Automatica, Vol. 5, pp. 15–26. Pergamon Press, 1969. Printed in Great Britain.

Applications of Dynamic Programming to the Control of Water Resource Systems*

R. E. LARSON† and W. G. KECKLER‡

Several dynamic programming computational procedures may aid in operating water resource systems: daily as in a pumped-storage facility and a four-reservoir system, and annually with stochastic inflows and in long term planning of system additions.

Summary—The complexity and expense of water system projects have made optimum operation and design by computer-based techniques of increasing interest in recent years. Dynamic programming offers a powerful approach to a wide variety of these problems.

Most water system problems can be classed as one of the following three types:

(1) Optimum operation during a short period, such as 24 hours, when all quantities are known;
(2) Monthly or yearly policy optimization when some system parameters, such as stream inflows, are not known exactly;
(3) Long-range planning or resource allocation when demands may or may not be known exactly.

Realistic water resource problems have many decision and state variable constraints. There are also nonlinearities or stochastic variations in both the state equations and the return function. This paper describes how dynamic programming can handle these difficulties.

Several specialized dynamic programming techniques applicable to water system problems are also introduced. These include successive approximations, forward dynamic programming, dynamic programming for stochastic control, and iteration in policy space.

Four examples are solved and discussed—short-term optimization of a two-reservoir system is solved with forward dynamic programming; short-term optimization of a four-reservoir system is treated by successive approximations; optimum operation over a year, when stream-flows are stochastic variables, is found by iteration in policy spaces; and optimum long-term planning of system additions given projected demand is treated by forward dynamic programming.

I. INTRODUCTION

As WATER resource systems have grown larger and more complex, the importance of optimum operation and planning of these systems has increased. The investment costs and operating expenses of projects are so large that even small improvements in system utilization can involve substantial amounts of money. Also, the various control points—power generators, irrigation outlets, pumping stations, etc.—interact in such a complicated manner that it is difficult to obtain an optimum design or operating policy using an intuitive approach. Thus, the potential benefits of using optimization techniques in these problems are very great indeed.

Dynamic programming provides an extremely powerful and general approach for solving these optimization problems. Nonlinearities in the system equations and performance criterion can easily be handled. Constraints on both decision and state variables introduce no difficulties. Stochastic effects can be explicitly taken into account.

In section II the basic equations of dynamic programming are briefly reviewed to introduce the terminology to be used as well as a number of computational procedures: the standard computational algorithm, successive approximations, forward dynamic programming, the standard computational algorithm for stochastic control problems, and iteration in policy space.

In section III some specific problems to which dynamic programming has been applied are discussed. Most water problems fall into one of the following three categories:

(1) Optimum operation during a short period, such as 24 hr, when all quantities are deterministic;
(2) Monthly or yearly policy optimization when some system parameters, such as stream inflows, must be treated as stochastic variables.
(3) Long-range planning or resource allocation where demands may or may not be treated as deterministic quantities.

Four illustrative examples are discussed, including at least one from each of the above categories. The first problem is the optimum short-term operation of a combined pumped hydro and irrigation storage facility involving two reservoirs; forward dynamic programming was used for this example. The second problem is the optimum short-term operation of a multipurpose four-reservoir system,

* Submitted to IFAC Haifa Symposium on Automatic Control of Natural Resources, September 1967. Manuscript received 15 April 1968 and in revised form 6 August 1968. Recommended for possible publication by associate editor P. Dorato.

† Wolf Management Services, Palo Alto, California.
‡ Stanford Research Institute, Menlo Park, California.

where power generation, irrigation, flood control, and recreation are all considered; the technique of successive approximations was applied in this case. The third problem is the optimum management of a single reservoir over a 1-year period, where stochastic variations of input stream-flows are considered; iteration in policy space was applied here. The fourth problem is the optimum planning of additions to a system over a 30-year period; forward dynamic programming was again used for this example.

II. BASIC CONCEPTS IN DYNAMIC PROGRAMMING

Most of the problems for which dynamic programming has been used to obtain numerical solutions can be formulated as deterministic discrete-time variational control problems [1–3]. The general case of this problem is formulated as follows

Given:

(i) A system described by the nonlinear difference equation

$$\mathbf{x}(k+1)=\mathbf{\Phi}[\mathbf{x}(k),\,\mathbf{u}(k),\,k], \tag{1}$$

where

$\mathbf{x}$ = state vector, n-dimensional
$\mathbf{u}$ = control vector, m-dimensional
k = index for stage variable
$\mathbf{\Phi}$ = n-dimensional vector functional;

(ii) A variational performance criterion

$$J=\sum_{k=0}^{K} L[\mathbf{x}(k),\,\mathbf{u}(k),\,k], \tag{2}$$

where

J = total cost
L = cost for a single stage;

(iii) Constraints

$$\mathbf{x}\in X(k) \tag{3}$$

$$\mathbf{u}\in U(\mathbf{x},\,k) \tag{4}$$

where

$X(k)$ = set of admissible states at stage k
$U(\mathbf{x}, k)$ = set of admissible controls at state $\mathbf{x}$, stage k;

(iv) An initial state

$$\mathbf{x}(0)=\mathbf{c}. \tag{5}$$

Find:

The control sequence $\mathbf{u}(0), \ldots, \mathbf{u}(K)$ such that J in equation (2) is minimized subject to the system equation (1), the constraint equations (3) and (4), and the initial condition (5).

The dynamic programming solution to the above problem is obtained by using an iterative functional equation that determines the optimal control for any admissible state at any stage. The minimum-cost function is defined for all $\mathbf{x}\in X$ and all $k, = 0, 1, \ldots K$, as

$$I(\mathbf{x},\,k)=\min_{j=k,\,k+1,\,\ldots,\,K}\left\{\sum_{j=k}^{K} L[\mathbf{x}(j),\,\mathbf{u}(j),\,j]\right\}, \tag{6}$$

where

$$\mathbf{x}(k)=\mathbf{x}.$$

Abbreviating $\mathbf{u}(k)$ as $\mathbf{u}$, the iterative functional equation becomes

$$I(\mathbf{x},\,k)=\min_{\mathbf{u}}\{L(\mathbf{x},\,\mathbf{u},\,k)+I[\mathbf{\Phi}(\mathbf{x},\,\mathbf{u},\,k),\,k+1]\}. \tag{7}$$

This equation is a mathematical statement of Bellman's principle of optimality [1–3]. The optimal control at state $\mathbf{x}$ and stage K, denoted as $\hat{\mathbf{u}}(\mathbf{x}, k)$, is directly obtained as the value of $\mathbf{u}$ for which the minimum in equation (7) is attained.

