An aggregate stochastic dynamic programming model of multi-reservoir systems

T.W. Archibald, † K.I.M. McKinnon, L.C. Thomas

October 1996

MS 96-018

† Department of Mathematics and Statistics
Department of Business Studies

Report also appears as Dept of Business Studies Working paper 96/14
It is a significantly revised version of report MS95-028 and BS 95/12
It will appear in Water Resources Research
Supported by EPSRC research grant GRH43700

Department of Mathematics and Statistics
University of Edinburgh, The King's Buildings, Edinburgh EH9 3JZ
Tel. (33) 131 650 5042 E-Mail : ken@maths.ed.ac.uk

An aggregate stochastic dynamic programming model of multi-reservoir systems

T. W. Archibald
Department of Business Studies, University of Edinburgh, William Robertson Building, 50 George Square, Edinburgh, EH8 9JY, UK

K. I. M. McKinnon
Department of Mathematics and Statistics, University of Edinburgh, The King's Buildings, Mayfield Road, Edinburgh, EH9 3JZ, UK

L. C. Thomas
Department of Business Studies, University of Edinburgh, William Robertson Building, 50 George Square, Edinburgh, EH8 9JY, UK

Abstract: We present a new method of determining an operating policy for a multi-reservoir system in which the operating policy for a reservoir is determined by solving a stochastic dynamic programming model consisting of that reservoir and a 2-dimensional representation of the rest of the system. The method is practical for systems with many reservoirs because the time required to determine an operating policy only increases quadratically with the number of reservoirs in the system and the operating policy for a reservoir is a function of few variables. We apply the method to examples of multi-reservoir systems with between 3 and 17 reservoirs and show that the operating policies determined are very close to optimal.

Keywords: Multi-reservoir systems, stochastic dynamic programming, aggregation.

American Geophysical Union index terms: Reservoirs (surface), Stochastic processes, System operation and management, Decision making under uncertainty.

1 Introduction

Stochastic dynamic programming models are attractive for multi-reservoir control problems because they allow the non-convex features of the problem, such as head effects, to be incorporated and the stochastic features of the problem to be modeled as Markov processes [3]. However with the exception of the simplest cases, these models are computationally intractable because of the large state and action spaces involved. Several methods of reducing the computational burden have been considered. [2] used high-order piecewise polynomial functions to approximate the value function so that a coarse discretization of the state space can be used. This approach proved successful in reducing the solution time for systems with 2–5 reservoirs, but its usefulness in general is limited because solution time still increases exponentially with the number of reservoirs in the system. Another approach is to combine the reservoirs in the system into one aggregate reservoir (see for example [4]). Although the resulting model can be solved easily, most of the detail of the original model is lost and applying the solution to the actual problem is a complex task. [6] proposes a method in which the control problem for a system of M reservoirs in series is decomposed into M subproblems each with two reservoirs — one a reservoir from the original problem and the other an aggregate representation of the reservoirs downstream of that reservoir. In this case solution time increases linearly with the number of reservoirs in the system and so the approach is practical for large systems. [5] proposes a different two reservoir decomposition method for a system of reservoirs in parallel. The disadvantage of these methods is that they can only be applied to a very restricted class of reservoir network.

This paper proposes a method of determining an operating policy for a broad class of reservoir network that is practical for systems with many reservoirs. The operating policy for a particular reservoir is determined by solving a stochastic dynamic programming model with a 3-dimensional representation of the volumes of water stored in the reservoirs. This representation consists of: a detailed model of the particular reservoir, an approximate model of reservoirs whose releases can reach that reservoir, and an approximate model of the remainder of the system. The justification for this decomposition is the belief that the factors which most influence the decision as to how much water to release from a reservoir are: the volume of water stored in that reservoir, the volume of water to take from reservoirs above that reservoir, the volume of water to pass on to reservoirs below that reservoir, and the effect of these on immediate and future rewards. There are several advantages to this approach to the problem: the solution time increases quadratically with the number of reservoirs in the system, so large problems can be tackled; the operating policy for a reservoir can be determined independently, so decision making can be decentralized and parallel processing can be used; the operating policy for a reservoir is a function of few variables, making it easy to implement; and the method can be applied to any acyclic network of reservoirs in which the release from a reservoir enters at most one other reservoir.

Section 2 describes a general model of a multi-reservoir system. Section 3 indicates how this model can be solved by discrete dynamic programming. This solution method will be referred to as the *full method*. In section 4 we present an alternative solution method for the general model which uses decomposition and aggregation techniques.

This method will be referred to as the *aggregate method*. Section 5 presents a comparison of the solution methods we consider.

2 A general model of a multi-reservoir system

This section describes a model of a multi-reservoir system which is appropriate for a broad class of reservoir networks and many applications including irrigation, water supply and hydro-electric power generation.

As presented the model applies to multi-reservoir systems for which the following assumptions are valid.

1. The reservoir network is connected.
2. The reservoir network is acyclic.
3. The water released from a reservoir directly enters at most one other reservoir.
4. The flow of water between reservoirs is instantaneous.
5. There are no evaporation or seepage losses.

The full and aggregate methods can be applied to reservoir networks that are not connected, but in such cases it is likely that there is a decomposition approach that explicitly exploits the physical separations in the network and so leads to a more efficient solution. Assumptions 2 and 3 rule out the possibility of pump storage in a hydro-electric power generation system. These assumptions can be removed at the expense of increasing the number of decision variables, but this greatly increases the time required to solve the model. When the time period being considered is large (for example a month), the time required for water to flow between reservoirs is often insignificant and so assumption 4 is reasonable. The solution methods can be modified for models that include evaporation or seepage losses for cases in which these are thought to be significant.

There are T time periods in the model. Revenue from the next period is discounted by a factor β. There are M reservoirs in the system. The set of reservoirs whose releases flow directly into reservoir i is denoted by I_i. Since the reservoir network is acyclic, the reservoirs can be labeled so that $I_i \subset \{1, 2, \ldots, i-1\}$. The volume in reservoir i at the beginning of each period is constrained by lower and upper limits, $\underline{H}_i$ and $\overline{H}_i$ respectively.

