Proceedings of the 1993 Winter Simulation Conference
G.W. Evans, M. Mollaghasemi, E.C. Russell, W.E. Biles (eds.)

WATERFRONT CAPACITY-PLANNING SIMULATIONS

William Heath

Shell International Petroleum Company Ltd
Shell Centre
York Road
London SE1 7NA, UNITED KINGDOM

ABSTRACT

This paper describes the purpose-built simulations of waterfront systems of some of the Shell Operating Companies. It examines the assumptions behind each of the simulations, especially with respect to the vessel movements within the port. It also looks at the modelling of fluid flows upstream and the ship movements downstream. Finally, the paper summarises the current thinking within Shell Companies of whether purpose-built simulations are still a valid approach.

1 INTRODUCTION

As part of the Shell Group's business, a vast amount of materials is shipped around the world. Of this, by far the largest proportion is made up of gas, crude oil and their refined products. The shipping element of each project involves a great deal of planning. Most important is the planning of the waterfront, to ensure that the capacity is correct. There must be enough berths to allow for the free movement of the liquids, and the storage facilities must be large enough. The capacity planning is so important that, if there is either over- or under-investment in the infrastructure, the entire project may no longer be profitable.

This paper describes the capacity-planning models built by Shell International Petroleum Company for some of the Operating Companies involved in the large-scale shipping of fuel liquids. There are currently six models being used. The main characteristics of a waterfront system and the similarities and the differences between the models are examined. Finally the paper discusses whether the current method of developing purpose-built models for each operation is still a valid approach.

2 OVERVIEW OF WATERFRONT SIMULATIONS

Waterfront capacity planning is required whenever there is a substantial change to the existing system or a new waterfront layout is being proposed. In either case, there are a number of factors that are of great interest to the planners: jetty occupancy, volume of product shipped and demurrage.

The definition of **jetty occupancy** varies from model to model. For clarity, the total time a berth is occupied is broken down into different time categories. These can include:

- steaming in and steaming out times;
- times waiting for tides and weather;
- the time waiting for enough stock be available;
- the time spent loading;
- the time spent on any administration.

Some or all of these are reported in all six of the simulations.

Demurrage is the penalty paid of delaying a vessel. When a ship is chartered for a voyage, a time is agreed between the charterer and ship owner. This represents the time required for the vessel to queue at the loading port, sail in, load and sail out. If the vessel is delayed at the port for any reason and it spends longer than the agreed time, the harbour authorities must pay a penalty, usually a price per hour over the agreed time. These penalties are high.

Five of the models expand the system to include the production and transportation of the materials to the loading port. Their outputs include flow rates through the different parts of the system and the amount of time a portion of the upstream is broken or undergoing maintenance. However, by far the most common outputs from this part of the model are the frequency and duration of the tanks being full. (This is called **shut-in**.)

One model also simulates the voyages of the vessels. There is also the possibility that a second model, which currently just looks at the port movements, may have to be expanded to model the voyages. When the voyage movements have been examined, the outputs of most interest have been the ship movements, delays on the voyages and volumes delivered.

One final complication occurs in two of the models, which simulate waterfront for alumina plants. Here, the waterfront is used both for unloading raw materials and for shipping alumina out. In these cases, the chief concern of the planners is shipping movements on the waterfront. This situation would also be applicable to refineries, where crude oil is delivered and the refined products are exported.

3 THE PORT

The main centre of attention for capacity planners is the port itself and the movements within it. It is within this area that any over- and under-capacity will become apparent. Figure 1 gives a pictorial view of the elements comprising a typical port. Outside the port is an anchorage point, at which the vessels will wait until given permission to enter. Between the anchor and the berths there is the sea lane. Usually only one ship can be moving on the sea lane at any time. In one of our simulations, this is made more complicated as there are four legs to the sea lane, different for in-coming and out-going vessels.

The berths themselves may be either alongside the shore or may extend kilometres out to sea. This very much depends upon the products that are passing across them. Two of the simulations model more than one product being shipped out of the port. In these cases, berths may not be equipped to load all the products. This is often something that the planner will wish to experiment with, since the more flexible the berth is, the more expensive it is to build and maintain. In some cases, the loading equipment may be shared between berths.

The products themselves are in storage tanks. The method of filling the tanks does vary, and is dealt with in Section 4.

The activities within the port are illustrated in Figure 2. Ships arrive at the loading port and join the entry queue. When the schedulers give permission for the vessel to enter and the circumstances are correct, i.e. the weather is fine, the tides are right and there is no other ship in the vicinity, it will steam in. In one simulation, an added constraint is the availability of tugs.

