Interfaces

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Improving the Maritime Transshipment Operations of the Noble Group

Ioannis Fragkos, Bert De Reyck

Interfaces
Vol. 46, No. 3, May–June 2016, pp. 203–217
ISSN 0092-2102 (print) | ISSN 1526-551X (online)

http://dx.doi.org/10.1287/inte.2015.0841

Improving the Maritime Transshipment Operations of the Noble Group

Ioannis Fragkos
Department of Technology and Operations Management, Rotterdam School of Management, Rotterdam 3062 PA, Netherlands, fragkos@rsm.nl

Bert De Reyck
UCL School of Management, University College London, London WC 1E 6BT, United Kingdom,
bdereyck@ucl.ac.uk

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. Although the company can hire additional resources (such as barges and floating cranes) in advance to minimize the impact of delays, the economic benefit of doing so is often unclear. 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, and resulting in monthly savings exceeding $1 million.

Keywords: scheduling; maritime; transshipment operations; decision support system; heuristics; decomposition.
History: This paper was refereed. Published online in *Articles in Advance* March 24, 2016.

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 an energy sector that accounts for more than 60 percent 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 the transportation of coal from diverse coal-mining supply sources to ocean vessels. Indonesia, the top exporter of energy coal globally (United Nations Conference on Trade and Development 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 and Muara Kaman, located in the South and the East of the island, respectively. Noble manages the first part of the supply chain; operating on behalf of its customers, it transports the coal from the mines to large ocean-bound vessels. 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, which 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. Figure 2 shows the image of a barge and tugboat.

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 previously agreed prices. 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 the contract specifies, 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 seven days. Table 1 shows the cost details, the potential penalties, and the operational restrictions of each barge type.

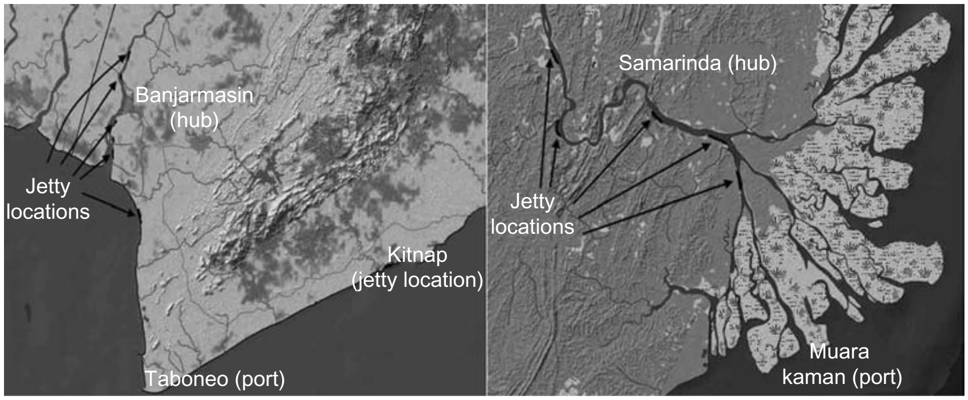

Figure 1: Noble transports coal from mines in Borneo to river jetties, where it is loaded onto barges, which move it to ports, where it is transferred onto ocean-bound vessels.

Figure 2: Coal transport operations utilize barges that carry between 3,000 and 10,000 tonnes of coal.

Barge Logistics

As soon as they are directed to start a voyage, barges sail from a Noble-owned hub, located centrally on one of several major river paths, to the river jetty locations of coal suppliers. Each supplier gives Noble a monthly schedule, which indicates the time windows that it can use for loading. Even when Noble barges arrive within the specified window, however, 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; clearing typically requires one working day, and can therefore span four days because of weekends. After the documents have been 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, each carrying between 20,000

Barge type	Cost structure	Potential penalties	Operational restrictions
Owned	Fuel and other variable costs	None	None
Leased	Price per tonne, which 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
Spot market	Price per tonne, which 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.

and 120,000 tonnes, and may require between three and 16 barges. Smaller vessels are typically equipped with on-board cranes that allow simultaneous loading from both sides. To load 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 (1) the vessel may not have arrived yet; (2) pther barges are being discharged onto the vessel; or (3) a floating crane, which is sometimes required to discharge the coal, might 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; Figure 4 shows discharging operations with on-board and floating cranes.

Demurrage and Despatch

Overall, a voyage (from hub to hub) can require between three and 10 days; therefore, a barge voyage needs to start well in advance of the anticipated arrival of the customer vessel to be able to start discharging on time. Each vessel has an estimated time of arrival (ETA); see Figure 5. Noble and the customer mutually agree on the ETA at least two weeks in advance. Also mutually agreed upon is the laytime, the maximum time window in which the vessel must be loaded.

