Proceedings of the 1983
Winter Simulation Conference
S. Roberts, J. Banks, B. Schmeiser (eds.)

THE "TELL-US-THE-ANSWER-YOU-WANT" PROBLEM

David E. Barnett
The Standard Oil Company (Ohio)
663 Guildhall Building
Cleveland, Ohio 44115

This paper describes a simulation model of the Trans-Panama Pipeline. The pipeline is used to transport crude oil between tankers in the Pacific and Atlantic Oceans. The model represents in many senses a classical queuing system, with vessel delay times at the Pacific and Atlantic terminal facilities the critical responses of interest. The paper focuses on the modeling of vessel arrivals. Actual data on vessel arrivals show that interarrival times are exponentially distributed. However, due to human scheduling the number of arrivals within any 15-day planning period equals the expected number. Since the number of arrivals generated by the exponential distribution can deviate substantially from the expected number, an alternative modeling technique is needed. The technique used to model this pattern of arrivals is described in this paper.

1. INTRODUCTION

The Trans-Panama Pipeline opened in 1982 to transport oil from tankers in the Pacific Ocean to tankers in the Atlantic Ocean. Before the pipeline was built, the smaller Atlantic tankers travelled through the Panama Canal to the Pacific Ocean to draw oil from the larger Pacific tankers via a transshipment facility. The pipeline was built in response to vessel delays caused by congestion in the canal.

The physical configuration of the pipeline system is shown in Figure 1. Pacific tankers travel from Valdez, Alaska to Puerto Armuelles, Panama, where they discharge their oil into three storage tanks which feed the pipeline. After settling and measurement, the oil is pumped across Panama into three more storage tanks at the Atlantic terminal (Chiriqui Grande). Atlantic lifters arrive at Chiriqui Grande, draw oil from the storage tanks (after it settles again), and then deliver the oil to ports along the U.S. Gulf Coast.

Pacific tankers arriving at Puerto Armuelles must first capture one of two berths. The berths are available on a first-come, first-serve basis. Once the tanker has successfully captured a berth, a storage tank must be captured for the tanker to discharge its oil. The tank which is free and has the least amount of oil in it is selected. A free tank is defined as one which is not receiving oil from another tanker, is not feeding the pipeline, and is notsettling. If no free tank is available, the tanker may preempt the pipeline. This is done if the tank the pipeline is drawing from has enough room to take all of the oil left on the tanker. Otherwise, the tanker must wait until a tank becomes free. A typical tanker will require more than one tank (it may capture only one at a time, however) since it carries more oil than one tank can hold. Once the tanker has discharged all of its oil, it leaves the system.

Atlantic tankers arriving at Chiriqui Grande must capture one of two sea buoys. These buoys are also available on a first-come, first-serve basis. Once the tanker has successfully captured a buoy, a storage tank must be captured for the tanker to draw oil. The tank which is free and has the most oil in it is selected. A free tank is defined as one which is not servicing another tanker, is not receiving oil from the pipeline, and is not settling. If no free tank is available, the tanker may preempt the pipeline. This is done if the tank the pipeline is filling has enough oil to fill the tanker. Otherwise, the tanker must wait until a tank becomes free. Atlantic tankers have less capacity than a storage tank, so it is possible for a tanker to draw all of its oil from one tank. They may, however, draw oil from more than one tank if necessary (only one tank at a time, however). Once the tanker has drawn all of the oil it needs, it leaves the system.

Atlantic tankers also have the option of travelling through the canal to Puerto Armuelles. These vessels draw directly from the Pacific tanks, thus bypassing the pipeline system. A tanker will take this option if the expected cost of a voyage to Puerto Armuelles is less than the expected cost of a voyage to Chiriqui Grande. There is one berth at Puerto Armuelles dedicated to Atlantic tankers, and it is also possible to use the smaller of the two Pacific tanker berths. The tanker must capture a berth and a tank, as described above for tankers at Chiriqui Grande. These Atlantic tankers may not preempt the pipeline, however; rather, the pipeline may preempt them.

2. MODEL DEVELOPMENT

The physical system of tankers arriving at a port to "discharge" or "lift" oil is a simple queuing problem, similar to the African Port Case Study presented by Schriber (1974). In addition to modeling the arrival and servicing of tankers however, the pipeline flow must also be modeled. The pipeline carries a continuous flow of oil (unless the pipeline must be shut down) from the Pacific tanks to the Atlantic tanks. The pipeline must choose a tank to draw from, or flow into, subject to tank availability. Once the pipeline selects a tank, it will draw the tank empty (Pacific) or fill it full (Atlantic) unless a tanker preempts the pipeline from that tank.