During a period water from external sources (for example rain or melting snow) flows into each reservoir. These inflows are not generally independent and to capture the correlation between them, Q inflow patterns are considered in each time period. The inflow to reservoir i from external sources during period t in pattern j is denoted by $q_{i,j}^t$.

The state of the system during a period is comprised of the volumes in the reservoirs and the hydrological state at the beginning of the period. Describing the hydrological state as a Markov process allows the persistence of the inflow patterns to be modeled. The volume in reservoir i at the beginning of period t is denoted by h_i^t and the hydrological state at the beginning of period t is denoted by s^t. The hydrological state in a period is represented by one of the S discrete values in the set $\{1, 2, \ldots, S\}$.

The probability that inflow pattern j occurs during period t depends upon the hydrological state, s^t, at the beginning of period t, and is denoted by $\Pr(j|s^t)$. The probability of the system being in hydrological state s^{t+1} at the beginning of period $t+1$ depends upon the hydrological state, s^t, and the inflow pattern, j, at period t, and is denoted by $\Pr(s^{t+1}|s^t, j)$.

A decision has to be made as to the amount of water to release from each reservoir during each period. We assume that the inflow pattern is known before this decision is made, so the decisions taken at period t depend upon the volumes in the reservoirs and the hydrological state at the beginning of period t and the inflow pattern during period t. The release from reservoir i during period t is denoted by $x_i^t(\underline{h}, s, j)$ where $\underline{h}$ and s are respectively the volumes in the reservoirs and the hydrological state at the beginning of period t, and j is the inflow pattern during period t. To simplify notation the variables upon which the releases depend will be omitted from this point on.

The immediate reward in period t when the volumes in the reservoirs at the beginning of period t are $\underline{h}$ and the releases from the reservoirs during period t are $\underline{x}$ is denoted by $r^t(\underline{h}, \underline{x})$ and can be any function of $\underline{h}$ and $\underline{x}$. In our test problems the immediate reward functions are of the following form.

$$r^t(\underline{h}, \underline{x}) = R^t\left(\sum_{i=1}^{M} G_i^t(h_i, x_i)\right) - \sum_{i=1}^{M} P_i^t(x_i)$$

This allows the reward function to depend upon the volume of water in each reservoir, so that the model can be applied to hydro-electric power generation systems. For such systems the functions appearing in the expression above can be interpreted as follows. The function G_i^t expresses the amount of electricity generated at reservoir i during period t in terms of the volume in reservoir i at the beginning of period t and the total release from reservoir i during period t. The dependency on the volume in the reservoir allows the efficiency of the generator to vary with the head of water in the reservoir. The dependency on the total release allows for spillage when the flow through the generator is greater than the generation capacity. The function R^t expresses the revenue obtained during period t in terms of the total amount of electricity generated in period t. The function P_i^t expresses the penalty cost for flooding caused by the release from reservoir i during period t in terms of the total release from reservoir i during period t.

The future value of the water stored in reservoirs at the end of the planning horizon depends on the prevailing hydrological state and can be any function of the volumes in the reservoirs $\underline{h}$ and the hydrological state s. We envisage the model being used with a rolling planning horizon where an instance of the model would be solved each time a decision had to be made. In this case only the first period decision is ever implemented and, provided the planning horizon is sufficiently long, the terminal value function will not have a significant effect on this. In our test problems the terminal value function is of the form $\sum_{i=1}^{M} \tau_i(h_i, s)$, where τ_i expresses the future value of the water in reservoir i as a function of the volume in reservoir i and the hydrological state at the end of period T.

The maximum discounted return from operating the reservoir system from the beginning of period t until the end of period T is denoted by $v^t(\underline{h}, s)$ where $\underline{h}$ and s are respectively the volumes in the reservoirs and the hydrological state at the beginning of period t. It can be seen from the above that for $t \leq T$, $v^t(\underline{h}, s)$ satisfies the following optimality equation.

$$v^t(\underline{h}^t, s^t) = \sum_{j=1}^{Q} \Pr(j|s^t) w^t(\underline{h}^t, s^t, j) \tag{1}$$

where

$$w^t(\underline{h}^t, s^t, j) = \max_{\underline{x}^t} \left\{ r^t(\underline{h}^t, \underline{x}^t) + \beta \sum_{s^{t+1}=1}^{S} \Pr(s^{t+1}|s^t, j) v^{t+1}(\underline{h}^{t+1}, s^{t+1}) \right\}$$

subject to

$$h_i^{t+1} = h_i^t + q_{i,j}^t + \sum_{k \in I_i} x_k^t - x_i^t, \tag{2}$$

$$\underline{H}_i \leq h_i^{t+1} \leq \overline{H}_i, \; x_i^t \geq 0;$$

$$v^{T+1}(\underline{h}, s) = \sum_{i=1}^{M} \tau_i(h_i, s).$$

3 Discrete dynamic programming

An approximate solution to (1) can be found by discretizing the state variables and decision variables and using the standard discrete dynamic programming recursion. We do this by discretizing the volumes in the reservoirs using a regular discretization in which N equally spaced values between $\underline{H}_k$ and $\overline{H}_k$ are considered for the volume in reservoir k. This means that there are N^M possible realizations of the state of the reservoirs in the system. We only consider those releases from reservoir k which ensure that the volume in reservoir k at the beginning of the next period is equal to one of the N values in the discretization of the volume in reservoir k. Hence we may have to consider as many as N releases from each reservoir, giving a total of N^M possible decisions in the worst case. Since we use the same discretization at each period, the decision not to change the reservoir volumes is always feasible. Hence discrete dynamic programming is well-defined for our choice of discretization.