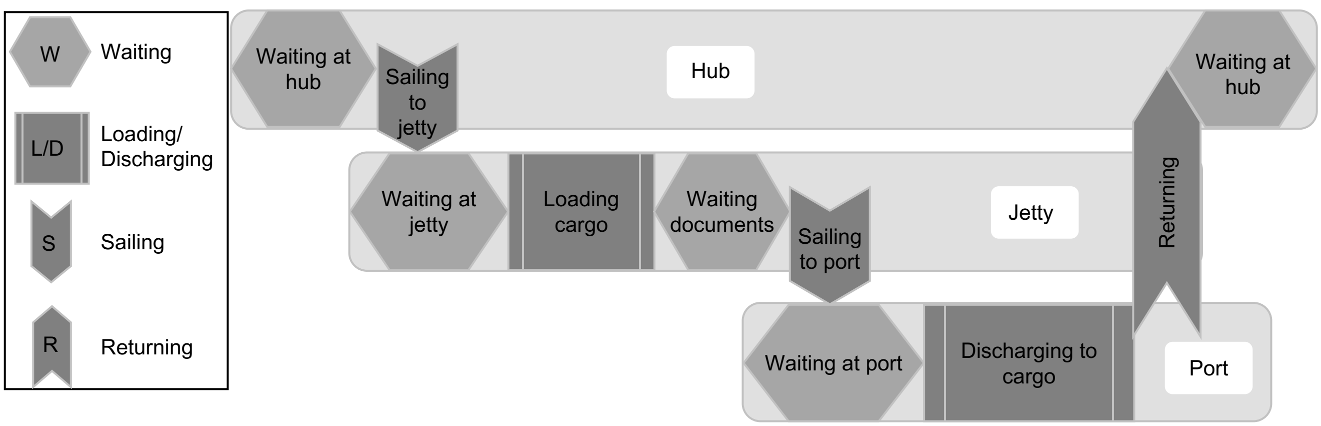

Figure 3: Barges rotate from Noble's hub to a supplier jetty where they are loaded, to the port where they are discharged onto a client vessel, and back to the hub. Waiting for available loading and discharging resources can happen at any stage.

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

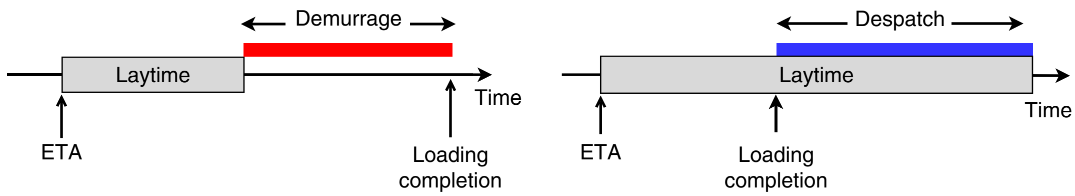

Figure 5: (Color online) A vessel's ETA denotes the start of the laytime. Overruns result in demurrage penalties (left), and early completion results in despatch bonuses (right).

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 occur frequently enough that demurrage penalties 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.

Maritime and Barging Problems

Although maritime problems similar to the one we consider appear sparingly in the operations research (OR) literature, their number has been increasing steadily during the past few years. Christiansen et al. (2007) give a comprehensive review of advances in maritime transportation modeling problems, and Vacca et al. (2010) give an overview of the berth allocation and crane assignment problems. To the best of our knowledge, O'Brien and Crane (1959), who use a simulation model to determine the allocation of tugboats and determine the optimal number of barge loads on the Ohio and Mississippi Rivers, is the first paper relevant to modeling barge transportation. Schwartz (1968) describes a transshipment-scheduling model that minimizes barge fleet costs; however, 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-Khayyal and Hwang 2007, Furman et al. 2011), with the exception of Agra et al. (2013), which study a short-sea transshipment problem.

This paper describes Noble's scheduling problem and the implementation of a decision support system the company adopted to schedule its 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), because the problem of sequencing barge voyages within a single vessel to minimize the loading completion time is a generalization of a two-stage hybrid flow-shop problem (Johnson 1954, Ruiz and Vázquez Rodríguez 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 interrelated decisions: (1) how many owned, leased, and spot barges to allocate to each customer vessel; (2) when to dispatch each barge; and (3) whether to hire a floating crane. Clearly, leasing additional barges can reduce waiting times and demurrage, but with a price. Dispatching barges early can also reduce demurrage, but will result in the barge being tied up for a longer period of time, resulting in additional daily fees and a need for more barges (instead of reusing barges on shorter voyages). Finally, hiring a floating crane can reduce demurrage by speeding up the discharging of coal, but this also has 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. As a result of the complex interactions between the operations and the uncertainty affecting the arrival of client vessels and the availability of jetties and floating cranes, the manual procedure also made it impossible to take into account the complex trade-offs between delays and the cost of additional resources.

Whenever an unexpected event threatened to upset a schedule, Noble managers often quickly put in place additional resources using only rough calculations of anticipated benefits versus cost. It frequently used 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. Noble managers sometimes used more complex rules of thumb, such as allocating a spot barge to a shipment only if the estimated marginal reduction of demurrage penalties outweighed the marginal spot barge cost. Although this rough 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 at close time intervals. Sometimes, however, the rules of thumb were not only myopic, but also incorrect; for example, they might include fixed overhead and sunk costs. Noble often chose to hire a barge rather than use its own, because it deemed the hiring cost to be lower than the cost of operating its own barge. The latter, however, often included fixed costs that would be incurred regardless of whether the barge was used. When senior 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 often viewed these rules and guidelines as confusing and contradictory, resulting in inefficiencies.

