Edinburgh Research Explorer

Modelling the operation of multi-reservoir systems using decomposition and stochastic dynamic programming

Link:
[Link to publication record in Edinburgh Research Explorer](#)

Document Version:
Peer reviewed version

Published In:
Naval Research Logistics

Modelling the operation of multireservoir systems using decomposition and stochastic dynamic programming

T. W. Archibald

University of Edinburgh Management School

K. I. M. McKinnon

School of Mathematics, University of Edinburgh

L. C. Thomas

School of Management, University of Southampton

Abstract

Stochastic dynamic programming models are attractive for multireservoir control problems because they allow non-linear features to be incorporated and changes in hydrological conditions to be modelled as Markov processes. However with the exception of the simplest cases, these models are computationally intractable because of the high dimension of the state and action spaces involved. This paper proposes a new method of determining an operating policy for a multireservoir control problem that uses stochastic dynamic programming, but is practical for systems with many reservoirs. Decomposition is first used to reduce the problem to a number of independent subproblems. Each subproblem is formulated as a low dimensional stochastic dynamic program and solved to determine the operating policy for one of the reservoirs in the system.

Keywords: Multireservoir systems; stochastic dynamic programming; decomposition; principal components.

1   Introduction

The efficient operation of multireservoir systems is an important factor in hydroelectric power generation, domestic and industrial water supply, irrigation programmes and other applications. Stochastic dynamic programming models of such systems are useful because they allow non-linear features to be incorporated and changes in hydrological conditions to be modelled as Markov processes [16, 9]. However for many practical problems these models are computationally intractable due to the high-dimensional state and action spaces involved. The reason for this is that the state space includes variables to model the volume of water stored in each reservoir in the system.

Arvanitidis and Rosing [3], and many authors since, overcome this problem by combining all the reservoirs in the system into one aggregate reservoir and solving a stochastic dynamic programming model of the aggregate reservoir. With this approach it is not generally possible

to capture the individual characteristics of the reservoirs (for example different penalties for flooding at different reservoirs due to variation in land use). Further the relationship between the amount of water to release from the aggregate reservoir and the amount of water to release from the individual reservoirs is not straightforward. For example Turgeon and Charbonneau [13] use a two stage disaggregation method using the decomposition method of [11] and heuristics.

An alternative approach to aggregation is to use a low dimensional representation of the volume of water stored in each reservoir. The idea is that instead of the state space being the product of the discretisation of the volume of water in each reservoir separately, one identifies a number of typical system situations, each of which consists of a record of the volumes of water stored in the reservoirs in the system. These situations are the states of the approximate model and the decision in a period is which state to move to by the start of the next period. Saad and Turgeon [10] use principal component analysis to achieve this. The principal component representation is determined by analysing a sample of observations of the volumes of water stored in the reservoirs. The sample can either be determined from historical data or by simulating the operation of the reservoir system. Neither approach is ideal. Generally reliable historical data is only available for at most 50 years, so with this approach the size of the sample is limited. Also the sample will be biased by the operating policies used in the past and these policies may not be consistent with current objectives and hydrological conditions. Simulation can be used to generate a much larger sample, but again the sample will be biased by the operating policy used.

Pereira and Pinto [8] solve a stochastic linear programming model of a multireservoir system using Benders decomposition. The disadvantage of this approach is that problems with a large number of periods and an adequate model of the weather conditions are computationally intractable [1].

Another approach is to decompose the high dimensional stochastic dynamic program into a number of low dimensional subproblems. Each subproblem includes a detailed representation of one of the reservoirs in the system and an approximate representation of the rest of the system. Turgeon [11, 12] develops decomposition methods for systems of reservoirs in parallel and in series. These methods cannot be applied to a general system of reservoirs and use aggregation as the method of approximation in the subproblems. Archibald et al [2] propose a decomposition method that can be applied to any connected reservoir system. Each subproblem corresponds to a partitioning of the reservoir system into three subsystems and uses a heuristic which assumes that reservoirs in the same subsystem are always equally full. The choice of the partition of the reservoir system is fixed by the method and there may be cases where an alternative partition, perhaps including more than three subsystems, would give better results. Also one can think of situations where operating reservoirs according to the equally full heuristic might not lead to good performance. For example the amount of electricity generated by a hydroelectric station at a reservoir is often assumed to be proportional to the product of the volume stored in the reservoir and the volume of water released from the reservoir. The variation in the generation with the volume stored in the reservoir is known as the head effect. When the head effect

at a reservoir is significant, it might be optimal to keep that reservoir well-filled even when other reservoirs are almost empty. Another example is where, due to differences in land use, the penalty for flooding at one reservoir is much higher than at others. In this case it might be optimal to keep that reservoir relatively empty during periods in which heavy rainfall is expected.

