WATER RESOURCES RESEARCH, VOL. 26, NO. 3, PAGES 447–454, MARCH 1990

Sampling Stochastic Dynamic Programming Applied to Reservoir Operation

JERSON KELMAN

Centro de Pesquisas de Energia Elétrica, Electric Energy Research Center, Rio de Janeiro, Brazil

JERY R. STEDINGER

School of Civil and Environmental Engineering, Cornell University, Ithaca, New York

LISA A. COOPER, ERIC HSU, AND SUN-QUAN YUAN

Pacific Gas and Electric Company, San Francisco, California

Most models for reservoir operation optimization have employed either deterministic optimization or stochastic dynamic programming algorithms. This paper develops sampling stochastic dynamic programming (SSDP), a technique that captures the complex temporal and spatial structure of the streamflow process by using a large number of sample streamflow sequences. The best inflow forecast can be included as a hydrologic state variable to improve the reservoir operating policy. A case study using the hydroelectric system on the North Fork of the Feather River in California illustrates the SSDP approach and its performance.

INTRODUCTION

Planning the operation of a river basin with even a single major reservoir and several downstream run-of-river powerhouses can be a complex problem for the following reasons: (1) Future inflows are uncertain; (2) the optimal releases from the reservoirs depend not only on their own storage and inflows, but also on the local inflows to the downstream powerhouses; (3) spatial correlations among concurrent streamflows are often high, autocorrelations vary in magnitude, and neither the spatial nor the time correlations should be neglected; (4) streamflow forecasts are often available and should be considered by the operating policy; and (5) because of head effects and the daily and monthly variation in system loads and corresponding thermal operating costs, the value of energy produced in a hydroelectric powerhouse is not a linear function of the flow through the turbines.

Most models that deal with reservoir system operations planning assume that the streamflow forecast is error-free [e.g., *Ikura and Gross*, 1984; *Grygier and Stedinger*, 1985]. If this was the case, future flows would be known with certainty, and a deterministic optimization algorithm would be appropriate. Many mathematical models that employ such algorithms are available [*Yeh*, 1985]. Because one seldom has a perfect streamflow forecast, deterministic models are often used in an adaptive mode: whenever there is an update of the forecast (often once a month), the model is rerun, producing a new energy production schedule for the next few months [*Neto et al.*, 1985]. Only the decision pertinent to the immediate month is actually implemented. This approach has been called naive feedback control.

An alternative to the naive feedback-control approach is to use stochastic dynamic programming (SDP) [*Loucks et al.*, 1981; *Yakowitz*, 1982; *Bras et al.*, 1983; *Stedinger et al.*, 1984]. This ambitious approach generates an operation policy or release decision for every possible reservoir storage state in each month, rather than just a single schedule of reservoir releases. (Alternatively, the derived future value function can be used to calculate optimal release targets as they are required.) Unfortunately, the representation of the system must often be simplified to make the algorithm computationally feasible [e.g., *Saad and Turgeon*, 1988], though recent papers have proposed alternative interpolation schemes that may reduce the computational effort [*Foufoula-Georgiou and Kitanidis*, 1988; *Johnson et al.*, 1988].

This paper develops a variation of SDP, sampling stochastic dynamic programming (SSDP), that generates an operation policy taking into account all of these issues. Unlike a sophisticated SDP, the complex structure of the streamflow process is not explicitly modeled in our SSDP. Rather, the features of the process are implicitly captured with a large number of 12-month streamflow sequences, observed or stochastically generated, that are possible realizations of the annual streamflow process. These 12-month streamflow series are called "streamflow scenarios." Unlike the approach of *Young* [1967], SSDP derives optimal decisions considering all of the streamflow scenarios simultaneously, instead of processing a sequence of decisions that are only optimal if one has perfect foresight. The SSDP algorithm is introduced in the next section, with the reservoir storage and a streamflow forecast as state variables, as suggested by *Stedinger et al.* [1984]. For a unique streamflow scenario (perfect foresight), the SSDP model reduces to deterministic dynamic programming. The no-forecast case, in which all streamflow scenarios are equally likely, has been investigated by *Araujo and Terry* [1974] and *Dias et al.* [1985].

The third section discusses the assignment of conditional probabilities to streamflow scenarios, given a streamflow forecast. A historical time series of streamflow forecasts is employed to develop the required conditional distributions. If this time series corresponds to the actual historical forecasts, the methodology captures the forecaster's experience,

Paper number 89WR02754.
0043-1397/90/89WR-02754$05.00

which usually goes beyond statistical manipulation of available quantitative data. However, if historical forecasts are not available, or the forecasting methodology has changed, it is possible to use a simple regression equation to backcast the streamflow forecast time series [*Grygier et al.*, 1989].

The fourth section presents an SSDP case study using a hydroelectric system operated by Pacific Gas and Electric Company (PG&E) on the North Fork of the Feather River in California. This system has one major reservoir with storage capacity of 1143 × 10^3 acre-feet (1406 × 10^6 m^3), located upstream of six cascaded powerhouses. For the case of a unique streamflow scenario, the SSDP solution is compared to the solution of a deterministic optimization model. For the case of multiple streamflow scenarios but no forecast, the SSDP solution is compared to the outcome of the simulation model actually used for planning. For the case of multiple streamflow scenarios with forecasts, the SSDP solution is used to estimate the benefit of forecasts in system operation. *Yeh et al.* [1982] provide an analysis of the worth of inflow forecasts for the Oroville-Thermalito reservoir system located in the lower portion of the Feather River.

SAMPLING STOCHASTIC DYNAMIC PROGRAMMING

Reservoir operators and planners are interested in the best use of stored water but are faced with uncertain future streamflows. In particular, they need to have a strategy for how much water to release over a planning period. The system of interest here is a river basin with one major reservoir operated primarily for electricity generation, with other objectives incorporated as constraints on operation.

