OPTIMIZING DAILY RESERVOIR SCHEDULING AT TVA WITH RIVERWARE

Timothy M. Magee[1], Operations Research Analyst, University of Colorado,
H. Morgan Goranflo[2], Senior Consultant, Tennessee Valley Authority, and
Suzanne H. Biddle[2], Specialist, River Scheduling, Tennessee Valley Authority

[1]Center for Advanced Decision Support for Water and Environmental Systems, Campus Box 421, Boulder, CO 80309-0421, 303-492-2657, magee@colorado.edu
[2] River System Operations and Environment, 400 West Summit Hill Dr., Knoxville, TN, 37902

Abstract

RiverWare is a general software tool for modeling river basin operations developed at the Center for Advanced Decision Support for Water and Environmental Systems (CADSWES) at the University of Colorado, sponsored jointly by the Tennessee Valley Authority (TVA) and the U.S. Bureau of Reclamation (Reclamation). RiverWare has several solution methodologies, and TVA has implemented two models that use RiverWare's optimization method followed by simulation of the optimal outflows: a 9-day model with 6-hour time steps and a 2-day model with an hourly time step.

We will present highlights from the mathematical framework, the software interface, and the model’s role in the larger scheduling process. The mathematical framework includes preemptive goal programming with a linear programming engine and piecewise-linear approximation of nonlinear functions. The software interface allows the schedulers to import and view data, modify a prioritized list of policy constraints, view and alter the optimal solution, and analyze how both physical constraints and policy constraints influenced the solution. The models are only part of a larger process that includes forecasts of hydrologic inflows, forecasts of the power market, special operations at reservoirs and reaches, and past operating experience.

INTRODUCTION

Dam construction in the United States has slowed considerably, and with that change, emphasis has shifted to improved management of existing reservoirs, typically with multiple, and some times conflicting uses. RiverWare is one general modeling tool that has been developed to meet this need. RiverWare was developed by the Center for Advanced Decision Support for Water and Environmental Systems (CADSWES) at the University of Colorado, and sponsored jointly by the Tennessee Valley Authority (TVA) and the U.S. Bureau of Reclamation (USBR), and since then also has been adopted by other modelers. RiverWare was inspired by past modeling experience at both TVA and USBR and reflects both the successes and limitations of prior models.

The previous models were considered useful to the agencies, but they had two notable limitations. First, previous models were dedicated to particular river basins and had to be maintained individually by the agencies without the benefit of similar work being conducted on other river basins. Secondly, policy was imbedded in the code and largely inaccessible to reservoir operators. Thus, changing the policy in the models either to reflect changes in reality or to conduct policy studies was a large programming task that depended on a small number of highly trained individuals. Additionally, "hiding" policy in the code made it more likely that errors would go undetected.

RiverWare has improved modeling in terms of these limitations, of course, without entirely eliminating programming or policy errors. Adding new concepts to a model still requires programming, and RiverWare has continuously added new features over the last ten years, but now those new features become available to all of the RiverWare models. Thus, the programming and maintenance costs are shared rather than repeated at each basin. Policy errors still appear in models from time to time, but now they are visible to the end users and can be repaired without rewriting and recompiling software.

While many river basins share many attributes, one frequent difference is the nature of policies that control the operation of a river basin. For example, some systems focus on maximizing hydropower revenue within limitations placed by the other reservoir uses. The Tennessee River is such a system because of the original legislation creating TVA. In contrast, other systems are so driven primarily by laws, agreements, and historical operating rules that little flexibility remains to manipulate hydropower generation. Still other systems are driven largely by water rights. Consequently, RiverWare has several "controllers" for setting reservoir releases; each oriented towards meeting one of these policy frameworks. In addition, RiverWare has a simulation controller that acts as a sophisticated "calculator" for any policy that has sufficiently specified reservoir operations. Each of the policy-oriented controllers relies on this simulation capability in some way. For example, the optimization controller relies on a post-optimization simulation to remove small errors created by linear approximation of non-linear functions. The simulation is driven by the optimal reservoir releases. Zagona *et al.* (2001) describe the other RiverWare controllers in more detail as well as some notable previous approaches to reservoir modeling. Eschenbach et al. (2001) includes a summary of previous approaches to optimizing reservoir operations.

The rest of this paper is organized as follows. First, we will describe the mathematical framework of RiverWare's optimization controller, a preemptive linear goal program. RiverWare is designed for use by reservoir schedulers rather than experts in mathematical programming. Thus, the second section will describe the software interface that allows schedulers to build, control and interpret models at a high level rather than at a mathematical programming level. We will illustrate the use of these tools with the models in use at TVA: a 9-day model with 6-hour time steps and a 2-day model with an hourly time step. In the final section, we will describe the larger reservoir scheduling process at TVA and how the RiverWare models fit into this process.

GOAL PROGRAMMING IN RIVERWARE

The optimization controller in RiverWare is based on optimizing a series of linear functions over a set of linear constraints: preemptive linear goal programming. The essential decision variables in this optimization are the reservoir releases at each time period. However, other variables are included in the model to represent constraints and objectives. For example, reservoir variables include storage, pool elevation, operating head, backwater elevation, inflow, outflow, spill, turbine release, turbine increase/decrease, evaporation, and bank storage. Reach variables include inflow and outflow. Canal variables include flow and head. Variables for water users include diversion, depletion, return flow, water available for diversion, and intake elevation. This list isn't comprehensive and new variables are frequently added as new features are modeled. In addition, modelers can create new variables in their models without any revision of RiverWare.

