Proceedings of the 2003 Winter Simulation Conference
S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds.

MODELING SHIP ARRIVALS IN PORTS

Eelco van Asperen
Rommert Dekker
Mark Polman
Henk de Swaan Arons

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

ABSTRACT

Ports provide jetty facilities for ships to load and unload their cargo. Since ship delays are costly, 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 article considers three arrival processes: stock-controlled, equidistant per ship type, and Poisson. 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. Poisson, which is the chosen arrival process 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 ship delays and required storage capacity.

1 INTRODUCTION

In this article we investigate the impact of a number of different arrival processes for ships on the efficiency of the loading and unloading process in a harbor. This study was performed using some 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.

In the original case study, a simulation 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 tankfarm 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. Given a base stock level of a raw material or product, the time to fill up or empty the tanks and the tonnage of the next arriving ship, a planned arrival time of this ship is calculated. The second type of arrival process is based on equidistant arrivals of ships per ship type. In this article we add a third kind of arrival process that was not considered in the original case study: a Poisson process.

The subsequent arrival times are actually expected arrival times but ships will seldom meet this schedule. For this reason early and late arrivals are modeled by disturbances generated for the estimated times of arrival (ETA) resulting in the actual times of arrival (ATA).

The three arrival processes will be compared in this article, using data from the original case study. 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. Still, the resulting model is general enough to draw conclusions applicable to many jetty simulation studies.

After a 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. In this article we focus on the impact of various types of arrival processes on the efficiency of the loading and unloading process. We therefore 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 (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 (1999) in that volume. Well-known to insiders are the reports from UNCTAD (1978) on the design of jetties. They report results from both queuing theory and simulation applied to the capacity of jetties. The reports are however difficult to obtain and they give yardsticks for simple cases only. The other papers more or less describe that they have done a simulation study, without trying to generalize their results. We like to mention Philips (1976) and Andrews et al (1996), who describe the planning of a crude-oil terminal, Baunach et al (1985), who deal with a coal terminal, Van der Heyden and Ottjes (1985), Ottjes (1992), and Ottjes et al (1994), who deal with the setup of the simulation programs for terminals. None of the 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.

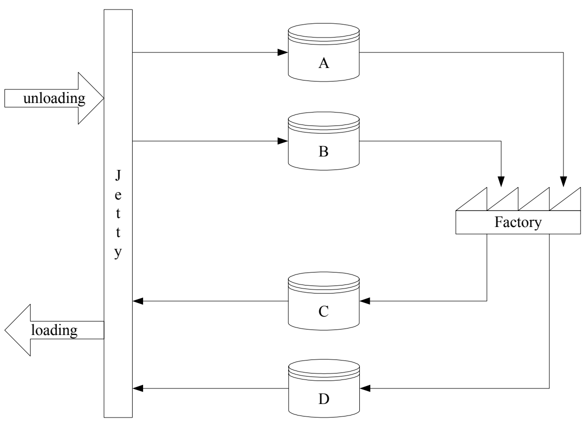

Figure 1: A Schematic Outline of the Loading and Unloading Process, with Jetty, Tanks and Factory

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. The mooring points also differ in their ability to load and/or unload different materials (raw materials A or B, and finished products C or D). For example, mooring point 1 can handle A, B and C, and mooring point 2 can only handle products C and D. In the original case study, several jetty layours were compared: here we consider just one layout.

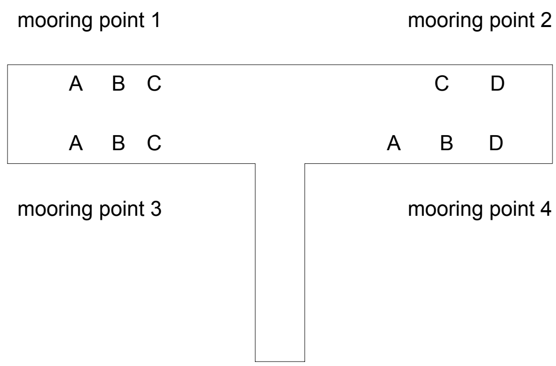

Figure 2: Jetty Layout

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 in our model, because they do not affect the comparison between the arrival processes (the original case study did model these restrictions). The same holds for stocks; for simplicity we allow 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 (sea-going vessels, short-sea shipping vessels, and inland barges) unload raw materials or load finished products. Each ship has four defining properties: the physical length (short or long), the cargo capacity (tonnage), the type of cargo it can handle (each ship can handle just one specific type of cargo), and (un)loading time (in hours).