This paper develops a new solution method for multireservoir control problems which combines decomposition and the approach of using a low dimensional representation of the reservoir volumes. Unlike the approach of Archibald et al [2], the new method can be applied to more general reservoir systems; allows arbitrary partitions of the reservoir system to be used; and does not rely on the equally full heuristic for the approximations of the subsystems. The stochastic dynamic programming model is decomposed into a number of independent subproblems — one for each reservoir in the system. For each subproblem the reservoir system is partitioned into a small number of subsystems, one of which consists of a single reservoir — the focus reservoir. The subproblems use a low dimensional representation of the volumes of water stored in the reservoirs in each subsystem. These representations can be determined using experience and expert judgment [15], an operating heuristic [2] or statistical modelling [4]. The solution to a subproblem is used to determine a target release for its focus reservoir.

This solution approach has a number of positive features. Most importantly the solution time only increases quadratically with the number of reservoirs in the system, so problems with many reservoirs can be solved. Individual reservoir volumes and releases are always available, so it is not necessary to approximate the reward generated by the reservoir system in terms of aggregate variables. The operating policy for a reservoir is determined by solving an independent subproblem which includes a detailed model of that reservoir. This approach supports decentralised planning and determines a policy that is sensitive to the individual characteristics of the reservoir (for example storage capacity and penalties for flooding). A subproblem can also be used to estimate the marginal value of the water stored in its focus reservoir. This information is useful to companies operating in a deregulated electricity market, as it helps them to decide when to sell and purchase electricity on spot markets. Finally the optimal policy for each subproblem is a feasible policy for the original model, so the solutions to the subproblems provide direct insight about the solution to the original model.

Section 2 describes a stochastic dynamic programming model of a multireservoir system. Section 3 presents the new solution method in its most general form. Sections 4 and 5 describe techniques that can be used to partition the system and form low dimensional representations of reservoir volumes for the new solution method. Section 6 presents observations based on our computational experiments with the proposed solution method.

2   Modelling multireservoir systems

This section describes a stochastic dynamic programming model of a multireservoir system that can be applied to a range of reservoir networks and many applications. The model is similar to

that studied in [2] and many other papers.

The model applies to a multireservoir system in which: (1) the reservoir network is acyclic; (2) the water released from a reservoir directly enters at most one reservoir; and (3) the flow of water between reservoirs is instantaneous. To allow for the possibility of pump storage in hydroelectric power generation systems, assumptions 1 and 2 would have to be removed. This can be done at the expense of an increase in the number of decision variables and hence the time required to solve the model. Assumption 3 is reasonable provided the time periods considered are large (for example a month), because then the time required for water to flow between reservoirs is often insignificant.

The model has a finite planning horizon consisting of T periods. The reward from the next period is discounted by a factor β where $0 < \beta \leq 1$. There are M reservoirs in the system and the volume of water stored in a reservoir at the beginning of each period must lie between known upper and lower bounds.

The state of the system in a period is comprised of the volumes in the reservoirs at the beginning of the period and the hydrological state. The hydrological state may involve recent rainfall patterns, short-term weather forecasts, long-term weather forecasts or snowpack height depending on the system to be modelled and is represented by a integer between 1 and S.

During each period water from external sources (for example rain or melting snow) flows into each reservoir. As the reservoirs in the system are likely to be geographically clustered, these runoffs are usually correlated with each other. The model captures this correlation by considering R runoff patterns at each period. A runoff pattern specifies the runoff to each reservoir during the period to which it applies. The runoff patterns can be generated from historical data using a stochastic streamflow model similar to those described in [14].

The probability that a particular runoff pattern occurs during a period depends on the hydrological state at the beginning of the period. The persistence of the hydrological state is modelled as a Markov process. The probability that the system is in a particular hydrological state at the start of a period depends on the hydrological state and the runoff pattern in the previous period.

The decision variables are the volumes of water to have stored in the reservoirs at the beginning of the next period. The decisions taken in a period depend on the reservoir volumes and the hydrological state at the beginning of the period and the runoff pattern during the period. The assumption that the decisions depend on the runoff pattern is reasonable because, in practice, the runoff pattern in a period will gradually be revealed as the release for that period is discharged. This means that by the time the release decisions in a period have to be finalised, a great deal is known about the runoff pattern for that period. As a result of assumptions 1 and 2, the volumes of water released from the reservoirs during a period can be determined explicitly from the state of the system and the decision variables in that period. If assumptions 1 and 2 were removed, a 1-period deterministic optimization problem would have to be solved to find the best way to effect the desired change in reservoir volumes, and hence the volumes of water released from the reservoirs.

The water released from the reservoirs generates revenue through the production of electricity, the supply of water or another means. The level of revenue generated is a function of the state of the system at the beginning of the period, the runoff pattern and reservoir releases during the period, and the period. The dependence on the period allows seasonal variations in the price of and demand for the commodity (for example electricity or water) to be modelled. The dependence on the hydrological state and runoff pattern allows uncertainty in the price and demand to be modelled. In this way the model can also capture the correlation between price, demand and weather conditions. The dependence on the reservoir volumes allows the efficiency of the generators to vary with the head of water in the reservoirs for hydroelectric power generation. If the quantity of water released exceeds the capacity of the channels used to discharge water from a reservoir, flooding of the surrounding land may result. The dependence on the individual reservoir releases and the period allows the penalties applied to reflect the use of the land surrounding each reservoir in different seasons.

