The Water Environment and Human Needs

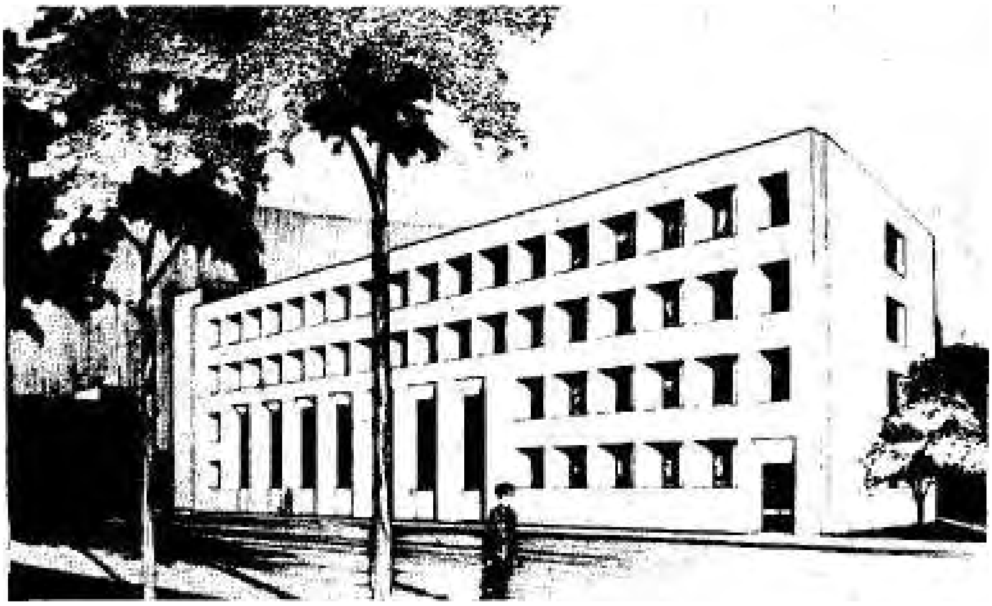

Ralph M. Parsons Laboratory for
Water Resources and
Hydrodynamics

Department of Civil Engineering

Massachusetts Institute of
Technology

Symposium October
1 and 2,1970

The Water Environment and Human Needs

A symposium held to Mark the Dedication
of the
Ralph M. Parsons Laboratory for
Water Resources and Hydrodynamics

October 1 and 2, 1970

Edited by Arthur T. Ippen
Institute Professor

Published by the Ralph M. Parsons Laboratory
for Water Resources and Hydrodynamics
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

A HIERARCHY OF MODELS FOR OPTIMIZING THE OPERATION OF WATER SYSTEMS

by

URI SHAMIR*

ABSTRACT

We consider the problem of optimizing the operation of large, complex water systems, comprising surface and/or ground reservoirs, a main grid and local distribution systems connected to it. Policies are sought for the operation over a wide range of levels in space and time: from the instantaneous operation of the local systems or the main grid for a few hours, to the yearly operation of the entire system for a period of one or possibly several decades.

It is proposed that the *way* in which this problem can be tackled is by a process of decomposition, in both space and time, which leads to the construction of a hierarchy of models, all embedded in a general framework, and related to each other via their constraints and objective functions.

Such a hierarchy of models is being developed in the Research Division of Mekoroth, Israel's water company, for optimizing the operation of Israel's water system. The structure of the individual models, and the way in which they are linked, is presented. Some of the achievements to date, and the difficulties which arose are discussed.

---

*Senior Lecturer, Faculty of Civil Engineering, Technion, Israel, and Consultant, Research Division, Mekoroth Water Company, Ltd., Israel

1. INTRODUCTION

The first, and probably most crucial, stage in analyzing a system is that of defining the system. In an optimization problem, once one has selected the decision variables and decided what the constraints and objective function are, obtaining the solution is merely a question of technique. The system being studied is usually but a small part of the "world", and one is then faced with the problem of defining the boundaries of the system under study, and of determining what the true conditions on these boundaries are. Moreover, one usually simplifies the description of the real system, with the aim of making its mathematical model manageable.