Since equation (7) determines $I(\mathbf{x}, k)$ and $\hat{\mathbf{u}}(\mathbf{x}, k)$ in terms of $I(\mathbf{x}, k+1)$, it must be solved backward in k. As a terminal boundary condition

$$I(\mathbf{x},\,K)=\min_{\mathbf{u}}[L(\mathbf{x},\,\mathbf{u},\,K)]. \tag{8}$$

The optimization over a sequence of controls is thus reduced to a sequence of optimizations over a single control vector.

An iterative equation analogous to equation (7) can be derived for continuous-time problems and for problems containing stochastic variables [1–4].

The standard computational procedure for solving equation (7) is to quantize admissible values of $\mathbf{x}$ and $\mathbf{u}$ to a finite number of discrete values and then to perform the minimization at any quantized value of $\mathbf{x}$ by a direct search over quantized values of $\mathbf{u}$. This procedure has a number of desirable properties: an absolute optimum is always determined; a feedback control policy is obtained; and considerable flexibility in handling constraints, nonlinearities, and stochastic effects is provided [1–4]. However, the procedure does have the drawback that computational requirements can become excessive in high-dimensional problems [1–4].

A number of procedures are available that reduce these computational requirements while retaining the desirable properties of the standard algorithm. These include successive approximations [1–4, 13, 14], forward dynamic programming [4], iteration in policy space [2–5], quasilinearization [6], iteration about a nominal using successively finer quantization increments [7], and state increment dynamic programming [4]. A comprehensive survey of these procedures appears in Ref. [8].

III. EXAMPLES

A. *Short-term optimization of a pumped-storage two-reservoir system*

1. *Problem statement.* In Ref. [9] a pump-storage system is described. The basis for the problem is the San Luis Reservoir and its forebay, a joint facility of the State of California and Bureau of Reclamation in the State Water Project. The solution of this problem utilized forward dynamic programming. The network configuration of the problem to be solved is shown in Fig. 1. The water

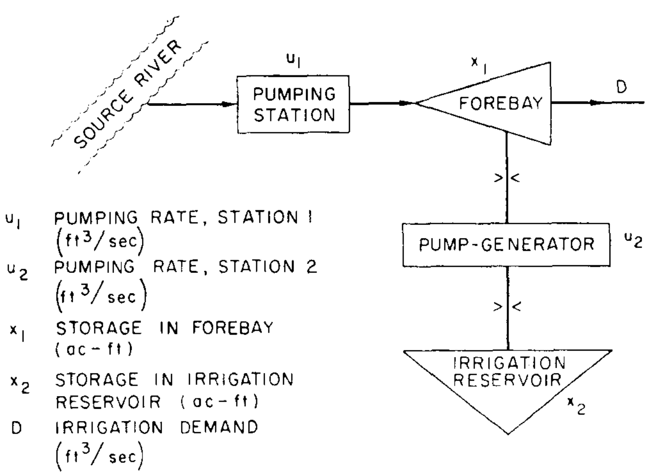

FIG. 1. Network configuration for two-reservoir example.

from the source river is pumped in the forebay x_1, from which it is either pumped into the storage reservoir x_2 or used to meet an irrigation demand. The pumps of the large reservoir can also function as generators of electrical power when water flows back to the forebay. The rate at which water can be removed from the source river has an upper limit, and the pumping plants have capacity limitations. Dollar values can be put on all costs and revenues. The problem is to operate within all constraints and to meet all demands on the system at minimum cost.

The quantized state variables are x_1 and x_2. The control variables u_1 and u_2 are allowed to vary continuously within certain upper and lower limits. Control u_1 varies from zero to an upper limit and control u_2 varies from some negative lower limit to some positive upper limit. A negative u_2 indicates that Pumping Station 2 is being used to generate electrical power. The irrigation demand D is limited to positive values and has the same units as the controls, a flow rate. Because of the dimensional differences in the u's and x's, a conversion factor is needed:

$$C = 12{\cdot}3 \frac{\text{acre-ft}}{\text{ft}^3/\text{sec}}$$

The operating procedure is to be computed one day in advance and is reconsidered every hour; therefore, time is quantized into increments of 1 hr.

The water balance equations* or the state equations are the following:

$$\begin{aligned} x_1[(k+1)\Delta t] &= x_1(k\Delta t) + C[u_1(k\Delta t)\Delta t \\ &\quad - u_2(k\Delta t)\Delta t - D(k\Delta t)\Delta t] \\ x_2[(k+1)\Delta t] &= x_2(k\Delta t) + C[u_2(k\Delta t)\Delta t] \end{aligned} \tag{9}$$

where in this problem

$$\Delta t = 1 \text{ hr}.$$

Thus, the equations become

$$\begin{aligned} x_1(k+1) &= x_1(k) + C[u_1(k) - u_2(k) - D(k)] \\ x_2(k+1) &= x_2(k) + Cu_2(k). \end{aligned} \tag{10}$$

However, operating the pumping stations for an hour incurs certain costs. The only pumping station operating cost considered in this problem is the cost of electrical power. This cost (K) is expressed as the cost of pumping at the rate of 1 ft^3/sec for 1 hr. The efficiencies of both pumping plants are the same, so the per-unit operating cost of each one is K if u_2 is positive. The efficiency of Station 2 changes when it is used as a generator; thus, there must be a different cost (benefit) K' when u_2 is negative. It is assumed that the electrical power that Station 2 generates can be sold at the same price that power can be purchased and that the power cost varies during a day. Thus, the cost of producing additional power is greater than the cost of producing the base level of power. Intuitively, it appears that in order to meet the irrigation demand, and minimize cost, there are times during the day when it is most profitable to release water. The solution to the problem verifies this supposition and determines when each policy should be followed.

The cost accrued during the kth time increment is

$$\begin{aligned} L(k) &= Ku_1(k) + K_2u_2(k) \\ K_2 &= \begin{cases} K, & u_2 \geq 0 \\ K', & u_2 < 0. \end{cases} \end{aligned} \tag{11}$$

* The time increment of 1 hr is long enough so that transient effects such as channel dynamics do not need to be included in this formulation. Therefore, although the resulting water balance equations may seem oversimplified they are actually very realistic models from which practical operating information can be obtained and optimum controls derived.

The total cost from the initial time to time k is thus

$$I(\mathbf{x}, k) = I(\mathbf{x}, k-1) + L(k)$$
$$I(\mathbf{x}, 0) = L(0) = 0 \text{ for all } \mathbf{x}. \quad (12)$$

The quantity $I(\mathbf{x}, N)$ is the cost of operating the system from initial to final time ($t = t_f = N\Delta t$ or $k = N$) and terminating in state $\mathbf{x}$. The problem becomes one of choosing the controls $\mathbf{u}(k)$ of equation (10) for all values of k such that all constraints are satisfied and $I(\mathbf{x}, N)$ is minimized for all $\mathbf{x}$.

