the estimated time of arrival of a vessel is a lower bound on the

start of discharging.

$$t \cdot y_{wstv}^o \leq x_{wstv}^o \leq (t+1-\epsilon)y_{wstv}^o + M(1-y_{wstv}^o), \quad \forall o \in \mathcal{O}, t \in T, w \in \mathcal{W}_{sv}, s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (17)$$

$$x_{wsv}^{load} + \sum_{\tau \in \mathcal{T}_s} t_{\tau s}^{load} y_{wst\tau v}^{voy} \leq x_{\bar{w}s\bar{v}}^{load} + M(1-y_{wv\bar{w}\bar{v}s}^{seq}), \quad \forall (w, \bar{w}, s, v, \bar{v}) \in \mathcal{P}^{load}. \quad (18)$$

$$y_{wv\bar{w}\bar{v}s}^{seq} + y_{\bar{w}\bar{v}wvs}^{seq} = 1, \quad \forall (w, \bar{w}, s, v, \bar{v}) \in \mathcal{P}^{load}. \quad (19)$$

$$x_{wsv}^{dis} + \sum_{\tau \in \mathcal{T}_s} t_{sv}^{dis} y_{wst\tau v}^{voy} \leq x_{\bar{w}s\bar{v}}^{dis} + M(1-y_{ws\bar{w}\bar{s}v}^{seqd}), \quad \forall (w, \bar{w}, s, \bar{s}, v) \in \mathcal{P}^{dis}. \quad (20)$$

$$y_{ws\bar{w}\bar{s}v}^{seqd} + y_{\bar{w}\bar{s}wsv}^{seqd} = 1, \quad \forall (w, \bar{w}, s, \bar{s}, v) \in \mathcal{P}^{dis} : v \in \mathcal{F}. \quad (21)$$

$$y_{ws\bar{w}\bar{s}v}^{seqd} + y_{\bar{w}\bar{s}wsv}^{seqd} + y_{ws\hat{w}\hat{s}v}^{seqd} + y_{\hat{w}\hat{s}wsv}^{seqd} + y_{\bar{w}\bar{s}\hat{w}\hat{s}v}^{seqd} + y_{\hat{w}\hat{s}\bar{w}\bar{s}v}^{seqd} \geq 1, \quad \forall (w, \bar{w}, \hat{w}) \in \mathcal{W}_{sv} \times \mathcal{W}_{\bar{s}v} \times \mathcal{W}_{\hat{s}v}, s, \bar{s}, \hat{s} \in \mathcal{S}_v, v \in \mathcal{V}. \quad (22)$$

$$x_{wsv}^{dis} + \sum_{\tau \in \mathcal{T}_s} t_{\tau v}^{dis} y_{wst\tau v}^{voy} \leq x_v^{com}, \quad \forall w \in \mathcal{W}_{sv}, s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (23)$$

$$x_{wsv}^{dis} \geq x_{wsv}^{load} + \sum_{\tau \in \mathcal{T}_s} t_{\tau s}^{load} y_{wst\tau v}^{voy} + t^{docs} + t_s^{sail}, \quad \forall w \in \mathcal{W}_{sv}, s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (24)$$

$$x_{wsv}^{dis} + \sum_{\tau \in \mathcal{T}_s} t_{\tau v}^{dis} y_{wst\tau v}^{voy} + t^{ret} - M(1-y_{wsv\bar{w}\bar{s}\bar{v}}^{seqv}) \leq x_{\bar{w}\bar{s}\bar{v}}^{load} - t_{\bar{s}}^{sail}, \quad \forall (w, \bar{w}, s, \bar{s}, v, \bar{v}) \in \mathcal{P}^{voy}. \quad (25)$$

$$eta_v \sum_{\tau \in \mathcal{T}_s} y_{wst\tau v}^{voy} \leq x_{wsv}^{dis}, \quad \forall w \in \mathcal{W}_{sv}, s \in \mathcal{S}_v, v \in \mathcal{V}. \quad (26)$$

Sequencing and Blocking Restrictions

The last part of the model, Constraints (27)–(31), describes sequencing and blocking restrictions. Concretely, Constraint (27) expresses that if the same regular barge is allocated to two voyages, then these voyages must not overlap, Constraint (28) expresses that no discharging can start with periods in which a floating crane is not available, while Constraints (29) show that no discharging can overlap with such periods. We note that Constraints (28) depend on the type of barge that makes the voyage, because the discharging duration, which indicates the periods in which discharging cannot start, depends on the type of each barge. Finally, Constraints (30)

and (31) are similar to Constraints (28) and (29), respectively, and express jetty availabilities.

