

Short-Term, Single, Multiple-Purpose Reservoir Operation: Importance of Loss Functions and Forecast Errors

BITHIN DATTA AND STEPHEN J. BURGESS

Department of Civil Engineering, University of Washington

Short-term operation policy for multipurpose reservoirs can be derived from an optimization model with the objective of minimizing short-term losses (opportunity costs). Construction of such loss functions requires the definition of target values for the decision variables, assessment of reliabilities with which inflows can be predicted, and an explicit statement of operational objectives. Formulation and evaluation of a model is complicated by the uncertainties inherent in the prediction of future streamflows and by controversies about the criteria of evaluation. We discuss these issues and illustrate some of our arguments with simple numerical experiments. A series of synthetic short-term forecasted values (which satisfy a specified distribution of forecast errors) is used to examine operation of a single reservoir. The quality of forecasted values is represented by the mean and variance of these errors or the coefficient of prediction (C_p). The objective function of the operation model is assumed to be the best possible tradeoff between probable deviations from two operation targets: release and storage volume. Reservoir release was effected according to the solution of the optimization model conditioned upon the forecasted streamflow volumes for a given time increment. The storage volume was then corrected to reflect actual streamflow for the forecasted period. This became the initial storage for the next forecast period. Actual losses, deviations between actual and forecasted losses, the variance of storage and release volumes, and operational performance measures, including reliability, resiliency, and vulnerability, were found to be sensitive to the relative importance given to deviations from release or storage targets and the quality of forecasts. The performance of an operation policy based on a model that uses predicted streamflows as deterministic inputs cannot be correlated directly with the shape of the assumed loss function.

INTRODUCTION

Most work reported in the literature concerning reservoir operation has used deterministic streamflow forecasts or has involved schemes for developing operating rules that are based on a single historical streamflow record. This means that the adopted policies are for a perfect flow forecast situation; however, operation policies should be developed on the basis of system operation being subject to uncertain forecasts of future inputs and demands. Some conclusions other investigators have reached are based on a single objective, e.g., meeting a release target. We show here that consideration of a second objective, the actual storage state, as well as noisy forecasts of future streamflow are important issues that substantially influence mathematically derived reservoir operation policies.

To illustrate the incorporation of a second objective and noisy forecasts into the problem of operating the simplest possible single reservoir, it is necessary to discuss multiple objectives, economic loss functions, the expected value decision-making criterion, the nature of decision variable targets, decision model structures, and aspects of short-term streamflow forecasts. As an outgrowth of a two-objective numerical example we found that earlier discussions of recently formalized terms: resiliency, robustness, and vulnerability [Hashimoto *et al.*, 1982a, b; Fiering, 1982a, b, c, d] need to be expanded.

Our work focused on developing reservoir operation policy in real time conditioned on uncertain future short-term streamflow into the reservoir. (Real time operations require decisions on flow release rates for short time increments, typically on the order of 1 hour. We have used release volumes corresponding to periods on the order of one week). Results

from models that do not reflect this aspect of reality, no matter how elegantly structured and solved, will give results that may have little utility. We have avoided dogmatic definitions; hence our use of "short-term" covers periods on the order of hours to weeks. In our numerical example a 5-day forecast increment was chosen arbitrarily for illustrative purposes only.

The objectives of operation are of prime importance in developing an operation policy, whether they are stated explicitly in the objective function or incorporated implicitly as binding constraints in the model. An example is provided by the chance-constrained models proposed for planning purposes [ReVelle *et al.*, 1969; Loucks and Dorfman, 1975; Houck, 1979] or for operation purposes [Datta, 1981; Datta and Houck, 1984]. All these models have to satisfy some implicit objectives, for example, meeting minimum and maximum storage and release bounds with specified reliabilities. These multiple implicit objectives stated in the form of constraints are, in addition to the objective, stated explicitly in the objective function. When only a single value of streamflow for a given period (which may be the actual value or a forecast) is used as an input to the model, the multiplicity of these implicit objectives is reduced. The optimum solution obtained by using such limited streamflow information may be far from satisfactory. When such models use an externally forecasted streamflow value, their performance depends on the quality of the forecast as characterized by the forecast error distribution.

The forms of loss functions and definitions of decision variable targets, as well as the expected value criterion, hydrologic and economic uncertainties, and storage release tradeoffs, are central to any operation model development. There are conflicting definitions and conflicting opinions in the literature; therefore we devote the next section to these issues. For short-term reservoir operation, where the time horizon for the single solution of the operation model may range between hours and a few months, the impact of economic uncertainty in the form

Copyright 1984 by the American Geophysical Union.

Paper number 4W0693.
0043-1397/84/004W-0693\$05.00

of future discount rates is minimal. Therefore, we chose not to discuss economic uncertainty in any detail.

COMMENTS ON LOSS FUNCTIONS

Long-term reservoir operation policy is derived typically from an optimization model with an objective of maximizing long-term benefits or minimizing long-term losses. For short-term operations the objective is to maximize short-term benefits or minimize the opportunity cost of a decision.

Losses are negative benefits; actual loss or benefit functions are assessed or specified objectively and reflect the actual benefits or extent of damages corresponding to a given state. On the other hand, penalty functions reflect opportunity costs. Any decision that causes a decrease from the maximum possible benefit (or any increase from the minimum possible loss) incurs a penalty.

While it is quite difficult to construct an accurate benefit or loss function, it is more difficult to assess the opportunity cost of a decision. In many cases the shape of the penalty function is imposed by the decision makers to reflect a policy designed deliberately to achieve specific results. In short-term reservoir operation the penalty function is obtained from the long-term benefit function, the short-term benefit function, a planning target (a release guaranteed with high reliability), and an operational target (corresponding to the maximum of short-term benefits for a planned release). Several authors [Hashimoto *et al.*, 1982a; Klemes, 1978] have described the penalty function as a loss function (for deviation from a target value). Therefore, we continue to describe the penalty function for short-term reservoir operation as a loss function.