Ships are categorized based on the type of cargo they can carry. 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 (e.g. if several suitable mooring points are available, select one of these randomly), including a priority scheme (see Section 5.3).

4 THE ARRIVAL PROCESS

Arrivals in client-oriented processes are often assumed to be based on a Poisson process. The underlying assumption is that the process cannot be controlled. Simulation languages and environments tend to offer this 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. Arrivals according to a Poisson process.

Stock-controlled arrivals aim at maintaining a target base stock level of raw material and finished product in the tanks. For the loading process, this 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.

For each product, the order in which the ships of different types arrive can be determined in several ways. Here we made a random selection with stratification for each ship type to make sure that a fixed number of ships arrive per year.

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.

Both strategies 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: 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.)

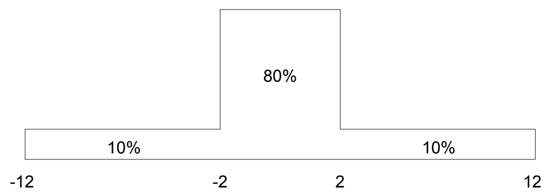

Figure 3: Distribution of Disturbances to Expected Times of Arrival

The third arrival process considered in this article is a straightforward Poisson process: within each cargo type, ships arrive with exponentially distributed interarrival times.

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. This is further explained in Section 5.3.

5 THE IMPLEMENTATION MODEL

The simulation model has been implemented in Enterprise Dynamics, a simulation package for discrete-event simulation. The implementation model, see Figure 4, comprises various types of atoms (atoms are the Enterprise Dynamics equivalents of objects). Some of the atoms implement the simulation's logic, others hold the simulation's data (tables), define the types of experiments or store the output (e.g., graphs).

5.1 The Logic

The Generator atom is responsible for generating ship arrivals. Upon arrival a ship proceeds along the atom Arrival Route (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 in one of the queues (Queue 1, 2, 3 or 4).

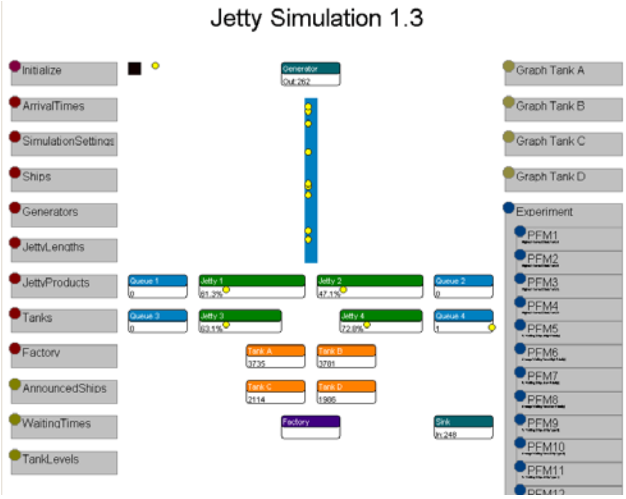

Figure 4: Implementation of the Simulation Model

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 and D, from which they are withdrawn to be loaded into ships. After loading or unloading the ships leave the system via the Sink atom (Sink is the standard Enterprise Dynamics atom at which entities leave a simulation model).

It is worthwhile 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.

Global parameters for the simulation experiment are set using the Experiment atom. This atom contains several PFM atoms (Performance Measure), each defining one output variable of interest. (Most of the PFM atoms are not shown in the figure.) For example, PFM1 till PFM4 provide the differences between the highest and lowest stock data of the tanks.

5.2 Data

The remaining atoms are mainly tables providing data for the simulation process. An important reason for using tables is that they can easily import input data from an external resource such as a spreadsheet into the model, and export the simulation results to another spreadsheet for later analysis. Spreadsheets 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).

An important table is ArrivalTimes, which, for each category of ships, provides the ETA and the ATA (which is derived by disturbing the ETA according to the distribution function outlined in Figure 3). The Generators table contains the arrival process to be used for each ship type.

The Ships table contains specific ship data such as type, size, length, type of cargo, loading time, and the number of ships of this type arriving annually.

The table JettyProducts describes which type of products can be handled by which mooring points, whereas the table JettyLengths holds the lengths of the mooring points. The base stock levels of the tanks are stored in table Tanks. Table Factory specifies the number of tons of raw materials to be processed, and finished products to be produced, both on a yearly basis. Table SimulationSettings holds some data concerning the distribution function used for disturbing ETAs.