This paper deals with the problem of operating a complex water system. We are seeking operating policies for time periods ranging from a few hours and up to several years. One can obviously not provide operating rules for the smallest components of the system, such as pumps and valves, for more than a few hours. On the other hand, in providing long range operating rules for reservoirs, the instantaneous hydraulics of the system cannot be included explicitly in the model. The traditional approach has been to deal with one specific time horizon, making some assumptions about the performance of the system for much longer and much shorter periods of time. The results of such models depend strongly on these assumptions. The approach we have taken is to deal with all time horizons, with each in a separate model, but in such a way that the models interact and provide one another with some results which are much better than mere assumptions. The result is a "family" of models, all embedded in a unified structure. Because each model deals with a different time horizon, we view the models as being arranged in a hierarchy. The longer the time horizon of a model, the higher up it is in this hierarchy.

2. THE HIERARCHY

The hierarchy of models treated here is being developed for optimizing the operation of Israel's water system. We shall describe in some detail this specific hierarchy, the models in it and the links between them. The concept is general, however, and should be applicable to other systems and conditions.

Israel's water system contains one main surface reservoir - Lake Kinnereth, several aquifers located throughout the country, and several operational surface reservoirs. There is a main line from Lake Kinnereth in the north to the south of the country, covering a total distance of approximately 200 km. It is made of an open channel, a 108" pipe, branching into two parallel pipes and then joining again in the south. Some 25 local projects are connected to the main grid, each serving one or several communities. Each such project draws its water either from the main grid or from local ground or surface sources, and at times may operate as a supplier to the main grid.

The hierarchy contains the following models: a D (decades)-Model of the main system, a Y (yearly)-Model of the main system, and I (instantaneous)-Models, one for the main grid (IM-Model), and one for each local project (IL-Models). The Y-Model and the IM-Model are in an advanced stage of development. One IL-Model, for a large local system, has been completed and fully tested. The D-Model has been formulated, work on it is in progress, but there is still much detailing to do.

The characteristics of each model will be discussed in accordance with the following points, following their order and numbering:

1. The part of the system it deals with, for example, the entire system, the main distribution grid only, or some local system.
2. The time horizon and the time steps. For example, a year divided into months, or a day divided into hours.
3. The decision variables. For example, reservoir levels, flows, quantities to be pumped, etc.
4. The constraints relevant to the model itself, such as maximum and/or minimum levels in reservoirs, demands to be supplied, hydraulic constraints, etc.
5. Constraints provided by the model below - the model which deals with the shorter time horizon and with a more detailed description of the system or part of it. These constraints,which cannot be included explicitly in the model under consideration, are questions of the form "Can the results of the present model be actually carried out?" The

answers to these questions have to be given by the model below because the model itself, the one under consideration, cannot be detailed enough to answer them, or it becomes too complex and unwieldy. The model below will determine whether the policy is feasible or not. If it is, the model below will also return the benefit or cost of carrying the policy out.

6. The directives and long range economic functions provided by the model above, the model dealing with the longer time horizon.
7. The objective function. It now contains the economic functions of the activities dictated by the long range consideration, as well as the immediate benefits and costs of carrying out any policy.
8. The frequency at which a new solution of the model is to be computed, and the urgency with which the solution is needed. A yearly model may be run once a year, or, if unforseen events take place, a new solution may be needed before a whole year has gone by since the last solution. In either case, several hours of computer time required for a single solution are not unreasonable. Moreover, a delay of several days or even weeks in obtaining the solution may be tolerated. In contrast, the hourly operation has to be determined at least a few times a week, and a delay of a few hours in obtaining the solution may make it obsolete. Thus, the efficiency of this model is of prime importance.
9. The data. One has to consider the types of data, their sources, methods of acquisition, the volume of fixed and variable data, the frequency at which the latter type is changing, the admissible time lag between acquisition and use, the required accuracy, etc. All these, and some other special considerations concerning available computer hardware and software, determine the data flow and data structure for each model.
10. The optimization algorithm. Selection of the algorithm is based on all the above characteristics of each model, and on the specific form and number of equations describing it.
11. The parameters appearing in the model, their ranges, and the sensitivity of the optimal solution to variations in their values. We use

