An aggregate stochastic dynamic programming model of multiple reservoir systems

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

June 1995

MS-95-028

Also appeared in Department of Business Studies Working Paper Series as No 94/11

For other papers in this series see http://www.maths.ed.ac.uk/preprints

An aggregate stochastic dynamic programming model of multiple reservoir systems

T. W. Archibald

Department of Mathematics and Statistics, University of Edinburgh

K. I. M. McKinnon

Department of Mathematics and Statistics, University of Edinburgh

L. C. Thomas

Department of Business Studies, University of Edinburgh

1 Introduction

Stochastic dynamic programming models are attractive for multiple reservoir control problems because they allow the non-linear features of the problem to be incorporated and the stochastic features of the problem to be modelled as Markov processes. 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. Johnson *et al* [1] have used high-order piecewise polynomial functions to approximate the value function so that a coarse discretisation of the state space can be used. This approach has 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 [2]. 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. Turgeon [4] proposes a method in which the control problem for a system of n reservoirs in series is decomposed into n 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. Turgeon [3] proposes a similar decomposition method for a system of reservoirs in parallel. The disadvantage of these methods is that they cannot be applied to a general network of reservoirs.

This paper proposes a method of determining an operating policy for a general network of reservoirs that is practical for systems with many reservoirs. The operating policy for a particular reservoir is determined by solving a stochastic dynamic programming model of a three reservoir system consisting of that reservoir, an aggregate representation of the reservoirs whose release can reach that reservoir and an aggregate representation of the remainder of the system. The justification for this decomposition is that the optimal decision as to how much water to release from a reservoir is influenced by the amount of water stored in that reservoir, the amount of water flowing into that reservoir, the state of reservoirs affected by the release from that reservoir and the revenue generated from other reservoirs in the system. There are several advantages to this approach to the problem: the solution time increases linearly 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 decentralised and parallel processing can be used; the operating policy for a reservoir is a function of few variables making it easy to implement.

Section 2 describes a general stochastic dynamic programming model of a reservoir system. Section 3 explains the decomposition method we propose. Section 4 deals with the computational aspects of the decomposition method and presents the results of a comparison of the operating policies generated by the full model and the aggregate model.

2 A general model for reservoir control problems

This section describes a finite-horizon discounted stochastic dynamic programming model of a general multiple reservoir control problem. It is assumed that

the reservoir network is connected;

the water released from a reservoir can flow directly into at most one reservoir;

the water released from one of the reservoirs leaves the system;

the flow of water between reservoirs is instantaneous;

there are no evaporation losses.

The reward function can be made to depend upon the level of each reservoir, so that the model can be applied to hydro-electric power generation problems. The details of the model are as follows.

Planning horizon

There are T time periods in the model.

Discount factor

Revenue from the next time period is discounted by a factor β.

Reservoir network

There are M reservoirs in the system. I_i denotes the set of reservoirs whose releases flow directly into reservoir i. The level of reservoir i at the beginning of each period is constrained by lower and upper limits — $\underline{H}_i$ and $\overline{H}_i$ respectively.

Inflow scenarios

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 scenarios are considered. $q_{i,j}^t$ denotes the inflow to reservoir i from external sources during period t in scenario j. The probability of a particular scenario occurring during a period depends upon the hydrological state at the beginning of the period.

State variables

The state of the system during a period is comprised of the reservoir levels and the hydrological state at the beginning of the period. h_i^t denotes the level of reservoir i at the beginning of period t and $\underline{s}^t$ denotes the hydrological state at the beginning of period t. The probability of a particular hydrological state occurring at the next period depends upon the hydrological state and the inflow scenario occurring at this period.

Actions

The decision on the amount of water to release from a reservoir during a period depends upon the reservoir levels and hydrological state at the beginning of the period and the inflow scenario during the period. $x_i^t(\underline{h}, \underline{s}, j)$ denotes the release from reservoir i during period t when the reservoir levels and hydrological state at the beginning of period t are $\underline{h}$ and $\underline{s}$ respectively and the inflow scenario during period t is j. For convenience the variables upon which the releases depend will be omitted from this point on.