The future value of the water stored in reservoirs at the end of the planning horizon is a function of the reservoir volumes and the hydrological state at that time. The role of the terminal value function in the model is to prevent the optimal policy from spuriously running the reservoir volumes down at the end of the planning horizon. If the planning horizon is sufficiently long, the terminal value function will not have a significant effect on the first period decisions. Since it is intended that an instance of the model be solved each time a decision has to be made, only the first period decisions will ever be implemented. For this reason the estimate of the terminal value function does not need to be highly accurate.

Notation

Throughout the paper the symbol $'$ is used to denote the state at the end of the period.

β — factor by which rewards are discounted per period.

T, M, S, R — number of periods in the planning horizon, reservoirs in the system, hydrological states and runoff patterns respectively.

t, i, j — index for time periods, reservoirs and runoff patterns respectively.

s, s' — indexes for hydrological states.

I_i — set of reservoirs whose releases flow directly into reservoir i.

$\underline{H}_i^t$, $\overline{H}_i^t$ — lower and upper bounds respectively on the volume in reservoir i at the beginning of period t.

h_i, h_i' — volume of water stored in reservoir i.

$\underline{h}$, $\underline{h}'$ — vectors of reservoir volumes.

$q_{j,i}^t$ — runoff to reservoir i from external sources in period t when runoff pattern j occurs.

$\underline{q}_j^t$ — vector of reservoir runoffs in period t when runoff pattern j occurs.

x_i — amount of water released from reservoir i during a period.

$\underline{x}$ — vector of reservoir releases during a period.

$P_{s,j}^t$ — probability of runoff pattern j occurring during period t when the hydrological state in period t is s.

$Q^t_{(s,j),s'}$ — probability of a transition from hydrological state s and runoff pattern j in period t to hydrological state s' in period $t+1$.

$r^t(\underline{h}, s, j, \underline{x})$ — reward in period t when the reservoir volumes and hydrological state at the beginning of the period are $\underline{h}$ and s respectively, and the runoff pattern and reservoir releases during the period are j and $\underline{x}$ respectively.

$\tau(\underline{h}, s)$ — value of the water stored in the system at the end of the planning horizon when the reservoir volumes are $\underline{h}$ and the hydrological state is s.

$v^t(\underline{h}, s)$ — maximum expected discounted reward from the beginning of period t until the end of the planning horizon when, at the beginning of period t, the reservoir volumes are $\underline{h}$ and the hydrological state is s.

It follows from the definitions above that $v^{T+1}(\underline{h}, s) = \tau(\underline{h}, s)$ and, for $t \leq T$, $v^t(\underline{h}, s)$ satisfies the following optimality equation.

$$v^t(\underline{h}, s) = \sum_{j=1}^{R} P^t_{s,j} \max_{\underline{h}'} \left\{ r^t(\underline{h}, s, j, \underline{x}) + \beta \sum_{s'=1}^{S} Q^t_{(s,j),s'} v^{t+1}(\underline{h}', s') \right\}$$

$$\text{subject to } x_i = h_i + q^t_{j,i} + \sum_{k \in I_i} x_k - h'_i, \tag{1}$$

$$\underline{H}^{t+1}_i \leq h'_i \leq \overline{H}^{t+1}_i \text{ and } x_i \geq 0 \text{ for } 1 \leq i \leq M.$$

The model above can easily be modified to include evaporation or seepage losses for problems in which these are thought to be important. The loss of water due to evaporation or seepage from a reservoir during a period could be any function of the state of the system at the beginning of the period, the runoff pattern and reservoir releases during the period, and the period. The loss of water from reservoir i during period t would have to be allowed for on the right hand side of the water balance equation (1) in the optimality equation above.

The volume of water in each reservoir is a continuous variable so, to solve this optimality equation, discretisation has to be used to determine a finite number of possible values for the volumes of water stored in the reservoirs at the beginning of each period. A common approach (see for example [5, 6]) is to apply a regular discretisation in which Δ equally spaced values between $\underline{H}^t_i$ and $\overline{H}^t_i$ are considered for the volume of water stored in reservoir i at the beginning of period t. All possible combinations of these values are considered to give Δ^M possible values for the volumes of water in the reservoirs at the beginning of period t. The model then has $\Delta^M S$ states per period and Δ^M possible actions in the decision problem for each state, runoff pattern and period. For a reservoir system with 6 or more reservoirs, it is not practical to use the model to determine an operating policy. Chen et al [4] propose a method of overcoming this difficulty using orthogonal array experimental designs. Another approach is to apply decomposition techniques to find an approximate solution to the problem by solving a number of low dimensional subproblems. The next section describes a new solution method based on such an approach.