The constraints that must be included to model reservoir releases are mass balance and continuity. However, many other constraints have been included to define all of the variables listed above. For example, reservoirs have constraints on turbine capacity, spill, and sloped storage.

One way to model the multi-purpose aspect of reservoir management is with goal programming. Each prioritized goal can advance one (or more) of the purposes. The goal program formulates a series of linear programs, each of the form:

$$\begin{aligned} &\mathit{Max}\ \mathbf{cx} \\ &s.t. \quad \mathbf{Ax} \le \mathbf{b} \\ &\mathit{where}\ \mathbf{c}, \mathbf{x} \in \mathbf{R}^{n}, \mathbf{b} \in \mathbf{R}^{\mathbf{m}}, \text{and } \mathbf{x} \text{ is unknown} \end{aligned}$$

The goal program can be defined by the following pseudo-code which steps through n prioritized objective functions indexed by p, where z_p is the p^{th} objective function, $z_p = \mathbf{c}^p\mathbf{x}$, and z^*_p is the optimal value of that objective:

For p = 1 to n
{ $\mathit{Max}\ \mathbf{c}^{p}\mathbf{x} \quad s.t. \quad \mathbf{Ax} \le \mathbf{b}$
Add a new constraint preserving the optimal objective function value:
$z_p = z^*_p$
}

(In practice, rather than using this formulation literally, a series of equivalent linear programs are formulated for a more efficient algorithm.)

One recent extension to this framework is to relax the constraints that the previous objectives must exactly their optimal values. For example, a previous objective might be constrained to be within 5% of optimality:

$$z_p \ge 0.95 z^*_p$$

By introducing a little flexibility into prior objectives, subsequent objectives may be dramatically improved. TVA expects to incorporate this feature in their models in the near future.

This framework can be further extended to allow modelers to specify a set of constraints in place of a prioritized objective. Of course, ideally the constraints would be fully satisfied. If this were the case, the constraints could just be added to the other constraints. However, if all the constraints cannot be simultaneously satisfied, a reasonable objective would be to satisfy the constraints as closely as possible. If more than one constraint exists, the constraints may be in conflict and some balanced means of resolving the conflict must be achieved.

RiverWare achieves such a balance by translating the satisfaction of the constraints back to an objective function. RiverWare requires that each variable is required to be bounded and hence any linear expression is bounded. For example, for the expression, 3w + 2y, the bounds are:

$$3LB(w) + 2LB(y) \le 3w + 2y \qquad \text{(lb), and}$$

$$3w + 2y \le 3UB(w) + 2UB(y) \qquad \text{(ub)}$$

where LB() and UB() represent the lower and upper bounds respectively on individual variables. Thus, a constraint, i, such as

$$3w + 2y \le \mathrm{b} \qquad \text{(i)}$$

could be written with a satisfaction variable, z_i, as

$$3w + 2y \le 3UB(w) + 2UB(y) - (3UB(w) + 2UB(y) - \mathrm{b})z_i$$

$$0 \le z_i \le 1$$

Notice, that when $z_i = 0$ the constraint repeats the upper bound on the expression, (ub), and when $z_i = 1$ the original constraint, (i), is fully satisfied. RiverWare provides two methods to combine the individual z_i into an objective.

