ALLOCATION OF SHIPS IN A PORT SIMULATION

Eelco van Asperen
Rommert Dekker
Mark Polman
Henk de Swaan Arons
PO Box 1738
Erasmus University Rotterdam
Faculty of Economics and Business
3000 DR Rotterdam, THE NETHERLANDS

KEYWORDS
Port Simulation, Discrete-Event Simulation, Arrival Processes, Port Logistics.

ABSTRACT

The limited jetty capacity of ports causes costly ship delays. This is a particular concern for large ocean-going vessels. Terminal operators attempt to reduce ship delays both in number and duration but have to take the number of jetties and their functionality and layout as a given. In such a setting, the arrival process of ships determines the delays in the loading and unloading process. Ships can arrive according to a schedule, for example based on stock levels or regular intervals, unscheduled, or even uncontrolled which is the case in a Poisson process. Priority rules in the processing of ships further impact the efficiency, both for stock-controlled and equidistant arrivals. Based on data from a real case study, this paper describes a number of simulation experiments to assess the impact of the arrival process on ship delays and, and to show the beneficial effect of the application of priority rules on the efficiency of loading and unloading.

1. INTRODUCTION

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 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, (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.

In this paper, we focus on the analysis of ship waiting statistics and stock fluctuations under different arrival processes using a simulation model which is fed with data (types and number of ships handled per year) from a confidential case study in the Port of Rotterdam. The case study concerns 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. We consider three types of arrival processes. 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 in time of ships carrying the same product type. The last type of arrival process is an uncontrolled process, derived from a Poisson process.

Within each arrival process type a further distinction can be made between prioritized and non-prioritized queues in front of the jetty's mooring points. In this paper the various arrival processes will be compared with and without the application of priority rules. In the simulation model, some details concerning the diversity of ships and their numbers have been omitted. 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.

Section 2 briefly describes the model of the loading and unloading process. The various arrival processes are discussed in more detail in Section 3 with a focus on the application of priority rules in processing ships. The implementation model is the subject of Section 4 and the experiments carried out with it and their results are discussed in Section 5. The conclusions are presented in Section 6.

2. THE MODEL

A detailed description of the model can be found in (Asperen et al. 2003b). The model comprises the arrivals in time of ships, a jetty with a number of mooring points, storage tanks and a factory.

Proceedings 15th European Simulation Symposium
Alexander Verbraeck, Vlatka Hlupic (Eds.)

ISBN 3-936150-28-1 (book) / 3-936150-29-X (CD)

The Jetty. This is the loading and unloading facility with a number of mooring points. In this case there are four mooring points (mooring point 1 to 4) in a T-shaped layout. They differ in a number of aspects such as the length of the ships they can handle and the materials (raw materials A or B, and finished products C or D) they can load and/or unload.

Raw Materials, Finished Products, Tanks and Stocks. 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. Tanks can be used for only one type of raw material or finished product. 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. 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).

Ships. There are ocean-going vessels, short-sea shipping vessels and inland barges which unload raw materials or load finished products. Each ship has properties relevant for the model such as size (tonnage), length (a distinction between long or short suffices), product (each ship handles just one specific type of cargo) and the (un)loading time. When a ship has arrived in the port, a suitable mooring point is selected according to specified rules, which are discussed below.

3. 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, this paper considers three scenarios to capture the ship arrival process.

Types Of Arrivals

Stock-controlled arrivals. These types of arrivals aim at maintaining a target base stock level 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, 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 late arrivals (or from ships unable to instantly find an unoccupied mooring point).

Equidistant arrivals. With equidistant arrivals, arrivals of ships within the same ship type are assumed to be evenly spread over the year. For example, per year, 12 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.

Uncontrolled arrivals. The third arrival process considered in this paper is an uncontrolled process: within each cargo type, ships arrive uniformly distributed over the year. This process is obtained by specifying the number of arrivals in a Poisson process. See (Asperen et al. 2003b) for details.

The stock-controlled and equidistant arrival processes actually yield a series of *expected times of arrival* (ETAs). However, in reality ships 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. See (Asperen et al. 2003b) for more details on these disturbances.

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 each carrying 4000 tons of product B, with a loading time of 26 hours (for the meaning of the priority column, see below).

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.

This implies 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.)

Priorities

In reality, the arrival time of a ship is known, sometimes days beforehand, to the plant. This information 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

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

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.

In this paper, 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 in front of the mooring point with the smallest workload, or, in case of equal workloads, the shortest queue so far. Here, 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 in front of it.

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 in front of 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 in front of 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. THE IMPLEMENTATION MODEL

The implementation model is based on the model outlined in Section 3. A detailed description can be found in (Asperen et al. 2003b). However, for a better understanding of the present paper, a few highlights are presented here.

