

# Arrival processes in port modeling: insights from a case study

Eelco van Asperen<sup>a</sup>, Rommert Dekker<sup>b</sup>,  
Mark Polman<sup>a</sup>, Henk de Swaan Arons<sup>a</sup>

*a Department of Informatics, b Department of Econometric and Management Science, Erasmus School of Economics,  
Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands*

---

## Abstract

This paper investigates the impact of arrival processes on the ship handling process. Two types of arrival processes are considered: controlled and uncontrolled. Simulation results show that uncontrolled arrivals of ships perform 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. The combination of stock-controlled arrivals for large vessels and equidistant arrivals for barges also performs better than the uncontrolled process. Careful allocation of ships to the mooring points of a jetty further improves the efficiency.

*Keywords:* Logistics; Supply Chain Management; Simulation; Transportation; Case study.

---

## 1. Introduction

This paper investigates the impact of ship arrival processes and jetty allocation schemes on the efficiency of the loading and unloading process in a port simulation. An arrival process is a formal specification of how entity arrivals in a system are scheduled. In our case, it determines, among others, the likelihood of several ships arriving simultaneously, which is an important aspect in, for example, determining the required jetty capacity. Our research was triggered by work done on a confidential case study with the objective to help determine the optimal layout of the jetty owned by a new chemical plant in the port of Rotterdam. The original tender of that case study provided detailed data on the types and numbers of ships to be handled annually, but failed to specify their arrival process. However, an initial simulation model described in (Van Asperen et al., 2003) demonstrated a considerable impact of the type of arrival process on system performance, in terms of both waiting times of ships waiting to load or unload at the jetty, and stock fluctuations in the tanks on the chemical plant's facilities. In this paper, we further develop and analyze the arrival processes themselves, evaluate their impact on system performance, and evaluate several jetty allocation schemes for additional performance enhancement.

A basic distinction can be made between uncontrolled and controlled arrival processes. Uncontrolled arrivals are typically modeled by a Poisson process, a common assumption, for example, in modeling incoming telephone calls in call center simulations. Controlled arrivals concern scheduled arrivals, such as scheduled airline flight arrivals to an airport (Banks, 2000). For our port system we distinguish two types of controlled arrivals. 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 plant's tanks. The second type is based on equidistant arrivals per ship type and relates to contracts prescribing product supply and pick-up at regular time intervals, e.g., once a month. We compare model outcomes based on four different arrival processes: uncontrolled, stock-controlled, equidistant and a blend of stock controlled arrivals for the larger ships and equidistant arrivals for the smaller ones. Furthermore, for all four types of arrival process, it is investigated to what extent careful allocation of ships to the jetty's mooring points enhances system performance.

Apart from some very scattered material, little practice with the simulation of port facilities can be drawn from existing literature. (Van Nunen and Verspui, 1999) provides insight in simulation and logistics in ports, but it is in Dutch only. Here, 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 in studies on jetty capacity. 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: (Philips, 1976) and (Andrews, 1996) describe the planning of a crude-oil terminal; (Baunach et al, 1985) deal with a coal terminal; (Heyden and Ottjes, 1985), (Ottjes, 1992) and (Ottjes et al., 1994) deal with the set-up of the simulation programs for terminals. None of these papers however, deals explicitly with arrival processes. (Kia et al., 2002) do mention the arrival process in the context of a port simulation: they assume a Poisson process.

In section 2 we provide a detailed description of the conceptual model of the system. In section 3 the various types of arrival processes are discussed in detail. Three schemes for the allocation of ships to the jetty's mooring points are given in section 4. Section 5 provides a brief discussion of how the simulation models have been implemented. The experiments conducted with the model and their results are discussed in section 6, and the conclusions are presented in section 7.

## **2. The conceptual model**

The system considered in this paper involves a chemical plant with a continuous production process. Both the supply of raw materials and the export of finished products occur through ships loading and unloading at a plant-owned jetty. Since disruptions in the plant's production process are very expensive, buffer tank capacity is required for sustained production and tolerance towards variations in ship arrivals and overseas exports through large ships. 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.

