Journal of Water Resources Planning and Management, Vol. 125, No. 3, pp. 143-153, May/June 1999

Some Derived Operating Rules for Reservoirs in Series or in Parallel

Jay R. Lund, Professor
Department of Civil and Environmental Engineering
University of California, Davis, CA 95616; jrlund@ucdavis.edu

Joel Guzman
former undergraduate student, Department of Civil and Environmental Engineering,
University of California, Davis

Abstract

This paper reviews a variety of derived single-purpose operating policies for reservoirs in series and in parallel for water supply, flood control, hydropower, water quality, and recreation. Such rules are useful for real-time operations, conducting reservoir simulation studies for real-time, seasonal, and long-term operations, and for understanding the workings of multi-reservoir systems. For reservoirs in series, several additional new policies are derived for special cases of optimal short-term operation for hydropower production and energy storage. For reservoirs in parallel, additional new special-case rules are derived for water quality, water supply, and hydropower production. New operating policies also are derived for reservoir recreation.

Introduction

Despite the development and growing use of optimization models (Labadie 1997), the vast majority of reservoir planning and operation studies are based predominantly or exclusively on simulation modeling, and thus require intelligent specification of operating rules. Practical real-time operations also usually require the specification of reservoir operating rules. These rules determine the release and storage decisions for each reservoir at each time-step during the simulation and help guide reservoir operators (Bower et al. 1966; Hufschmidt and Fiering 1966). This paper reviews a variety of common and new derived operating rules for single-purpose reservoirs in series and in parallel. These derived rules can all be supported by conceptual or mathematical deduction from principles of engineering optimization for special cases. These rules can be contrasted with the many and often highly effective empirically-based rules common in practice, such as various pool-based rules and balancing rules (Wurbs 1996; USACE 1977; Nalbantis and Koutsoyiannia 1997). The rules examined here are intended mostly for seasonal and long-term studies. Real-time studies, with an hourly or daily time-step, often have more detailed safety, habitat, and facility limitations not usually important for studies using coarser time-steps or longer operating horizons. The rules presented here offer some guidance for real-time operation, but are more generally applicable to seasonal and long-term operations planning and modeling studies. Desireable operating rules for reservoir systems with mixed purposes might have very different forms from those presented here. Previous reviews of reservoir operating rules include Sheer's fine concise review (1986) and Loucks and Sigvaldason (1982). Several Corps of Engineers reports (cited below) also develop reservoir operating rules for specific purposes. The work presented here is drawn from work done for the U.S. Army Corps of Engineers (Lund and Guzman 1996).

For those developing operating policies for real reservoir systems, the rules presented are part of a bag of tricks with some theoretical or practical basis to recommend them. The rules discussed here are organized by reservoir configuration, in series or in parallel, and by various operating purposes. Operation for multiple purposes, i.e., most real systems, requires some combination of these rules or use of other operating rule forms (Lund and Guzman 1996). The presumption in many of these rules is that a system of reservoirs can be operated to produce greater benefits than operating the individual reservoirs independently. While this is often the case (Palmer et al 1982), it is not always the case (Needham 1998).

The paper begins with examination of rules for reservoirs in series for specific single purposes. Various single-purpose rules for reservoirs in parallel are then reviewed. A general storage-effectiveness-based rule for recreation storage is then presented, followed by a short general discussion of specification of operating rules by stochastic dynamic programming. The conceptual rules are summarized in Tables 1 and 2. A commentary and conclusions end the paper. Derivations of selected rules appear in appendices. Pseudo-code for implementing many of these rules appear in Lund and Guzman (1996).

Rules for Reservoirs in Series

The paper begins with examination of operating rules for reservoirs in series, illustrated in Figure 1, for water supply storage, flood control, energy storage, and hydropower production. Most rules presented are assembled from earlier cited work. This work is extended in some cases. The conceptual rules for reservoirs in series are summarized in Table 1.

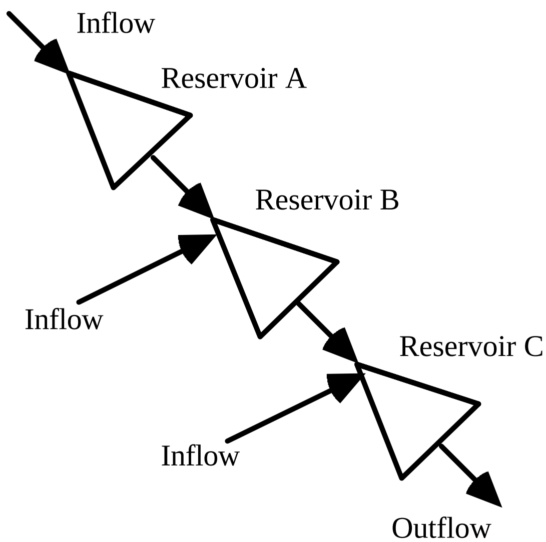

Figure 1: Reservoirs in Series

Table 1: Conceptual Rules for Reservoirs in Series*

	Season/Period	
Purpose	**Refill**	**Drawdown**
Water Supply	Fill upper reservoirs first	Empty lower reservoirs first
Flood Control	Fill upper reservoirs first	Empty lower reservoirs first
Energy Storage	Fill upper reservoirs first	Empty lower reservoirs first
Hydropower Production	Maximize storage in reservoirs with greatest energy production	
Recreation	Equalize marginal recreation improvement of additional storage among reservoirs	

* Exceptions and refinements are discussed in the text.

Water Storage Rules

For reservoirs in series providing water supply, a reasonable objective is to maximize the amount of water available, which is the same as minimizing spilled water. The resulting rule for single-purpose water supply reservoirs in series is simply to fill the higher reservoirs first, and the lowest last (Sheer, 1986).

The likelihood and severity of shortages is reduced by preventing any water from leaving the system as uncontrolled and unproductive spills. For reservoirs in series, with intermediate inflows, the probability of spill from the system is minimized by first filling the uppermost reservoirs and retaining storage capacity in the lower reservoirs to capture potentially large flows and reduce the likelihood of spills from the system. Spillage from any but the lowest reservoir can then be captured by a lower reservoir. During the draw-down season, where system inflows are less than demands, the system should be drawn down in order of the downstream reservoirs first, to provide storage to accommodate potentially excess intermediate inflows or an early onset of the refill season.

For reservoirs in series serving water supply as a sole purpose, the above rule seems universal. An exception might be where higher reservoirs suffer higher rates of water loss from evaporation and seepage (Kelley 1986). In this case, any increased evaporation or seepage from higher reservoirs would have to be weighed against the increased potential for loss due to spill from concentrating storage at lower elevations.

Flood Control Rules

For reservoirs in series with intermediate inflows and storage serving solely for flood control downstream, it is optimal to regulate floods by filling the upper reservoirs first and emptying the lower reservoirs first. The objective is to maintain as much control over flows entering the system above a critical flood-prone reach as possible. Flood storage at the reservoir closest upstream from a critical flood control reach always provides greater flood controllability than for any other reservoir (Mariën et al. 1994). These are typically the lowest reservoirs in the series. Thus, for single-purpose flood control storage in a series of reservoirs, it is best to fill the higher reservoirs first and empty the lower reservoirs first. An illustration of this rule can be found in a large Brazilian multi-reservoir system (Kelman et al. 1989).

An exception to this rule can be where the outflow capacity of the lower reservoir is restricted. Here, it can be better to fill the lower reservoir first to increase head on the outlet, thereby increasing release capacity from the entire system to the downstream channel capacity (USACE 1976). This allows higher pre-releases to increase total storage available for a coming flood event.