3   Decomposition method for a multireservoir system

This section describes a decomposition method which creates an independent subproblem for each reservoir in the system. In the subproblem for reservoir f, reservoir f is referred to as the focus reservoir and the set of reservoirs in the system is partitioned into N groups, one of which contains only reservoir f. The subproblem is formulated as a stochastic dynamic program in which the dimensions of the state and action spaces are proportional to the number of groups in the partition. The number of groups in the partition should therefore be limited to at most 4 so that the subproblem can be solved in a reasonable time. However this practical restriction does not have a great impact on the solution method we propose, as we argue in section 4 that the key influences on the operations of the focus reservoir can be modelled using a partition consisting of at most 4 groups.

It is clear that the reservoirs whose releases could flow into the focus reservoir have an important effect on the operation of the focus reservoir. These reservoirs form an obvious group as it is the volume of water entering the focus reservoir that is important for the operation of that reservoir, rather than which reservoir the water came from. So one decomposition would be into three groups consisting of the focus reservoir, the group of reservoirs that feed it and a group consisting of the remaining reservoirs. The partitioning of a system of 6 reservoirs is illustrated in figure 1. The focus reservoir is reservoir 4 which is the only reservoir in the group G_2. The reservoirs in the group G_1 feed the focus reservoir. The other reservoirs in the system are in group G_3. Section 4 describes two systematic approaches to partitioning the reservoir system which can be applied to any reservoir system and create at most 4 groups.

Each group in the partition is examined to determine a number of "typical" situations for the volumes of water stored in the reservoirs in the group at the beginning of each time period. For example we might observe that the reservoirs in a group are always approximately equally full and choose 5 typical situations in which the reservoirs in the group are all empty, all 25% full, all 50% full, all 75% full and all completely full. Alternatively, for a group consisting of two reservoirs, we might observe that in certain periods one reservoir is kept full while the volume of water stored in the other is allowed to vary. We could then choose 5 typical situations for these periods in which the first reservoir is full and the second reservoir is empty, 25% full, 50% full, 75% full and completely full respectively. Section 5 describes two systematic methods of determining typical situations for a group of reservoirs.

Each subproblem can be formulated as a stochastic dynamic program as follows.

Additional notation

N — number of groups in the partition of the reservoir system.

$G_1, \ldots, G_N$ — the groups in the partition of the reservoir system.

g — index for reservoir groups.

A_g^t — number of typical situations for the volumes of water in the reservoirs in G_g at the beginning of period t.

α_g, α'_g — indexes for typical situations in G_g.

$\underline{\alpha}$, $\underline{\alpha}'$ — vectors of indexes for typical situations in all groups.

$L^t(\underline{\alpha})$ — vector of reservoir volumes corresponding to the typical situations defined by $\underline{\alpha}$.

$\nu^t(\underline{\alpha}, s)$ — maximum expected discounted reward from the beginning of period t until the end of the planning horizon when, at the beginning of period t, the reservoir groups are in situations defined by $\underline{\alpha}$ and the hydrological state is s.

It follows that $\nu^{T+1}(\underline{\alpha}, s) = \tau(L^t(\underline{\alpha}), s)$ and, for $t \leq T$, $\nu^t(\underline{\alpha}, s)$ satisfies the following optimality equation.

$$\nu^t(\underline{\alpha}, s) = \sum_{j=1}^{R} P_{s,j}^t \max_{\underline{\alpha}'} \left\{ r^t(\underline{h}, s, j, \underline{x}) + \beta \sum_{s'=1}^{S} Q_{(s,j),s'}^t \nu^{t+1}(\underline{\alpha}', s') \right\}$$

subject to $\underline{h} = L^t(\underline{\alpha})$, $\underline{h}' = L^{t+1}(\underline{\alpha}'), x_i = h_i + q_{j,i}^t + \sum_{k \in I_i} x_k - h_i'$ and $x_i \geq 0$ for $1 \leq i \leq M$,

$$1 \leq \alpha_g \leq A_g^t \text{ and } 1 \leq \alpha_g' \leq A_g^{t+1} \text{ for } 1 \leq g \leq G.$$

The dimensions of the state and action spaces in the formulation of a subproblem are equal to $N + 1$ and N respectively. The actions considered in the formulation of a subproblem are all feasible actions for the original model of the reservoir system. Hence the optimal expected rewards for a subproblem is a lower bound on the optimal expected reward for the original model.

There is a great deal of flexibility with the proposed discretisation approach and care has to be taken to avoid creating subproblems that are computationally intractable or infeasible. For example we could apply a regular discretisation to identify typical situations for the values of water in the reservoirs in a group. However if the group contains a large number of reservoirs, this could result in more typical situations than it is practical to consider when solving the subproblem. In fact if a regular discretisation is employed in every group in the partition, the subproblem is equivalent to the original model with a regular discretisation. If the typical situations do not vary with the period, it is always possible to arrange for the groups to be in the same situations at the next period. However if we allow the typical situations for a group to vary with the period, there may be no way to effect a change from a typical situation at one period to any of the typical situations at the next period. In such cases there could be no feasible solution to the subproblem. This is more likely to happen when the number of typical situations per period is small. Ways of determining typical situations that avoid these problems are discussed in section 5.

