

On the effect of ship arrival processes on jetty and storage capacity

Eelco van Asperen Rommert Dekker Mark Polman Henk de Swaan Arons*
vanasperen@few.eur.nl dekker@few.eur.nl polman@few.eur.nl deswaanarons@few.eur.nl

PO Box 1738
Erasmus University Rotterdam
Rotterdam School of Economics
3000 DR Rotterdam, THE NETHERLANDS

Abstract

Ports provide jetty facilities for ships to load and unload their cargo. Jetty capacity is costly and therefore limited, causing delays for arriving ships. However, ship delays are also costly, so terminal operators attempt to minimize their number and duration. Here, simulation has proved to be a very suitable tool. However, in port simulation models, the impact of the arrival process of ships on the model outcomes tends to be underestimated. This report considers three arrival processes: stock-controlled, equidistant, and uncontrolled. We assess how their deployment in a port simulation model, based on data from a real case study, affects the efficiency of the loading and unloading process, making a case for careful modeling of arrival processes in port simulations. Uncontrolled, which is an assumed arrival process property in many client-oriented simulations, actually performs worst in terms of both ship delays and required storage capacity. Stock-controlled arrivals perform best with regard to large vessel delays and storage capacity. Additional control of the arrival process through the application of a priority scheme in processing ships further impacts efficiency in all three cases.

Keywords: Port Simulation, Discrete-Event Simulation, Arrival Process, Port Logistics.

Topic area: Maritime Transport and Ports.

1. Introduction

In this report we investigate the importance of arrival process modeling in a port simulation. This is done by measuring the impact of the selected arrival process for ships on the efficiency of the loading and unloading process. This study was performed using data from a confidential case study in the Port of Rotterdam. The tender of that case study provided detailed data on the types and numbers of ships to be handled per year, but did not specify their timing, hereafter referred to as the *arrival process*. The engineering firm responsible for the tender evidently did not realize its importance.

* Corresponding author.

The case study model was used to optimize and evaluate various scenarios for the jetty and tank layout for the loading and unloading process of raw materials and finished products. Due to unforeseen business events (including a takeover of the company), the plant was built six years later, and no feedback on how the results were used has been given.

The model used in this report focuses on the analysis of ship waiting statistics and stock fluctuations under different arrival processes. However, the basic outline is the same: central to both models are a jetty and accompanying tank farm facilities belonging to a new chemical plant in the Port of Rotterdam. Both the supply of raw materials and the export of finished products occur through ships loading and unloading at the jetty. Since disruptions in the plant's production process are very expensive, buffer stock is needed to allow for variations in ship arrivals and overseas exports through large ships.

In the case study two types of arrival processes were considered. The first type are the so-called stock-controlled arrivals, i.e., ship arrivals are scheduled in such a way, that a base stock level is maintained in the tanks. The second type of arrival process is based on equidistant arrivals per ship type. A third kind of arrival process was not considered previously: an uncontrolled process, derived from a Poisson process. Furthermore, within each arrival process type, a further distinction can be made between prioritized and non-prioritized queues before the jetty's mooring points. In this report, all resulting arrival processes will be compared.

With respect to the original case study, some simplifications apply. For reasons of confidentiality, the diversity of ships has been skewed down, and their numbers modified. Also, details concerning tank operation, tank farm layout, and inland transport have been abstracted from. Still, the resulting model is general enough to draw conclusions applicable to many jetty simulation studies.

After a brief literature review in Section 2 we continue in Section 3 with a detailed discussion on the loading and unloading process: the layout of the jetty where ships unload raw materials or load finished products, the factory which converts raw materials into products, the tanks that hold raw materials or finished products, and the arrival of ships. We discuss the various arrival processes in more detail in Section 4. The implementation model is the subject of Section 5, the experiments carried out with it and their results are discussed in Section 6, and the conclusions are presented in Section 7.

2. A Literature Review

Little has been published on the simulation of port facilities, apart from some very scattered material. There is a nice book edited by Van Nunen and Verspui (Nunen and Verspui 1999) on simulation and logistics in the port, but it is in Dutch only. We briefly recapitulate the literature review on jetty design from Dekker (Dekker 1999) in that volume. Well-known to insiders are the reports from UNCTAD (UNCTAD 1978) on the design of jetties. They report results from both queuing theory and simulation applied to the capacity of jetties. However, the reports are difficult to obtain and they give yardsticks for simple cases only. Other papers more or less describe particular simulation studies, without trying to generalize their results. We like to mention (Philips 1976) and (Andrews 1996) who describe the planning of a crude-oil terminal, (Baunach et al 1985), who deal with a coal terminal, (Heyden and Ottjes 1985), (Ottjes et al. 1992) and (Ottjes et al. 1994), who deal with the set-up of the simulation programs for terminals. None of these papers however, deals explicitly with the arrival process.