The existing manual system also lacked crucial information, such as the current state of operations; for example, it did not include the location of barges and cost information, which were only recorded ex-post, because a proper cost-estimation system was not in place. Additionally, much of the required information (such as updates on supplier availability, prices of spot barges, cost of fuel, and availability of floating cranes) was often held and maintained by other Noble divisions. The logistics managers could not always access these data 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. Given the current start time of each barge, they then had to estimate 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 result in 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 *barge rotation system*, integrates large amounts of information 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 daily, are located in different sheets. This makes the system more ergonomic and facilitates the data entry task. Figures 6–8 show data entry tables containing

Anchorage Name:	Muara Berau	Time Muara Berau - Pulau Buaya:	1.0 Days	Bunker Price ($/Kl):	1045.00	
Hub Name:	Pulau Buaya		0.3 Days			
Location	**COA ($/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: (Color online) The location specifies the cost per tonne and detention penalty (for utilizing a barge for more than seven days) for leased barges of various sizes for each supplier (location).

information on supplier locations (updated monthly), shipments (updated daily), and barges (updated 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.

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 were not readily available, and managers were not comfortable with assigning probabilities to uncertain events, and a stochastic version of our system would

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*, updated daily, includes customer data (top) and quantities sourced from suppliers (bottom). A previously agreed upon contracted loading rate (tonnes per day) and the stowage plan determine the laytime; the *expected* loading rate, however, depends on the loading infrastructure.

File Home Insert Page Layout Formulas Data Review View Developer Add-Ins PowerPivot Team XY Chart Labels

Barges	Barge Type	Size	Current Vessel	Current Jetty	Current Operation	End of Curre
Barge 1	BB	270 ft	MV Vessel 1	KBM	Discharging at MV	25/02/20
Barge 2	BB	270 ft	MV Vessel 1	BMSA	Waits to discharge at Muara Berau	25/02/20
Barge 3	Spot	300 ft	MV Vessel 1	ORBIS	Waits to discharge at Muara Berau	25/02/20
Barge 4	Spot	300 ft	MV Vessel 1	ORBIS		25/02/20
Barge 5	BB	270 ft	MV Vessel 1	KBM		25/02/20
Barge 6	BB	270 ft	MV Vessel 1	PEI - ALH		25/02/20
Barge 7	BB	270 ft	MV Vessel 1	PEI - ALH		26/02/20
Barge 8	CoA	300 ft	MV Vessel 2	BJA		25/02/20
Barge 9	BB	300 ft	MV Vessel 2	BJA	Sailing from Jetty to Muara Berau	25/02/20
Barge 10	BB	300 ft	MV Vessel 2	BKS	Sailing to Jetty	25/02/20
Barge 11	BB	270 ft	MV Vessel 2	BKS	Sailing to Jetty	25/02/20
Barge 12	BB	270 ft	MV Vessel 2	BKS	Not Started Yet	25/02/20
Barge 13	BB	270 ft	MV Vessel 2	BKS	Not Started Yet	25/02/20
Barge 14	BB	270 ft	MV Vessel 3	BKS	Waits to discharge at Muara Berau	24/02/20
Barge 15	BB	270 ft	MV Vessel 4	BBE	Sailing to Jetty	25/02/20
Barge 16	CoA	300 ft	MV Vessel 4	BJA	Waits to discharge at Muara Berau	26/02/20
Barge 17	BB	270 ft	MV Vessel 4	BJA	Sailing to Jetty	26/02/20
Barge 18	FAS	300 ft	MV Vessel 4	ORBIS	Sailing from Jetty to Muara Berau	26/02/20
Barge 19	FAS	300 ft	MV Vessel 4	ORBIS	Sailing from Jetty to Muara Berau	26/02/20
Barge 20	BB	270 ft	MV Vessel 6	BBE	Sailing to Jetty	26/02/20
Barge 21	FAS	300 ft	MV Vessel 6	ORBIS	Discharging at MV	24/02/20

Awaiting Loading
Loading
Waiting for Documents
Processing Documents
Sailing from Jetty to Muara Berau
Waits to discharge at Muara Berau
Discharging at MV
Returns to Pulau Buaya

Figure 8: (Color online) The barge list, updated daily, describes the current state of each barge. Some data are anonymous for confidentiality purposes.

be computationally intractable. Therefore, because rescheduling takes place frequently (often multiple times per day), we decided to build a reactive deterministic model to incorporate unforeseen changes and new information. A reactive system works well in practice, because uncertainty in the short term is not high; 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 also inflated some nominal operation times, such as loading and sailing times, based on feedback from the operators, to create buffers. The appendix includes a detailed mathematical programming formulation of the voyage allocation problem. 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 two subproblems— voyage allocation and voyage scheduling—which are invoked initially to generate a feasible schedule, and are then 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 (1) the quantity that must be sourced from each supplier is covered, and (2) the number of voyages allocated to each vessel does not exceed the maximum number of available barges of each type and size, including leased and spot barges. Using leased or spot barges may be cheaper than using owned barges, because leased or spot barges are often larger and can combine shipments that would otherwise 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 nonincreasing demurrage penalties, and the schedule for each vessel considers restrictions on the availability of barges, floating cranes, and jetties imposed by vessels already scheduled.

