REGULAR ARTICLE

A quantitative analysis of container vessel arrival planning strategies

Niels Lang · Albert Veenstra

Abstract   In maritime container transport, the random nature of vessel arrival and terminal service processes often lead to significant handling delays and/or resource underutilization. Arrival planning strategies (APS) promise to mitigate such undesirable effects by managing the variance of the terminal arrival process, taking different cost components and situational dynamics into account. We present a quantitative arrival scheduling simulation to analyze contrasting APS, in order to identify promising strategy design directions. Results are presented for central and decentral strategies, under high and low fuel price regimes. The analysis results in significant quantitative and qualitative differences between the strategies.

Keywords   Maritime container handling · Vessel arrival planning · Vessel-terminal coordination · Optimized arrival scheduling

1 Introduction

Delays of ships are a reality for today's maritime container terminal operations. Notteboom (2006) reports fully loaded vessels being delayed for up to 10 days in southern California during the summer of 2004. UNCTAD (2005, ch. 5B) reports queues of up to 33 deep sea vessels for the same region and period and maximum waiting times of up to 24 h for deep sea vessels in the port of Rotterdam for the same period.

N. Lang (✉)
Imtech ICT Technical Systems, Luchthavenweg 18b, 5657 EB Eindhoven, The Netherlands
e-mail: niels.lang@imtech.nl

A. Veenstra
Department of Information and Decision Sciences, Rotterdam School of Management,
Burgemeester Oudlaan 50, 3062 PA Rotterdam, The Netherlands
e-mail: aveenstra@rsm.nl

Besides terminal delays, liner delays also occur regularly. Analysis of maritime consultant Drewry shows that there is a group of carriers, that is almost never on time, and may easily have delays of 2 or 3 days. The top-20 carriers are mostly in the category that is on time 50–60% of the time (Drewry 2007).

The container shipping system is characterized by high capital-intensity and both shipping companies and terminal operators are under considerable pressure to utilize their resources in an optimal way. This can only be achieved through a delicate balance between terminal resources and liner shipping interests. Terminals would like to have a regular arrival pattern of ships, in line with what was previously agreed or announced by shipping companies. These companies in turn want to minimize their vessels' waiting time, even if they are delayed.

The reason why this balance is currently difficult to attain is that terminals and shipping lines only plan their activities jointly in the contract negotiation stage. At that point in time, the shipping line requests a certain handling capacity over time, and the terminal investigates the time slots in which this capacity can be made available. In daily operations, there is little joint planning, and shipping lines just announce their expected arrival in regular time intervals without reference to the pre-agreed time schedule.

A possible solution is to integrate the operational planning of ship sailing and terminal operations. This paper will investigate one possible approach that is based on controlling the speed of approaching vessels. Such an analysis would have to take the financial consequences of speeding up and slowing down into account, since the fuel bill is one of the main cost components for shipping lines.

The present paper describes an integrated analysis of ship arrivals and terminal operations by means of a simulation model, in which the arrival behavior of ships is controlled by an optimization routine. We measure the impact of our centralized vessel arrival planning strategy by means of a generalized cost approach and compare its performance with an information-poor decentralized arrival planning strategy.

The remainder of this paper is structured as follows. First, we present a theoretical discussion on the relationship between ships and terminals. We then describe the design of our experiment. In subsequent sections, we present more details of the simulation model and the optimization approach. We then report the results, and finish with some concluding remarks.

2 Theoretical framework

The starting point for the present analysis is the idea that ports (read: terminals) should not be planned in isolation, but that their dimensions and operational strategy should take into account the operational efficiency of their users, the shipping companies. This idea was articulated in many port development manuals (see for instance, UNCTAD (2005, p. 9ff)), and is illustrated in Fig. 1.

Port costs consist of a fixed cost component that decreases with increase in traffic in the port, and a variable cost component that is essentially constant across different levels of traffic, but may increase slightly for very high traffic due to the use of more elaborate or expensive cargo handling equipment. The result is a U-shaped total port
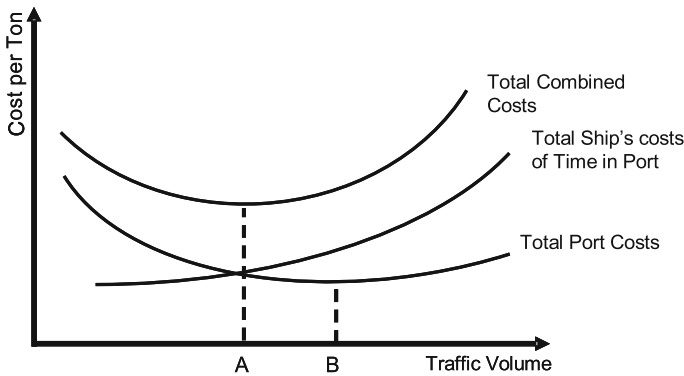
Fig. 1 Total costs in port

cost curve. Ship's cost of time represents the combination of ship's time waiting before berthing and ship's time at berth. Ship's waiting time per ton increases with increasing traffic, while the ship's time at berth is, in principle, independent of traffic, but may decrease slightly for high levels of traffic.

The total combined costs per ton curve is the summation of total port costs and total ship's cost of time per ton. The minimum of this curve can, and will, be different from the minimum of the total port cost per ton curve by itself. This means that port planners have to include the interests of ship owners, in order to develop an efficient port for its users.

The figure illustrates that situations may occur where the minimum of the port costs curve is maintained, rather than the minimum of the joint port–ship cost curve. Examples can be found among fully privatized ports, where the port capacity often significantly lags traffic growth. As a result, liners are often faced with ports operating at relatively high levels of traffic volume, which imposes significant costs on them. This paper investigates joint arrival planning as an instrument to lower the liner's cost curve at such levels of traffic.

The quantitative tool to study ports and shipping operations jointly is queuing theory, where the relation between the service arrangement and the queue in a system are studied in a systematic way. Queuing theory results (Kingman 1970) show that for a system with more or less random arrivals, substantial service delays can be the outcome. More structured or even deterministic arrival patterns, e.g. as obtained by joint arrival planning, are therefore expected to result in reduced service delays and queue sizes.

In the shipping industry, the joint analysis of shipping and terminal operations has been studied surprisingly little. See Stahlbock and Voß (2008) for an elaborate review of OR applications in container terminals. Notteboom (2006) presents a largely qualitative analysis of delays in shipping and ports, but does not offer an integrated analysis. Kia et al. (2002) offer one of a number of contributions on simulation studies of terminals, but they only consider a standard ship arrival process (that follows the negative exponential distribution). Daganzo (1990) studies the interaction in the arrival process on terminal operations of scheduled liners and unscheduled tramp vessels.

He presented approximations for queueing conditions under particular heuristic assignment rules. Imai et al. (2001) focus on the berth allocation problem (allocating a set of ships to a given length of berth, while minimizing waiting and service time). They propose a dynamic formulation of this problem, but they do not take into account arrivals out of schedule. Outside container transport, contributions have been made for bulk terminals by for instance, Zrnic et al. (1999), and Radmilovich (1992). These contributions are more integral in their approach (looking at both ships and port), but describe transport systems that are distinct from container shipping. Finally, agent-based or distributed planning approaches for inland barge traffic in the port of Rotterdam have been suggested (Moonen et al. 2005). This research shows that distributed planning, based on more elaborate information sharing between terminals and inland barges does lead to a better rotation of inland barges in a port. The information shared is about berth availability, and the reaction of the barge operators is to develop a feasible schedule that minimizes total waiting time during the rotation. However, in the study of Moonen et al. (2005), barge operators are not asked to speed up or slow down to make a specific loading/unloading slot.