Immediate reward

$r^t(\underline{h}, \underline{x})$ denotes the immediate reward in period t when the reservoir levels at the beginning of period t are $\underline{h}$ and the releases from the reservoirs during period t are $\underline{x}$.

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

For a problem involving hydro-electric power generation the functions appearing in the above expression could be interpreted as follows. Z_i^t expresses the head of water in reservoir i during period t as a function of the level in reservoir i at the beginning of period t. G_i^t expresses the amount of water used to generate electricity at reservoir i during period t as a function of the total release from reservoir i during period t. R^t expresses the revenue obtained during period t as a function of the total amount of electricity generated in period t. S_i^t expresses the penalty cost for spillage from reservoir i during period t as a function of the total release from reservoir i during period t.

Terminal value

τ_i expresses the future value of reservoir i as a function of the level of reservoir i and the hydrological state at the end of period T. Hence if $v^{T+1}(\underline{h}, \underline{s})$ denotes the value of having reservoir levels $\underline{h}$ at the end of period T when the hydrological state is $\underline{s}$, then

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

Optimality equation

$v^t(\underline{h}, \underline{s})$, the maximum discounted return from operating the reservoir system from the beginning of period t until the end of period T when the reservoir levels and the hydrological state at the beginning of period t are $\underline{h}$ and $\underline{s}$ respectively, satisfies the following optimality equation.

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

where

$$\begin{aligned} h_i^{t+1} &= h_i^t + q_{i,j}^t + \sum_{k \in I_i} x_k^t - x_i^t \\ \max\left\{0, h_i^t + q_{i,j}^t + \sum_{k \in I_i} x_k^t - \overline{H}_i\right\} &\leq x_i^t \leq h_i^t + q_{i,j}^t + \sum_{k \in I_i} x_k^t - \underline{H}_i \end{aligned}$$

3 An aggregate model for reservoir control problems

This section proposes a method of decomposing a general multiple reservoir control problem into a number of small-scale aggregate subproblems. There is one subproblem for each reservoir in the system and the solution to a subproblem determines an operating policy for the corresponding reservoir. A subproblem consists of three reservoirs in series. These reservoirs will be referred to as the focus, upstream and non-feeder reservoirs. The focus reservoir is one of the reservoirs from the original problem. The upstream reservoir is an aggregate representation of all reservoirs whose releases can reach (directly or indirectly) the focus reservoir. The non-feeder reservoir is an aggregate representation of the rest of the system. The release from the upstream reservoir is discharged into the focus reservoir, the release from the focus reservoir is discharged into the non-feeder reservoir and the release from the non-feeder reservoir leaves the system. Within the non-feeder reservoir a distinction is made between reservoirs that the release from the focus reservoir can reach (downstream reservoirs) and other reservoirs. An assumption about the distribution of water stored in an aggregate reservoir is made so that individual reservoir levels and releases can be calculated from the aggregate level and release. The details of the subproblem corresponding to reservoir i are as follows.

Planning horizon and discount factor

As before there are T time periods in the model and revenue from the next time period is discounted by a factor β.

Aggregate reservoirs

The focus reservoir is reservoir i from the original problem. Let $U_i = I_i \cup \{\cup_{k \in I_i} U_k\}$ denote the set of reservoirs in the upstream reservoir. The assumptions about the reservoir network ensure that there exists an order in which the sets U_i $1 \leq i \leq M$ can be calculated explicitly. Let N_i denote the set of reservoirs in the non-feeder reservoir. Let $D_i = \{k : i \in U_k\}$ denote the set of reservoirs that the release from the focus reservoir can reach. If $U_i = \emptyset$ or $N_i = \emptyset$ then the subproblem consists of only two reservoirs in series. Let U and N be subscripts denoting the upstream and non-feeder reservoirs respectively.

Figure 1 illustrates the application of the aggregation technique to a six reservoir system. The focus reservoir is 4, $U_4 = \{1, 2\}$, $N_4 = \{3, 5, 6\}$ and $U_4 = \{6\}$.

