SOME PROBLEMS IN THE SIMULATION OF VERY LARGE HYDORDYNAMIC SYSTEMS

By

A. B. Nelson
The Rand Corporation
1700 Main Street
Santa Monica, California 90405

ABSTRACT

Large scale continuous simulation systems place greater demands on computer resources than most computer applications. The need for extremely large amounts of high speed core and CPU time can tax even the largest and fastest machines. In addition, the high cost of each simulation experiment places a great responsibility on the people involved in the design, implementation and use of such systems.

This paper describes approaches to dealing with the difficulties inherent to the operation of several simulation programs including both 2-D and 3-D estuary models (used for the study of tidal dynamics and water quality) and global models of ocean and atmospheric circulation. The computer requirements are discussed as well as some of the data management problems. Various procedures for simplifying the use of the programs and for validating and analyzing simulation results are considered, including graphic display.

INTRODUCTION

Computer simulation is valuable in identifying or isolating problems in various types of real systems and in formulating solutions to the problems. The computer also has an important role in the study of the effects of proposed solutions to avoid unanticipated negative effects.

Systems developed by humans such as supply, communications, or transportation are usually well enough defined and understood that a change will not have far-reaching, unexpected side effects. Natural systems, such as atmospheric circulation or tidal dynamics, are usually far more complex; an attempt to improve one part of the system may result in disastrous changes to another part. For example, a breakwater intended to protect a harbor may result in undesirable silting or other damage. On a global scale, a variation in large scale winds or ocean currents might cause widespread changes in climate with resulting disruption of expected rainfall and other vital weather factors.

It is, therefore, of considerable interest to study the implications of a change that may cause irreversible damage to an area, whether it is a small bay or the entire world. For small areas such as bays and estuaries, physical (hydraulic) models such as the nine acre, $15 million enclosed model of Chesapeake Bay constructed by the Army Corps of Engineers may be used. Another approach to relatively small areas, and one that is necessary for large studies of ocean dynamics and global atmospheric circulation, is computer simulation.

This paper will refer specifically to experience gained from atmospheric, oceanic and estuarine models; however, many of the observations are applicable to other models that use continuous as well as discrete simulation.

EXPERIENCE AT RAND

Rand has been involved in environmental simulation since about 1964. The simulations have included a global atmospheric circulation model (at one time modified for the Martian atmosphere), three global ocean models, and both two-dimensional and three-dimensional estuary models. The modeling effort has used a variety of computers including, among others, the IBM 7094 and 360/91, CDC 6600 and 7600, Philips P1400, and the PDP-10.

Rand is currently running a general circulation model on the ILLIAC IV located at Ames Research Center, Moffett Field. We have remote access to the ILLIAC IV over the ARPANET. Not only is the use of the simulation controlled over the network, but graphic output is retrieved over it (see Fig. 1). We have been using the ARPANET to communicate with computers at UCLA, ISI (Information Sciences Institute), and IAC (NASA Ames) for some time and have found that the network provides an invaluable link to remote installations. Since this paper is concerned with simulation per se, I will not go into the difficulties encountered in using the variety of computers and systems interfaces necessary for remote computing.

In studying changes in large scale phenomena such as weather patterns, the experimenter may be concerned with seasonal or even annual variations. This requires large amounts of computer time. Some of the Rand experiments using global models have required dozens of hours of CPU time on an IBM 360/91 at UCLA. Even with the use of the ILLIAC IV, the time for computation will be great since longer term effects will be examined.

In contrast, models of bays and estuaries study a considerably smaller area than the global models and the time to be simulated is more likely to be in terms of days rather than months. Even so, the detail need to study tidal dynamics and water quality is such that the model uses considerable time and, ordinarily, more data are required than for the global models. Some experiments using the Rand three-dimensional estuary model used 2200K eight bit bytes of high speed core. Virtual storage has not been available on the computers we have used and programmed data swapping is quite expensive, so that the models have been run entirely contained in high speed core whenever possible.

At present, the estuary models can provide extremely detailed information concerning the subject area. The global models, on the other hand, are used to show the gross effects of major changes, such as the modified atmospheric circulation resulting from a change in ocean temperature over a large area.