Unresolved questions regarding the best choice of a loss function for reservoir operation include the issues of convexity, concavity, or symmetry of the loss functions assumed. Stedinger [1978] has argued that penalization of releases in excess of the target value is unrealistic. Klemes [1978] disagreed with this argument and pointed out that it depends on the definition of "target" and on the shape of the loss functions for deviations from the target value. According to his proposition, a release target may be defined either as a "scale of development" (as done, for instance, in Fiering [1967], i.e., the release in excess of this target value generates benefits; failure to meet this target value is associated with severe economic penalties) or as a value that causes no losses (the value corresponding to the minimum of the loss function or maximum of the benefit function as used by Klemes [1977]). When the second definition of a target value is accepted, it is possible to assume the loss to be zero (or a constant) in the vicinity of the target, implying no losses for small deviations from this particular value and progressively higher losses for larger deviations in both directions.

If the only objective of operating a reservoir or reservoir system is to ensure a dependable flow during dry periods, and other objectives are ignored, it is possible to adopt a loss function that constitutes only the dry branch of a two-sided generalized loss function. A two-sided loss function may be necessary when multiple objectives, e.g., recreation, flood damage mitigation, navigation, water supply, and hydropower are important.

Klemes [1979] pointed out that considering a loss function $L(y) = y^a$, where y is the outflow from a single reservoir (with mean $E(y)$; $E(\cdot)$ is the mathematical expectation), for a convex loss function ($a < 0$ or $a > 1$) a constant release equal to $E(y)$ is economically superior to a sequence of variable releases. For a concave loss function ($0 < a < 1$) a variable release, or a

constant release equal to the target until the reservoir is full and then a release equal to the capacity of the reservoir, is superior to a constant release. For a linear or a constant loss function ($a = 1$ or $a = 0$) the overall economic effect is independent of the outflow pattern. Accordingly, it appears that no general optimization is possible for the last two cases. It should also be noted, however, that these conclusions are based on restrictive assumptions. Klemes's results are valid for an unconstrained (semiinfinite) reservoir and for a zero discount rate for future returns. For completeness we include here relationships between finite and infinite reservoirs. The essence of these relationships, contained in the next three paragraphs, were provided to us by Vit Klemes (personal communication, 1983).

For a finite reservoir with a storage capacity K the optimum operation policy must be rephrased as "release equal to $E(y)_{\Delta t}$ or to $E(y)_{\Delta t} \pm K/\Delta t$," where $E(y)_{\Delta t}$ is the mean inflow during one reservoir "working cycle" Δt , i.e., during a period between two successive instants when the reservoir is either empty or full. Thus the above operation policy requires perfect knowledge about future inflows for one working cycle ahead; given the fact that the length of this cycle increases with reservoir storage capacity, for an infinite reservoir this knowledge must extend over the entire future [Klemes, 1979].

Since a zero discount rate was taken as implying a complete absence of economic uncertainty [Klemes, 1977], the above policy can be regarded only as an "ideal" optimum policy, i.e., as a limit corresponding to "absolute certainty," both hydrological and economical [Klemes, 1979]. On the other hand the standard operation policy (release equals target, whenever possible) resulted in the opposite limit, i.e., as an optimal policy under conditions of "absolute uncertainty" (infinite discount rate and complete ignorance about future inflows [Klemes, 1977]).

It is also important to note that all of the above propositions were based on the assumption that the minimum expected value of losses is an appropriate criterion for optimal operation. However, the validity of this criterion, especially for optimization of a single project has been questioned [Klemes, 1978, 1979].

It is appropriate, therefore, to suggest that it is not possible to guarantee the optimality of the "ideal" operation criterion of making a release equal an estimate of the input mean for a given period, without perfect hydrologic information and/or building a large (semiinfinite) reservoir. On the other hand the standard operation policy suffers from severe limitations that will be discussed in a later section.

Expected Value of Losses

One of the drawbacks of justifying any result according to the expectation criterion is that actual values and expected values differ. In the past, many arguments have been made for and against using expected returns from a particular policy as the sole criterion for decision making. The objection generally cited is that "the expected value criterion does not take into consideration the variations in return" [Nemhauser, 1966].

The above reasoning for not using expected returns as the sole decision making criterion is only superficially valid. When utility measures (according to the von Neumann-Morgenstern Cardinal Utility Theory), not monetary values, are considered, the criterion of maximizing expected utilities accounts for the risk associated with probable returns and the decision maker's preferences for a combination of values. In practice the major difficulty is in constructing a proper utility function. Even with

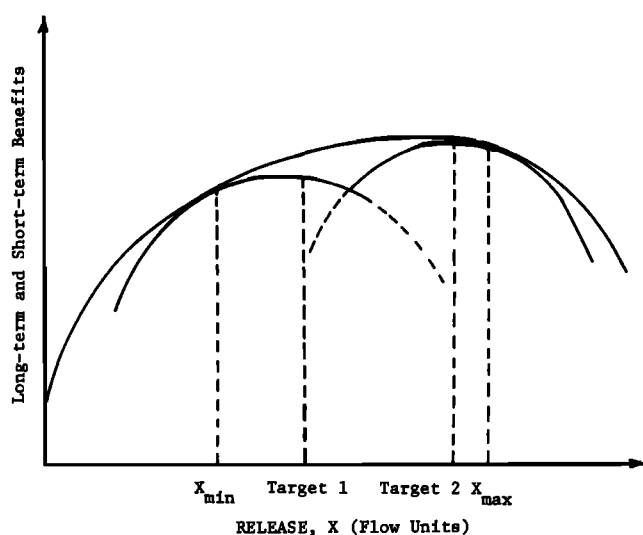


Fig. 1. Long-term and short-term benefit functions for reservoir release.

proper measures for returns, an expected utility criterion is impractical for water resource or other socioeconomic systems because usually more than one decision maker (representing different constituents) is involved. A single many-to-one mapping that relates all measures of importance to a utility scale is not obtainable in such cases. Therefore, expected return as a criterion for decision making may be unsuitable for water resources systems optimization problems.