The solution to each subproblem specifies an optimal release for each reservoir in each time period. There are several ways in which this data could be used to construct an operating policy for the reservoir system. We propose a method that only uses the optimal release from the focus reservoir in the first period. We refer to this as the target release for the reservoir, denoted by x_i^{target} for reservoir i. Since the target releases are determined independently of each other, they do not necessarily form a feasible first period operating policy for the reservoir system. We propose the following simple method to find actual releases for each reservoir that do form a

feasible first period operating policy and are as close as possible to the target releases.

1. Find a reservoir i for which the actual release has not been set, but actual releases have been set for all reservoirs in I_i.

2. Calculate h_i^{target} = level of reservoir i at the beginning of period 2 assuming the target release for reservoir i and the actual releases for reservoirs in I_i are applied.

3. Set the actual release for reservoir i to be

$$x_i^{target} + \max\{0, h_i^{target} - \overline{H}_i^2\} - \min\{0, \underline{H}_i^2 - h_i^{target}\}.$$

4. Repeat the above steps until actual releases have been set for all reservoirs.

The model would not be used to determine operating policies for subsequent periods. Instead a new instance of the model would be generated to reflect what is known about the reservoir system at the time the decision has to be taken.

4   General methods of partitioning a reservoir system

This section describes two methods of partitioning a reservoir system which are systematic, yield groupings that have a physical interpretation, and can be applied to any reservoir system. An important difference between the two methods is the maximum number of groups in the partitions formed. The first method forms partitions with at most 3 groups while partitions formed by the second method can have up to 4 groups. As the dimension of a subproblem is proportional to the number of groups in the partition, this means that subproblems created using the first method are generally less demanding computationally. However the actions considered in a subproblem created using the first method are all feasible actions in the subproblem with the same focus reservoir created using the second method. Hence the solution to a subproblem created using the first method is a lower bound on the solution of the corresponding subproblem created using the second method, and the solution to the original model derived using the second method will generally be better.

Method 1 — 3-group partition

The partition consists of one group containing the focus reservoir, one group containing all the reservoirs whose releases can reach the focus reservoir, either directly or indirectly, and one group containing all other reservoirs. These groups are referred to as the focus reservoir, the feeder reservoirs and the non-feeder reservoirs respectively. The releases from feeder reservoirs have a significant effect on the amount of water available for release from the focus reservoir. It therefore seems intuitive that it is useful to be able to model the feeder reservoirs separately from the rest of the system. This is the method used in [2] to decompose the problem. Figure 2(a) shows an example of this method for a system of 6 reservoirs. In some cases the partition formed will only have 2 non-empty groups. For example the partition for the subproblem for reservoir 1 in the system in figure 2.

Method 2 — 4-group partition

The partition consists of one group containing the focus reservoir, one group containing all the reservoirs whose releases can reach the focus reservoir, one group containing all the reservoirs that the release from the focus reservoir can reach, and one group containing all other reservoirs. These groups can be thought of as the focus reservoir, the feeder reservoirs, the downstream reservoirs and the unrelated reservoirs respectively. The amount of water available for release from the downstream reservoirs depends on how much water is released from the focus reservoir. It therefore seems likely that the ability to model the downstream reservoirs separately from the rest of the system would be an advantage. This method also allows for the feeder reservoirs to be modelled separately. Figure 2(b) shows an example of this method for a system of 6 reservoirs. In some cases the partition formed will only have 2 or 3 non-empty groups. For example the partitions for the subproblems for reservoirs 6 and 3 respectively in figure 2.

We argue that, in terms of the water storage and releases in the original model, the most significant factors affecting the operations of a reservoir during a period are: the amount of water stored in the reservoir; the amount of water released into the reservoir; and the impact of the water released from the reservoir on the volume of water stored downstream. We have shown that these effects can be modelled using partitions consisting of at most 4 groups. We believe that to improve the operating policies that the method finds, one would be better to use a finer discretisation in the 4-group partition rather than add further groups to the partitions used.

5   Modelling the storage of water in a group of reservoirs

This section describes two methods of determining typical situations to describe the volumes of water in all the reservoirs of a group at the beginning of a period. The aim is to capture enough detail about the variability in the reservoir volumes to give a good approximation without making the subproblems computationally intractable. The methods are also designed to limit the risk of choosing typical situations that result in infeasible subproblems. The idea is that by restricting the state space in this way, we also restrict the action space since we must end up in one of the typical situations at the start of the next time period. The first method uses the equally full heuristic proposed by [2] and the second method uses principal component analysis.

Equally full heuristic

Let $\underline{h}$ denote the volumes of water stored in the reservoirs in $G = \{1, 2, \ldots, m\}$ at the beginning of period t. The reservoirs in G are said to satisfy the equally full heuristic if the ratio of $h_i - \underline{H}_i^t$ to $\overline{H}_i^t - \underline{H}_i^t$ equals a constant α ($0 \leq \alpha \leq 1$) for all $i \in G$. The equally full heuristic uses a one-dimensional approximation to the volumes in the reservoirs of the form $h_i = (1 - \alpha)\underline{H}_i^t + \alpha\overline{H}_i^t$, i.e. all the reservoirs are a fraction α full. We use Δ equally spaced values between 0 and 1 for α to identify typical situations for the volumes in the reservoirs. It is always possible to increase the amount of water released from a reservoir to bring the volume in the reservoir down to its

lower bound. Hence including the case $\alpha = 0$ will eliminate the risk of creating an infeasible subproblem.