We conclude that the analysis of ship arrival patterns with out-of-schedule arrivals, and the use of an optimizing control mechanism to improve arrivals at the terminal is a new contribution to the literature on container transport logistics. This analysis builds on a preliminary analysis that was presented by Veenstra and Lang (2005).

3 Case: a fictitious Western European container terminal

We consider a fictitious Rotterdam-based container terminal, serving as first port of call in the Hamburg–Le Havre range for a substantial liner volume. The liner traffic is assumed to consist of large (10 kTEU range) container vessels, serving inter-continental routes. A terminal with a capacity of about 1 million moves per year is assumed, serving about a kilometer of quay with nine ship-to-shore (STS) cranes. The fictitious container terminal will be fully operational 24 h a day, a common situation in the Port of Rotterdam. The terminal will be able to berth up to three vessels concurrently. Four STS cranes, with an assumed constant crane productivity of 30 moves per hour, can handle a vessel's call in about 25 h. It is assumed that remaining terminal operations have enough capacity to support the STS crane operations. Under these assumptions, terminal operations can be planned deterministically.

Our analysis considers arrival patterns of scheduled ships that generally arrive outside their scheduled arrival slot. To generate these arrivals, we use a simulation model. This model generates an initial schedule of ships for arrival at the terminal, taking random delays into account. Additionally, the simulation models a ship's sailing trajectory under varying speeds.

Our berth planning approach assumes it is possible to communicate with ships about the availability of berths. It is assumed that communication infrastructure (e.g. maritime satellite communication) is already in place and that variable communication costs are negligible. Our model suggests the ships to speed up or delay their arrival to arrive just in time for an available berth. For the ship owner, the considerations involved are fuel consumption savings or spending, waiting time reduction and service time. The terminal will benefit from more predictable arrivals.

Speeding up and slowing down have considerable consequences for the fuel bill, since the relationship between ship's speed and fuel consumption is a third power polynomial. In addition, the potential of speed variation for ships is limited due to the construction of the ship and the minimization of stresses and vibration in the hull.

Planning options   If the ship arrives at the terminal out of schedule, there are various options: the first is that the ship has to wait until a berth is free. This waiting time comes at a considerable cost, due to the high cost of investing and operating a ship, and the opportunity cost of carrying other cargo.

A second scenario is that a delayed ship may be served immediately, but is asked to leave early to make room for another scheduled ship. Ship owners sometimes prefer this option to get the ship back into its schedule. This means, however that unloading and loading of all planned containers may not be completed, which also leads to a loss of revenue. This are called 'call gap costs' below.

A third option is serving the ship partly in the slot of another ship, which results in waiting time for that other ship. Such an option can be preferable from a system cost optimization perspective.

Planning costs   The impact of dealing with delayed ships can be measured with a generalized cost function that includes fuel costs, waiting time, service time, and terminal costs. Of these cost elements, Veenstra and Lang (2005) have found that in situations of relatively high traffic volumes, in fact terminal costs are almost constant, whatever the arrival pattern of ships (we assume fixed crane productivity and a finite vessel stream requiring a constant number of crane hours for handling, regardless of the exact arrival and allocation decisions). We will, therefore, ignore terminal costs in the generalized cost analysis. In addition, service times are also almost constant, given our assumptions on call size (see below). In the generalized cost analysis, we therefore concentrate on vessel delay costs, fuel costs, and call gap costs, all from the perspective of a sailing vessel planning its next terminal visit.

Vessel delay costs come into existence when vessel and terminal fail to finish a scheduled terminal call in time. Vessel delay costs consist of various components, such as good depreciation costs, lost vessel opportunity costs, and terminal out-of-window handling penalties. In this case, only goods depreciation costs are considered, which is assumed to amount up to a 100,000 € per day per vessel. This corresponds with annual good depreciation costs of about 3,500 € per year per TEU. We assume an annual depreciation rate of about 35% and 10 ton load per TEU, valued at 1 € per kilogram. We note that delay costs in practice may vary significantly, depending on the exact nature of the goods involved. We charge delay costs for the full vessel load under the assumption that delays are carried until the end of the vessel's (cyclic) schedule. It is assumed that the vessel-delay induced *extra* depreciation costs are fully attributed to the liner operator, e.g. in the form of penalty clauses in the respective transport contracts. We implicitly assume the schedule contains enough slack time to start the next cycle delay-free (such schedule slack time can be considered to be a fixed vessel opportunity cost). In actual business practice, arrangements concerning delay propagation handling and penalty pricing may, of course, vary.

Fuel costs are largely determined by vessel speed. Therefore, fuel prices directly impact planning flexibility. At high fuel price levels, the cost of vessel speed increments will be prohibitive, yielding fewer options for vessel speed adjustments. At low fuel prices, the reverse is expected. From practice, it is well known that fuel prices can fluctuate heavily. For example, between July 2008 until February 2009, the bunker price of Rotterdam sourced IFO380 has fluctuated between 200 and 670 dollar per ton (BunkerWorld 2009).

Call gap costs are costs related to the required rerouting of containers that either could not be loaded on the vessel, or containers that could not be unloaded, since the vessel was running out of time in a particular terminal. We will assume a fixed cost for such containers in the analysis.

In practice, container terminals are characterized by the joint usage by deepsea vessels, short sea vessels and inland barges. We concentrate our analysis on deepsea vessels only, for the sake of clarity of analysis.

Liner behavior Liner traffic is assumed to enter the terminal's planning horizon at a distance of about 5,400 NM (10,000 km). This roughly corresponds to the sailing distance Rotterdam–Santos (Brazil) on the Europe–South America trade route, Rotterdam–Angola on the North-South trade route and Rotterdam–Indian Ocean (between India and the Gulf of Aden) on the East–West trade route. The distance can be sailed in about 10 days at 22.5 knots. We assume that the vessels can sail at speeds ranging from 20 to 25 knots (although possibly at prohibitive fuel costs), which yields a terminal arrival freedom of about 2 days. The vessels are assumed to follow an evenly distributed 4-call schedule in Europe, yielding a Rotterdam call size of 2,500 TEU, for both loading and unloading. Assuming a TEU factor of about 1.7, we obtain an average total call size of about 3,000 moves per vessel. To fill the terminal's capacity of 1 million moves per year, on average about 1 vessel per day will arrive. It is assumed that liner traffic will enter the planning horizon randomly, i.e. vessels will arrive according to a Poisson process.

Planning strategies In order to analyze the potential benefit of joint arrival planning, we contrast *central* and *decentral* arrival planning strategies (APS).

For the central APS analysis, we assume full visibility of relevant vessel attributes, such as speed, call size, and distance. Terminal conditions (resource planning and congestions) are similarly visible. It is assumed that vessels always can and will obey any (reasonable) speed direction issued by the planning entity, even when significant (fuel) costs for the liner are involved. Since a new vessel arrival may alter optimal arrival conditions, the planning should be re-optimized on every vessel arrival. The objective of the planning process is to minimize generalized costs.