Target Values for Decision Variables

Theoretical target reservoir releases, shown schematically in Figure 1, can be defined in two ways. The first concerns a value, X_{min} , guaranteed with high reliability for short-term allocation. For a release equal to X_{min} , short-term benefits equal long-term benefits, and no penalties are incurred. Releases in excess of X_{min} give rise to increased benefits, although these benefits are lower than could have been attained (from the long-term benefit function) had the guaranteed amount been higher. The long-term benefit function should start dipping down at some value of reservoir release X^* , when an incremental release will cause a problem such as flooding, damage to aquatic life, or loss of recreational opportunity. The short-term benefit function will also dip down, but for a corresponding release, $X < X^*$, except for the unusual case of X_{min} coinciding with X^* , in which case $X = X^*$. (The long-term benefit function is the envelope of short-term benefit functions.)

If the main concern of reservoir operation is to mitigate the impacts of high flows, i.e., flooding, similar arguments can be made. Now the maximum outflow that may not be exceeded with a high reliability is defined as X_{max} ; any decrease in release volumes from X_{max} might increase actual benefits. At a certain point, however, disbenefits occur when the amount planned for was X_{max} , and the actual outflow is appreciably smaller.

A second definition of target release can be developed by using the release magnitudes (target 1 and target 2) shown in Figure 1. These values are treated as variables where short-term benefits are a maximum, i.e., penalties (opportunity cost for given X_{min} or X_{max}) are zero, as defined by Klemes [1978]. (Short-term losses are at a minimum for releases target 1 and target 2 where the penalties are zero.)

Because target 1 and target 2 are conditioned on X_{min} and X_{max} , respectively, some readers may be tempted to define X_{min} and X_{max} as "planning targets" and targets 1 and target 2 as operational targets. However, we refrain from such rigid definition. It is more useful to emphasize that X_{min} or X_{max} are releases guaranteed with high reliability (a decision variable), and target 1 and target 2 are based on these values. (It is possible to compute all these quantities from an operation policy optimization model.) Selection of target 1 or target 2 is conditioned on the objectives of operation. Target 1 is important for water supply objectives; target 2 is associated typically with flood flow management.

Reservoir Release Targets and Optimal Operating Policy

When the target value is defined as X_{min} (or X_{max}), so that an increase (or decrease) of release from this value is not penalized, some investigators claim the Standard (Linear) Operating Policy (SLOP), an acronym suggested by colleague Rick Palmer, is the optimal operation criterion. This policy requires release of the (fixed) target (X_{min}) whenever enough water is available or release of as much as is available when the storage plus inflow is lower than X_{min} ; when the reservoir is full, the release is equal to X_{min} plus the spill. Klemes [1977] questioned the optimality of such a policy and claimed it leads to inferior release decisions. Stedinger [1978] contested this argument, claiming that Klemes used an unrealistic objective function that penalized both positive and negative deviations from the target value. It is now clear that Stedinger's argument is true only when X_{min} is assumed as the target.

Some important points must be noted here. The main drawback of the SLOP is that, while it is simple to specify a release policy ("release the target volume whenever enough water is available,"), at any time a decision to release is made, this volume can only be estimated; therefore uncertainty must be considered explicitly. This uncertainty introduces the need for denoting X_{min} (or X_{max}), which might not be the same as target 1 (or target 2), as shown in Figure 1. If there were no uncertainties involved, there would be no need to differentiate between X_{min} and target 1 (or X_{max} and target 2). (X_{min} and X_{max} are typically used in chance-constrained models as lower and upper bounds corresponding to specific reliability levels.) Further complications arise when the actual rate of release on an hour-to-hour basis is to be ascertained once these release volumes have been specified. A detailed discussion of this aspect is beyond the scope of this paper.

The target value should be interpreted as that volume of water for which penalties are a minimum; any deviation from the vicinity of this value is penalized to whatever extent is appropriate. The exact shape of the benefit function will vary from basin to basin and may also be modified according to the perceptions of decision makers. Hashimoto *et al.* [1982a] wrote that loss functions of the type $L(X) = [(T - X)/T]^\beta$ for $X < T$ and $L(X) = 0$ for $X \geq T$ (X is the release and T the target release; β is a constant), when incorporated in an optimization model that minimizes the expected value of losses subject to some physical constraints, result in different types of policies, depending on the value of β . They reported that the operation policy specified by the loss function for $\beta > 1$ results in hedging from the target release, even if enough water was available (as evident from hindsight). This does not occur for $\beta < 1$. This result can be visualized conveniently by examining the loss function $(T - X)^\beta$ shown in Figure 2. Let p denote a possible deviation from the target release at the end of operating period 1, q the corresponding possible deviation during




Fig. 2. Shapes of loss functions for different values of β .

period 2, and $p + q$ the combined deviation at the end of the two operating periods.

It is now evident that for an expected deviation equal to p at the end of period 1 and q during period 2 the losses can be minimized if the combined deviation is actually postponed up to the end of period 2 for $\beta < 1$. If $\beta = 1$, the delay does not affect the losses and therefore may result in alternative optimum solutions. In the case of $\beta > 1$, because $p^\beta + q^\beta < (p + q)^\beta$, it is always costlier to postpone the deviations and preferable to incur maximum losses at the first period rather than adding them up for the second period. This conclusion is valid when using the expected loss criterion in the objective function because the transition probabilities (from one discrete flow state to another) used are assumed stationary and may be considered as constant multipliers to the loss magnitudes.

Hedging from the SLOP is only a consequence of different β values, as indicated by Hashimoto *et al.* [1982a]. The extent of hedging depends on the state value and the transition probabilities from one discrete state to another. By assuming a one-sided loss function, when additional releases in excess of the target value are not penalized, it is immaterial how much is released once the target is met as a lower bound.