Aggregate limits and inflow scenarios

Lower and upper limits on the levels of the aggregate reservoirs and inflows to the aggregate reservoirs are defined as follows.

$$\underline{H}_{\mathrm{U}} = \sum_{k \in U_i} \underline{H}_k, \quad \overline{H}_{\mathrm{U}} = \sum_{k \in U_i} \overline{H}_k, \quad q^t_{\mathrm{U},j} = \sum_{k \in U_i} q^t_{k,j} \quad j = 1, \ldots, Q$$
$$\underline{H}_{\mathrm{N}} = \sum_{k \in N_i} \underline{H}_k, \quad \overline{H}_{\mathrm{N}} = \sum_{k \in N_i} \overline{H}_k, \quad q^t_{\mathrm{N},j} = \sum_{k \in N_i} q^t_{k,j} \quad j = 1, \ldots, Q$$

State variables

The state of the system during a period is comprised of the level of the upstream, focus and non-feeder reservoirs and the hydrological state at the beginning of the period. For period t these are denoted by h^t_{U}, h^t_i, h^t_{N}, and $\underline{s}^t$ respectively.

Actions

The decision to be taken during a period is how much water to release from the upstream, focus and non-feeder reservoirs. This decision depends on the level of the three reservoirs and the hydrological state at the beginning of the period, and on the inflow scenario during the period. For period t these actions are denoted by x^t_{U}, x^t_i and x^t_{N} respectively.

Distribution of water within aggregate reservoirs

It is assumed that at the beginning of each period the working water (i.e. $h_k - \underline{H}_k$) in an aggregate reservoir is distributed between the reservoirs in the aggregation in proportion to the working capacity (i.e. $\overline{H}_k - \underline{H}_k$) of the individual reservoirs. Define the following constants.

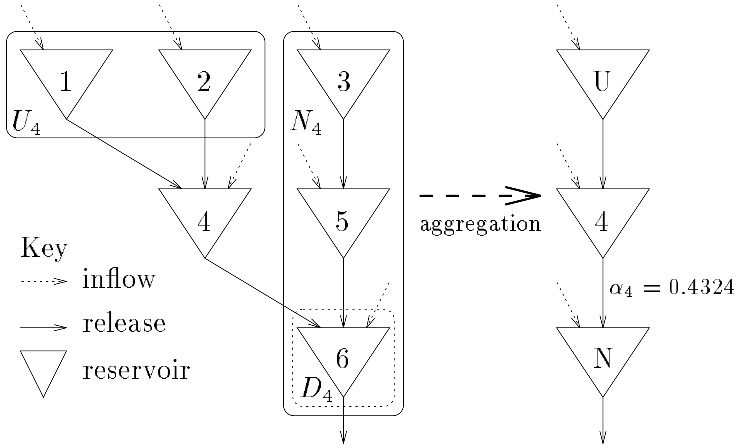

k	$\overline{H}_k$	$\overline{H}_k$	f_k	F_k
1	300	50	0.5000	0.5000
2	250	0	0.5000	0.5000
3	500	0	0.6250	0.6250
4	75	25	—	—
5	50	0	0.0625	0.6875
6	300	50	0.3125	1.0000
U	550	50	—	—
N	850	50	—	—

Figure 1: An example of the aggregation technique

$$f_k = \frac{\overline{H}_k - \underline{H}_k}{\overline{H}_{\mathrm{U}} - \underline{H}_{\mathrm{U}}} \text{ if } k \in U_i, \quad f_k = \frac{\overline{H}_k - \underline{H}_k}{\overline{H}_{\mathrm{N}} - \underline{H}_{\mathrm{N}}} \text{ if } k \in N_i, \quad F_k = f_k + \sum_{\ell \in I_k \setminus \{i\}} F_\ell \text{ for } k \neq i$$

f_k is the proportion of working water in an aggregate reservoir that is stored in reservoir k.

For $k \in U_i$, $F_k x_{\mathrm{U}}$ is the amount of water that must be released from reservoir k so that the release from the upstream reservoir is x_{U} and the distribution of working water in the upstream reservoir at the beginning of the next period is as required. Hence a unit release from the upstream reservoir gives rise to a total release of $\sum_{k \in U_i} F_k$ from individual reservoirs.