GENERAL DEVELOPMENT

The development of a computer simulation for research is frequently an open ended project. It continues, with improvements and additions to the model, until the developers give it up or the funding stops. During the evolution of the model, a cycle of design (redesign), test, verification, and improvement is established. This cycle is repeated many times as the sensitivity and accuracy of the model is refined.

Occasionally, a version of the model will satisfy the requirements for applications. At this time, the model can be "frozen" and should be reasonably well documented so that it can be put to practical use. The design cycle will continue while the "frozen" version of the model is used for experiments. When several "frozen" and developmental versions are in use, it is very important to avoid confusing one with another.

As may be expected, the cost of each experiment is high. This places a great responsibility on anyone involved in the development and use of the system of programs used to simulate the physical phenomena.

DESIGN AND IMPLEMENTATION

The design of the system includes the development of differential equations describing what are believed to be the dominant factors influencing the physical system. Occasionally, potentially significant factors are not well-defined in nature so that, because of cost and the uncertain importance of their contribution to the system, they are not incorporated in the model. The careful selection of appropriate physical characteristics is crucial.

The decision as to what numerical techniques to use in the solution is also very important since the stability, speed, and accuracy of integration are dependent on the numerical process.

In order to test the validity of the basic design, test runs must be made and compared with known patterns of behavior of the model environment. In the case of relatively well-defined systems such as estuaries, simulated results can be compared with observed velocities and water levels at specific locations. For example, in modeling Jamaica Bay, Long Island, New York, the standard deviation of the difference between observed and computed water levels at test stations was about .05 feet during several tidal cycles. The variation in water level during each cycle was over five feet.

There is considerably less precision in testing global models since observed global data are always incomplete and frequently inaccurate. We must use a more general approach and first eliminate obvious anomalies and then apply tests for reasonableness. For example, among other questions we may ask if an ocean model demonstrates some known characteristics of water build-up in certain areas, and whether major currents such as the Gulf Stream and Japanese Current (Kuroshio) evolve during the run. In an atmospheric circulation model we look at rainfall, the location of typical high and low pressure areas, and for the development of cyclone tracks along their known paths. I would like to emphasize that there is no 'correct" weather pattern, except in a very general sense, as can be attested to by any weather forecaster.

In addition to checking the development of weather or current patterns for qualitative reasonableness, we may apply some quantitative tests to global models. Numerical comparisons of computed and observed values of such indicators as global and zonal means (as illustrated in Figs. 2,3) may show trends and detect problems. Various statistical analyses may also be used to assist in calibrating the model.

Graphic display is generally accepted as necessary for the examination of the enormous amounts of data generated by large simulations; however, use of graphic display for design validation is frequently overlooked. When we limit graphics to the final presentation of experimental results, we neglect a powerful tool for the identification of design and programming problems. A graph or chart, drawn using data obtained during checkout, may be examined much more easily and quickly than arrays of thousands of printed numbers. Figs. 2-5 show the use of graphics to compare observed and computed temperatures and pressures from the Rand global atmospheric model.

The design of the system of programs for the simulation must not only satisfy the physical requirements of the system to be modeled and assist the human designers and users, but must also consider the program's own environment--the computer. Computers are, after all, electro-mechanical devices given to occasional failure. For this reason, we should include safeguards such as restart capability. Preferably, we should provide options to either restart at a given time (by permanently saving the necessary data at predetermined times) or to automatically restart at the last available checkpoint by simply reloading without user intervention. The automatic restart capability should use two sets of checkpoint data which are rewritten alternately, so that the program can be restarted even if a failure occurs while one set of data is being written.

As well as providing for restart in the case of catastrophic machine failure, we should check the computed data during the course of the experiment. Variables may be tested to make sure that they are within a reasonable range and that successive values are compatible with a realistic rate of change. These tests will show possible computer failures (for instance in data transfer) and will help in detecting input errors.

DATA PREPARATION

The preparation of data for a model that may require tens of thousands of numbers is an awesome project. The necessity to provide terrain depths or elevations, as well as coefficients of various types at each of perhaps ten or twenty thousand points, can stun the potential user. Not only will selecting the values be overwhelming, but the realization that one small mistake may result in the loss of thousands of dollars worth of computing may panic the experimenter. No matter how elegant or accurate a model may be, it is worthless if one's initial contact with it results in confusion followed by a hasty retreat.