For the decentral APS analysis, we assume that vessel speed is determined by vessels themselves, without knowledge of actual terminal conditions. It is assumed that vessels base their decision solely on *estimated* terminal resource availability. Under these rules, vessels select a constant sailing speed when entering the planning horizon and subsequently stick to it. Since we consider a rather homogenous stream of vessels, the 'decentral' speed decisions are likely to be rather similar. It is, therefore, expected that the exponential arrival pattern at the planning horizon will also be seen at the
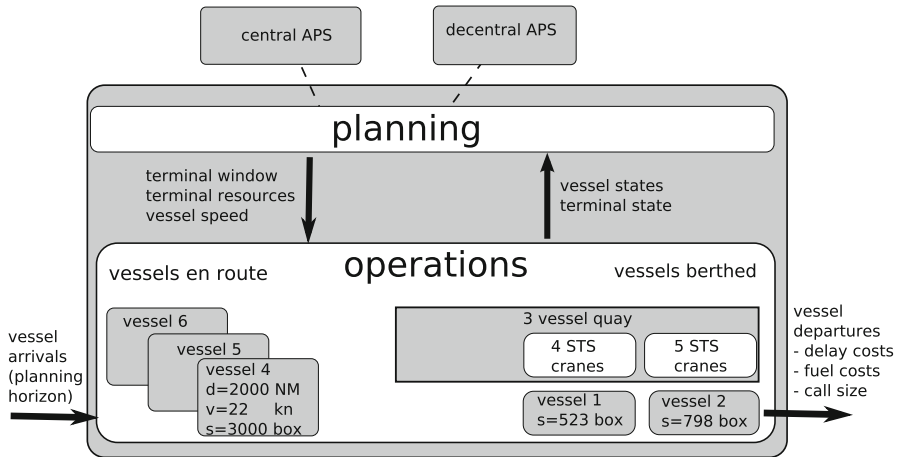
Fig. 2 Arrival planning system

terminal. This high-variance arrival pattern is likely to result in substantial congestion effects at the terminal.

Since the central APS performance is expected to depend heavily on available planning flexibility, we have performed the analysis under a high and a low-fuel price regime. In the high fuel price scenario, a bunker price of $260 €$ per ton is assumed. In the low fuel price scenario, the fuel price is taken to be 35% less.

Summary   Figure 2 graphically summarizes the case described above. The operational system illustrates the simplified terminal and vessel sailing operations assumed in the case. The planning system, driven by one of the APS, is assumed to have access to the relevant vessel and terminal properties, such as distance (d), speed (v) and remaining call sizes (s). Based on the planning strategy, the planning system dynamically reissues vessel speeds and terminal resource allocations. In the case of a decentral APS, most dynamic information will simply be ignored by the planning system. The systemic description of Fig. 2 is the blueprint for a simulation model we developed to assess APS impact. We present this simulation model in the next section.

4 Model description

We simulate the arrival patterns and behavior of vessels. The vessel's speed and terminal allocation decisions are determined by a particular planning model and are recalculated over time. Both vessels and terminal are assumed to strictly adhere to directions issued by the planning model. We implemented planning models for a central and a decentral APS, whose formulations are presented in the next sections. Since our formulation for the central APS requires time discretization in order to remain tractable, the simulation model has been implemented as a periodic step-based simulator. A step size of 2 h has been selected, which fits the time resolution of the APS models. The simulation process behaves as illustrated in Fig. 3.
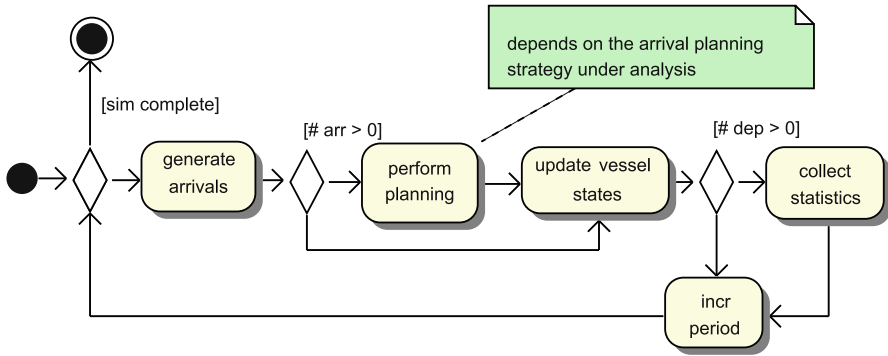
Fig. 3 The simulation process, in UML activity diagram format

Initialization   The simulation initializes cost parameters and selects a planning model based on the scenario being modeled.

Arrival generation   At the start of every period, a number of arrivals are generated according to a Poisson process. Vessels enter the system a the planning horizon, at a distance of 10,000 kilometer from the terminal.[1] Call sizes are normally distributed, with an average and variance of 3,000 boxes. A norm departure time is assigned to each vessel by calculating the earliest possible terminal departure time under nominal sailing and terminal handling conditions. Delay times will be calculated relative to this norm departure time.

Planning   If arrivals occurred in the current period, the simulation will trigger replanning. Otherwise, the previous planning is still deemed optimal. The planning process is performed by a planning model, which implements a particular APS. The planning models are described in detail in the next sections. The outcome of the planning process is a complete (re)assignment of vessel speeds and terminal resource allocations.

Vessel state updating   The latest planning model directions are used to update all vessel states. A vessel's speed is updated such that the terminal will be reached at precisely the assigned time. The vessel's remaining distance is reduced with the distance sailed during the period, and remaining call size is reduced with the number of boxes handled by the terminal during the period.

Statistics collection   Vessels depart from the system once their call at the terminal has been completed. At that time, several statistics are collected for analysis, including the vessel's realized sailing speeds, departure time, and realized call size. To obtain statistically relevant results, statistics for a substantial number of vessels need to be collected. Our analyses were therefore conducted with a warm-up period of 1,000 h, followed by

---

[1] Our implementation is actually somewhat more complicated than described here. Vessels first enter in 'preplan' mode, where they are already visible, but not controllable by the planning model. This allows the planning model to anticipate on upcoming congestions.
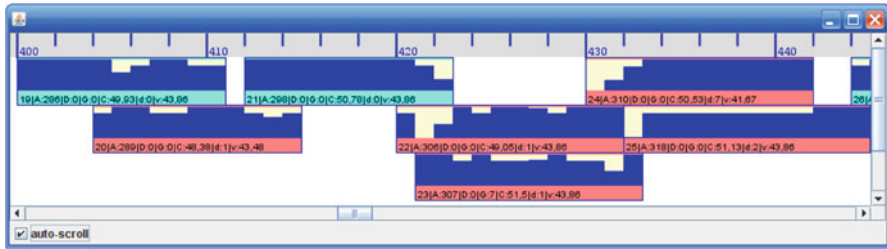
Fig. 4 Some realized vessel operations in the fuel low/speed fixed scenario

10 replications of 2,000 h each. For a particular simulation run, this resulted in about 800 replan operations for about 875 vessel arrivals.

The simulation model is implemented using the Java-based DSOL libraries (Jacobs et al., 2002). Visualization tooling has been developed to support dynamic monitoring of the simulation process.

Figure 4 presents a snapshot of the visualization. The vertical dimension represents the berthing locations at the terminal. The horizontal bar indicates time in 2-h steps. Both past times (i.e. realized visits) as well as future times (i.e. planned visits) can be displayed. In this particular snapshot, some vessel visits realized in the past are displayed.