Once a feasible schedule is determined for each vessel, the voyage improvement algorithm begins (blocks III–IX). This procedure is necessary because the initial allocation of voyages tends to overuse owned barges, which are typically less expensive, but which might create excessive delays, thus resulting in high

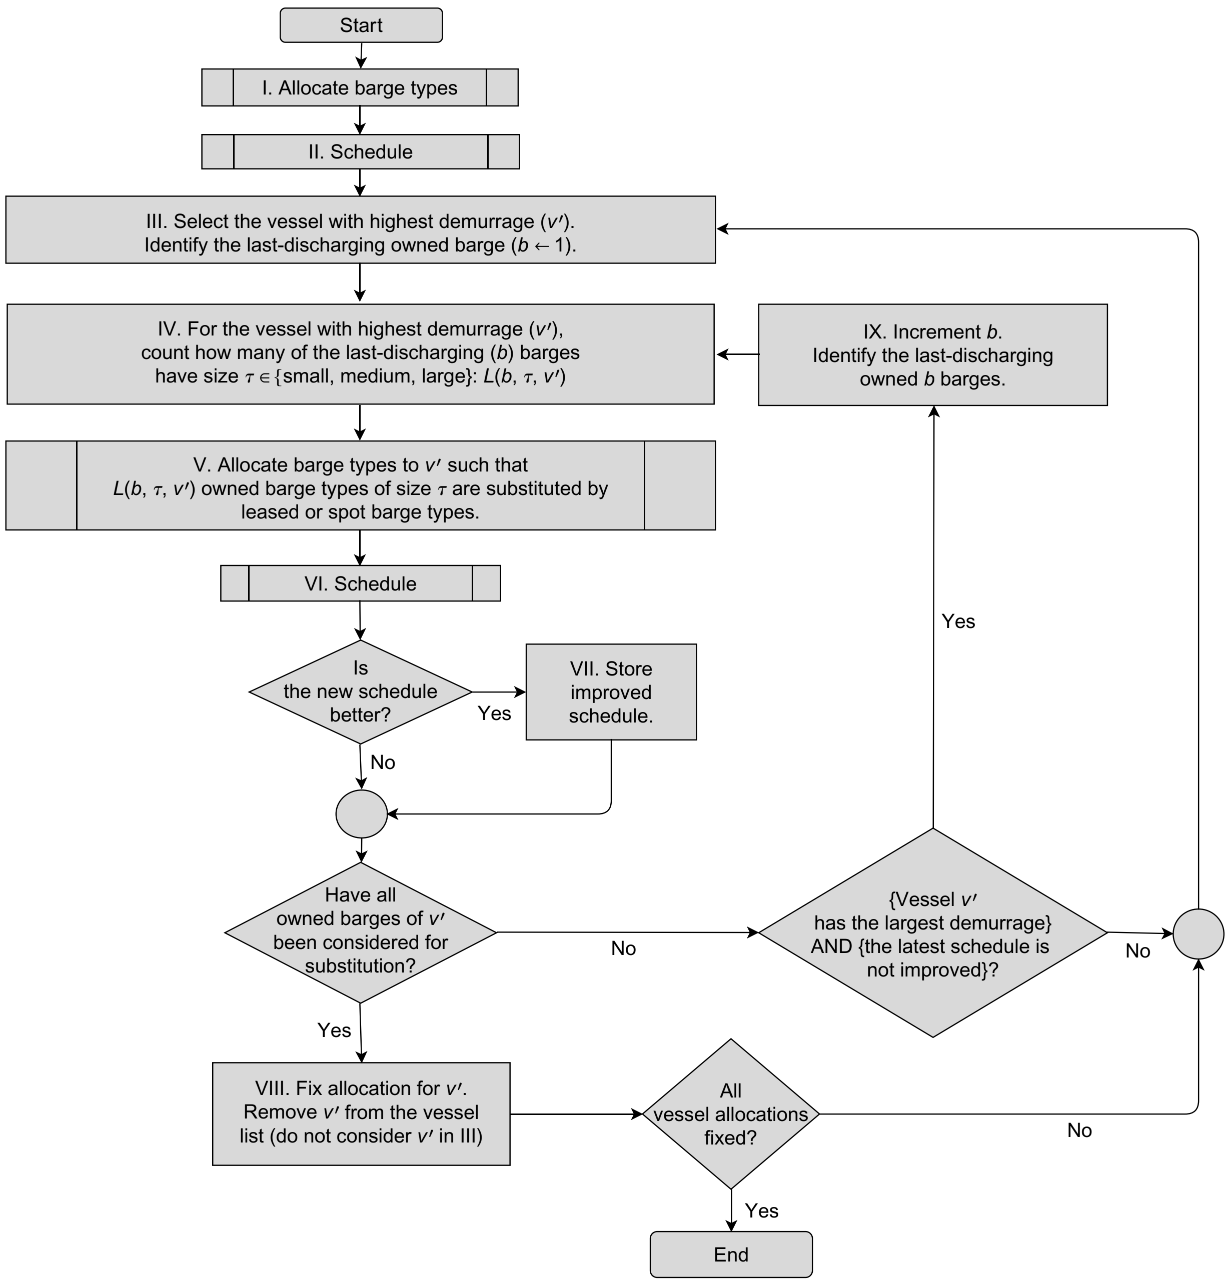

Figure 9: The barge rotation algorithm has a voyage allocation (block V) and a voyage scheduling part (block VI). Blocks III–IX iterate between the two parts, to generate improved feasible solutions.