2. *A typical problem.* A FORTRAN program using forward dynamic programming has been implemented for the two-reservoir, two-pump station facility. The control is not quantized, but allowed to vary continuously between certain limits. The computed trajectories can therefore be forced to go from one quantized state to a quantized state at the next stage of the process. Thus, no interpolation is required and one has continuous, piecewise-linear trajectories in the state space.

Figure 2 shows the demand curve of irrigation water and the incremental power cost curve. The incremental power cost is the cost of the last

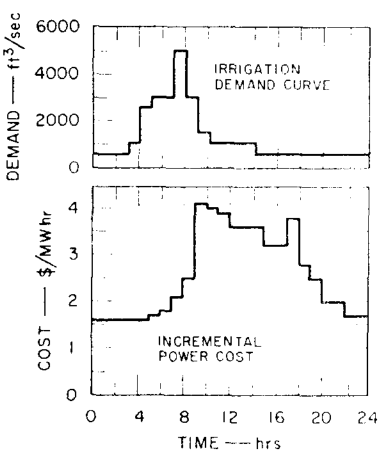

FIG. 2. Input quantities for the numerical example: irrigation demand and incremental cost as a function of time.

megawatt-hour produced during each hour. Since the whole system analyzed here operates as an additional load or source to the electric power grid, it will either have to buy power at the incremental power cost or replace power which costs this much. This curve was derived from information given in Ref. [10]. The irrigation demand curve was assumed to be shown, but dynamic programming could include many other formulations of this demand. The initial value of the reservoir levels are the $k = 0$ values shown in Fig. 3.

The optimum cost for each terminal state varies considerably, and the one which is the overall optimum depends on the penalty assessed for

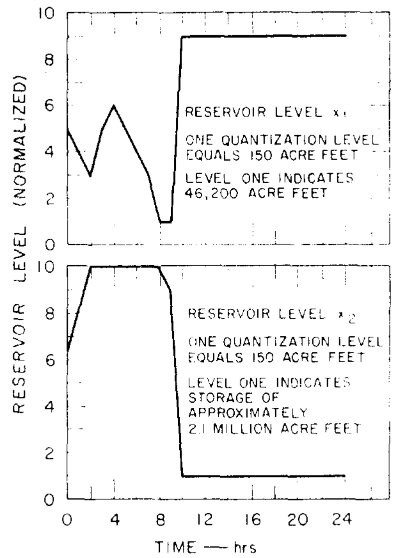

FIG. 3. State variables as a function of time.

arriving at each of these final states. If there is no penalty assessed for arriving at different terminal states, but a terminal constraint is imposed that the total amount of water in the two reservoirs must be 10 units, then the optimum terminal state is $x_1 = 9$, $x_2 = 1$. The minimum cost for this state is $I(\mathbf{x}, N)$ \$611.93. The optimal policy corresponding to this state is shown in Figs. 3–5; the reservoir levels as function of time are shown in Fig. 3, the optimum controls are shown in Fig. 4, and the cumulative operating cost is shown in Fig. 5.

In this case, the best policy is to fill Reservoir 2 early in the day when power is least costly and to drain all that is possible during the period, the tenth hour, when the return is greatest. Reservoir

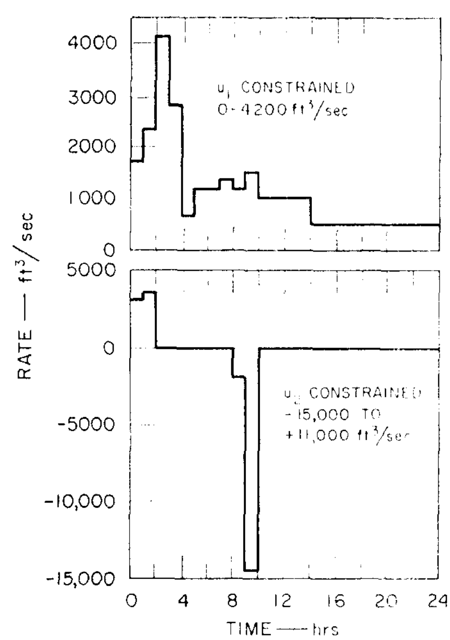

FIG. 4. Pumping rates of the numerical example as a function of time.

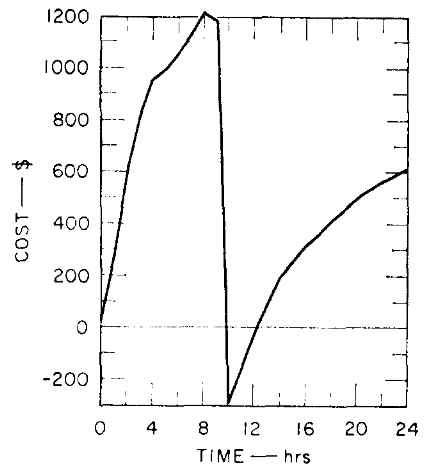

FIG. 5. Cumulative operating cost.

1 is first operated to ensure that it is at its lowest allowable level in the tenth hour and is thus able to receive the water released from Reservoir 2. In the tenth hour the reservoir rises to the ninth quantization level and remains there. No change in the level of the reservoir is possible because the irrigation demand requires less than one quantization level of water and the additional demand is met by u_1. The cumulative operating cost (Fig. 4) reflects the pumping policy shown in Fig. 5. It shows high cost as the pumps fill Reservoir 2 early in the day, but the return for this policy is high during the tenth hour.

3. *Extensions.* Dynamic programming is able to handle a wide variety of constraints that result from physical situations. Some of these which can be expressed in the context of this example are cited below. One constraint is a limitation on the amount of water that can be pumped from the source river during a 24-hr period. This is a very real problem in California; the Sacremento River Delta could be contaminated by salt water if the flow of the river were disturbed too much. As a result, the irrigation requirement often is also expressed as the amount to be delivered during a 24-hr period.

Many pump-generator stations have already been built and integrated into a power system. These stations provide a "spinning reserve" during certain hours of the day. This responsibility requires that u_2 be constrained to be less than some negative value during these hours. Other contractual requirements would be imposed on a realistic system. These include penalties for not exceeding minimum levels of irrigation or electric power demand. If too much water or electric power is produced, the return for the excess may be less than for the basic deliveries. Since the short-term control situation is embedded in a longer-term operation, the final values of the two reservoir levels are confined to certain regions of the state space. A penalty cost is assessed for not reaching the desired final state and bonus given if this value is exceeded.