GPSS was chosen for this simulation, mainly because it was readily available on Sohio's CDC computer and it is the simulation language the author is most familiar and comfortable with. (FORTRAN HELP blocks are used to handle some of the more complex decision logic.) The tanker arrival, queuing, and service sequence can be modeled quite nicely with GPSS.

The storage tanks and berths can be represented by GPSS storages and facilities, respectively. The continuous flow of oil through the pipeline, however, is not a natural candidate for a discrete-event simulation language. The author realized that a language such as SLAM, which allows both discrete and continuous event processing, was more appropriate. However, since time constraints did not allow us the luxury of acquiring and learning a new language, the models were developed in GPSS. The pipeline flow is modeled by generating a transaction each hour which removes one hour's worth of oil from a Pacific tank and enters the oil into an Atlantic tank.

Since the physical system being modeled can be thought of as two separate queuing systems (Atlantic side and Pacific side) tied together by the pipeline, the model was developed in phases. First, a stand-alone Atlantic model was developed, followed by a stand-alone Pacific model, and finally the two models were tied together. Since the Pacific side is essentially a mirror image of the Atlantic side, development of the Pacific model was just a matter of reversing some of the Atlantic logic and adding a few concepts unique to the Pacific side. This strategy offered two major advantages over trying to build the whole model at once. First, it provided Sohio's Crude Trading and Transportation Department with a usable model of the Atlantic operations within a time frame acceptable to them. Second, this strategy facilitated debugging and model validation. Logic errors were easier to pinpoint in the separate models than in the combined model where there are many more events interacting.

Rather than describing the GPSS models in detail, this paper will focus on one particular implementation problem encountered when the two models were combined to form the overall system. For the most part, combining the models was straightforward. The main task was to connect the pipeline flow segments of the two models. The separate models had assumed an uninterrupted flow of oil through the pipeline. For example, the Pacific model assumed there was always room in the Atlantic tanks for the oil to flow into, and the Atlantic model assumed that there was always oil in the Pacific tanks to draw. In the combined model, however, these assumptions are not valid. If there is not room in the Atlantic tanks, the pipeline must shut down, and therefore no oil can be drawn from the Pacific tanks. Similarly, if there is no oil available in the Pacific tanks, the pipeline must shut down and no oil flows into the Atlantic tanks.

Connecting the two pipeline flow segments was accomplished with a minor amount of effort. The combined model was, however, producing extremely high vessel delay times, much higher than the stand-alone models and much higher than the "gut feeling" of the sponsor. Initially it was felt that high delay times were caused by the pipeline shutting down frequently because of inadequate tank capacity on either or both sides. As many as three additional tanks were added to the model on each side, however, and still the waiting times were too high. Finally, while closely examining the arrival patterns of Atlantic tankers over a period of 30 simulated days, the sponsor commented with surprise that 23 tankers arrived rather than the expected 18. The author responded that the exponential distribution of interarrival times could easily generate from 10 to 26 arrivals (roughly a 95% confidence interval) in any given month. The sponsor maintained that due to scheduling, the number of arrivals within any 30-day period would be within one tanker of the expected number. It was now apparent that the total capacity of tankers arriving on the Pacific side in any month has to match the total capacity of tankers arriving on the Atlantic side for the system to function properly.

This led us to re-examine the assumption that interarrival times are exponentially distributed. A distribution of the interarrival times of dischargers for the three months preceeding the development of the model is shown in Figure 2. Clearly these interarrival times resemble the theoretical exponential distribution. So here is a situation where interarrival times appear to be exponentially distributed, while the number of arrivals in a given unit of time equals the expected number of arrivals.

An initial attempt to solve this problem was to pre-generate one month's tankers at a time, on both the Atlantic and Pacific sides. This was done using exponential interarrival times but then scaling those times so that exactly the expected number of tankers arrived during the month. For example, if the unscaled interarrival times sum to 37 days, each interarrival time is multiplied by 30/37=0.81. The interarrival times are still exponentially distributed, and the desired number of tankers arrive each month. However, delay times were still too high after this scaling procedure was implemented. Even though tanker capacities were matched every 30 days, there was still too much variation in arrivals at any given point <u>during</u> the 30-day period. We then tried matching arrivals on <u>both</u> sides every 15 days and this produced results in line with the sponsor's gut feeling and the historical delay times experienced by tankers using the Pacific transshipment facility.