3. The Model

The model comprises the arrivals of ships, a jetty with a number of mooring points, storage tanks and a factory. These are briefly described in this section. Figure 1 provides a schematic outline of the model as a whole.

3.1. The Jetty

Central in the loading and unloading facility to be simulated are a number of mooring points. In this case there are four mooring points (mooring point 1 to 4) in a T-shaped layout (Figure 2). They differ in a number of aspects. One of these is the length of the ships that the mooring point can handle. Mooring points 1 and 2 are suited to long ships; mooring points 3 and 4 can handle only short ships (see also Table 1).




Figure 1: A schematic outline of the loading and unloading process, including jetty, tanks and factory




Figure 2: A jetty layout

Pipes facilitate the transport of chemicals to and from the tanks, but cost considerations are a limiting factor on their construction. Therefore, each mooring point can load and unload a subset of the chemicals used in this system, where A and B denote raw materials,

whereas C and D denote finished products. For example, mooring point 1 can handle A, B and C, whereas mooring point 2 can only handle products C and D.

3.2. Raw Materials and Finished Products

After being unloaded, raw materials are stored in tanks A and B, from where they are withdrawn by the factory. Finished products are transferred to tanks C and D, to be loaded into ships.

3.3. Tanks and Stocks

Tanks can be used for just one type of raw material or finished product. The transfer of products from ships into tanks, from tanks to the factory, and from the factory into the tanks are continuous processes. In reality, there are several restrictions that affect actual tank operations, e.g. no simultaneous pumping and running into and out of a tank. We ignore these restrictions, because they do not affect the comparison between the arrival processes. On the same grounds, we allow for the stocks to take on any value (including negative values), and neglect ship delays because of stock outs or lack of ullage (available tank space).

3.4. Ships

Ships (ocean-going vessels, short-sea shipping vessels and inland barges) unload raw materials or load finished products. Each ship has four defining properties relevant for our model:

- size (tonnage);
- length (a distinction between long or short suffices);
- product (each ship handles just one specific type of cargo);
- (un)loading time (in hours).

Loading or unloading can only be done at a mooring point that can handle a ship's length and product. When a ship has arrived in the port, a suitable mooring point is selected according to specified rules, which are discussed below.

4. The Arrival Process

In many simulation studies, the assumption is made that arrivals in client-oriented processes cannot be controlled. Consequently, simulation languages and environments tend to offer Poisson as a first-choice option for the specification of arrival processes. As mentioned above, we have looked at three scenarios:

1. Stock-controlled arrivals;
2. Equidistant arrivals;
3. Uncontrolled arrivals.

4.1. Stock-Controlled Arrivals

Stock-controlled arrivals occur in a situation where arrivals can be planned by the factory. The factory's aim is to maintain a target base stock level in the tanks. In our model, this is implemented as follows. For the loading process, it implies that the arrival time of the next ship is planned to coincide with the moment that, through production, there is sufficient stock in the tank to load the ship without dropping below base stock level. In this calculation, the parameters are the loading time of the present ship, the cargo capacity and loading time of the next ship, and the production capacity of the factory. Setting the appropriate

base stock level for a tank involves an estimation of the tendency of ships to arrive ahead of schedule (see below), this being the only threat to maintaining base stock level.

For the unloading process, maintaining base stock levels in the raw materials tanks is achieved by planning the next ship's arrival to coincide with the moment that, through extraction of raw material during production, base stock level will be reached. In this calculation, the parameters are the cargo capacity of the present ship, and the rate at which the factory extracts material from the tank. Here, the danger of stock dropping below base stock level comes from ships arriving late (or from ships unable to instantly find an unoccupied mooring point).

