Modeling Hydropower in RiverWare

Edith A. Zagona[1] and Timothy M. Magee[2]

Abstract

RiverWare™ is a river basin modeling tool that provides flexibility to model a range of timesteps with multiple solvers including simulation, rulebased simulation and optimization. RiverWare also provides a selection of methods for modeling physical processes. RiverWare's basic structure and solution approaches are described. RiverWare provides various hydropower modeling capabilities to effectively represent hydropower objectives in a broad range of operating and planning models. Several methods of modeling hydropower in simulation include simple equations for the entire plant, a peak/base method for longer timesteps, a detailed plant characteristics method and a method for modeling each individual generator. The economic value of hydropower is calculated in terms of the thermal replacement value in a thermal-hydro power mix. For optimization, hydropower linearization techniques are described, as well as an economic objective to maximize the thermal replacement value of hydropower, trading off the current value of generation against the future value of generation.

Introduction

RiverWare, a collaborative effort by the Tennessee Valley Authority (TVA), the U.S. Bureau of Reclamation (USBR), and the University of Colorado Center for Advanced Decision Support for Water and Environmental Systems (CADSWES), is a general multiobjective river basin modeling tool. RiverWare provides both simulation and optimization capabilities. RiverWare provides various hydropower modeling capabilities to effectively represent hydropower objectives in a broad range of operating and planning models. This paper describes RiverWare, emphasizing the representation of hydropower in simulation and optimization.

River Basin Objects and Their Methods

The basic building blocks of a RiverWare model are *objects,* which represent the features of the river basin. The objects are represented by icons on the graphical work-

[1]Senior Researcher, Water Resources, CADSWES, University of Colorado, Boulder, CO 80309-0421, ASCE Member
[2]Operations Research Analyst, CADSWES, University of Colorado, Boulder, CO 80309-0421

space which open to show *slots*, the data structures for variables and parameters associated with the physical process equations for that feature. For example, all reservoirs have slots called Inflow, Outflow, Storage and Pool Elevation, among others.

A model is constructed by dragging objects from a palette onto the workspace, naming and populating them with data, and linking them together to form the basin topology. Specifically, a slot on one object is linked to a slot on another object, permitting propagation of the information among objects via the links. Currently, the RiverWare palette contains the following objects and their primary water quantity processes modeled:

Storage Reservoir - mass balance, evaporation, bank storage, spill.
Level Power Reservoir- mass balance, hydropower, energy, tailwater, operating head.
Sloped Power Reservoir - Level Power Reservoir plus wedge storage for long reservoirs.
Pumped Storage Reservoir - Level Power Reservoir plus pumped inflow.
Reach - routing in a river reach, gains and losses; diversions and return flows.
Aggregate Reach - many Reach objects aggregated to save space on the workspace.
Confluence- brings together two Inflows to a single Outflow as in a river confluence.
Canal - bidirectional flow in a canal between two reservoirs.
Diversion - diversion structure with gravity or pumped diversion.
Water User - consumption and return flow from a user of water.
Aggregate Water User- multiple water users; diversion from river or reservoir.
Aggregate Distribution Canal - models supplies to off-line water users.
Groundwater Storage Object - stores and returns water from seepage or return flows.
River Gage - specified flows imposed at a river node.
Thermal Object - economics of thermal power system and value of hydropower.
Data Object - user specified data: expression slots or data for policy statements.

Each object contains various *methods* to model the physical processes of the river basin feature it represents. Objects have two types of methods: dispatch and user-selectable. *Dispatch Methods* map the input/output configuration specified by the user to the correct basic solution algorithm. For example, reservoir objects have many dispatch methods for solving the mass balance equations. If inflow and pool elevation are known, the dispatch method for solving for outflow and storage is invoked.

User Selectable Methods are alternative calculations for particular details of the physical process, selected through the graphical user interface. For example, there are several hydropower generation methods, each of which calculates power as part of the mass balance calculation. Examples are shown in Table 1. The object-oriented modeling approach permits this flexibility and has the added benefit of extensibility; new methods can be added and integrated quickly and easily. Because the methods are easy to add, it is possible to have a large selection. The methods may include ones that are particular to an agency or site, but necessary to that agency for institutional reasons.