While it may not be possible to make data preparation pleasant, steps should be taken to simplify the job. A data preprocessor can assist in organizing and understanding the input data. For instance, it can perform such simple tasks as assigning, whenever possible, reasonable default values to arrays. Obviously, variations must be allowed in the field of data, but should only be required as input at the specific points where values differ from the default.

Just as in the simulation itself, the preprocessor should have built-in tests to warn the user if some values are outside of a reasonable range or if a seemingly excessive variation occurs in an array of data. A complete listing of all input values, along with as complete an analysis as possible, should be provided to the experimenter, who should carefully study it before the simulation is begun. It can save thousands of dollars and much embarrassment. Again, I strongly suggest the use of graphic display. Charts that show isocontours of terrain and values of coefficients along with graphs of time varying data are extremely valuable to the user before the run.

DATA MANAGEMENT

Now that the design and implementation of the model have progressed to a point that we have a useful program and it is possible for us to use it with a minimum of trauma, what happens to the huge amount of data generated?

The experimenter seldom knows beforehand exactly what information will be required to adequately analyze a run. A study of one of the variables may show the need for an examination of other parameters. For this reason, large amounts of data are usually saved; for example, the atmospheric circulation model at Rand used eighty IBM 2314 disk packs to store the output data from about twenty experiments. Access to this amount of data would present a problem even if it were to be used in some consistent order.

During the simulation, the data for a given time are computed so that each variable is contained in an array that represents its value at every point. From these arrays, data may be saved at selected points; however, a more general post-processing capability requires saving the full arrays. This ordering of data (grouped by time and sub-grouped into each variable given at all locations) is fine for charts of a variable over the whole field at some time; however, for an analysis involving several variables at each point or requiring a time sequence of variables, a different ordering may be more effective. In addition, for comparisons of different experiments (contrasting some combination of type of variable, location and time span), we may prefer still other data sequencing.

The data management problems are magnified by the fact that not all experiments will save the same data. Because of model changes or a difference in the importance of some parameters to a study, the data handling procedures must allow for different formats for each experiment. To compare experiments, we may need separate programs to locate the data and then to manipulate them to provide the desired values for each experiment. Confusion between different data sets and processing programs can occur just as it can between various versions of the simulation itself.

There are, of course, cases in which large amounts of data cannot be saved. With the use of some extraordinary computers, it may be cheaper to save only what is currently needed and then rerun the simulation if more data are required. This certainly simplifies the data management, but may mean rerunning several experiments to make some necessary comparisons.

DATA DISPLAY

Examining stacks of computer printout several feet high is enough to cool the enthusiasm of the most dedicated researcher. To avoid being the instant benefactor of the paper drive (and in order to save trees), the researcher should judiciously select the printed output. With a usable data management system, only those data of immediate interest need be displayed. Other data may be retrieved later as desired. The use of time histories of selected variables at a few sensitive points is more concise than full arrays and can be extremely informative.

Comparisons of the effects of different program changes and environmental modifications will need to be made. The resulting differences in such characterisics as velocities of winds or currents, pressure patterns and pollutant dispersion should be used to compare the results of each change with the control run (the study area as it exists) and to show the difference in the effects of alternate modifications. In addition to simple differences, studies may be made to determine the statistical significance of the changes.

Graphic display presents the results in the most useful way, whether full array maps or time histories. We have found that generally the printed numbers are used only as an adjunct to the graphs and charts and are referred to when a question arises from an examination of the pictures. Fig. 6 shows an example of an array showing velocity vectors and isocontours. Fig. 7 illustrates both a time history and a comparison of alternate environments.

COM (Computer Output on Microfilm) devices seem most suitable for the very complex charts we generate. Pen plotters (flatbed or drum) work well for the simpler time histories; however, COM equipment is usually faster and shows better definition in detailed drawings. The microfilm can also be used directly for slides and possibly for motion pictures.

Animated films of the simulation are both interesting and very expensive. For example, charts showing velocity vectors and isolines of water level or the concentration of some substance may be made at frequent intervals and connected (usually with each frame repeated using an optical printer) to form a sequence showing movement in the system. This technique is mainly useful as a dramatic means of presentation, since a few charts showing activity at especially interesting times will usually provide adequate information and cost much less.