The simulation model has been implemented in Enterprise Dynamics, a simulation package for discrete-event simulation (Enterprise Dynamics 2003). The implementation model outlined in Figure 1, comprises various types of atoms, the Enterprise Dynamics equivalents of objects. Some of the atoms implement the simulation'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 which have entered the port thus far (262). Nine ships are on their way to the jetty. The utilization of mooring points 1 through 4 up to now has been 61.3%, 47.1%, 63.1% and 72.8%, respectively. At present, all four are occupied. Queues 1, 2, and 3 are empty, whereas Queue 4 contains one waiting ship. The actual contents of tanks A through D

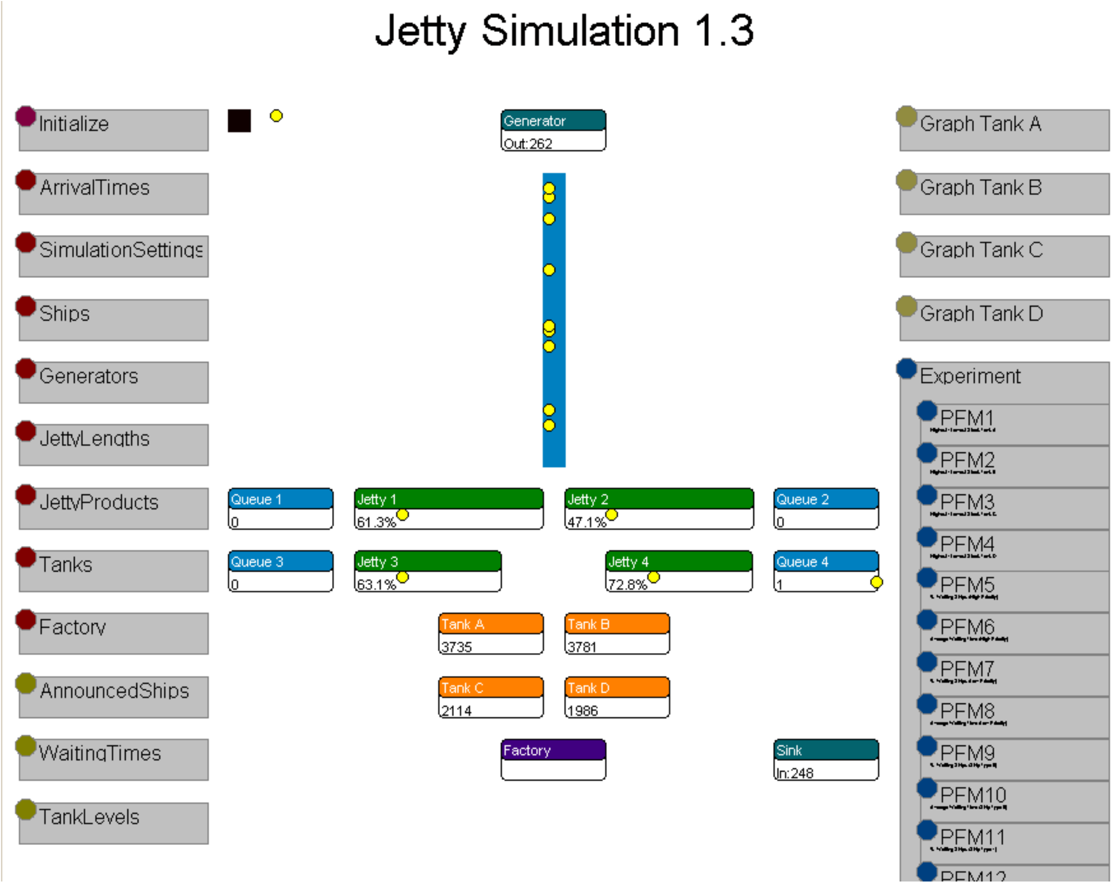

Figure 1: Implementation of the Simulation Model

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.

Logic

The Generator atom 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 waits 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. The tanks are assumed to have unlimited capacity and the possibility to contain negative stock. This simplification does not affect the simulation's objective.

Data

The atoms on the left side represent tables providing data for the simulation process. All but the Initialize atom, which contains some code to be executed at the beginning of each run, are actually tables. The top seven of these are filled from text files at the beginning of each run, and contain data concerning the arrival times (both ETA, ATA, including disturbances); some initializing data in the simulation settings; specific ship data such as type and size; the lengths of the mooring points and the products they can handle; the base stock levels of the various tanks; and the annual amounts of raw material processed and finished products produced by the plant.

The bottom three tables on the left are filled with data during simulation runs. They contain the data concerning the allocation of a ship to a mooring point, the waiting times statistics for all ship types and the stock level movements for each tank.

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 more than 30 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 remaining PFMs are used to collect similar data per individual ship type.

5. 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: with stock-controlled arrivals, equidistant arrivals per ship type and uncontrolled arrivals, each with or without the use of priority rules.

