

MULTIOBJECTIVE OPERATIONS OF RESERVOIR SYSTEMS VIA GOAL PROGRAMMING

By Elizabeth A. Eschenbach¹, Member ASCE, Timothy Magee², Edith Zagona³, Member ASCE, Morgan Goranflo⁴, Member ASCE, Richard Shane⁵, Member ASCE

ABSTRACT

Many authors have documented the minimal use of optimization in practical day to day multipurpose reservoir operations. The RiverWare decision support system RiverWare is a flexible general river basin modeling tool that allows water resources engineers to both simulate and optimize the management of multipurpose reservoir systems for daily operations. This paper describes RiverWare's optimization capabilities and its use by Tennessee Valley Authority operations schedulers. Input data requirements include: (1) physical and economic characteristics of the system, (2) prioritized policy goals, and (3) parameters for automatic linearization. RiverWare automatically generates and efficiently solves a multiobjective, pre-emptive linear goal programming formulation of a reservoir system. An advanced feature of RiverWare is that both the physical model of the river basin and the operating policy are defined and easily modified by the modeler through an interactive graphical user interface. Any modifications are automatically incorporated into the linear preemptive goal program. RiverWare's combination of detailed system representation, policy expression flexibility, and computational speed make it suitable for use in routine daily scheduling of large complex multiobjective reservoir systems.

1. Assistant Professor, Environmental Resources Engineering, Humboldt State University, Arcata CA, 95521,
Phone: 707-826-4348; Fax 707-826-3616; email: eae1@axe.humboldt.edu

2. Operations Research Analyst, Center for Advanced Decision Support for Water and Environmental Systems (CADSWES), University of Colorado, Boulder, CO 80309

3. Water Resources Engineer, CADSWES, University of Colorado, Boulder, CO 80309

4. Senior Consultant, River System Operations, Tennessee Valley Authority, 400 West Summit Hill Dr.,
Knoxville, TN 37918

5. Technical Specialist, River System Operations, Tennessee Valley Authority, 400 West Summit Hill Dr.,
Knoxville, TN 37918

Key Words: Water Resources Management, Multiple Objective Analysis, Optimization, Operations, Hydroelectric Resources, Decision Support System, Tennessee Valley Authority.

INTRODUCTION

River basin management has become increasingly more complex and more dynamic due to growing conflicts among competing objectives. The challenges include water quality, endangered species habitat preservation, and various recreational uses in addition to the traditional objectives of flood control, water supply, navigation, and hydropower production. Furthermore, the deregulation of the power industry sharpens the need for maximizing hydropower benefits. Reservoir managers must improve the management of existing resources and must adapt quickly to changing objectives and requirements.

These management requirements translate into a need for river basin modeling tools that provide prescriptive results as well as allow the decision maker to easily modify operating policy and physical and economic characteristics of a river basin. Due to the complexity of representing a multiobjective reservoir system, most models in use have been descriptive (Wurbs 1996). In addition, models developed for a particular river basin have been "hard-wired" to represent operating policies and physical and economic characteristics. These hard-wired models have become costly and difficult to maintain, as they are not easily adapted to changing objectives and requirements.

Optimization models can often provide prescriptive results. Summaries and reviews of multiobjective reservoir optimization have been published by several authors, such as Yeh (1985), Wurbs (1993) and Labadie (1997), all of whom have commented on the continuing gap between research and practical application. This gap has been the topic of entire symposia (Loucks and Shamir 1989) and is often attributed to models being overly simplified, rendering them unrealistic for operational scheduling (Wurbs 1996). In addition, optimization may be difficult to use in operations because it requires optimization experts to setup, run, and/or interpret model results.

General and modular optimization tools are currently used in daily operations of multipurpose reservoir systems, but these tools have limited prescriptive capability in terms of the complexity of policy objectives that they can incorporate and in their ability to optimize non-hydropower objectives. Hydrosoft (Robitaille et al. 1995) and VISTA (Allen et al. 1996) are modular tools that must be adapted by programmers in order to model new basins. These modular tools optimize hydropower using successive linear programming, but do not optimize non-hydropower objectives. Hydrosoft, VISTA and general river basin modeling tools, for example HEC-5 (HEC-5 1998), MODSIM (Labadie 1995), WEAP (WEAP 1998) and AGUATOOL (Andreu et al 1996), allow predefined policy constraints, such as the preferred sources of water supply or the relative importance of rule curves to be prioritized. However, these constraints cannot be formulated as objectives; hence, decision-makers cannot get a "best" solution when these policy constraints are not met.

Pre-emptive goal programming (Can and Houck 1984; Loganathan and Bhattacharya 1990) is a multiobjective optimization method that allows the flexible expression of policy constraints as objectives. Unlike many multiobjective optimization approaches, this method avoids the practical problems associated with assigning and justifying values of relative weights (Schultz 1989). The Tennessee Valley Authority's (TVA) previous experience with pre-emptive goal programming as part of a weekly planning model (Shane et al. 1988) also indicates the potential for practical success of the goal programming approach.

RiverWare optimization uses pre-emptive goal programming and is part of a general basin modeling package (Zagona et al. 1998). Through an interactive graphical user interface, operations modelers can easily express and change a wide range of policy constraints and objectives. The model can be run and the results can be understood without optimization expertise. RiverWare automatically translates the constraints and objectives into a pre-emptive goal program-linear program formulation (GP/LP). Appropriate physical constraints for the basin model are automatically added to the formulation and nonlinear functions are translated to linear or piecewise linear approximations. Automating the optimization formulation process necessitates that the analysis and debugging of the model also be automated for the non-expert modeler. RiverWare provides checks and utilities to facilitate this process, such as infeasibility prevention and verifying convexity of functions.