CONCLUSIONS

Computer modeling is very useful, although time consuming and expensive. The alternatives, such as enormous physical models or inadvertent damage to the environment, are more expensive.

Larger and faster computers are needed for simulation programs. The use of networks to gain access to remote "super-computers" seems inevitable for large models. The present requirements may be approached by such machines as the ILLIAC IV; however, the requirements for computer power will always expand to be more than is available.

Data Management of the great volume of output is a major consideration in post-processing as is selection of data to be displayed. The ordering of data and the file format (perhaps self-defining files) will affect the ease of retrieval and the time spent in post-processing, especially if the data are to be used frequently for various comparisons and analyses. The development of large archival storage (such as the UNICON laser memory system) and large random access storage devices will help solve the data management problems.

The simulation designer should reproduce the physical system being modeled, should anticipate computer limitations and--of great importance--should consider the interface that will exist with a human user. Extensive use of graphic display will assist in all phases of model development and use. The system of programs used in a simulation must be carefully designed from initial formulation through to the final presentation of results.

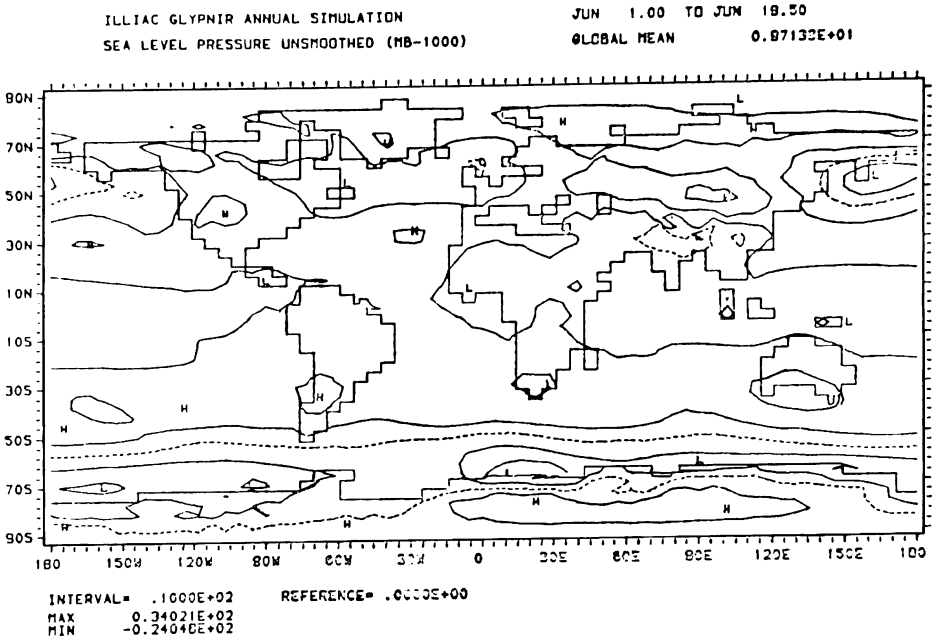

Figure 1 -- Sea level pressure map drawn using data computed by the Rand atmospheric circulation model on the ILLIAC IV. The map was returned to a Tektronix display device at Rand over the ARPANET.