A further constraint upon the entry of the vessel is the amount of the product in storage. In Figure 2, the constraint is bracketed. This indicates that the stock levels are either considered before the vessel sails in or after it is moored. The latter is common in simulations where there is only one product leaving the port.

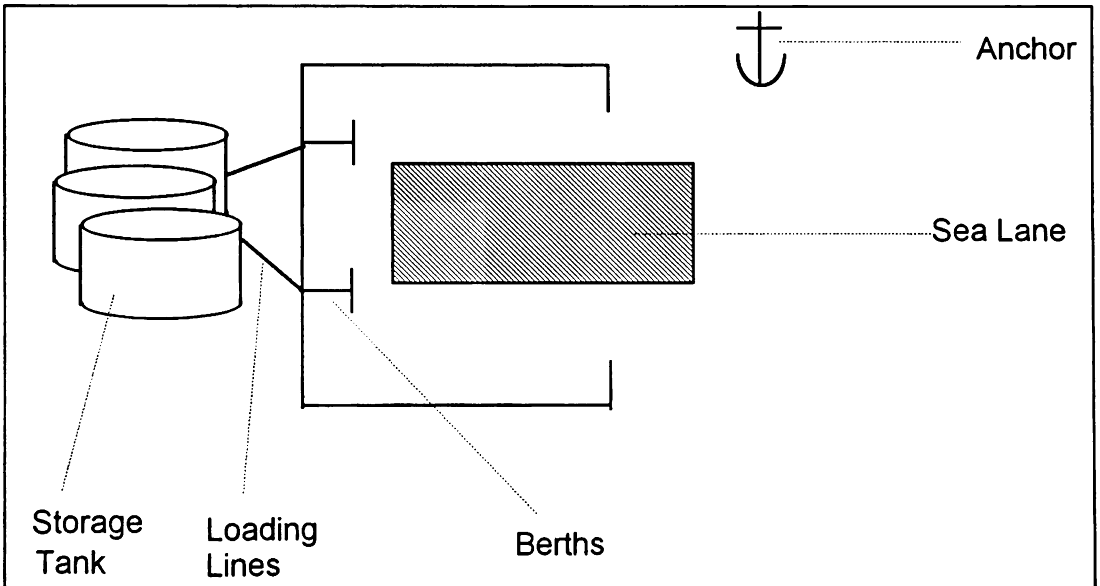

Figure 1: The Loading Port

The ship then moors and any administration is undertaken. The administration is usually modelled as part of the mooring activities, as it is assumed to last a constant time. It would only be modelled separately when other activities are dependent upon its completion. For example, in one of the models with two jetties loading the same product, the emergency valves must be checked before a ship can be loaded. Any other ship being loaded at that time must stop until the valves are checked.

The loading line is then attached to the vessel and loading starts. If the weather becomes too severe, loading must stop. For safety reasons, the ship must unmoor and move away from the jetty until the weather is calm enough for loading to resume.

On completion, the vessel waits for the weather and tides to be right and the sea lane to be free, before departing.