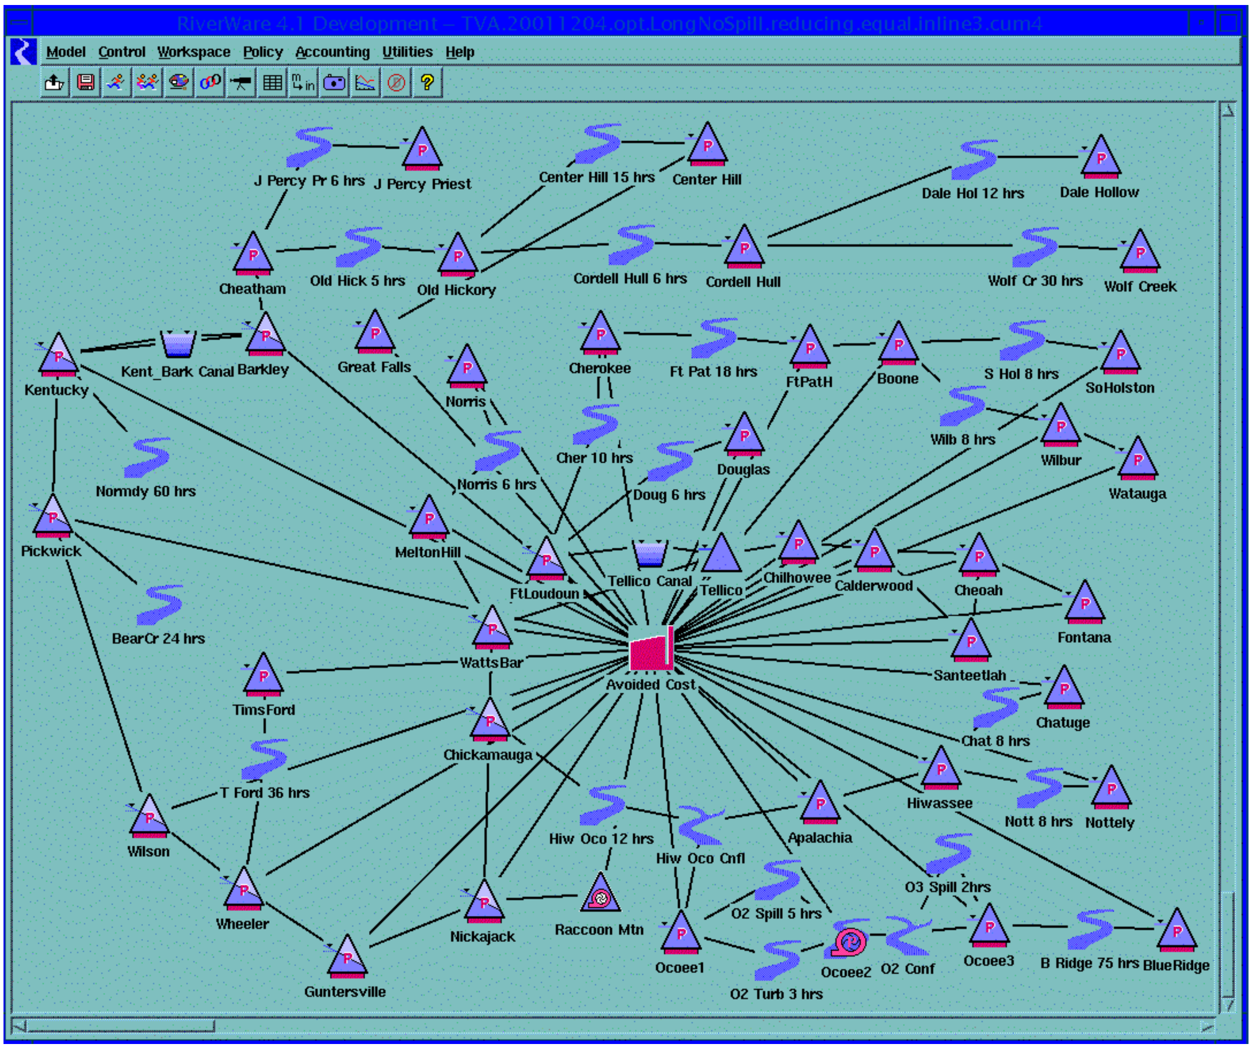

Figure 1 : TVA system modeled in RiverWare

One approach is to define

$$z_p = \sum_i z_i$$

which is called “Summation” in RiverWare. The other approach is to define

$$z_p \le z_i$$

which is called “MaxiMin” in RiverWare because it maximizes the satisfaction of the least satisfied constraints. A slight twist to this objective is that there may be alternate optimal solutions with some of them leaving room to improve the satisfaction of the other constraints. Just as in the larger goal program, we could “freeze” the optimal solution for the least satisfied constraint(s) and optimize over the remaining constraints. This is exactly what RiverWare does. (Technically, this repeated optimization procedure maximizes the “L_∞ norm” of the satisfaction variables, while the summation approach maximizes the “L_1 norm”.) The vast majority of TVA’s priority levels use the MaxiMin approach on a set of constraints.

RIVERWARE SOFTWARE INTERFACE

RiverWare uses an object-oriented approach for modeling river basins. Different parts of a river basin are modeled as different classes objects. For example, TVA’s model (`Figure 1`) contains Storage Reservoirs, Level Power Reservoirs, Sloped Power Reservoirs, Pumped Storage Reservoirs, Inline Power Objects, Reaches, Canals, Data Objects, and a Thermal object. A full description is beyond the scope of this paper. See Zagona et al. (2001), for more details on these objects and others.