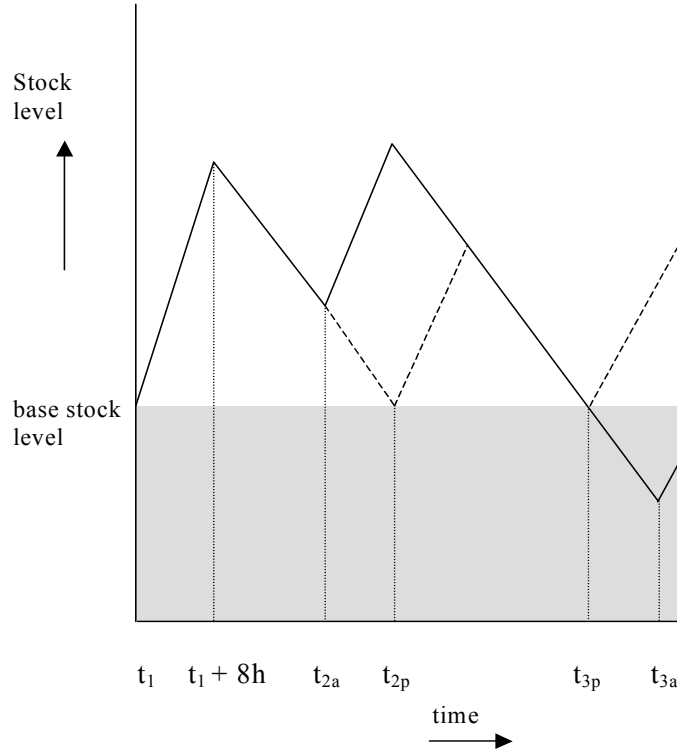


Figure 3: Stock level fluctuations in raw material tank with stock-controlled arrivals.

To illustrate the above, Figure 3 shows stock level fluctuations in raw material tank A over time with stock-controlled arrivals. At time t_1 , when the tank contents is at base stock level, a 1000 ton barge arrives, unloading its cargo into the tank over an 8 hour period. This implies that 8 hours later, the tank will contain an extra 1000 tons of raw material, minus the volume of raw material pumped out of the tank by the factory. After this point, the tank's contents will steadily decrease back to base stock level. The next ship's arrival is planned to coincide with this moment t_{2p} ('p' for 'planned'). However, this ship could arrive ahead of time, for example at t_{2a} ('a' for 'actual'), causing stock to start rising again before reaching base level. The dotted line shows how stock level would develop if all ships arrived exactly as planned. The solid line shows actual stock level development. After the last ship's early arrival, the next ship is again scheduled to arrive when stock reaches base level (t_{3p}). However, it arrives late at time t_{3a} , causing stock to drop below base level.

4.2. Equidistant Arrivals

Equidistant arrivals model a situation in which loading and unloading ships arrive at regular intervals. This regularity could be the consequence of year-based contracts specifying, for example, annual amounts of raw product to be delivered in equal batches every n weeks.

In our model, equidistant arrivals imply that arrivals of ships within a ship type are assumed to be evenly spread over the year. For example, per year, twelve vessels carrying 6000 ton of product B arrive (see Table 1). With equidistant arrivals, this means a 1-month inter-arrival period between such ships.

4.3. Uncontrolled Arrivals

With both stock-controlled and equidistant arrivals, the assumption is that there is some sort of control over the times at which ships arrive. If this is not the case, opting for a Poisson process is the logical choice. In our model, however, the number of arrivals per year within each ship type is fixed. We therefore assume a uniform distribution of these arrivals per ship type over the year, which yields a similar arrival process.

4.4. Disturbances to Expected Arrival Times

The stock-controlled and equidistant arrival processes actually yield a series of *expected times of arrival* (ETAs). However, in reality ships will seldom meet this schedule. For this reason disturbances to the ETAs are generated, modeling early and late arrivals resulting in the actual time of arrival (ATA) of each ship. Figure 3 shows the distribution of disturbances to the ETA of a ship as used in all of our experiments: all ATAs are within a margin of twelve hours before and twelve hours after the corresponding ETA. Eighty percent of these are within a margin of 2 hours before and 2 hours after the corresponding ETA (these values were set together with shipping experts).




Figure 4: Distribution of disturbances to expected times of arrival.

More specific, sea-going vessels arrive mutually independent according to the following statistical distribution around their ETA. If x is the deviation of the ATA in hours from the ETA, then (U stands for the uniform distribution):

$$\begin{aligned} x &= U(-12, -2) \quad \text{with } p = 0.1; \\ &U(-2, +2) \quad \text{with } p = 0.8; \\ &U(2, 12) \quad \text{with } p = 0.1. \end{aligned}$$

4.5. Ship types and arrival rates

In order to be able to compare model outcomes over multiple years and among multiple arrival processes, the annual total number of arriving ships of each type is fixed, and identical for stock-controlled, equidistant, and uncontrolled arrivals. Table 1 shows which ship types are distinguished, and how many arrive per year. For example, every year, a total of 14 short vessels arrive carrying 4000 tons of product B, with a loading time of 26 hours (for the meaning of the priority column, see below).