The pre-emptive goal programming-linear programming approach was chosen for three main reasons: (1) Deterministic optimization is acceptable given the relatively short time horizon for operational modeling, (2) GP/LP can model the multiple objectives and physical aspects of reservoir systems in a sufficiently realistic manner, and (3) GP/LP is sufficiently efficient and robust to be used in daily operations. Other multiobjective optimization approaches could improve on one of these aspects at the expense of the others, but GP/LP strikes a balance appropriate for operational scheduling.

This paper describes the flexible modeling characteristics of RiverWare's optimization component as well as outlines the practical modeling benefits experienced by TVA. First, an overview of the RiverWare optimization DSS is presented, focusing on the graphical user interface (GUI). The automatic generation of the goal program is then explained, focusing on how the policy expressions are "translated" into code for the linear goal program. Next, a description is provided of how TVA has used RiverWare optimization in daily scheduling of its multipurpose reservoir system. Finally, several hypothetical examples illustrate how flexible policy expression can be used to study operational alternatives.

RIVERWARE OPTIMIZATION DECISION SUPPORT SYSTEM

Decision support systems “help decision makers utilize data and models to solve unstructured problems” (Sprague and Carlson 1982) by creating a software environment that allows them to manage models and data through a user interface. This section provides an overview of RiverWare’s optimization DSS by reviewing the GUI tools that facilitate: (1) specifying the physical and economic models, (2) describing operating policies in terms of pre-emptive goals, and (3) entering parameters for translating constraints and objectives into linear expressions.

Physical Process Model

A model of the reservoir system is created using the GUI. Features of the river basin are represented by icons, as shown in Figure 1, and are modeled by corresponding objects in RiverWare. Icons, selected from a palette, create storage reservoirs, level power reservoirs, sloped storage power reservoirs, pumped storage reservoirs, river reaches, canals, etc. The topology of the system is established by graphically linking appropriate data structures on the objects; for example, the outflow of an upstream river reach is linked to the inflow of a downstream reservoir.

Each river basin object contains algorithms to model the physical processes for that object. Data required for the physical process models are stored on the object and accessed through the icon. Thus, by clicking on a power reservoir icon, the data structure containing the table of the storage-elevation relationship for that reservoir can be viewed or edited. Data not required by the process models on the object, such as operational guide curves used in policy statements, are stored on data objects. Data can be easily imported to and exported from the RiverWare DSS through a data management interface (DMI), which tailors communication with an organization’s databases, external models, reporting facilities, etc. Zagona et al. (1998) provide additional information on building models using RiverWare.

RiverWare provides both optimization and simulation solvers for the basin network. The simulator models all processes in complete nonlinear form. Specific methods for modeling processes—such as hydropower and tailwater elevation—are selected by the modeler, while general processes—such as mass balance of reservoirs—are included in the object’s default behavior. For optimization, each object automatically adds physical process constraints to the GP formulation. Processes that are modeled include mass balance for all reservoirs, reservoir routing in sloped reservoirs, regulated and unregulated spill, turbine capacity and efficiency, river reach routing, and bi-directional gravity canal flow. These physical constraints are summarized in Appendix III. Some processes are represented in the optimization only when introduced by the operating policy constraints and objectives. For example, if power is expressed in a policy constraint, then the appropriate constraints are automatically added to the formulation.

In this case, the object knows how to translate the policy into linear expressions. This translation or linearization process is described below.

Constraint Editor and Expression Language

One innovative component of RiverWare is the ability to easily express and modify operations policy. Constraints and objectives are expressed in terms of physical variables—such as pool elevation, flows or spill, or in terms of economic variables—such as net replacement value, future value of used energy, spill cost, and the cost of alternative power resources. Figure 2 shows three policy goals as they appear in the RiverWare optimization constraint editor. The first policy goal is to maintain the pool elevation level within the daily operations levels for Reservoir 4 and Reservoir 5. The second policy goal is to maintain a minimum flow from Reservoirs 3, 4, 6 and 8. The third policy goal is to maximize net avoided cost. Note that the first two goals are constraints on the system, while the third goal is an objective.

Operating policies are input using the optimization constraint editor (Figure 2) and the expression editor (Figure 3). Through the graphical constraint editor, the modeler creates, names, and prioritizes goals and selects the type of each goal, MaxMin, Summation, or Objective. The MaxMin goal minimizes the maximum deviations from a given set of constraints. The Summation goal minimizes the sum of deviations from a set of constraints. The Objective goal optimizes the stated objective. Each goal and each constraint in a goal can be toggled On/Off, indicating whether it is to be included in the optimization formulation. Comments can be added to each goal for user reference.

In Figure 2, the first priority goal is selected “On” and is solved using the MaxMin method. The first constraint within the priority 1 goal is

$$\forall [t \text{ in "Time"} ("Res 4.Pool Elevation" [@ t] \geq "Res 4 Con. Bottom of Daily Operation Zone"[@t])] \quad (1)$$

This constraint requires that the pool elevation at Reservoir 4 be greater than the daily operating zone for all time periods that are indicated in the data object called “Res 4 Con.” This one expression is used to generate a constraint for every timestep. For example, for a 7-day run with 6-hour timesteps, 28 constraints would be generated with this expression. The constraint in Eq. 1 can be represented mathematically for any reservoir r as

$$E_{r,t} \geq BOZ_{r,t} \quad \forall t \quad (2)$$

where E_t is the pool elevation for reservoir r at time period t , $BOZ_{r,t}$ is the pool elevation for the bottom of the operating zone for reservoir r at time period t .