The operation of reservoirs in series for flood control is fairly complementary with water supply operations, at least in regard to the preferred location of storage. The maintenance of water supply storage still preys on the absolute flood control capability of a system, and vice versa.

USACE (1976) presents methods for allocating flood control space in reservoirs in series with flood control locations both downstream of the system and between reservoirs. In such cases, except for small flood events, there are likely to be inherent trade-offs in protection of these sites from different storage allocation and operations decisions.

Hydropower Rules

Hydropower rules for reservoirs in series vary between refill and drawdown seasons or periods. During a refill period, the problem usually is to maximize the storage of energy at the end of the period. During a drawdown period, the objective is to maximize hydropower production for a given total storage amount. Different rules are employed for each period. A difficult problem is the transition between seasons or periods.

Energy Storage Rules

The objective of the energy storage rule for reservoirs in series is to maximize the total energy stored at the end of a refill season or period. Here, the refill season is defined as the season when system inflows exceed those needed to meet water supply or hydropower production demands. The energy storage rule for reservoirs in series is to always fill the upper reservoirs first.

To maximize the energy stored for a future time, water storage is preferred in upstream reservoirs. Water stored at higher elevations has a higher energy content (kilowatt-hours/unit volume of water stored) than water stored at lower elevations. This is particularly true for water stored in reservoirs in series, where water eventually released from upper reservoirs generates hydropower at the lower reservoirs as well. Any spills from upper reservoirs are available for capture in space available in lower reservoirs. Kelman, et al. (1989) mathematically examine the allocation of energy storage and flood control storage capacity in complex multi-reservoir systems. Their results will often indicate the compatibility of the desirable distribution of energy and flood control storages in such two-purpose systems.

Fortunately, energy storage and water supply storage rules for reservoirs in series are quite compatible for the refill season, at least in terms of where storage is preferred in the system and their general intent to accumulate the maximum amount of water. However, with the coming of the drawdown season, hydropower production rules are required.

Hydropower Production Rule

When it comes time to produce energy, rules to maximize hydropower production may be employed. Upper reservoirs generate hydropower by releases which consequently also increase downstream power generation by increasing heads, if stored downstream, or by subsequent turbine releases downstream.

The steady-state hydropower production rule attempts to maximize hydropower generation during a single time step, given a total storage target for the system, primarily during the drawdown season. This problem involves allocating a given total storage to maximize hydropower production. In general, the rule favors allocation of storage to those reservoirs which create a higher head per unit volume of storage, have higher generation efficiencies, and have higher releases, since hydropower production is the product of head, efficiency, and release. A variant of this rule is the Corps' "storage effectiveness index" presented in Appendix I (USACE 1985; Lund and Guzman 1996). This rule typically applies to a drawdown period or season.

The maximum amount of power that can be generated in a reservoir system occurs when head levels in all reservoirs are at their highest. Where the total amount of water stored in the system is limited, perhaps due to other operating purposes or variations in hydrology and demands, then the problem becomes one of allocating a limited amount of storage among the individual reservoirs to maximize hydropower production. This hydropower maximizing water storage allocation depends on reservoir capacities, inflows, efficiencies of energy production, and the total amount of water (or energy) to be stored.

Water often is stored first in smaller reservoirs, where head usually increases more per unit volume of additional storage than in most large reservoirs. This is illustrated in Figure 2, where the volume of water needed to increase head by an amount x in the lower reservoir, V2, is much less than that needed to achieve an equivalent increase in head for the upper reservoir, V1.

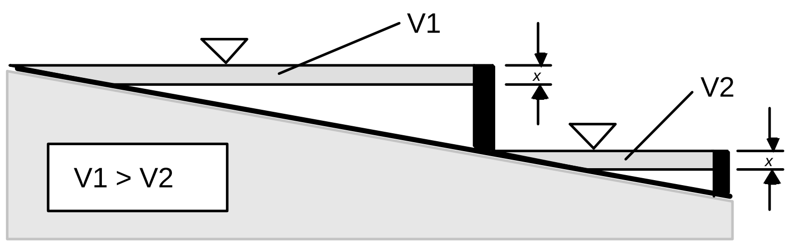

Figure 2: Change in Head With Varying Capacities

Another consideration is release flow rate. All else being equal, hydropower production is maximized by allocating available stored water to reservoirs with the greatest release rates. Thus, for reservoirs in series with intermediate inflows, it is often desirable to maximize storage in the lower reservoirs first. Typically, downstream reservoirs receive more direct and indirect inflows than upstream reservoirs. Storage in downstream reservoirs is therefore often kept at high levels to take advantage of increased flows (with some increased chance of energy spills).

Lastly reservoirs with higher generation efficiencies should be maintained at higher levels of storage at the expense of reservoirs with lower efficiencies. The combination of reservoir capacity, amount of total inflows, and power generation efficiencies determines the overall potential of the reservoir to produce power, assuming all reservoirs have adequate turbine capacity. When reductions in storage are necessary, they are made from reservoirs with the least ability to produce power. Conversely, if an increase in storage can be made, it should be in reservoirs with the greatest ability to produce power.

The interaction of these factors is examined mathematically in Appendix II. The result is to calculate the following ratio for each reservoir i at each simulation time-step:

$$V_i = a_i e_i \left(\sum_{j=1}^{i} I_j \right), \quad (1)$$

where

- V_i increased power production per unit increase in storage
- a_i the unit change in hydropower head per unit change in storage (the slope of the head-storage curve),
- e_i the power generation efficiency of reservoir i, and
- I_j the direct inflows and releases into reservoir j, for all reservoirs upstream of reservoir i.

Here reservoir 1 is the uppermost reservoir in the series of reservoirs. For this steady-state hydropower production rule, reservoirs are ordered in terms of their values of V_i, and are filled from highest to lowest value of V_i until the total water storage target is met.

For drawdown periods, these rules can be employed for hydropower production systems of reservoirs in parallel, in series, as well as mixed systems. Where the total storage constraint is desired to be in terms of energy storage, rather than water storage, then the more elaborate linear programming approach presented in Appendix II is required.

Many systems of reservoirs in series maintain their lower, smaller reservoirs full for hydropower production. This is the case for the Missouri River System and the Columbia River System. Often, maintaining high storage levels in the lower reservoirs also aids navigation through the lower reaches of the reservoir system, as in the lower Columbia River System.

The hydropower production rule for reservoirs in series can be implemented directly using the values of V_i presented above, or by using the more elaborate linear programming formulations of the problem presented in Appendix II. This rule should work best for individual time-steps under nearly steady-state conditions, where small changes in storage are anticipated.

A more dynamic and general version of this simple rule is discussed by Lund (submitted), based on a concept by Dan Sheer (1986).

Rules for Reservoirs in Parallel

The operation of reservoirs in parallel, Figure 3, differs from reservoirs in series in that downstream reservoirs cannot be used to capture additional water from underestimated flows or benefit from the transfer of water stored upstream if flows are overestimated. The following sections present balancing rules for water supply, energy storage, water quality, and flood control. These rules typically apply to the reservoir system's refill season. Computational studies suggest that these rules tend to work rather well over a wide range of conditions, perhaps because the performance response surface is flat for such storage allocation decisions (Sand 1984). Several hydropower production rules also are presented. Conceptual rules for operating parallel reservoirs are summarized in Table 2.