Container vessel visits at a particular berthing location are represented by the rectangular shapes at particular rows. The left and right side of the rectangles indicate terminal arrival and departure times. Each rectangle consists of a descriptive line (descriptor) and a trajectory indicating the number of cranes allocated over time. The descriptor starts with the vessel ID (e.g. '19'). The color of the descriptor indicates the 'quality' of the terminal visit. A cyan color indicates a 'successful' call, where the vessel was able to handle the complete call size and could depart on time (i.e. before or on the norm departure time). A reddish color (here: vessels 20 and 22–25) indicates the vessel had to reduce the call size and/or departed with delay. In addition, the descriptor shows the vessel's arrival time at the planning horizon (A), the current distance from the terminal (D, in km), the gap between desired and realized call size (G, in boxes), the total crane capacity allocated for the call (C, in crane-periods), the delay on departure (d, in periods) and the average sailing speed (v, in km/h).

In this particular case, the decentral planning model was used. Multiple vessels (20, 22–24) departed with delay even though berthing capacity was available. This can in part be explained by the reduced average crane availability in the presence of multiple vessels. In our model, five cranes can be assigned to a single vessel. However, during full berth utilization, only three cranes are available per vessel. This is why vessel 24 was planned with a delay of 7 periods. It could have arrived earlier, but would have found no cranes available. For vessel 25, the planner in addition had to deal with limited berth capacity.

Most vessels were planned to accept delay rather than to reduce call size. A notable exception is vessel 23, that was planned to skip 7 containers in order to forego another period of delay.

This example raises the question whether these delays could have been prevented under a different APS. Vessel 20 experiences delay since the nominal crane availability assumed was not available in practice. Had the actual terminal situation been taken into account in a dynamic way, the vessel might have opted to speed up in order to prevent the delay. Vessel 22 seems to hinder a number of subsequent vessels. Under a central APS, it might have proven beneficial to speed up this vessel (taking a fuel cost hit) in order to be able to accommodate subsequent vessels better. These examples clearly illustrate potential benefits of a more dynamic, joint APS. In the next sections, the central and decentral APS planning models are explained in detail. Thereafter, analysis results are presented.

5 MIP formulation for the central APS

The planning model implementing the central APS assumes full information visibility and a dynamic planning responsibility for vessel speeds and terminal resource allocation. It minimizes system costs, consisting of delay, fuel and call gap costs, and takes feasibility constraints regarding vessel sailing and terminal handling operations into account. It has been formulated as a Mixed Integer Program (MIP, see e.g. Nemhauser and Wolsey (1988)), which allows the use of common of-the-shelf optimization software to conduct the optimization.

The input for the planning model consists of all vessel states and the terminal state. Vessel state includes in particular *speed*, remaining *distance*, and remaining *call size*. Terminal state consists of the number of cranes and quay positions available and the crane productivity.

The output of the planning model consists of an optimal set of future vessel speeds and terminal resource allocations.

The number of decision variables in our formulation depends linearly on the number of time periods covered by the model. For reasons of tractability, the model therefore discretizes the planning period into 2-h periods.

5.1 Variable definitions

Tables 1, 2 and 3 define time-related, allocation-related, and remaining variables, respectively. In addition, Table 1 defines the subscript i to generally refer to vessel i in the set V of vessels.

Time-related variables   The planning period is divided into $|T|$ time periods t. The period series starts with $t = 0$, which represents the current planning moment. All other periods represent moments in the (planned) future. Each period has a duration of π hours. The time distance relative to the current moment, therefore, equals $t * \pi$. T is bound by the length of the *planning period*. We use a planning period of 200 periods (i.e. more than 16 days). This proved sufficient to accommodate the largest occurring delays (less than 2 days, see Sect. 7).

The time variables $b_{t,i}$, $q_{t,i}$ and $a_{t,i}$ reflect the process of some vessel i. First, the vessel sails to the terminal. These periods are marked with positive $b_{t,i}$. Then, the

Table 1 Time related variables

Variable	Range	Definition (unit,value)
$i \in V$	$\{0, 1, \ldots\}$	Vessel i in V, the collection of vessels under consideration for planning purposes
t	$\{0, 1, \ldots\}$	A time period $\in T$. T is an ordered set of equal-length, consecutive time-periods
π	$(0, \rightarrow)$	Constant: duration of any period t (hour, 2)
t_i^s, t_i^e	$\{-1, 0, 1, \ldots\}$	Terminal handling window decisions, start (exclusive) / end (inclusive) respectively (periods)
τ_i	$\{0, 1, \ldots\}$	Constant: norm period of completion for vessel i
t_i^d	$[0, \rightarrow)$	The total delay time incurred by vessel i, $= max(0, t_i^e - \tau_i)\pi$ (hours)
$b_{t,i}$	$\{0, 1\}$	Binary variable indicating whether t occurs *before* the handling period of vessel i
$a_{t,i}$	$\{0, 1\}$	Binary variable indicating whether t occurs *after* the handling period of vessel i
$q_{t,i}$	$\{0, 1\}$	Binary variable indicating whether t occurs *during* the handling window of vessel i
w_i	$\{0, 1, \ldots\}^{t_i^e - t_i^s}$	Set of consecutive time periods constituting the terminal handling window of vessel i

Table 2 Allocation related variables

Variable	Range	Definition (unit,value)
Q	$\{0, 1, \ldots\}$	Constant: quay length (number of vessels, 3)
$c_{t,i}$	$[0, \rightarrow)$	The number of cranes allocated in period t to vessel i
C	$\{0, 1, \ldots\}$	Constant: total number of STS cranes available (number of cranes, 9)
C_i	$\{0, 1, \ldots\}$	Constant: maximum number of STS cranes able to handle vessel i concurrently (number of cranes, 5)
p_c	$(0, \rightarrow)$	Constant: productivity of a single crane (boxes/hour, 30)
s_i	$[0, \rightarrow)$	*Realized* or, equivalently, *allocated* call size of vessel i (number of boxes)
σ_i	$[0, \rightarrow)$	Constant: *original* call size for vessel i (number of boxes)
v_i^{mn}	$(0, \rightarrow)$	Constant: minimum vessel cruise speed (km/h, 40)
v_i^{mx}	$(0, \rightarrow)$	Constant: maximum vessel speed (km/h, 50)
δ_i	$[0, \rightarrow)$	Constant: momentaneous distance of vessel i from the terminal (km)
F_i	$[0, \rightarrow)$	Trip fuel expenditure of vessel i, relative to minimum trip fuel costs (euro)

vessel enters the terminal at t_i^s. From $t_i^s + 1$ onwards, the vessel is available for terminal handling, up to and including period t_i^e. During this handling period, periods are marked with positive $q_{t,i}$. The window length in hours is equal to $(t_i^e - t_i^s) * \pi$. After handling, the vessel departs at $t_i^e + 1$. These periods are marked with positive $a_{t,i}$. t_i^s may be assigned the special value -1 which allows resource allocation for in-port vessels right from $t = 0$ onwards.

Table 3 Other variables

Variable	Range	Definition (unit,value)
L	$(0, \rightarrow)$	Constant: a sufficiently large, positive constant (used in conditional constraints)
γ_x	$[0, \rightarrow)$	Constant: cost parameter, for some factor x

The norm period of completion for a vessel (τ_i) is the benchmark to determine whether a vessel departs with delay. In our analysis, τ_i is calculated when the vessel enters the system, assuming homogenous nominal sailing and handling conditions, i.e. a sailing speed of 45 km/h and (continuous) availability of four STS cranes. However, the formulation supports vessel arrivals with more significantly varying completion norms (e.g. some vessels may already enter the planning horizon with substantial delay, some liners may operate under more strict service levels than others).