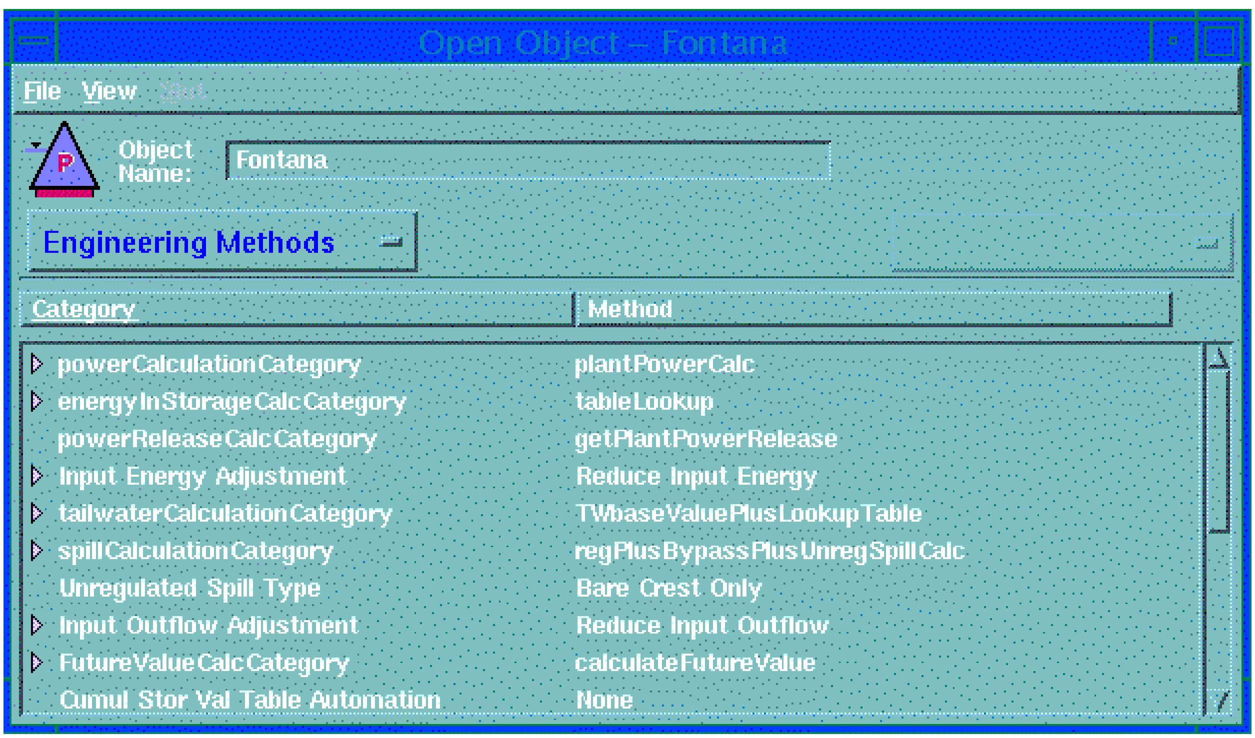

Figure 2: Method Categories and Selection on a Reservoir

Each class of object has knowledge of its physical processes, modeling features, and data. Modelers configure individual objects by selecting "methods". The available method choices depend on the controller, object class, and other method selections. An important group of methods in the optimization controller controls how nonlinear functions are converted to linear and piecewise linear functions. While modelers can select which method to use, RiverWare also provides default methods that work well in most cases. `Figure 2` illustrates some of the method categories and their current selected methods for a reservoir.

For a given selection of methods, a set of applicable data values, "slots", is defined. Slots consist of time series, "series slots", such as inflow and outflow, and data tables, "table slots", that typically hold functional relationships such as elevation-volume or elevation-spill. Series slots may be input by users, output by a controller, or some combination of both. Slots may be linked so that two slots on different objects have equal values or one slot may equal the sum of several others. A common use of links is to link inflows on one object to outflows from other objects. `Figure 3` illustrates the slots on a reservoir.

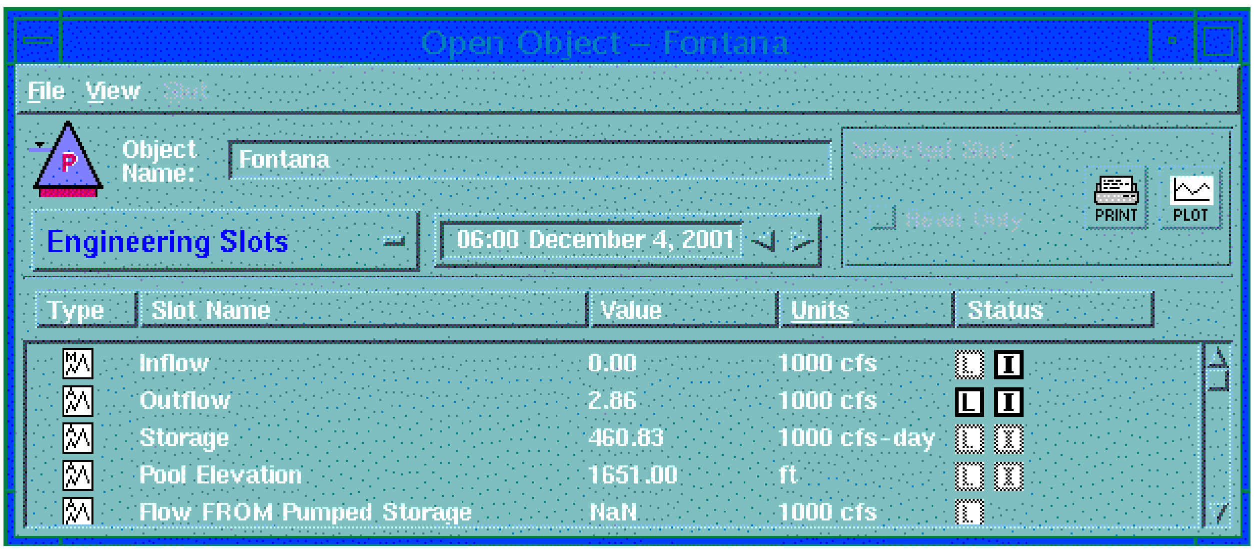

Figure 3: Slots on a Reservoir

In optimization, many of the series slots correspond to variables that are automatically added to the linear program. In addition, RiverWare automatically adds constraints on these variables to the linear program. These constraints reflect physical constraints on the variables, auxiliary constraints for modeling piecewise-linear functions, and policy constraints written by modelers. Physical constraints include mass balance for each object, links between serial slots, turbine capacity, elevation-spill curves, canal flow as a function of elevations, lagged reach flows, etc. Modelers can control some of the variables and constraints that are included in a model through method selections. For example, evaporation may be omitted for reservoirs with inconsequential evaporation. More details about the automatic generation of the goal programming formulation can be found in Eschenbach et al. (2001).