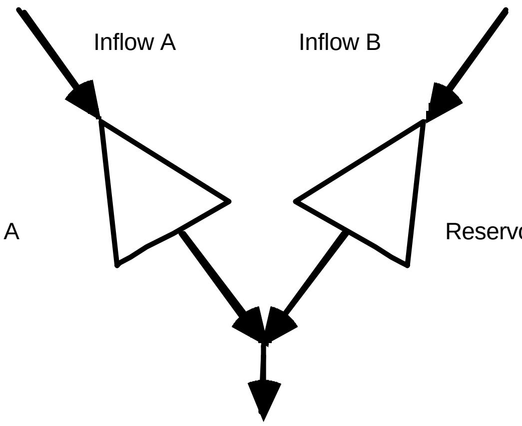

Figure 3: Reservoirs in Parallel

Table 2: Conceptual Rules for Reservoirs in Parallel*

	Season/Period	
Purpose	**Refill**	**Drawdown**
Water Supply	Equalize probability of seasonal spill among reservoirs	Equalize probability of emptying among reservoirs
Flood Control	Leave more storage space in reservoirs subject to flooding	N.A.
Energy Storage	Equalize expected value (EV) of seasonal energy spill among reservoirs	For last time-step, equalize expected value (EV) of refill season energy spill among reservoirs
Water Quality	Equalize EV of marginal seasonal water quality spill among reservoirs	For last time-step, equalize EV of refill season water quality spill among reservoirs
Hydropower Production	Maximize storage in reservoirs with greatest energy production	
Recreation	Equalize marginal recreation improvement of additional storage among reservoirs	

* Exceptions and refinements are discussed in the text.

Water Supply, Energy Storage, and Water Quality Rules

Several types of rules have been developed for water supply, energy storage, and water quality operations for parallel reservoirs during refill periods. For reservoir systems providing either water supply or energy production, a reasonable objective is to minimize expected shortages. The severity of shortages is reduced by avoiding any water leaving the system as uncontrolled and unproductive spills (Sand 1984). These rules prescribe ideal release or storage levels for reservoirs in parallel to avoid the inefficient condition of having some reservoirs full and spilling, while other reservoirs have unused storage capacity (Bower et al. 1966). These rules are derived in Appendix III.

New York City Rules (NYC rules)

The NYC rules use the probability of spills rather than the direct amounts of physical spill in the minimization of expected shortages. When the probabilities of spilling at the end of the refill season are the same for each reservoir, it follows that physical spill also is minimized (See Appendix III). Water supply shortfall is consequently minimized as well.

The NYC rule was first stated by Clark (1950) for the New York City water supply system, "In operating this system an attempt is made to have the storage in each of the watersheds, at all times, fall on the same percentage year." The draw from each reservoir is adjusted to equalize the probability of refill by the end of the refill season, about June 1 (Clark 1956).

There are three requirements for rigorous application of the NYC rule (Sand 1984):

1. The system contains reservoirs operating in parallel;
2. The system provides for a single demand downstream of all reservoirs;
3. Expected shortages are to be avoided or minimized.

Modified forms of the NYC rule also can handle situations where the unit value of water varies between reservoirs but is constant in any individual reservoir. In terms of water supply, the quality of water might affect its unit value. For example, higher total suspended solids (TSS) concentrations correspond to greater treatment costs and lower desirability as a water supply source. For energy production the value of water is based on the maximum head of each reservoir, so water contained in reservoirs with greater head has proportionally greater value.

Application of the NYC rule depends on predicted inflows. Thus greater accuracy in these predictions should yield better results (fewer spills). Since releases are recalculated at each period, a high degree of accuracy in predicted inflows is not critical in the early periods of the refill season. Towards the end of the refill season, reliable flow forecasts become more important as the chances of spill increase (Bower et al. 1966). Therefore, it is important to have enough historical and watershed data for probabilistic streamflow forecasts.

Optimality also depends on the coefficient of variation of mean monthly flows and the correlation between flows on adjacent streams (Bower et al. 1966; Sand 1984). The NYC rule has been found to behave optimally or near-optimally for a wide variety of operating conditions and system configurations (Sand 1984).

The general form of the NYC rule equates the probabilities of spill at the end of the refill season adjusted by the unit value of water for each reservoir:

(2) $$h_i \Pr[CQ_i \geq K_i - S_{fi}] = l \quad \text{for all } i \quad \text{where}$$

- h_i unit value of water in reservoir i,
- CQ_i cumulative inflow to reservoir i from the end of the current period to the end of the refill season,
- K_i storage capacity of reservoir i, assumed to be the same in every period,
- S_{fi} end-of-period storage for the current period for reservoir i, and
- l a constant across all reservoirs in parallel.

The values of h_i depend on water quality and energy storage issues as described later. Historical data are typically used to establish the cumulative inflows, CQ_i. Individual reservoir releases for the current period, R_i, are found by taking the initial storage, S_{oi}, plus expected inflow for the current period, $E[Q_i]$, and subtracting the end-of-period storage, S_{fi}, that satisfies equation (2):

(3) $$R_i = S_{oi} + E[Q_i] - S_{fi}$$

such that the sum of releases equals a current-period total downstream release target R_{OT},

(4) $$\sum_{i=1}^{n} R_i = R_{OT}.$$

This often requires a search over λ to find the proper total release.

Water Supply

When the unit value of water, h_i, is the same among reservoirs providing water supply, h_i is incorporated into the constant λ and thus drops from the equation:

(5) $$\Pr[CQ_i \quad K_i - S_{fi}] = l \text{ for all } i$$

Water Quality

When the quality of water varies between reservoirs, such as varying total dissolved or suspended solids concentrations the probabilities of spill are adjusted by h_i, the marginal value of water use minus its treatment cost for each reservoir. Thus, if the marginal value of treated water use downstream is \$500/ac-ft and the cost of treatment for water from Reservoir 1 is \$50/ac-ft, h_1 = \$450. Such a formulation also might be applied to provide a balancing rule that preferentially stores water for downstream fish flows.

Energy Storage

For energy storage applications, the probabilities of the potential energy of spill are equated. Therefore the probabilities of spill are adjusted by h_i in Equation 2, the full head level of each reservoir, representing the marginal value of energy lost from spill.

The NYC space rules apply to the refill season of systems of parallel reservoirs and attempt to minimize the expected value of spilled water. The primary difficulties are specification of inflow probabilities, computational implementation of the rule (now a minor problem), and potentially the absence of considering future refill season demands on the inflows into the system.

Space Rule

The space rule seeks to leave more space in reservoirs where greater inflows are expected, or where greater potential energy of inflows are expected in the case of energy storage (Bower et al. 1966). It is a special case of the NYC rule, seeking to minimize the total volume of spills. The same conditions for applicability and optimality apply. The NYC rule becomes the space rule when the distributional forms of inflows into each parallel reservoir are the same, with distributions scaled by their expected value, i.e., the distribution $f_i(CQ_i/EV(CQ_i))$, where EV() is the expected value operator, is identical for all reservoirs (Sand 1984). This derivation appears in Appendix III. The advantage of the space rule over the NYC rule is its direct computation. Like the NYC rule, the spill minimizing objective implies that this rule is applicable to the system's refill season. Like the NYC rule, the space rule has been found to behave optimally or near-optimally for a wide variety of operating conditions and system configurations (Sand 1984; Wu 1988).

Johnson et al. (1991) examined the application of space rules for operating the Central Valley Project in Northern California. In this system, power output is maximized while maintaining high levels of water supply reliability. This application appeared to offer improvements over simulated operation of the system.