Operational costs of such a facility increase when ships have to wait to (un)load, or amplitudes in stock level fluctuations widen (tankage is costly as well). Causing factors of such events include the shape of the jetty, the number of mooring points it has and their restrictions with respect to the types of ships and cargo they can handle, and whether the port is an open port or has locks. Other possibly relevant factors – and that is the key subject of this paper – are the arrival processes of the various types of incoming ships and the allocation of ships to the jetty's mooring points. Figure 1 provides a schematic outline of the model as a whole. Apart from the arrivals of ships, it comprises a jetty with a number of mooring points, several storage tanks and a chemical plant, which are described in sequence below.




Figure 1. A schematic outline of the loading and unloading process, including jetty, tanks and plant.

### 2.1. The jetty

The jetty provides four mooring points (numbered 1 to 4) in a T-shaped layout (Figure 2). Ships arriving at the jetty to load or unload cargo dock at one of these. Mooring points 1 and 2 are suited to handle ships of all sizes; mooring points 3 and 4 can handle only short ships (see also Table 1).




Figure 2. The jetty layout.

Incoming ships unload raw materials (A or B), or load finished products (C or D). Pipes facilitate the transport of all chemicals to and from the ships. Since cost considerations are a limiting factor on their construction, not every type of raw material and finished product can be (un-)loaded at every mooring point. For example, mooring point 1 can handle A, B, and C, whereas mooring point 2 can only handle products C and D.

## 2.2. Tanks and stocks

After unloading, raw materials are stored in tanks A and B, for later extraction and processing by the plant. Finished products are transferred to tanks C and D, to be loaded into ships.

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 plant, and from the plant into the tanks are continuous processes, which, in reality, are subject to several restrictions. One restriction prescribes that there shall be no simultaneous pumping and running into and out of a tank. Another restriction is that stocks are limited. However, for simplicity we allow them to take on any value, and neglect ship delays because of stock outs or lack of ullage (available tank space). We ignore all these restrictions, because they do not affect the comparison between the arrival processes.

## 2.3. Ships

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

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

When a ship has arrived in the port, a suitable mooring point is selected according to a set of rules, which are discussed below. Table 1 shows all types of ships loading and unloading at the jetty along with their values for the aforementioned properties. For example, every year, a total of fourteen short vessels arrive carrying 4,000 tons of product B, with a loading time of 26 hours. Columns “Ships per year”, “Priority” and “Tons per year” are discussed in more detail later.

Table 1  
Ship types, properties, and arrival rates

| Ship type | barge/<br>vessel | Size<br>(metric tons) | Length | Product | Loading<br>time<br>(hours) | Ships<br>per year | Priority | Tons per<br>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        |

### 3. The arrival process

In many simulation studies it is assumed that arrivals in client-oriented processes cannot be controlled. Simulation languages and environments acknowledge this and tend to offer Poisson as a first-choice option for the specification of arrival processes. However, in some port situations, a definite measure of control over the arrival process can be observed. This suggests that care should be taken in settling on a process to feed the simulation with ship arrivals.

In order to understand how such considerations affect our simulation study, one should first analyze the planning process and organizational structure. Usually, every month or few months, depending on the company, the sales/marketing department sets up tactical sales plans, including contract sales over a long period, new contract sales and spot sales. In order to see whether the production required for sales fulfillment can be achieved, possible bottlenecks in the production process need to be identified. Sales plans and bottleneck analysis together constitute the primary building blocks for a tactical production/sales plan. For our chemical plant this plan ultimately determines the required production level for the coming period, and provides direction for the logistics department to plan order pickups and deliveries.

However, many long-term contracts in the bulk oil and chemical sector, while including detailed price specifications (to avoid uncertainties as a consequence of market fluctuations), are considerably less rigid about the exact delivery dates. It is up to the waterfront part of the logistics department to agree with clients and suppliers on pickup and delivery schedules. Furthermore, short term sales and purchases require additional planning effort, since these often depend on ad hoc opportunities as short term traders tend to focus on prices, disregarding logistical feasibility. The logistics department is now faced with the challenge of accommodating this type of deals as well.

Finally, it should be noted, that during the design of a new plant, it is often unclear to both the logistics department and the construction engineers what purchasing/sales contracts will be used by the marketing/sales department in the future, and in what the ratio of short term deals and long term contracts will be.