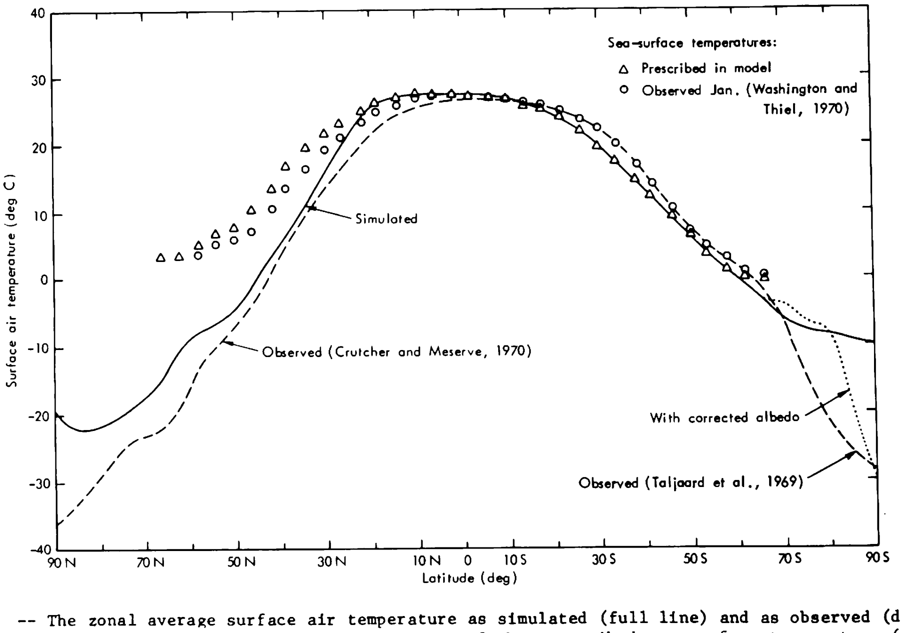

Fig. 2 -- **The zonal average surface air temperature as simulated (full line) and as observed (dashed line)**, Also shown are the zonal averages of the prescribed sea-surface temperature (as used in the simulation) and of the observed sea-surface temperature for January.

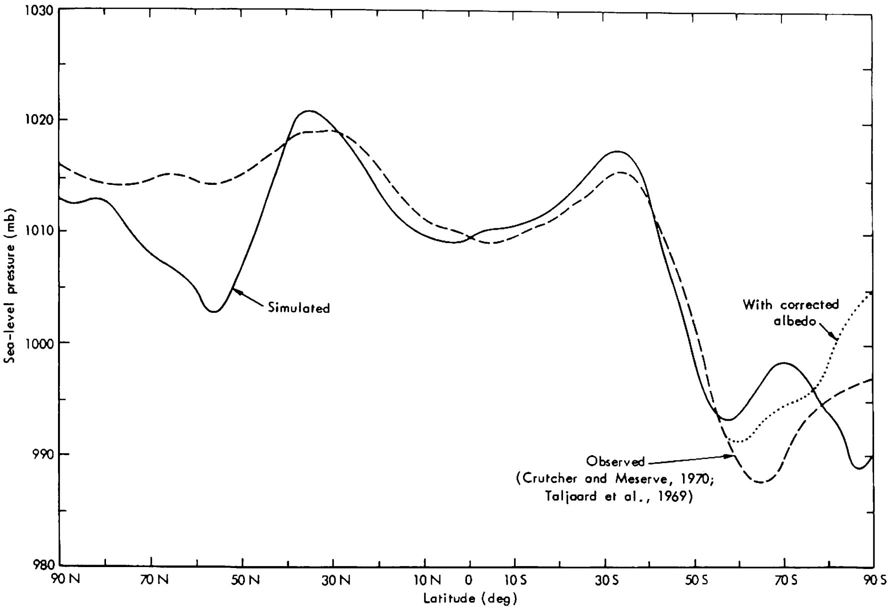

Fig. 3 -- The zonal average sea-level pressure as simulated (full line) and as observed (dashed line). The dotted curve at high sourthern latitudes is the simulation with a corrected albedo for snow and ice, as obtained from a subsequent January integration.