The AnnouncedShips, WaitingTimes, and TankLevel atoms are used to store data collected during the simulation run. The Graph Tank atoms visualize the tank levels over time.

5.3 Priority of Ships

The priority scheme used in the selection of a suitable mooring point for a ship is based on the following:

- There are only two priorities (high and low);
- Long ships have high priority, short ships have low priority.

Ships with low priority do not get assigned to a mooring point when a high priority ship is known to arrive within the next 48 hours and would be assigned to the same mooring point within the time frame that the low priority ship will still be busy with (un)loading.

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 ships' waiting times and stock levels. In doing so, some level of detail is maintained, in that we make a distinction between waiting times for high priority and for low priority ships.

6.1 Simulation Parameters

For each of the three arrival processes, a ten-year simulation run is conducted. With the equidistant and stock-planned arrival processes, year-based stratification is applied to ship arrivals (for the Poisson process this is not possible). This means that the total number of ships of each type is fixed per year (see Table 1), and aligned with the factory's production capacity, yielding a constant base stock level over time (i.e. stock does not structurally increase or decrease over one year, hence not over ten years). Year-based stratification is consistent with reality, in the sense that many contracts concerning transport of raw material and finished products are based on specified quantities per year (often to be shipped in, for example, monthly batches).

The simulation 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.

Table 1: An Example of the Variety of Ships

Shiptype		Size (tons)	Length	Product (type)	(Un)load (hours)	Number per year
1	Barge	1,500	short	A	6	196
2	Vessel	2,000	short	A	8	48
	…	…	…	…	…	…
8	Vessel	6,000	short	B	26	12
	…	…	…	…	…	…
10	Vessel	2,000	long	C	14	126
	…					
13	Vessel	10,000	long	D	44	14
14	Vessel	20,000	long	D	56	8

6.2 Results and Analysis

Tables 2 and 3 show the relevant simulation outcomes. Table 2 contains the waiting statistics for ships, per arrival process, each divided into separate results for high and low priority ships. Table 3 reports on the maximum and minimum stock levels reached for each of the arrival processes, both in raw material and finished product tanks. Standard deviations values are based on a comparison of the outcomes for each of the ten years.

6.2.1 Waiting Times

From Table 2, 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 Poisson 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.

Table 2 reveals a noticeable difference between the outcomes of equidistant arrivals and stock-controlled arri-

Table 2: Ship Statistics per Arrival Process (Means over a 10-Year Period)

	Ship Priority: Low		Ship Priority: High	
	Mean	St. dev.	Mean	St. dev.
Poisson Arrivals:				
Total number of ships	1,174	31	205	16
Percentage of ships that had to wait	39.5%	3.2%	18.9%	3.4%
Average waiting time of ships that had to wait (hours)	9.1	1.5	14.1	1.6
Equidistant Arrivals:				
Total number of ships	1,163	0	208	0
Percentage of ships that had to wait	28.7%	0.7%	9.2%	1.5%
Average waiting time of ships that had to wait (hours)	7.2	0.2	9.8	1.2
Stock-controlled Arrivals:				
Total number of ships	1,163	0	208	0
Percentage of ships that had to wait	14.2%	0.9%	8.5%	1.4%
Average waiting time of ships that had to wait (hours)	3.8	0.3	10.0	2.3

Table 3: Stock Levels Ranges per Arrival Process (Means in Tons over a 10-Year Period)

	Tank A		Tank B		Tank C		Tank D	
	Mean	St. dev.	Mean	St. dev.	Mean	St. dev.	Mean	St. dev.
Poisson	88,342	21,937	46,741	13,252	41,470	15,041	116,812	39,839
Equidistant	10,756	273	11,265	342	3,381	283	27,474	574
Stock-controlled	6,702	474	5,893	296	2,945	340	15,552	682

vals. 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 equidistant 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 throwing in a Poisson process, is simply insufficient, rendering any subsequent decision (for example on expensive alternative jetty layout to reduce waiting times) ill founded.

6.2.2 Stock Levels

Table 3 shows ten-year stock level statistics in terms of the difference between minimum and maximum levels reached. Poisson arrivals allow for the broadest stock fluctuations. In fact, since Poisson arrivals constitute an uncontrolled process, stock range values are theoretically unbounded. This is not the case for stratified equidistant and stock-controlled arrivals. However, with equidistant arrivals, considerable fluctuations are still observed, necessitating high base stock levels to avoid stock outs.