The particular form of the equal ratio space rule depends on the reservoir purpose being examined for the system. Space rules have been developed for water supply storage and energy storage purposes.

Water Supply

For water supply purposes, implementing the space rule consists of setting target storages in each reservoir so that the ratio of space remaining at the end of the current period to the expected value of remaining refill season inflow for each reservoir is identical (Johnson et al 1991). This is expressed mathematically as,

$$\frac{K_i - S_{fi}}{EV(CQ_i)} = \frac{\sum_{i=1}^{n} K_i - V}{\sum_{i=1}^{n} EV(CQ_i)}, \forall i, \tag{6}$$

with

$$V = \sum_{i=1}^{n} (S_{oi} + EV(Q_i)) - R_{OT} \tag{7}$$

where V = the total water storage of the system at the end of the current time-step and all other terms are as defined in the NYC rule. Using the above equation, releases in a parallel system of reservoirs containing equally valued units of water are determined similarly to implementing the NYC rule,

$$S_{fi} = K_i - \left(\frac{\sum_{i=1}^{n} K_i - V}{\sum_{i=1}^{n} EV(CQ_i)} \right) EV(CQ_i). \tag{8}$$

Energy Storage

When preventing energy spills, the space rule equation used for water supply is modified by replacing reservoir capacities, available storage, and expected inflows with their potential energy counterparts. These consist of maximum potential energy that can be stored or the capacity of the turbines, available potential energy storage, and potential energy of expected inflows, respectively (Johnson et al. 1991). The substitution of these elements yields the following equation:

$$\frac{KE_i - E_{fi}}{EV(CE_i)} = \frac{\sum_{i=1}^{n} KE_i - E}{\sum_{i=1}^{n} EV(CE_i)}, \forall i, \tag{9}$$

where

KE_i the maximum energy content of reservoir i,
E_{fi} the target energy content of reservoir i at the end of the current time step,

CE_i the cumulative energy inflow to reservoir i,

E the total target energy content of the reservoir system at the end of the current time step, and

EV() the expected value operator.

If releases or storages fall outside their permitted upper or lower bounds, the decision variables can be set to those bounds while the remaining variables are balanced according to the space rule (Stedinger et al. 1983). Johnson et al. (1991) use a quadratic program to implement their energy space rule for each time step for their California Central Valley Project simulation model.

Space rules are simpler to implement than NYC rules. However, they rest upon distributional assumptions which might not always hold. The importance of these distributional assumptions can be tested for particular situations using long-term simulation modeling.

Linear Program NYC Rule

Like the two previous balancing rule forms, the linear program (LP) NYC rule also relies on historical data to determine expected inflows. However, rather than producing cumulative inflow distributions for CI_i, the inflow data are entered directly into a linear program. Spill, or the value of spill, is minimized by considering all individual cumulative inflows from each period to the end of the refill cycle in past years. Compared to the NYC rules, the primary advantage of LP-NYC rules is its ability to incorporate other (linear) short-term reservoir operation constraints into the rule. Such additional constraints might include minimum or maximum flows downstream of each reservoir or required diversions below a subset of reservoirs.

The same conditions required for the NYC rule apply to the LP-NYC rule. LP-NYC rules are slightly more general than NYC rules in that they also can incorporate other linear operating constraints in the setting of short-term storage targets for each reservoir and could also incorporate other operating purposes in the objective function (Johnson et al. 1991). However, implementation of LP-NYC rules require greater computational effort. For each time-step, the linear program described in Appendix III would be solved. The values of h_i in the formulation depend on whether water supply, water quality, or energy storage issues are being considered. Thus, water supply, water quality, and energy storage versions of the LP-NYC rule can be developed.

Flood Control Balancing Rule

The approach taken for flood control in parallel reservoirs is to maintain a balance between reservoirs in terms of occupied capacities and flood runoff from drainage areas. If a reduction in outflows is required, it is made from the reservoir with the least percentage occupancy or smallest flood runoff. When an increase in releases is possible, it is made from the reservoir with the greatest capacity occupied or where relatively higher flood runoff is occurring. Higher releases from reservoirs receiving greater flood-runoff may thus be counterbalanced by reducing releases from reservoirs receiving lesser runoff (Ghosh 1986).

The intent of the flood control balancing rule is to operate the parallel reservoirs to balance the amount of flood control storage available, while maximizing undamaging releases from the system. While the principle of balancing flood control storage on parallel reservoirs should be clear, operation to meet this objective is not exact. If the objective were to minimize the expected value of damaging spills above the downstream channel capacity, then flood control rules could be developed analogous to the New York City rules. Unfortunately, the objective of flood control is more likely to be minimization of peak downstream flood flows during the refill season, where peak inflows to the system can arrive during a very short time. This situation is less rigorously represented by the NYC rule approach.

USACE (1976) suggests the following method for allocating flood control space between two parallel reservoirs (A and B) with a single downstream flood damage location.

1) Route the project design flood and other observed floods with a maximum amount of runoff occurring above reservoir A and with maximum non-damaging releases from reservoir A. Allow reservoir B to make the remaining releases, up to the maximum non-damaging level. Plot the space required at reservoir A versus total space required.

2) Perform the same exercise for reservoir B, with the maximum design and observed flood flows entering above reservoir B. Plot the space required at reservoir B versus total space required.

3) The ratio for balancing flood storage between the two reservoirs should lie between these two curves.

While the concept of a flood control balancing rule appears conceptually sound, exact formulations remain unclear.

Water Supply Drawdown

For drawdown of water supply reservoirs in parallel, Wu (1988) suggests a rule which equalizes the probability of each reservoir being empty at the end of the drawdown season. For systems of reservoirs in parallel with side demands, depending on releases from a specific individual reservoir, such operation is meant to avoid the possibility of having to short a side demand when water is available to meet other demands. Wu fashions a drawdown rule along these lines similar to the space rule. In simulation tests, he finds a combination of space rule and this drawdown rule to be simple to implement and providing near-optimal performance relative to other rule forms.

Hydropower Production Rules for Reservoirs in Parallel

Steady-State Hydropower Production Rule

For steady-state hydropower production at reservoirs in parallel, the hydropower production rule for small time-steps derived in Appendix II reduces to:

(16) $$\text{Max } P = \gamma \sum_{j=1}^{m} e_j a_j S_j I_j.$$

where the subscript j refers to an individual parallel reservoir, and variables are defined as for Equation 1. Taking the first derivative, $\partial P/\partial S_j = e_j a_j I_j = V_j$, or the hydropower storage effectiveness of parallel reservoir j. The resulting rule is: Empty parallel reservoirs sequentially, beginning with those with the smallest V_j. Fill in the reverse order.

Power Production and Energy Drawdown Rules for Parallel Reservoirs

Sheer (1986) provides a rule for drawing down reservoirs with the minimum impact on long-term hydropower production (e.g., lost potential energy) for reservoirs which will fill before they empty. Reservoirs for which withdrawal results in the smallest reduction in potential energy should be drawn down first. This rule is derived and elaborated by Lund (submitted). Some of these results for parallel reservoirs appear below.