The planning period is divided into T stages, and the state of the system at each stage is described by a state vector. One of the components of the state vector is the reservoir storage level. Transition between stages is constrained by the water continuity equation at the reservoir. A benefit function is associated with each stage t and can depend on (1) the actual release from the reservoir, R_t; (2) the reservoir storage levels at the beginning and at the end of the stage, S_t and S_{t+1}; and (3) the local inflows throughout the basin, described by the vector $\mathbf{Q}_t$. The benefit function $B_t(R_t, \mathbf{Q}_t, S_t, S_{t+1})$ translates these variables into a dollar value for hydropower production that includes avoided thermal costs and capacity benefits, recognizing how much of the release would pass through the turbines and how much would spill or go for other purposes.

A target release from the reservoir for stage t, R_t^* (the decision variable), is obtained by optimizing the objective function $f_t = B_t + B_{t+1} + B_{t+2} + \cdots + B_T + f_{T+1}$, where f_{T+1} describes the value of water at the end of stage T, the last stage in the planning period.

General Dynamic Programming Formulation

If the matrix of local inflows $\{\mathbf{Q}_t, \mathbf{Q}_{t+1}, \cdots, \mathbf{Q}_T\}$, which includes the inflows to the reservoir $\{Q_t, Q_{t+1}, \cdots, Q_T\}$, is known, the actual release should equal the target release and an optimal trajectory for the reservoir can be found by deterministic dynamic programming using the recursive equation

$$f_t(S_t) = \max_{R_t} \{B_t(\ ) + \alpha f_{t+1}(S_{t+1})\} \tag{1}$$

together with the continuity equation

$$S_{t+1} = S_t + Q_t - R_t - e_t(S_t, S_{t+1}) \tag{2}$$

where $e_t(S_t, S_{t+1})$ is the evaporation loss, α is the discount factor, and S_{t+1} is subject to the constraint $S_{\min} \le S_{t+1} \le S_{\max}$. $S_{\min}$ and $S_{\max}$ are the lower and upper bounds, respectively, on storage.

Equations (1) and (2) cannot be used if the inflows that determine the reservoir's evolution are unknown. Instead, stochastic dynamic programming (SDP) can be employed to identify a policy that maximizes the expected value of the objective function.

To take into account the persistence of streamflows in an SDP model, a hydrologic state variable X_t is generally added. The previous month's flow has been the most common choice, though *Gal* [1979] used two preceding flows, and Loucks has often used the current month's value [*Loucks et al.*, 1981; *Stedinger et al.*, 1984]. Recently, *Stedinger et al.* [1984] suggested that using the best forecast of the current period's inflow as a hydrologic state variable can be advantageous. In their case, $X_t = \hat{Q}_t$. In northern California, several sources of information, including snowpack, are used to forecast the snowmelt season's (January–July) runoff. For this situation it seems natural to let the hydrologic state variable X_t be the forecast of the remainder of the seasonal runoff

$$X_t = \sum_{\tau = t}^{\text{July}} \hat{Q}_\tau$$

made at month t (where t may be any month from January through July).

The recursive equation that yields the optimal policy becomes

$$f_t(S_t, X_t) = \max_{R_t^*} \underset{\mathbf{Q}_t|X_t}{E} \left\{ B_t(\ ) + \alpha \underset{X_{t+1}|\mathbf{Q}_t,X_t}{E} [f_{t+1}(S_{t+1}, X_{t+1})] \right\} \tag{3}$$

where, for a target release R_t^* and an inflow Q_t, the actual release R_t is given by

$$R_t = \min \left\{ \begin{array}{l} \max [R_t^*, S_t + Q_t - S_{\max} - e_t(S_t, S_{\max})], \\ S_t + Q_t - S_{\min} - e_t(S_t, S_{\min}) \end{array} \right\} \tag{4}$$

so that the minimum and maximum storage bounds are honored by equation (2). In equation (3), $E_{\mathbf{Q}_t|X_t}$ represents the conditional expectation of the inflow vector $\mathbf{Q}_t$ given X_t. Note that the actual release R_t equals the target release R_t^* when physically possible. In any case, any function of R_t is also a function of R_t^*.

SSDP Approach

In practical applications of SDP, representation of the multivariate distribution of $\mathbf{Q}_t$ given a univariate hydrologic state variable X_t poses a serious problem. It is particularly difficult to represent concisely the high but less than perfect cross correlation among flows at many sites, as well as the autocorrelation of flows at each site.

Alternatively one can use the SSDP approach, first used by *Araujo and Terry* [1974] for the operation of a hydro system and by *Dias et al.* [1985] for the optimization of flood control and power generation requirements in a multipurpose reservoir. With SSDP, one selects M possible streamflow scenarios for the system to describe the joint distribution of reservoir inflows and local inflows. In our example, $T = 12$ and each scenario is a year of observed monthly streamflow data representing one 12-month realization of the corresponding stochastic process. These streamflow scenarios are used to simulate the reservoir's operation and river basin energy production for all possible combinations of storage and hydrologic state in each month.

Let

- $S_t(k)$ reservoir storage at stage t, discretized into K values ($k = 1, \cdots, K$) with $S_t(1) = S_{min}$ and $S_t(K) = S_{max}$;
- $X_t(l)$ streamflow forecast stage t, discretized into L values ($l = 1, \cdots, L$) with $X_t(1)$ corresponding to a "dry" forecast and $X_t(L)$ a "wet" forecast;
- $\mathbf{Q}_t(i)$ vector of inflows throughout the basin at stage t for the ith scenario ($i = 1, \cdots, M$);
- $R_t^*(k, l)$ target release in state (k, l) at stage t;
- R actual release at any stage/state;
- B_t return at stage t due to the release R, given the initial and final storages;
- $f_t(k, l, i)$ benefit of reservoir operation from t through T, when the state is (k, l) and the ith scenario occurs;
- $P_t(i|l)$ probability of the ith scenario at stage t, given streamflow forecast $X_t(l)$;
- $PX_t(v|l, i)$ transition probability from forecast $X_t(l)$ and inflow $\mathbf{Q}_t(i)$ to forecast $X_{t+1}(v)$;
- α monthly discount factor.