B. *Short-term optimization of a multipurpose four-reservoir system*

In this section, the optimum operation over 24 hr of a multipurpose four-reservoir system is determined. The reservoir network, which contains both series and parallel connections, is shown in Fig. 6. In this optimization, use of water for power generation, irrigation, flood control and recreation is considered. Interaction of the short-term optimization with longer-term operating policies is also taken into account.

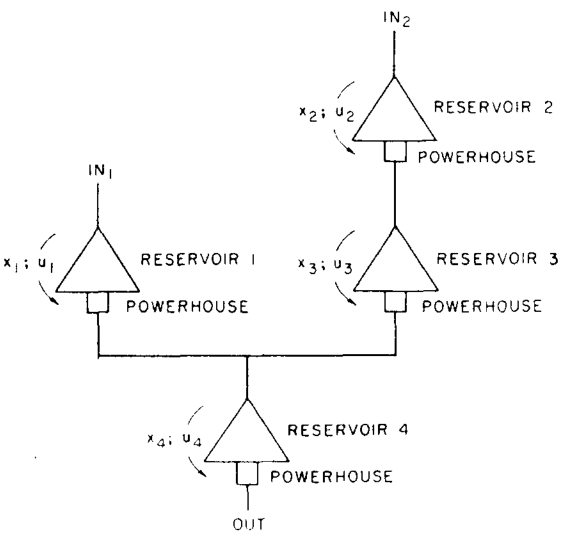

FIG. 6. Network configuration of four-reservoir problem.

The amount of water in the ith reservoir is denoted as x_i, $i=1, 2, 3, 4$, where each x_i is expressed in normalized units.

On the basis of potential use of the reservoir for recreation purposes, a minimum water level for each reservoir is specified: the amount of water needed to achieve this level is arbitrarily set as $x_i=0$, and a constraint is imposed that the amount of water in each reservoir cannot drop below this value.

On the basis of flood control considerations, a maximum water level for each reservoir is established. The amount of water needed to raise the level from the minimum to the maximum value is then expressed in terms of the normalized units, and a constraint is imposed that each x_i cannot exceed this level.

The particular constraints considered in this example are expressed as:

$$0 \leq x_1 \leq 10$$

$$0 \leq x_2 \leq 10$$

$$0 \leq x_3 \leq 10$$

$$0 \leq x_4 \leq 15. \qquad (13)$$

The flow of water between reservoirs is also expressed in the same normalized units; the control variables $u_i(k)$, $i=1, 2, \ldots 4$, specify the amount of water released from the ith reservoir over the kth time interval. In this example each time interval is 2 hr. For each reservoir a maximum flow is determined by the capacity of the power generators, and a minimum flow is determined by considering the use of the downstream river beds for navigation, conservation, and municipal and industrial water supplies. For this example the constraints were

$$\begin{aligned} &0 \le u_1 \le 3\\ &0 \le u_2 \le 4\\ &0 \le u_3 \le 4\\ &0 \le u_4 \le 7. \end{aligned} \tag{14}$$

The system equations express how the water flows between the reservoirs. They are:

$$\begin{aligned} &x_1(k+1)=x_1(k)-u_1(k)+IN_1\\ &x_2(k+1)=x_2(k)-u_2(k)+IN_2\\ &x_3(k+1)=x_3(k)-u_3(k)+u_2(k)\\ &x_4(k+1)=x_4(k)-u_4(k)+u_3(k)+u_1(k)\\ &\qquad k=0, 1, \ldots 11. \end{aligned} \tag{15}$$

The inflows IN_1 and IN_2 are assumed constant over the day as

$$\begin{aligned} &IN_1=2\\ &IN_2=3. \end{aligned} \tag{16}$$

The performance criterion considers the use of water for both power generation and irrigation. It is assumed that there is a power generation station at each reservoir outflow. The benefit from the flow over a given 2-hr period is assumed to be a linear function of the flow, i.e. the benefit from a flow out of reservoir at time k is $c_i(k)u_i(k)$. The function $c_i(k)$ is based on the power curve in Part A of this section. The values of $c_i(k)$ are shifted in k with respect to each other to account for the transport delay of water between reservoirs. This delay is 4 hr from Reservoir 1 to Reservoir 4, 4 hr from Reservoir 2 to Reservoir 3, and 2 hr from Reservoir 3 to Reservoir 4. The values of $c_i(k)$, $i=1, 2, 3, 4$ are shown in Table 1.

Irrigation benefits are considered only for the outflow from Reservoir 4. The benefit is again linear with flow—i.e. the benefit from flow $u_4(k)$ is $c_5(k)u_4(k)$. The function $c_5(k)$ is shown in Table 1.

TABLE 1. CONSTANTS IN PERFORMANCE CRITERION

k	$c_1(k)$	$c_2(k)$	$c_3(k)$	$c_4(k)$	$c_5(k)$
0	1·1	1·4	1·0	1·0	1·6
1	1·0	1·1	1·0	1·2	1·7
2	1·0	1·0	1·2	1·8	1·8
3	1·2	1·0	1·8	2·5	1·9
4	1·8	1·2	2·5	2·2	2·0
5	2·5	1·8	2·2	2·0	2·0
6	2·2	2·5	2·0	1·8	2·0
7	2·0	2·2	1·8	2·2	1·9
8	1·8	2·0	2·2	1·8	1·8
9	2·2	1·8	1·8	1·4	1·7
10	1·8	2·2	1·4	1·1	1·6
11	1·4	1·8	1·1	1·0	1·5

The benefit function also includes a terminal cost for failing to reach a specified level for each reservoir at the end of the day. This function accounts for the long-term policy of filling or emptying the reservoir during a particular season. This function assesses a heavy penalty for having less than the specified amount of water at the end of the day, but gives no credit for having more than this amount. The particular function used was

$$\psi_i[x_i(12), m_i]=\begin{cases}-40[x_i(12)-m_i]^2, & x_i(12)\le m_i\\ 0, \text{ otherwise} & \end{cases} \tag{17}$$

where m_i = desired level of reservoir i at the end of the day ($k=12$).