Table 1: Ship types, properties, and arrival rates

Ship type	barge/vessel	Size (tons)	Length	Product	Loading time (hours)	Ships per year	Priority	Tons per year
1	barge	1,500	short	A	8	196	low	294,000
2	vessel	2,000	short	A	8	48	low	96,000
3	vessel	4,000	short	A	20	80	low	320,000
4	vessel	6,000	long	A	26	60	high	360,000
								1,070,000
5	barge	1,000	short	B	10	38	low	38,000
6	vessel	2,000	short	B	11	161	low	322,000
7	vessel	4,000	short	B	26	14	low	56,000
8	vessel	6,000	short	B	26	12	low	72,000
								488,000
9	barge	1,000	short	C	10	180	low	180,000
10	vessel	2,000	long	C	14	126	high	252,000
								432,000
11	barge	1,500	short	D	8	134	low	201,000
12	vessel	2,000	short	D	8	300	low	600,000
13	vessel	10,000	long	D	44	14	high	140,000
14	vessel	20,000	long	D	56	8	high	160,000
								1,101,000

For each product/cargo type, the number of ships carrying it is chosen so that the total amount of cargo transported matches the factory's capacity. For instance, per year, the factory processes 1,070,000 tons of raw material A. Therefore, the total cargo capacity of ships carrying product A into the port needs to be 1,070,000 tons, which can be verified from the table.

The implication of this is that among simulation runs, only the mutual order of arriving ships and their interarrival times are variable. Thus comparisons regarding port efficiency among arrival processes are kept clean (i.e. devoid of other circumstantial factors such as random fluctuations in production.)

4.6. Priorities

In reality, the arrival of a ship is known, sometimes days beforehand, to the plant. This can be used in a mooring point allocation system based on priorities. The general idea is to incorporate all ships within an n -hour horizon into the choice of mooring point for an incoming ship, in order to reduce costs induced by waiting for available mooring points, given the fact that for some ship types, waiting is more expensive than for others (e.g. dependent on the type of cargo, the capacity, or the crew size.)

This general idea can be implemented in many ways.

4.6.1 A simple scheme

In this report, we use a simple priority scheme, with two priority classes (high and low), in which long ships get high priority, and short ones get low priority. The allocation of a mooring point to a ship can now proceed as follows. A high-priority ship entering the port is in principle assigned to a free mooring point suitable for its cargo type and length. If all suitable mooring points are occupied, the ship is placed in a queue before the mooring point with the smallest workload^{*}, or, in case of equal workloads, the shortest queue so far.

For low-priority ships, the situation is similar, apart from an additional condition. To explain this, let s be a low-priority ship, let t be the current time, let $Wi(t)$ be the workload of mooring point i at time t , and let $Di(s)$ be the time that ship s needs if serviced at mooring point i . Then mooring point i is considered reserved if a high-priority ship arriving within a 48-hour horizon will need mooring point i between t and $t + Wi(t) + Di(s)$. If this is the case, s is not assigned to i , or enqueued before i . Note, that the shorter mooring points at the jetty are never reserved by high-priority ships, since all high-priority ships are too long for these mooring points. Hence, a low-priority ship will always either be assigned to a mooring point directly or placed in a queue before one.

In the presentation of the results in Section 6, we will make a distinction between model outcomes with and without priority-based mooring point allocation, so that the impact of incorporating such allocation is clearly visible.

4.6.2 An advanced scheme

A more sophisticated choice would be a cost-based approach requiring complex calculations of the waiting costs of the various types of ships. Also an enumeration algorithm may be applied: select the optimal allocation schedule of all possible schedules given a look-ahead time window and a cost function. In general, this is a time-consuming approach. The various scenarios are currently investigated in more detail with promising results, but in this paper we restrict ourselves to the cases of equal and simple priorities.

5. The Implementation Model

Based on what is outlined in Section 3, a simulation model has been implemented in Enterprise Dynamics (Enterprise Dynamics 2003), a simulation package for discrete-event simulation. This implementation model, see Figure 5, comprises various types of atoms, the Enterprise Dynamics equivalents of objects. Some of the atoms implement the simula-

* The workload of a mooring point at instant t is defined as the total time from t that the mooring point will be occupied by the ship currently using it, and the ships currently in the queue before it.

tion's logic, others hold the simulation data (tables), define the types of experiments or provide the desired output (e.g., graphs).

The figure shows the number of ships that have entered the port thus far (262). Nine ships are on their way to the jetty. All mooring points are occupied. Their utilization up until now has been 61.3%, 47.1%, 63.1% and 72.8%. Queues 1, 2, and 3 are empty, whereas Queue 4 contains one waiting ship. The actual contents of the tanks are 3735, 3781, 2114 and 1986 tons, respectively. The total number of ships that have been processed is 248, which, added to the nine approaching ships and the 5 at the mooring points, matches the number of ships generated thus far.