The release from the focus reservoir is discharged into the non-feeder reservoir. In the original problem water released from the focus reservoir can be used a further $|D_i|$ times (where $|D_i|$ denotes the number of elements in D_i). If we apply the method used in the upstream reservoir to the non-feeder reservoir, a unit release from the non-feeder reservoir would give rise to a total release of $\sum_{k \in N_i} F_k$ from individual reservoirs. Hence, when $\sum_{k \in N_i} F_k > |D_i|$, water released from the focus reservoir could subsequently be used more times than in the original problem and, when $\sum_{k \in N_i} F_k < |D_i|$, the full potential of water released from the focus reservoir could not be realised. The first problem is addressed by scaling the focus release by

$$\alpha_i = \min\left\{1, \frac{|D_i|}{\sum_{k \in N_i} F_k}\right\}.$$

(We do not consider $\alpha_i > 1$ because then the amount of water discharged into the non-feeder reservoir would be greater than the amount released from the focus reservoir.) The second problem is addressed by allowing the focus release to cascade through the downstream reservoirs as follows. If $x_{\mathrm{N}}^t \geq \alpha_i x_i^t$ then the focus release is allowed to cascade through the downstream reservoirs and in addition $x_{\mathrm{N}}^t - \alpha_i x_i^t$ is released from the water stored in the non-feeder reservoir. F_k relates this additional release to releases from the individual reservoirs in the non-feeder reservoir. If $x_{\mathrm{N}}^t < \alpha_i x_i^t$ then part of the focus release is allowed to cascade through the downstream reservoirs and no water is released from the other reservoirs in the non-feeder reservoir. Hence the release from the non-feeder reservoir can be divided into two components as follows.

$$x^t_{\mathrm{N},a} = \min\left\{x^t_i, \frac{x^t_{\mathrm{N}}}{\alpha_i}\right\}, \quad x^t_{\mathrm{N},b} = \max\left\{0, x^t_{\mathrm{N}} - \alpha_i x^t_i\right\}$$

The table in figure 1 shows the values of f_k and F_k for one problem. Figure 1 can also be used to illustrate the importance of the scaling applied to the focus release. Suppose no scaling is applied. A unit of water that is released from the focus reservoir at this period may be released from the non-feeder reservoir at the next period, giving rise to releases of 0.6250 from reservoir 3, 0.6875 from reservoir 5 and 1.0000 from reservoir 6. Hence the water released from the reservoir 4 has been used a further 2.3125 times. In the original problem water released from reservoir 4 can subsequently be used at most once (via a release from reservoir 6). If the focus release is scaled by $\alpha_4 = 0.4324$ and released from the non-feeder reservoir at the next period, it gives rise to releases of 0.2703 from reservoir 3, 0.2973 from reservoir 5 and 0.4324 from reservoir 6. Hence the water released from reservoir 4 is only used one more time.

Immediate reward

Under our assumptions, the immediate reward in period t when the reservoir levels at the beginning period t are h_{U}, h_i and h_{N}, and the releases from the reservoirs during period t are x_{U}, x_i and x_{N} can be calculated as follows.

$$\begin{aligned}
r^t(h_{\mathrm{U}}, h_i, h_{\mathrm{N}}, x_{\mathrm{U}}, x_i, x_{\mathrm{N}}) = R^t \Bigg(& \sum_{k\in U_i} Z^t_k \left(\underline{H}_{\mathrm{k}} + f_k(h_{\mathrm{U}} - \underline{H}_{\mathrm{U}})\right) G^t_k(F_k x_{\mathrm{U}}) \\
& + Z^t_i(h_i)G^t_i(x_i) + \sum_{k\in N_i\setminus D_i} Z^t_k \left(\underline{H}_{\mathrm{k}} + f_k(h_{\mathrm{N}} - \underline{H}_{\mathrm{N}})\right) G^t_k(F_k x_{\mathrm{N},b}) \\
& + \sum_{k\in D_i} Z^t_k \left(\underline{H}_{\mathrm{k}} + f_k(h_{\mathrm{N}} - \underline{H}_{\mathrm{N}})\right) G^t_k(x_{\mathrm{N},a} + F_k x_{\mathrm{N},b})\Bigg) \\
& - \sum_{k\in U_i} S^t_k(F_k x_{\mathrm{U}}) - S^t_i(x_i) - \sum_{k\in N_i\setminus D_i} S^t_k(F_k x_{\mathrm{N},b}) - \sum_{k\in D_i} S^t_k(x_{\mathrm{N},a} + F_k x_{\mathrm{N},b})
\end{aligned}$$