demurrage. The key concept of the voyage improvement algorithm is to identify the vessel with the highest demurrage (block III) and determine if substituting owned barges with leased or spot barges would result in a lower total cost (blocks IV–VIII). We evaluate both the substitution of sets of barges and 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; this value is then implemented in block V, which reoptimizes the voyage allocation in a way similar to block I. Each 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. For a given allocation, the barge transportation costs are fixed, and the voyage scheduling algorithm aims to minimize any penalties caused by delays (i.e., demurrage and detention). In the mathematical programming formulation in the appendix, the voyage scheduling algorithm minimizes (heuristically) the objective components (1) and (3), subject to constraints (2)–(10) and (13)–(31), where the barge-type allocation variables in constraint (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 multimachine hybrid flow shop problem (Ruiz and Vázquez Rodríguez 2010), with (1) loading as the first-stage operation and discharging as the second stage; (2) the suppliers as the machines; and (3) the barge voyages as individual jobs, with the objective of minimizing 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 (1) the earliest availability of each barge and (2) the loading restrictions of each supplier. We transform the sequence of voyages into a schedule for each voyage by starting from the discharging operations and propagating backwards, using shifting operations to accommodate any loading, discharging, and barge availability restrictions. Finally, we use a swapping operation to determine 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 a particular customer vessel, and visualizes the schedules for each barge allocated to that vessel. The system also breaks down the total cost into vessel- and barge-related costs, and calculates a cost per tonne, an important performance indicator of a schedule's economic efficiency.

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

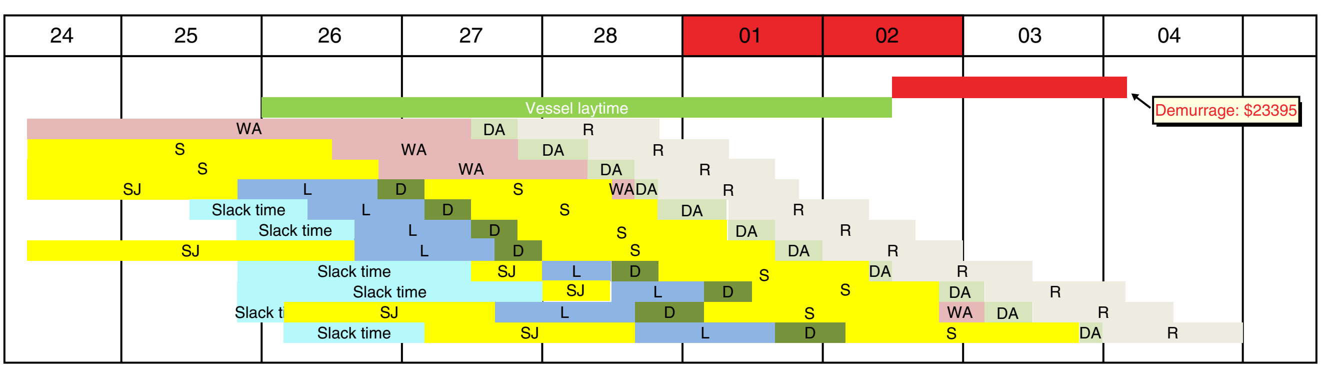

Figure 10: (Color online) The graph shows the timing and duration of the barge operations for each vessel. The top bar is a timeline; the vessel laytime and an estimate of demurrage are below it. The voyages of each barge are indicated by their states (SJ = sailing to jetty, WL = waiting for loading, 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.

the authors were instructed to design and implement the system in the South port only; however, because a prototype in June 2012 showed great promise, this was extended to the East port in February 2013. Designing, developing, testing, and refining the system lasted approximately a year.

A design constraint imposed by Noble's senior management was that the barge rotation system should run in a spreadsheet environment, such as Microsoft Excel, and users should be able to run it without installing any additional software. This was a rigid requirement because management wanted to be able to circulate the spreadsheet via internal email so that managers from other divisions could review, modify, and invoke the scheduling process. 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 imposed by the standard Excel Solver.

We also developed a procedure for monitoring the quality of the solutions generated by the system. We compared them to a lower bound for the total cost computed using a column-generation approach based on a Dantzig-Wolfe decomposition (Dantzig and Wolfe 1960) of the mathematical formulation; 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 using an advanced solver. This means, however, that the column-generation part is not portable; therefore, only one dedicated logistics manager, who is responsible for ensuring the quality of the schedules, currently uses it. We expect that over time, confidence will grow in the capabilities of the system, such that checking the solution quality will no longer be required. As of this writing, we are continuing to support Noble in maintaining and refining the barge rotation system.