Principal component analysis

The concept of principal component analysis is fully described in [7]. It was first used to solve a reservoir control problem by Saad and Turgeon [10] who applied the method to a complete system to reduce the dimension of the stochastic dynamic programming model from 10 to 4. Although Saad and Turgeon [10] note that principal component analysis could be applied separately to subsystems, the method developed in this paper is the first to do so. Basically given a sample of observations of the volumes of water in a group of m reservoirs at the beginning of a period, principal component analysis can be used to find the one-dimensional representation of the reservoir volumes that captures the maximum of the variability in the sample. The idea is equivalent to plotting points in m dimensions, each of which represents typical volumes found in the m reservoirs. One then finds the line which most closely fits these points and the states of the system are points on this line. It seems likely that if it is beneficial to have both full and empty reservoirs in the group at the same time, this representation will yield a better solution than the equally full heuristic.

Principal component analysis can be applied to the group of reservoirs $G = \{1, 2, \ldots, m\}$ as follows. Suppose there are n observations of the volumes of water in the reservoirs in G at the beginning of the period. Let $\underline{z}_i$ be the vector of observations of the volume in reservoir $i \in G$; μ_i be the mean of the observations of the volume in reservoir $i \in G$; $c_{i,j} = \underline{z}_i.\underline{z}_j/n - \mu_i\mu_j$ be the covariance of the observations of the volumes in reservoirs i and $j \in G$; and C be the covariance matrix of the observations of the volumes in reservoirs in G. We then determine the eigenvalues and eigenvectors of C. Let $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_m$ denote these eigenvalues and $\underline{b}_1, \underline{b}_2, \ldots, \underline{b}_m$ denote the corresponding eigenvectors, normalised so that $\underline{b}_i.\underline{b}_i = 1$.

Define B to be the matrix $(\underline{b}_1, \underline{b}_2, \cdots, \underline{b}_m)$ and let $\underline{h}$ denote the volumes in the reservoirs in G at the beginning of the period. If $\underline{\tilde{h}}$ is defined to be $B^{\mathrm{T}}\underline{h}$ (where $^{\mathrm{T}}$ denotes transpose) then, from the orthogonality of the eigenvectors of C, $\underline{h} = B\underline{\tilde{h}}$ — a linear combination of the eigenvectors. Hence both $\underline{h}$ and $\underline{\tilde{h}}$ describe the volumes in the reservoirs in G. It can be shown that the total variation in the reservoir volumes explained by the eigenvector $\underline{b}_i$ is proportional to λ_i. Thus if we allow ourselves only one dimension to describe the volumes in the reservoirs in G, we should choose a linear combination involving the eigenvector $\underline{b}_1$. The variables $\tilde{h}_2, \ldots, \tilde{h}_m$ can be approximated by their mean values $\tilde{\mu}_2, \ldots, \tilde{\mu}_m$ where $\underline{\tilde{\mu}} = B^{\mathrm{T}}\underline{\mu}$. A one-dimensional approximation to the volumes in the reservoirs in G is then given by

$$\underline{h} = \tilde{h}_1\underline{b}_1 + \sum_{i=2}^{m} \tilde{\mu}_i\underline{b}_i. \tag{2}$$

To identify Δ typical situations for the volumes in the reservoirs, we need to choose Δ values for $\tilde{h}_1$ in (2). A simple way to do this is to calculate the maximum and minimum values that $\tilde{h}_1$ can take without violating the upper and lower bounds on the volumes in the reservoirs,

and use Δ equally spaced values between these extremes. Including the extreme values in the discretisation will reduce the risk of the subproblem being infeasible. This approach can be refined by using the two extreme values and $\Delta - 2$ equally spaced values from the range of observed values for $\tilde{h}_1$ in the sample. This range is determined by calculating $\min_{1 \leq j \leq n} \sum_{i=1}^{m} \underline{b}_1(i)\underline{z}_i(j)$ and $\max_{1 \leq j \leq n} \sum_{i=1}^{m} \underline{b}_1(i)\underline{z}_i(j)$ where $\underline{b}_1(i)$ denotes the ith element of the eigenvector $\underline{b}_1$ and $\underline{z}_i(j)$ denotes the jth element of the vector of observations of the volume in reservoir i. This method includes the extreme values to reduce the risk of an infeasible subproblem, but allows a finer discretisation to be used in the region where we anticipate the reservoir volumes to be.

Following the approach of [10], we have used the covariance matrix in the principal component analysis because in reservoir control problems it is important to distinguish between large and small reservoirs. Using the correlation matrix would effectively give all reservoirs the same weight. If the first principal component does not capture a large proportion of the variability of the sample then, subject to the subproblem remaining computationally tractable, further principal components can be used. The possibility that the sample is biased by the operating policies used in its generation has less impact when principal component analysis is applied to subsystems rather than the complete system.