For every state (k, l) and stage t, the target release $R_t^*(k, l)$ is obtained by solving

$$\max_{R_t^*} \sum_{i=1}^{M} P_t(i|l)\left[B_t(R_t, i, k, S_{t+1}) + \alpha \sum_{v=1}^{L} PX_t(v|l, i) f_{t+1}(S_{t+1}, v, i)\right] \quad (5)$$

where for each R_t^*, the actual release R_t and ending storage S_{t+1} are given by (4) and (2), respectively. Once $R_t^*(k, l)$ is found, the benefit functions f_t are updated separately for each sequence using the equation

$$f_t(k, l, i) = B_t(R_t, i, k, S_{t+1}) + \alpha \sum_{v=1}^{L} PX_t(v|l, i) f_{t+1}(S_{t+1}, v, i) \quad (6)$$

to reflect the value of the release decision $R_t^*(k, l)$ with each streamflow scenario. Again, the appropriate R_t and S_{t+1} are given by (4) and (2). This procedure is repeated for $t = T, T - 1, \cdots, 1$, and for each t for reservoir storages $k = 1, \cdots, K$, and forecasts $l = 1, \cdots, L$.

These calculations yield an optimal operating policy for the reservoir in the sense that it maximizes the average across the M scenarios of the operating benefits achievable by a release policy that depends on S_t and X_t. SSDP (equation (5)) uses the M streamflow scenarios to describe the distribution of flows over time and space. SDP (equation (3)), on the other hand, generally employs a discrete approximation of a continuous distribution of Q_t given X_t, presumedly based on the observed historical record. Both methods employ a Markov chain model to describe the probabilistic evolution of the hydrologic state variable X_t. The Markov model for X_t in SSDP, while conditioned on the scenario i, is necessary to avoid a one-to-one relationship between the forecast state variable and the scenarios.

Boundary Condition

Solution of equations (5) and (6) requires that the boundary condition $f_{T+1}(S_{T+1}, v, i)$ be given. Initially, one can assume that this function is identically zero. In other words, $f_{T+1}(S_{T+1}, v, i) = 0$ for all S, v, and i. Alternatively, $f_{T+1}(\ )$ can be approximated by the potential energy of stored water S_{T+1} times the average value of energy, for cases where energy values predominate. After an initial solution is obtained for $t = T, T - 1, \cdots, 1$, a second iteration can be performed using the boundary condition

$$f_{T+1}(S_{T+1}, v, i) = \sum_{j=1}^{M} P_{\text{scenario}}(j|i) f_1^o(S, v, j) \qquad \forall\ S, v, \text{ and } i$$

where o identifies the "old" value of f_1 (from the first iteration) and $P_{\text{scenario}}(j|i)$ is the conditional probability of streamflow scenario j following scenario i. Use of appropriate $P_{\text{scenario}}(j|i)$ values can correspond to use of a lag 1 Markov model of either annual or monthly flows. An annual flow model in conjunction with the historical or synthetically generated 12-month annual flow sequences could correspond to adoption of a lag 1 Markov model of annual flows with disaggregation to monthly values. Such synthetic streamflow models are often used for river basin planning [*Salas et al.*, 1980; *Loucks et al.*, 1981; *Pereira et al.*, 1984]. Alternatively, $P_{\text{scenario}}(j|i)$ could be based on the month-to-month correlations. In our example, little year-to-year annual correlation was observed and $P_{\text{scenario}}(j|i)$ was just $1/M$, so that all streamflow scenarios were assigned the same probability. The iterative process proceeds until the operating policy converges.

SSDP for Other Cases

Equation (5) is applicable when the current month's inflow is uncertain, which makes the actual release uncertain. When the inflow vector $\mathbf{Q}_t$ is assumed to be known at the beginning of stage t, as is often done in simulation studies and some SDP studies [*Stedinger et al.*, 1984], the actual release is equal to the target and equation (5) becomes

$$\max_{R_t} \left[B_t(R_t, \mathbf{Q}_t, k, S_{t+1}) + \alpha \sum_{v=1}^{L} PX_t(v|l, Q_t) \sum_{j=1}^{M} P_{t+1}(j|v) f_{t+1}(S_{t+1}, v, j) \right] \quad (7)$$

An interesting case is when $M = 1$ so that there is only a single streamflow scenario. Then both equations (5) and (7) reduce to deterministic dynamic programming and produce the optimal release sequence for that streamflow scenario and the given boundary condition $f_{T+1}(\ )$.

Another interesting case is when $L = 1$ so that the streamflow forecasting capability is neglected and all scenarios are equally likely. In this case, R_t^* is found by solving the simpler problem

$$\max_{R_t^*} \left[\frac{1}{M} \sum_{i=1}^{M} [B_t(R_t, i, k, S_{t+1}) + \alpha f_{t+1}(S_{t+1}, i)] \right] \quad (8)$$

with R_t given by (4), while the cost functions should be updated using

$$f_t(k, i) = B_t(R_t, i, k, S_{t+1}) + \alpha f_{t+1}(S_{t+1}, i)$$

Again, if the inflow vector $\mathbf{Q}_t$ is known at the beginning of stage t, equation (8) should be replaced by

$$\max_{R_t} \left[B_t(R_t, \mathbf{Q}_t, k, S_{t+1}) + \frac{\alpha}{M} \sum_{i=1}^{M} f_{t+1}(S_{t+1}, i) \right] \quad (9)$$

with S_{t+1} subject to $S_{\min} \le S_{t+1} \le S_{\max}$.