Run Options

RiverWare may be run using pure simulation, rulebased simulation or optimization solutions. In addition, simulation and rulebased simulation can be run in concert with water quality calculations or water accounting, which are not described in this paper.

Table 1: Selected User Methods in RiverWare

Object Type	User Method Category	User Methods	
Reservoirs	Spill	Unregulated Spill Regulated Spill	Regulated Plus Bypass Unregulated Plus Regulated
Power Reservoirs	Power	Plant Power Unit Generator Power	Peak Base Power Simple Power
	Tailwater	Tailwater Base Value Only Tailwater Base Value Plus Lookup Table	Tailwater Stage Flow Lookup Table Hoover Tailwater
Reaches	Routing	No Routing Time Lag Routing Variable Time Lag Routing Storage Routing	Muskingum Kinematic Wave Muskingum-Cunge MacCormack

The optimization run is typically followed by a post-optimization simulation, automatically configured with the optimal reservoir releases as inputs to drive the simulation.

The start and end date/time of the run and the computational timestep are selected on the Run Control dialog. Timestep options are 1 hour, 6 hours, 12 hours, 1 day, 1 week, 1 month and 1 year.

Simulation and Rulebased Simulation Solutions

Pure simulation solves a uniquely and completely specified problem. Each object executes a dispatch method as soon as it has the correct combination of data. The dispatch method solves for the unknown slots on the object, and information is propagated across links to other objects. Too much (conflicting) information results in an error state and termination of the run. Not enough information results in parts of the model left unsolved. These situations are detected easily in messaging and run analysis utilities.

In cases where there are multiple links between objects, i.e., the boundary conditions are solved mutually by the two objects, the objects iterate until a solution meets the convergence criteria or the maximum iteration count is exceeded. An example is two reservoirs linked outflow to inflow and upstream tailwater base value to downstream pool elevation. The release capacity of the upstream reservoir depends on the downstream elevation, which in turn depends on the upstream release.

Objects solve at any timestep for which they receive new information, allowing flexibility in specifying models in which the solution does not propagate from upstream to downstream and forward in time. River reaches with time lags may solve for inflow given outflow, setting the inflow value at a previous timestep and propagating that value upstream. In addition, target operations on reservoirs meet a future target storage or elevation by adjusting the reservoir's outflow over a specified period.

Whereas in pure simulation the model is exactly specified, a rulebased simulation run does not have enough input data to solve the system. Additional information is provided by prioritized policy statements (rules) specified by the user. The rules are if-then constructs which examine the state of system (functions of values of slots on the objects) in the antecedent, and then set slot values depending on that state. A rule set is global and each rule has a unique priority, although it may pertain to only one or a few objects.

Rules are created in a graphical editor which guides the formulation of syntactically correct expressions. The rule language permits reusable functions which perform complicated calculations used in the rules's logic. Rule statements are parsed and interpreted, and then the instructions are executed by the Rule Processor. The simulator propagates the effects of the rules in the model and does not allow lower priority rules to override higher priority rules.

Hydropower Modeling in Simulation

The Power Reservoir objects (Level Power Reservoir, Sloped Power Reservoir and Pumped Storage Reservoir) contain four selectable methods for simulating hydropower. Hydropower is calculated along with turbine release in the reservoir's mass balance calculations according to the user method selected.

Simple Power Method

The Simple Power method models power, P, according to the relationship

$$P = \alpha(OH,QT) \cdot QT \cdot OH \qquad \text{(Eq. 1)}$$

where α is an empirical coefficient which captures the properties of water and the plant efficiency, QT is turbine flow, and OH is operating head, given by headwater elevation minus tailwater elevation.