4 Using decomposition and aggregation techniques

4.1 Introduction

This section proposes a new method of decomposing the model described in section 2 into a number of independent small-scale subproblems. There is one subproblem for each reservoir in the system and each subproblem is solved by discrete dynamic programming. The solution to a subproblem determines target releases for the corresponding reservoir. Since the target releases are determined independently, they do not necessarily constitute a feasible set of releases for the original model. A procedure is proposed to check the feasibility of the target releases and modify them if required. Often a first period decision is all that is required, because decisions for subsequent periods will be determined by reapplying the method with initial conditions appropriate to these periods. The alternative approach, which avoids the overhead of resolving for each time period, is to use the solutions to the subproblems to generate a T-period operating policy for the original model. Interpolation is required to determine target releases for periods $2, 3, \ldots, T$ as is explained at the end of section 4.3. The expected value of this alternative approach can be approximated using simulation on a sample of the future weather conditions over the T periods.

In each subproblem an operating policy that is feasible for the original model is evaluated using the immediate reward and terminal value functions from the original model. This is one of the advantages of this method over methods which aggregate reservoirs into a single reservoir. The immediate reward functions are complex functions involving spill and non-linear head effects and so are difficult to approximate as functions of an aggregate volume and release. Also since the optimal policy for each subproblem is a feasible policy for the original model, the value functions for the subproblems are lower bounds on the value function for the original model.

Each subproblem corresponds to a unique partitioning of the set of reservoirs in the system into three sets. These sets will be referred to as the focus reservoir, the upstream reservoirs and the non-upstream reservoirs. The focus reservoir is one of the reservoirs from the original problem. Other reservoirs are classified as upstream or non-upstream depending on whether their releases can or cannot reach the focus reservoir. Releases from an upstream reservoir can reach (either directly or indirectly) the focus reservoir, and releases from a non-upstream reservoir cannot reach the focus reservoir. As a consequence of assumptions 2 and 3, water stored in the upstream reservoirs must pass through the focus reservoir before leaving the system.

In each subproblem the state of the reservoirs in the system is represented by a vector with (no more than) three dimensions. This representation is comprised of the volume in the focus reservoir, the state of the upstream reservoirs and the state of the non-upstream reservoirs. The state of the upstream reservoirs describes a set of volumes in the upstream reservoirs and similarly the state of the non-upstream reservoirs describes a set of volumes in the non-upstream reservoirs. The aim is to characterize the states of the upstream and non-upstream reservoirs in a manner that captures sufficient detail and yields a computationally tractable model. For example if we include all combinations of discrete reservoir volumes in our characterizations, then each subproblem would be equivalent to the full method, but no subproblem could be solved in a reasonable time.

With our decomposition method the change in the release from a reservoir when the total volume in the upstream reservoirs is increased can be different from the change when the total volume in the non-upstream reservoirs is increased. Methods that use a single aggregate reservoir cannot distinguish between the two situations.

4.2 Constructing the subproblem for reservoir f

This section gives details of the subproblem corresponding to reservoir f.

The focus reservoir is reservoir f from the original problem. Let $U_f = I_f \cup \left\{\cup_{i \in I_f} U_i\right\}$ denote the set of upstream reservoirs for reservoir f. As a consequence of the labeling of the reservoirs we have chosen to use (see section 2), the sets U_i, $1 \leq i \leq M$, can be calculated explicitly in the order of increasing i. Let N_f denote the set of non-upstream

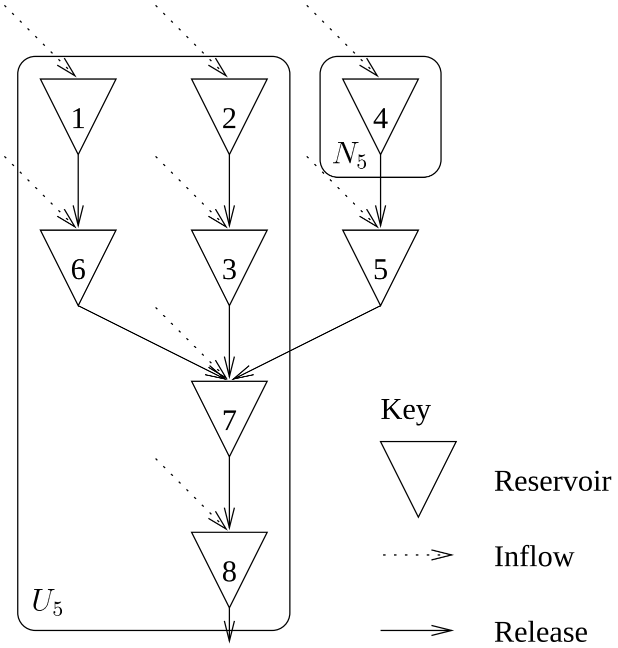

Figure 1: The partitioning of the reservoir network for L08 with reservoir 5 as the focus

reservoirs for reservoir f. By definition N_f is the set of reservoirs not included in $U_f \cup \{f\}$. If $U_f = \emptyset$ or $N_f = \emptyset$ then the representation of the state of the reservoirs in the system will only have two dimensions.

Figure 1 illustrates the partitioning of an eight reservoir system corresponding to focus reservoir 5. For this case $U_5 = \{4\}$ and $N_5 = \{1, 2, 3, 6, 7, 8\}$.

The state of the system at a period is composed of the state of the upstream reservoirs, the volume in the focus reservoir, the state of the non-upstream reservoirs and the hydrological state at the beginning of the period. For period t these are denoted by $\underline{a}^t_f$, h^t_f, $\underline{b}^t_f$, and s^t respectively. The volume in reservoir $i \neq f$ at the beginning of period t, is given by $\underline{a}^t_f$ or $\underline{b}^t_f$ as follows:

$$a^t_{f,i} = h^t_i \text{ if } i \in U_f$$

and is undefined otherwise, and

$$b^t_{f,i} = h^t_i \text{ if } i \in N_f$$

and is undefined otherwise. The possible states that the upstream (the "A-bove") reservoirs and non-upstream (the "B-elow") reservoirs can be in at the beginning of period t are denoted by the sets A^t_f and B^t_f respectively.