Policy

RiverWare has a "constraint editor" that modelers can use to specify constraints and objectives. Modelers build these expressions by selecting menu items. The menu choices are restricted by the context of a partially constructed expression. In addition, the constraint editor allows modelers to activate (check mark), deactivate (X mark), and prioritize policies. Figure 4 illustrates the constraint editor with TVA's policy. Each of the priority levels shown in the figure can be opened to reveal the individual constraints by selecting the triangle.

Modelers can access all of the defined variables when creating policy as well as data values by selecting slots from the menus. In addition to the variables that are directly represented in the optimization problem, the modeler can just as easily access functions that will be automatically replaced with linear or piece-wise linear functions when the optimization problem is formulated. This substitution process is recursive: by referencing a single variable, a modeler may create an expression that has thousands of variables after all of the substitutions are made. For example, the economic objective is stated in TVA's models with reference to a single variable, Net Avoided Cost, and yet through substitution the objective ends up being expressed in terms of all of the individual reservoir flows, the final reservoir storages, and more. Of course, model users can control the details of the substitution through the selection of methods. More details about the linearization methods can be found in Eschenbach et al. (2001) and more details on policy expression can be found in Magee et al. (2001).

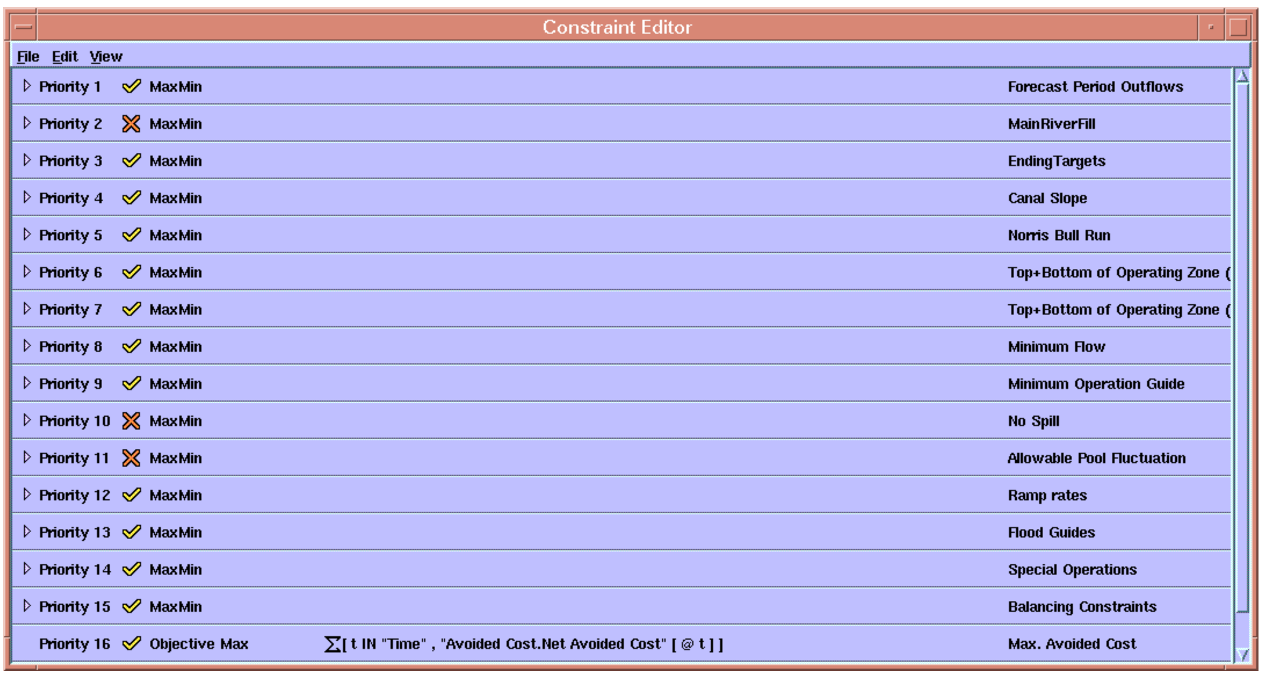

Figure 4: TVA policy in the Constraint Editor

Policy in TVA's 6-hour model can be described as a combination of flood control, navigation, recreation, water quality, power generation, and special operations. Many of these policies are described by individual reservoir guide curves that define the allowable flow based on reservoir elevations.

Other kinds of constraints in TVA's 6-hour model include:

meeting an ending target elevation,
ramping rates - setting maximum change in power generation,

restricting increases in reservoir storage to natural inflows,
capping the head across a canal,
constraining the total storage in a subbasin,
forcing fluctuations to limit mosquito population,
balancing elevations at comparable reservoirs,
limiting spill, and
meeting regional power demands before using power for economic benefit.

Some of these constraints are activated or adjusted on a seasonal basis. In addition, TVA occasionally activates "special operations" constraints that limit power or flow based on temporary circumstances such as plant maintenance or recreational activities.

Finally, this model contains an economic objective to maximize the value of hydropower. The details of hydropower modeling are beyond the scope of this paper, but are described in detail in Zagona and Magee (1999). TVA makes two similar runs with this model that differ largely in terms of the objective. Two runs are required because more detailed information on the value of power is available in the short term than the long term.