Goals 1 and 2 in Figure 2 are both solved using MaxMin and demonstrate that a number of different constraints can be associated with a single priority. Goal 3 is an objective (maximize net avoided cost for all time periods) and has only one expression associated with it. This goal is solved using the Objective Max method.

The interactive expression editor, as shown in Figure 3, facilitates syntactically correct mathematical specifications of constraints and objectives. A menu provides a selection of common variables and equation prototypes for valid expressions. Question marks represent parts of the expression, which have yet to be built. Table 1 lists RiverWare variables that are commonly used in the Expression Editor. Constants may be included in any part of the expression; RiverWare automatically moves constants to the right-hand side of each constraint before submitting the problem to the LP solver. A single expression can represent a set of constraints by using symbols representing summation (Σ), average (\square) and for-all operations (\forall).

Automatic Linearization

As seen in Table 1, two types of variables can be used in policy expressions: implicit and explicit. Implicit variables are physical or economic variables that may be expressed in a constraint or an objective, but are translated into linear functions of explicit variables before the constraint or objective can be part of the linear pre-emptive goal program model. Thus, the LP solver only works with linear expressions of explicit variables, while policy may be expressed with implicit or explicit variables.

Automatic linearization by RiverWare simplifies the optimization process so that non-optimization experts can use the modeling tool. Each linearization performed by RiverWare is based on a user-designated preferred method. Linearization choices are restricted based on the convexity of a given constraint, the number of terms and the engineering and/or optimization appropriateness of the linearization. For example, tangent is not allowed in power, because it often provides a positive power generation with a zero turbine release. In general, a constraint can either have a single term (ST) or multiple terms (MT) on the left-hand side of the constraint and the operand of the expression is either = (LE) or \geq (GE), providing four possible cases (STLE, STGE, MTLE, and MTGE). Equality constraints are also translated; such constraints are automatically treated as the more restrictive of the GE and LE cases.

Table 2 lists the linearization methods and summarizes the required data for each. Figure 4 shows the user interface for parameterizing the linearization method for pool elevation STLE. In the figure, the substitution method has been selected for this case. For methods such as tangent, line or piecewise, the LP Param table must be com-

pleted. Note that for each implicit variable listed in Table 1, data is entered into a linearization parameter table and a data table relating that implicit variable to another variable.

The substitution method can be used when a constraint contains only one term, which is an implicit variable on the left-hand side (LHS) and an upper or lower bound on the right-hand side (RHS). The substitution method translates the constraint to an equivalent constraint on an explicit variable without introducing any linearization error. For example, Eq. 2 is a constraint on pool elevation, an implicit variable. This constraint can be linearized by substituting the LHS with storage (an explicit variable) and substituting the RHS with the associated storage value, which is interpolated from the object's storage-pool elevation table. Thus, the pool elevation constraint is substituted for a constraint on storage for the same reservoir without any linearization error.

As shown in Appendix III, backwater profiles are modeled in RiverWare with a λ -method (Williams 1990), which takes a user-supplied list of operating points (storage, inflow, outflow, pool elevation and backwater levels) and from these, generates a set of valid operating points. The optimization solution is forced to be a convex combination of these operation points. The λ -method is also one of the methods available for power linearization. In this case, the operating points are valid combinations of power, turbine release, spill, pool elevation, and tail water.

Economic Model for Hydropower

The user can track the economics of hydropower and other power sources by adding a Thermal Object from the palette to the river basin model. If a Thermal Object is linked to the reservoir system, then the user has the option to add policy goals that contain economic variables, such as the objective to maximize the economic benefit of hydropower in a mixed thermal and hydropower system. For example, in the constraint editor the modeler can choose to maximize Net Avoided Cost as shown in goal 3 of Figure 2. RiverWare automatically translates Net Avoided Cost into linear functions of explicit variables. Below is a summary of the components of the Thermal Object. The mathematical formulation is presented in Appendix IV.

Net Avoided Cost is the difference between Avoided Operating Cost in the thermal system (as a result of using hydropower) and the long-term value of water used for power generation and the value of spilled water. There are two options to model the Avoided Operating Cost in the thermal system. Either a piecewise linear replacement value function of hydropower is specified for each timestep or simplified thermal units are directly included in the optimization model. Under the second option, the modeler specifies the cost and capacity of each unit, and for each timestep, specifies the unit availability and the system load that is to be met by all power sources. External power

sources are characterized by supplying the total energy to be used over some period, typically a day, and the maximum and minimum power levels for each timestep. The optimization allocates this energy in concert with the other power sources so as to maximize the Net Avoided Cost.

AUTOMATIC GOAL PROGRAM FORMULATION AND SOLUTION

Formulation of a model for an optimization run includes constructing the basin model, entering the policy constraints, selecting and parameterizing the linearization methods, setting initial conditions, and fixing values of variables at desired timesteps. When the optimization run is executed, RiverWare optimization formulates a sequence of problems for the LP solver. Each problem corresponds to one policy goal expressed in the constraint editor. This automatically generated sequence of problems comprises the goal program.