The conclusion that a target release value equal to the mean inflow is optimum is not acceptable if, for example, variations in storage levels are of any concern in determining operation policy. It may be possible to release the target volume(s) whenever there is enough water available, but this may cause great variations in the reservoir level(s), which may be undesirable for recreational or hydropower generation purposes. Penalties not only for deviations from release targets but also from prespecified storage targets are needed for even the simplest reservoir systems.

Multiple Targets

The previous discussion has considered a release target. A storage target is now considered in addition to a release target, making the operating problem a two decision variable situation for a single time step. The target storage is not the same as a minimum storage because it is reasonable to assume that in many possible objectives of reservoir operation a positive or a negative deviation from the storage target may be undesirable. Once a storage target is introduced into decision-making considerations, many of the observations discussed

above are rendered inadequate. Now, more than one loss function, one for deviations from the target release and one for the deviation from target storage, must be considered. An increase in release causes a decrease in storage and vice versa. Therefore these two objectives become conflicting; preference ordering of the objectives (release and storage) becomes necessary to reach a single solution.

Cohon and Marks [1975] pointed out that the above vector optimization problem cannot be optimized in a strict sense. In the absence of preference information an optimal solution cannot be found to the problem because all feasible solutions are not ordered. However, an incomplete ordering may be obtained from the set of noninferior solutions where alternatives are eliminated only on the basis of inferiority. A basis for more complete ordering is the articulation of preferences prior to the solution. Therefore we have adopted the goal-programming technique in the examples given later for optimizing operation of a single reservoir where two objectives are to be satisfied.

While goal programming provides some help, the reservoir operation problem is complicated further by uncertainties introduced by the unknown future inflows to the reservoir. We explore this issue explicitly by examining the conditional distributions of future inflows where the future inflows are assumed to be forecasted with different levels of precision.

QUALITY OF INFLOW VOLUME FORECASTS

Lettenmaier [1984] proposed a simple index for measuring the quality of forecasted streamflow volumes. This index, the coefficient of prediction C_p , is defined as

$$C_p = 1 - \frac{E(\bar{X}_t - X_t)^2}{\sigma_{\Delta t}^2} \quad (1)$$

where \bar{X}_t is forecasted streamflow volume for the time t , X_t is actual streamflow volume for the time t , $\sigma_{\Delta t}^2$ is unconditioned variance of streamflow for the same time interval, $E(\cdot)$ is the expectation operator, and Δt is the length of each time period. The coefficient of prediction has a useful range from 1.0 for a perfect forecast (forecast error variance = 0) to 0.0 for a forecast containing no information other than the mean of the forecast period streamflow. Negative values of C_p mean that the forecast contains negative information (forecasts that are more variable than the historically recorded runoff volume for the same time steps). This index can be used for forecasts generated by stochastic models or by conceptual or physics-based models, so long as the errors are considered stationary for a limited time horizon.

NUMERICAL EXPERIMENTS

The importance of the issues introduced above is explored for a single hypothetical reservoir operated to derive benefits from both release and the volume of water stored. We deliberately kept the model simple to ensure that attention remains focused on tradeoffs between conflicting objectives in the presence of uncertain inflows.

Operation Model

The objective was minimization of losses from deviations from two targets: release and storage volumes. The nonlinear loss functions used (Figures 3a and 3b) were piecewise linearized and incorporated into a deterministic linear (goal) programming formulation:

Minimize

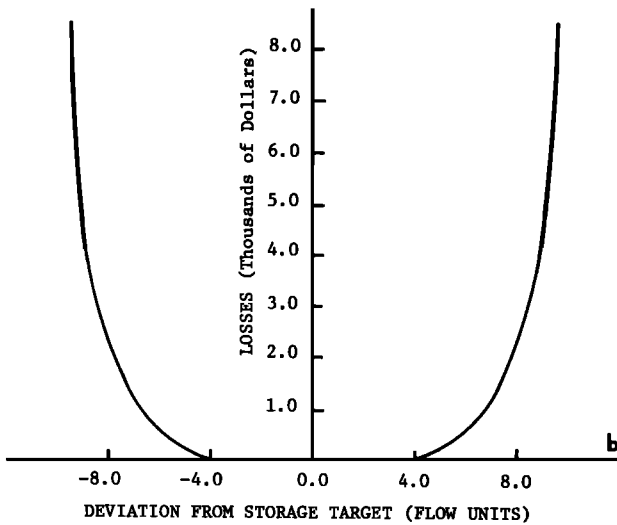
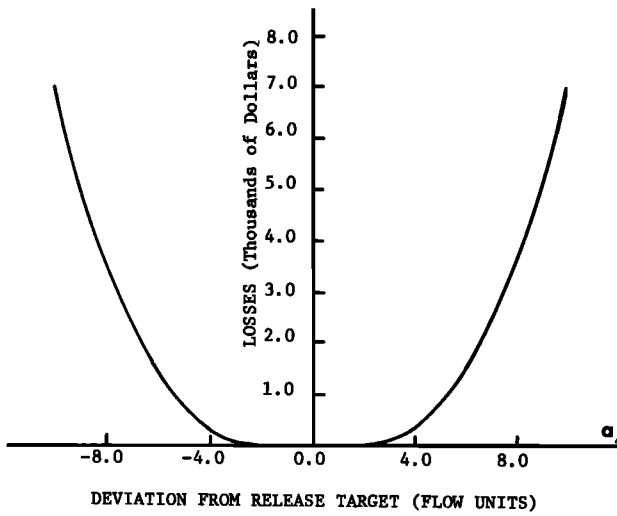


Fig. 3. Loss function for deviations from (a) release target and (b) storage target.