In designing a simulation model for such a logistical process, one cannot but make some assumptions about the level of control that the logistics department maintains over ship arrivals. Several possibilities for modeling such control (or lack thereof) are described below. Their impact on simulation outcomes is this paper's main subject.

#### 3.1. *Expected Times of Arrival (ETA)*

Modeling control over ship arrivals involves the notions of Expected Times of Arrival (ETAs) and Actual Times of Arrival (ATAs). Here, the time of arrival is the time at which a ship arrives before the jetty. Let us start with the ETA. We consider two major types of controlled arrival processes yielding ETAs: stock-controlled and equidistant arrivals, and a third type, hybrid, which is a blend of these two types.

##### *Stock-controlled arrivals*

The plant management's aim is to achieve efficient production, avoiding costly interruptions such as those caused by stock-outs in the raw materials tanks. Further efficiency can be attained through prevention of stock-outs in the finished products tanks. These would cause ships to have to wait around for cargo, which is also costly. In case ship arrivals can be planned by plant management, stock-controlled arrivals can be used to maintain a target base stock level in the tanks as a buffer for production (raw materials) and transport (finished products). 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 plant. 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 plant 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).

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 1,000 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 1,000 tons of raw material, minus the volume of raw material pumped out of the tank by the plant. 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  $t_p$  ('p' for 'planned'). However, this ship could arrive ahead of time (see section 3.2), for example at  $t_a$  ('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.




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

### *Equidistant arrivals*

Equidistant arrivals model situations in which loading and unloading ships arrive at regular intervals. This regularity could, for example, be the consequence of year-based contracts specifying 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 the same ship type are assumed to be evenly spread over the year. For example, per year, twelve vessels carrying 6,000 ton of product B arrive (see Table 1). With equidistant arrivals, this means a 1-month inter-arrival period between such ships. Note that ships from different ship types may still arrive simultaneously.

### *Hybrid arrivals*

In a hybrid arrival process, the total population of ships is partitioned along some criterion, after which each type is assigned an arrival process for scheduling the arrivals of its members. In this paper, we consider one hybrid process, in which the smaller ships (below 6,000 tons) arrive equidistantly, whereas the larger ones are subject to stock-controlled scheduling. The arrivals of all larger ships are scheduled on a per-product basis whereas the smaller ships are scheduled per ship type.

The underlying assumption is that contracts with clients and suppliers are such that alignment of the corresponding shipments with the production process is, in principle, hard. Hence, the majority of deals results in equidistant pick-ups and deliveries, partly due to transportation-related clauses in the contracts, and partly due to a client/supplier (especially those transporting many smaller shipments) preference for regularity in their logistical processes. Under such circumstances, the logistics department's focus will be on aligning the larger shipments with the production process. This is feasible for two reasons. First of all the number of large shipments is limited. Second, the plant operator and clients and suppliers requiring large shipments have a shared interest in coordinating ship arrivals and thus reducing waiting times. From the plant's point of view, large shipments are most likely to cause stock-outs or lack of available tankage, and from the client/supplier's point of view, avoiding delays for their large ships pays off (waiting by large ships is relatively costly).

Obviously, when simulating with this hybrid arrival process one implicitly assumes that stock-controlling large ship arrivals is feasible.

### *3.2. Actual Times of Arrival (ATA)*

In reality ships will seldom exactly meet the schedule as defined by the ETAs. Most ships arrive within a relatively short interval around their expected time of arrival, while some arrive significantly earlier or later. Such deviations are modeled by a disturbance to the ETA. An ETA together with a disturbance yields the actual time of arrival (ATA) of a ship. The parameters of the disturbances were set together with shipping experts, taking into account the fact that the Port of Rotterdam is an open port, with relatively stable weather conditions.

The distribution function of the deviation in hours from the ETA can be described as follows. If the deviation in hours is denoted as  $x$ , then:

$$\begin{aligned}
 x &= U(-12,-2) && \text{with } p = 0.1 \\
 &U(-2,2) && \text{with } p = 0.8 \\
 &U(2,12) && \text{with } p = 0.1
 \end{aligned}
 \tag{1}$$

where  $U$  is the uniform distribution function. This means that 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 two hours before and two hours after the corresponding ETA, in all cases with constant probability density (see Figure 4).