The decision variable in equations (3), (5), (7), (8), and (9) was the target release. Similarly, it could be the target storage at the end of the month S_{t+1}^*. In this case, the actual release would not be calculated through equation (4) but through

$$R_t = \min \{R_{\max}, \max [R_{\min}, S_t + Q_t - S_{t+1}^* - e_t(S_t, S_{t+1}^*)]\}$$

where $R_{\min}$ and $R_{\max}$ are bounds on the reservoir releases. The actual end-of-month storage would still be given by equation (2).

Advantages of SSDP

The advantages of the SSDP methodology include the following:

1. The traditional SDP approach fits marginal distributions to the key flows in each month or period, and models the autocorrelation structure of those flows by a Markov process. The resulting conditional distributions are discretized, and the Markov process is then described by the corresponding Markov chain. Use of SSDP avoids the imposition of a specific streamflow model structure; in particular, the joint distribution of flows used in the computation of the benefit-to-go $f_t(\ )$ is described by streamflow scenarios, which can represent not only the empirical marginal distribution but also the empirical joint distribution of the within-year flows. SSDP avoids the distortion that results from discretizing the fitted continuous inflow distributions and then describing the persistence of those discrete values by a Markov chain, a process almost certain to underestimate the severity of droughts. Moreover, the SSDP approach can incorporate a lag 1 Markov model of the year-to-year persistence of annual flows; the implicit Thomas-Fiering monthly streamflow model incorporated in most SDP models yields very little correlation between successive annual flows.

2. In situations where flows at a number of sites in a basin are important, the SSDP methodology readily incorporates the flows for all such sites if their values are available for the historical record period. Moreover, without increasing the dimension of the hydrologic state variable (or using none at all), it captures the entire joint empirical distribution of the within-year flows at all sites.

3. The SSDP methodology attempts to derive the optimal policy for the selected streamflow scenarios using the specified state variables. Thus, in theory, trial-and-error derivation of such a policy through simulation of the same scenarios could not yield a better result.

Probability Distribution of the Streamflow Scenarios

To implement the SSDP algorithm, the following questions must be answered:

1. Given the forecast for the remainder of the snowmelt season X_t, what is the probability of $P_t(i|l)$ of the ith scenario, i.e., that the inflow will equal $Q_t(i)$?

2. What is the probability distribution of the forecast X_{t+1} that would be made next month given that inflow $Q_t(i)$ occurred in month t when the forecast X_t was made?

These probabilities can be precomputed and stored so that they need not be recalculated in each SSDP iteration.

A multivariate distribution for X_t, $Q_t(i)$, and X_{t+1} will be employed, at least implicitly, to find the above probabilities. Let streamflow scenario i have unconditional probability $p(i)$. If the M scenarios correspond to all of those in an M-year historical record, then $p(i) = 1/M$. The computational requirements of the algorithm could be reduced, however, by deleting scenarios that are very similar to others, and giving the remaining scenarios the combined unconditional probability. For example, the scenarios might be ranked by snowmelt season runoff; then, many of the scenarios near the median might be deleted while neighboring scenarios are given the combined probability.

The conditional probability of scenario i given X_t can be calculated using Bayes theorem based on the actual inflow that occurs with sequence i between month t and July

$$Y_t(i) = \sum_{\tau=t}^{\text{July}} Q_\tau(i)$$

The probability assigned to scenario i is

$$P_t(i|X_t) = \frac{p[X_t|Y_t(i)]p(i)}{\sum_{j=1}^{M} p[X_t|Y_t(j)]p(j)} \quad (10)$$

where $p[X_t|Y_t]$ is the probability density function (pdf) of the forecast given the actual inflow Y_t; this pdf can be calculated directly by regressing X_t on Y_t and assuming normal residuals, or from the bivariate distribution of X_t and Y_t.

Also needed is the conditional distribution of X_{t+1} given both X_t and Q_t. (Conditioning on Q_t is especially important in California in months such as May and June when the inflow will likely have a critical impact on the value of the residual snowmelt forecast X_{t+1} made at the beginning of the following month.) Let $X_{t+1}(v)$ be the L discrete values of the forecast in month $t + 1$ and $pX_{t+1}(v)$ the unconditional probability assigned to each. Then, for each sequence i, using Bayes theorem, the conditional probability of each $X_{t+1}(v)$ would be

$$PX_t(v|X_t, Q_t(i)) = \frac{p[X_t, Q_t(i)|X_{t+1}(v)]pX_{t+1}(v)}{\sum_{j=1}^{L} p[X_t, Q_t(i)|X_{t+1}(j)]pX_{t+1}(j)} \quad (11)$$

Here $p[X_t, Q_t|X_{t+1}]$ is the conditional joint pdf for X_t and Q_t given the future forecast X_{t+1}; it can be calculated from the multivariate density function for all three [*Pegram et al.*, 1988]. Alternately, the probabilities $PX_t(v|X_t, Q_t)$ of the discrete forecasts $X_{t+1}(v)$ could be calculated by deriving a conditional pdf for X_{t+1} by regressing on X_t and Q_t and integrating over intervals associated with the selected discrete values of X_{t+1}. *Born* [1988] has shown that there is relatively little difference between these two approaches. Such an option is not attractive for the computation of $P_t(i|X_t)$ because the values of $Q_t(i)$ are those historically observed and thus are irregularly spaced.