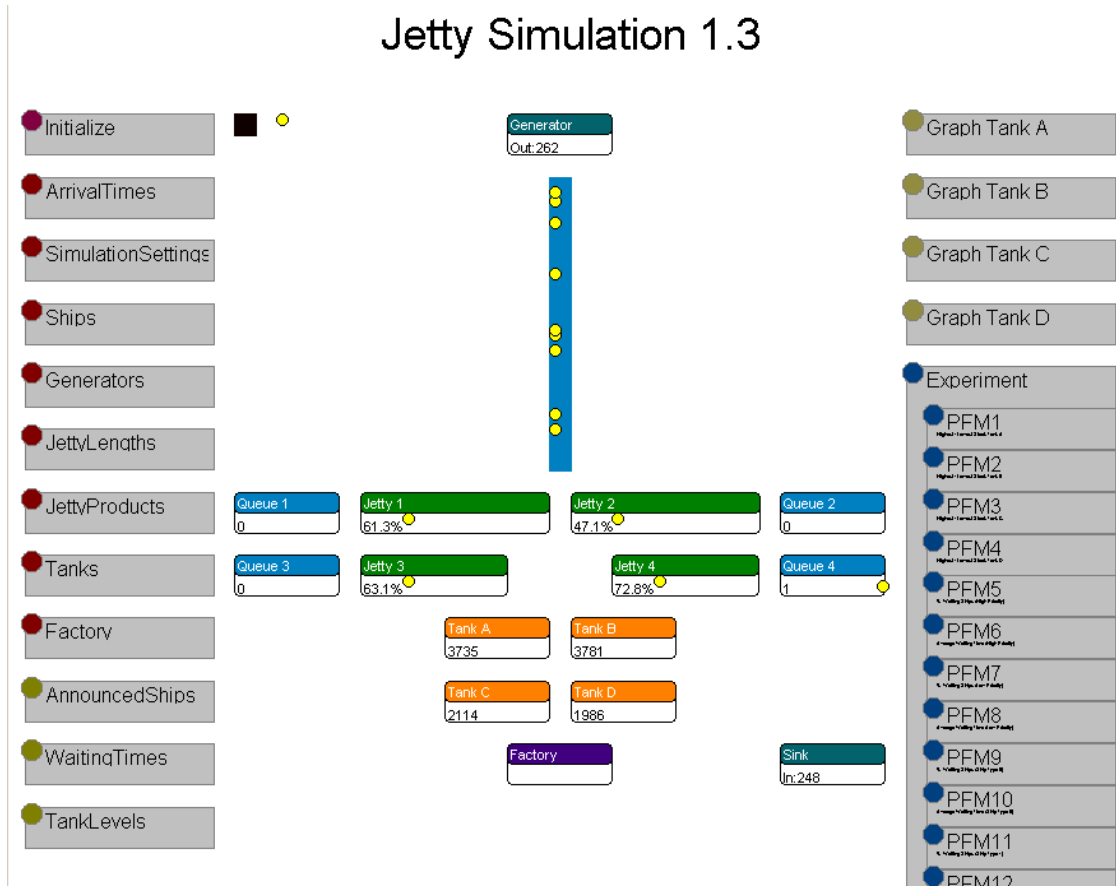


Figure 5: Implementation of the simulation model

5.1. Logic

One atom (Generator) is responsible for generating ship arrivals.

After arrival a ship proceeds along the atom ArrivalRoute (the vertical atom in the middle) to one of the four mooring points that suits its length and cargo type (see Section 3.4). If all suitable mooring points are occupied, the ship will wait (see also paragraph 4.6) in one of the queues (Queue 1, 2, 3 or 4).

Raw materials are unloaded and transferred to either Tank A or B, from which they are withdrawn by the Factory atom. The factory stores finished products in Tank C or D, from which they are withdrawn to be loaded into ships. After loading or unloading the ships leave the system.

It is worth mentioning that the stock of the tanks is not modeled as a continuous variable, but is updated at discrete intervals (every two hours). As stated before, for this study we assume that the process is not limited by the capacities of the tanks. As a consequence, we can model storage by using tanks with unlimited capacity and with the possibility to contain negative stock. This simplification does not affect the simulation's objective.

The arrival and queue atoms contain specific programming code refining their default (i.e. as defined in Enterprise Dynamics) logic. The others are custom developed to perform dedicated tasks. Finally, the atom Initialize contains code to be executed prior to each simulation run.