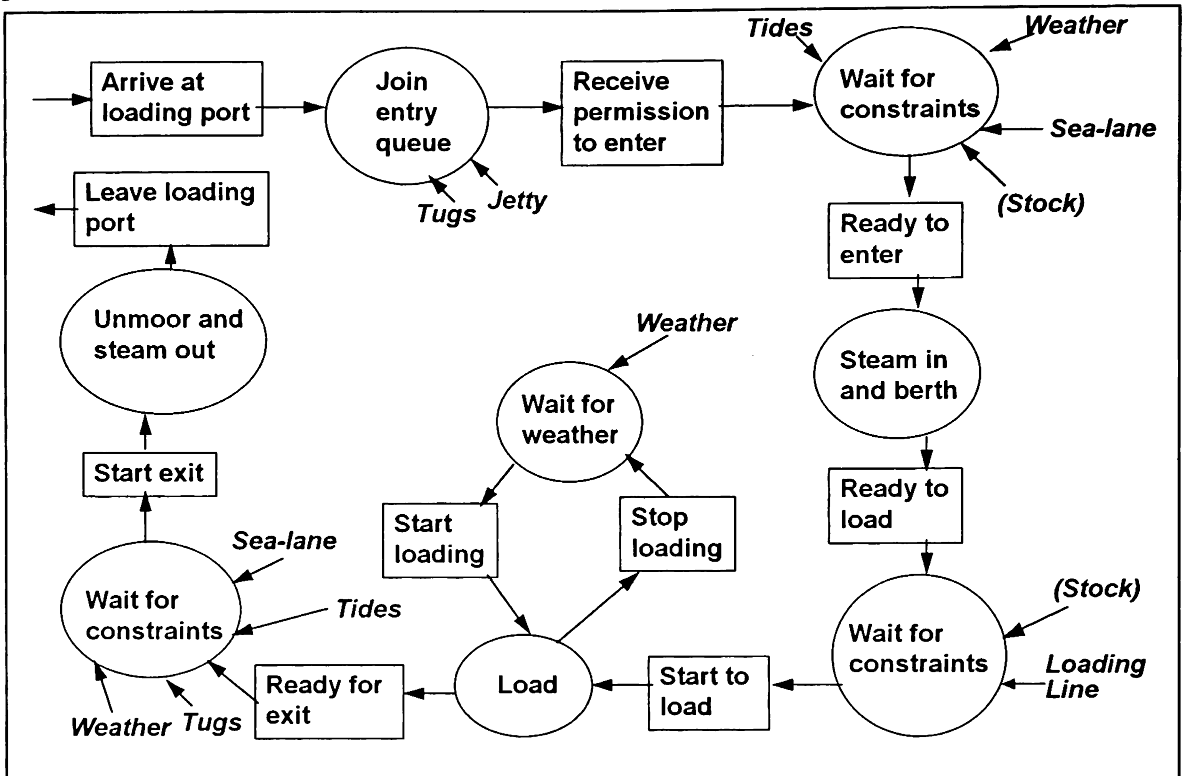

Figure 2: Activities Within The Port

3.1 Tides

In all of the simulations, tides are modelled on a 12.5-hour cycle between high water and high water. Ships are then assigned 'windows' within the cycle, during which time the ship must moor and unmoor. These windows are user inputs. The size of the vessel is the important factor in deciding the windows, with the larger vessels often having to berth at an exact moment on a rising tide.

In cases where the tides play a critical part in the vessel movements, accurate information on the high- and low-water times is obtained from the Maritime Office, a British government institution.

3.2 Weather

The models vary most in the handling of the weather and the impact of different weather patterns upon the activities. We have used three methods to model weather:

- a Markov Chain;
- a certain amount of bad weather per day or voyage;
- a distribution of delay and mean time to bad weather.

The large number of methods of simulating weather patterns reflects the variations in data that are recorded. For instance, in the North Sea, the data records wind speed and the length of time it lasted.

The North-Sea data lends itself to a Markov Chain. There are three weather conditions: calm, rough and stormy. How each weather state affects the activities depends very much upon the individual waterfront. An example may be: during calm weather all activities may occur; during rough weather no ships may move, but loading may continue; during stormy weather, nothing may happen.

In one model, where the data available are not as detailed as those used in the Markov Chain, a weather delay is assigned to each vessel cycle in the port. However, the most common way of modelling is to use a mean time to bad weather and a delay distribution. Both the mean time and the distribution would be derived from available data. These data are often shown in the form of weather delay by vessel.

In all cases, the time spent delayed because of bad weather is recorded. Within the port boundaries, this usually means that the vessel is either on, approaching or leaving a berth and this time is included in the breakdown of the jetty occupancy statistics.

3.3 Scheduling Entry To Port

The scheduling of the vessel's entry into the port depends very much upon the number of products being handled. In the simplest cases, where there is only one product leaving the port, ships have always been scheduled on a first-in, first-out (FIFO) basis.

The level of complexity rises when several products are being shipped. This is a feature of five of the six simulations. Here again, the vessels are scheduled on a FIFO basis, but usually there is a constraint that a product can only be loaded on one jetty at a time.

However, there are two factors that may override the FIFO rule. One is demurrage and the other is stock level. If the stock level in the tanks is nearing capacity, then it is urgent that the stock is loaded as soon as possible, avoiding a shut-down in production.

The other factor, demurrage, is far more complicated and will affect the actual scheduling rules. For instance, if a vessel is late, then the demurrage 'clock' will not start until permission is given for it to enter. Similarly if the vessel is early, the demurrage 'clock' starts either when it is given permission to enter, or at its scheduled arrival time, whichever is sooner.

Owing to the complexity of demurrage rules, the scheduling process is usually overlooked for capacity-planning purposes, but included in the calculation of the demurrage charges. This does not effect the activities significantly and reduces the complexity of analysis.