Figures 5, 6 and 7 show example stock behavior over time for product D over a one-year period (notice that the scale of the vertical axis varies). Figure 6 shows that fluctuations are such, that the initial stock level for product D (2000 tons) does not suffice to avoid stock outs.

The stratified arrival processes 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. This does not hold for Poisson arrivals, as is evident from Figure 5.

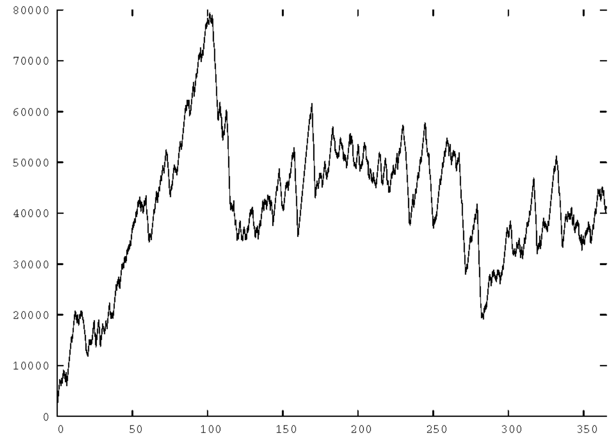

Figure 5: Level of Tank D During One Year with Poisson Process

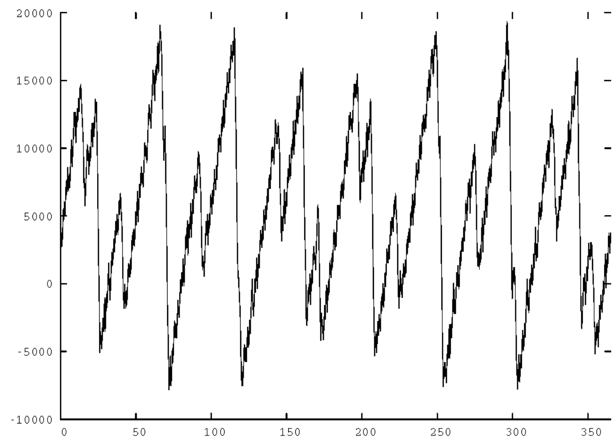

Figure 6: Level of Tank D During One Year with Equi-Distant Process

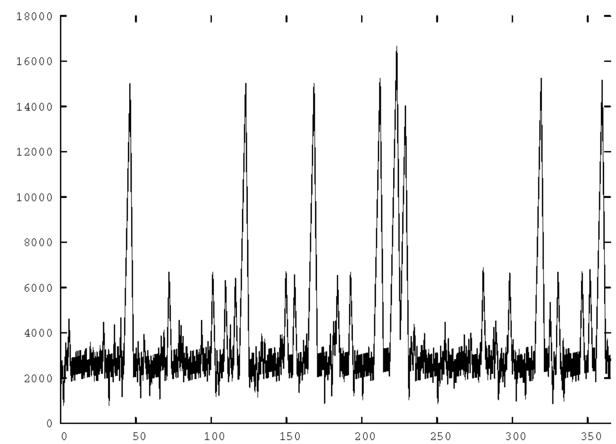

Figure 7: Level of Tank D During One Year with Stock-Control Process

Figure 7 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.

6.2.3 Jetty Utilization

The jetty utilization varies little over the three arrival processes. This is due to the fact that with all three processes, roughly the same number of ships is generated. In fact, year-based stratification with equidistant and stock-controlled arrivals causes ships to be generated in identical numbers and types. Differences in jetty utilization follow from differences in end-of-year situations among simulation runs.

7 CONCLUSIONS AND FURTHER RESEARCH

The importance of carefully modeling the arrival processes is clearly demonstrated in this article. The often-used Poisson process has by far the worst performance of the three processes discussed, both in terms of the 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.

Obviously, the challenge in managing logistical processes will be to determine which arrival processes can be actually realized. This requires close collaboration between production, logistics and the sales or marketing functions within a company. If such cooperation is lacking, a marketing department might buy or sell large quantities to meet sales targets, causing serious disruptions in planned arrivals, yielding costly delays.

There are various directions in which future research is planned. First, the role of the jetty's layout needs to be explored, specifically the impact of limited length of the individual mooring points, and the restrictions on the availability of piping for specific products.

Also, the effects of using the allocation scheme for assigning ships to mooring points, including the priority scheme based on 48-hour lookahead, requires further study.