$$k_1 * [L_r'(D_t^r + E_t^r)] + k_2 * [L_s^s(D_t^s + E_t^s)] \quad (2)$$

Subject to

$$S_{t+1} = S_t + \bar{X}_t - R_t \quad (3)$$

$$S_{t+1} \leq CAP \quad (4)$$

$$S_{t+1} \geq S_{min} \quad (5)$$

$$S_{t+1} + D_t^s - E_t^s \geq T^s \quad (6)$$

$$R_t + D_t^r - E_t^r \geq T^r \quad (7)$$

$$D_t^s, E_t^s, D_t^r, E_t^r \geq 0 \quad (8)$$

where

k_1 a constant weight given to the losses due to deviation from the flow release target;

k_2 a constant weight given to the losses due to deviation from the storage volume target;

$L_s^s(\cdot)$ loss function for storage deviation from target;

$L_r^r(\cdot)$ loss function for release deviation from target;

- S_t storage at the beginning of time t ;
- \bar{X}_t forecasted streamflow volume at time t ;
- R_t release during the time interval t to $t + 1$;
- T^s storage target;
- T^r release target;
- CAP capacity of the reservoir;
- S_{min} minimum allowable storage;
- D_t^s, E_t^s deviations from storage target;
- D_t^r, E_t^r deviations from release target.

The variables D and E are included in the formulation to facilitate calculation of losses. The loss functions (Figures 3a and 3b) are positive for both positive and negative deviations from the targets. D and E are nonnegative; D covers the branch of the loss functions corresponding to negative deviations; E covers the branch of the loss functions corresponding to positive deviations. For any solution of the model, if D is positive, E is forced to zero and vice versa. The weights k_1 and k_2 sum to unity.

Flow Forecast Procedure

For convenience we chose to generate each forecast sequence as the sum of the recorded streamflow for a 5-day period plus an error term. Other procedures would be used in practice, but this was sufficient for our purpose. The forecasted flow \bar{X}_t is given by

$$\bar{X}_t = X_t + \epsilon_t \quad (9)$$

where ϵ_t is the error term.

Lettenmaier [1984] gives complete details for this scheme. The correlation structure of daily flows (the shortest time step for available recorded data) was assumed to be Markovian, with an implied correlation coefficient ρ for each specified value of C_p .

The error term ϵ_t is assumed to have a Markovian structure; with this assumption the variance and correlation coefficient for ϵ_t can be computed for a specified value of C_p and the corresponding value of ρ .

The relevant equations are

$$\sigma_\epsilon = (1 - C_p)^{0.5} \sigma_{\Delta t} \quad (10)$$

where σ_ϵ is the standard deviation of ϵ_t and $\sigma_{\Delta t}$ is the standard deviation of historical data corresponding to a time increment Δt . (We used $\Delta t = 5$ days and 2 years of flow data; hence $\sigma_{\Delta t}$ was computed from 146 consecutive 5-day flow volumes.)

The random (Gaussian) component in the Markovian error term has zero mean and variance

$$(1 - \rho_{\epsilon_t, \epsilon_{t-1}})^2 \sigma_\epsilon^2$$

where for $\Delta t = 5$ and using results from Lettenmaier [1984],

$$\rho_{\epsilon_t, \epsilon_{t-1}} = \rho(1 - \rho^5)^2 [5(1 - \rho)(1 + \rho) + \rho^4]^{-1} \quad (11)$$

System Data

Five-day flow increments were adopted for the experiments. Streamflow data for a 2-year period (October 1972 to September 1974; 146 5-day increments) from the Skykomish River, Washington (USGS gauge 12123400) were used. The average 5-day flow volume was 2.1 units, and the average annual flow μ was 153.3 units. The reservoir had a capacity (CAP) equal to 85 units; $S_{min} = 10$ units. The storage target was $T^s = 40$ units. The initial storage $S_0 = 40$ units; the beginning time was taken at the start of a water year. The release target was held constant at $T^r = 2.5$ flow units. This high target was chosen deliberately to stress the system heavily. (The mean

flow per 5-day period was 2.1 flow units, hence the annual target release to mean flow ratio was 1.19. This target cannot be met physically in the long-term; it is infeasible to satisfy this target at all times). The reservoir was relatively large ($CAP/\mu = 0.55$).

Other capacity to mean flow ratios could have been chosen, and results probably would differ from those we report below. We emphasize, however, that our primary concern is to show operation strategies for conflicting objectives in the presence of noisy forecasts.

Three models, models 1, 2, and 3, were explored. Each model used the loss functions shown in Figures 3a and 3b. (These loss functions were arbitrarily chosen. We recognize that construction of each function is a nontrivial task and that data for their construction are difficult to obtain and estimate. Fuzzy, time-varying functions rather than the crisp functions used here may be all that could be hoped for in practice.) The only difference between the models resulted from changes in k_1 and k_2 —the weights placed on the release and storage objectives, respectively. The models were (model number, k_1 , k_2): model 1, 0.5, 0.5; model 2, 0.8, 0.2; model 3, 0.2, 0.8.

The loss functions shown in Figures 3a and 3b were so constructed that for equal weights on storage and release deviations, one unit of storage deviation is penalized less than a unit of release deviation. However, at a particular level of deviation (9.5 units) the deviations from storage become more costly than equivalent deviations from the release target. This implies that at a certain point it becomes more important to meet the storage target than the release target. We chose this construction because meeting a release target is often more important than meeting a storage target. Although it is not evident from Figure 3a (because of the linear scale used), any nonzero deviation from the release target was penalized. This also prevented the LP algorithm from finding alternative optima.

The range of C_p values used was 0.1 to 0.9. The daily correlation coefficient was forced to assume values such that for a given unconditioned variance $\sigma_{\Delta t}^2$ (obtained from recorded data of length 2 years) the resulting coefficient of prediction (C_p) was as specified. Once ρ was computed for a given C_p , $\rho_{e_t, e_{t-1}}$ could be determined. For example, for $C_p = 0.1$, $\rho = 0.5737$; for $C_p = 0.9$, $\rho = 0.9785$. Also, for $C_p = 0.1$, $\rho_{e_t, e_{t-1}} = 0.0201$, and for $C_p = 0.9$, $\rho_{e_t, e_{t-1}} = 0.0940$.