4 UPSTREAM MODELLING

There are various options open to the modeller on how complex the simulation upstream of the port should be. Below are outlined the three methods used in these simulations.

In reality, storage systems are complicated. There are usually a number of small tanks for each type of product, which can either be filling or discharging, but never both simultaneously. For the simulations, we have been able to assume that there is only one large storage tank, which both fills and discharges.

4.1 Constant Flows

The simplest method of modelling the flows into the port storage tanks is to use a flow rate, which may be either 0 or a constant. The product will flow into the tank until it is shut-in, when the flow stops. Input into the tank resumes when the level has dropped to a certain level below the tank capacity. This is a common approach when several products are being exported from the same port.

In more complex cases, all tanks stop filling if one tank shuts in. This is a more realistic approach, as usually all the products being exported come from the same source, which would have to stop producing if any one product's tank shut in. Figure 3 illustrates this set-up.

4.2 Flows And Processes

Figure 4 is an illustration of a possible upstream configuration, where the flows and processes are modelled in greater detail. This type of configuration is used in two of the simulations. The system is typical for natural gas. The gas is extracted at a number of points in a field, transported to a central point, where it is liquefied

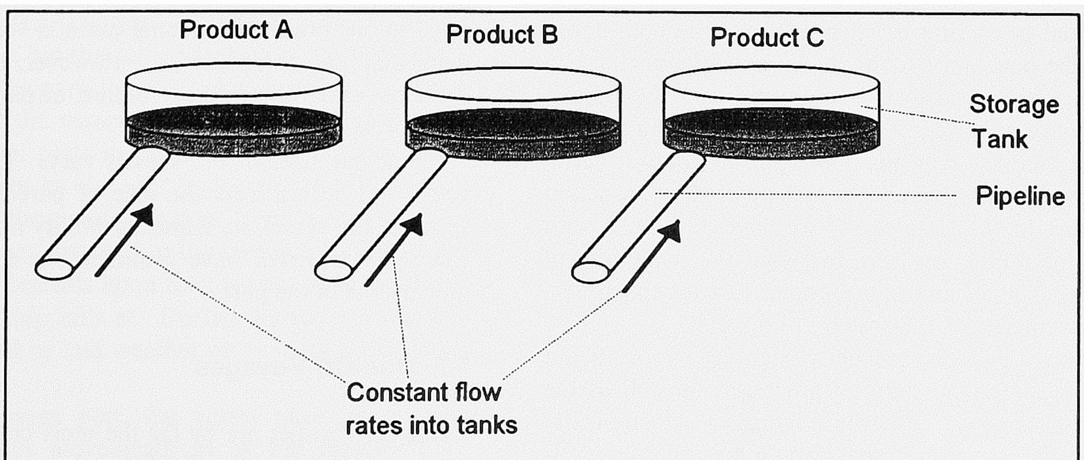

Figure 3 - Constant Flow

and then transported to storage. In Figure 4, two different processes are illustrated, the products of which are fed into common storage. If the processes delivered different products then different storage areas would be necessary.

With this configuration each part, or element, of the system has some likelihood of failing. This may lead to a partial reduction or complete stoppage of flow through the element. With a large number of elements, each affecting the flow rate into the final storage tank, the process of capacity planning becomes a complicated one.