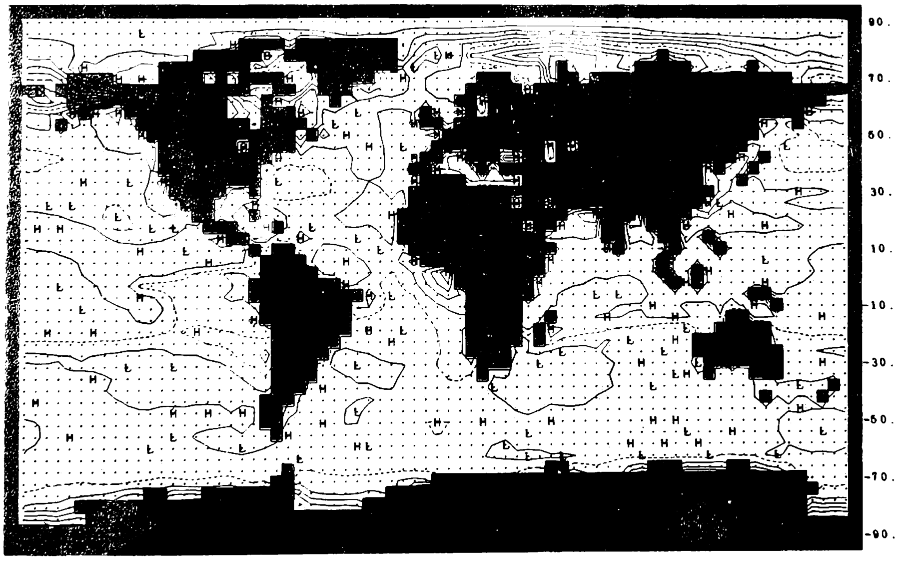

Fig. 4 -- The difference between the simulated and observed surface air temperature. The isoline interval is 3 deg C with the zero line dashed.

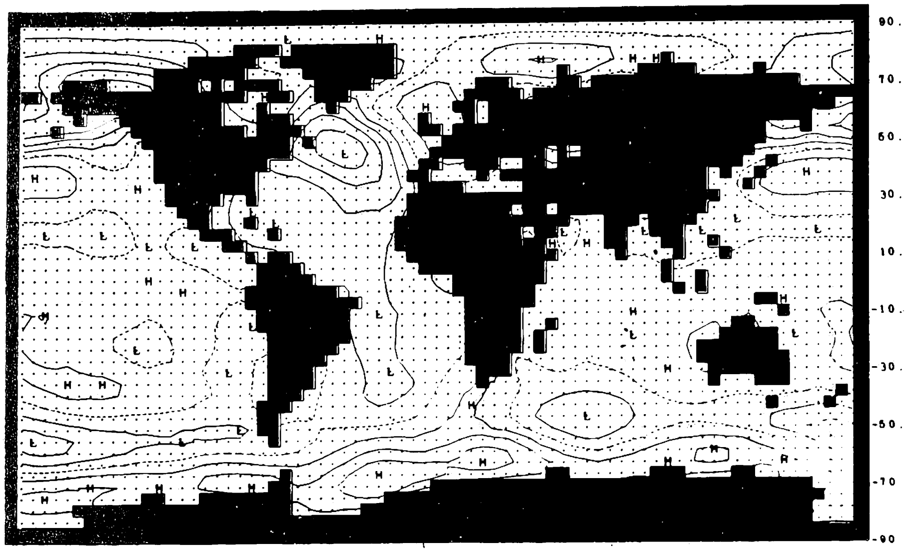

Fig. 5 -- The difference between the simulated and observed sea-level pressure. The isobar interval is 5 mb with the zero line dashed.