the word "parameters" to denote scalars, entire vectors or single variables, whose values are either unknown with sufficient accuracy, or unknown at the time the solution is being computed. In the latter group there appear stochastic variables, such as rainfall and consumptions. There is also another type of unknown variables. For example, while developing one model the adjacent models, the ones above and below it, may still be under development themselves (if they have been tackled at all). Thus, some constraints and certain components of the objective function of the model under consideration are still unknown. This leads one to estimating the values of the unknown variables, using judgement and the best information available. These estimates, together with expected values of the stochastic variables, are used in a "basic run" of the model. Then one performs a sensitivity analysis with respect to these variables, covering their probable range of variation.

The parametric runs performed during the sensitivity analysis have a great value in themselves, as they enable the operator to examine the operation of the system under a variety of conditions.

3. THE DECADES MODEL (D-MODEL)

Work on the D-Model has covered only the design phase. Actual implementation is still only at its beginning. Therefore, what follows is not considered to be the final form of the model, as it will probably undergo changes during the detailed implementation.

3.1 The D-Model deals with the entire system in its present form, as well as with its development and expansion. The system is described in a schematic form only. For example, a regional water project, or even several such projects, may appear as a single consumer or contributor of water.

3.2 Time horizon: several years, up to a few decades. Time steps: at least one year, possibly several years. At present it is thought that the Y-Model will be the building block for the D-Model.

3.3 The decision variables are divided in two:

3.3.1 Variables which determine the development and expansion of the system, for example, exploitation of new sources (conventional or man-made), augmentation of ground storage, or additions to the distribution system.

3.3.2 Variables which determine the operation of the system, for example, amounts to be drawn from each source or put into storage into each ground reservoir in any given year. The answers from the D-Model are given in policy rules for operation under all possible conditions.

3.4 The constraints relevant to this model are:

3.4.1 The hydrology limits the amounts of water which are available for use during any period of time. The hydrology is stochastic, and has to be taken as such in the model, either explicitly or by simulation and sensitivity analysis. In semi-arid regions the frequency and severity of draughts is probably one of the most important characteristics of the hydrology, although the entire sequence of stochastic yearly inputs into all water sources has to be considered,

3.4.2 Economic constraints, determined by the availability of capital, of labor, etc., or given as directives by those at the helm of the national economy.

3.5 The model below is the Y-Model, to be described in detail later. Given the state of the system in any year, the Y-Model can determine whether any policy dictated by the D-Model can indeed be carried out. If so, the Y-Model will also yield the cost (or benefit) of carrying it out, in the optimal manner.

3.6 The D-Model being at the top of the hierarchy, there is no model above it. Actually, the long range economic constraints mentioned in 3.4.2 can appear either as constraints in the D-Model, as explained above, or appear in its objective function.

3.7 There are two possible ways to formulate the D-Model, which differ primarily in their objective functions. These are:

3.7.1 To consider the water system as a production system in itself, designed to "produce" water and sell it. In this case the objective is to maximize the net benefit of the owner of the water system. This is the easier of the two approaches, as the water system becomes a separate "industry" in the national economy, although tied to other industries, such as agriculture (which in Israel is the consumer of some 85% of the water) by the prices it can get for the water sold. This does not represent accurately the conditions in Israel, where national authorities control the water supply company on the one hand, and agriculture (via regulations and subsidies) on the other. Thus we are led to the other alternative:

3.7.2 To consider water as one resource in the national economy, and to take as objective of the D-Model the optimization of some long range national economic goal.

For either of the two alternate ways of defining the objective function, we will have to construct curves of supply and demand versus the cost and the price of water, respectively.