This problem has been solved by successive approximations. The initial state was taken to be

$$\begin{aligned} &x_1(0)=5\\ &x_2(0)=5\\ &x_3(0)=5\\ &x_4(0)=5. \end{aligned} \tag{18}$$

The desired final state was

$$\begin{aligned} &m_1=5\\ &m_2=5\\ &m_3=5\\ &m_4=7. \end{aligned} \tag{19}$$

The system dynamic equations are as in equations (15) and (16). The constraints are expressed in equations (13) and (14). The performance criterion is

$$J=\sum_{k=0}^{11}\sum_{i=1}^{4} c_i(k)u_i(k)+\sum_{k=0}^{11} c_5(k)u_4(k) + \sum_{i=1}^{4}\psi_i[x_i(12), m_i] \tag{20}$$

where $c_i(k)$, $i=1, 2, \ldots 5$ is specified in Table 1, $\psi_i[x_i(12), m_i]$ is as shown in equation (17) and m_i, $i=1, 2, 3, 4$, are given in equation (19).

The initial policy chosen is shown in Table 2. Basically, this policy consists of setting outflow equal to inflow at every time period, so that the water level in each reservoir remains constant. The only exception to this policy occurs at the end of the day, when the terminal cost function is taken into account.

TABLE 2. INITIAL POLICY

k	$x_1(k)$	$x_2(k)$	$x_3(k)$	$x_4(k)$	$u_1(k)$	$u_2(k)$	$u_3(k)$	$u_4(k)$
0	5	5	5	5	2	3	3	5
1	5	5	5	5	2	3	3	5
2	5	5	5	5	2	3	3	5
3	5	5	5	5	2	3	3	5
4	5	5	5	5	2	3	3	5
5	5	5	5	5	2	3	3	5
6	5	5	5	5	2	3	3	5
7	5	5	5	5	2	3	3	5
8	5	5	5	5	2	3	3	5
9	5	5	5	5	2	3	3	5
10	5	5	5	5	2	3	3	5
11	5	5	5	5	2	3	3	3
12	5	5	5	7				

Total benefit = 362·5

The optimum policy is shown in Table 3. The improvement in benefit was from 362·5 units to 401·3 units. The amount of computer time required for convergence to the optimum policy was about 30 sec in the B5500.

TABLE 3. OPTIMUM POLICY

k	$x_1(k)$	$x_2(k)$	$x_3(k)$	$x_4(k)$	$u_1(k)$	$u_2(k)$	$u_3(k)$	$u_4(k)$
0	5	5	5	5	1	4	0	0
1	6	4	8	7	0	1	0	2
2	8	5	10	5	0	2	4	7
3	10	7	8	1	2	0	4	7
4	10	10	4	0	3	3	4	7
5	9	10	3	0	3	4	4	7
6	8	9	3	0	3	4	4	7
7	7	8	3	0	3	4	4	7
8	6	7	3	0	3	4	4	7
9	5	6	3	0	3	4	4	7
10	4	5	3	0	3	4	4	0
11	3	4	3	7	0	2	0	0
12	5	5	5	7				

Total benefit = 401·3

The extension of this approach to larger systems is clearly feasible. Time-varying constraints and more general types of performance criteria can easily be handled. Furthermore, the problem formulation can be modified to perform optimization over time periods other than 24 hr. Convergence to the true optimum can be proved in many cases [14]. At this time it appears that optimization of 20-reservoir systems is well within the capability of present-day computers.

C. *Optimization in the presence of stochastic inflows*

1. *Problem statement.* The following example [11] shows how dynamic programming can be applied to an annual scheduling problem with stochastic inputs. The problem posed can be solved by means of iteration in policy space to yield a series of optimum policies for the management of one reservoir.

The problem is expressed in terms of a transaction between two businessmen—one the manager of a reservoir and one the owner of a hydro-electric plant fed by this reservoir. A similar problem could be posed even if both facilities were operated by the same group. The manager of a water storage reservoir wishes to maximize the average return from his reservoir over many years. The reservoir has three sources of income.

(1) An annual payment from agricultural users of water which is released during the growing season—April through September.
(2) An annual return from recreational use which is a function of the reservoir level on 30 September—the end of the water year.
(3) A return for each acre-foot of water released during the winter months between 1 October and 31 March. This revenue comes from the owner of a hydro-electric power generator downstream.

During the winter, when much of the precipitation falls as snow and thus is not immediately available, this power facility is faced with a severe water shortage. Thus, the owner is willing to pay well for each acre-foot of water guaranteed to be delivered, less for each acre-foot delivered in excess of this guarantee, and invokes a penalty for each promised acre-foot that is not delivered.

Each October the manager is faced with the problem of deciding how much water should be promised to the hydroelectric company. He knows the distribution of annual inflows into the reservoir and that 65 per cent of this inflow will occur during the duration of the contract. He must meet the terms of the contract unless there is no water in the reservoir. He will deliver no more water than is specified in the contract unless the reservoir will otherwise overflow.

The operating policy during the summer months is already specified. The reservoir is operated during these months to yield the best immediate return. The only trade-off is between use of the reservoir for recreational purposes and sale of water for irrigation. There is no return for release of irrigation water above a certain specified amount.

In the spring the inflow for the rest of the water year is much better defined. The manager knows how much snow fell in the watershed of his reservoir and thus is able to specify the remaining inflow more closely, but not with certainty.

Under the assumptions outlined here the whole operating policy for a year is specified once the contract with the hydroelectric company is signed. The manager's problem thus becomes to decide how much water he should promise to deliver and how this amount should vary with the level of the reservoir in October. This example answers this question for the reservoir whose characteristics are described in the next section.

2. *The source of data.* The problem described in Part 1 is typical of many faced by water planners, and the procedure used below can be combined with successive approximations or other decomposition techniques [15] to analyse large systems. The Tables and graphs described below were obtained from the records of several government agencies concerned with water resources. A representative reservoir is assigned a storage capacity of 50,000 acre-ft, which is discretized into eleven values ranging from 0 to 50,000 acre-ft in 5000-acre-ft increments.

Figure 7 relates the annual benefit to the annual delivery of irrigation water. Negative value is given to small deliveries of water because the agricultural investment of the users is not utilized.

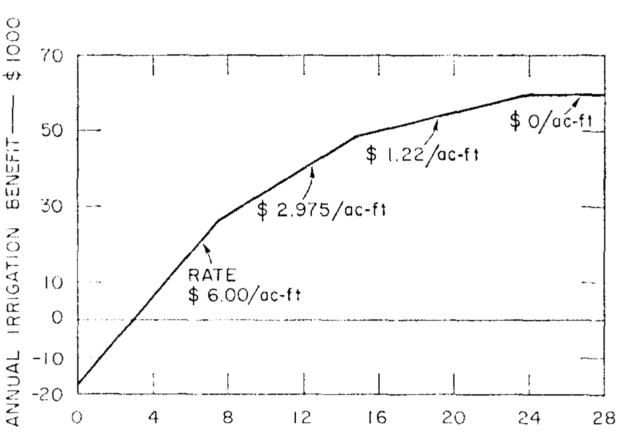

FIG. 7. Annual irrigation benefit.