Terminal value

Define

$$\begin{aligned}
&\nu^{T+1}(h_{\mathrm{U}}, h_i, h_{\mathrm{N}}, \underline{s}) = \\
&\quad \sum_{k\in U_i} \tau_k \left(\underline{H}_{\mathrm{k}} + f_k(h_{\mathrm{U}} - \underline{H}_{\mathrm{U}}), \underline{s}\right) + \tau_i(f_k h_{\mathrm{N}}, \underline{s}) + \sum_{k\in N_i} \tau_k \left(\underline{H}_{\mathrm{k}} + f_k(h_{\mathrm{N}} - \underline{H}_{\mathrm{N}}), \underline{s}\right)
\end{aligned}$$

to be the value of having levels h_{U}, h_i and h_{N} in the upstream, focus and non-feeder reservoirs at the end of period T when the hydrological state is $\underline{s}$.

Optimality equation

$\nu^t(h_{\mathrm{U}}, h_i, h_{\mathrm{N}}, \underline{s})$, the maximum discounted return from operating the aggregate model from the beginning of period t until the end of period T when the levels in the upstream, focus and non-feeder reservoirs and the hydrological state at the beginning of period t are h_{U}, h_i, h_{N} and $\underline{s}$ respectively, satisfies the following optimality equation.

$$\nu^t(h^t_{\mathrm{U}}, h^t_i, h^t_{\mathrm{N}}, \underline{s}^t) =$$
$$\sum_{j=1}^{Q} \Pr(j|\underline{s}^t) \max_{x^t_{\mathrm{U}}, x^t_i, x^t_{\mathrm{N}}} \left\{ r^t(h_{\mathrm{U}}, h_i, h_{\mathrm{N}}, x_{\mathrm{U}}, x_i, x_{\mathrm{N}}) + \beta \sum_{\underline{s}^{t+1}} \Pr(\underline{s}^{t+1}|\underline{s}^t, j)\nu^{t+1}(h^{t+1}_{\mathrm{U}}, h^{t+1}_i, h^{t+1}_{\mathrm{N}}, \underline{s}^{t+1}) \right\}$$

subject to

$$\begin{aligned}
h^{t+1}_{\mathrm{U}} &= h^t_{\mathrm{U}} + q^t_{\mathrm{U},j} - x^t_{\mathrm{U}} \\
h^{t+1}_i &= h^t_i + q^t_{i,j} + x^t_{\mathrm{U}} - x^t_i \\
h^{t+1}_{\mathrm{N}} &= h^t_{\mathrm{N}} + q^t_{\mathrm{N},j} + \alpha_i x^t_i - x^t_{\mathrm{N}} \\
\max\left\{0, h^t_{\mathrm{U}} + q^t_{\mathrm{U},j} - \overline{H}_{\mathrm{U}}\right\} &\leq x^t_{\mathrm{U}} \leq h^t_{\mathrm{U}} + q^t_{\mathrm{U},j} - \underline{H}_{\mathrm{U}} \\
\max\left\{0, h^t_i + q^t_{i,j} + x^t_{\mathrm{U}} - \overline{H}_i\right\} &\leq x^t_i \leq h^t_i + q^t_{i,j} + x^t_{\mathrm{U}} - \underline{H}_i \\
\max\left\{0, h^t_{\mathrm{N}} + q^t_{\mathrm{N},j} + \alpha_i x^t_i - \overline{H}_{\mathrm{N}}\right\} &\leq x^t_{\mathrm{N}} \leq h^t_{\mathrm{N}} + q^t_{\mathrm{N},j} + \alpha_i x^t_i - \underline{H}_{\mathrm{N}}
\end{aligned}$$

An operating policy for the original problem

We propose a straightforward method to determine an operating policy for the original problem from the solutions of the aggregate subproblems. Given the level of water in all M reservoirs and the hydrological state at the beginning of a period, and the inflow scenario occurring during the period, 'target' releases for the reservoirs can be determined from the solutions of the subproblems. The algorithm to determine the actual reservoir releases is as follows.