3.8 The D-Model will be run probably once every few years, but possibly as frequently as several times a year, when examination of new conditions is called for.

3.9 The data base for the D-Model will be developed from the data structures of the more detailed models, selecting only the necessary data for the schematic definition of the system.

In addition, hydrologic and meteorologic data will be obtained, incorporated into the data structure and updated once a year or once every few months.

3.10 Considering the stochastic nature of the hydrology, meteorology and of the demand, we anticipate that no mathematical optimization algorithm will be suitable for solving the D-Model. Therefore, simulation will probably be used.

3.11 Unknown parameters of the D-Model will be the stochastic yearly hydrology, demand patterns, technological developments (such as progress in desalination) and long range economic trends. The Y-Model will have been completed and tested when the D-Model is detailed. Therefore, the feasibility of carrying out each year's operation, and the associated cost or benefit will be readily available.

It is commonly accepted that Israel is already using over 90% of its conventional water potential, although there is some difference of opinions concerning the definition of "potential" (sometimes termed "safe yield"), and in fixing its numerical value. Thus, the development of new, unconventional sources of supply has to be considered. This makes it difficult to formulate the D-Model in a definite and final form. Rather, it will have to be used in various forms, to test possible developments or proposed technologies.

4. THE YEARLY MODEL (Y-MODEL)

4.1 The Y-Model deals with the main system. The configuration is fixed in advance for each year, although it may change from one year to the next, as modifications and additions are implemented. The description of the system is schematic, same as that of the D-Model or somewhat more detailed.

The main system contains one over-year surface reservoir (Lake Kinnereth), some 15 ground reservoirs (each an entire aquifer or a well defined cell of an aquifer), five operational surface reservoirs used for hourly and daily regulation, and some 25 local projects, in each of which we define a demand, transfers from or to the main grid, pumping from a ground reservoir and recharge into it.

4.2 Time horizon: one year. Time steps: months.

Actually, the building block of the Y-Model is a monthly model. This latter model, not defined separately in the hierarchy, deals with the operation during one month. Twelve such blocks are linked to form the entire Y-Model.

4.3 The decision variables are quantities of water to be transferred during each month from point to point in the system. These include the amounts to be taken from the surface and ground reservoirs, amounts to be recharged into the ground, amounts to be transferred from or to all local projects, and the monthly quantities passing through each line in the main grid.

4.4 The constraints are divided into several categories:

4.4.1 Hydrology determines the yearly maximal amounts to be drawn from each source.

4.4.2 Hydrologic, engineering and institutional considerations determine maximum and minimum admissible levels in each surface and ground reservoir.

4.4.3 Demands have to be met. Demands are predicted for each month of the year for all consumers, and are assumed to be deterministic.

4.4.4 At certain points in the system a given salinity of the water is not to be exceeded. The sources have different salinities, some above the admissible limit, and waters have to be mixed to meet the salinity constraints.

4.4.5 Monthly quantities flowing through the main grid have to be hydraulically feasible.

4.5 The models below are the IM-Model and the IL-Models. These should determine whether a policy selected by the Y-Model for a given month can indeed be implemented. Both the salinity (4.4.4) and hydraulic (4.4.5) constraints are non-linear. Including these constraints explicitly in the Y-Model makes an otherwise linear model into a non-linear one. Therefore salinity constraints are dealt with through successive linearizations. Simplified hydraulic constraints are incorporated into the Y-Model itself, but the detailed check will have to be performed in the IM-Model.

4.6 The D-Model will eventually provide the Y-Model with the long range economic functions. In the D-Model's absence these have to be examined parametrically (see sections 4.7 and 4.11 below).