The graph approximates a smooth curve by a series of linear functions to simplify digital computer use of this data. Figure 8 relates the annual recreational benefit to the end of the water year (30 September) storage in the reservoir. A negative value is given to zero storage because the investment in recreation facilities is not utilized. The values of this benefit at the quantized values of reservoir storage are used in computations. An analysis of the history of water year inflows yielded the discrete probability distribution function which is shown

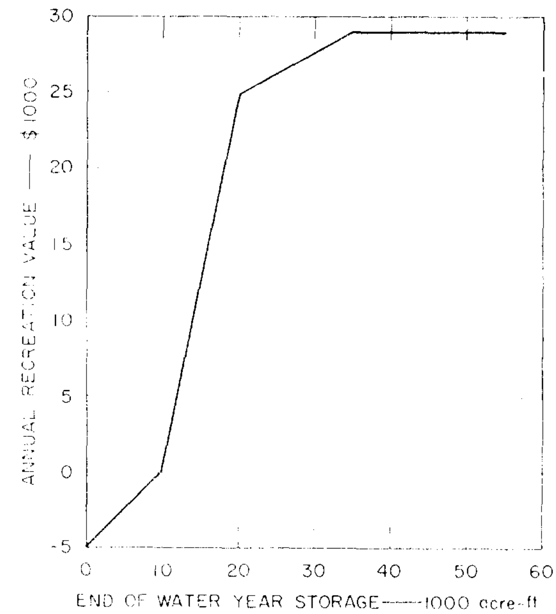

FIG. 8. Recreation benefit.

in Table 4. Figures 7 and 8 and Table 4 describe the state of knowledge of the manager of the reservoir in the autumn when the contract with the electric power company is negotiated.

TABLE 4. PROBABILITY DISTRIBUTION OF INFLOW

Probability of occurrence	Volume (acre-ft)
0·02	1715
0·08	3920
0·10	6550
0·10	9300
0·10	12,200
0·10	15,200
0·10	18,800
0·10	23,500
0·10	29,400
0·10	39,100
0·08	55,800
0·02	92,000

By comparing records of the predicted and actual inflows into a number of reservoirs during the summer months, it is possible to obtain a probability distribution function of the expected flow during the summer given the amount that flowed during the winter. Table 5 shows this result. To ease computations it is assumed that 65 per cent of the annual inflow (Table 4) occurs during the winter months. However, once this period is past the rest of the inflow is described by a "more packed" probability distribution function than the one available at the beginning of the water year. The decrease in uncertainty results from knowledge of the amount of water stored in the mountains in the form of snow and from the shorter prediction interval necessary to predict to the end of the water year.

TABLE 5. SUMMER PROBABILITY DISTRIBUTION OF INFLOW

Probability	Annual predicted rainfall which occurs in the summer (%)
0·10	20
0·20	30
0·40	35
0·20	40
0·10	50

3. *Problem formulation and indicated solution.* Application of the technique of iteration in policy space [5] requires that quantized state variables, transition probabilities, and rewards be defined. The level of the reservoir at the beginning of each water year is a continuous variable which can be quantized into a relatively small number of discrete values. Neither the transition probabilities nor the rewards for these transitions which are required for iteration in policy space are specified directly in Part C-2 of this section, but both can be computed from this data and the conditions specified in Part C-1. The remainder of the problem becomes a straight-forward application of iteration in policy space. The optimal schedule of release levels that the manager should promise as a function of reservoir level corresponds to the optimal policy obtained from the iterative procedure.

4. *Results.* The technique described in section III-C-3 above has been implemented in a FORTRAN computer program. This program has been machine translated to ALGOL and run on the Burroughs B5500 at Stanford Research Institute. It requires about 1 min to complete the computation of one case and usually converges to a solution in three or four policy iterations.

A number of cases have been run [11] using the data outlined in part C-2 of this section. The only quantities that were allowed to vary were the outflows that could be promised the hydroelectric company and the charges associated with this contract ($S1$, $S2$, and $S3$). These charges are chosen so that the return per acre-foot is comparable to that obtained from the scale of irrigation water. The charge under contract ($S1$) is chosen in the range of the slopes shown in Fig. 7 so there is a conflict between various policies. Two cases and the results are shown in Table 6.

In these two cases, the decision options and the contractual penalties remain the same, but the contract payment ($S1$) changes. When the contract price drops from \$10 to \$6 per acre-foot, the expected annual income declines from \$110,284 to \$72,839. The policy changes indicate that the reservoir management should not risk losing irrigation and recreation revenue at the lower contract price. The relative values show the long-term value of being in a given state $x^{(i)}$ at the present time compared to the value of being in a state $x^{(1)}$ at the present time. Notice that the percentage loss in value of a full reservoir is even greater than the percentage decline in contract price.

TABLE 6. EXAMPLE FOR ITERATION IN POLICY SPACE

Alternative policies (acre-ft)	
0	15,000
2000	20,000
5000	27,000
8000	35,000
11,000	45,000

Results

	Case 1		Case 2	
	$S1 = 10$\$/acre-ft, $S2 = 3$\$/acre-ft, $S3 = 15$\$/acre-ft		$S1 = 6$\$/acre-ft, $S2 = 3$\$/acre-ft, $S3 = 15$\$/acre-ft	
State	Optimal policy (acre-ft)	Relative value (\$)	Optimal policy (acre-ft)	Relative value (\$)
1	8000	0	5000	0
2	15,000	51,108	11,000	29,287
3	20,000	101,108	15,000	60,000
4	27,000	148,312	20,000	90,000
5	27,000	202,875	27,000	118,443
6	35,000	251,108	27,000	148,787
7	35,000	297,460	35,000	180,000
8	45,000	351,108	35,000	206,744
9	45,000	297,460	45,000	240,000
10	45,000	428,952	45,000	266,744
11	45,000	461,914	45,000	290,872
Gain	\$110,284		\$72,839	

$S1$—payment for each acre-ft delivered under contract
$S2$—payment for each acre-ft delivered in excess of the contract
$S3$—penalty for each acre-ft contracted for but not delivered

D. *Optimum planning of system additions*

The benefit of a large natural resource project is dependent on the timing of its construction. If it is completed too early, years may pass before its benefits can be fully utilized. However, it if is completed too late, there will be a long period when system users are denied its benefits or forced to pay higher costs than necessary. Thus, the decision of when to commit capital to a large project becomes critical. Unused investment is waste, and so is underdevelopment.

Dynamic programming is one way to optimally schedule when additional investment should be made when a long range solution is desired. KORSAK of SRI [13] has worked out the example given below of planning expansion of a power facility

30 years into the future. The problem may be stated as follows.