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 and 3 show the relevant simulation outcomes. Table 2 contains the waiting statistics for ships for the three arrival processes without priority rules, each divided into separate results for high and low-priority ships (this distinction is made to facilitate a comparison with the results of simulation runs that *do* include a priority scheme, as described below.) Table 3 reports on the maximum and minimum stock levels reached for each of these arrival processes, both in raw material and finished product tanks. Table 4 shows the differences for each arrival process between using and not using priority rules for mooring point allocation.

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 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.

Table 2 also reveals 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 equidistant arrivals keep ships of identical types apart. With multiple ship types per cargo type this is an advantage. However, the arrival rates of the individual ship types (which is something very particular to this simulation) have an impact as well. Consider, for example, the 126 type 10 vessels, and the 14 type 13 vessels from Table 1. If, with equidistant arrival times, the first ships of both types have identical expected times of arrival, every arrival of a ship of the latter type coincides with one of the former. The observed differences in waiting time statistics among arrival processes, and their causative factors, clearly demonstrate the need for careful arrival process modeling, which is this paper's primary objective. Obviously, arrival process modeling requires a careful look

Table 2: Ship statistics for the various arrival processes 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 (%)				
Stock-controlled	21.1	3.7	12.0	1.0
Equidistant	34.7	1.8	23.5	0.8
Uncontrolled	45.7	2.1	35.2	2.0
Average waiting time of ships that had to wait (hrs)				
Stock-controlled	7.9	1.1	3.5	0.2
Equidistant	9.5	0.6	6.2	0.2
Uncontrolled	12.3	1.8	7.5	0.9

Table 3: Stock level ranges for the various arrival types 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

Table 4: Ship statistics for the various arrival processes, priority rules vs. no priority rules (means over a 10 year period)

	Ship Priority			
	High		Low	
	No priority rules	Priority rules	No priority rules	Priority rules
Percentage of ships that had to wait (%)				
Stock-controlled	21.1	8.5	12.0	14.2
Equidistant	34.7	9.2	23.5	28.7
Uncontrolled	45.7	18.3	35.2	40.5
Average waiting time of ships that had to wait (hrs)				
Stock-controlled	7.9	10.0	3.5	3.8
Equidistant	9.5	9.8	6.2	7.2
Uncontrolled	12.3	14.6	7.5	9.4

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 insufficient, rendering any subsequent decision (for example on expensive alternative jetty layout to reduce waiting times) ill-founded.

Stock Levels

Table 3 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. It is clear that the choice of arrival process is an important factor in simulation outcomes. More information about the stock fluctuation patterns over time can be found in (Asperen et al. 2003b).

The Effect Of Using Priority Rules

In section 4.6 it was explained that priority rules are 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.

Table 4 shows the ship waiting statistics over a ten-year simulation period for each arrival process, both with and without (copied from Table 2) priority rules. Standard deviations have been omitted for brevity.

In all cases, applying priority rules indeed reduces the percentage of high-priority ships, while increasing the percentage of low-priority ships that have to wait. 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 applying priority rules, 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.

6. CONCLUSIONS AND FURTHER RESEARCH

In (Asperen 2003a) it was already concluded that careful arrival process modeling is very important with respect to ship and stock statistics. 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.

In this paper the emphasis was put on how priority rules affect these results. It may be concluded that priority rules have a positive effect on the ship statistics. The percentage of the high-priority ships that had to wait was reduced with a factor of about 3, at the cost of low-priority ships. In the less frequent cases that high-priority ships had to wait, priority rules slightly pushed up average waiting times.

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 more sophisticated allocation strategies than a two-class priority scheme for assigning ships to mooring points, requires further study.

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

ACKNOWLEDGEMENTS

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.

WEB

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

REFERENCES

Andrews, S., F.H. Murphy, X.P. Wang, and S. Welch. 1996. "Modeling crude oil lightering in Delaware Bay." *Interfaces* 26, No. 6, 68-78.

Asperen, E. van, R. Dekker, M. Polman and H. de Swaan Arons. 2003a. "Modeling Ship Arrivals in Ports." In *Proceedings of the Winter Simulation Conference 2003*, S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice (Eds.). New Orleans. Accepted for publication.

Asperen, E. van, R. Dekker, M. Polman, H. de Swaan Arons, and L. Waltman. 2003b. "Arrival Processes for Vessels in a Port Simulation." ERIM Report Series Research in Management, ERS-2003-067-LIS.

Baunach, G.R., E.S. Wibberley, and B.R. Wood. 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 *SimLog, Simulation and logistics in the harbor*, J. Van Nunen and L. Verspui (Eds.). Eburon Delft, Netherlands, 105-116 (in Dutch).

Enterprise Dynamics. 2003. *Documentation material*. Incontrol Enterprise Dynamics. Contact http://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).* 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 (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 1994 Conference on Modelling and Simulation*, Guasch and Huber (Eds.), 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.

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/`.