Finally, w_i defines the set of all periods constituting the handling window of vessel i. This definition merely serves to simplify the formulation of some constraints.

Allocation related variables   The terminal resources that can be allocated to vessels are the quay positions and STS cranes. With respect to quay allocation, we assume that 10,000 TEU vessels are rather homogeneously sized, somewhere between 350 and 400 m in length. This implies that a quay of around 1,200 m always offers exactly three vessel slots, which is the value selected for Q. The berth location may be located at various distances from the terminal-side container stack section supporting the vessel's call. We assume that quay-stack transport is organized such that STS cranes are efficiently serviced, even in the presence of a significant vessel-stack distance. We therefore do not consider actual berthing positions in our formulation, but only guard the berth's maximum capacity.

STS crane allocation is constrained by the total number of cranes available (C) and the maximum number of cranes concurrently assignable to a single vessel (C_i). Vessel differentiated crane assignment limits can reflect differences in liner-terminal service level agreements. Crane productivity is assumed constant.

A concrete crane allocation planning results in a realized number of boxes handled during the vessel's call (s_i). This may be less than the original call size desired by the liner (σ_i). In that case, a call size gap, sized $\sigma_i - s_i$ boxes will result.

Variables δ_i and υ_i^{mx} relate to vessel speed allocations. υ_i^{mx} should be defined in a relatively small band around nominal speed. We assume in our formulation that a margin of about 10% relative to a nominal speed of 45 km/h (24.3 knots) is possible, albeit at possibly prohibitive fuel expenses (F_i).

Other variables   Other variables include L, an auxiliary constant used to express various logical constraints involving binary variables. The various cost parameters are clarified in Sect. 5.2.

5.2 Objective function

Table 4 summarizes the cost components that constitute the planning model's objective function. As was discussed above, terminal costs components are not included.

Table 4 Overview of cost components (to be minimized)

Costs	Formula	Remarks
Trip fuel	$\sum_i F_i$	Based on a linearized relation between sailing speed and trip fuel costs
Call gap	$\gamma_{cs} \sum_i (\sigma_i - s_i)$	Call gaps imply diversion costs for non-handled containers
Delay	$\gamma_d \sum_i t_i^d$	Costs of cargo depreciation

Fuel costs In general, the required propulsion power and associated fuel consumption rate for a vessel is a function of speed raised to the third power (Hughes 1996). However, including such relations in the MIP formulation would result in a non-linear formulation which is hard to solve to optimality. We therefore decided to linearize fuel consumption in the speed regime of interest for our analysis, i.e. between 40 and 50 km/h. We based this linearization on propulsion power data reported by MAN (2005) for large (container) vessels.

Equation 1 presents how the quantity of interest for the model formulation, $\frac{dF_{ct}}{dT}$, the marginal impact of a change in sailing time on trip fuel expenses, can be estimated from the speed-fuel consumption linearization.

$$
\begin{aligned}
F_{ct}(V) &= C(V)\frac{D}{V} \\
&\sim (\alpha V + \beta)C_N \frac{D}{V} \\
&= C_N(\alpha D + \beta T) = F_{\tilde{ct}}(T) \Rightarrow \\
\frac{dF_{\tilde{ct}}}{dT} &= C_N \beta
\end{aligned}
\tag{1}
$$

where V denotes constant vessel speed, D sailing distance, T trip time, i.e. sailing period, $F_{ct}(V)$ true trip fuel costs, as a function of vessel speed, $F_{\tilde{ct}}(T)$ estimated trip fuel costs, as a *linear* function of sailing time, $C(V)$ hourly fuel costs, as a function of vessel speed, α, β regression coefficients in the regressed $C(V)$ relation, and C_N norm hourly fuel costs, at some norm speed V_N

We consider trip fuel costs (TFC) on a sailing distance D. When a constant speed is maintained, this can be expressed as $C(V)\frac{D}{V}$. The function $C(V)$ is subsequently linearized with respect to V. This results in an estimated linear TFC function $F_{\tilde{ct}}(V)$, which can be rewritten into $F_{\tilde{ct}}(T)$, an estimated linear TFC function of T. Differentiation yields $\frac{dF_{\tilde{ct}}}{dT}$ (the marginal TFC effect). This estimate only makes sense in the V (and corresponding T) regime for which the linearization was performed. The quantity expresses the TFC impact of a change in sailing time and is used to relate a vessel's TFC to her terminal arrival time (see Sect. 5.3).

The linearized fuel cost relation yields a constant value for the marginal TFC effect. For a simple *cubic* speed–power relation, the (absolute) marginal TFC effect would *increase* in the speed interval considered. We estimated by regression modeling that at 50 km/h, the effect is about twice as large as at 40 km/h. The linearized constant value is within 1% of the maximum marginal TFC effect found for the cubic speed–power relation. This implies that we significantly overestimate the marginal TFC effect at low speeds. More elaborate fuel cost models, incorporating detailed propulsion and hydrodynamic models, should be considered for further research.

At an estimated norm hourly fuel consumption rate of about 2.3 k€ at 24 knots , we arrived at a marginal TFC effect of 11.3 k€ per period. That is, at the sailing distance considered in the model (10,000 km), every period gained by increased sailing speed will be accompanied by 11.3 k€ of extra fuel expenses. For the low fuel cost scenario, a 35% lower fuel price was assumed, resulting in a corresponding marginal TFC effect of 7.4 k€ per period.

Call gap costs   Costs for call diversion include extra handling and transport, which are set at 200 € per box, based on an estimated container transfer cost from Rotterdam to Antwerp. The exception are empty containers, which are patient and have no fixed destination. However, skipping empty containers may impact subsequent vessel call sizes as global equilibrium must be maintained. Empty container skipping is not considered in our formulation.

Delay costs   Based on the assumption presented in the case description (see Sect. 3), we arrive at estimated vessel delay costs of 100 k€ per day, which translates to about 4 k€ per hour delay.

5.3 Constraints

Tables 5 and 6 summarize all time and allocation related constraints included in our formulation. Unless indicated otherwise, the constraints apply to all values t and i.

Time related constraints   The 'before indicator', 'after indicator' and 'exclusive periods' constraints define the phases passed by a vessel. The 'before indicator' constraint marks all vessel periods occurring on or before t_i^s as 'before (handling)' periods. The 'after indicator' constraint marks all vessel periods occurring after vessel departure as 'after (handling)' periods. The 'exclusive periods' constraints express that for each vessel, a period is either a 'before', 'after' or a 'quay' period. As explained previously, 'quay' periods $q_{t,i}$ constitute the vessel's handling window, during which the vessel is berthed and can be handled by STS cranes. Together, these constraints ensure that each vessel moves from a 'before' period (during which sailing occurs) to a 'quay' period (during which handling occurs) and finally to an 'after' period (which occurs after the vessel's departure from the terminal).