A power system has a current hydro capacity of 200 MW and a current demand of 500 MW. The hydro-generated energy remains constant and no cost is associated with this type of generation. The power demand is assumed to grow at a rate of 7 per cent per year. To simplify the model, it is assumed that the power demand makes discrete jumps of 7 per cent at the beginning of a given year and remains constant until the beginning of the next year.

The difference between the power demand and the hydro capacity can be made up each year in one of two ways. One may use plants that were purchased in preceding years and buy the remaining power at the rate of 12 mills per kWh or one may buy and install either a 250 MW or a 500 MW nuclear plant (but not one of each) at the beginning of the year. The operating cost of a nuclear plant is 3 mills per kWh. A 250-MW plant costs $\$3.45 \times 10^7$ dollars and a 500-MW plant costs $\$5.60 \times 10^7$ dollars. Since buying a plant might not meet the power requirement in a given year, one might still need to buy some power.

An interest rate of 12 per cent per annum on initial capital is used in determining the costs of any power plants.

The cost incurred during the kth year can thus be expressed as follows:

$$C(k)=\frac{V(k)}{1{\cdot}12^{k-1}}+\frac{1{\cdot}27\times 10^7}{1{\cdot}12^k}x(k)$$
$$+5{\cdot}08\times 10^7[1{\cdot}07^k-h(k)-x(k)] \qquad (21)$$

where

$C(k)$ = cost incurred during the kth year adjusted to beginning of first year with interest of 12 per cent.

$V(k)$ = cost of purchasing a plant in the kth year

$h(k)$ = hydro capacity in the kth year

$x(k)$ = total number of 500-MW units installed at time k (a 250-MW unit is considered half a 500-MW unit).

If no power is purchased,

$$C(k)=\frac{V(k)}{1{\cdot}12^{k-1}}+\frac{1{\cdot}27\times 10^7}{1{\cdot}12^k}x(k). \qquad (22)$$

The optimization is to minimize

$$\sum_{i=0}^{29} C(k)$$

subject to specified terminal constraints. The state variable $x(k)$ satisfies the equation

$$x(k)=x(k-1)+u(k)$$

where $u(k)$ is the control variable representing the decision to add 0, 250 MW, or 500 MW.

This problem was solved by forward dynamic programming. The optimum policy for reaching any feasible final state is determined. If the final state is selected as the one for which least cost is incurred, the optimal policy is as shown in Fig. 9. In this case the system capacity at the end is less than the demand; however, it can also be seen that in early years the capacity exceeds demand. These results reflect the assumption that no consideration is made of system operation after the final year. If the final state is constrained so that system capacity at the end exceeds demand, the optimal policy is exactly the same as in Fig. 9, except that during the last year before termination a 500-MW unit is purchased.

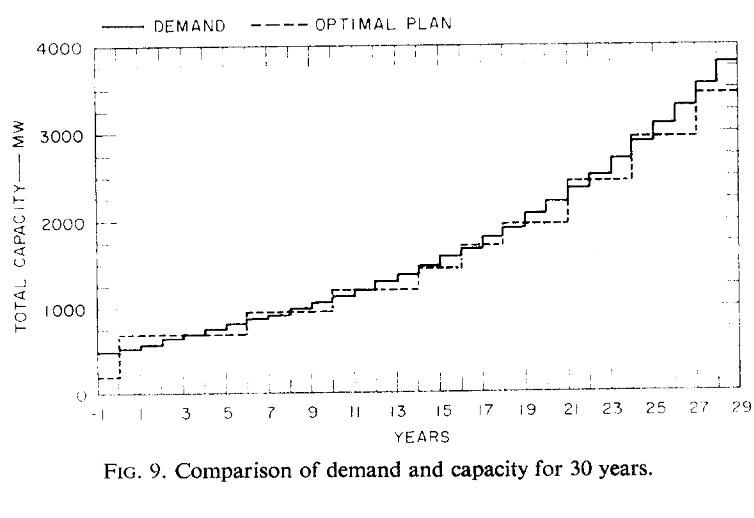

FIG. 9. Comparison of demand and capacity for 30 years.

IV. CONCLUSIONS

In this paper it has been shown that dynamic programming provides a powerful approach to many of the optimization problems that occur in water resource systems. Extremely general system equations and performance criteria can be handled, multiple constraints of a wide variety present no difficulties, an absolute optimum solution is obtained, the results are in a feedback control form, and stochastic variations can be explicitly taken into account. The major difficulty in applying it to practical problems has been the computational requirements associated with the standard computational algorithm. Sophisticated computational procedures are available that retain these desirable properties while substantially reducing computational requirements from those of the standard algorithm. The successive approximations technique applied in section III-B provides a particularly promising approach to high-dimensional problems.

The four problems discussed at length in section III show the breadth of the water resource problems that can be solved. These range from hourly control of a system involving hydroelectric power, water storage, and irrigation to long-range optimum investment planning. The stochastic character of nature is considered in the example of section III-C. None of these examples are the most difficult of their type that can now be solved, but they do demonstrate the principles and power of dynamic programming. Much more complicated problems are being solved now, and further research in computer technology and dynamic programming techniques will allow an even greater range of applications.

REFERENCES

[1] R. Bellman: *Dynamic Programming.* Princeton University Press (1957).
[2] R. Bellman: *Adaptive Control Processes.* Princeton University Press (1961).
[3] R. Bellman and S. Dreyfus: *Applied Dynamic Programming.* Princeton University Press (1962).
[4] R. E. Larson: *State Increment Dynamic Programming.* American Elsevier, New York (1968).
American Elsevier, New York (1967).
[5] R. A. Howard: *Dynamic Programming and Markov Processes.* Wiley, New York (1960).
[6] R. Bellman and R. E. Kalaba: *Quasilinearization and Nonlinear Boundary-Value Problems.* American Elsevier, New York (1965).
[7] W. G. Keckler: Optimization about a single trajectory by means of dynamic programming. Presented at *SIAM 2nd Int. Conf. Computing Methods in Optimization Problems,* San Remo, Italy, September (1968).
[8] R. E. Larson: A Survey of Dynamic Programming Computational Procedures. *IEEE Trans. A.C.,* **AC-12,** No. 6, 767–774, December (1967).
[8] R. E. Larson: Computational aspects of dynamic programming. 1967 *IEEE Int. Con. Rec.* Part 3, March (1967).
[9] W. G. Keckler: Optimum management of a combined pumped hydro and irrigation storage facility. Technical Memorandum, Stanford Research Institute, Menlo Park, California, March (1966).
[10] P. J. Bernard, J. F. Dopazo and G. W. Stagg: A method for economic scheduling of a combined pumped hydro and steam generating system. *IEEE Trans. Power Apparat. Sys.* **PAS-83,** No. 1, 23–29, January (1964).
[11] W. G. Keckler: Reservoir management and policy iteration. Technical Memorandum, Stanford Research Institute, Menlo Park, California, June (1966).
[12] M. Cuénod and A. J. Korsak: Optimal planning of power systems. Technical Memorandum, Stanford Research Institute, Menlo Park, California, June (1966).
[13] R. E. Larson: A dynamic programming successive approximations technique. Presented at 1968 *JACC.* University of Michigan, June (1968).
[14] A. J. Korsak and R. E. Larson: Convergence proofs for a dynamic programming successive approximations technique. Submitted to *4th IFAC Congr.,* Warsaw, Poland, June (1969).
[15] W. A. Hall and R. W. Shephard: Optimum operations for planning of a complex water resource system. Contribution No. 122, University of California Water Resources Center, Los Angeles, California, October (1967).