$$y_{wsv\bar{w}\bar{s}\bar{v}}^{seqv} + y_{\bar{w}\bar{s}\bar{v}wsv}^{seqv} \geq y_{wsbv}^{bar} + y_{\bar{w}\bar{s}\bar{b}\bar{v}}^{bar} - 1, \quad \forall (w, \bar{w}, s, \bar{s}, v, \bar{v}) \in \mathcal{P}^{voy}, b \in \mathcal{B}. \quad (27)$$

$$y_{w,s,t-u,v}^{dis} \leq 1 - y_{ws\tau v}^{voy},$$

$$\forall w \in \mathcal{W}_{sv}, \tau \in \mathcal{T}_s, s \in \mathcal{S}_v, v \in \mathcal{V}_t, t \in T, u \in \{0, \dots, \min\{t-1, \lceil t_{\tau v}^{dis} \rceil - 1\}\}. \quad (28)$$

$$x_{wsv}^{dis} \leq t - t_{\tau v}^{dis} y_{w,s,t-\lceil t_{\tau v}^{dis} \rceil, v}^{dis} + M(2 - y_{w,s,t-\lceil t_{\tau v}^{dis} \rceil, v}^{dis} - y_{ws\tau v}^{voy}),$$

$$\forall w \in \mathcal{W}_{sv}, \tau \in \mathcal{T}_s, s \in \mathcal{S}_v, v \in \mathcal{V}_t, t \in T. \quad (29)$$

$$y_{w,s,t-u,v}^{load} \leq 1 - y_{ws\tau v}^{voy},$$

$$\forall w \in \mathcal{W}_{sv}, \tau \in \mathcal{T}_s, s \in \mathcal{S}_v \cap \mathcal{S}_t, v \in \mathcal{V}, t \in T, u \in \{0, \dots, \min\{t-1, \lceil t_{\tau s}^{load} \rceil - 1\}\}. \quad (30)$$

$$x_{wsv}^{load} \leq t - t_{\tau v}^{load} y_{w,s,t-\lceil t_{\tau s}^{load} \rceil, v}^{load} + M(2 - y_{w,s,t-\lceil t_{\tau s}^{load} \rceil, v}^{load} - y_{ws\tau v}^{voy}),$$

$$\forall w \in \mathcal{W}_{sv}, \tau \in \mathcal{T}_s, s \in \mathcal{S}_v \cap \mathcal{S}_t, v \in \mathcal{V}, t \in T. \quad (31)$$

References

- Agra, A., M. Christiansen, A. Delgado. 2013. Mixed Integer formulations for a short sea fuel oil distribution problem. *Transportation Science* **47**(1) 108-124.
- Al-Khayyal, F., S.J. Hwang. 2007. Inventory constrained maritime routing and scheduling for multi-commodity liquid bulk, part i: Applications and model. *European Journal of Operational Research* **176**(1) 106-130.
- Christiansen, M., K. Fagerholt, B. Nygreen, D. Ronen. §*Handbooks in Operations Research and Management Science 14*, Transportation. Elsevier B.V., 2007.
- Dantzig, G.B., P. Wolfe. 1960. Decomposition principle for linear programs. *Operations Research* **8**(1) 101-111.
- Floudas, C.A., L. Xiaoxia. 2005. Mixed integer linear programming in process scheduling: Modeling, algorithms, and applications. *Annals of Operations Research*, **139**(1).
- Furman, K.C., J. Song, G.R. Kocis, M.K. McDonald, P.H. Warrick. 2011. Feedstock routing in the ExxonMobil downstream sector. *Interfaces* **41**(2) 149-163.
- Geoffrion, A.M. 1974. Lagrangian Relaxation for Integer Programming. *Mathematical Programming Study* **2** 82-114.
- Johnson, S.M. 1954. Optimal two-and three-stage production schedules with setup times included. *Naval Research Logistics Quarterly* **1**(1) 61-68.
- O'Brien, G., R. Crane. 1959. The scheduling of a barge line. *Operations Research* **7**(5) 561-570.
- Persson, J.A., M. Gothe-Lundgren. 2005. Shipment planning at oil refineries using column generation and valid inequalities. *European Journal of Operational Research* **205**(1) 1-18.
- Taylor, G.D., T.C. Whyte, G.W. DePuy, D.J. Drosos. 2005. A simulation-based software system for barge dispatching and boat assignment in inland waterways. *Simulation Modeling Practice and Theory* **13**(7) 550-565.
- Richetta, O., R.C., Larson. 1997. Modeling the increase complexity of New York City's refuse marine transport system. *Transportation Science* **31**(3) 272-293.

Ruiz, R., J.A. Vázquez Rodríguez. 2010. The hybrid flow job scheduling problem. *European Journal of Operational Research* **205**(1) 1-18.

Schwartz, N. L. 1968. Discrete programs for moving known cargos from origins to destinations on time at minimum bargeline fleet cost. *Transportation Science* **2**(2) 134-145.

UNCTAD. Review of maritime transport 2013. Technical report, UNCTAD, 2013.

U.S. Energy Information Administration 2014. Accessed 22 June 2014
<http://www.eia.gov/coal/transportationrates/trend-coal.cfm#fig3>

Vacca, I., Salani, M., & Bierlaire, M. 2010. Optimization of operations in container terminals: hierarchical vs integrated approaches. *Proceeding of STRC*.

Varelas, T., S. Archontaki, J. Dimotikalis, O. Turan, I. Lazarakis, O. Varelas. 2013. Optimizing ship routing to maximize fleet revenue at Danaos. *Interfaces* **43**(1) 37-47.

Wagner, M.R., Z. Radovilsky. 2012. Optimizing boat resources at the U.S. Coast Guard: Deterministic and Stochastic Models. *Operations Research* **60**(5) 1035-1049.