The 'arrival' constraint expresses the fact that vessels need a minimum sailing time to travel a certain distance, which is determined by her maximum sailing speed. This prevents the planning model from scheduling impossible vessel speeds. Note that

Table 5 Overview of time-related constraints

id	Formula	Description
before indicator	$t_i^s - L\, b_{t,i} \leq t - 1$	'before' indicator is true whenever t occurs before/on t_i^s
	$t_i^s - L\, b_{t,i} \geq t - L$	'before' indicator is false whenever t occurs after t_i^s
after indicator	$t_i^e + L\, a_{t,i} \geq t$	'after' indicator is true whenever t occurs after t_i^e
	$t_i^e + L\, a_{t,i} \leq L + t - 1$	'after' indicator is false whenever t occurs on/before t_i^e
exclusive periods	$b_{t,i} + q_{t,i} + a_{t,i} = 1$	each period is uniquely qualified as either a 'before', 'quay/handling' or 'after' period
arrival	$t_i^s + 1 \geq \dfrac{\delta_i}{\pi\, v_i^{mx}}$	a vessel requires a minimum sailing time, determined by her maximum sailing speed v_i^{mx} and her remaining sailing distance δ_i
departure	$t_i^s - t_i^e \geq 0$	a vessel cannot depart before arrival.
delay	$\pi\, t_i^e - t_i^d \leq \pi\, \tau_i$	late vessel departures produce delay
no-takeover	$t_i^s \leq t_{i+1}^s$	vessels do not 'take-over' previous
	$t_i^e \leq t_{i+1}^e$	vessels
trip fuel costs	$F_i + \gamma_f t_i^s \geq \gamma_f t_{i,v_i^{mn}}^s$	sailing time decrements increment trip fuel expenses

Table 6 Overview of allocation-related constraints

id	Formula	Description
quay limit	$\displaystyle\sum_i q_{t,i} \leq Q$	No berthing beyond quay limit
terminal crane limit	$\displaystyle\sum_i c_{t,i} \leq C$	No crane allocation beyond the number of terminal cranes available
vessel crane limit	$c_{t,i} \leq C_i$	Limits concurrent crane allocation
vessel presence	$\displaystyle\sum_{t \notin w_i} c_{t,i} = 0 \iff$	No production in absence of the vessel
	$c_{t,i} - L\, q_{t,i} \leq 0$	
call balance	$\pi\, p_c \displaystyle\sum_t c_{t,i} \geq s_i$	The total allocated productivity during a vessel's call at a minimum equals the allocated call size
call size	$\sigma_i - s_i >= 0$	The allocated (i.e. realized) call size cannot exceed the original call size

according to our variable definitions, a vessel has $t_i^s + 1$ periods available for sailing. The 'departure' constraint ensures that vessels do not depart before arrival. Note that vessels can bypass quay allocation & handling if $t_i^s = t_i^e$.

The 'non-takeover' constraint was added to improve the tractability of the formulation. For the homogenous fleet composition assumed in our case, with similar speed and call size characteristics, this constraint is not expected to significantly alter the system costs realized by the planning model.

Finally, the 'trip fuel costs' constraint ensures that vessel TFCs increases with increasing sailing speed (and consequently, earlier arrivals). The constant γ_f equals the estimated TFC impact of a sailing time decrement, $\dfrac{\mathrm{d}F_{\tilde{c}t}}{\mathrm{d}T}$, in € per period sailing time gained (see Sect. 5.2). The right-hand side of the constraint is chosen such that the TFC effect equals zero for a vessel sailing at minimum sailing speed υ_i^{mn}. With this formulation, decreasing vessel speed below υ_i^{mn} will not yield additional TFC benefits. This formulation was selected since our fuel cost linearization does not have validity below υ_i^{mn} (40 km/h). The actual third order relation between speed and fuel consumption is likely to flatten significantly towards lower speeds. Additionally, speeds below υ_i^{mn} are moving farther away from the hull's optimal design speed. As a result, hull inefficiencies will reduce the fuel cost impact of speed reductions below υ_i^{mn} even further.

Allocation related constraints   The 'quay limit' constraint reduces to a simple maximum concurrent presence formulation, given our stated assumptions on homogenous vessel size distribution and berthing position independent terminal handling performance. The 'terminal crane limit' and 'vessel crane limit' constraints limit the number of cranes simultaneously allocated, from a terminal and vessel perspective respectively.

The 'vessel presence' constraint expresses the impossibility of handling a vessel which has not yet arrived at the terminal.

Finally, 'the call balance' and 'call size' constraints ensure that a vessel's allocated call size actually gets handled and does not exceed the vessel's original call size.

5.4 Model implementation

The optimization problem formulation described above is dynamically generated during a simulation whenever a new vessel is generated. The exact variables and constants of the dynamically generated problem are fully determined by the state of the vessels present in the simulation model.

The number of vessels included in the optimization problem is an important driver for the problem's tractability. On average, about 16 vessels will be present in the planning period, which yields a planning problem requiring significant computational effort to solve to optimality. We therefore attempted to minimize the number of vessels considered in the problem formulation. First, a tentative optimization is conducted for the newly arrived vessel *in isolation*. If the resulting solution does not violate the previous planning, the solution is accepted. Otherwise, the problem is regenerated, with the conflicting vessels included. This procedure occasionally still results in large problem formulations, especially at times of congestion, when vessel schedules are most likely to be inter-dependent. In more than half of the planning instances, however, a tractable model containing less than five vessels proved sufficient.

Occasionally, a problem proved hard to solve. Such problems often involved more than 15 vessels. In order to guarantee the continuity of the simulation, we used an automatic stop condition, based on a dynamic maximum gap condition, a maximum node count and the moment of last improvement. A maximum of one million potential
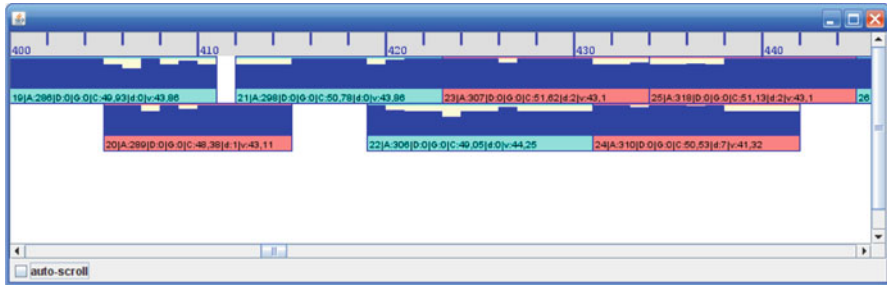
Fig. 5 Some realized vessel operations in the fuel low/speed optimized scenario

solution nodes were allowed. The dynamic maximum gap condition was set to a thousand euros per hour of computation. Finally, the optimization would *not* be stopped if a solution improvement was found less than a hundred seconds before, to prevent a model stop during a fruitful period.

Computational performance The planning problems were solved using ILOG's CPLEX solver on an Intel Core2Duo 2.4 GHz based machine with 2GB of RAM available. The average solution time was about 15 min, with a maximum problem computation time of over 7.5 h. Although we do not rule out the possibility of significant computational improvement using careful optimizer tuning, we deem the computational performance sufficient for the case under study. By far, the major part of the problems were solved to optimality. The average remaining optimality gap was less than 2,000 € per vessel.

Planning performance Figure 5 illustrates the planning model outcomes for the same situation presented earlier in Sect. 4.

When comparing the outcomes, some differences can be noted. Although vessels are still delayed, most vessels (20, 23, 24, 25) are handled faster than before. This allows them to arrive later, thereby saving fuel costs. The planning model made this possible by allocating less vessels concurrently at the quay, which makes more cranes available per vessel. The speeding up of vessel 22 enabled this rearrangement.