RiverWare builds the goal program in an object-oriented fashion, directed by the “controller.” First, the controller directs each basin object to generate its physical constraints according to its own data and its knowledge of its own physical processes. Most objects generate a mass balance equation. Continuity equations establishing flow between objects are generated using topological relationships defined by the links. Any modifications made to the topology will result in corresponding changes to the physical constraints generated. Appendix III lists the physical constraints generated by each of the RiverWare objects.

Next, the controller processes the physical and policy constraints and policy objectives, translating them, when necessary, into linear combinations of explicit variables. The context (STLE, STGE, MTLE, and MTGE) of each expression of implicit variables is determined, and the appropriate selected linearization method is applied. For example, if the substitution method has been selected for STGE expressions of pool elevation, then the Priority 1 goal in Figure 2 would be translated from constraints on pool elevation, to constraints on storage on the Res 4 and Res 5 reservoir objects. Mathematically, Eq. 2 is changed to

$$S_{r,t} \geq BOZS_{r,t} \quad \forall \quad t \tag{3}$$

where $S_{r,t}$ is the pool storage for reservoir r at time period t , and $BOZS_{r,t}$ is the storage value that corresponds to the pool elevation level for the bottom of the operating zone for reservoir r at time period t .

The Priority 2 goal in Figure 3 requires no linearization, as outflow is already an explicit variable. Priority 3 goal requires automatic linearization and additional constraints, given it is an economic objective. The mathematical formulation of the economic objective is presented in Appendix IV.

*As with all expressions, RiverWare applies the appropriate selected linearization method to translate the economic objective and its associated constraints into linear functions of explicit variables. In general, some constraints require multiple replacements of implicit variables. For example, energy is replaced by ($\Delta t * power$), and power is in turn expressed as a linear function of turbine release. Once all expressions are only functions of explicit variables, RiverWare is ready to build the pre-emptive goal program.*

Goal Programming and Satisfaction Variables

In traditional formulations of goal programming, policy goals are incorporated into the GP by adding deviation variables to the constraints and minimizing the deviation (Can and Houck 1983; Loganathan and Bhattacharya 1990). Rather than minimizing deviation variables, RiverWare maximizes satisfaction variables to enhance performance. The implementation of the satisfaction variable is outlined below.

Pre-emptive GP insures the optimal solution of a higher priority goal is not sacrificed in order to optimize a lower priority goal. A satisfaction variable, Z_p , is assigned for each goal or priority level p . For each goal, Z_p is maximized, while requiring that all higher priority satisfaction levels be maintained as hard constraints. Let Z'_i be the maximal level that is achieved for any priority $i < p$. Then the goal program solves the following problem for the p^{th} goal

$$\text{Max } Z_p \tag{4}$$

subject to

$$Z_i = Z'_i \quad \text{for } i = 1 \text{ to } p-1 \text{ and} \tag{5}$$

Physical constraints (e.g., mass balance, physical bounds, turbine capacity, etc.)

Incorporating Satisfaction Variables Into Summation, MaxMin, and Objective Policy Goals

The policy constraints (e.g., Priorities 1 and 2 in Figure 2) and policy objectives (e.g., Priority 3 in Figure 2) are translated in RiverWare to include satisfaction variables so that they can be solved as part of the goal program. Several constraints of equal importance can be assigned the same priority p . Frequently, these N_p constraints cannot be simultaneously satisfied to the same satisfaction level, so a relative satisfaction variable ($Z_{p,r,t}$) is defined for these constraints relative to a previous attainment level, $PL_{r,t}$, or the maximum possible value for each constraint. For example, the previous attainment level for an elevation constraint might be a guide curve more relaxed than the oper-

ating zone. Let $Z_{p,r,t} = 1$ represent a fully satisfied constraint and $Z_{p,r,t} = 0$ represent a fully unsatisfied constraint.

Then Eq. 2 is translated into

$$S_{r,t} \geq PL_{r,t} - (PL_{r,t} - BUZ_{r,t})Z_{p,r,t} \quad (6)$$

When $Z_{p,r,t} = 0$, Eq. 6 is equivalent to bounding the storage to a previously attained bound ($PL_{r,t}$), of $S_{r,t}$. As $Z_{p,r,t}$ increases, the constraint tightens. When $Z_{p,r,t} = 1$, then Eq. 6 is equivalent to Eq. 3, the desired policy constraint for timestep t and reservoir r . This procedure is also used for maximum level guide curves.

Once the relative satisfaction variables, $Z_{p,r,t}$, are added to each of the constraints, the $Z_{p,r,t}$ are linked to the satisfaction variable for the entire priority, Z_p . Modelers can select between two metrics for this linkage. The Summation method is shown in Eq. 7. The method maximizes the total satisfaction, which is the sum of the relative satisfaction variables for each priority level.

Max Z_p

subject to

$$Z_p = \sum_{r=1}^R \sum_{t=1}^T Z_{p,r,t}$$

$$Z_i = Z'_i \quad \text{for } i = 1 \text{ to } p - 1$$

and all other physical constraints.

(7)

The Summation metric is used when the marginal value of satisfaction does not depend on the level of satisfaction. For example, the Summation metric can be applied to constrain the spill to zero for all time periods. The optimization values a solution of spill of 5 cms at timestep 1 and 15 cms at timestep 2 as equivalent to a solution of a spill of 10 cms at both timesteps.