APPLICATION OF SSDP TO THE FEATHER RIVER

River Basin Representation

Figure 1 shows a schematic representation of the North Fork Feather River hydroelectric system. The energy generated by the 10 powerhouses in any month is a function of the local inflows (11 inflow points in Figure 1) and of the releases from the nine reservoirs. Table 1 gives the storage range for each reservoir and their useful storage as a percentage of Almanor's; the full useful storage in Lake Almanor is approximately 1.5 times the mean annual flow into the reservoir.

The storage capacity of Lake Almanor overwhelms that of the other reservoirs. In our application of SSDP the primary state variable S_t is the Lake Almanor storage and the decision variable is the target release from Lake Almanor in each month. Mount Meadows, Butt Valley, and Bucks reservoir storages are not state variables. It is assumed that the monthly carryover in each of them, divided by the corresponding useful storage (denoted Δ_i), is equal to the corresponding Δ for Lake Almanor.

Given the release from Lake Almanor and the 11 local inflows, it is possible to estimate the energy produced in the Feather River system by routing the flows from upstream to downstream, maximizing the value of the power generated when there are alternate routings. A simple simulation model performs this calculation, taking into account the water duty (energy generated per unit of water) associated with each powerhouse as well as the upper and lower bounds on the monthly flows due to the capacity of the hydraulic conveyances and to fish releases. The simulation model also takes into account the constant ratio of flows through Caribou 1 and Caribou 2 powerhouses and the variability of the water

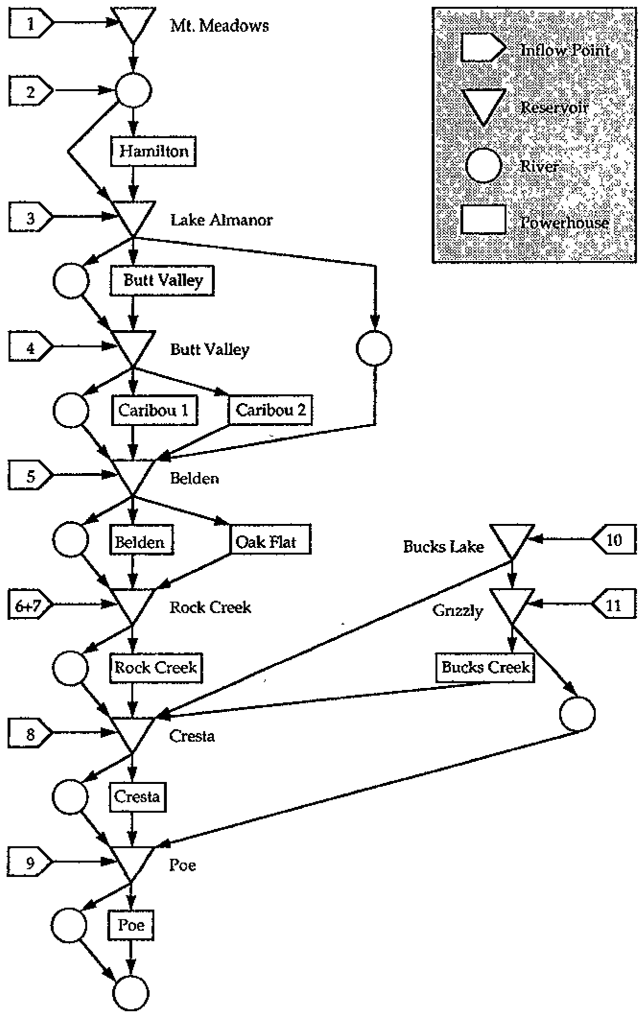

Fig. 1. Schematic representation of the Feather River.

duty of the Butt Valley Powerhouse as a function of the storage level in Lake Almanor. The other powerhouses have constant heads.

The simulation model also checks (and informs the optimization algorithm) whether any powerhouse fails to meet the minimum energy requirement. This requirement depends on the installed capacity and on the minimum capacity factor of each powerhouse. The minimum capacity factor for the

TABLE 1. Reservoir Storages

Reservoir	Maximum	Minimum	Useful	% of Almanor
Mount Meadows	24	2	22	2
Lake Almanor	1143	40*	1103	100
Butt Valley	50	2	48	4
Belden	2	2	0	0
Rock Creek	4	4	0	0
Cresta	4	4	0	0
Poe	1	1	0	0
Bucks	105	35	70	6
Grizzly	1	1	0	0

In thousands of acre-feet (TAF).

*The operational minimum storage of Lake Almanor is more than 40 TAF because of recreational concerns.

TABLE 2. Subjective Coefficients Used in Comparison of SSDP With HYSS Solutions

Month	c_t
January	1.25
February	1.05
March	0.90
April	0.94
May	0.98
June	1.20
July	1.32
August	1.34
September	1.30
October	1.15
November	1.00
December	1.10

summer months (about 10%) is different from the factor for the other months (about 8%) because of the shape of the electric system's load-duration curve.

Comparison With Deterministic Optimization

Deterministic nonlinear programming models such as PG&E's HYSS [*Ikura and Gross*, 1984] need not adopt the Δ simplification, which is the price paid by the stochastic model to avoid the "curse of dimensionality" (the high computational cost of models with multidimensional state variables). The following experiment was done to evaluate the effect of the Δ simplification:

1. The HYSS model was run with the mean monthly inflows for the Feather River. The objective function was simply $\Sigma_t c_t G_t$, where G_t is the total energy generated and c_t are the subjective coefficients in Table 2, which represent relative energy values in different months.

2. The SSDP model was run for the same 12-month streamflow scenario, in which case the SSDP model reduces to an approximation of the deterministic model. The approximation arises because of the Δ simplification.