In total, this 4-day planning increases net sailing time with 5 periods, while the net amount of incurred delay remains unchanged. This translates to a net fuel consumption saving of about 37 k€ , under the low fuel cost scenario.

6 MIP formulation for the decentral APS

The main assumption of the decentral APS is the limited exchange of vessel speed information with the terminal. We assume that, upon arriving at the planning horizon, vessels announce their current sailing speed to the terminal. This initial speed is based upon an estimated available terminal capacity (since the actual available terminal capacity is not communicated to the vessel).

We further assume vessels will always maintain their initial speed, therewith undoing any potential fuel savings related to delay anticipation. This corresponds to a situation where liners do not actively provide dynamic sailing speed instructions to vessel captains. This formulation fully prevents fuel costs savings by active delay anticipation.

6.1 Model implementation

The decentral APS formulation is identical to the central APS formulation described in Sect. 5 except for the following differences. First, fuel cost considerations are not explicitly optimized, which leads to removal of the 'trip fuel' objective component and the 'trip fuel costs' constraint from the planning problem formulation. Second, the planning model is explicitly disallowed from assigning vessel speed increases, by setting a vessel's maximum speed, υ_i^{mx}, equal to a vessel's initially selected speed.

The quality of the initial speed selection rule may be highly sensitive to the expected terminal capacity available. We therefore performed an empirical optimization on the parameter expected crane availability (ECA). The optimal value was found to be nearly 5 under the low fuel cost scenario. Compared with a value of 4.5, the difference in averaged total vessel costs was less than 5%.

7 Results

We have analyzed the performance of the central and decentral APS under a low and a high fuel price scenario, resulting in four simulation runs. The simulation outcomes are presented in Figs. 6, 7 and 8 below. The simulation runs are labeled by a four character code of the form $F\{L|h\}V\{f|o\}$. The first two characters denote the fuel price scenario, either low (FL) or high (Fh). The last two characters denote the APS, either Vf (fixed speed: decentral APS) or Vo (optimized speed: central APS).

The charts provide information on several variable properties. The 95% reliability interval for a variable's mean value (computed following Law and Kelton (1982)) is indicated by the lined box. The estimated standard deviation $\tilde{\sigma}$ (for which we have not computed a reliability interval) is indicated by a shaded box, which has a length equal to $4\tilde{\sigma}$ and is centered around the variable's estimated mean. This shaded box gives an impression of the variance of the variable's distribution. Finally, the variable's minimum and maximum values (as encountered during any replication) are illustrated by horizontal error lines.

We note that the simulation treatment yielded sufficiently precise estimates for the mean values of the variables of interest. In the following paragraphs, the outcomes are discussed for the low and high fuel price scenario, respectively.

Low fuel price Figure 6 shows that the central APS ($FLVo$) results in significantly lower costs per vessel. The difference between the mean values is about 9.8 k€ per vessel. Figure 7 reveals that this difference can *not* be ascribed to improved delay management under the central APS, since the mean reliability intervals overlap for the average vessel delay. Differences in call size costs were comparatively negligible.
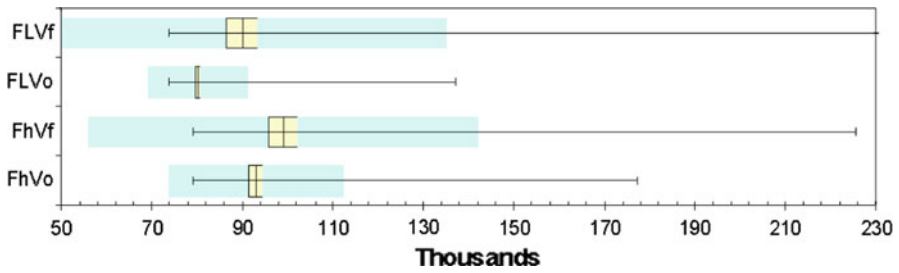
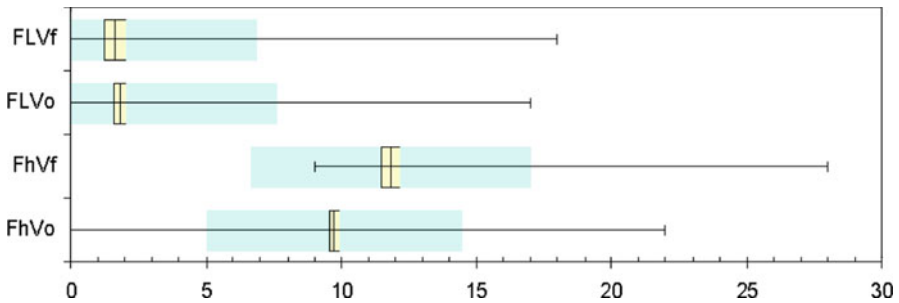**Fig. 6** Results for average vessel costs (€)

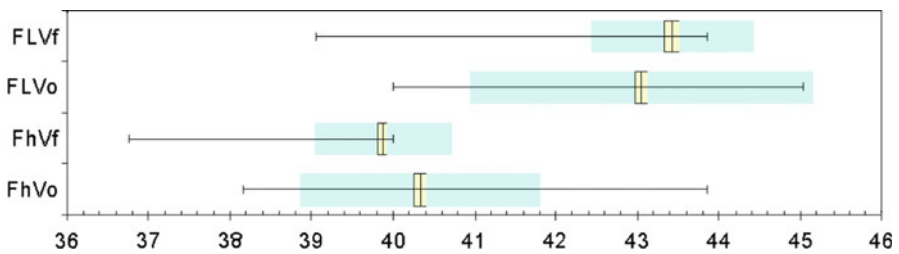**Fig. 7** Results for average vessel delay (h)

Fig. 8 Results for average vessel speed (km/h)

Vessel cost differences are, therefore, fully explained by fuel cost savings. This is indicated in Fig. 8, which shows vessels under the central APS to be able to sail at lower speeds than under the decentral APS.

The average speed reported for the decentral APS is the *planned* vessel speed. Actually, vessels sailed at their *initial* speed, which is somewhat higher. The actual speed difference between vessels in the respective APS was about 0.59 km/h, or little more than a quarter knot. When such a small speed difference can be sustained for a sailing time of about 10 days, a fuel cost saving in the order of 10 k€ can nonetheless be realized, according to the fuel cost assumptions described previously.

The figures also reveal that under the central APS, vessel speeds are actively manipulated, as can be observed from the relatively high $\tilde{\sigma}$ for vessel speed in the $FLVo$ simulation run. The maximum realized speed is also larger in the $FLVo$ run, which indicates that a vessel speed-up was issued at least once.

Finally, we note that the cost spread appears to be significantly lower under the central APS (see Fig. 8). Apparently, the central APS trades operational dynamics for financial stability.

High fuel price The outcomes under a High Fuel Price scenario again indicate a small performance benefit of the central APS ($FHVo$) over the decentral APS ($FHVf$). The difference between average vessel prices amounts to about 6.0 k€ per vessel. Interestingly, the factor explaining the difference is here delay management, not fuel cost savings (see Figs. 7 and 8). Differences in call size costs were again negligible.