The decisions to be taken at a period are the state in which to leave the upstream reservoirs for the next period, the amount of water to release from the focus reservoir during the period and the state in which to leave the non-upstream reservoirs for the next period. For period t these decisions are denoted by $\underline{a}^{t+1}_f$, x^t_f and $\underline{b}^{t+1}_f$. The decisions depend on the values of the state variables at the beginning of the period and on the inflow pattern during the period. As a consequence of assumptions 2 and 3, the amount of water released from reservoir $i \neq f$ during period t, can be determined directly from the water balance equation (2) as follows.

$$x_i^t = h_i^t + q_{i,j}^t + \sum_{k \in I_i} x_k^t - h_i^{t+1} \text{ when inflow pattern } j \text{ occurs in period } t. \tag{3}$$

Due to the labeling of the reservoirs we have chosen to use (see section 2), x_i^t can be calculated in order of increasing i. If $x_i^t < 0$ for any i, then the corresponding decision is infeasible for that state and inflow pattern at period t and so is not considered. The method can be extended to deal with problems for which assumptions 2 and 3 do not hold. The subproblems would then take longer to solve, because an optimization would be required to determine the best way to release water from reservoirs in order to achieve a given change in the volumes in the reservoirs. When many reservoirs can release water directly into more than one reservoir, this optimization cannot be performed in a reasonable time.

Let $\nu_f^t(\underline{a}_f^t, h_f^t, \underline{b}_f^t, s^t)$ denote the maximum discounted return from operating the reservoir system from the beginning of period t until the end of period T when (at the beginning of period t) the state of the upstream reservoirs is $\underline{a}_f^t$, the volume in the focus reservoir is h_f^t, the state of the non-upstream reservoirs is $\underline{b}_f^t$ and the hydrological state is s^t. It can be seen from the above that for $t \leq T$, $\nu_f^t(\underline{a}_f^t, h_f^t, \underline{b}_f^t, s^t)$ satisfies the following optimality equation.

$$\nu_f^t(\underline{a}_f^t, h_f^t, \underline{b}_f^t, s^t) = \sum_{j=1}^{Q} \Pr(j|s^t)\omega_f^t(\underline{a}_f^t, h_f^t, \underline{b}_f^t, s^t, j)$$

where

$$\omega_f^t(\underline{a}_f^t, h_f^t, \underline{b}_f^t, s^t, j)$$

$$= \max_{\underline{a}_f^{t+1}, x_f^t, \underline{b}_f^{t+1}} \left\{ r^t(\underline{h}^t, \underline{x}^t) + \beta \sum_{s^{t+1}=1}^{S} \Pr(s^{t+1}|s^t, j)\nu_f^{t+1}(\underline{a}_f^{t+1}, h_f^{t+1}, \underline{b}_f^{t+1}, s^{t+1}) \right\}$$

subject to

$$h_i^{t+1} = h_i^t + q_{i,j}^t + \sum_{k \in I_i} x_k^t - x_i^t,$$

$$\underline{a}_f^{t+1} \in A_f^{t+1},\ \underline{b}_f^{t+1} \in B_f^{t+1},\ \underline{H}_f \leq h_f^{t+1} \leq \overline{H}_f,\ x_i^t \geq 0,$$

where by definition $h_i^\ell = a_{f,i}^\ell$ if $i \in U_f$, $h_i^\ell = b_{f,i}^\ell$ if $i \in N_f$;

$$\nu_f^{T+1}(\underline{a}_f^{T+1}, h_f^{T+1}, \underline{b}_f^{T+1}, s^{T+1}) = \sum_{i=1}^{M} \tau_i(h_i^{T+1}, s^{T+1}).$$

In most cases the aim is to determine an optimal decision for given initial conditions. The initial conditions specify the volumes in the reservoirs, $\underline{h}^1$, the hydrological state and the inflow pattern for period 1. As these are the only volumes of interest to us in period 1, $\underline{h}^1$ gives complete characterizations of the upstream and non-upstream reservoirs in period 1.

$$A_f^1 = \{\underline{h}^1\}, \text{ and } B_f^1 = \{\underline{h}^1\}$$

For period $t > 1$ we need a method of identifying typical states for the upstream and non-upstream reservoirs. Define the capacity of a reservoir to be the difference between the upper and lower limits on the volume in the reservoir. Define the working volume of a reservoir to be the difference between the volume in the reservoir and the lower limit on the volume in the reservoir. We say that a set of reservoir volumes satisfies the *equally full heuristic* if the ratio of working volume to capacity is the same for all reservoirs in the set. For a wide range of conditions, operating the reservoirs according to the equally full heuristic yields good performance. We use the equally full heuristic to characterize the states of the upstream and non-upstream reservoirs as follows.

$$\begin{aligned} a_{f,i}^t &= \underline{H}_i + \hat{a}_f^t \left(\overline{H}_i - \underline{H}_i\right) \text{ if } i \in U_f,\ \hat{a}_f^t \in [0,1] \\ b_{f,i}^t &= \underline{H}_i + \hat{b}_f^t \left(\overline{H}_i - \underline{H}_i\right) \text{ if } i \in N_f,\ \hat{b}_f^t \in [0,1] \end{aligned} \tag{4}$$

For computational purposes we have to discretize the ranges of $\hat{a}_f^t$, h_f^t and $\hat{b}_f^t$. As for the full method, we use N equally spaced values between $\underline{H}_f$ and $\overline{H}_f$ in the discretization of h_f^t. To discretize $\hat{a}_f^t$ and $\hat{b}_f^t$ we use N equally spaced values between 0 and 1. Hence there are N^3 possible realizations of the state of the reservoirs in the system and, in the worst case, N^3 possible decisions in each state.

4.3 An operating policy for the original problem

Each subproblem provides a first period target release for its focus reservoir for the given initial conditions. Since the target releases are determined independently, they do not necessarily constitute a feasible set of releases for the original model. We propose the *top-down correction* method to identify a feasible set of first period releases for the original model from the target releases. This is only one of several possible methods to determine feasible first period releases from the solutions of the subproblems. We propose the top-down correction method because it is straightforward and has given good results with all of our test problems.