5.2. Data

The remaining atoms on the left side represent tables. The top seven of these are filled from text files at the beginning of each run, and provide data for the simulation process. They contain the following data:

Table	Contains
ArrivalTimes	Data concerning the expected arrival times (ETA), and the actual arrival times (ATA, expected arrival times disturbed according to the distribution function outlined in Fig. 4) for each category of ship for a number of years.
SimulationSettings	Some initialization data (e.g. the parameters of the code used to disturb expected times of arrival.)
Ships	Specific ship data such as type (barge or vessel), size in tons, length, loading time, the number of ships arriving annually, and so on.
Generators	The generator to be used for each ship type.
JettyLengths	The lengths of the mooring points.
JettyProducts	The products that can be handled per mooring point.
Tanks	Base stock levels of the various tanks.
Factory	Yearly amounts of raw material processed and finished product produced by the plant.

It should be noted that the text file acting as the source for the ArrivalTimes table is filled by running an external Java program generating arrival times for the arrival process of choice. Seed, number of years, and other aspects are parameters to this program, and should be supplied by its user.

The bottom three tables on the left are filled with data during simulation runs. They contain the following data:

Table	Contains
AnnouncedShips	Intermediate data used in the allocation of a ship to a mooring point.
WaitingTimes	Waiting times statistics for all ship types.
TankLevel	Stock level movements for each tank.

An important reason for using tables is that they can easily be used to import data from external resources (e.g. a text file or csv file) into the model, and to export the simulation results for later analysis. External files as a source of input data and storage mechanism for simulation results are easy to maintain and provide more flexibility (e.g., in modeling the arrival processes and in converting simulation data into graphs).

The Graph atoms on the right side (Graph Tank A to B) convert simulation results into the necessary graphs. The other atom (Experiment) on the right allows the user to define general preferences of a simulation experiment. In this case the Experiment atom also contains a number of PFM atoms (Performance Measure), each defining one output variable of interest. The atoms PFM1 till PFM4 provide the differences between the highest and lowest stock data of the tanks; PFM5 provides the percentage of the high priority waiting ships and PFM6 their average waiting times; PFM7 and PFM8 do the same for the low-priority ships. The other PFMs (most of which are not visible in the figure) are used to collect similar data per individual ship type.

6. Experiments and Results

The implementation of the model outlined in the previous section has been used to carry out experiments. While it is capable of generating results on a variety of topics, and on many levels of detail, we focus on the ones relevant to our objective: assessing the impact of using different arrival processes on stock levels and ships' waiting times. All in all, a total of six ten-year simulation runs are conducted:

- Stock-controlled arrivals (with and without a priority scheme);
- Equidistant arrivals per ship type (with and without a priority scheme);
- Uncontrolled arrivals (with and without a priority scheme).

Each run starts in a steady-state situation, with the tanks filled to base stock level. This eliminates the need for a warm-up period, which has consequently been omitted.

Tables 2 through 8 show the relevant simulation outcomes. Tables 2, 3 and 4 contain waiting statistics for ships, one table for each arrival process *without* the priority scheme outlined in Section 4.6, each table in turn divided into separate columns for high and low-priority ships*. Table 5 reports on the maximum and minimum stock levels reached for each of these arrival processes, both in raw material and finished product tanks. Tables 6, 7, and 8 show the differences for each arrival process between using and not using the priority scheme for mooring point allocation.

6.1. Waiting times

From Tables 2, 3 and 4, it can be observed that the choice for a particular arrival process has significant impact on the number of waiting ships and the number of hours spent waiting by these ships. With uncontrolled arrivals both numbers are higher than those observed with equidistant and stock-controlled arrivals. This holds for both high and low-priority ships.

Clearly, the lack of a mechanism to keep ships apart, whether it be equidistant or stock-controlled arrival planning, allows for clusters of ships arriving within a small time frame, causing queues.

Tables 2 and 3 reveal a noticeable difference between the outcomes of equidistant arrivals and stock-controlled arrivals. For both low and high-priority ships, the stock-controlled arrival process 'outperforms' the equidistant arrival process.

The explanation for this is manifold. For one, stock-controlled arrivals are more efficient overall since they tend to keep ships of identical cargo types apart, whereas equidis-

* The distinction between high and low-priority ships is made here to facilitate a comparison with the results of simulation runs that *do* include a priority scheme.

tant arrivals keep ships of identical types apart. With multiple ship types per cargo type this is an advantage.

Furthermore, simulation-specific factors have to be taken into account. Consider the arrival rates of the individual ship types. Here, care has been taken to avoid introducing unrealistic queuing situations. With equidistant arrivals, for example, special measurements seek to prevent the scheduling of arrivals for multiple ship types in such a way, that they all coincide several times a year. Not all such mechanisms are that obvious though, especially when related to another simulation-specific aspect: the jetty layout. The combined effects of these factors are still subject to further research.