If $I_i = \emptyset$ then the actual release from reservoir i is set equal to the target release.

While there exists a reservoir for which the actual release has not been set:

> select a reservoir k such that the actual releases from all reservoirs in I_k have been set (this is always possible because of the assumptions about the reservoir network);
>
> calculate h^+_k, the level of reservoir k at the next period assuming that the target release for reservoir k is applied;
>
> set the actual release for reservoir k equal to
>
> $$\text{target release} + \max\{0, h^+_k - \overline{H}_k\} - \min\{0, \underline{H}_k - h^+_k\}.$$

Table 1: The computational complexity of the full and the aggregate models

Model	Problems	States per problem	Inflow scenarios	Actions per state	Action evaluations
Full	1	$K\Delta^M$	Q	up to $(\Delta+1)^M$	$KQ\Delta^M(\Delta+1)^M$
Aggregate	M	up to $K\Delta^3$	Q	up to $(\Delta+1)^3$	$KQ\Delta^3(\Delta+1)^3$

4 A comparison of the full and aggregate models

This section compares both the computational complexity and the quality of the solution of the aggregate model with those of the full model. The comparison of the computational complexity is described for a general reservoir system. The comparison of the quality of the solution uses the test problem from the paper by Turgeon [4].

Discretisation

The state space and action spaces are discretised for computational purposes in the following way. Δ equally spaced values between $\underline{H}_k$ and $\overline{H}_k$ are considered for the level of water in reservoir k at the beginning of a period. The only releases from reservoir k considered are a zero release (in which case linear interpolation is used to approximate the expected future return) and those releases which ensure that the level in reservoir k at the beginning of the next period is equal to one of the Δ values in the discretisation of the level of reservoir k. Thus as many as $\Delta + 1$ releases from a reservoir can be considered. This method of discretisation is applied to both the full and the aggregate models.

Computational complexity

We assume that the computational cost of evaluating an action in the full model is approximately the same as in the aggregate model. This is a reasonable assumption because, while the calculation of the immediate reward requires more work in the aggregate model, the calculation of the expected future cost requires more work in the full model. Let K denote the number of hydrological states. Table 1 shows the complexity of the state and action spaces in the two models. The final column shows that while the complexity of the full model grows exponentially with the number of reservoirs, the complexity of the aggregate model grows linearly. As a result if $\Delta = 10$, the time required to solve a 4 reservoir problem using the full model is sufficient to solve a problem with up to 110 reservoirs using the aggregate model.

Quality of solution

In [4] a hydro-electric power generation model is used as a test problem in a comparison between a decomposition method and stochastic dynamic programming. The planning horizon is one year and the periods are months. The problem has four reservoirs connected in series. There is one hydrological state and five inflow scenarios. The limits on the levels of the

Table 2: Lower and upper limits on reservoir levels

i	1	2	3	4
$\underline{H}_i$	0	0	0	0
$\overline{H}_i$	9628	570	50	3420

reservoirs are shown in table 2. Further details can be found in [4]. We have solved this problem with $\Delta = 11$ using both the full model and the aggregate model. We tested the solutions using 60 sets of 12 monthly inflows with a range of initial conditions and found that the difference in the discounted 12 period return was never greater than 1%. While the policy determined by the full model is optimal in terms of the expected discounted 12 period return, a different policy may result in better discounted 12 period returns for some future scenarios. In roughly 30% of the future scenarios we considered, the policy determined by the aggregate model yields a better discounted 12 period return than the policy determined by the full model. This is a further indication that the policy determined by the aggregate model is a good alternative to the policy determined by the full model. In many cases only the first period actions will be implemented. Future actions will be determined by solving the model for a data set that has been updated to reflect changes in the conditions. For this reason the first period actions are very important. In more than 90% of the future scenarios we considered, the full and aggregate models choose the same actions for the first period. Tables 3, 4 and 5 show the actions chosen by the full and aggregate models in each period for 3 possible future scenarios. In each case there is a very good agreement between the two models throughout the year. The examples shown in tables 3 and 4 were designed to test the robustness of the aggregate model to the distribution of water between the reservoirs. In both cases one reservoir

Table 3: The actions chosen by the full and aggregate models when the reservoir inflows are low throughout the year and the initial reservoir levels are 9628, 380, 17 and 0.