Top-down correction

1. For the top reservoirs (reservoir i is a top reservoir if $I_i = \emptyset$) the target releases are always feasible, so set the actual releases equal to the target releases.

2. Select any reservoir for which the actual releases have been set for all upstream reservoirs and apply the target release at that reservoir.

   If the target release is feasible, set the actual release equal to the target release.

   If the lower limit on the volume in the reservoir is violated, set the actual release so that the volume in the reservoir at the beginning of the next period is at its lower limit.

   If the upper limit on the volume in the reservoir is violated, set the actual release so that the volume in the reservoir at the beginning of the next period is at its upper limit.

3. Repeat step 2 until actual releases have been set for all reservoirs.

When the first period releases are applied, the volumes in the reservoirs at the beginning of the next period do not necessarily correspond to the characterizations of the upstream and non-upstream reservoirs in any of the subproblems. For example if the resulting volumes do not satisfy the equally full heuristic, there will not always be a corresponding state in the characterizations described in (4). Therefore to determine an operating policy for periods 2 onwards, we need to use interpolation.

We use a 3-dimensional linear interpolation scheme involving the potential of the upstream reservoirs, the volume in the focus reservoir and the potential of the non-upstream reservoirs. The potential of a set of reservoirs is the total amount of water released if every reservoir in the set is emptied. If the volumes in the reservoirs at the beginning of period t are $\underline{h}^t$ and the focus reservoir is f, then the potential of the upstream and non-upstream reservoirs at period t is given by

$$\sum_{i \in U_f} \left(h_i^t + \sum_{k \in U_i} h_k^t \right) \text{ and } \sum_{i \in N_f} \left(h_i^t + \sum_{k \in U_i \cap N_f} h_k^t \right) \tag{5}$$

respectively. We calculate the potential of the upstream and non-upstream reservoirs for the reservoir volumes described by the states in A_f^t and B_f^t to give a discrete set of potentials for the upstream and non-upstream reservoirs. The volume in the focus reservoir is discretized in the same way as in the corresponding subproblem. The solution to a subproblem gives target releases for the focus reservoir for all combinations of the discrete values of the potential of the upstream reservoirs, the volume in the focus reservoir and the potential of the non-upstream reservoirs. We use linear interpolation of these target releases to find target releases at other points. (This approach is also suitable for characterizations of the upstream and non-upstream reservoirs that do not use the equally full heuristic.)

For hydro-electric power generation problems, one could define the potential of a set of reservoirs to be the total amount of electricity that could be generated if every reservoir in the set is emptied. In (5) it would then be necessary to weight the total amount of water released from a reservoir by the efficiency of the generator at that reservoir. We have found that using the total amount of water released gives sufficient accuracy.

4.4 Evaluating the operating policy

The objective function values for the subproblems are lower bounds on the objective function value for the original problem, because they are determined using releases that are feasible for that problem. However there is no guarantee that these values are good approximations to the expected value of applying the operating policy constructed by combining the solutions to all the subproblems.

We use simulation to estimate the expected value of applying the combined operating policy. A sample of possible future weather conditions (ie the hydrological state and inflow pattern occurring in each period t, $1 \leq t \leq T$) is generated and the actions that would be selected by the combined operating policy are determined (using interpolation and top-down correction) and evaluated for each element in the sample. We can then compare this expected value with that for the operating policy determined by other solution approaches evaluated on the same sample.

5 A comparison of the full and aggregate methods

This section compares both the computational complexity and the quality of the solution of the full method with those of the aggregate method.

Table 1 compares the complexity of the full and aggregate methods in terms of the number of action evaluations required per time period. The second last column shows that for the full method the number of action evaluations required increases exponentially with the number of reservoirs, while for the aggregate method this number only increases linearly with the number of reservoirs. For both methods the time required to perform an action evaluation is proportional to the number of reservoirs. Hence the solution time for the aggregate method increases quadratically with the number of reservoirs.

Five test problems based on hydroelectric power generation systems have been used to assess the performance of the aggregate method described in section 4. T04 has 4 reservoirs in series and is described in detail in [6]. H03, H04 and L17 have 3, 4 and 17 reservoirs respectively and the connectivity of the reservoir networks is shown in figure 2. L08 has 8 reservoirs and the connectivity of the reservoir network is shown in figure 1. These 4 problems cannot be solved by the method proposed in [6] because the reservoirs are not connected in series. We used a two stage procedure to determine suitable terminal value functions for these 4 problems. We first solved the problems using terminal value functions based on the average efficiency of the reservoirs and the average price of electricity. Secondly we used the resulting optimal value functions to construct separable piecewise linear terminal value functions for the problems. Full details of H03, H04, L08 and L17 are given in the appendix. In all cases we used $N = 11$ points in our discretizations of the state variables.

Although H03 only has 3 reservoirs, the aggregate method is still quicker than the full method because all the subproblems use a 2-dimensional representation of the state of the reservoirs in the system. This occurs because in each subproblem either the set of upstream reservoirs or the set of non-upstream reservoirs is empty.

For H03, H04, L08 and L17, which all have $T = 5$ time periods, $S = 2$ hydrological states and $Q = 3$ inflow patterns, we are able to consider all possible future weather conditions (1296 in total) when evaluating the operating policy determined by the aggregate method. T04 has $T = 3$ time periods, $S = 1$ hydrological state and $Q = 5$ inflow patterns, so we only consider a sample of 3000 of the 5^{11} possible future weather

Table 1: Computational complexity of the full and aggregate methods

Method	Problems	States per problem	Inflow patterns	Actions per state	Action evaluations	Cost of action evaluations
Full	1	SN^M	Q	up to N^M	SQN^{2M}	$\propto M$
Aggregate	M	up to SN^3	Q	up to N^3	$MSQN^6$	$\propto M$

Table 2: Quality of solution of the full and aggregate methods

Problem	Full method Expected value	Aggregate method Difference in expected value	Difference in worst case	Difference in best case
T04[a]	24637373	−0.3%	−0.6%	0.0%
H03	19662882	−0.3%	−1.8%	+1.4%
H04	19729363	−0.2%	−2.2%	+1.4%

[a]Figures for T04 are based on a sample of 3000 future weather conditions

Table 3: Comparison of the exact solution and the aggregate method

Problem	Maximum expected value[a]	Aggregate method Difference in expected value	Difference in worst case	Difference in best case
L08	18721904	−2.2%	−10.2%	+8.7%
L17	18691236	−3.1%	−11.0%	+8.3%

[a]Determined by linear programming

conditions.