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

### 3.3. Uncontrolled arrivals

The assumption underlying uncontrolled arrivals is that – in contrast to both stock-controlled and equidistant arrivals – there is no control by plant management over the intervals at which ships arrive. In that case, opting for a Poisson process is the logical choice. This does imply that the number of arrivals per year can vary. In the process industry, however, annual throughput is more or less fixed. As a consequence, in our model, the total number of arrivals per year within each ship type is fixed across all arrival processes. This implies that, if the distribution function of interarrival times is exponential, the arrival times are uniformly distributed (Banks, 2000). When simulating uncontrolled arrivals, we therefore draw the arrival times per ship type from a uniform distribution over the year.

### 3.4. Ship arrival rates

Table 1 shows how many ships of each type arrive per year. For each product/cargo type, the number of ships carrying it is chosen such that the total amount of cargo transported matches the plant's capacity. For instance, per year, the plant 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).

With constant loading and unloading times per ship type, fixing the number of ships implies that the utilization rate of the jetty will be the same for all arrival processes. In our case, the utilization rate is 61%. According to industry norms, this is considered to be busy but not overloaded.

### 3.5. Input analysis

As was mentioned in the introduction, the case study's original tender did not specify the ships' arrival process, providing only the estimated numbers of ships arriving annually per ship type. This is a quite common phenomenon in simulation studies: the distribution functions of the various stochastic processes governing a system, such as interarrival times, service times etc., are often unavailable. In the case of arrivals, a Poisson process has proven to be a reliable choice when arrivals appear to be random. As a consequence, many simulation development environments present the Poisson arrival process as a first option for configuring simulation entity sources, see e.g., (Enterprise Dynamics, 2003), (Kelton, 2004) and (Arena, 2003). If historical arrival data is available, one may attempt to fit a distri-



bution function onto the dataset, and use it in the simulation model to generate arrivals. However, this strategy can easily lead to serious errors.




Figure 5. Result of Arena's Input Analyzer fit of a hybrid arrival process.

To illustrate, suppose that the actual system is fed by a hybrid arrival process as outlined before (ships of 6,000 tons and up arrive stock-controlled and ships of less than 6,000 ton arrive equidistantly), but the modeler is not aware of this. He may then use a data fitting program on the collected historical arrival data to help select the distribution function for his model. The results are displayed in Figure 5.

The figure was conceived as follows. Arrival data from a hybrid arrival process as generated by our own simulation model were fed to a data fitting program (Arena's Input Analyzer (Arena, 2003)). According to this data 13,708 ship arrivals occurred over a ten year period. The figure displays their interarrival times divided over 40 intervals. The distribution function that fits best based on criteria such as the square error is the beta distribution function  $-0.001 + 2010 * \text{beta}(0.957, 4.04)$ , but the exponential distribution function is still quite close (see (Law and Kelton, 2000) for both distribution functions). Based on these results the beta or even the exponential function appear suitable candidates for modeling the arrival process. However, as can be learned from Table 4 (shown and discussed in section 6) experiments show dramatic differences in simulation outcomes among the various arrival processes considered. This suggests that a data fitting strategy should be preceded by a thorough arrival process analysis to eliminate the possibility of the process being controlled instead of truly random.

#### 4. Jetty scheduling

The arrival process determines when a ship arrives at the port. Next, a scheduling algorithm can be used to control how the ship will be handled in the port. A ship entering the port will eventually be assigned a free mooring point which suits the ship's cargo type and length. The simplest mooring point allocation scheme we consider is one in which the ship is assigned the shortest suitable and available mooring point. If all suitable mooring points are occupied, the ship is placed in a queue before the mooring point with the smallest workload<sup>1</sup>, or, in case of equal workloads, the shortest queue so far. Such a scheme disregards any information on future ship arrivals that might be available.

However, in reality, the ATA of a ship is known to plant management, sometimes days beforehand, by a so-called *pre-arrival notice*, which can be used in more advanced mooring point allocation algo-

<sup>1</sup> 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.

rithms. The general idea is to incorporate all ships within an  $n$ -hour horizon into the choice of a mooring point for an incoming ship. Given the fact that for some ship types waiting is more expensive than for others (e.g., dependent on the type of cargo, the ship's capacity or crew size), adequate priority rules might reduce total costs induced by waiting for available mooring points. Also an enumeration algorithm may be applied to select the optimal allocation schedule of all possible schedules within the look-ahead time window. In general, this is a time-consuming approach.