Résumé—La complexité et le coût des projets de systèmes hydrauliques ont apporté un interêt accru, durant ces dernières années, à l'optimalisation du fonctionnement et de l'étude à l'aride de techniques basées sur les calculateurs. La programmation dynamique constitue un moyen puissant d'approcher une grande varieté de ces problèmes.

La plupart des problèmes de systèmes hydrauliques peut être classée dans une des catégories suivantes;

(1) Fonctionnement optimal pendant une courte période, telle que 24 heures, lorsque toutes les quantités sont connues;

(2) Optimalisation de la politique mensuelle ou annuelle, lorsque certains paramètres du système, tels que les débits affluents, ne sont pas connus avec précision;

(3) Planification ou allocation des ressources a long terme lorsque les demandes peuvent être connues avec précision ou non.

Les problèmes réalistes de ressources hydrauliques possèdent de nombreuses contraintes de décisions et de variables d'état. Il existe également des non-linéarités ou des variations aléatoires à la fors dans les équations d'état et dans la fonction de reaction l'article explique comment la programmation dynamique peut traiter ces difficultés.

Plusiers techniques spécialisées de programmation dynamique applicables aux problèmes hydrauliques sont également introduites. Celles-ci comprennent les approximations successives, la programmation dynamique directe, la programmation dynamique pour commande aléatoire et l'itération dans l'espace des politiques.

Quatre exemples sont résolus et discutés—l'optimalisation à court terme d'un système à 2 réservoirs est résolue à l'aide de la programmation dynamiques directe; l'optimalisation à court terme d'un système à 4 réservoirs est traitée par approximations successives; le fonctionnement optimal sur une année, lorsque les débits affluents sont des variables aléatoires, est trouvée par iteration dans les espaces de politiques; et la planification optimale à long terme des additions au système, étant donnée une demande planifrée, est traitée par la programmation dynamique directe.

Zusammenfassung—Wegen der Kompliziertheit und der Kosten von Bewässerungsprojekten brachte man in den letzten Jahren der optimalen Operation und dem Entwurf durch den Einsatz der Rechentechnik steigendes Interesse entgegen.

Die meisten Wasserversorgungsprobleme können in eine der drei folgenden Klassen eingeordnet werden:

(1) Optimale Operation während einer kurzen Periode, etwa von 24 Stunden, wenn alle Großen bekannt sind.

(2) Monatliche oder jährliche praktische Optimierung, wenn einige Systemparameter, wie der Zufluß, nicht genau bekannt sind.

(3) Langfristige Planung oder Hilfsmittelzuteilung, wenn die Forderungen bekannt oder unbekannt sein können.

Realistische Bewässerungsprojekte besitzen viele Beschränkungen in den Entscheidungs- und Zustandsvariablen. Vorhanden sind überdies Nichtlinearitäten oder stochastische Veränderungen und zwar sowohl in den Zustandsgleichungen als auch in der Rückkehrfunktion. Die Arbeit beschreibt, wie diese Schwierigkeiten durch die dynamische Programmierung behoben werden können.

Mehrere spezielle Verfahren der Anwendung der dynamischen Programmierung auf Bewässerungsprobleme werden vorgestellt. Sie schließen sukzessive Approximationen ein, dynamische Vorwärtsprogrammierung, dynamische Programmierung für stochastische Kontrolle und Iteration im Raum der Strategie.

Vier Beispiele werden gelöst und diskutiert.—Eine kurzfristige Optimierung eines Zwei-Speicher-Systems wurde gelöst, eine kurzfristige Optimierung eines Vier-Speicher-Systems mit sukzessiver Approximation behandelt. Optimale Operation über ein Jahr, wobei die Stromlinienflüsse stochastische Variable sind, wird durch Iteration in Strategie-Räumen gefunden. Optimale langfristige Planung von Systemerweiterungen gegebenen projektierten Bedarfs wird durch dynamische Vorwärtsprogrammierung behandelt.

Резюме—Сложность и затраты проектов водяных систем принесли увеличивающийся интерес, в течении последних лет, к оптимизации функционирования и рассчета с помощью техник основанных на вычслительных машинах. Динамическое програмирование представляет собой мощный способ подхода к большой разнообразности таких проблем.

Большинство проблем водяных систем может бытв отнесено к одной из следующих категорий:

(1) Оптимальное функционирование в течении оротког периода, как например 24 часа, когда все количества известны;

(2) Оптимизация ежемесячной или ежегодной политики, когда некоторые параметры системы, как например притоки, точно не известны;

(3) Долгосрочное планирование или назначение рессурсов, когда требования могут быть точно известны или нет.

Реалистические проблемы водяных рессурсов имеют многочисленные ограничения решений и переменных состояния. Существуют также нелинейности или случайные изменения как в уравнениях состояния так и в функции обратной связи. Статья объясняет как динамическое программирование может обращаться с этими затруднениями.

Введено также несколько специализированных техник динамического программирования применяемых к водяным проблемам. Эти техники включают последовательные приближения, прямое динамическое программирование, динамическое программирование для стохастического управления и повторение в пространстве политик.

Решены и обсуждены четыре примера- краткосрочная оптимизация системы с двумя резервуарами решена с помощью прямого динамического программирования; краткосрочная оптимизация системы с четырьмя резеруарами рассмотрена с помощью последовательных приближений; годичное оптимальное функцирование, когда притоки представляют собой случайные переменные, найдено с помощью повторения в пространствах политик; и долгосрочное оптимальное планирование прибавлений к системе, с заранее запланированными требованиями, рассмотрено с помощью прямого динамическог опрограммирования.