But what was so magical about 15 days? Why not 20 days or 10 days? The sponsor was getting a little nervous about the ease with which we could alter vessel delay times by altering the "exact-match-period" (see Figure 3). If the model were to have any credibility at all, some real-world basis for deciding the was "exact-match-period" needed. Conversations with the sponsor led us to think of the model's "exact-match-period" as a proxy for the skill-levels of the human schedulers in the real system. The better the scheduling, the better the balance between arrivals on both sides and therefore the more frequently the arrivals (in terms of cumulative tanker capacity) would be in exact balance. The real-world data on ship arrivals were examined to check the frequency with which the Atlantic and Pacific tankers were evenly matched in terms of cumulative capacity. This data showed that on the average, cumulative capacities came into balance every 14.5 days (see Figure 4). With this solid and defendable real-world basis, we were able to confidently proceed with an exact-match-period of 15 days. Vessel delay times fell into line with expectations, and the model began to help Sohio analyze the impact of the various operating procedures and physical configurations.

3. SUMMARY

This paper has described a real-world system in which interarrival times appear to be exponentially distributed yet the expected arrival rate in any given planning period must be precisely met. This situation was not detected until the two separate models were tied together. In the separate models, a theoretical exponential distribution worked fine. The problem occurred when these two models had to interact with each other. The technique used to model this situation is by no means sophisticated, but it seems to model the real-world arrival pattern extremely well. The author suspects that this arrival pattern occurs in many real-world systems, and feels that the technique described in this paper can be used to model these patterns accurately.

ACKNOWLEDGEMENTS

The author would like to acknowledge Dr. Robert Obee and Mr. James Kantelis for their efforts in the development of this model. Dr. Obee, supervisor of Sohio's Management Science Group at the time the model was developed, provided invaluable insight and suggestions concerning the GPSS model. Mr. Kantelis, the primary sponsor in Sohio's Crude Trading and Transportation Department, contributed several ideas without which the model would not have been successful.

REFERENCES

Schriber, T. J. (1974), <u>Simulation Using GPSS</u>, John Wiley & Sons, Inc., New York, New York.