4.7 The objective of the Y-Model is to minimize yearly operating costs. This would lead to a tendency to deplete the ground storage which is the cheapest way to meet this year's demand. The reason is that storing water this year for use in future years incurrs costs and results in no immediate benefits. To force the optimal solution to consider the long range benefits of storage one can do one of two things; either

4.7.1 Introduce constraints giving the amounts to be put into storage, or

4.7.2 Introduce a "forcing function" into the objective function, which makes it beneficial to store.

The latter method has been tried, by adding a large penalty for not pumping water out of Lake Kinnereth. The objective function is still the minimization of yearly expenditures, now including this penalty as well. The optimal solution thus contains pumping from the Lake which exceeds the demand. The excess is recharged into ground storage, divided among the aquifers according to a given priority list. The first method is being tested now, to see whether it is easier to work with. Its advantage is that it gives the shadow prices, a most useful tool in evaluating the effect of these constraints.

4.8 The Y-Model will probably be run a few times a year. Whenever the actual hydrologic or meterorologic events differ significantly from what was assumed at the time the last solution was computed, a new solution will have to be obtained, taking into account the new conditions.

4.9 The data include a definition of the distribution system, characteristics of all the reservoirs and the hydraulic constraints. Much of the data are permanent, some change infrequently and some are periodic. All data can be acquired by simple means and no on-line acquisitopm facilities are contemplated for use with this model alone.

4.10 The Y-Model was formulated as a linear programming problem. Some linearization had to be performed to include hydraulic constraints, and salinities are taken care of by successive solutions with improved approximate salinity functions. A standard computer program is being used. The computation of the solution is rather efficient and comprehensive. Parametric runs are easily performed, and shadow prices can be used to interpret the results.

4.11 The objective function and some of the constraints have to be examined parametrically. Using the penalty method (4.7.2), the effect of changing the relative weight of the penalty and of the operating costs was determined by parametric runs. The same can be done for the other method (4.7.1) of fixing the amounts to be recharged into the aquifers. Using the penalty function method, one encounters a difficulty. The penalty, which is imposed for not pumping from Lake Kinnereth, is a measure of the economic loss which will be incurred as a result. The water not pumped from the Lake, and therefore not put into ground storage, will be needed at some future date. Its absence will either require that a new source be developed, or else that some loss will result. It is not possible to determine the cost of developing new sources in the future with any accuracy, nor is it possible to compute the loss due to reduction in agricultural production as a result of water shortage. Therefore, the penalty has to be guessed at and examined parametrically.

5. THE INSTANTANEOUS MODEL OF THE MAIN SYSTEM (IM-MODEL)

5.1 The IM-Model deals with the main system. All physical components are fixed for the duration of the model's time horizon.

5.2 The time horizon is several hours, and the time step is one hour.

5.3 The decision variables are those components of the system which can be controlled: starting and shutting down pumps, reservoir levels, valve settings, etc. These result in discharges throughout the system and in hourly water transfers.

5.4 The constraints are:

5.4.1 Hourly demands to direct consumers, as predicted by a separate model, are to be met.

5.4.2 The hydraulics of the system determine the flows and pressures. Some constraints on admissible flows and pressures have also to be satisfied.

5.4.3 Salinity cannot exceed a given maximum at certain points in the system.

5.4.4 Monthly transfers of water to local projects, as dictated by the Y-Model have to be carried out. The distribution within the month is to be determined by the IM-Model itself.

5.5 There is no model below. However, the IL-Models should be used to determine whether hourly transfers from and to the main grid are hydraulically feasible in the local projects. Alternatively, it is possible to include in the IM-Model itself constraints representing the hydraulics of the connections to the local projects. Both approaches can be used, selecting for each connection to a local project the more suitable one.

5.6 The Y-Model yields the monthly quantities to be put in or taken out of storage, mainly in the ground reservoirs. These become constraints on the IM-Model. It is obvious that these transfers incur immediate costs, whereas benefits from them will accrue only in the distant future. Therefore forcing these transfers upon the IM-Model must be based on weighing their immediate costs against their long range benefits. It is not possible to do this within the IM-Model itself, because it would require running it for a period of one year, or even several years, a period long enough so that at its end the economic value of the resulting state of the reservoirs could be evaluated. Thus we introduce the long range benefits through the constraints forced upon the IM-Model by the Y-Model. This method is justified because the long range benefits of stored water are very much greater than the costs incurred in transferring it to storage. A complicating factor is the fact that some water is lost through underground flow out of the aquifer.