The second metric, MaxMin, maximizes the relative satisfaction variable of the least satisfied constraint. In contrast to the Summation metric, the MaxMin metric is interested only in the least satisfied constraint. The MaxMin method is iterative: only after maximizing the satisfaction of the least satisfied constraint does the method attempt to improve other constraints. Once the least satisfied constraint has attained its maximal satisfaction, then that constraint's relative satisfaction variable is fixed and the method goes to the next iteration and re-optimizes over the remaining constraints forcing the second least satisfied constraint to attain its maximal satisfaction. The iterations continue until all $Z_{p,r,t}$ are fixed or are equal to 1. Each of these iterations is a subgoal of priority p and is denoted Z_p^j .

The steps in generating and solving priority p using the MaxMin formulation are

Do until all $Z_{p,r,t} = 1$ or are fixed :

1. Max Z_p^j

subject to

$$Z_p^j \leq Z_{p,r,t} \quad \forall \quad r = 1, K, R \quad t = 1, K, T \quad \text{such that } Z_{p,r,t} \text{ is not fixed}$$

$$Z_i = Z_i^j \quad \text{for } i = 1 \text{ to } p - 1$$

$$Z_p^k = Z_p^k \quad \text{for } k = 1 \text{ to } j - 1$$

and all other physical constraints.

2. Fix the value of any $Z_{p,r,t}$ which is restricting Z_p^j in the optimal solution : $Z_p^j = Z_{p,r,t}$

3. Replace Z_p^j with Z_p^{j+1} in all remaining constraints

4. Go to next subgoal ($j = j + 1$)

The repeated use of the MaxMin objective on a constraint such as Eq. 6, transforms the discrete flood guide curves to a continuum of guide curves between the flood guide and the previously attained guide curve and is called a "shrinking envelope"(Magee 1999). Shrinking envelopes implicitly generate intermediate guide curves, which are consistent with the original guide curves without the modeler having to specify a series of intermediate guide curves. These implicit intermediate guide curves force the optimization to provide a solution that is balanced over reservoirs, which is generally a desired operating policy.

While the above example discusses the transformation of pool elevation policy constraints to GP objectives, the approach can be applied to any policy constraint. For example, the total storage on one branch of a river or the average flow in a river reach over the last 24 hours can be constrained. The Summation and MaxMin metrics can be applied at the user's discretion.

Goals that are expressed as policy objectives (e.g., Priority 3 goal in Figure 2) can be directly incorporated into the GP formulation. First, the objective is translated into a linear expression (as described earlier), then a satisfaction variable for the p^{th} goal (Z_p) is set equal to the objective. Given the modeler's selection, Z_p is then maximized or minimized.

The constraints enforcing the attainment of previous goals are not explicitly written, because it is both inefficient and leads to numerical instability when performed on a large scale. Instead, the solution can be enforced implicitly by manipulating non-basic variables (Magee 1999).

CPLEX Interaction

RiverWare transfers the problems or subproblems to solve each goal to a commercial optimization library (CPLEX 1994). RiverWare's graphical Run Control dialog allows the user to interrupt CPLEX in the solution of each of these goals and to receive updates on the status of the solution. In addition, an advanced modeler can fine-tune performance by changing CPLEX and RiverWare performance parameters through the RiverWare optimization GUI.

After all the goals have been processed and optimized by CPLEX, the values of the explicit variables are propagated from CPLEX back into the objects' data structures and can be viewed in the GUI. Values of implicit variables are not returned from CPLEX because multiple linearization methods may have been in effect for a single implicit variable, resulting in multiple values for one implicit variable. The values of implicit variables are determined during the post optimization simulation.

Post Optimization Simulation

Once RiverWare has completed an optimization run, schedulers can directly use the output from the optimization as simulation input, or they can modify the optimization output before submitting it as input for a simulation run. The post optimization simulation results can be used to determine how much error was introduced into the optimization model due to nonlinear function approximations. The simulation tools include models of the complex nonlinear processes such as reach routing methods, hydropower generation, and backwater elevations on sloped reservoirs.

Data Consistency Checking and Access to Input and Output

An essential feature of RiverWare is its ability to trap errors and identify data inconsistencies. In some cases, RiverWare automatically populates data tables or uses appropriate default values for missing or out-of-range data. Data is checked to see that it follows a consistent pattern: e.g., table values are in increasing order, with no duplicates; or table data cover the minimum and maximum values specified elsewhere; or required tables are convex or concave. In cases where it is possible to clearly indicate the nature of a data inconsistency, RiverWare performs limited feasibility testing for subsystems of equations and variables. These checks catch most potential infeasible problems.

Schedulers have a number of choices of methods for viewing the input and output data from any model run. The data can be presented in units selected by schedulers. Input or output from a model run for a particular reservoir can be viewed by opening the reservoir icon. Alternatively, the user-configured System Control Table (SCT) displays

data from many objects simultaneously in a spreadsheet-style grid. This tool also allows interactive data entering, and editing in a system-wide view. Input and output data can be graphed or exported to external programs, reports or analysis tools through the DMI.

TVA APPLICATION EXPERIENCE

The RiverWare optimization module was developed through a collaborative effort involving a core group of TVA modelers and researchers from the Center for Advanced Decision Support for Water and Environmental Systems (CADSWES) at the University of Colorado at Boulder. TVA began using RiverWare optimization in June 1998 as a guide for developing optimal hydropower release schedules. This effort insured that RiverWare optimization has the functionality and user-friendly features required for daily scheduling. As the RiverWare optimization features were implemented by CADSWES, the TVA modelers performed many test runs and provided feedback to the development team regarding functionality, error trapping, and additional features needed for production use. The TVA team was comprised of senior operations staff with experience in mathematical modeling and other staff involved in the day-to-day tasks associated with scheduling and monitoring the river system. The TVA operations expertise was instrumental in building the production models used by schedulers. The models are “over built” with many goals turned off, so that when special operations arise, schedulers can make the appropriate adjustments using the constraint editor.