Assuming the reservoirs do not empty before they refill, the most efficient drawdown of water volume from the system would minimize total reduction in annual hydropower production. The marginal economic value of hydropower release is calculated for each reservoir,

$$(17) \qquad \partial z/\partial T_i = e_i \left(P_0 H_i(S_{i0}) + P_r R_i \; \partial \bar{H}_i(S_{i0})/\partial T_i - P_f H_i(K_i)\alpha_i \right).$$

where,

e_i is the efficiency of power generation at reservoir i,

T_i is the volume of storage to be released from reservoir i in the present time-step,

$\bar{H}_i(S_{i0})$ is the expected flow-weighted hydropower head for reservoir i until refill with present reservoir storage S_{i0},

$H_i(K_i)$ is the hydropower head with storage at refill capacity K_i,

P_0 is the present price of energy,

P_r is the flow-weighted average price of energy expected until the reservoir refills,

P_f is the expected price of energy when the reservoir is filled, and

α_i is the marginal proportion of additional storage in the present which would not be spilled during the refill season for reservoir i (if = 0.9, 10% of any additional storage now is expected to be spilled).

The rule then is to draw down reservoirs with the greatest values of $\partial z/\partial T_i$ first, and refill them in the reverse order. (Note that the last two terms are negative and a positive $\partial z/\partial T_i$ indicates increased hydropower value with increased release.) This rule is particularly applicable where the withdrawals are being made to supply some downstream water supply volume requirement.

If the current drawdown is intended to supply an energy demand or contract, then the above rule is modified somewhat. The most efficient drawdown of potential energy from the system would minimize total reduction in annual hydropower production. The marginal economic value of energy release from each reservoir is estimated, $\partial z/\partial E_i$, where $E_i = H_i(S_{i0})e_iT_i$. This leads to equation 18:

$$(18) \qquad \partial z/\partial E_i = P_0 + \frac{P_r R_i}{H_i(S_{i0})} \frac{\partial \bar{H}_i(S_{i0})}{\partial T_i} - \frac{P_f H_i(K_i)}{H_i(S_{i0})} \alpha_i.$$

The rule then is to draw down reservoirs with the greatest values of $\partial z/\partial E_i$ first, and refill them in the reverse order.

Both hydropower production rules should apply well where the reservoirs refill in most years and do not empty. Under these circumstances, energy spills might be common unless sufficient turbine flow capacity exists to pass common high refill-season flows. Thus, the coefficient α_i can be important. Alternative rules are developed for systems where reservoirs are expected to empty before they refill (Lund, submitted). Since the value of hydropower production often varies seasonally, these formulations also allow consideration of relative energy prices in different periods.

Storage Allocation for Reservoir Recreation

Where reservoir recreation is the predominant purpose for operations during a time-step, how should a given total storage be allocated among reservoirs? For individual reservoirs, recreation potential usually varies discontinuously around storage levels corresponding to the elevations of docks, boat-ramps, and beaches. However, within these ranges reservoir recreation potential is likely to be roughly proportional to reservoir surface area (A_i), perhaps weighted by some constant representing accessibility or recreational facilities (r_i) at each reservoir. Thus, the systemwide reservoir recreation objective is to maximize total weighted surface area over all n reservoirs:

(19) $$\text{Max } A_T = \sum_{i=1}^{n} r_i \, A_i(S_i)$$

Subject to:

(20) $$\sum_{i=1}^{n} S_i = S$$

Solving this problem using Lagrange multipliers gives the condition that for any two reservoirs i and j,

(21) $$r_i \quad A_i(S_i)/ \quad S_i = r_j \quad A_j(S_j)/ \quad S_j,$$

or, each reservoir's marginal storage contribution to recreation should be equal. This becomes a storage allocation rule for an arbitrary system of reservoirs predominantly used for reservoir recreation. Since reservoir surface area is usually a concave function of storage, the rule usually should roughly balance storage among all reservoirs, skewed by the weighting factor r_i. This type of storage effectiveness rule is likely to be of greatest use during the drawdown season, when storage is decreasing and the distribution of remaining storage becomes an important recreational issue.

During the drawdown season, operations for hydropower and recreation might be made more compatible by varying the development of reservoir recreation facilities and access. This would vary r_i to make the outcomes of hydropower and recreation rules more closely agree.

Stochastic Dynamic Programming-Based Rules

Rules for operation of multi-reservoir systems also can be developed by stochastic dynamic programming (SDP). Here, an explicit characterization of streamflow probabilities is used together with an explicit loss function and definition of system configuration and constraints to numerically derive optimal reservoir operating policies. This approach has long been explored and developed (Little 1955; Tejada-Guibert et al. 1995; Kim and Palmer 1997). Various formulations have been developed to include probabilistic streamflow forecasts and improvements and approximations to improve computational speed.

SDP approaches to deriving operating rules are rigorous and conceptually flexible. However, they suffer from extreme computational demands for large problems and require explicit probabilistic characterization of unimpaired streamflows. Attempts to reduce computational requirements by increasing the coarseness of storage and flow discretizations lead to results being more approximate (Klemes 1977). In most cases, it also is difficult to have confidence in a specific explicit probabilistic characterization of a region's inflows and there seem to be a limited variety of approches available for representing streamflow relationships within a SDP format. SDP-based rules continue to show increasing promise, but remain somewhat to prohibitively difficult to apply in practice.

A few tests of SDP versus other operating rules have been performed (Johnson et al 1991; Wu 1988; Karamouz and Houck 1987). These tests show that well-crafted conventional operating rules typically perform nearly as well, and sometimes better, than SDP-based rules.

Commentary

Refill versus Drawdown Periods

The rules presented here apply mostly to either periods of refill or drawdown. NYC and Space rules for reservoirs in parallel and water supply, energy storage, and flood control rules for reservoirs in series typically apply to refill seasons. Hydropower production rules apply mostly to drawdown periods. Within each type of period, it is relatively easy to determine desirable single-purpose operating policies for the system. It is far more difficult to determine operationally when these periods begin and end, and thus how operations should make the transition between these periods. Rules for determining which season applies can be based on time of year or streamflow forecasts, and evaluated using historical probabilities, explicit optimization, implicit optimization, and/or historically-based simulation studies. Forecasting is likely to be important here.

Multi-Purpose and Generic Multi-Reservoir Rules

Most reservoir systems are multi-purpose. Fortunately, many of the operating rules presented here show a compatibility between different reservoir purposes, such as flood control, water supply, and energy storage for refill periods on reservoirs in series. In complex cases, optimization approaches and traditional simulation may be used to calibrate the rule forms identified above or parameterized simplifications of such multi-reservoir rules. A wide variety of parameterized forms of these and other multi-reservoir operating rules are available, including the Corps' highly flexible "index-level method" for allocation of total storage among multiple reservoirs (USACE 1977; Nabantis and Koutsoyiannis 1997).

Traditionally, iterative simulation methods are used to calibrate operating rules to perform well for multiple purposes and remain the final analysis method used to refine and test operating rules (UASCE 1977). Performance-based optimization models also can help sort through the operations of complex multi-purpose systems. Where implicit stochastic optimization is used for computational convenience (Lund and Ferreira 1996), some of the rule forms suggested here might be useful for disentangling the results. Some of these rule forms also may be useful for other simulation-based reservoir optimization studies, such as genetic algorithms (Oliveira and Loucks 1997).

Reservoir Aggregation in Multi-Reservoir Systems

For systems consisting of many reservoirs, it is common to aggregate some of the reservoirs to reduce the computational or data storage demands of large simulation and optimization models (Saad, et al. 1994). A final use of the rules presented here is to help aggregate reservoirs and understand how to disaggregate operations of aggregated reservoirs. In this regard, reservoirs in series are typically easier to aggregate and disaggregate than reservoirs in parallel. A related use of these rules is to derive improved starting-point solutions for application of large scale optimization methods or more refined simulation studies.