5.7 The objective function is the minimization of operating costs. Long range benefits are implicit in the constraints giving transfers to storage.

5.8 Ultimately this model could be used on-line, in real time, to control the operation of the main system, and be re-run every few hours. Alternatively, the model can be used first to study in detail the operation under normal conditions, and later be used to deal only with exceptional or emergency situations. In this mode it does not have to be used on-line. The decision on how to use it will depend on the benefits which can be derived from using it in either mode. Once completed, the model will yield the savings which can be realized by its use. Also, we will know how long it takes to run the model on the computer and how much each solution costs. With this information one should be able to reach a decision on the best mode of using the model.

5.9 On-Line operation of the model in real-time will require an extensive and very expensive data acquisition system. The data to be collected and relayed to a control center, where the computer is located, include reservoir levels, indications of pumps in operation, pressures, discharges to consumers, etc. A data base, giving the full description of the system, has to be resident in the computer, or in secondary storage, to be used together with the variable data in the optimization model.

If, on the other hand, the model is used off-line, to study the normal operation of the system and for dealing with exceptional conditions, then the data acquisition and dispatch problems are greatly reduced. It is expected that the IM-Model will have a major impact on the process of evaluating the desirability of installing a data acquisition and dispatch system, used to relay data to a control center which is in charge of operating the system.

5.10 The large number of decision variables and of constraints, the complexity of the system, the number of feasible solutions, the large number of times the model is to be run and the urgency with which the solution is needed — all put severe limitations on the selection of an optimization algorithm.

We have broken the network into four parts. For each, we use a dynamic programming formulation, together with some heuristics aimed at reducing the amount of computation. The individual solutions are then combined, to obtain the solution for the entire system.

5.11 The parameters in the IM-Model are first the transfer amounts dictated by the Y-Model. Next, the relative cost of various operations has to be considered parametrically, as they may change with time. This will be included in a parametric examination of the entire objective function.

6. THE INSTANTANEOUS MODEL OF A LOCAL SYSTEM (IL-MODEL)

There are some 25 local systems connected to the main grid. These vary in size, complexity, types of consumers, local storage, etc. We hope to deal with some, though probably not all, local projects, and have a library of programs — one for each such project. The investment, in both time and money, needed for developing a model for each project will have to be weighed against the benefits to be derived from it. Each model is useful in itself, as a tool for the operators to study the system in depth, and learn how to Improve its operation. It should also be a part of the hierarchy, and therefore be made compatible with the other models.

Only one local model has been completed. It deals with a relatively large system, called the Kfar-Baruch project. The consumption is primarily agricultural. There is an 8 million cubic meter surface reservoir in which water from both the main grid and local sources is stored and then pumped for use. The connection with the main grid, as defined in the IM-Model and the Y-Model is via the water transfers. This project always appears as a consumer, as it never supplies water to the main grid.

6.1 This IL-Model deals with the Kfar-Baruch project. Its description contains the main reservoir, 9 operational surface reservoirs used for daily regulation and 10 booster stations.

6.2 Time horizon: one day. Time steps: variable - given by the time to the next change in operating conditions.

6.3 Pumps are controlled in this project by set points in the reservoirs, and possibly by clock settings. In each booster station there are defined groups of pumps, called pumping configurations, or simply configurations. Each configuration has a given capacity-head curve. The configuration is switched on when the water level in the controlling reservoir drops to a given set-point. It is stopped when the level rises and reaches another point. An over-ride of these points can be provided by clock settings. The settings of all these points for all the configurations are the decision variables of the IL-Model.

6.4 Constraints are:

6.4.1 All demands have to be met. The demands are given, and assumed to be deterministic.

6.4.2 Flows and heads in the network, under any operating conditions, have to satisfy the hydraulic laws.

6.4.3 Maximum and minimum water levels in all reservoirs are not to be exceeded.

6.4.4 Certain limitations on the locations of set points have to be satisfied, such as: minimum distance between adjacent points, the highest starting point has to be below the lowest stopping point, etc.