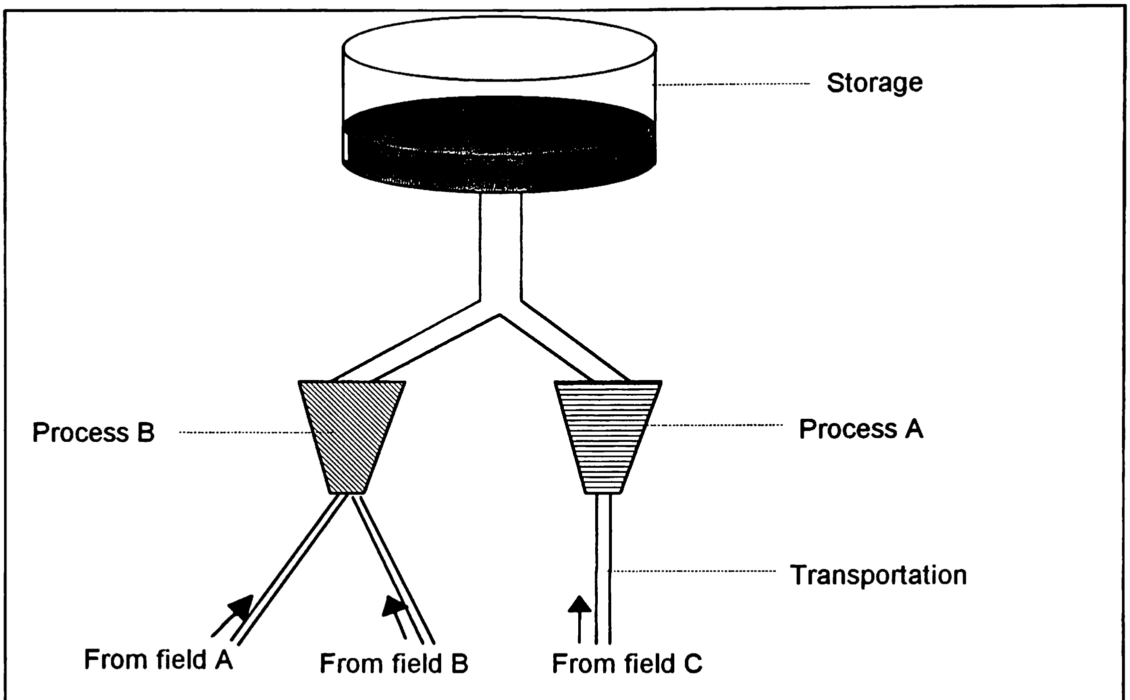

Figure 4: Flows And Processes

4.3 Inputs And Outputs

The final possibility is when products are both delivered to the port and exported. An example of this is a refinery, where feedstock is delivered and the refined products are exported. The rate of production of the exported products depends both upon the rate of delivery of the imported products and upon the production processes. In the simulations we have developed with this configuration, we have bypassed the modelling of the refining itself and concentrated just on the waterfront movements, as illustrated in Figure 5.

There have, however, recently been discussions about the modelling of the refining processes, as well as the waterfront systems. This would increase the complexity of the model, but would give a more accurate representation of the ship movements.

5 DOWNSTREAM

Whereas the purpose of modelling upstream is to achieve a more accurate picture of how the product arrives at the waterfront, modelling the downstream system is done to get an accurate understanding of vessel arrival patterns. In Shell International Petroleum Company, we have used two methods to simulate the arrival patterns: one-off voyages and round voyages.

5.1 Vessels

As the majority of the products shipped by Shell Companies are fluids, the vessels used are usually dedicated to one product. In some cases, a vessel may have more than one-compartment. However, each compartment can only carry a fixed volume, or parcel, of a dedicated product.

The actual sizes of the vessels vary. The size of the vessel not only affects the size of parcel but also the speed of the vessel, its manoeuvrability and its draught. The latter two will have ramifications for the vessel's movements in the port.

5.2 One-Off Voyages

One-off voyages are by far the most common method used to model the arrival of vessels. Vessels arrive at the port at a pre-defined time, go through the loading process and depart. This assumes that vessels will be chartered only for that voyage. As the voyages of the ships are not modelled, the only relevant attributes of the ships are the parcel size(s) and the product(s) it can carry.