Prior to the first RiverWare run, TVA runs their internally developed program to estimate the long-term (60-90 days) value of remaining project storage at the end of the RiverWare planning horizon. The inputs to this program, Value of Project Storage (VPS), include: ending elevations from previous RiverWare runs, future target elevations, and expected hydrologic inflows to predict the percentage of time each hydro plant is operating. The model then assumes that each plant will operate only during the most valuable hours. The least valuable hour used under this assumption defines the marginal value of an increase in generation. This model assumes that each hour has a fixed value of generation that does not depend on the quantity of power generated, sometimes referred to as a "system lambda." The data used at TVA is called "POSE" after a previous model with an hourly value of power. By running the VPS program with alternate ending RiverWare elevations, the marginal value of stored water can be computed across a range of storage levels for each hydro plant. These marginal values are passed to the first RiverWare run. The first RiverWare run uses the same POSE data for valuing short-term power generation and trades this off against the long-term value of leaving the water in the reservoirs as predicted by the VPS model. The solution to this run includes an optimal ending elevation for each reservoir.

In the second run, these ending elevations become constraints, replacing the economic value provided by the VPS program. Thus, power values required are only required for the short-term and a more detailed power valuation can be used. Specifically, each time period has a piecewise linear value of generation with decreasing returns to scale. Using this detailed data allows for a more informed optimization of short-term power generation.

The main purpose of TVA's hourly model is to allocate the daily releases from the 6-hour model to individual hours. One approach would be to require the hourly model to meet the 6-hour releases exactly. Instead, this condition is relaxed to require only that the 24-hour totals match. These "daily volume" constraints remove the need for many, but not all, of the constraints in the

6-hour model. In addition, some other constraints are added at the hourly level. For example, some small plants are required to operate in tandem with adjacent larger plants.

Post-optimization

RiverWare has an Optimization Analysis Tool (OAT) that assists modelers in explaining why the optimal solution is optimal. Ideally, schedulers and those affected by the schedule would prefer a series of cause and effect explanations for the solution. Unfortunately, in the worst case, the only explanation for a linear programming solution is that it is the solution of n equations and n unknowns with the other variables set to their bounds. Fortunately, OAT can bridge this gap most of the time by taking advantage of the special structure of many constraints. OAT provides a high level grid view (`Figure 5`) identifying the priority level that fixed the value of key decision variables, such as reservoir storage and release, for each object at each time period.