Simulation and Optimization

In almost every case, simulation modeling is the standard by which operating rules are refined and tested. Simulation models can provide more realistic and detailed representation of reservoir system operations and much lower computational demands than optimization models for all but the most straight-forward cases. Simulation models also are more common in practice, and therefore are more likely to be trusted as a standard of comparison.

Comparison of proposed operating rules by simulation modeling provides many benefits. In many cases, simulation results will show that several sets of operating procedures will provide approximately equivalent performance (Wu 1988). In other cases, simulation results will demonstrate the tradeoffs of multiple performance objectives with different operating policies. In a few cases, the ability to demonstrate tradeoffs will aid in negotiations over how the system should best be operated. For all these purposes, it is useful to have a wide variety of potential operating rule forms available for examination.

Conclusions

Derived operating rules available for reservoirs in series and in parallel have been reviewed. These rules are summarized in Tables 1 and 2. A wide variety of operating rules are available for simulation modeling of reservoirs in series and in parallel. These provide a great deal of flexibility in the specification of system operations under various flow, storage, and demand conditions. Fortunately, for many single-purpose cases, a derived basis exists to help engineers narrow the search for appropriate operating rules.

A particular set of operating rules can be supported technically in a number of ways. Some rules are based on simple engineering principles for reservoir operations, such as keeping reservoirs full for water supply or empty for flood control. Several rules are derived from more formal optimization principles, such as the New York City space rules and hydropower production and energy storage rules. However, many rules in practice are based largely on empirical or experimental successes, either from actual operational performance, performance in simulation studies, or optimization results. These experimentally-supported rules are common for large multi-purpose projects.

Many opportunities exist for the use of formal optimization methods within reservoir simulation models. Examples include implementing storage allocation rules for reservoirs in series and in parallel, as well as general penalty-minimizing operations and allocating water among uses within a given time-step. Effective means of employing and hybridizing these rules are likely to be required for multi-purpose systems. Similarly, rules for shifting operating rule sets according to drawdown or refill conditions are an important area for further work.

Acknowledgments: This work was supported by the US Army Corps of Engineers Hydrologic Engineering Center and benefited considerably from discussions with Dan Barcellos, Richard Hayes, Morris Israel, Ken Kirby, Dan Sheer, Jery Stedinger, and Mimi Jenkins. Thanks also go to two reviewers for their helpful comments.

References

Bower, B.T., M.M. Hufschmidt, and W.W. Reedy (1966), "Operating Procedures: Their Role in the Design of Water-Resource Systems by Simulation Analyses," in A. Maass, et al., *Design of Water-Resource Systems*, Harvard University Press, Cambridge, MA, pp. 443-458.

Clark, E. J. (1950), “New York control curves,” *Journal of American Water Works Association*, Vol. 42, No. 9, pp. 823-827.

Clark, E.J. (1956), "Impounding Reservoirs," *Journal of the American Water Works Association*, Vol. 48, No. 4, April, pp. 349-354.

Ghosh, S. N (1986), *Flood Control and Drainage Engineering*, Oxford & IBH Publishing Co., pp. 158 - 159.

Hufschmidt, M.M. and M.B. Fiering (1966), *Simulation Techniques for Design of Water-Resource Systems*, Harvard University Press, Cambridge, MA.

Johnson, S.A., J.R. Stedinger, and K. Staschus (1991), "Heuristic Operating Policies for Reservoir System Simulation," *Water Resources Research*, Vol. 27, No. 6, May, pp. 673-685.

Karamouz, M. and M.H. Houck (1987), "Comparison of stochastic and deterministic dynamic programming for reservoir operating rule generation," *Water Resources Bulletin*, Vol. 23, No. 1, pp. 1-9.

Kelley, K.F. (1986), *Reservoir Operation During Drought: Case Studies*, Research Document No. 25, Hydrologic Engineering Center, Davis, CA.

Kelman, J., J.M. Damazio, J.L. Mariën, and J.P. daCosta (1989), "The Determination of Flood Control Volumes in a Multireservoir System", *Water Resources Research*, Vol. 25, No. 3, March, pp. 337-344.

Kim, Y.-O. and R.N. Palmer (1997), "Value of seasonal flow forecasts in Bayesian stochastic programming," *Journal of Water Resources Planning and Management*, ASCE, Vol. 123, No. 6, pp. 327-335.

Klemes, V. (1977), Discrete representation of storage for stochastic reservoir operation," *Water Resources Research*, Vol. 13, No. 1, pp. 149-158.

Labadie, J. (1997), "Reservoir system optimization models," *Water Resources Update*, University Council on Water Resources, Number 108, Summer, pp. 83-110.

Little, J.D.C. (1955), "The use of storage in a hydroelectric system," *Operations Research*, Vol. 3, pp. 187-197.

Loucks, D.P. and O.T. Sigvaldason (1982), "Multiple-reservoir operation in North America," in *The Operation of Multiple Reservoir Systems*, Z. Kaczmarek and J. Kindler (eds.), International Institute for Applied Systems Analysis, Laxenburg, Austria.

Lund, J.R. (submitted), "Power Production And Energy Drawdown Rules For Reservoirs," *Journal of Water Resources Planning and Management*, ASCE.

Lund, J.R. and I. Ferreira (1996), "Operating Rule Optimization for the Missouri River Reservoir System," *Journal of Water Resources Planning and Management*, ASCE, Vol. 122, No. 4, July.

Lund, J.R. and J. Guzman, (1996), "Developing Seasonal and Long-term Reservoir System Operation Plans using HEC-PRM", Technical Report RD-40, Hydrologic Engineering Center, U.S. Army Corps of Engineers, Davis, CA, June.

Mariën, J.L., J.M. Damázio, and F.S. Costa (1994), "Building flood control rule curves for multipurpose multireservoir systems using controlability conditions," *Water Resources Research*, Vol. 30, No. 4, April, pp. 1135-1144.

Nalbantis, I. and D. Koutsoyiannis (1997), "A parametric rule for planning and management of multiple-reservoir systems," *Water Resources Research*, Vol. 33, No. 9, September, pp. 2165-2178.

Needham, J.T. (1998), *Application of Linear Programming for Flood Control Operations on the Iowa and Des Moines Rivers*, MS Thesis, University of California, Davis.

Oliveira, R. and D.P. Loucks (1997), "Operating rules for multireservoir systems," *Water Resources Research*, Vol. 33, No. 4, April, pp. 839-852.

Palmer, R.N., J. Smith, J. Cohon, and C. ReVelle (1982), "Reservoir management in the Potomac River Basin," *Journal of Water Resources Planning and Management*, ASCE, Vol. 108, No. 1, pp. 47-66.

Saad, M., A. Turgeon, P. Bigras, and R. Duquette (1994), "Learning Disaggregation Technique for the Operation of Long-Term Hydroelectric Power Systems," *Water Resources Research*, Vol. 30, No. 11, November, pp. 3195-3203.

Sand, G.M. (1984), *An Analytical Investigation of Operating Policies for Water-Supply Reservoirs in Parallel*, Ph.D. Dissertation, Cornell University, Ithaca, NY.

Sheer, D. (1986), "Dan Sheer's Reservoir Operating Guidelines," as interpreted by J.R. Stedinger, School of Civil and Environmental Engineering, Cornell University, June/July, 1 page.

Stedinger, J. R., B. F. Sule, and D. Pei (1983), "Multiple reservoir system screening models," *Water Resources Research*, Vol. 19, No. 6, pp. 1383-1393.