In this paper we use the ATA information gained from the pre-arrival notices to implement 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 time horizon is 36 hours, i.e., the pre-arrival notice is received 36 hours before the ship's ATA. The priority scheme makes reservations for the high-priority ships based on their ATA. The assignment of a ship to a mooring point can be done as follows. A high-priority ship entering the port is in principle assigned to a free mooring point that suits 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.

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  $W_i(t)$  be the workload of mooring point  $i$  at time  $t$ , and let  $D_i(s)$  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 36-hour horizon will need mooring point  $i$  between  $t$  and  $t + W_i(t) + D_i(s)$ . 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. We will also consider an enumeration algorithm to find the optimal allocation schedule within a 36 hour window.

## 5. The implementation model

The model outlined in section 2 has been implemented in Enterprise Dynamics (Enterprise Dynamics, 2003), a simulation environment for discrete-event simulation. With this implementation, the experiments in (Van Asperen et al., 2003) were carried out. Later the model has been implemented in Java using a simulation library. The results presented in this paper are based on both implementations.

Simulation environments are generally easy to use, and allow for quick model construction. Also they provide built-in animation, generate statistics, and form well-tested simulation environments. Unfortunately, they also have their weak points. Relevant in this context is that, generally speaking, their programming facilities are poor and communication with other programming languages such as Java usually is laborious. General purpose programming languages such as Java or C++ lack the inherent advantages of the simulation environments. On the other hand, they provide a powerful, flexible and fast programming environment. This quality may be indispensable for solving some specific modeling problems, such as complex jetty allocation algorithms.

The initial simulation model was constructed fairly quickly using the Enterprise Dynamics (ED) environment. This implementation provides animation, which facilitates debugging and communication about the simulation model. However, ED's scripting language proved to be too limited for the implementation of complex issues, most notably stock-controlled arrivals. Hence, we implemented the arrival processes in an external (Java) program. The resulting list of interarrival times was used by a custom-built ED object to generate ship arrival events.

Due to more implementation problems concerning the mooring point allocation (e.g. using priorities) and the need for increased runtime speed, the second simulation model was developed in the Java programming language, using the DESMO-J library (Desmo-J, 2003). This discrete-event simulation framework has been a sound platform for our work.

## 6. Experiments and results

The Java 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.

We consider four arrival processes: a Poisson process as described in section 3.3, equidistant arrivals per ship type, stock-controlled arrivals per product type, and arrivals modeled using the hybrid process described in section 3.1. Each run starts in a steady-state situation, with the tanks partly filled.

Table 2 through Table 4 show the relevant simulation outcomes. Table 2 contains waiting statistics for ships with the simplest mooring point allocation scheme as outlined in section 4, divided into separate columns for high and low-priority ships<sup>2</sup>. 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. Table 4 adds the results of using the simple priority scheme outlined in section 4 and an enumeration algorithm to determine the mooring point allocation that yields the least waiting by ships within a 36 hour planning horizon. This is further discussed in section 6.4.

### 6.1. Waiting times

From Table 2 it can be observed that the choice for an equidistant, stock-controlled or hybrid arrival process shows a significant difference in terms of the number of waiting ships and the number of hours spent waiting by these ships when compared to the uncontrolled arrival process. This holds for both high and low-priority ships.

Clearly, a mechanism to keep ships apart, whether it be equidistant or stock-controlled arrival planning, prevents clusters of ships arriving within a small time frame, causing queues. For both low and high-priority ships, the stock-controlled arrival process 'outperforms' the equidistant arrival process. The results of the hybrid arrival process are in between those of the equidistant arrival process and those of the stock-controlled process.

Table 2

Waiting times per arrival process.