Figure 2 shows a comparison of the results. The evolution of storage in Lake Almanor is nearly the same, indicating that the two policies are practically equivalent. Indeed, the objective function values for the HYSS and SSDP policies respectively were equal to 4395 and 4391, a difference of just 0.09%. This experiment demonstrates that the Δ simplification should not cause any noticeable degradation in the quality of the derived operating policy.

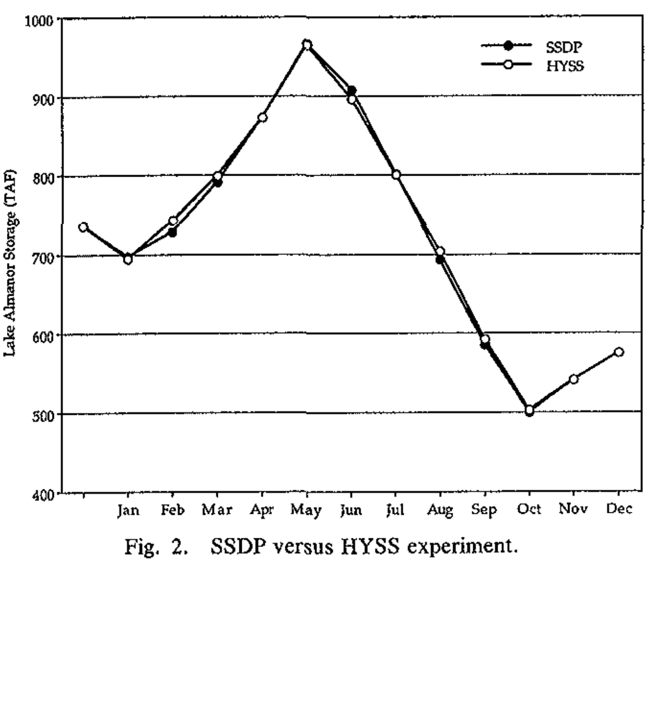

Fig. 2. SSDP versus HYSS experiment.

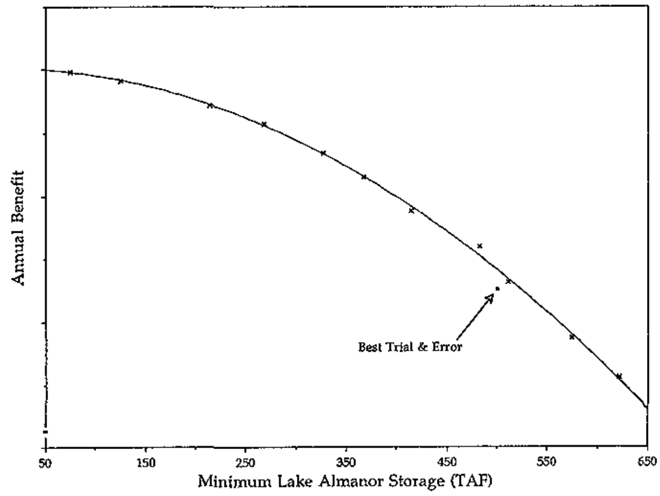

Fig. 3. Sensitivity study results.

No-Forecast Case

The SSDP algorithm was applied to the case of multiple streamflow scenarios but no-forecast hydrologic state variable using equation (8). It was assumed that the monthly inflow to Lake Almanor is not known until the end of the month. The target release is selected at the beginning of each month, based on the state of the reservoir and that target is assumed not to change, regardless of the actual inflow in the month. The minimum capacity factor in each month was checked to prevent any Lake Almanor release that was not sufficient to meet the minimum capacity factor requirement for all powerhouses and streamflow scenarios.

The model was run iteratively using the 57 available historical 12-month streamflow scenarios and 35 storage states. In the first iteration the boundary condition $f_{T+1}(S_{T+1}, i)$ was the potential energy of the stored water (storage times an estimated total of the average water duties). In subsequent iterations the boundary condition was the expected future benefit at the beginning of January from the previous iteration. The value function and release policy converged after the second iteration.

The model was run several times to perform a sensitivity study on an operational constraint at Lake Almanor. In general, Lake Almanor is not drawn below an assumed target minimum level. Obviously, the higher this target level is, the smaller the reservoir operating range. This change in flexibility has a monetary value, because if load is not met by hydropower production, it must be met by thermal generation.

Sets of optimal release rules were generated with the SSDP model for various minimum storage levels. These sets of rules were then applied in the daily simulation model actually used for planning, which incorporates detailed legal requirements. Figure 3 shows the average annual generation benefits calculated by this simulation model for a typical set of power values.

Also shown in Figure 3 is the point corresponding to the

TABLE 3. Correlations Between Seasonal Flows Y_t, Forecasts X_t, and Monthly flows Q_t; State Variables; and Distributions Used to Compute $P_t(i|X_t)$

Month	$\rho(Y_t,X_t)$	$\rho(Q_t,X_t)$	$\rho(Q_t,Q_{t-1})$	State Variable	$P_t(i\|X_t)$ Based On
January	0.72	0.65	0.60	X_t	$Y_t\|X_t$
February	0.75	0.50	0.59	X_t	$Y_t\|X_t$
March	0.82	0.62	0.59	X_t	$Y_t\|X_t$
April	0.89	0.75	0.55	X_t	$Y_t\|X_t$
May	0.83	0.89	0.79	X_t	$Q_t\|Q_{t-1}$
June	0.94	0.93	0.91	X_t	$Y_t\|X_t$
July	···	···	0.91	Q_{t-1}	$Q_t\|Q_{t-1}$
August	···	···	0.91	Q_{t-1}	$Q_t\|Q_{t-1}$
September	···	···	0.90	Q_{t-1}	$Q_t\|Q_{t-1}$
October	···	···	0.62	none	···
November	···	···	0.50	none	···
December	···	···	0.48	none	···

original rule curve used in the simulation model. This rule curve was developed through a difficult trial-and-error process. The SSDP model was able to develop a slightly better rule, as well as rules for other minimum storage levels, with much less time and effort.