We have solved the 3 small problems (T04, H03 and H04) which have non-linear head effects using the aggregate method and the full method, and our findings are summarized in table 2. The difference between the expected values of the operating policies determined by the full and aggregate methods is at most 0.3%. Hence we have achieved a great reduction in solution time (a factor of 60 in the case of T04) without a significant reduction in the quality of the operating policy.

For T04 we used the initial conditions from [6], so that initially the top reservoir was 70% full and the others were 98% full. Hence the set of initial reservoir volumes does not satisfy the equally full heuristic. Despite this the aggregate method still gives very good results. This shows that the success of the method does not rely on the validity of the equally full heuristic.

As noted in section 4.1, in many cases only the first period decisions will be implemented. Future decisions will be determined by solving the model for initial conditions that reflect the changes in the volumes in reservoirs and the weather conditions. For this reason the first period decision is very important. The aggregate method and the full method frequently select the same first period decision. This suggests that in practice the difference between the two methods may be even less than the small margin mentioned above.

One of the features of the aggregate method is that it can be used to solve problems which are much larger than those that can be solved by the full method. We used L08 and L17 to assess the performance of the aggregate method for large multi-reservoir systems. In these problems the reservoir heads are constant, so the problems can be formulated as linear programming problems and solved exactly. Table 3 compares the expected value of the policy determined by the aggregate method with the exact maximum expected value found by linear programming. Given that we are using very crude discretizations of the high dimensional state spaces, it is very encouraging that the expected value of the operating policy determined by the aggregate method is within 3.1% of the maximum expected value. A more detailed comparison of linear programming and dynamic programming based solution methods can be found in [1].

The optimal objective function values in the subproblems can be within 1% of the expected value of the operating policy determined by the aggregate method. However there is a large variation and the subproblem which will give the best approximation cannot be determined beforehand. In the case of L08 the best approximation is given by the subproblem corresponding to focus reservoir 3.

While the policy determined by the full method (or in the case of L08 and L17 by linear programming) is optimal in terms of expected return, a different policy may result in better returns for some future weather conditions. This is often the case for the policy determined by the aggregate method and can be seen by the positive entries in the columns headed "difference in best case" in tables 2 and 3. The probability that the policy determined by the aggregate method leads to a higher return is 0.35, 0.46, 0.41 and 0.37 for H03, H04, L08 and L17 respectively.

Both methods offer considerable scope for parallel computing and the implementations we used exploited this. In the implementation of the full method, the optimal decisions for the inflow patterns in a period were determined in parallel. In the implementation of the aggregate method, the subproblems were solved in parallel and, within each subproblem, the optimal decisions for the inflow patterns in a period were determined in parallel.

6 Conclusions

We have presented a new method of determining an operating policy for multi-reservoir problems with head effects which reduces the original problem to a series of independent, low dimensional subproblems. A subproblem considers one of the reservoirs in detail and a 2-dimensional aggregate representation of the rest of the system. The subproblems can be solved routinely by dynamic programming. This aggregate method can be applied to acyclic reservoir networks in which each reservoir can release water directly into at most one other reservoir. We have reported the results of tests in which the aggregate method gives objective function values that are within 3.1% of optimum, which is highly satisfactory given the size and complexity of the problems we have considered.

Acknowledgments

The authors are grateful to the Engineering and Physical Sciences Research Council for their support of this work under grant GRH43700 and to Crawford Buchanan for computing the linear programming solutions.

References

[1] T. W. Archibald, C. S. Buchanan, K. I. M. McKinnon, and L. C. Thomas. Nested benders decomposition and dynamic programming for reservoir optimisation. Technical Report MS 96-001 and BS96/2, Department of Mathematics and Statistics, University of Edinburgh, 1996.

[2] S. A. Johnson, J. R. Stedinger, C. A. Shoemaker, Y. Li, and J. A. Tejada-Guibert. Numerical solution of continuous-state dynamic programs using linear and spline interpolation. *Operations Research*, 41:484–500, 1993.

[3] K. Reznicek and T. C. E. Cheng. Stochastic modelling of reservoir operations. *European Journal of Operational Research*, 50:235–248, 1991.

[4] M. Saad, A. Turgeon, P. Bigras, and R. Duquette. Learning disaggregation technique for the operation of long-term hydroelectric power systems. *Water Resources Research*, 30:3195–3202, 1994.

[5] A. Turgeon. Optimal operation of multireservoir power systems with stochastic inflows. *Water Resources Research*, 16:275–283, 1980.

[6] A. Turgeon. A decomposition method for the long term scheduling of reservoirs in series. *Water Resources Research*, 17:1565–1570, 1981.

Appendix

This appendix gives details of the 4 new test problems (H03, H04, L08 and L17) which have been used in this paper. All four problems have $T = 5$ time periods and use a discount factor $\beta = 1$. The connectivity of the reservoir networks in the problems is shown in figures 1 and 2. Each problem has three inflow patterns in each time period, $j =$ Low, Medium and High. The total rainfall, F_j^t, in inflow pattern j in time period t is given in table 4. A proportion f_i of the total rainfall in a period flows into reservoir i. Hence $q_{i,j}^t = f_i F_j^t$. These proportions and the upper limits on the volume of water that can be stored in the reservoirs are given in tables 6, 7, 8 and 9 for problems H03, H04, L08 and L17 respectively. The lower limit on reservoir storage is zero in all cases. Each problem has 2 hydrological states in each period, $s =$ Dry and Wet. The transition probabilities are given in table 5.