Solution Procedure

For fixed k_1 and k_2 a value of C_p was chosen (between 0.1 and 0.9), and the specified model was solved as follows. A 5-day forecast was made and the optimization model solved. At the end of the 5-day period the actual flow (first 5 days of water year 1973) was used to compute the actual storage in the reservoir rather than the value that resulted from use of the forecasted 5-day volume. The storage penalty was recomputed with the updated actual storage. The process was repeated for each of the remaining 5-day time increments in the 2-year time period. The initial storage for each 5-day period was the updated actual storage at the end of the previous 5-day period.

The procedure was repeated until operation for 10 forecasted sequences, each of length 2 years, had been completed. We ran experiments with 50 sets to be sure that results were not overly sensitive to the number of scenarios examined. For low C_p , results for 50 sequences and 10 sequences differed by about 10%. For $C_p > 0.5$ there was little difference in results. Consequently, all results presented here are for 10 sequences

of length 146 5-day periods. (The computational expense was significant, hence our preference for small forecast sequence sample size.)

RESULTS

All significant results are given in graphical form, starting with Figure 4. The abscissa is either the coefficient of prediction C_p or the weight applied to the release objective k_1 . Data on the graphs are connected by straight lines for simplicity in illustrating trends; much greater computational effort than is appropriate would be needed to obtain data suitable for construction of continuous curves.

It is evident from Figure 4 that the actual losses incurred decrease substantially as the coefficient of prediction increases. Also, at various ranges of C_p one of the models seems to be more efficient in reducing the losses. Consideration of seasonal target values (rather than the fixed value for storage used here) may modify these results, but we expect that the relative conclusions will still be true. Figure 4 also shows that model 1 may not be a desirable alternative if the forecasts are not reliable. The objective mix presented by model 2 may not be a desirable alternative for $0.5 \leq C_p < 1.0$. (i.e., increasing accuracy tending to a perfect forecast); model 3 may be more efficient.

If the variances of release and storage are criteria for evaluation, other conclusions can be obtained from Figures 5 and 6. As seen from Figure 5, variations in the priorities given to




Fig. 4. Total losses incurred in operation.




Fig. 5. Variance of release volumes; each value is the average of the variances computed from 10 time series of releases, each series included 146 time increments.

meeting the storage and release targets have distinct effects on the variances of storage and release volumes for each period. It is seen from Figure 5 that as the priority for meeting the release target increases the variance of the release decreases. In general, model 2 results in minimum variance, and model 3 results in maximum variability. Similarly for the storage variability, as the priority for meeting the storage target increases, the variability of storage decreases; therefore model 3 shows the lowest variability. However, this trend is not evident when models 1 and 2 are compared. This is because, while the releases made depend completely on the operation policy specified by the optimization model once the initial storage is determined, the end storage for each period is determined, not by the storage that results from using forecasted inflow but by the actual flow that occurs. Therefore storage volumes can be expressed as

$$S_{t+1}(\text{actual}) = S_{t+1}(\text{forecasted}) + (\text{error term}) \quad (12)$$

This error term depends upon the forecast accuracy and determines the deviation between the storage predicted by the solution of the optimization model (knowing only the forecasted streamflow volume) and the end storage determined by the actual streamflow volume.

Another important point to be noted here is that these figures only indicate comparative trends. The absolute magnitudes shown in the figures depend on the storage target chosen and especially the release target used, as well as the shape of the loss functions. Additional trials with different target release volumes showed that, as the target release approached the mean value of streamflow during the time horizon of this evaluation, the variance of the releases decreased as anticipated.

RELIABILITY, RESILIENCY, AND VULNERABILITY OF OPERATION

In addition to the commonly accepted measure of reliability as a system performance criterion for evaluating an operation policy, other criteria, including resiliency and vulnerability of operation policies, have been suggested recently [Hashimoto et al., 1982a]. Reliability of a system performance is defined as the probability that a state of the system X_t is in a satisfactory state S :

$$\alpha = \text{Prob}[X_t \in S] \quad (13)$$

The resiliency of a system operation may be defined as the probability of a system's recovery to a satisfactory state S in time period $(t + 1)$, given that the system was in failure state F at time period t [Hashimoto et al., 1982a]. Therefore, resiliency of a system operation policy was defined as

$$\text{Probability}\{X_{t+1} \in S | X_t \in F\} \quad (14)$$

This is equivalent to the inverse of the expected value of the length of time the system is in a failure or unsatisfactory state. Thus it may be defined as




Fig. 6. Variance of storage volumes; each value is the average of the variances computed from 10 time series of storages, each series included 146 time increments.




Fig. 7. Release and storage reliability and resiliency as a function of k_1 ($C_p = 0.7$).

$$E \left[\frac{\text{number of sojourns into an satisfactory state } F}{\text{total number of periods in unsatisfactory state } F} \right] \quad (15)$$

where $E[\cdot]$ is the expectation operator.

Vulnerability may be defined as the expected severity of a failure when in a failure state F . This definition does not, however, take into account actual variability in this criterion.

These system performance criteria were evaluated for our dual-objective single reservoir operated under conditions of imperfect inflow forecasts. The failure state of our system was specified arbitrarily as that state when the absolute deviation of the actual release was greater than 0.2 times the target release or when the actual storage (corrected for forecast errors) at the end of a period had an absolute deviation greater than 0.125 times the storage target. Every time the release or storage volumes entered these failure states, a failure was counted for release or storage, respectively. The choice of the failure states is subjective; one choice may be no better than another. When all deviations from the release target are penalized, is it appropriate to consider all deviations from the target release as failures? We chose the failure states arbitrarily, reflecting the relative weights given to storage and release deviations. This choice illustrates one possible situation.