Month	Releases from full model				Releases from aggregate model			
Oct	437	409	430	0	437	409	480	0
Nov	376	858	905	0	376	733	780	6
Dec	1287	1064	1105	686	1287	1189	1229	953
Jan	1175	1482	1508	1819	1175	1482	1458	1547
Feb	1093	1426	1492	2710	1093	1426	1442	2755
Mar	2035	1562	1576	3103	1072	1169	1216	3038
Apr	1010	1426	1432	2526	1010	1027	1049	1754
May	0	216	230	406	120	165	180	356
Jun	0	182	243	963	0	182	193	378
Jul	0	0	31	0	0	0	84	0
Aug	0	0	7	0	0	0	4	0
Sep	0	0	33	0	0	0	33	0

Table 4: The actions chosen by the full and aggregate models when the reservoir inflows are moderate throughout the year and the initial reservoir levels are 0, 190, 33 and 3420.

Month	Releases from full model				Releases from aggregate model			
Oct	0	0	94	1392	0	0	94	1392
Nov	0	212	301	1349	0	212	301	1349
Dec	0	185	247	979	0	185	247	979
Jan	1286	1462	1503	1976	1286	1462	1503	1976
Feb	1168	1301	1376	2703	1168	1301	1345	2622
Mar	208	691	673	3150	359	842	845	2905
Apr	180	290	354	814	187	297	322	1219
May	0	265	354	1402	0	265	352	1079
Jun	0	446	596	2360	0	446	596	1775
Jul	0	90	217	0	0	0	176	0
Aug	0	20	73	0	0	111	163	552
Sep	0	305	408	1180	0	305	408	1615

Table 5: The actions chosen by the full and aggregate models when the reservoir inflows are high throughout the year and the initial reservoir levels are 4814, 285, 25 and 1710.

Month	Releases from full model				Releases from aggregate model			
Oct	0	209	350	597	0	209	350	597
Nov	0	456	587	2130	0	456	587	2130
Dec	675	859	943	1932	675	859	943	1932
Jan	1397	1557	1611	2247	1397	1557	1611	2247
Feb	1243	1346	1416	2852	1243	1346	1416	2830
Mar	1222	1545	1592	2998	1222	1545	1543	3113
Apr	1289	1410	1450	2955	1289	1410	1500	2959
May	0	770	898	3160	0	827	940	2866
Jun	0	710	933	3059	0	710	948	2865
Jul	0	291	487	68	0	286	482	512
Aug	205	639	851	2894	205	498	671	2714
Sep	1107	1483	1617	3647	1107	1572	1744	3774

is initially empty and another is initially full.

Parallel processing

Both models offer considerable scope for parallel computing and the implementations we used exploited this to some extent. In the implementation of the full model, the optimal releases for each inflow scenario in a period were determined in parallel. In the implementation of the aggregate model, the subproblems were solved in parallel, and within each subproblem, the optimal releases for each inflow scenario in a period were determined in parallel. The elapsed times for these implementations of the full and aggregate models were around 22 hours and 6 minutes respectively.

References

1. S.A. Johnson, J.R. Stedinger, C.A. Shoemaker, Y. Li and A. Tejada-Guibert (1993), "Numerical solution of continuous state dynamic programs using linear and spline interpolation", *Operations Research* **41(3)** 484–500.

2. M. Saad, A. Turgeon, P. Bigras and R. Duquette (1994), "Learning disaggregation technique for the operation of long-term hydroelectric power systems", *Water Resources Research* **30(11)** 3195–3202.

3. A. Turgeon (1980), "Optimal operation of multireservoir power systems with stochastic inflows", *Water Resources Research* **16(2)** 275–283.

4. A. Turgeon (1981), "A decomposition method for the long term scheduling of reservoirs in series", *Water Resources Research* **17(6)** 1565–1570.