The functions R^t, G_i^t and P_i^t, which together comprise the immediate reward function, have the following forms.

$$\begin{aligned}
R^t(y) &= \rho_1^t \min\{1432779.5, y\} + \rho_2^t \max\{0, y - 1432779.5\} \\
G_i^t(h,x) &= \min\{g_i, x\} \left(z_{i,0} + z_{i,1} \min\{0.5\overline{H}_i, h\} + z_{i,2} \max\{0, h - 0.5\overline{H}_i\} \right) \\
P_i^t(x) &= 0
\end{aligned}$$

The constants ρ_n^t are given in table 4. The other constants are problem specific and are given in tables 6, 7, 8 and 9 for problems H03, H04, L08 and L17 respectively.

The functions τ_i, which define the terminal value of water stored in the system, have the following form.

$$\tau_i(h,s) = \begin{cases} \alpha_{i,s,1}h \text{ if } 0 \leq h \leq 0.3\overline{H}_i \\ 0.3\alpha_{i,s,1}\overline{H}_i + \alpha_{i,s,2}(h - 0.3\overline{H}_i) \text{ if } 0.3\overline{H}_i \leq h \leq 0.7\overline{H}_i \\ 0.3\alpha_{i,s,1}\overline{H}_i + 0.4\alpha_{i,s,2}\overline{H}_i + \alpha_{i,s,3}(h - 0.7\overline{H}_i) \text{ if } h \geq 0.7\overline{H}_i \end{cases}$$

The constants $\alpha_{i,s,n}$ are given in tables 10, 11, 12 and 13 for problems H03, H04, L08 and L17 respectively.

Table 4: Revenue function parameters and total rainfall for all 4 problems

t	F^t_{Low}	F^t_{Medium}	F^t_{High}	ρ^t_1	ρ^t_2
1	4240.23	19327.56	39739.83	0.85	0.51
2	4079.50	16166.12	32518.62	0.91	0.55
3	1660.59	5549.77	10811.60	1.35	0.81
4	2269.02	15730.27	33942.55	1.00	0.60
5	4313.20	20272.24	41863.89	0.73	0.44

Table 5: Transition probabilities for all problems and periods

j	s	$\Pr(j\|s)$	$\Pr(\text{Dry}\|s,j)$	$\Pr(\text{Wet}\|s,j)$
Low	Dry	0.50	0.95	0.05
Medium	Dry	0.30	0.75	0.25
High	Dry	0.20	0.55	0.45
Low	Wet	0.10	0.60	0.40
Medium	Wet	0.30	0.30	0.70
High	Wet	0.60	0.10	0.90

Table 6: Inflow distribution, head functions and storage and generation limits for H03

i	f_i	$\overline{H}_i$	$z_{i,0}$	$z_{i,1}$	$z_{i,2}$	g_i
1	0.2814	31030	177.90	3.5462E−4	2.6597E−4	4532.25
2	0.2776	31049	169.87	3.3841E−4	2.5381E−4	3541.80
3	0.4410	29363	172.60	3.6360E−4	2.7270E−4	11973.38

Table 7: Inflow distribution, head functions and storage and generation limits for H04

i	f_i	$\overline{H}_i$	$z_{i,0}$	$z_{i,1}$	$z_{i,2}$	g_i
1	0.2814	31030	177.90	3.5462E−4	2.6597E−4	4532.25
2	0.2776	31049	169.87	3.3841E−4	2.5381E−4	3541.80
3	0.1255	12866	43.36	2.0846E−4	1.5634E−4	8301.90
4	0.3155	16497	166.96	6.2600E−4	4.6950E−4	10222.23

Table 8: Inflow distribution, head functions and storage and generation limits for L08

i	f_i	$\overline{H}_i$	$z_{i,0}$	$z_{i,1}$	$z_{i,2}$	g_i
1	0.1802	20097	120.79	0.00	0.00	3919.93
2	0.0045	504	66.36	0.00	0.00	384.10
3	0.0967	10429	84.11	0.00	0.00	3949.62
4	0.1155	13015	125.90	0.00	0.00	920.10
5	0.1621	18034	114.27	0.00	0.00	4413.92
6	0.0419	4694	36.33	0.00	0.00	3141.01
7	0.1881	12866	67.40	0.00	0.00	7882.02
8	0.2110	11803	170.16	0.00	0.00	8727.84

Table 9: Inflow distribution, head functions and storage and generation limits for L17

i	f_i	$\overline{H}_i$	$z_{i,0}$	$z_{i,1}$	$z_{i,2}$	g_i
1	0.0060	675	24.87	0.00	0.00	480.40
2	0.1542	17217	87.90	0.00	0.00	1709.60
3	0.0200	2205	41.70	0.00	0.00	7464.30
4	0.0045	504	66.36	0.00	0.00	384.10
5	0.0486	5260	60.70	0.00	0.00	2839.60
6	0.0482	5169	53.50	0.00	0.00	2988.00
7	0.1155	13015	125.90	0.00	0.00	920.10
8	0.1116	12454	62.23	0.00	0.00	5013.50
9	0.0505	5580	71.30	0.00	0.00	2698.70
10	0.0228	2566	20.00	0.00	0.00	2439.10
11	0.0191	2128	25.40	0.00	0.00	2572.40
12	0.1255	12866	44.70	0.00	0.00	8301.90
13	0.0626	6934	22.70	0.00	0.00	7055.20
14	0.0270	3038	3.90	0.00	0.00	351.90
15	0.0273	3041	11.96	0.00	0.00	1132.20
16	0.0518	5724	45.00	0.00	0.00	1647.50
17	0.1050	11632	118.00	0.00	0.00	11831.30