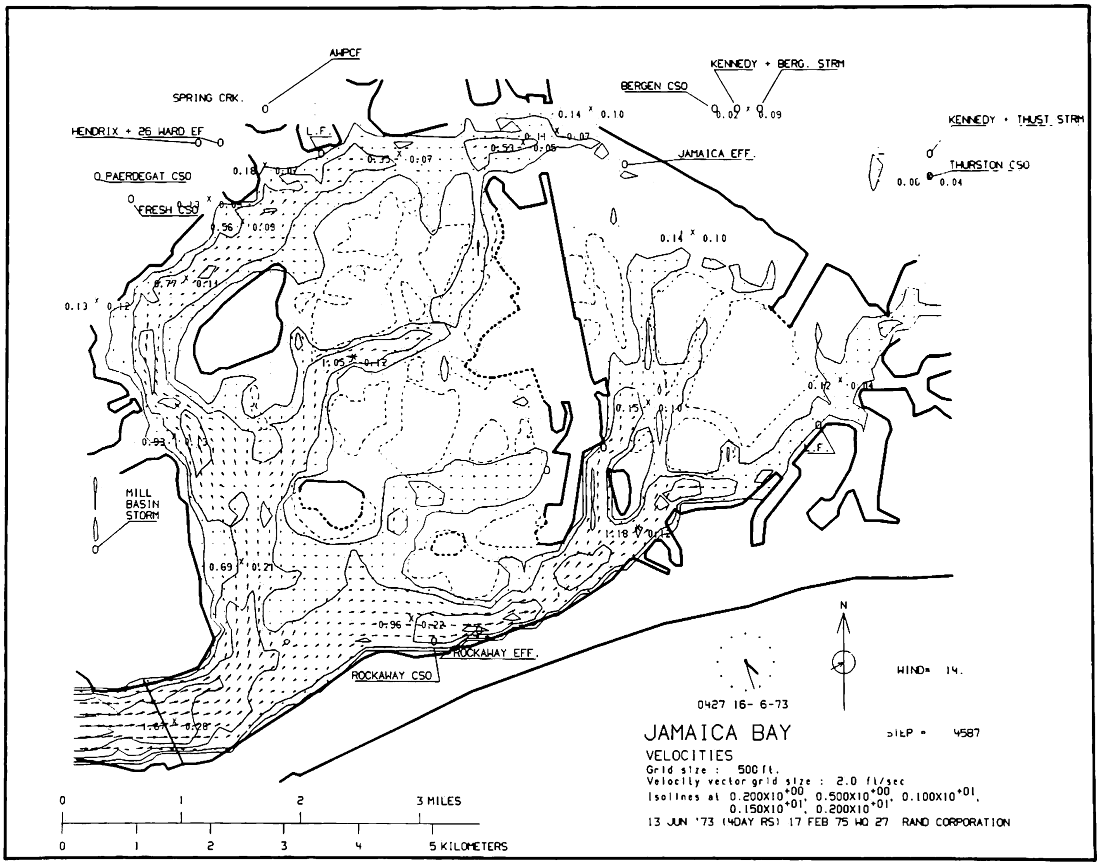

Fig. 6 -- Velocity vectors and isolines computed for Jamaica Bay, New York. Flooding and drying of tidal flats is indicated by the presence and absence of dots.

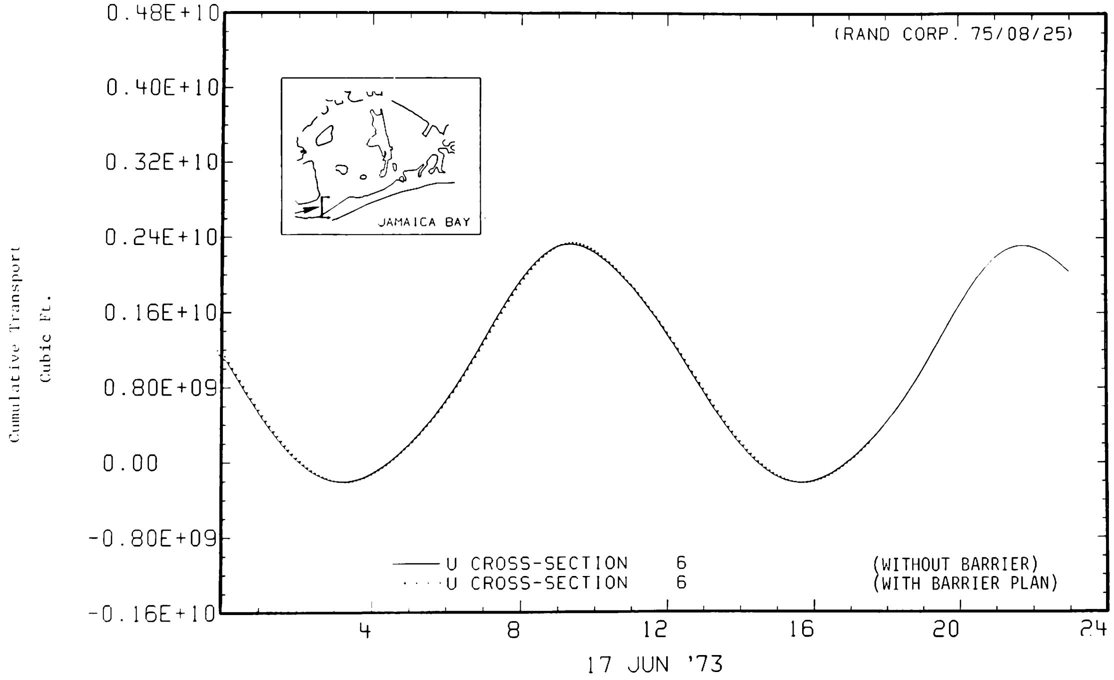

Fig. 7 -- Comparison of east-west transport of water through a given cross-section of Jamaica Bay, New York, with and without construction of hurricane barrier