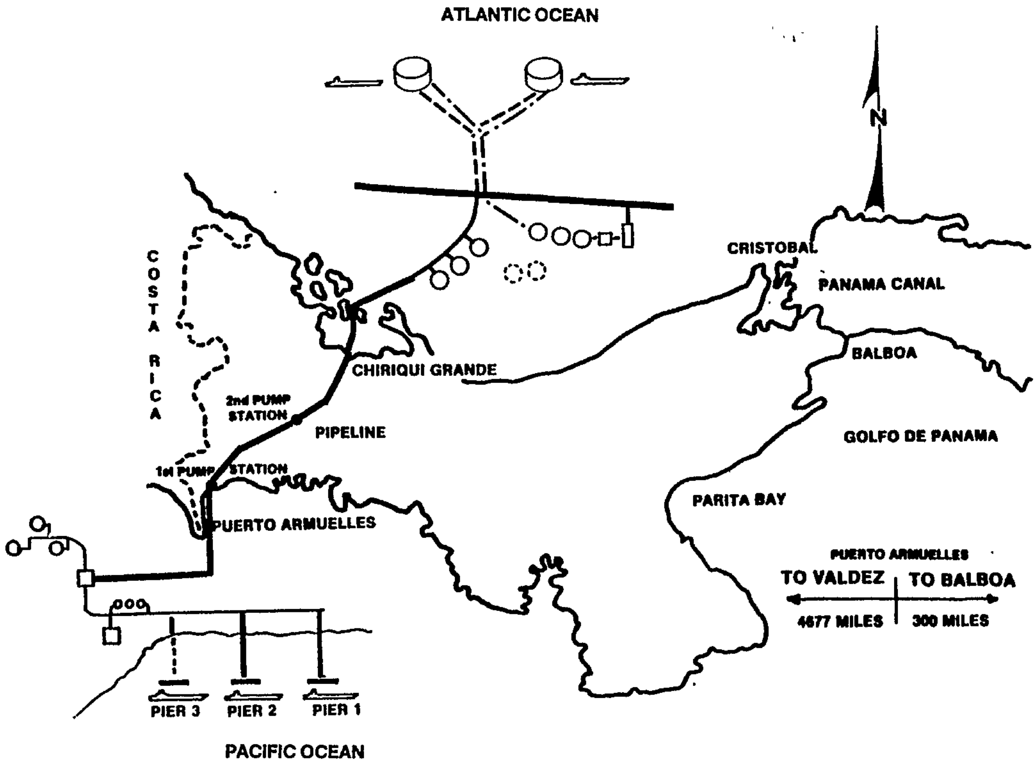

Figure 1: The Physical System

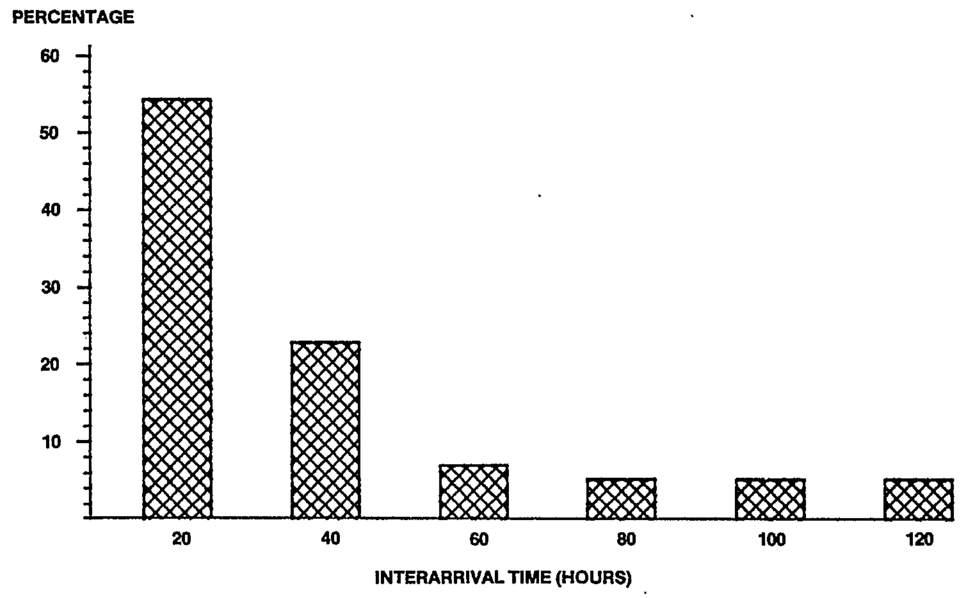

Figure 2: FREQUENCY DISTRIBUTION OF VESSEL INTERARRIVAL TIMES

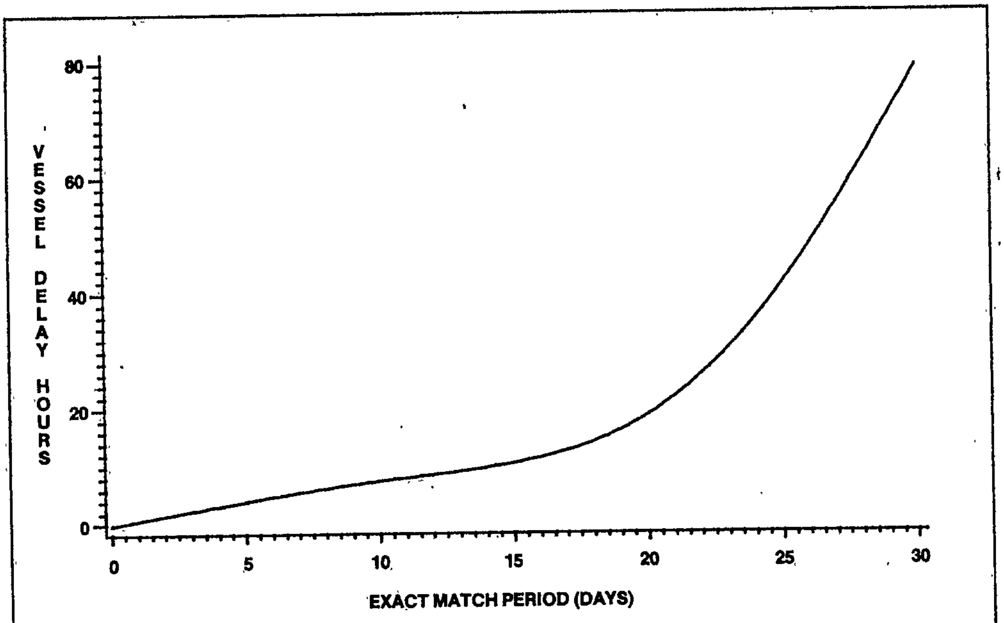

Figure 3: The TELL-US-THE-ANSWER-YOU-WANT Problem

VESSEL ARRIVALS (MBBLS.)

Date	Cumulative Pacific Capacity	Cumulative Atlantic Capacity	Net Capacity
12/31	—	—	0
1/01	1000	—	1000
1/03	1000	600	400
1/04	2000	600	1400
.	.	.	.
.	.	.	.
.	.	.	.
1/11	4000	3400	600
1/13	4000	4200	-200
.	.	.	
.	.	.	
.	.	.	
.	.	.	

13 days

$\bar{X}$ = 14.5 days

Figure 4: Exact Match Period