Finally, we intend to consider yet another arrival process, a hybrid one, with planned arrivals for the larger vessels and equidistant or Poisson arrivals for the smaller barges.

More information on this study is available online via http://www.few.eur.nl/few/research/eurfew21/m&s/article/jetty/. The website contains graphs showing the levels of all tanks over a one year period and a video that shows a simulation run.

ACKNOWLEDGMENTS

The authors would like to thank the students Nees Jan van Eck, Arthur Oink, Gerard Seedorf and Ludo Waltman 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.

REFERENCES

Andrews, S., F.H. Murphy, X.P. Wang, and S. Welch. 1996. Modelling crude oil lightering in Delaware Bay, *Interfaces* 26(6): 68-76.

Baunach, G.R., E.S. Wibberley, and B.R. Wood. 1985. Simulation of a coal transhipment terminal: Batam Island, Indonesia. *Math. Comp. Simul*. 27: 115-120.

Dekker R. 1999. Simulation of jetty and storage activities for oil and chemicals. In: *SimLog, Simulation and logistics in the harbor*, ed. J. Van Nunen and L. Verspui, 105-116. Eburon. Delft (in Dutch).

Enterprise Dynamics. 2003. Documentation material. Incontrol Enterprise Dynamics. More information is available online at <www.enterprisedynamics.com>.

Heyden, W.P.A. van der, and J.A. Ottjes. 1985. A decision support system for the planning of the workload on a grain terminal. *Decision Support Systems*. 1: 293-297.

Nunen, J. van, and L. Verspui, 1999. *SimLog, Simulatie en logistiek rond de haven*, (translation: *SimLog, Simulation and logistics in the harbor*), Delft, Eburon (in Dutch).

Ottjes, J.A. 1992. *Modelvorming en simulatie van logistieke systemen*. Manual i76C, Delft University (in Dutch).

Ottjes, J.A., S. Hengst, and W.H. Tutuarima. 1994. A simulation model of a sailing container terminal service in the port of Rotterdam. In: *Proceedings of the Conference on Modeling and Simulation*, ed. Guasch and Huber, 876-880.

Philips, O.O. 1976. *Optimization models for a crude oil storage export system*, Ph.D. thesis, Penn. State Univ.

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. p. 108-128.

AUTHOR BIOGRAPHIES

EELCO VAN ASPEREN graduated in Business Computer Science at Erasmus University Rotterdam in 1993. From 1991 to 2000 he was a member of the IT support group at the Erasmus University Rotterdam, for the department of Computer Science and from 1995 for the Faculty of Economics, focusing on the design and implementation of large scale computer facilities. Since January 2000 he is an assistant professor at the Department of Computer Science of the Faculty of Economics at Erasmus University Rotterdam. His research focuses on simulation with applications in logistics. You can reach him by e-mail at <vanasperen@few.eur.nl>; and his web address is <http://www.few.eur.nl/few/people/vanasperen/>.

ROMMERT DEKKER is a full professor in operations research at the Econometric Institute of Erasmus University Rotterdam. He obtained his Ph.D. in operations research at the State University of Leiden, and his M.Sc. degree in industrial engineering from Twente University of Technology. He worked with Shell for seven years on reliability and refinery logistics. His current research interests are: maintenance and logistics (inventory control, spare parts, ports, containers and reverse logistics). He has applied simulation models in various logistical problems. His e-mail address is <rdekker@few.eur.nl>; his web address is <http://www.few.eur.nl/few/people/rdekker>.

MARK POLMAN graduated in Business Computer Science at Erasmus University Rotterdam in 1993. Since January 2000 he is an assistant professor at the Department of Computer Science of the Faculty of Economics at Erasmus University Rotterdam. Previous research areas include machine learning and communication modeling in distributed systems. His research focuses on discrete-event simulation. His e-mail address is <polman@few.eur.nl>; his web address is <http://www.few.eur.nl/few/people/polman>.

HENK DE SWAAN ARONS is an associate professor at the Department of Computer Science of the Faculty of Economics at Erasmus University Rotterdam. He graduated in Applied Mathematics at Delft University of Technology in 1972. In 1991 he obtained his Ph.D. degree in computer science at Delft University of Technology. The thesis was mainly concerned with the design, applicability and applications of expert system tools. His research focuses on discrete-event and continuous simulation, with the emphasis on economical applications. His e-mail address is <deswaanarons@few.eur.nl>; his web address is <http://www.few.eur.nl/few/people/deswaanarons/>.