Peak Base Power Method

The Peak Base Power method computes the power and energy generated by the entire plant based on the fractions of each timestep operated at peak flow and base flow. It is a long timestep method, modeled after the U. S. Bureau of Reclamation's monthly CRSS peak-base power calculation. A peaking flow value is first determined from outflow, tailwater elevation and best efficiency turbine flow, QT_{best}. A minimum base flow, to meet environmental or other needs, is calculated as a function of total outflow. Maximum turbine flow, QT_{max}, is used if additional flow is available after best efficiency peaking. The number of hours to operate at peak power is calculated from the remaining volume of water released during that timestep; maximum turbine discharged is used if outflow does not. Peaking production and base production are then added to determine the total energy and plant efficiency.

Plant Power Method

The Plant Power method uses data tables giving QT_{best} and QT_{max} vs. OH, and tables relating power coefficient to operating head for the best efficiency (PC_{best}) and maximum capacity (PC_{max}) points. The power coefficient, PC_t, is the ratio of a unit of power generated to a unit of flow through the turbines.

For a given or calculated total reservoir outflow at time t, the method calculates the turbine release, QT_t, and power coefficient, PC_t. If outflow is greater than QT_{max}, QT_{max} is used and the remainder is spilled. PC_{max} is used to calculate power. If outflow is less than QT_{best}, best efficiency is assumed, and PC_{best} is used to calculate power, assuming turbines will run for the portion of the timestep needed at maximum efficiency. If outflow is between QT_{best} and QT_{max}, an interpolated PC_t is calculated. Power is then calculated as

$$P_t = PC_t \times QT_t \qquad \text{(Eq. 2)}$$

Plant capacity fraction is a timeseries of values between 0.0 and 1.0 which allows the plant capacity to be reduced. If the Plant Capacity is less than 1.0, QT_{max} and QT_{best} are reduced proportionally in calculating PC_t.

Experience indicates that plant characterization data is often not accurate, so the method provides a way to "calibrate" the calculation to current reported values from the plant. The user may input PC_t, in which case the input value is used. Another parameter that can be input is a power limit. If a power limit is specified and exceeded in the calculations described above, the QT_t, PC_t, and power are scaled down accordingly.

Instead of specifying outflow directly, the user may specify an energy value or input a flag indicating that the plant operates at its best efficiency or maximum capacity points.

Other variables are defined for use in evaluating the economic benefit of the hydropower. Part of the economic objective is direct loss to spill and the opportunity cost of current generation. The opportunity cost is modeled in RiverWare by penalizing water withdrawals with an estimated future value of power at each project. Typically, the value of power depends on the system energy in storage. RiverWare tracks the energy in storage at each project and calculates the system total.

Spilled energy is the energy that could have been produced by the spill, QS.

$$\text{SpilledEnergy} = PC_t \times QS_t \qquad \text{(Eq. 3)}$$

Energy in storage for each reservoir is an estimate of the stored water's potential for producing energy when it is released from that reservoir as well as downstream projects. It is calculated at each timestep using a table of storage vs. energy in storage values.

Future value of water data is input by the user as a single value per unit of energy. It is multiplied by the energy produced at each timestep to calculate the future value of used energy. Spill cost is the future value of water times the spilled energy.

Unit Generator Power Method

The Unit Generator Power Method provides for specification of *n* unit types, and *m* units of each type. For each unit type, data tables of QT_{best} and QT_{max} vs. operating head, and $Power_{best}$ and $Power_{max}$ vs. operating head are required. Each generator has an availability index between 0 and 1, indicating the portion of the timestep that the unit is available, and a power limit that caps the power generated by the unit.

The method assumes or uses an input QT_t, then, starting with the most efficient unit type at that operating head, adds units operating at best efficiency. Units in order of decreasing efficiency are added until the turbine release is met. If all units are used and QT_t is not achieved at best efficiency, a similar calculation commits units at maximum capacity. A plant PC_t is calculated and displayed. The economic evaluation variables are also calculated for this method.

Thermal Replacement Value of Hydropower in Simulation

The Thermal Object holds data pertaining to the cost of producing power by thermal sources in the system, and finds the economic value of hydropower in terms of the marginal effects on the thermal plants. Some power sources, such as pumped storage and contracts to allocate power from other utilities, can reasonably be considered either as part of the thermal unit commitment problem or part of the hydropower scheduling problem. RiverWare allows these sources to be modeled in either way. Any combination of the "alternate power sources" including conventional hydropower, pumped storage hydropower, and other (fixed) allocated sources can be evaluated in terms of the thermal energy it replaces in meeting the system load. The hydropower and economic evaluation variables of all power reservoirs and pumped storage reservoirs in the model are linked to, and summed on, the thermal object. Two methods are provided for characterizing the cost of producing thermal power.