Table 10: Terminal value functions for H03

i	$\alpha_{i,\text{Dry},1}$	$\alpha_{i,\text{Dry},2}$	$\alpha_{i,\text{Dry},3}$	$\alpha_{i,\text{Wet},1}$	$\alpha_{i,\text{Wet},2}$	$\alpha_{i,\text{Wet},3}$
1	203.86	143.35	102.69	156.76	103.26	45.66
2	185.42	123.53	82.52	138.78	77.80	34.49
3	65.40	47.68	32.38	31.12	17.70	6.03

Table 11: Terminal value functions for H04

i	$\alpha_{i,\text{Dry},1}$	$\alpha_{i,\text{Dry},2}$	$\alpha_{i,\text{Dry},3}$	$\alpha_{i,\text{Wet},1}$	$\alpha_{i,\text{Wet},2}$	$\alpha_{i,\text{Wet},3}$
1	203.51	140.22	100.87	157.33	102.67	42.62
2	188.61	119.00	73.78	139.95	77.68	30.74
3	57.21	51.48	44.77	29.29	18.98	12.77
4	42.23	32.97	21.20	14.23	7.56	3.01

Table 12: Terminal value functions for L08

i	$\alpha_{i,\text{Dry},1}$	$\alpha_{i,\text{Dry},2}$	$\alpha_{i,\text{Dry},3}$	$\alpha_{i,\text{Wet},1}$	$\alpha_{i,\text{Wet},2}$	$\alpha_{i,\text{Wet},3}$
1	210.36	160.65	116.09	169.14	118.78	52.39
2	134.05	95.75	69.19	107.78	76.99	62.11
3	100.48	71.77	51.86	65.03	37.88	20.80
4	225.12	138.83	94.41	160.77	83.14	39.67
5	117.05	104.91	93.38	87.93	66.57	33.28
6	86.64	64.25	46.26	59.45	43.61	33.60
7	62.87	31.44	24.08	21.67	9.61	6.23
8	19.00	15.26	10.99	5.45	3.89	3.16

Table 13: Terminal value functions for L17

i	$\alpha_{i,\text{Dry},1}$	$\alpha_{i,\text{Dry},2}$	$\alpha_{i,\text{Dry},3}$	$\alpha_{i,\text{Wet},1}$	$\alpha_{i,\text{Wet},2}$	$\alpha_{i,\text{Wet},3}$
1	116.00	116.00	116.00	109.72	83.60	69.53
2	219.43	161.12	111.07	195.00	105.50	54.00
3	94.46	94.46	94.46	71.77	61.50	58.66
4	93.39	93.39	93.39	93.40	85.16	69.54
5	110.35	81.52	81.52	68.99	46.89	33.04
6	59.08	59.08	59.08	55.37	30.41	21.04
7	219.33	140.99	109.82	165.38	95.21	45.69
8	148.92	106.25	99.72	103.69	76.93	39.99
9	62.31	62.31	62.31	61.96	29.68	22.34
10	82.92	68.99	68.99	85.74	52.23	45.49
11	49.65	49.65	49.65	40.08	25.78	25.78
12	57.48	49.17	39.14	39.45	25.95	5.39
13	29.91	23.91	23.91	17.50	6.31	2.39
14	51.56	51.56	51.56	60.41	43.23	33.12
15	49.25	49.25	49.25	49.50	34.69	30.37
16	40.31	36.59	36.59	31.52	31.52	17.79
17	40.28	22.71	11.15	21.91	6.14	1.05

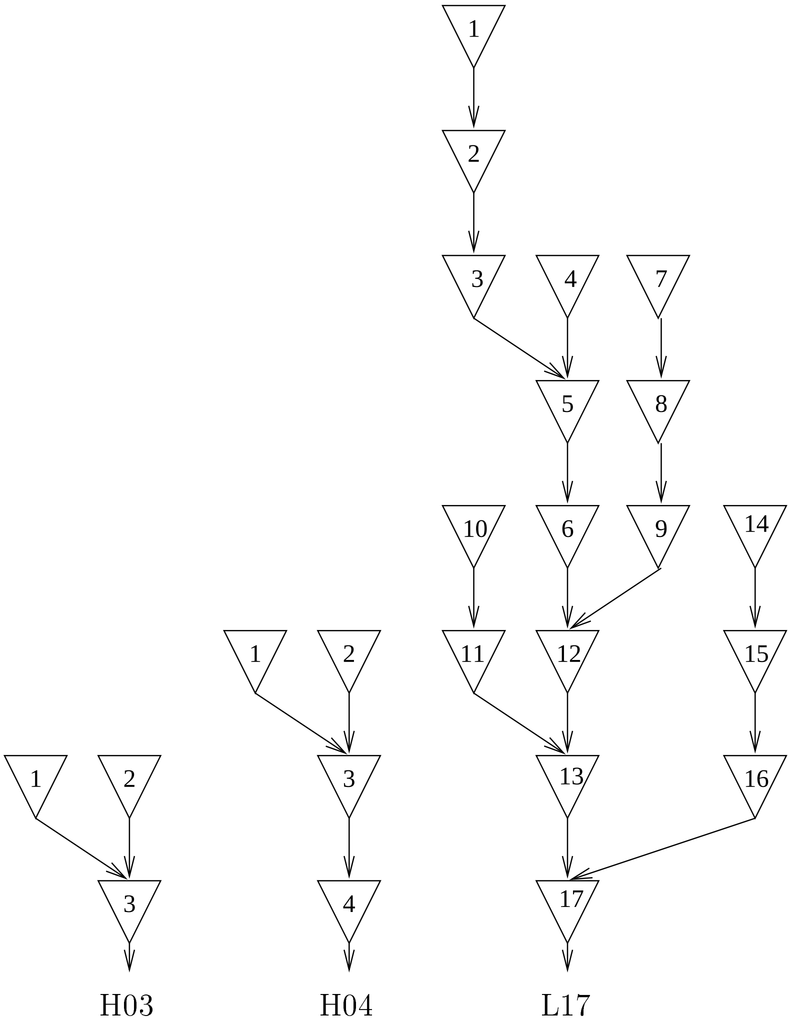

Figure 2: Connectivity of the reservoir networks in test problems H03, H04 and L17