For TVA, RiverWare optimization provides multiple benefits in the form of optimal timing of turbine releases and automation of operating policy documentation, which is explained below. Table 3 summarizes the features of the RiverWare optimization model used by TVA. Table 4 summarizes a typical list of TVA prioritized policy goals for the summer season. The two tables reflect the complex nature of TVA’s operating constraints that include pool elevation guides, channel capacities, spill avoidance, minimum flow requirements, ramp rates, and special operations for 35 reservoirs.

Each day, scheduling staff loads the RiverWare model with initial conditions, current inflow forecasts, hydrogeneration unit availabilities, and hydropower values. If necessary, the priority of policy goals is modified. The RiverWare optimization is then executed, followed by one or more interactive RiverWare simulation runs. As seen in Table 3, most scheduling runs take about 5 minutes and use a 6-hour timestep for a weeklong run. The senior reservoir operations staff considers linearization approximation errors to be well within acceptable limits. The RiverWare optimization solution model accurately characterizes the system and produces a realistic schedule for TVA's demanding operational environment.

The use of RiverWare optimization provides forecasters with explicit economic information about the trade-offs between using water in the near term vs. saving it for future use. Previously, daily operations were based primarily on projections of water use to meet long-term future elevation targets. While this information is still available and provides valuable information to schedulers, the RiverWare optimization economic analysis shows the magnitude of the costs of alternate operations. Not only is this economic information useful in setting the routine daily use, but it is useful in quantifying the cost of requested special operations for civic events and project maintenance.

The RiverWare model file, saved daily, provides an easy means of policy documentation. It contains the physical representation of the system, the hydrologic inflow and forecast data, the policy objectives and constraints, and the economic data. Prior to using RiverWare, there was little historical record of daily operating policy other than the senior operations engineer's notes. To the extent that TVA follows the optimization guidance in setting schedules, it is now possible to reconstruct the operating policy on a past date from the saved RiverWare model file.

EXAMPLE APPLICATIONS

The following five examples illustrate how RiverWare can be used as a tool to optimize hydropower releases and explore daily operations alternatives. The river system used in the examples has been constructed by drawing from characteristics of various portions of the TVA system. However, the examples and data are sufficiently different from the TVA system that no conclusions about TVA operations should be inferred from the examples.

Figure 1 shows the example basin. Table 5 summarizes some of the physical and economic characteristics of the system. A six-hour timestep is used to optimize the generation schedule over a typical 7-day period during the summer. The operation policy goals are similar to those in Table 4. At this time of year, with the elevations reported in Table 5, the solution is driven by the economic goal of the maximization of avoided power cost after satisfying other multipurpose goals. The economic data used by the model includes a piecewise representation of hydro value by hour during the 7-day optimization and data representing future water value beyond the 7-day period for each project. Table 6 is a description of the five examples, and Table 7 outlines the results of the five model runs.

A reasonable reservoir management policy is to have target elevations for various times of the year. Table 5 indicates that August 1 target elevations exist for only four of the eight reservoirs. Three of the eight reservoirs are already at this target level, so the model is set to constrain these reservoirs at these elevations. As shown in Table 6, the final objective value for the Base Case run is \$5,215,000.

The Scheduled Draw Down example is similar to the Base Case in all respects except target elevations are

fixed for all reservoirs at the last time period of the run, as shown in Table 5. The target values reflect a linear draw down rate from beginning of the run until August 1. This policy is a rational approach to meeting the elevation targets for August 1. Table 7 shows that the final elevations of this run are lower than the base case for reservoirs that were unconstrained in the Base Case example and that the objective value, as shown in Table 6 for this run, reflects the loss of \$469,000. Forcing the uneconomic drawdown at numerous reservoirs causes this loss. Given that the future value of water is higher toward the end of the summer, it makes economic sense to hold the water back. In this example, there is more generation occurring at the headwater reservoirs compared to the Base Case example. With this type of analysis, RiverWare allows a daily operations scheduler to reevaluate the cost of retaining a historic drawdown policy.

The Unit Outage I and Unit Outage II examples consider two different approaches to scheduling a turbine maintenance outage at Reservoir 5. Without an optimization model, daily operations schedulers may simulate a few outage scenarios in order to meet the demand in the best way possible. When an optimization model is used to schedule a maintenance outage, the entire reservoir system can make up the difference for the outage. Unit Outage I schedules the outage at Reservoir 5 by constraining the releases to zero for three days of the week, while Unit Outage II constrains the releases to half the capacity for the entire week. As seen in Table 6, there is a significant savings (\$20,000) by requiring the outage to occur at 50% over a 7-day period, rather than a 100% outage over a 3-day period. Of course, other costs associated with scheduling maintenance would need to be weighed against the scheduling impact to determine the most appropriate alternative, but both examples adjust outflows throughout the system to accommodate the outage. Without an optimization tool, such a system-wide adjustment would be a time-consuming exercise. Therefore, the use of RiverWare optimization can lead to more economic maintenance planning.