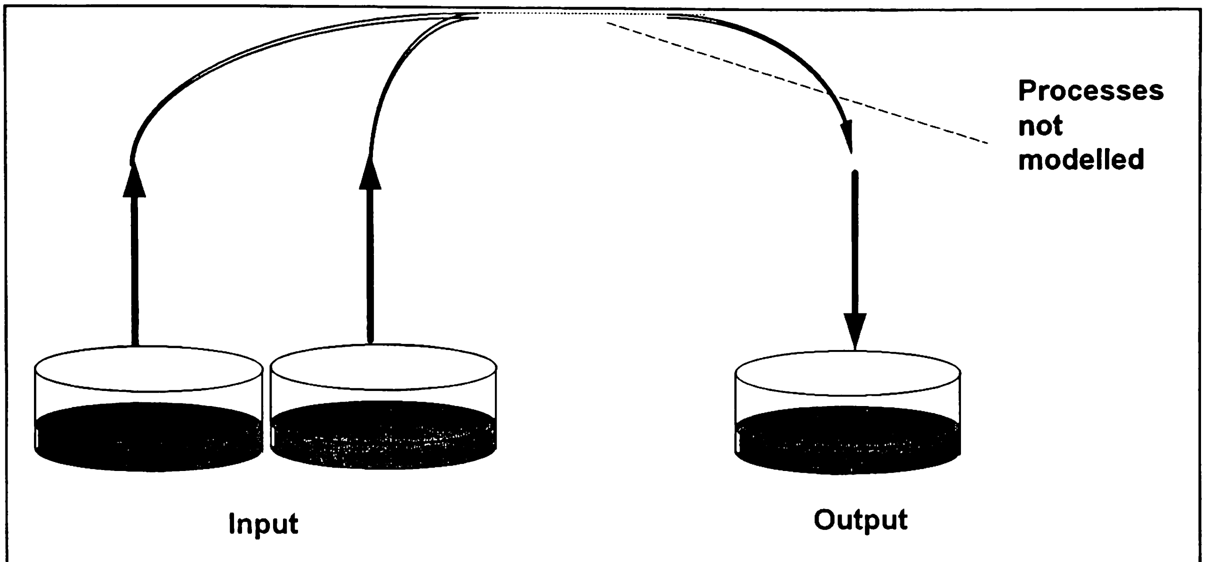

Figure 5: Inputs And Outputs

Ships are generated by a number of algorithms, each algorithm creating a vessel that can carry a certain product(s). The inter-ship time depends upon the flow rate into the tanks and the number of vessels of a similar type already in the queue. The parcel size of the ship is then drawn from the distribution of different parcel sizes. Finally, it is given an expected time of arrival (ETA) and an actual time of arrival.

5.3 Round Voyages

Round voyages need a dedicated fleet, sailing from loading port to customer port and back. In these cases, the vessels are either owned by the Operating Company or are on long-term charter. Therefore, demurrage is no longer an issue, but the movements of the vessels are.

On leaving the port, the schedulers must decide which customer port the vessel should sail to. Usually there are a number of ports to choose among. The customer demand is input by the user as a number of cargoes a month, which the simulation must try to meet.

The vessel then sets sail to the customer. The length of the voyage depends upon the distance to the port and the speed the ship sails at. During the voyage, the ship may be delayed by bad weather or by mechanical faults on the ship.

At the discharge port, the vessel must go through much the same movements as at the loading port. However, there is one notable difference. At the loading port, ships are delayed by the movements of other ships. At the customer port, no other vessels are modelled; so a delay is automatically included into the time each ship is alongside the berth.

After the customer port, the vessel either sails back to the port or goes onto another discharge port to deliver more of its cargo.

If this method of modelling vessels is used, then the port movements are only one category of a number of ship movements that would be of interest to an analyst.

6 CONCLUSIONS AND DISCUSSION: PURPOSE-BUILT VS. GENERIC MODELS

The majority of the waterfront simulations developed by the Business Consultancy department have been purpose-built, specially designed for one specific waterfront system. In a couple of cases, the simulations have been developed for one type of system, such as oil or gas transportation, but even then the models often need adjusting to allow for the peculiarities of the specific system being modelled.

The cost of a waterfront simulation is high and a substantial amount of the money is spent on the actual development of the model. The Business Consultancy department is now finding that it is dedicating a significant amount of its resources to maintaining and developing a relatively small number of simulations, which are being used to plan very large operations.

Business Consultancy could add significant value to Shell Group's business by making its skills more accessible to the capacity planners around the world. However, increased demand would lead to a resourcing problem. In order to overcome this problem, a method of producing solutions and advice, which is less time-consuming, needs to be found. Currently, the possibility of building a development environment with a graphical user interface (GUI) is being investigated. Behind the GUI would be a 'generic' waterfront simulation tool.

Our experience with such systems has mainly been in the Systems Dynamics field, using packages such as Stella and iThink for the Apple Macintosh. Using these packages, the amount of time spent explaining the model to the clients and validating it with the system experts was greatly reduced.

However, a similar application for waterfront capacity planning would need a great many more 'objects' than the System Dynamic packages can handle. There is commonality between waterfront simulations, which could be exploited, but the idiosyncrasies of each individual system would mean that any 'generic' model would need tailoring. Maintenance would still be required for each model, but the lead-time from the initial definition of the model to having a working simulation would be reduced.

AUTHOR BIOGRAPHY

WILLIAM HEATH received a Masters degree in Mechanical Engineering and Management from Birmingham University, England. He now works in the Business Consultancy department of Shell International Petroleum Company in London, working closely with clients in the natural gas and oil industries, examining the feasibility of proposed green-field and expansion projects. The work includes financial modelling, strategy development and risk analysis, as well as capacity planning. He is interested in the interface between clients, model and modeller and how the building large models can be made into a more efficient process.