However, the observed differences in waiting time statistics among arrival processes, whatever their causing factors, clearly demonstrate the need for careful arrival process modeling, which is this article's primary objective. Obviously, arrival process modeling requires a careful look at the real situation, involving expert input on many subjects. Only then are simulation results valid, and can they be used in corporate decision-making. Alternatively stated, providing only the numerical data from Table 1, and simply assuming an uncontrolled process, is not sufficient, rendering any subsequent decision (for example on expensive alternative jetty layout to reduce waiting times) ill founded.

Table 2: Ship statistics, Stock-controlled arrivals without priorities
(means over a 10-year period)

	Ship Priority			
	High		Low	
	Mean	St.dev.	Mean	St.dev.
Percentage of ships that had to wait	21.1%	3.7%	12.0%	1.0%
Average waiting time of ships that had to wait (hours)	7.9	1.1	3.5	0.2

Table 3: Ship statistics, Equidistant arrivals per ship type without priorities
(means over a 10-year period)

	Ship Priority			
	High		Low	
	Mean	St.dev.	Mean	St.dev.
Percentage of ships that had to wait	34.7%	1.8%	23.5%	0.8%
Average waiting time of ships that had to wait (hours)	9.5	0.6	6.2	0.2

Table 4: Ship statistics, Uncontrolled arrivals without priorities
(means over a 10-year period)

	Ship Priority			
	High		Low	
	Mean	St.dev.	Mean	St.dev.
Percentage of ships that had to wait	45.7%	2.1%	35.2%	2.0%
Average waiting time of ships that had to wait (hours)	12.3	1.8	7.5	0.9

6.2. Stock Levels

Table 5 shows 10-year stock level statistics in terms of the difference between minimum and maximum levels reached. As could be expected, stock fluctuations are smallest with stock-controlled arrivals, whereas uncontrolled arrivals allow for the largest. Also, with equidistant arrivals, considerable fluctuations are observed.

Table 5: Stock levels ranges, Stock-controlled, Equidistant per ship type, and Uncontrolled arrivals without priorities (means in tons over a 10-year period)

	Tank							
	A		B		C		D	
	Mean	St.dev.	Mean	St.dev.	Mean	St.dev.	Mean	St.dev.
Stock-controlled	6970	468	5890	294	3011	320	15982	578
Equidistant	10756	273	11245	312	3381	283	27474	574
Uncontrolled	74396	18333	48058	11789	32045	9112	89177	15112

Figure 5 shows example stock behavior over time for product D over a one-year* period (notice that the scale of the vertical axis varies). Figures 5 (left and middle) show that fluctuations are such, that the initial stock level for product D (2000 tons) does not suffice to avoid stock outs. Figure 5 (right) clearly shows the typical stock fluctuation pattern for stock-controlled arrivals. Peak levels are reached whenever large ships are scheduled to arrive for loading. In fact, the largest available vessel (see Table 1) comes in to load product D eight times a year, which explains the eight peaks in the Figure. Notice that in the case of product D, stock fluctuation is almost completely determined by the size of this large vessel, which makes it easy to determine the required tank capacity.

So, again, the choice of arrival process is an important factor in simulation outcomes. For example, should the simulation be part of a cost-benefit analysis to the acquisition of additional tankage, then its results are of no value without realistic arrival process modeling.

6.3. The effect of using a priority scheme

In section 4.6 it was explained that a priority scheme is expected to reduce the waiting costs of high-priority ships. A simple priority scheme was considered with two priority classes (high and low), where long ships get high priority, and short ones low priority.

Tables 6, 7, and 8 show ship waiting statistics over a ten-year simulation period for each arrival process, both with and without (copied from Tables 2, 3, and 4) a priority scheme. Standard deviations have been omitted for brevity.

* As stated before, arrivals are aligned with production in such a way, that stock does not structurally grow or shrink over a one-year period. Any difference between stock levels at the start or the end of a year are due to ships still being loaded and unloaded at the end.




Figure 5: Level of Tank D during one year with uncontrolled (top), equidistant (middle) and stock-controlled process (bottom)

In all cases, applying priorities indeed reduces the percentage of high-priority ships, while increasing the percentage of low-priority ships waiting. All waiting time means go up, for which there are, again, multiple causing factors. One seemingly obvious mechanism is that high-priority ships are now very rarely blocked from suitable mooring points by low-priority ships. Hence, if a high-priority ship has to wait, it is probably for another high-priority ship, which takes longer to (un)load, causing longer delays.