Benefits Realized

The adoption of the barge rotation system brought multiple benefits, both quantitative and qualitative, to the Noble Group. Results collected during a period before

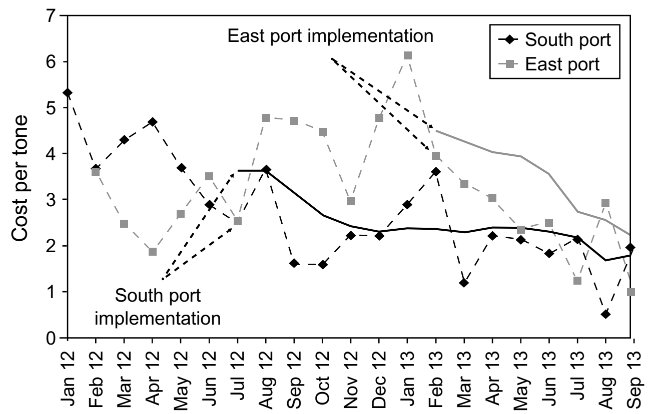

Figure 11: The system implementation significantly reduced the average cost per tonne (dotted lines) in both ports. The bold lines denote six-month moving averages.

and after the system implementation indicate that the realized benefits are approximately \$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, transport cost of owned barges (before and after the implementation of the barge rotation model in the two ports), and a six-month moving average. In particular, it shows that the six-month average cost per tonne was reduced from \$3.7 to \$2.2 in the east port and by \$1.8 in the south port. Given the number of tonnes transported during the observation period, this represents savings of approximately \$1.3 million per month or \$15 million per year. Despite the high volatility in the monthly average cost, a statistical analysis confirms that the reduction in average cost per tonne is significant at a five percent level.

Other factors, such as supplier availability, the amount of vessel traffic intensity, the availability of floating cranes, and the price of oil, may also have influenced the transportation costs. To isolate the effect of the model implementation and control for the impact of supplier availability, traffic intensity, and the other factors, we did 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. Because delays can propagate to subsequent months, we also tested a version in which carried tonnage lagged by one month. In addition, 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 poor 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 (i.e., backward, forward, stepwise, and best-subsets regression) to see which combination of explanatory variables yielded the best outcome, as measured by the adjusted R^2 of each model. The data set we used has a balanced number of observations before and after the barge rotation system implementation.

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

In addition to the improved operational efficiency and associated cost reductions, the system's barge hiring recommendation yielded important qualitative insights. For example, for cases in which many vessels arrive within a short 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, depending on the corresponding demurrage penalties and hiring costs. Solutions that use barge hiring and also incur demurrage are typically not optimal. This was an unexpected result, and scheduling managers did not anticipate that batch hiring can be optimal in busy periods; however, when they realized that this practice could be beneficial, they started adopting it. Using mixed-integer programming, we were able to verify the optimality of this extreme-hiring structure in examples with up to four vessels.

Finally, the system's portability enabled management to circulate it to other departments that cooperate with the logistics department. In particular, the marketing team now also uses it when it negotiates the arrival dates of new shipments and the floating crane management team (which needs to know when the floating cranes can be made available to external customers) uses it to generate additional income. An important factor that contributed to this wider adoption is the user-friendly graphical interface, which facilitates observing and amending the barge, cargo, and floating crane availability.

Explanatory variable	Best subset	Forward selection
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

Table 2: A regression analysis for the South Kalimantan port reveals a significant decline of cost per tonne as a result of implementing our system.
Notes. **Values represent dollars per tonne. Standard errors of regression coefficients appear in parenthesis. Stepwise regression and backward regression selected the same model as the best subset.**
$^{*}p < 0.05$; $^{**}p < 0.01$.

Challenges and Opportunities in the Maritime Industry

Although a significant number of OR applications in the maritime industry 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 manually make complex operational decisions based by using intuition and limited data. Significant advances that have been made in optimization mean that a broader class of problems can now be addressed successfully, although customization is often still required. The maritime environment is a promising area for OR applications, and OR can provide tangible benefits to maritime businesses; therefore, 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 later) 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 four vessels and seven 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 (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.
$o \in \mathcal{O} := \{\text{load}, \text{dis}\}$: 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

$n^{\max}_{\tau v}$: Maximum number of barges of type τ that can be allocated to vessel v.
$c^{\text{ton}}_{\tau s}$: Tonnage cost of barge type τ when directed to supplier s [\$/*tonne*].
$c^{\text{det}}_{\tau s}$: Detention cost of barge type τ when directed to supplier s [\$/(*tonne* · *day*)].
$t^{\text{det}}_{\tau s}$: Detention time window of barge type τ when directed to supplier s [days].
r^{dem}_v: Agreed demurrage rate of vessel v [\$/*day*].
r^{des}_v: 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^{sail}_s: Sailing time to supplier s [days].
$t^{\text{load}}_{\tau s}$: Duration of the loading operation of barge type τ at supplier s [days].
$t^{\text{dis}}_{\tau v}$: 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 $v \in \mathcal{V}$:

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

where $\lceil q_{sv}/(\min_\tau cap_\tau) \rceil$ 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 later. 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:

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

Decision Variables

$z_{\tau sv} \in \mathbb{N}$: Number of barges of type τ allocated to supplier s for vessel v.
$x^{\text{dem}}_v \geq 0$: Amount of demurrage for vessel v, [\$].
$x^{\text{des}}_v \geq 0$: Amount of despatch for vessel v, [\$].
$x^{\text{det}}_{ws\tau v} \geq 0$: Amount of detention of voyage (w, s, v) that uses a type τ barge, [days].
$x^{\text{load}}_{wsv} \geq 0$: Start of loading for voyage (w, s, v), [time].
$x^{\text{dis}}_{wsv} \geq eta_v$: Start of discharge for voyage (w, s, v), [time].
$x^{\text{com}}_v \geq eta_v$: Loading completion time for vessel v.
$y^{\text{voy}}_{ws\tau v} \in \{0, 1\}$: =1 if voyage (w, s, v) uses a type τ barge, 0 otherwise.
$y^{\text{dem}}_v \in \{0, 1\}$: =1 if the vessel is in demurrage ($x^{\text{dem}}_v \geq 0$), 0 otherwise.
$y^{\text{bar}}_{wsbv} \in \{0, 1\}$: =1 when voyage (w, s, v) is allocated to barge b, 0 otherwise.
$y^{\text{load}}_{wstv} \in \{0, 1\}$: =1 if loading for voyage (w, s, v) starts on day t, 0 otherwise.
$y^{\text{dis}}_{wstv} \in \{0, 1\}$: =1 if discharge for voyage (w, s, v) starts on day t, 0 otherwise.
$y^{\text{seq}\,l}_{wv\bar{w}\bar{v}s} \in \{0, 1\}$: =1 if voyage (w, s, v) loads cargo before voyage $(\bar{w}, s, \bar{v})$, 0 otherwise.
$y^{\text{seq}\,d}_{ws\bar{w}\bar{s}v} \in \{0, 1\}$: =1 if voyage (w, s, v) discharges cargo before voyage $(\bar{w}, \bar{s}, v)$, 0 otherwise.
$y^{\text{seq}\,v}_{wsv\bar{w}\bar{s}\bar{v}} \in \{0, 1\}$: =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: (i) the joint demurrage cost or despatch bonus for all vessels; (ii) the total transportation cost, which depends on the barge type (i.e., its size and its contract structure);

and (iii) the penalty detention, which occurs if the voyage transshipment operations exceed a predefined time window.

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

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.

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

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

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

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

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

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

$$x_v^{\text{com}}-\text{eta}_v-l_v\le My_v^{\text{dem}},\quad \forall\, v\in\mathcal{V}, \quad (8)$$

$$x_v^{\text{com}}-\text{eta}_v-l_v\ge M(y_v^{\text{dem}}-1),\quad \forall\, v\in\mathcal{V}, \quad (9)$$

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

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 $o\in\mathcal{O}:=\{\text{load},\text{dis}\}$ 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.

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

$$\sum_{\tau\in\mathcal{T}_s}\text{cap}_\tau z_{\tau s v}\ge q_{sv},\quad \forall\, s\in\mathcal{S}_v,\ v\in\mathcal{V}, \quad (12)$$

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

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

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

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

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. Constraints (23)–(26) model 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.

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

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

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

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

$$y_{ws\bar{w}\bar{s}v}^{\text{seq}\,d} + y_{\bar{w}\bar{s}wsv}^{\text{seq}\,d} = 1, \quad \forall\,(w,\bar{w},s,\bar{s},v)\in\mathcal{P}^{\text{dis}}\colon v\in\mathcal{F}, \quad (21)$$

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

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

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

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

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

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.

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

$$y_{w,s,t-u,v}^{\text{dis}} \le 1 - y_{ws\tau v}^{\text{voy}},\quad \forall\,w\in\mathcal{W}_{sv},\ \tau\in\mathcal{T}_s,\ s\in\mathcal{S}_v,\ v\in V_t,\ t\in T,\ u\in\{0,\dots,\min\{t-1,\lceil t_{\tau v}^{dis}\rceil - 1\}\}, \quad (28)$$

$$x_{wsv}^{\text{dis}} \le t - t_{\tau v}^{\text{dis}} y_{w,s,t-\lceil t_{\tau v}^{\text{dis}}\rceil,v}^{\text{dis}} + M\big(2 - y_{w,s,t-\lceil t_{\tau v}^{\text{dis}}\rceil,v}^{\text{dis}} - y_{ws\tau v}^{\text{voy}}\big),\quad \forall\,w\in\mathcal{W}_{sv},\ \tau\in\mathcal{T}_s,\ s\in\mathcal{S}_v,\ v\in\mathcal{V}_t,\ t\in T, \quad (29)$$

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

$$x_{wsv}^{\text{load}} \le t - t_{\tau v}^{\text{load}} y_{w,s,t-\lceil t_{\tau s}^{\text{load}}\rceil,v}^{\text{load}} + M\big(2 - y_{w,s,t-\lceil t_{\tau s}^{\text{load}}\rceil,v}^{\text{load}} - y_{ws\tau v}^{\text{voy}}\big)\quad \forall\,w\in\mathcal{W}_{sv},\ \tau\in\mathcal{T}_s,\ s\in\mathcal{S}_v\cap\mathcal{S}_t,\ v\in\mathcal{V},\ t\in T. \quad (31)$$

References

Agra A, Christiansen M, Delgado A (2013) Mixed integer formulations for a short sea fuel oil distribution problem. *Transportation Sci.* 47(1):108–124.

Al-Khayyal F, Hwang SJ (2007) Inventory constrained maritime routing and scheduling for multi-commodity liquid bulk, part I: Applications and model. *Eur. J. Oper. Res.* 176(1):106–130.

Christiansen MK, Fagerholt K, Nygreen B, Ronen D (2007) Marine transportation. Barnhart C, Laporte G, eds. *Handbooks in Operations Research and Management Science*, Vol. 14 (Elsevier B.V., Oxford, UK), 189–284.