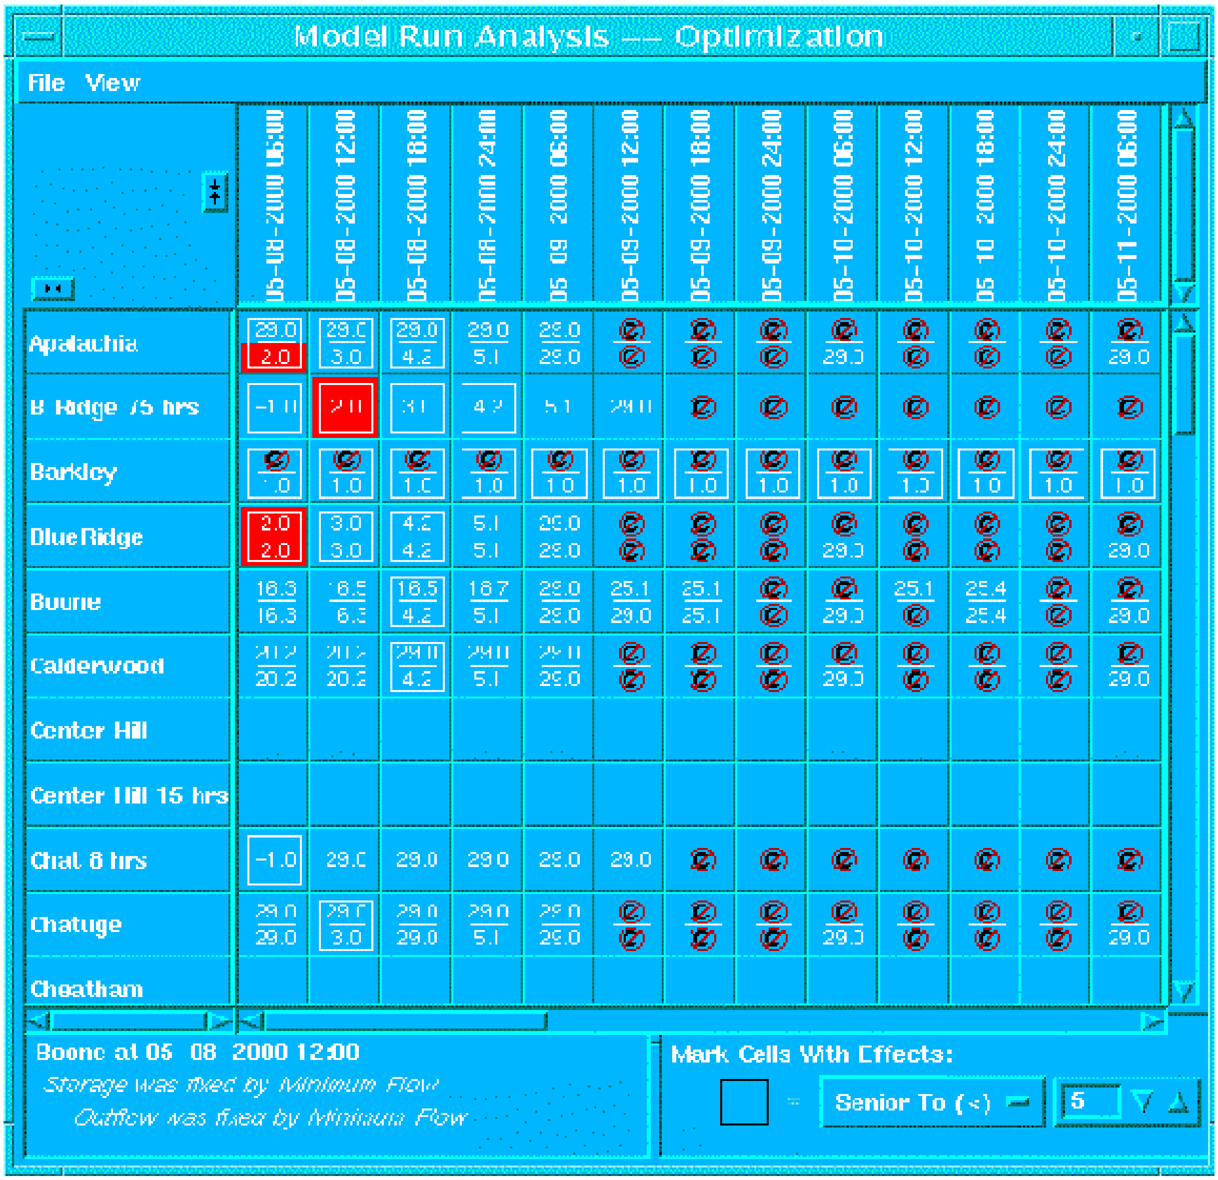

Figure 5: Top-level view of Optimization Analysis Tool

In addition, OAT has a more detailed view (`Figure 6`) for each object – time period combination that suggests the constraints that might be "responsible" for the value of each variable on that object at that time. More details on OAT can be found in Magee et al. (2001).

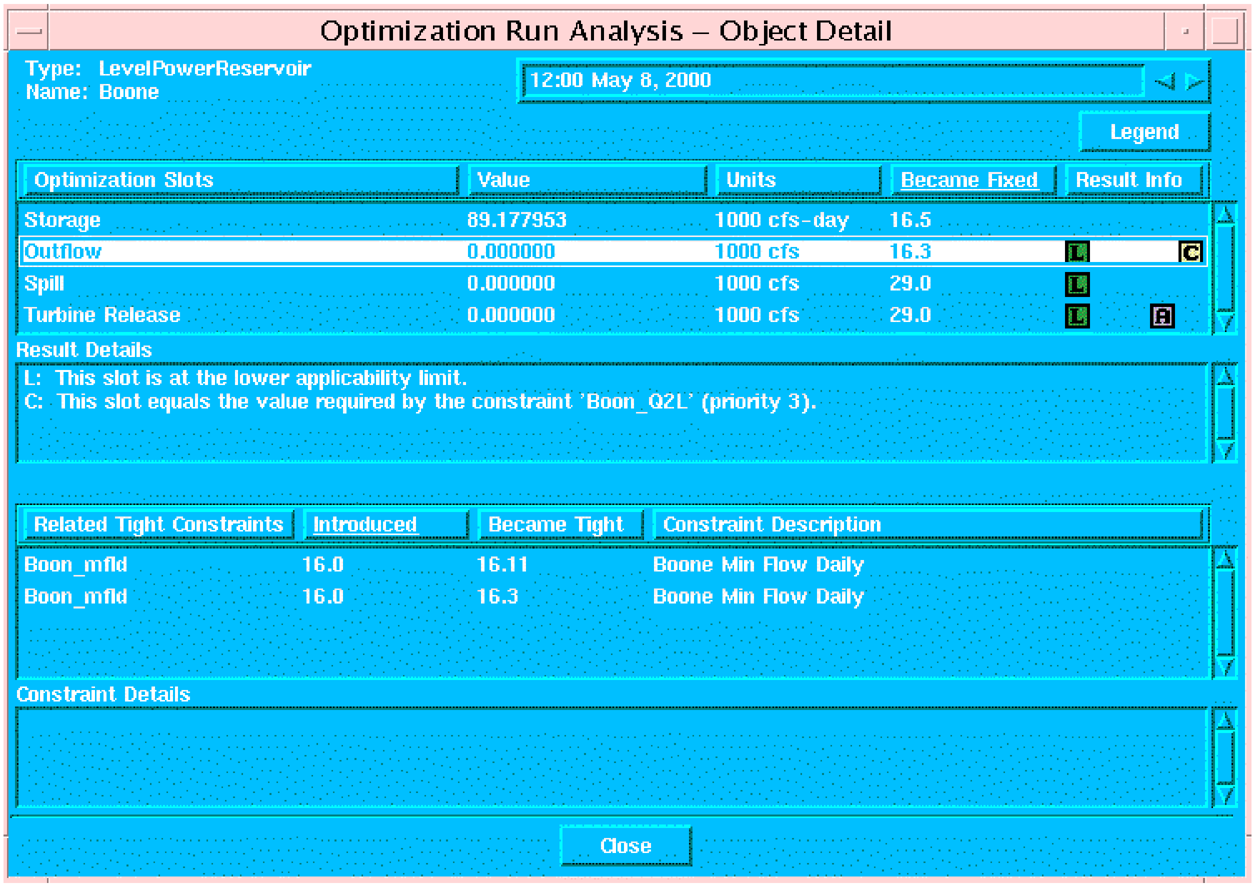

Figure 6: Detailed View in Optimization Analysis Tool

Simulation runs follow all of the optimization runs. These runs serve several purposes. Simulation eliminates small linearization errors in optimization. Additional processes are modeled in simulation that are not included in optimization. Schedulers can manually adjust any part of the optimization solution.