The question as to whether total waiting costs are reduced by incorporating priorities, or to what extent, depends on how much more expensive an idle high-priority ship is over a low-priority ship. The tender of the original case study did not provide a cost function.

Table 6: Ship statistics, Stock-controlled arrivals, priority scheme vs. no priority scheme (means over a 10 year period)

	Ship Priority			
	High		Low	
	No priority scheme	Priority scheme	No priority scheme	Priority scheme
Percentage of ships that had to wait	21.1%	8.5%	12.0%	14.2%
Average waiting time of ships that had to wait (hours)	7.9	10.0	3.5	3.8

Table 7: Ship statistics, Equidistant arrivals per ship type, priority scheme vs. no priority scheme (means over a 10 year period)

	Ship Priority			
	High		Low	
	No priority scheme	Priority scheme	No priority scheme	Priority scheme
Percentage of ships that had to wait	34.7%	9.2%	23.5%	28.7%
Average waiting time of ships that had to wait (hours)	9.5	9.8	6.2	7.2

Table 8: Ship statistics, Uncontrolled arrivals, priority scheme vs. no priority scheme (means over a 10 year period)

	Ship Priority			
	High		Low	
	No priority scheme	Priority scheme	No priority scheme	Priority scheme
Percentage of ships that had to wait	45.7%	18.3%	35.2%	40.5%
Average waiting time of ships that had to wait (hours)	12.3	14.6	7.5	9.4

7. Conclusions

The importance of careful arrival process modeling is clearly demonstrated in this report. Model outcomes over various arrival processes vary significantly, e.g. the uncontrolled process has by far the worst performance of the three processes discussed, both in terms of waiting times and in terms of the required storage capacity, whereas the stock-controlled process performs best overall. Although these results were obtained in a specific case, we think that they are general enough to be appropriate for many port and jetty simulation studies. As soon as there is some sort of control over arrivals, it should be incorporated in the model.

Acknowledgements

The authors would like to thank the students Ludo Waltman, Nees Jan van Eck, Arthur Oink, and Gerard Seedorf for their solid implementation of the simulation model in Enterprise Dynamics and for carrying out the experiments. They are also grateful to Stef Kurstjens for his work in the original case study.

Web

More information on this study can be found on the website:

<http://www.few.eur.nl/few/research/eurfew21/m&s/article/jetty/>.

References

- Andrews, S., F.H. Murphy, Wang X.P., and Welch S., 1996. Modeling crude oil lightering in Delaware Bay. *Interfaces* 26 (6) 68-78.
- Asperen, E. van., Dekker, R., Polman M., Swaan Arons, H. de, 2003. Modeling Ship Arrivals in Ports. In: Chick, S., Sánchez, P.J., Ferrin, D., and Morrice, D.J. (Eds.), *Proceedings of the Winter Simulation Conference 2003*, New Orleans.
- Baunach, G.R., Wibberley E.S., Wood B.R., 1985. Simulation of a coal transshipment terminal: Batam Island, Indonesia, *Math. Comp. Simul.* 27 115-120.
- Dekker R. 1999. Simulation of jetty and storage activities for oil and chemicals. In: Nunen, J. van, Verspui, L., (Eds.), *SimLog, Simulation and logistics in the harbor*, Eburon Delft, Netherlands, 105-116 (in Dutch).
- Enterprise Dynamics. 2003. Documentation material. Incontrol Enterprise Dynamics. (Downloadable from website <http://www.enterprisedynamics.com>).
- Heyden, W.P.A. van der, Ottjes, J.A., 1985. A decision support system for the planning of the workload on a grain terminal. *Decision Support Systems*, 1, 293-297.
- Nunen, J. van, Verspui, L., 1999. *SimLog, Simulation and logistics in the harbor*, Eburon Delft, Netherlands (in Dutch).
- Philips, O.O., 1976. Optimization models for a crude oil storage export system. Ph.D. thesis. Penn. State University.
- Ottjes, J.A. 1992. *Modelvorming en simulatie van logistieke systemen*. Manual i76C, Delft University of Technology (in Dutch).
- Ottjes, J.A, Hengst, S. and Tutuarima, W.H., 1994 A simulation model of a sailing container terminal service in the port of Rotterdam. In: Guasch and Huber (Eds.), *Proceedings of the 1994 Conference on Modelling and Simulation*, 876-880.
- UNCTAD. 1978. *Port development: a handbook for planners in developing countries*. New York UN, TD/B/C.U/175, chapter II: the break-bulk berth group, 108-128.