Dantzig GB, Wolfe P (1960) Decomposition principle for linear programs. *Oper. Res.* 8(1):101–111.

Floudas CA, Lin X (2005) Mixed integer linear programming in process scheduling: Modeling, algorithms, and applications. *Ann. Oper. Res.* 139(1):131–162.

Furman KC, Song J-H, Kocis GR, McDonald MK, Warrick PH (2011) Feedstock routing in the ExxonMobil downstream sector. *Interfaces* 41(2):149–163.

Johnson SM (1954) Optimal two-and three-stage production schedules with setup times included. *Naval Res. Logist. Quart.* 1(1): 61–68.

O'Brien GB, Crane RF (1959) The scheduling of a barge line. *Oper. Res.* 7(5):561–570.

Persson JA, Gothe-Lundgren M (2005) Shipment planning at oil refineries using column generation and valid inequalities. *Eur. J. Oper. Res.* 205(1):1–18.

Richetta O, Larson RC (1997) Modeling the increased complexity of New York City's refuse marine transport system. *Transportation Sci.* 31(3):272–293.

Ruiz R, Vázquez Rodríguez JA (2010) The hybrid flow job scheduling problem. *Eur. J. Oper. Res.* 205(1):1–18.

Schwartz NL (1968) Discrete programs for moving known cargos from origins to destinations on time at minimum bargeline fleet cost. *Transportation Sci.* 2(2):134–145.

Taylor GD, Whyte TC, DePuy GW, Drosos DJ (2005) A simulation-based software system for barge dispatching and boat assignment in inland waterways. *Simulation Model. Practice Theory* 13(7): 550–565.

United Nations Conference on Trade and Development (2013) Review of maritime transport 2013. Accessed November 23, 2015, http://unctad.org/en/PublicationsLibrary/rmt2013_en.pdf.

Vacca I, Salani M, Bierlaire M (2010) Optimization of operations in container terminals: Hierarchical vs integrated approaches. Accessed November 23, 2015, http://www.strc.ch/conferences/2010/Vacca.pdf.

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

Wagner MR, Radovilsky Z (2012) Optimizing boat resources at the U.S. Coast Guard: Deterministic and stochastic models. *Oper. Res.* 60(5):1035–1049.

Verification Letter

Tim Gazzard, Executive Director, Noble Group, Noble Resources Pte Ltd., 60 Anson Road, Singapore 079914, writes:

"I understand that Professor Bert De Reyck and Dr. Ioannis Fragkos have submitted research that is based on the project they completed for PT PELAYARAN NASIONAL TANJUNGRIAU SERVIS ("PNTS"), a Noble Group subsidiary that manages logistics in Indonesia. It is a pleasure to write a statement of support for them as the financial impact of their work is outstanding.

"In May 2012 I contacted Professor Bert De Reyck with a request to investigate the efficiency of the operations of our Indonesian coal logistics division, which at the time I was responsible for. At that time we were facing difficulties with managing our resources and infrastructure, namely barges and floating cranes, and we were experiencing a substantially high logistics cost, resulting from high demurrage penalties, inefficient allocation of barges to vessels and often poor decisions of hiring spot barges. As a result, our cost base was extremely high and we were uncompetitive in the market.

"The system implementation for both areas was carried out between July 2012 and July 2013. Since implementation, we were able to reduce our logistics costs by approximately $1.0 per tonne. These savings amount to around $1.5–$2 million per month, using the most conservative estimates. This tremendous improvement has an unprecedented impact on the financial results of the logistics division. As a result, we are considering how we can apply this system to other parts of our business around the world where logistics run under a similar mode for other commodities.

"An important attribute of the implemented system is its graphical interface, which captures the current state of the operations and shows their future evolution. Managers are able to foresee potential issues well in advance, such as the impact of high vessel traffic, to communicate their concerns to other divisions, and, importantly, to picture the impact of external resources, such as supplier and floating crane availability on the logistics costs. The graphical interface is in a friendly format that is easily understandable. Other divisions, such as marketing and product supply, are using the schedule graphs to support their decision making process. For example, marketing checks that logistics can deliver efficiently future shipments, and product supply evaluates if switching to a spot supplier, in case of missing cargo, could be profitable overall.

"In its totality, the implemented system has not only improved the financial performance of my division, but has also helped the planners improve their understanding of operations, and the other departments to appreciate the complexity of logistics, and understand better how their decisions influence the logistics operations."

Ioannis Fragkos is an assistant professor in the Department of Technology and Operations Management at the Rotterdam School of Management. He has a PhD from the University College London School of Management, was a researcher at the Department of Logistics and Operations Management at HEC Montreal, and was an instructor at London Business School and the London School of Economics and Political Science. His research focuses on decision analytics, decision support systems, and the development of large-scale optimization models.

Bert De Reyck is the director of the University College London School of Management, where he is also a professor in the Analytics and Operations Department. Before joining University College London, Bert held positions at London Business School, the Kellogg School of Management at Northwestern University, and the Rotterdam School of Management. Bert's research focuses on decision analytics, decision support systems, project management, and project portfolio management. He has worked with a wide variety of organisations, including the European Commission, the UK Department for Transport, Network Rail, Eurocontrol, Heathrow Airport, the New York Transit Authority, The Noble Group, Great Ormond Street Hospital, Dunlop Aerospace, and Eaton Aerospace.