6   Computational Example

The decomposition method presented in this paper has been applied to a hypothetical example of a multireservoir system consisting of 35 reservoirs in series. The planning horizon in this problem consists of 5 periods, and there are 2 possible hydrological states and 3 possible runoff patterns in each period. The reward generated during a period is a non-linear function of the water stored in the reservoirs at the beginning of the period and the water released from the reservoirs during the period. A complete description of the problem can be found at http://homepages.ed.ac.uk/twa/problemData/H35.pdf.

We first used our decomposition method with the 3-group partition method of section 4 and the equally full heuristic approach of section 5 to obtain an approximate solution to the problem. This is equivalent to the solution method proposed by Archibald et al [2]. We then simulated this solution for the 1296 possible future scenarios of hydrological state and runoff pattern. The resulting expected discounted reward was £11,560,405. In this simulation we determined target releases in periods 2, 3, 4 and 5 by interpolating the solutions to the subproblems. In practice we envisage that a new instance of the problem will be solved each time a decision has to be taken. However the computational burden of this is too great for the simulation because it would involve solving 6^{t-1} problems to obtain the target releases during period t, for $t = 2, 3, 4$ and 5.

The simulation was also used to generate a sample of observations of the volumes of water in the reservoirs at the beginning of each period. We used this sample to apply the principal component analysis approach of section 5 with the 3-group partition in our decomposition

method. This gave an expected discounted reward of £11,982,156, representing an improvement of 3.6% on the decomposition method of [2]. As before the simulation was used to generate a sample of observations of the volumes of water in the reservoirs at the beginning of each period, so that the principal component analysis approach could be applied again. When these steps were applied repeatedly, the highest recorded expected discounted reward was £12,449,074, an improvement of 7.7% on the initial solution obtained. We note that during this procedure the expected discounted reward does not increase monotonically with the number of iterations performed.

As a comparison we used the first sample of observations with the principal component analysis approach of [10]. We used the first 3 principal components in the approximation so that the dimension of the stochastic dynamic programming model was the same as in our decomposition method. The expected discounted reward was £12,291,713, approximately 1.3% lower than with our decomposition method. Note that the sample of observations that could be generated by simulation in this case would not yield a significantly different approximation, so this method cannot be applied iteratively.

When the reservoirs are connected in series, as in the problem above, the two methods of partitioning the reservoir network described in section 4 are equivalent. Experiments with differently structured reservoir networks have shown that small improvements can be obtained using the 4-group partition method.

7   Conclusions

We have presented a new solution method for a multireservoir control problem which uses decomposition to reduce the problem to a number of independent subproblems. Each subproblem includes a detailed model of one of the reservoirs in the system and an approximate model of the rest of the system. The approximation partitions the rest of the system into a small number of subsystems and uses low dimensional representations of the volumes of water stored in the subsystems. The choice of the partition and the representations of reservoir volumes in the subsystems is arbitrary, and this creates the opportunity for the insight of experts to be incorporated into the method. Since the approximation does not use aggregation, all decisions considered in the subproblems are feasible in the original model. This is important because the revenue generated by the system depends on the characteristics of individual reservoirs and hence is difficult to approximate in terms of aggregate volumes and decisions.

The number of subproblems that have to be solved is equal to the number of reservoirs in the system. In a subproblem the number of calculations required to determine individual reservoir releases and the revenue generated by the system for each period is proportional to the number of reservoirs in the system. Hence the solution time only increases quadratically with the number of reservoirs in the system, and problems with many reservoirs can be solved in a reasonable time.

A subproblem can provide an estimate of the marginal value of water stored in its focus

reservoir. This is done by solving the subproblem twice — once starting with the actual reservoir volumes and once with a higher or lower quantity of water in the focus reservoir initially. The marginal value of water in the focus reservoir is estimated by the difference between the optimal expected discounted rewards in the two cases. In many practical situations, accurate estimates of the marginal value of water are important. For example comparing the difference in the marginal value of water stored in two reservoirs with the cost of pumping water between the reservoirs, can identify situations where adjusting the distribution of water stored in the system is worthwhile.

It is possible to use the method in an iterative manner to find better solutions to the original model. We can start the process by generating principal component representations of the subsystems using a sample of observations of reservoir volumes generated by simulating the operating policy determined by the equally full heuristic method. (Note that because the releases from the reservoirs are determined independently, the operating policy determined by the equally full heuristic method does not necessarily satisfy the heuristic.) We can generate another sample of observations by simulating the operating policy that these representations yield, and use this to refine the principal component representations of the subsystems. This process can be repeated to find better operating policies for the multireservoir system.

References

[1] T. W. Archibald, C. S. Buchanan, K. I. M. McKinnon, and L. C. Thomas, Nested benders decomposition and dynamic programming for reservoir optimization, *Journal of the Operational Research Society* 50 (1999), 468–479.

[2] T. W. Archibald, K. I. M. McKinnon, and L. C. Thomas, An aggregate stochastic dynamic programming model of multiple reservoir systems, *Water Resources Research* 33 (1997), 333–340.