Figure 7 contains results obtained from simulation of actual operating conditions once a release had been made according to the solution of the optimization model. All these results are average values for two criteria (reliability and resiliency) obtained from 10 sets of generated forecasts for $C_p = 0.7$. This particular value of C_p was chosen as representative of forecasts that were neither exceptionally erratic nor accurate, and they reflect the state of practice. It is seen from Figure 7 that the resiliency of the operation policy specified by models 1, 2, and 3 differs for storage values and release values. While the absolute magnitude of this criterion depends on reservoir size

and flow variability, the relative values are of concern. These results show that specifying very high priority to meeting the storage or release target may actually decrease the resiliency of the policy. However, under these conditions, the reliability of satisfying the storage and release targets may be increased. When the model is required to meet the release target with a high priority, the reliability of attaining the release target is maximum, achieved at a cost of reducing the resiliency. A similar conclusion is true for the storage volumes.

It is also evident that the reliability or resiliency curves do not follow any trend that might have been inferred from results obtained from methods that used perfect forecasts. For perfect forecasts the reliability of meeting the storage or release target should increase as the priorities on meeting these targets increase. Figure 8 may explain the reason for such counterintuitive results. When the optimization model is solved, the model calculates an end storage based upon the forecasted streamflow volume. The correct storage at the end of the period is recomputed with the actual flow volume. The model has effected an optimal policy by using an uncertain flow forecast. Losses are corrected for the updated (true) storage before a solution is sought for the next time period. If the forecasted value of the end storage (storage at the end of the period when release for that period has been made) was the same as the actual value (the perfect forecast case), the reliability and resiliency curves would be as shown in Figure 8. Reliability and resiliency depend upon the priorities used in the model to satisfy the storage target. Due to forecast errors, however, these trends are modified as indicated in Figure 7. Figures 7 and 8 illustrate the relevance of forecasting errors and of prespecified priorities in a multiobjective model.

Vulnerability was defined as a measure of the severity of a failure when in a failure state. Hashimoto *et al.* [1982a] used the average of the maximum deficit that occurred in each mathematical run of failures within an operating time series as an index of vulnerability. In addition to the limitations of this measure that result from using a one-sided loss function (a




Fig. 8. Reliability and resiliency of anticipated storage volumes obtained from forecasted inflows for $C_p = 0.7$.

single objective, release, was considered; deficits were evaluated and excesses ignored) their approach has another potential drawback. Consider, for example, two alternative time series for release from a reservoir. The deficits for a common period of eight time increments are 0.0, 0.2, 0.6, 2.8, 0.5, 0.0, 0.6, 0.0 and 0.0, 0.2, 0.8, 2.6, 0.4, 0.0, 0.8, 0.0, respectively. The maximum deficit for each run of deficits is in italics. Both series yield vulnerability equal to 1.7 (series 1— $(2.8 + 0.6)/2$; series 2— $(2.6 + 0.8)/2$). However, the first sequence may actually be worse when, the loss function is convex, i.e., of the type (target-release) ^{β} with $\beta > 1$. This is not reflected in the above definition of vulnerability. A different picture may be obtained if the vulnerability is defined in terms of the losses incurred.

We defined vulnerability as the total losses incurred by a sequence of operations that follow optimal policy. If this definition is acceptable, Figure 4 will also denote the vulnerability of the system operation both in terms of storage and release deviations because the total losses were computed on the basis of deviations. As noted before, vulnerability now becomes a function of the forecasting accuracy C_p and weights k_1 and k_2 . Model 3 indicates a desirable combination of k_1 and k_2 for forecasts for most C_p values from the consideration of minimum vulnerability.

Robustness is described as a measure of overall economic performance. For a given positive constant, robustness may be defined as the probability that the ratio of economic regret to the least cost is less than or equal to that constant [Hashimoto *et al.*, 1982b]. For an operation policy the least cost or minimum penalties of operation may be considered as that constant. (Robustness is a more appropriate measure for evaluating a long-term policy. However, economic regret is certainly relevant to any decision based on economic considerations.) Therefore, economic regret, which may be defined as actual penalties minus minimum possible penalties, should be a useful index in evaluating on operation policy. Regret is defined here as the difference between actual total penalties and the possible minimum penalties as a function of forecast accuracy. This model finds a policy that results in minimum losses for an assumed perfect forecast for each forecast period. The actual loss incurred will be different and will depend on the difference between forecasted flow and actual flow. The ordinate in Figure 9 denotes the sum of deviations, averaged over 10 forecast scenarios, between forecasted losses predicted by the model (knowing only the forecasted values) and the actual losses (from observed inflow) that occur during each time period. It is clear from Figure 9 that, if the given definition of regret is accepted, the regret for the operation policy is a function of C_p . As forecasts become more accurate, i.e., the C_p values increase, regret also decreases. For high C_p values (~ 0.9) the results are relatively insensitive to the model adopted.

SUMMARY AND CONCLUSIONS

The aim of this work was to explore the sensitivity of various performance criteria for reservoir operation to the accuracy of forecasted streamflow volumes. For a single reservoir the sensitivity of these criteria to meeting conflicting storage and release targets was examined. When a tradeoff is made between incurring one unit of storage deviation and one unit of release deviation from respective target values, the compromise solution depends on uncertain future streamflow as well as the shapes of the loss functions. Therefore the relative shapes of the storage or release loss function should incor-




Fig. 9. Difference between the sum of actual losses and the sum of losses conditioned on forecasted flow.

porate economic information applicable to the operation. Loss functions for short-term operation represent opportunity costs, which depend on future hydrologic and economic uncertainties. It is unlikely that such functions can be constructed with objective information only.