6.5 The IL-Model is the most detailed possible. There is no model below it and therefore there are no constraints imposed "from below".

6.6 The IM-Model uses the various IL-Models to deal with the instantaneous operation. The compatibility between them is through the water transfers from the main grid to the local system or back. Thus, the IL-Models will receive from the IM-Model the amounts to be transferred during each period of time. No long range economic functions will be provided as directives to the IL-Models. Rather, they will be implicitly present in the constraints provided "from above".

6.7 The objective of each IL-Model is to minimize operating costs. At present we measure these by the energy expended during a period of 24 hours. Each pumping configuration takes a certain amount of energy to operate, and the sum over 24 hours of the energy spent is the measure of the efficiency of the present values of the set-points.

6.8 Each IL-Model can be used in one of two possible ways. It can either be used to study in depth the operation of the system, with the aim of obtaining operating rules for normal times, and kept on hand for dealing with emergency conditions. Or else it can be used routinely, run once every day or every few days, to take care of changing conditions. At present it is thought that the IL-Models will be used in the former mode. This means they will not be run on-line, and therefore one can tolerate long running times (if necessary), and a relatively long response time (from initiation of a run to its completion).

6.9 The data contain a full description of the hydraulic system, hourly demands of all consumers (as predicted by a prediction model), pumping configuration data (capacity - head curve, energy consumption, etc.), and the set-points and their limitations.

6.10 The complexity of the model precludes the application of any of the better known mathematical optimization techniques. We have resorted to a search technique, based on experience and some heuristics, which is essentially a strategy for changing one or a few decision variables at a time, and moving from one feasible point to a better one. There is no guarantee that the global optimum is indeed obtained by this strategy, but experience shows that the solution thus obtained is satisfactory.

6.11 Sensitivity of the model was examined with respect to the following:

6.11.1 Demand pattern. Changes in peak magnitude, duration and timing during the day.

6.11.2 Storage volume in reservoirs.

6.11.3 Number of pumps and their sizes, in various booster stations.

6.11.4 Initial and required final conditions of the system, as indicated by reservoir levels.

7. CONCLUSION

Development of the hierarchy as a concept and as a working tool is carried on at the same time the individual models in it are being worked on. The general structure and some of the operating rules of the hierarchy were set in advance, to make sure each model is structured in accordance with them. But things keep changing in the progress of the work. This is a normal and sound procedure. It does cause difficulties, however, as each model has to be adjusted to "meet" the neighboring models on the boundaries. Furthermore, each model needs some results of the other models. These are often not available at the moment they are needed, and one has to assume certain things, proceed with the development, and provide enough flexibility to incorporate new results as they become available.

A further difficulty arises from the fact that we are dealing with concrete, real systems. The models cannot be abstract, but rather have to incorporate all practical "real world" considerations of operating the systems. Such things as the lag between occurence of events and acquisition of the data describing them, the operators' level of competence to execute complicated policy rules, etc., have all to be considered.

Some immediate benefits are being derived from the work. The individual models are being used by the operators of the system to study it in detail under various conditions. The model is used to test proposed operating rules and their physical and economic consequences.

We are also learning from this experience how to attack complex problems in water resources. The progress to date encourages us in thinking that the approach we have taken is a promising one. The idea is to decompose the system in a physical, rather than mathematical sense, to deal with each sub-problem separately after taking due consideration of its particular structure, and to build the individual models so they fit together into a general structure, in our case a hierarchy.

ACKNOWLEDGEMENTS

The work presented here is the result of a team effort by the staff and consultants of the Research Division of Mekoroth Water Company, Ltd., Israel, headed by Dr. Nathan Arad.

I have been intimately involved in the formulation of the general framework and of some of the models. On other models I merely keep in the picture by periodic briefings. Whatever merit the work deserves is due to the entire team. For any shortcomings of the presentation I alone am responsible.