SSDP With a Hydrologic State Variable

The SSDP algorithm with a hydrologic state variable was also tested. At PG&E and other western U.S. water management agencies, it is customary to use linear regression models to develop forecasts X_t in month t of the t-through-July seasonal runoff Y_t. These forecasts often make use of previous streamflow and precipitation values, current snowpack water content, and other information. Historical forecasts were available for the Lake Almanor inflow for 1950–1984, allowing the use of equation (10) to develop conditional probabilities $P_t(i|X_t)$ for each scenario i for which historical forecasts were available in the months of January–April and June. For May, equation (10) was employed with May's flow Q_t replacing Y_t because a higher correlation was observed between Q_t and X_t than between Y_t and X_t. For the months of July–September, the previous month's flow Q_{t-1} served as the hydrologic state variable; so, Q_t and Q_{t-1} replace Y_t and X_t in equation (10). Finally, for the months of October–December, since forecasts are not developed and the month-to-month inflow correlations were more modest, no hydrologic state variable was employed. These decisions are summarized in Table 3.

The marginal (unconditional) probabilities computed for the various streamflow scenarios, using the conditional probabilities from (10) with the given $pX_t(v)$, deviated appreciably from the expected value of $1/M$, particularly for the January–June period. This deviation was found to be caused primarily by the historical streamflow record being substantially longer than the historical forecast record [*Born*, 1988]. To correct the anomalies, the prior probabilities $p(i)$ used in equation (10) were adjusted using a first-order Newton scheme until the resultant marginal probabilities for the scenarios closely approximated $1/M$. As a result, the unconditional mean and variance of the flows, and the month-to-month correlations, very closely matched the historical values.

TABLE 5. SSDP Results With and Without Forecasting

Minimum Storage (TAF)	Value of Energy With	Value of Energy Without	Value of forecasts, %
	Set 1		
500	3109.3	3103.6	0.2
900	2927.5	2916.1	0.4
	Set 2		
500	7961.4	7844.4	1.5
700	7753.3	7587.3	2.2
800	7485.6	7311.4	2.4
900	6850.0	6680.9	2.5

In all months when the seasonal forecast X_{t+1} served as the next month's hydrologic state variable, both X_t and Q_t were found to be significant at the 5% level in a predictive linear model for X_{t+1}; so both were employed in equation (11) to calculate the transition probabilities $PX_t(v|l, i)$. When Q_t served as the next month's hydrologic state variable, and $Q_t(i)$ had been observed in month t, then the transition probabilities $PX_t(v|l, i)$ corresponded to interpolation between the two discrete values of the Q_t state variable that bracket $Q_t(i)$ because $f_{t+1}(\ \ , v,\ \ )$ was only computed for $v = 1, \cdots, L$ (rather than $1, \cdots, M$). (With traditional SDP formulations such interpolation is generally not necessary because Q_t is restricted to only the allowable discrete values.) In the test case, $L = 5$, corresponding to standard normal quantiles of ±1.72, ±0.76, and 0 with probabilities of 0.107, 0.245, and 0.296, respectively.

The SSDP model was run with and without forecasts for two sets of cases, using the subjective objective function coefficients c_t in Table 4. The results are summarized in Table 5. In the first set of cases the c_t coefficients were all set to one; thus the objective was to maximize the energy generated by the Feather River system. The difference between the energy production levels obtained with and without forecasting is small, less than 0.4%. In both cases the SSDP release rules generally kept Lake Almanor's storage level low so that large winter and spring inflows could be stored and used to generate energy when spills around the powerhouses are less likely. Moreover, about half of the inflow in the basin enters the system below Lake Almanor and hence is essentially uncontrollable on a monthly time scale.

In the second set of cases the objective function coefficients reflect projected future relative energy values in a year with average hydrologic conditions (neither wet nor unusually dry). As can be seen, energy in May is expected to have

TABLE 4. Subjective Coefficients Used to Investigate the Value of Forcasts

Set	Jan.	Feb.	March	April	May	June	July	Aug.	Sept.	Oct.	Nov.	Dec.
1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
2	2.12	1.56	2.07	1.83	0.58	1.66	3.11	3.09	3.09	3.24	3.17	2.88

relatively little value because of plentiful hydroelectric energy available in that month; summer and fall energy prices are relatively high, reflecting relatively low natural flow levels and relatively high energy demands. Thus it is advantageous to store spring inflows in Lake Almanor and generate energy in the summer and fall, provided that the increased value is not offset by lost value from (1) spills at Almanor that must bypass some powerhouses or (2) spills below Almanor when unregulated inflows are larger than anticipated. (Flood control for communities in California's Central Valley is provided by the state's large Oroville Reservoir downstream from PG&E's hydroelectric system.)

As can be seen in Table 5, with the realistic energy value function, the value of using forecasts ranged from 1.5 to 2.5% depending on the minimum Lake Almanor storage level. Over the range considered, the higher the minimum storage, the greater was the value of forecasts. (For minimum storage levels above 900 TAF the useful storage was insufficient to provide required flows to meet the minimum capacity constraint). In general, the additional expected value exceeded one million dollars per year for this basin.

CONCLUSIONS

This paper has explored the use of sampling stochastic dynamic programming and extended previous descriptions of the approach by the introduction of a hydrologic state variable and the associated conditional distributions for the various streamflow scenarios and values of future hydrologic state variables. Sampling stochastic dynamic programming as developed here employs the empirical multivariate temporal and spatial streamflow distribution for a basin, allowing the detailed simulation within the optimization model of PG&E's complex Feather River hydroelectric system. In this regard, sampling SDP has significant advantages over the traditional SDP approach. Because sampling SDP employs selected historical or synthetic streamflow traces, the actual multimonth persistence of streamflows can be captured in the calculation of the expected benefits.