Tejada-Guibert, J.A., S. Johnson, and J. Stedinger (1995), "The value of hydrologic information in stochastic dynamic programming models of multi-reservoir systems," *Water Resources Research*, Vol. 31, No. 10, pp. 2571-2579.

U.S. Army Corps of Engineers (1976), *Flood Control by Reservoirs,* Hydrologic Engineering Methods for Water Resources Development, Volume 7, Hydrologic Engineering Center, Davis, CA, February.

U.S. Army Corps of Engineers (1977), *Reservoir System Analysis for Conservation*, Hydrologic Engineering Methods for Water Resources Development, Volume 9, Hydrologic Engineering Center, Davis, CA, June.

U.S. Army Corps of Engineers (1985), *Engineering Manual: Engineering and Design, Hydropower*, EM 1110-2-1701, U.S. Army Corps of Engineers, Washington D.C.

Wu, Ray-Shyan (1988), *Derivation of Balancing Curves for Multiple Reservoir Operation*, M.S. Thesis, Civil and Environmental Engineering, Cornell University, Ithaca, NY.

Wurbs, R.A. (1966), *Modeling and Analysis of Reservoir System Operations*, Prentice-Hall PTR, Upper Saddle River, NJ.

Appendix I

"Storage Effectiveness Index" Rules for Hydropower Production

The "storage effectiveness index" has been developed by the U.S. Army Corps of Engineers for maximizing firm hydropower production during the drawdown season (USACE, 1985). For each reservoir, a "storage effectiveness index" is calculated for each time-step, using forecast inflows and power demands for the current time-step and remaining time-steps in the drawdown season. Reservoirs with a low index value are drawn down first.

Step 1: Find the firm energy requirement for the current time-step, E_f.

Step 2: Estimate the shortfall of firm hydropower production due to insufficient inflows to the system.

$$S_f = E_f - \frac{720}{11.81}\sum_{i=1}^{n} I_{Ui} H_i(S_i) e_i ,$$

where,

S_f = energy shortage for the current time-step,

I_{Ui} = inflow upstream of reservoir i during the current time-step,

$H_i(S_i)$ = hydropower head as a function of reservoir storage for reservoir i,

S_i = current reservoir storage for reservoir i,

e_i = the hydropower production efficiency of reservoir i, and

the constant is a conversion factor for I_{Ui} in cfs, H_i in feet, and S_f in Kwh.

This assumes that all flow can be utilized through the turbines.

Step 3: For each reservoir, estimate the drawdown required for that reservoir to individually eliminate the shortfall.

$$S_f = \frac{720}{11.81(59.5)} \Delta S_i H_i e_i ,$$

where,

$1/59.5$ = conversion of cfs to acre-ft draft per month,

ΔS_i = drawdown, in acre-ft., and

H_i = an average head corresponding to the drawdown (often found iteratively).

Solving for ΔS_i:

$$\Delta S_i = \frac{11.81(59.5)}{720} S_f / (H_i e_i) .$$

Step 4: For each reservoir, estimate the energy loss in the remainder of the drawdown season due to a drawdown of ΔS_i during this time-step (month).

$$E_{Li} = \frac{720}{11.81}(CI_{Ui} + V_{pi}) H_i (S_i - \Delta S_i) e_i / 59.5,$$

where,

E_{Li} = drawdown season power loss due to drawdown of reservoir i by ΔS_i,

CI_{Ui} = the cumulative natural inflow upstream of reservoir i for the remainder of the refill season, and

V_{pi} = the volume (acre-ft) of upstream storage to be emptied during the remainder of the drawdown season.

Step 5: Calculate the storage effectiveness ratio for each reservoir i:

$$SER_i = E_{Li} / S_f.$$

Reservoirs with the lowest ratios are to be drawn down first.

Appendix II: Derivation of Hydropower Production Rules for Reservoirs in Series

Linear Programming Short Term Storage Allocation

The objective is to find the allocation of a total storage volume which maximizes hydropower production for one period, subject to inflow forecasts for each reservoir, reservoir storage capacities, and a total storage target. As an equilibrium analysis, changes in reservoir storage are neglected. This is expressed mathematically below.

(1) Max P = $g \sum_{i=1}^{n} H_i(S_i) Q_i e_i$

Subject to:

(2) $Q_1 = I_1,$

(3) $Q_i = Q_{i-1} + I_i , \ \forall\, i > 1$

(4) $S_i \le K_i , \quad \forall\, i$

(5) $\sum_{i=1}^{n} S_i = S$

Definition of Variables:

P sum total of energy produced by all reservoirs,

n number of reservoirs in system,

H_i level of head in reservoir i,

S_i storage target for reservoir i,
S Total system storage target,
Q_i total inflow and release for reservoir i,
I_i direct inflows into reservoir i,
K_i storage capacity of reservoir i,
e_i efficiency of turbines in reservoir i,
V_i change in overall power production, P with change in storage in reservoir i.
γ unit weight of water
Note: reservoir i =1 is most upstream reservoir in series.

For short term allocation, the head-storage relationship can often be linearized, or $H_i(S_i) = a_i S_i$. The following linear program results:

(6) Max $P = g \sum_{i=1}^{n} a_i e_i S_i Q_i$

Subject to constraint Equations (2) - (5)

where a_i is a constant relating change in head with change of storage in reservoir i (for small changes in head). $H_i(S_i)$ is commonly non-linear, but may be piece-wise linearized and solved with a linear program, since head-storage relationships for most reservoirs are concave.

Derivation of Hydropower Production Rule

The above linear programming formulation can typically be simplified, where the head-storage relationship can be linearized. Using objective function (6) and substituting in equations (2) and (3) results in the simpler linear program.

(7) Max $\mathrm{P} = g \sum_{i=1}^{n} a_i e_i \left(\sum_{j=1}^{i} I_j \right) S_i$

Subject to constraint equations (4) and (5). This problem is solved by finding the slope of the objective function with respect to storage for each reservoir,

(8) $\dfrac{\partial P}{\partial S_i} a a_i e_i \left(\sum_{j=1}^{i} I_j \right) = V_i$.

Rule: Fill reservoirs in order of highest to lowest V_i until total storage, S, is filled. Empty in the reverse order.

Linear Programming Short Term Energy Storage Allocation

This problem is slightly modified from that above in that instead of seeking to maximize hydropower production subject to a given total water storage, it is desired to maximize hydropower production of reservoirs in series subject to a given total energy storage. For this problem, the objective and constraints in Equations 7 and 4 are re-stated below, and Equation 5 modified to an energy storage constraint.

(9) Max $\mathrm{P} = g \sum_{i=1}^{n} a_i e_i \left(\sum_{j=1}^{i} I_j \right) S_i$

Subject to:

(10) $S_i \quad K_i, \ " \ i$

(11) $$\sum_{i=1}^{n} E_i(S_i) = E,$$

where E is the total energy storage sought and $E_i(S_i)$ is the energy content of each reservoir as a function of it water storage.

For small changes in reservoir storage, the function $E_i(S_i)$ can probably be linearized into the form $b_i S_i$ (where b_i is a constant), allowing Equation 11 to be made linear and Equations 9-11 to be employed as a linear program to allocate storage among reservoirs in series to maximize hydropower production, subject to a total energy storage level. For large changes in storage, another solution method is likely to be required, since $E_i(S_i)$ is convex.

Appendix III
Derivations of Parallel Balancing Rules

These derivations of the New York City and Space rules are adapted from derivations presented by Sand (1984) and Johnson, et al. (1991). These derivations are further extended to examine energy storage and water quality applications.