[3] N. V. Arvanitidis and J. Rosing, Optimal operation of multireservoir systems using a composite representation, *IEEE Transactions on Power Apparatus and Systems* 89 (1970), 327–335.

[4] V. C. P. Chen, D. Ruppert, and C. A. Shoemaker, Applying experimental design and regression splines to high-dimensional continuous-state stochastic dynamic programming, *Operations Research* 47 (1999), 38–53.

[5] E. Foufoula-Georgiou and P. K. Kitanidis, Gradient dynamic programming for stochastic optimal control of multidimensional water resources systems, *Water Resources Research* 24 (1988), 1345–1359.

[6] 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 (1993), 484–500.

[7] K. V. Mardia, J. T. Kent, and J. M. Bibby, *Multivariate analysis*, Academic Press, London, 1979.

[8] M. V. F. Pereira and L. M. V. G. Pinto, Multi-stage stochastic optimization applied to energy planning, *Mathematical Programming* 52 (1991), 359–375.

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

[10] M. Saad and A. Turgeon, Application of principal component analysis to long term reservoir management, *Water Resources Research* 24 (1988), 907–912.

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

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

[13] A. Turgeon and R. Charbonneau, An aggregation/disaggregation approach to long term reservoir management, *Water Resources Research* 34 (1998), 3585–3594.

[14] R. M. Vogel and J. R. Stedinger, The value of stochastic streamflow models in overyear reservoir design applications, *Water Resources Research* 24 (1988), 1483–1490.

[15] S. Walker, P. D. Walsh, and T. Wyatt, Derivation and application of medium term operating policies for the Northern Command Zone system of North West Water (UK). In *Closing the gap between theory and practice*, D. P. Loucks and U. Shamir (Editors), International Association of Hydrological Sciences Publication No 180, 1989, pp. 53–62.

[16] W. W.-G. Yeh, Reservoir management and operations models: A state-of-the-art review, *Water Resources Research* 21 (1985), 1797–1818.
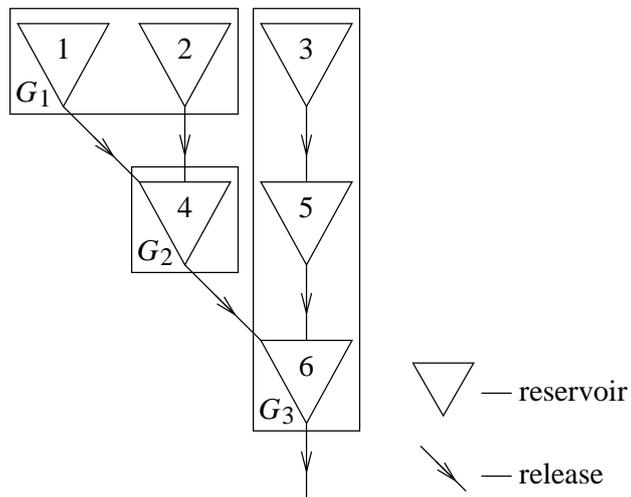
Figure 1: A possible partition for the subproblem for reservoir 4

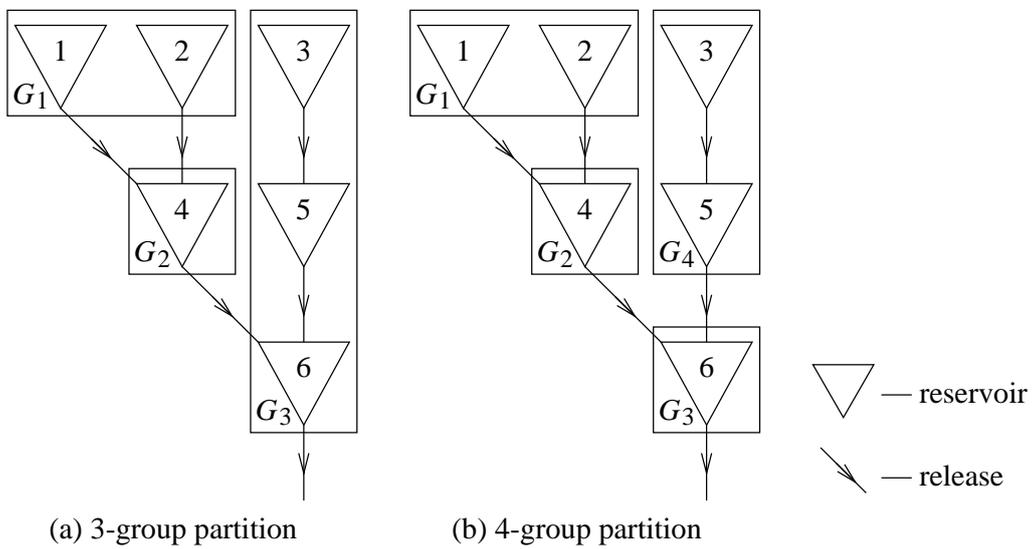(a) 3-group partition   (b) 4-group partition

Figure 2: Applying the methods to partition the system for the subproblem for reservoir 4