Means over a 10-year period; standard deviation is based on ten runs of one year.

|   | Ship Priority |          |       |          |
|---|---------------|----------|-------|----------|
|   | High          |          | Low   |          |
|   | Mean          | St. dev. | Mean  | St. dev. |
| <i>Percentage of ships that had to wait</i>                   |               |          |       |          |
| Uncontrolled  | 45.7%         | 2.1%     | 35.2% | 2.0%     |
| Equidistant   | 34.7%         | 1.8%     | 23.5% | 0.8%     |
| Stock-controlled  | 21.1%         | 3.7%     | 12.0% | 1.0%     |
| Hybrid  | 31.4%         | 3.6%     | 20.6% | 1.1%     |
| <i>Average waiting time of ships that had to wait (hours)</i> |               |          |       |          |
| Uncontrolled  | 12.3          | 1.8      | 7.5   | 0.9      |
| Equidistant   | 9.5           | 0.6      | 6.2   | 0.2      |
| Stock-controlled  | 7.9           | 1.1      | 3.5   | 0.2      |
| Hybrid  | 8.3           | 0.7      | 5.6   | 0.3      |

<sup>2</sup> 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.

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, spreading the arrivals of the first ship of each type, seeks to prevent the scheduling 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.

However, the observed differences in waiting time statistics among the arrival processes, whatever their causing 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 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 an expensive alternative jetty layout to reduce waiting times) ill founded.

## 6.2. 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. The results of the hybrid arrival process are again a blend of the equidistant and stock-controlled results.

Table 3

Stock levels ranges in tons per arrival process.

Results based on ten runs of one year.

|                  | Product |          |        |          |        |          |        |          |
|------------------|---------|----------|--------|----------|--------|----------|--------|----------|
|                  | A       |          | B      |          | C      |          | D      |          |
|                  | Mean    | St. dev. | Mean   | St. dev. | Mean   | St. dev. | Mean   | St. dev. |
| Uncontrolled     | 74,396  | 18,333   | 48,058 | 11,789   | 32,045 | 9,112    | 89,177 | 15,112   |
| Equidistant      | 10,756  | 273      | 11,245 | 312      | 3,381  | 283      | 27,474 | 574      |
| Stock-controlled | 6,970   | 468      | 5,890  | 294      | 3,012  | 320      | 15,982 | 578      |
| Hybrid           | 8,212   | 508      | 8,032  | 274      | 3,369  | 274      | 20,932 | 623      |

Figure 6 shows example stock behavior over time for product D over a one-year<sup>3</sup> period. The initial stock level for each arrival process was set to a value that would prevent stock-outs. Figure 6a shows the results of an uncontrolled arrival process. Note that the scale of figure 6a differs from the scales of the other three graphs: the uncontrolled nature of this arrival process causes large fluctuations in the stock level.

The largest available vessel (see Table 1) comes in to load product D eight times a year. This is clearly visible in the graph for the equidistant arrival process (Figure 6b). Figure 6c shows the typical stock fluctuation pattern for stock-controlled arrivals. Peak levels are reached whenever large ships are scheduled to arrive for loading. A late arrival around day 220 causes a larger peak due to continued production whereas the early arrival of the next ship makes the stock level drop below the base stock level. Figure 6d shows the stock level fluctuations for the hybrid arrival process, with stock-controlled arrivals for the larger ocean-going vessels and equidistant arrivals for all other vessels. Notice that in

<sup>3</sup> 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.

case of product D, stock fluctuation is almost completely determined by the size of the largest vessel, which makes it easy to determine the required tank capacity. The fluctuation patterns observed with the other products are similar in shape. However, their amplitude is considerably smaller, as product D is the only product transported by ships carrying as much as 10,000 and 20,000 tons of chemicals. So, again, the choice of an 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.