Basic Derivation of the New York City Rule

Definition of Variables:

- z value of objective function,
- h_i unit value of water in reservoir *i*,
- n number of reservoirs in the system,
- S_{fi} end-of-period storage for the current period for reservoir *i*,
- S_{0i} beginning of current period storage for reservoir i,
- CQ_i the cumulative inflow to reservoir *i* from the end of the current period to the end of the refill cycle,
- K_i storage capacity of reservoir *i*, assumed to be the same in every period,
- V total volume of water in storage at the end of the current period,
- I_i inflow to reservoir *i* for the current period,
- D demand for current period (release for current period), and
- $f_i(CQ_i)$ the probability density of CQ_i.

The objective of the New York City balancing rule is to minimize the expected value EV() of total cumulative spill from all of the parallel reservoirs at the end of the refill season. This is reflected in the objective function in Equation 1. In Equation 1, the term h_i represents the relative value of water stored in each reservoir. Variation in the value of water can reflect variation in pumping costs, treatment costs, or energy content between the various reservoirs, as discussed later in the derivation. To implement this rule, this optimization problem is solved for each time-step during the refill reason.

(1) Min z = $$EV\left(\sum_{i=1}^{n} h_i \min(0, S_{fi} + CQ_i - K_i)\right)$$

subject to the following constraints:

(2) $$\sum_{i=1}^{n} S_{fi} = V$$

(3) $$V = \sum_{i=1}^{n} (S_{0i} + I_i) - D$$

Constraint Equation 3 indicates that the total water available should equal the sum of available water (current storage plus current period inflows) minus downstream water demands for the current period.

Expanding the expected value function in the objective function (Equation 1) yields,

(4) $$\text{Min } z = \sum_{i=1}^{n} h_i \left(\int_{K_i - S_{fi}}^{\infty} (S_{fi} + CQ_i - K_i) f_i(CQ_i) dCQ_i \right)$$

subject to constraint Equations 2 and 3.

The Lagrangian for this problem is:

(5) $$L = \sum_{i=1}^{n} h_i \left(\int_{K_i - S_{fi}}^{\infty} (S_{fi} + CQ_i - K_i) f_i(CQ_i) dCQ_i \right) + \lambda \left(\sum_{i=1}^{n} S_{fi} - V \right)$$

$$= \sum_{i=1}^{n} h_i \left((S_{fi} - K_i) \int_{K_i - S_{fi}}^{\infty} f_i(CQ_i) dCQ_i + \int_{K_i - S_{fi}}^{\infty} CQ_i f_i(CQ_i) dCQ_i \right) + \lambda \left(\sum_{i=1}^{n} S_{fi} - V \right)$$

or

(6) $$L = \sum_{i=1}^{n} h_i \left((S_{fi} - K_i) \left(1 - \int_{0}^{K_i - S_{fi}} f_i(CQ_i) dCQ_i \right) + \overline{CQ}_i - \int_{0}^{K_i - S_{fi}} CQ_i f_i(CQ_i) dCQ_i \right) + \; l \left(\sum_{i=1}^{n} S_{fi} - V \right),$$

where $\overline{CQ}_i$ is the expected value of cumulative inflows for reservoir i during the remainder of the refill season.

The first order conditions for solving this problem are:

(7) $$\frac{\partial L}{\partial S_{fi}} = 0$$

$$= h_i \left(\left(1 - \int_{0}^{K_i - S_{fi}} f_i(CQ_i) dCQ_i \right) + (S_{fi} - K_i)(-f_i(CQ_i = K_i - S_{fi})) - (K_i - S_{fi}) f_i(CQ_i = K_i - S_{fi}) \right) + \lambda$$

or

(8) $$h_i \left(1 - \int_{0}^{K_i - S_{fi}} f_i(CQ_i) dCQ_i \right) = \lambda$$, or

(9) $h_i \, Pr(CQ_i > K_i - S_{fi}) = l$, $\forall$ i, or

(10) h_i Pr(any spill in reservoir i) = λ.

This general result indicates that the storage targets for all reservoirs should have the same probability of spill, weighted by the value of water for each reservoir, h_i. The use of the New York City Rule for water supply, water quality, and energy storage purposes all follow Equations 9 and 10, with different interpretations of h_i.

New York City Water Supply Rule

For simple water supply purposes, h_i has the same value for all i, so Equation (9) becomes:

(11) $\Pr(CQ_i > K_i - S_{fi}) = \lambda, \forall i$.

New York City Rule for Water Quality

For simple water supply purposes with important water quality differences (e.g., TDS or perhaps water temperature or TSS) between reservoirs, h_i varies between reservoirs and can be interpreted as the marginal value of water use minus its water treatment cost for reservoir i. In this case Equations 9 and 10 remain the same, but with this net water value varying with spills of different water qualities.

New York City Rule for Energy Storage

Here, the objective is to minimize the expected value of potential energy spilled rather than physical water spilled. Here, h_i has the interpretation of the energy content of water stored in reservoir i. Equations 9 and 10 remain the same and applied with this interpretation.

Derivation of Space Rules

Returning to Equation 9, the central result of the New York City Rule, the assumption is made that the distributions $f_i(CQ_i)$ have the same distributional form, except that they are scaled by the average cumulative flow of the basin, $\overline{CQ}_i$. Where this assumption holds, then the distributions,

(12) $$f_i\left(CQ_i / \overline{CQ_i}\right) = f_j\left(CQ_j / \overline{CQ}_j\right),$$

for any two reservoirs i and j.

Where this is the case, the ratio $(K_i - S_{fi})/\overline{CQ}_i$ becomes a standard deviate for all distributions, having the same probability of exceedence for all reservoirs. If this ratio is set so that it equals the same ratio at the basin-wide scale,

(13) $$\frac{\sum_{i=1}^{n} K_i - V}{\sum_{i=1}^{n} EV(CQ_i)},$$

then the reservoirs are all balanced in terms of minimizing expected value of spill and maximizing capture of current inflows.

By replacing water inflows, water storage capacities, and water storage levels with energy inflows, energy storage capacities, and energy storage levels, the space rule can be adapted to energy storage purposes much as the NYC rule can be adapted to other operating purposes (Johnson, et al., 1991).

Derivation of Linear Programming-NYC Rules

Derivation of the linear programming space rules begins with the Equations (1-3) used in the derivation of the NYC rules. Additional constraints can also be added to the LP-NYC rule problem, so long as the additional constraints are linear.

In this case the expected value operator in Equation 1 is replaced by use of the weighed summation of spill values that would result from each year of the historical record or the probabilities of several representative year-types. Given historical streamflows of equal record length >*m* for each of *n* reservoirs, *m* representative refill seasons can be inferred. This yields the following linear program:

(14) $$\text{Min } z = \sum_{j=1}^{m} p_j \sum_{i=1}^{n} h_i L_{ij}$$

Subject to:

(15) $$\sum_{i=1}^{n} S_{fi} = V$$

(16) $$L_{ij} - E_{ij} = CQ_{ij} + K_i - S_{fi} \quad \text{for all i and j,}$$

(17) $$V = \sum_{i=1}^{n} (S_{fi} + Q_i) - D$$

plus any other linear constraints on present-period operations

where

- m number of equally probable refill seasons;
- p_j probability of hydrologic year-type j;
- CQ_{ij} in hydrologic year *j* , the expected cumulative inflow to reservoir *i* from the end of the current period to the end of the refill cycle;
- L_{ij} spill from reservoir *i* under hydrologic year *j*; and
- E_{ij} empty storage capacity in reservoir *i* under hydrologic year *j*.