The final example, Remove End Target, releases the August 1 target elevation constraint for Reservoir 4, which, as shown in Table 5, has the highest future value of water. Removing this target allows an upstream reservoir (Reservoir 1) to increase generation even though it was not economical to generate at Reservoir 4: Reservoir 4 stores the excess water. Given the difference from the Base Case objective value (\$5,100), we can conclude that Reservoir 4 had been acting as an economic bottleneck in the Base Case example. Water was not used for generation upstream in the Base Case, because it was uneconomical to generate with it downstream. Table 7 shows the different ending elevations for the two runs. In the Remove End Target example, Reservoir 4 releases only during peak demand and only at best efficiency, and its elevation goes to the next higher guide curve.

CONCLUSION

RiverWare optimization is a general river basin decision support tool that allows water resources engineers

to solve a complicated optimization problem by specifying a physical and economic model of the system, listing prioritized policy goals, and indicating linearization parameters. A user-friendly graphical user interface and error trapping capabilities facilitate the specification and modification of this information. RiverWare is suitable for formulating a wide range of operating policies at a variety of large, multipurpose river basins. The policy can be changed easily, and RiverWare automatically generates an efficient and robust pre-emptive goal program to optimize the policy. The solution to the goal program automatically defines a simulation run, which can predict the exact consequences of the optimization solution.

ACKNOWLEDGMENTS

RiverWare has been developed as a joint collaboration between CADSWES at the University of Colorado at Boulder, the Tennessee Valley Authority, and the U.S. Bureau of Reclamation.

APPENDIX I. REFERENCES

- Allen, R.B., T. Olason and S.G. Bridgeman (1996). A Decision Support System for Power Systems Operations Management, Proc., ASCE Fifth Water Resources Operations Management Workshop, Aris P. Georgakakos and Quentin W. Martin, eds.
- Andreu, J., J. Capilla and E. Sanchis (1996). AQUATOOL, a generalized decision-support system for water-resources planning and operational management. *Journal of Hydrology* 177(3/4), 269 - 291.
- Can, E. and M. Houck (1983). "Real-Time Reservoir Operations by Goal Programming," *Journal of Water Resources Planning and Management*, ASCE, 110(3), 297-309.
- CPLEX (1994). CPLEX Optimization, Inc., Suite 279, 930 Tahoe Blvd., Bldg. 802, Incline Village, NV 89451.
- HEC-5 Simulation of Flood Control and Conservation Systems User's Manual; Version 8.0.* (1998). U.S. Army Corps of Engineers Hydrologic Engineering Center, Davis, CA.
- Labadie, J. (1995). River Basin Model for Water Rights Planning, MODSIM: Technical Manual, Department of Civil Engineering, Colorado State University, Ft. Collins, CO.
- Labadie, J. W. (1997). "Reservoir System Optimization Models" *Water Resources Update*, 108, Universities Council on Water Resources.
- Loganathan, G. V. and D. Bhattacharya (1990). "Goal-Programming Techniques for Optimal Reservoir Operations," *ASCE Journal of Water Resources Planning and Management* , 116(6), 820- 838.
- Loucks, D.P. and U. Shamir (1989). "Preface: Systems analysis for water resources management: closing the gap between theory and practice." *Proc., Third Scientific Assembly of the International Association of Hydrological Sciences (IAHS)*, Baltimore, Maryland. Daniel P. Loucks and Uri Shamir. (eds.). International Association of Hydrological Sciences, Wallingford, U.K.
- Magee, T., (1999). "Shrinking Envelopes," In preparation for submission to *Journal of Operations Research*, March.
- Robitaille A., F. Welt, and L. Lafond (1995). Development of a Real Time River Management System for Hydro-Quebec's Short Term Operation, Waterpower '95, John J. Cassidy, (ed.).
- Shane, R., D. Waffel, J. Parsly and H. M. Goranflo (1988). "TVA Weekly Scheduling Model Application Experi-

ence,” *Computerized Decision Support Systems for Water Managers*, J. Labadie et al., (eds.), ASCE, New York.

Schultz, G.A. (1989). "Ivory tower versus ghosts? -- or -- the interdependency between systems analysts and real-world decision makers in water management," *Systems analysis for water resources management: closing the gap between theory and practice. Proc., Third Scientific Assembly of the International Association of Hydrological Sciences (IAHS)*, Baltimore, Maryland. Daniel P. Loucks and Uri Shamir, (eds.), International Association of Hydrological Sciences, Wallingford, U.K. p.23-32.

Sprague, R. H. and Calson, E. D. (1982). *Building effective decision support systems*. Prentice-Hall, Inc., Englewood Cliffs, N.J.

WEAP Water Evaluation And Planning System Users Guide, (1997). Tellus Institute, Boston Mass.

Williams, H.P. (1990). *Model Building in Mathematical Programming, 3rd Edition*, Wiley & Sons, New York, NY.

Wurbs, R. (1993). "Reservoir-system Simulation and Optimization Models," *Journal of Water Resources Planning and Management*, ASCE, 119(4), 455-472.

Wurbs, R. (1996). *Modeling and Analysis Of Reservoir System Operations*, Prentice Hall, New Jersey.

Yeh, W. (1985). "Reservoir management and operations models: A state-of-the-art-review," *Water Resour. Res.*, 21(12), 1797-1818.

Zagona, E., T. Fulp, H. M. Goranflo and R. Shane (1998). "RiverWare: A General River and Reservoir Modeling Environment," *Proc., First Federal Interagency Hydrologic Modeling Conference*, Las Vegas, NV, 5-113.