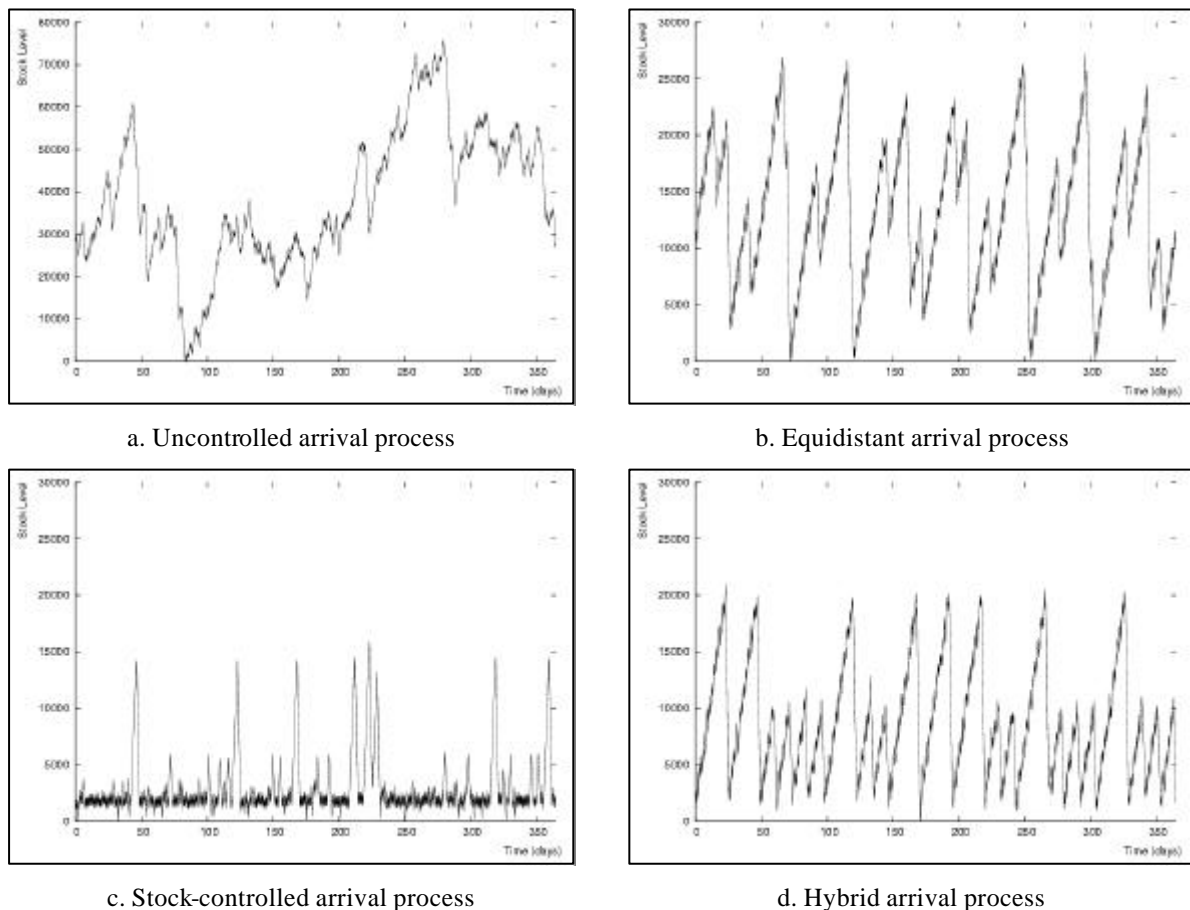


Figure 6. Level of tank D over a one year period.

### 6.3. The effect of using a priority scheme

In section 4 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.

Table 4 shows ship waiting statistics over a ten-year simulation period for the same types of arrival process, both with and without a priority scheme. Standard deviations have been omitted for brevity.

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.

Table 4

The effect of using a priority scheme and the optimal berthing sequence with a 36-hour horizon. Means over a 10-year period.

|   | Ship Priority      |                 |         |                    |                 |         |
|---|--------------------|-----------------|---------|--------------------|-----------------|---------|
|   | High               |                 |         | Low                |                 |         |
|   | No priority scheme | Priority scheme | Optimal | No priority scheme | Priority scheme | Optimal |
| <i>Percentage of ships that had to wait</i>                   |                    |                 |         |                    |                 |         |
| Uncontrolled  | 45.7%              | 18.5%           | 15.1%   | 35.2%              | 40.1%           | 34.9%   |
| Equidistant   | 34.7%              | 9.2%            | 3.5%    | 23.5%              | 28.7%           | 23.0%   |
| Stock-controlled  | 21.1%              | 8.5%            | 3.6%    | 12.0%              | 14.2%           | 7.3%    |
| Hybrid  | 31.4%              | 10.1%           | 3.7%    | 20.6%              | 24.6%           | 19.4%   |
| <i>Average waiting time of ships that had to wait (hours)</i> |                    |                 |         |                    |                 |         |
| Uncontrolled  | 12.3               | 14.5            | 13.7    | 7.5                | 9.3             | 7.7     |
| Equidistant   | 9.5                | 9.5             | 9.3     | 6.2                | 7.2             | 5.9     |
| Stock-controlled  | 7.9                | 10.0            | 8.7     | 3.5                | 3.8             | 2.6     |
| Hybrid  | 8.3                | 9.7             | 11.0    | 5.6                | 6.3             | 5.0     |