Most of the results given are averaged over 10 sets of forecasted operation sequences corresponding to a single 2-year historical sequence. When the results for a single set are compared with those of the averaged values, the difference between the expected trends and the actual trends that might occur becomes obvious. Figure 10 shows typical variability of the actual losses for two values of C_p . It is clear that a single bad forecast in a series of forecasts that are quite good may offset all the benefits derived from using a fairly accurate forecast model. Therefore these results must be viewed with proper understanding of the uncertainties that play the most important role, especially in short-term operation. The following are the conclusions of this work:

1. The Standard Linear Operation Policy (SLOP) is not a reasonable operation policy when the objectives of operation are multiple and conflicting in nature.
2. If the objective of operation is maintenance of both storage and release targets, the relative shapes of the loss or penalty functions, as well as the prespecified priorities of meeting the storage or release targets, should reflect the risk averseness and the subjective judgment of the decision makers about economic and hydrologic uncertainties.
3. Use of short-term forecasts appears to be desirable for short-term reservoir operation whenever the forecasts are good enough so that the variance of the distribution of actual streamflows, conditioned on the forecasted values, is smaller than the unconditioned variance. The coefficient of prediction (C_p) criterion may be useful for determining the worth of a forecast in terms of information added. While these forecasts are actually surrogates for feedback information, the C_p values




Fig. 10. Typical actual cumulative losses with two different sets of synthetic forecasts for coefficient of prediction equal to 0.7 and 0.3.

should aid the decision maker's subjective judgment about the utility of forecasts.

4. The accuracies of inflow forecasts, made by a model external to the optimization model, influence the system performance in various ways. Evaluation of any policy without accounting properly for the probable errors in these forecasts may lead to misleading conclusions.

5. An operation policy is based on the assumptions embedded in the model used to develop the policy. However, the performance of this policy depends on the actual operating conditions of the system. While the performance of a policy can be judged with respect to various operational criteria, including reliability, resiliency, and vulnerability, it is not possible to correlate such criteria to the economic assumptions used in the model when future uncertainties are not incorporated explicitly into the model.

6. The variances of the storage and release volume depends on the magnitudes of k_1 and k_2 .

A very simple model was used here. Usually, it is necessary to consider forecasts several time steps ahead; the optimum time horizon to be chosen is site specific and is influenced by the reservoir size to inflow volume ratio. A more complicated model than used here can be constructed to make use of a time series of forecasted inflows.

We have demonstrated that an operation policy derived from a model that uses future predictions as deterministic inputs has severe drawbacks. It cannot be overemphasized that when dealing with future uncertainties it is unreasonable to propose any method that may be universally accepted. A catastrophic event may offset all expectations in terms of losses and benefits. Any short-term operation policy not capable of incorporating some judgement of experienced operators or decision makers explicitly should be discarded. This work should caution practitioners about some limits of methods

used to develop optimum short-term reservoir operation policy, given ubiquitous future hydrologic uncertainty.

Acknowledgment. Dennis Lettenmaier provided helpful comments at various stages of this work. His recent work and initial suggestions motivated us to examine the impact of flow forecast quality on reservoir operation. Vit Klemes, National Hydrology Research Institute, Canada, provided thoughtful comments on an earlier draft. Computer support, provided by the Department of Civil Engineering, University of Washington, is appreciated. Our work was supported in part by a grant from the Office of Water Research and Technology, U.S. Department of the Interior.

REFERENCES

- Cohon, J. L., and D. H. Marks, A review and evaluation of multi-objective programming techniques, *Water Resour. Res.*, 11(2), 208-220, 1975.
- Datta, B., Stochastic optimization models for long term planning and real time operation of reservoir systems, Ph.D dissertation, 192 pp., Purdue Univ., West Lafayette, Ind., 1981.
- Datta, B., and M. H. Houck, A stochastic optimization model for real-time operation of reservoirs using uncertain forecasts, *Water Resour. Res.*, in press, 1984.
- Fiering, M. B., *Streamflow Synthesis*, Harvard University Press, Cambridge, Mass., 1967.
- Fiering, M. B., A screening model to qualify resilience, *Water Resour. Res.*, 18(1), 27-32, 1982a.
- Fiering, M. B., Alternative indices of resilience, *Water Resour. Res.*, 18(1), 33-40, 1982b.
- Fiering, M. B., Estimates of resilience indices by simulation, *Water Resour. Res.*, 18(1), 41-50, 1982c.
- Fiering, M. B., Estimates of resilience by canonical analysis, *Water Resour. Res.*, 18(1), 51-57, 1982d.
- Hashimoto, T., J. R. Stedinger, and D. P. Loucks, Reliability, resiliency, and vulnerability criteria for water resource system performance, *Water Resour. Res.*, 18(1), 14-20, 1982a.
- Hashimoto, T., D. P. Loucks, and J. R. Stedinger, Robustness of water resources systems, *Water Resour. Res.*, 18(1), 21-26, 1982b.
- Houck, M. H., A chance-constrained optimization model for reservoir design and operation, *Water Resour. Res.*, 15(5), 1011-1016, 1979.
- Klemes, V., Value of information in reservoir optimization, *Water Resour. Res.*, 13(5), 837-850, 1977.
- Klemes, V., Reply to comment on 'Value of information on reservoir optimization', by J. R. Stedinger, *Water Resour. Res.*, 14(5), 987-990, 1978.
- Klemes, V., Storage mass curve analysis in a system-analytic perspective, *Water Resour. Res.*, 15(2), 359-370, 1979.
- Lettenmaier, D. P., Synthetic streamflow forecast generation, *J. Hydraul. Eng. Am. Soc. Civil. Eng.*, 110(3), 277-289, 1984.
- Loucks, D. P., and P. J. Dorfman, An evaluation of some linear decision rules in chance-constrained models for reservoir planning and operation, *Water Resour. Res.*, 11(6), 777-782, 1975.
- Nemhauser, G. L., *Introduction to Dynamic Programming*, chap. 5, pp. 149-183, John Wiley New York, 1966.
- ReVelle, C. S., E. Joeres, and W. Kirby, The linear decision rule in reservoir management and design, 1, Development of the stochastic model, *Water Resour. Res.*, 5(9), 767-777, 1969.
- Stedinger, J. R., Comment on 'Value of information in reservoir optimization', *Water Resour. Res.*, 14(5), 984-986, 1978.

B. Datta and S. J. Burgess, Department of Civil Engineering, University of Washington, Seattle, WA 98195.

(Received August 4, 1983;
revised March 5, 1984;
accepted April 26, 1984.)