The relatively high delays under the decentral APS are a result of the initial speed selection rule, which trades fuel costs and delay costs. Under a High Fuel Price scenario, it turned out that the benefits of one period less delay would be outweighed by the extra fuel costs required to gain one period. Hence, under the decentral APS, all vessels selected and stayed at the minimum allowed vessel speed of 40 km/h (actual average speeds are sometimes lower, since the planning model discounts waiting time in the planned vessel speed). Figure 8 shows that under the central APS, the average speed also lingers near the minimum vessel speed. At the minimum vessel speed, reducing sailing speed will increase delay, but no longer reduce fuel costs. At times, the planning model under the central APS was therefore able to realize significant delay reductions with tactical speed increases. Whenever a single vessel's speed increase resulted in reduced delay for a whole block of vessels, vessel speed increasing proved beneficial, even in the presence of high fuel costs.

Again, it can be observed that under the central APS, the average vessel costs are not only lower, but also are less spread out than under the decentral APS.

8 Concluding remarks

We have presented an analytical approach to investigate the relationship between ship arrivals and terminal operations, by simulating ship arrivals and integrating this with an optimization routine that determines the optimal approach speed for the arriving vessels. We compared the impact of this centralized-information rich arrival planning strategy with a strategy where speed of the approaching vessels is not optimized. For the analysis, we looked at a generalized cost function that includes fuel costs and delay costs.

We demonstrated that for specific scenarios, a central APS can provide benefits compared to a decentral APS. The latter is a strategy encountered in current practice. For the particular scenarios analyzed, the potential gains are in the order of thousands of euros per vessel arrival. We found greater benefits for lower fuel prices. Furthermore, we observed qualitatively different outcomes under different fuel price scenarios.

For container terminals and shipping lines, this means that more coordination—in our case shipping lines depending on the terminal to give speed suggestions to approaching vessels—has the potential to generate real monetary and operational benefits for both parties. This is even the case, when the coordination results in vessels having to speed up to make their time slot at the terminal.

For terminal operators, managing congestion by coordination rather than by facility investment has the added benefit of flexibility. Coordination allows terminal operators to operate their facilities at a higher utilization level (for a given service level) and therewith enables them to increase returns on investments.

For shipping lines, coordination could be a viable instrument to manage delays in structurally congested ports. Our analysis can be extended in many ways. First, we have demonstrated the potential of coordination for a fictitious case only. Extending the result to business practice may require refinement of the cost and operation models presented in this work.

Second, our objective function only considered operational costs. However, a coordination strategy typically has cost components beyond operational costs. Coordination costs and transaction risks (Kumar and van Dissel 1996) may render a centrally coordinated APS unattractive in the presence of a 'cheating culture', despite the operational benefits. A principal-agent analysis of the relationship between shipping companies and terminals might shed some light on the optimal contractual arrangement between these two parties.

Third, in this work, we contrasted central and decentral coordination approaches. A promising alternative approach is agent-based coordination, based on information-rich decentralized information processing. Several indirect benefits may render such a strategy attractive from an overall business perspective. These include robustness, lower associated coordination costs and higher responsiveness to unexpected events. With some adjustment, the simulation model presented in this paper can be used as an analytical tool to assess whether such indirect benefits actually outweigh losses of planning performance relative to a centralized, 'optimal' strategy.

Finally, the scope of this work is limited to the coordination between shipping lines and a single terminal, being first port of call. Future work could extend the ideas presented in this paper to research the impact of coordination in a *cluster* of terminals, e.g. selected container terminals in the North-European Hamburg–Le Havre range. An interesting research question is whether the extra planning potential introduced by such a scope would justify the organizational costs of cluster-wide coordination.

Appendix: Key model parameters

This appendix summarizes the key parameters introduced in the text (see Tables 7, 8, 9 and 10).

Table 7 Simulation related parameters

Parameter	Value
Period length (π)	2 h
Warm-up length	500 periods
Replication length	1,000 periods
Replication count (per scenario)	10

Table 8 Vessel related parameters

Parameter	Value
Arrival process	Poisson
Average arrival rate	1 vessel per 24 h
Initial distance from terminal	10,000 km
Total call size distribution (σ_i)	Normal distribution, average and variance set to 3,000 moves
Maximum number of cranes supported (C_i)	5
Maximum vessel cruise speed (v_i^{mx})	50 km/h
Minimum vessel cruise speed (v_i^{mn})	40 km/h
Scheduled period of completion (τ_i)	Estimated port call completion time under nominal sailing (45 km/h) and handling conditions (4 STS cranes available)

Table 9 Terminal related parameters

Parameter	Value
Installed number of berthing slots (Q)	3
Installed number of STS cranes (C)	9
Single crane productivity (p_c)	30 moves/h

Table 10 Cost parameters

Parameter	Value
Delay costs	€ 8,332 per period delay
Marginal TFC effect (regular fuel price scenario)	€ 11,361 per period gained
Marginal TFC effect (low fuel price scenario)	€ 7,384 per period gained
Call gap costs	€ 200 per box

References

BunkerWorld (2009) Rotterdam. 30 day rolling average bunker prices. IFO380. http://www.bunkerworld.com/markets/prices/nl/rtm/30day. Accessed 17 August 2009

Daganzo C (1990) The productivity of multipurpose seaport terminals. Transp Sci 24(3):205

Drewry (2007) Container shipper insight 1st quarter 2007. Technical report. Drewry Shipping Consultants Limited

Hughes C (1996) Ship performance. LLP, London

Imai A, Nishimura E, Papadimitriou S (2001) The dynamic berth allocation problem for a container port. Transp Res B 35(4):401–417

Jacobs P, Lang N, Verbraeck A (2002) D-SOL; a distributed Java based discrete event simulation architecture. In: Yücesan E, Chen CH, Snowdon J, Charnes J (eds) Proceedings of the 2002 Winter Simulation Conference, pp 793–800

Kia M, Shayan E, Ghotb F (2002) Investigation of port capacity under a new approach by computer simulation. Comput Ind Eng 42(2–4):533–540

Kingman J (1970) Inequalities in the theory of queues. J R Stat Soc Ser B (Methodological) 32(1):102–110

Kumar K, van Dissel H (1996) Sustainable collaboration: managing conflict and cooperation in interorganizational systems. MIS Q 20(3):279–300

Law A, Kelton W (1982) Simulation modeling & analysis, 2nd edn. McGraw Hill, New York

MAN (2005) Propulsion trends in container vessels. Technical report. http://www.manbw.com/files/news/filesof4672/P9028.pdf. Accessed 17 August 2009

Moonen J, van de Rakt B, Miller I, van Nunen J, van Hillegersberg J (2005) Agent technology supports interorganizational planning in the port. Erim report series research in management. Erasmus Research Institute of Management (ERIM)

Nemhauser G, Wolsey A (1988) Integer and combinatorial optimization. Wiley-Interscience, New York

Notteboom T (2006) The time factor in liner shipping services. Marit Econ Logist 8(1):19–39

Radmilovich Z (1992) Ship-berth link as bulk queueing system in ports. J Waterw Port Coast Ocean Eng 118(5):474–495

Stahlbock R, Voß S (2008) Operations research at container terminals: a literature update. OR Spectrum 30:1–52

UNCTAD (2005) United nations conference on trade and development. Review of maritime transport. http://www.unctad.org/en/docs/rmt2005_en.pdf. Accessed 17 August 2009

Veenstra A, Lang N (2005) Simulating the impact of arrival patterns on container terminal performance. In: Proceedings of the first international conference on transportation logistics, 27–29 July 2005, Singapore

Zrnic D, Dragovic B, Radmilovic Z (1999) Anchorage-ship-berth link as multiple server queuing system. J Waterw Port Coast Ocean Eng 125(5):232–240