#### 6.4. Exhaustive search for the best berthing sequence

In addition to the simple priority scheme for mooring point allocation, we have implemented an enumeration algorithm. This algorithm uses the same information as the simple priority scheme: the pre-arrival notices that are available a number of hours before the actual time of arrival. Rather than looking at just the high-priority ships, the enumeration algorithm evaluates the waiting time for all ships. Every time a ship arrives or a mooring point becomes available, this algorithm determines the best berthing sequence by evaluating all possible berthing sequences. The sequence with the least amount of waiting is then selected as the best possible sequence (in this implementation, we do not distinguish among ship types).

As Table 4 shows, the application of the enumeration algorithm provides a clear improvement over the simple priority scheme, both in the percentage of ships that had to wait and in the average number of hours that were spent waiting. The process that best aligns the arrivals with production (the stock-controlled arrival process) achieves the best results. The percentage of larger vessels that have to wait can be reduced to around 3.5% by the application of this enumeration algorithm.

## 7. Conclusion

The importance of careful arrival process modeling is clearly demonstrated in this paper. 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. An optimization procedure for jetty allocation yields a substantial performance improvement over a first-come-first served allocation, especially in combination with the stock-controlled or hybrid arrival process. Although these results were obtained in a specific case with a relatively high jetty utilization, they are general enough to be appropriate for many port and jetty simulation studies, when the logistical process is directly linked to the production process. The stock-controlled arrival process works well in case

of a limited number of products and a large variety in ship sizes. It does however, not coordinate arrivals of ships for different products. The hybrid process provides an alternative in situations where only limited control over arrivals can be implemented. In any case, as soon as there is some sort of control over arrivals, it should be explicitly incorporated in the model.

Obviously, the challenge in shaping and managing these logistical processes is to realize the importance of arrival processes and to assess which one 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. In this case, brute overcapacity in terms of available jetty facilities, piping and tankage is the only alternative.

## Acknowledgments

The authors would like to thank Ludo Waltman for his solid implementations of the simulation model in both Enterprise Dynamics and DESMO-J and for carrying out the experiments. They also appreciate his many valuable suggestions.

## 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, X.P. Wang, and S. Welch (1996). Modeling crude oil lightering in Delaware Bay. *Interfaces* 26, No. 6, 68-78.
- Arena (2003). Documentation material. Rockwell Software. See also: <http://www.software.rockwell.com/arenasimulation/>
- Asperen, E. van, R. Dekker, M. Polman, and H. de Swaan Arons (2003). 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.
- 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.
- Banks, J., J.S. Carson II, B.L. Nelson, and D.M. Nicol (2000), *Discrete-Event System Simulation*, Prentice-Hall.
- Dekker R. (1999). Simulation of jetty and storage activities for oil and chemicals (1999). In *SimLog, Simulation and logistics in the harbor*, J. Van Nunen and L. Verspui (Eds.). Eburon Delft, Netherlands, 105-116 (in Dutch).
- DESMO-J (2003). Documentation material. Contact: <http://www.desmoj.de/>
- 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.
- Kelton, W.D., R.P. Sadowski, and D.T. Sturrock (2004). Simulation with Arena, 3<sup>rd</sup> edition. McGraw-Hill.
- Kia, M., E. Shayan, and F. Ghotb (2002). Investigation of port capacity under a new approach by computer simulation. *Computer & Industrial Engineering* 42, 533-540.
- Law, Averill M., and W. David Kelton (2000). Simulation, Modeling and Analysis, 3<sup>rd</sup> edition. McGraw-Hill.
- 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). *Modeling and Simulation of logistic systems*. 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.