APPENDIX II. NOTATION

The following symbols are used in this paper:

$AV_{u,t}$ = availability of unit u for timestep t

$A_{a,t}$ = allocated power from external source a for timestep t

BOZ = bottom of operating zone pool elevation

$BOZS$ = storage corresponding to bottom of operating zone pool elevation

b = generic coefficient

C = plant capacity fraction

E = reservoir elevation

$EA_{a,d}$ = total energy from external source a available for allocation on day d
(other timesteps possible).

EIS = energy in storage

$E1$ = canal elevation 1

$E2$ = canal elevation 2

EN = energy

$f()$ = a function of another variable, usually represented by a user selected linearization
method.

L_t = forecasted load.

$LTC_{r,t}$ = long term expected value of water,

(typically determined by a separate long term planning model.)

N_A = number of external energy sources to allocate.

N_b = number of blocks in the piecewise function.

N_{day} = number of days in the planning horizon.

N_u = number of units.

OH = operating head

P = power

PL = previous attained level for a given variable

PP = pump power

Q_c = canal flow

Q_d = diversion

Q_h = hydrologic inflow

Q_i = inflow to a reservoir, reach

Q_{i1} = first inflow to a confluence

Q_{i2} = second inflow to a confluence

Q_p = pump flow

Q_o = outflow from a reservoir, reach, or confluence

Q_s = spill

Q_{tr} = turbine release

R = number of reservoirs

S = storage

SE = spilled energy

$SEPC$ = spilled energy power coefficient

SLB = spill lower bond

SUB = spill upper bond

T = number of timesteps

TC_u = unit cost of operating thermal unit u .

TOC_t = cost of solution using only thermal units to meet the load in time period t ,
(determined by simulation)

TS = timestep, e.g. 6 hours.

$TU_{u,t}$ = variable, amount of thermal unit u to use in meeting the load.

TW = tailwater

$U_{b,t}$ = fraction of block b is generated by the optimal solution in time period t .

$V_{b,t}$ = the marginal value of hydropower for block b in time period t .
(slope of the piecewise linear function,

λ_j = the fraction of solution j used in a convex combination of trial solutions



= the value of any variable X in a trial solution j

Z_p = satisfaction variable for the p^{th} goal

Z_p' = maximum value obtained by the satisfaction variable for the p^{th} goal

Z_{prt} = relative satisfaction variable for the p^{th} goal, for the reservoir r and timestep t

Z_p^j = satisfaction variable for j^{th} iteration of subgoal of the p^{th} goal

$Z_p^{k'}$ = is the maximum value obtained by the satisfaction variable for k^{th} iteration of
subgoal of the p^{th} goal

Subscripts

lag = lag time

p = p^{th} policy goal

r = reservoirs

t = time index

APPENDIX III. AUTOMATICALLY GENERATED PHYSICAL CONSTRAINTS

The physical constraints that are generated by each object used in RiverWare optimization are presented below. Each of the symbols is defined in Appendix II.

Confluence

Mass Balance

River Reach

Mass Balance, lagged routing

Canal

Mass Balance

Reservoir

Mass Balance

Flow

$$Qo_t = Qs_t + Qtr_t$$

Elevation Volume Relationship

$$E_t = f(S_t)$$

Spill

Operating Head

Energy in Storage

Energy

Power Option 1: Independent Linearizations

Turbine Capacity

Power (using estimated or initial operating head)

Pump Capacity

Pump Power (using estimated or initial operating head)

Power Option 2: Convex Combination of Trial Points Provided by Modeler

Backwater Profile for Sloped Storage Reservoirs

Convex Combination of trial points

APPENDIX IV. ECONOMIC MODEL

The constraints presented below represent the economic model that is produced by the Thermal Object. The modeler can select to use either of the models presented below, or can use a combination of the economic variables to generate a different economic objective.

Net Avoided Cost, Piecewise Linear Formulation

$$\text{Max} \sum_{t=1}^T \left[\sum_{b=1}^N V_{b,t} U_{b,t} - \sum_{r=1}^R LTC_r Q_{r,t} \right] \quad (\text{net avoided cost})$$

Subject to

$$\sum_{b=1}^N U_{b,t} = \sum_{r=1}^K P_{r,t} - \sum_{r=1}^K PP_{r,t} + \sum_{a=1}^N A_{a,t} \quad \forall t = 1, K, T \quad (\text{total generation})$$

$$\sum_{t \in \text{Day } d} A_{a,t} = EA_{a,d} \quad \forall a = 1, K, N_a, d = 1, K, N_{\text{day}} \quad (\text{daily allocation})$$

Net Avoided Cost, Thermal Unit stack

$$\text{Max} \sum_{t=1}^T \left[TOC_t - \sum_{u=1}^N TC_{u,t} TU_{u,t} - \sum_{r=1}^R LTC_r Q_{r,t} \right] \quad (\text{net avoided cost})$$

Subject to

$$\sum_{u=1}^N TU_{u,t} + \sum_{r=1}^K P_{r,t} - \sum_{r=1}^K PP_{r,t} + \sum_{a=1}^N A_{a,t} = L_t \quad \forall t = 1, K, T \quad (\text{load})$$

$$\sum_{t \in \text{Day } d} A_{a,t} = EA_{a,d} \quad \forall a = 1, K, N_a, d = 1, K, N_{\text{day}} \quad (\text{daily allocation})$$

$$0 \leq \sum_{u=1}^N TU_{u,t} \leq AV_{u,t}$$