TVA SCHEDULING PROCESS

TVA's scheduling process has evolved as RiverWare has evolved. TVA was the first agency to adopt RiverWare's simulation controller in 1996. During this period, TVA's scheduling process was largely similar to manual scheduling enhanced with tools to improve efficiency. When the optimization controller was completed, it was tested and refined in side-by-side mode with simulation. In 1998, optimization initialized the scheduling process with schedulers making substantial changes during simulation. With time some of the scheduler's expertise has been incorporated in optimization and post-simulation changes have become smaller. In 2001, the hourly optimization model was put into production. More detailed modeling, particularly of economic value, has led to the current set of three optimization runs. The evolutionary aspect of this modeling has been important to its success. Model use has led to important feedback for the development of model capability. In particular, as scheduling needs have changed they have been incorporated in modeling capability.

Using RiverWare has meant considerable automation of a previously more manual schedule process. Under the previous process, most of the schedulers worked on different parts of the river in parallel during a single shift. Far fewer schedulers are required for a single run with RiverWare. TVA has used this advantage to convert to 24/7 scheduling and staffing. The same total number of schedulers is employed, but the schedulers are now divided into 6 separate teams working different shifts. Under the 24/7 process, the entire modeling sequence can be repeated as either hydrologic, power, or economic conditions change, typically 2-3 times per day.

The RiverWare models are only part of a larger process. The scheduling staff is also responsible for forecasting unregulated inflows into the system, collecting forecasts of power value based on alternative sources of generation, scheduling hydropower from other rivers, working in real-time with parties affected by reservoir operations, and general monitoring of the hydropower system.

TVA's current scheduling process can be summarized as follows. The 6-hour portion of scheduling includes:

1. build the data sets including forecasts of inflows based on weather and hydrology,
2. simulate day 1 with new hydrology to determine day 2 beginning reservoir elevations,
3. calculate the Value of Project Storage for day 9 from POSE and seasonal elevation targets,
4. 6-hour optimization of days 2-9 with POSE data to determine ending reservoir elevations,
5. 6-hour optimization of days 2-9 with piecewise value of power to determine 6-hour releases, and
6. simulate days 1-14 based on prior runs and day 14 targets set by senior water engineers.

The hourly portion of scheduling includes:

1. simulate day 1 based on yesterday's plan plus any changes that have been made,
2. optimize days 2-3 using elevations from day 1,
3. simulate days 2-3 and export hourly generations and ending elevations,
4. transfer data to "preschedule" tool,
5. automated adjustment for minimum flow pulses and plant set points, and
6. manual review and double check for errors.

Bear in mind that this description is a snapshot of an evolving process. Historically, some manual post-processing has been eventually incorporated into the automated process freeing schedulers to improve on the overall process with either additional runs or more detailed modeling. We expect this process to continue. For example, efforts are under way to modify the hourly optimization model to directly include some constraints that are currently handled with manual post-processing of the generation schedule.

SUMMARY

TVA has successfully used RiverWare's Goal Programming algorithm to model reservoir operations. Schedulers have used method selections to configure individual water objects to model the desired physical processes and create the corresponding optimization model. Once these methods are selected, RiverWare automatically generates variables, constraints and linearizations of non-linear functions. The schedulers have defined policy constraints and objectives through RiverWare's constraint editor. RiverWare translates this policy into a corresponding goal program. In addition to suggesting reservoir releases, RiverWare provides a tool to help visually explain why the system solution is optimal. The automation and tools in RiverWare reduce the technical optimization burden on schedulers allowing schedulers to focus more of their attention on the larger process of river scheduling.

Currently, three RiverWare runs are performed as part of a larger scheduling process that is repeated several times per day as water and power conditions change. The schedulers work in six teams to run a 24/7 river scheduling operation that includes forecasting reservoir inflows, incorporating power forecasts, estimating the long-term value of reservoir storage, and manual revision of the solutions. Both the scheduling process and RiverWare's capabilities continue to evolve in response to changing scheduling needs.

REFERENCES

Eschenbach, E., Magee, T., Zagona, E., Goranflo, M., and Shane, R., 2001, "Multiobjective Daily Operations of Reservoir Systems via Goal Programming*, "Journal of Water Resources Planing and Management*, Vol. 127, No. 2, pp. 108-120, ASCE.

Magee, T., Zagona, E., Frevert, D., 2001, "Operational Policy Expression and Analysis in the RiverWare Modeling Tool*," Proceedings of the ASCE World Water and Environmental Congress*, Orlando, FL, May 20-24, 2001, ASCE

Zagona, E.A., Fulp, T.J., Shane, R., Magee, T., and Goranflo, H.M., 2001, "RiverWare: A Generalized Tool for Complex Reservoir System Modeling*," Journal of the American Water Resources Association*, Vol. 37, No. 4, pp.913-929, AWRA.

Zagona, E. and Magee, T., 1999, "Modeling Hydropower in RiverWare," *Waterpower '99*, Proceedings of the International Conference on Hydropower, ASCE.