The sampling SDP approach has been useful in PG&E's hydropower planning because it can generate efficient operating policies for alternative and proposed hydroelectric system configurations faster than the traditional trial-and-error deviation of rule curves to prescribe reservoir operations. The inclusion of the use of forecasts in the prescription of reservoir releases can allow for the development of even more efficient operating policies. However, in our examples, the value of seasonal forecasts depended on the objective function employed and the system configuration. When we attempted simply to maximize the average energy produced, the optimal policy generated without forecasts operated almost as well as the policy that used forecasts. We also considered a case where energy generated in the summer and fall had increased value over spring generation, as it does in northern California. Then forecasts were of value in planning reservoir operations that attempted to transfer energy into the summer and fall without keeping Almanor so full that large inflows result in avoidable spills around some hydroelectric powerhouses.

REFERENCES

Araujo, A. R., and L. A. Terry, Operation of a hydrothermal system (in Portuguese), *Braz. J. Electr. Energy*, 1974.

Born, P. H. S., Hydrologic models and their representation in stochastic dynamic programming algorithms for reservoir optimization, M.S. thesis, Cornell Univ., Ithaca, N. Y., 1988.

Bras, R. L., R. Buchanan, and K. C. Curry, Real time adaptive closed loop control of reservoirs with the High Aswan Dam as a case study, *Water Resour. Res.*, *19*(1), 33–52, 1983.

Dias, N. L. C., M. V. F. Pereira, and J. Kelman, Optimization of flood control and power generation requirements in a multipurpose reservoir, paper presented at Symposium on Planning and Operation of Electric Energy Systems, Int. Fed. of Autom. Control, Rio de Janeiro, Brazil, 1985.

Foufoula-Georgiou, E., and P. K. Kitanidis, Gradient dynamic programming for stochastic optimal control of multidimensional water resources systems, *Water Resour. Res.*, *24*(8), 1345–1359, 1988.

Gal, S. Optimal management of a multireservoir water supply system, *Water Resour. Res.*, *15*(4), 737–749, 1979.

Grygier, J. C., and J. R. Stedinger, Algorithms for optimizing hydropower systems, *Water Resour. Res.*, *21*(1), 1–10, 1985.

Grygier, J. C., J. R. Stedinger, and H.-B. Yin, A generalized maintenance of variance extension procedure for extending correlated series, *Water Resour. Res.*, *25*(3), 345–349, 1989.

Ikura, Y., and G. Gross, Efficient large-scale hydro system scheduling with forced spill conditions, *IEEE Trans. Power Appar. Syst.*, *103*(12), 3502–3520, 1984.

Johnson, S. A., J. R. Stedinger, and C. A. Shoemaker, Computational improvements in dynamic programming, *Forefronts 4(7)*, pp. 3–7, Cent. for Theory and Simulation Sci. and Eng., Cornell Univ., 1988.

Loucks, D. P., J. R. Stedinger, and H. A. Haith, *Water Resource Systems Planning and Analysis*, Prentice-Hall, Englewood Cliffs, N. J., 1981.

Neto, T. A. A., C. B. Cotia, M. V. F. Pereira, and J. Kelman, Comparison of stochastic and deterministic approaches in hydrothermal generation scheduling, paper presented at IFAC Symposium on Planning and Operation of Electric Energy Systems, Int. Fed. of Autom. Control, Rio de Janeiro, Brazil, 1985.

Pegram, G., J. R. Stedinger, and P. H. S. Born, Statistical analysis program for SDP optimization, report, Dep. of Environ. Eng., Cornell Univ., Ithaca, N. Y., April 27, 1988.

Pereira, M. V. F., G. C. Oliveira, C. C. G. Costa, and J. Kelman, Stochastic streamflow models for hydroelectric systems, *Water Resour. Res.*, *20*(3), 379–390, 1984.

Saad, M., and A. Turgeon, Application of principal component analysis to long-term reservoir management, *Water Resour. Res.*, *24*(7), 907–912, 1988.

Salas, J. D., J. W. Delleur, Y. Yevjevich, and W. L. Lane, *Applied Modeling of Hydrologic Series*, Water Resources Publications, Littleton, Colo., 1980.

Stedinger, J. R., B. F. Sule, and D. P. Loucks, Stochastic dynamic programming models for reservoir operation optimization, *Water Resour. Res.*, *20*(11), 1499–1505, 1984.

Yakowitz, S., Dynamic programming applications in water resources, *Water Resour. Res.*, *18*(4), 673–696, 1982.

Yeh, W. W.-G., Reservoir management and operations models: A state-of-the-art review, *Water Resour. Res.*, *21*(1), 1797–1818, 1985.

Yeh, W. W.-G., L. Becker, and R. Zettlemoyer, Worth of inflow forecast for reservoir operation, *J. Water Resour., Plann. Manage. Div. Am. Soc. Civ. Eng.*, *108*(WR3), 257–269, 1982.

Young, G. K., Jr., Finding reservoir operating rules, *J. Hydraul. Div. Am. Soc. Civ. Eng.*, *93*(HY6), 297–321, 1967.

L. A. Cooper, E. Hsu, and S.-Q. Yuan, Pacific Gas and Electric Company, 77 Beale Street, San Francisco, CA 94106.

J. Kelman, CEPEL, Caixa Postal 2754, CEP 20001, Rio de Janeiro, RJ Brazil.

J. R. Stedinger, School of Civil and Environmental Engineering, Hollister Hall, Cornell University, Ithaca, NY 14853.

(Received December 9, 1988;
revised August 16, 1989;
accepted August 24, 1989.)