Thermal Unit Method

The Thermal Unit method characterizes the thermal system by thermal units with fixed capacities and costs. Each unit also has a timeseries of availability factors between 0.0 and 1.0 to capture unit commitment decisions in the thermal system. The system load varies in time to capture peak and off-peak variations. The gross replacement value (GRV) of each alternate power is the cost of producing power for the most expensive units of thermal power it can replace. The user selects which alternative energy sources are included in the evaluation and the order in which they replace the thermal units.

Hydro Block Method

The Hydro Block method addresses the problem of unit commitment by representing the thermal replacement value of hydropower and other alternative energy as a piecewise linear function of energy. This representation at each timestep implicitly reflects the load, the thermal unit commitments, and alternative energies for which the economic evaluation is not to be included. This approach extends the conventional "system lambda" approach which is linear rather than piecewise linear, hence assumes that hydrogeneration does not affect the marginal value of power. The GRV is found directly from the data for the given hydropower and/or other alternative energies produced.

Using either method, the net replacement value (NRV) of the designated non-thermal energy sources at each timestep is found by subtracting the opportunity cost of power used and spilled

$$NRV_t = GRV_t - \sum_r \text{FutureValueofEnergyUsed}_t - \sum_r \text{SpillCost}_t \qquad \text{(Eq. 4)}$$

Optimization

RiverWare utilizes pre-emptive goal programming using linear programming (LP) as an engine to optimize each of a set of prioritized policy goals input by the user. The goals are input through the graphical, syntax-directed Policy Editor. Each goal is either a simple objective or a set of constraints which is turned into an objective to minimize deviations from the constraints. The optimizer solves over all timesteps; the solution trades off current costs and benefits against future costs and benefits.

Objectives and constraints are formulated using *policy variables*, a subset of RiverWare's slots. Decision variables in the LP include outflow, turbine release, spill and storage slots on reservoirs. Policy variables such as hydropower, pool elevation, and backwater elevation on sloped reservoirs are nonlinear functions of the decision variables. RiverWare automatically reformulates objectives containing non-decision policy variables as linear expressions of the decision variables, using data provided by the modeler. Nonlinear relationships are linearized using specified approximation points.

RiverWare's optimization software also uses data and links between objects to automatically generate physical constraints on the system reflecting mass balance, continuity, and upper and lower bounds of the variables. These automatic features in RiverWare's optimization software allow the user to focus on expressing the policy in the goals, and make it possible for water resources engineers without an optimization background to generate and solve goal programming models.

In addition to policies governing flows, elevations, spill and other variables in the physically-based model, an economic power objective may be added to the analysis through methods developed on the Thermal Object. Typically, after all water policy goals have been satisfied as fully as possible, a water management agency or utility will optimize hydropower's economic benefit. A policy objective can be specified which maximizes the avoided cost of producing thermal power resulting from hydropower generation and the other sources of alternate power.

When the optimization run is executed, the physical constraints are generated and sent to the solver as the highest priority objective. Then, each user-specified goal is interpreted, linearized, and optimized. For each goal, the solutions of the higher priority objectives are maintained as constraints. The optimal solution, the values of the decision variables, are returned from the solver and entered into the slots on the object after each goal. After all goals are optimized, a post-optimization simulation run is automatically set up with the optimal reservoir release schedule as inputs. The simulation allows the modeler to check and refine the optimization solution.

Hydropower Linearization Methods

Hydropower on reservoir *r* at time *t*, P_{rt}, may be used as a policy variable in constraints and objectives. Three user-selectable methods are available for linearizing this variable on a plant level. Data for the linearizations are generated from the PC_{max}, PC_{best}, QT_{max} and QT_{best} data tables.

Simple Power Linearization

The simple power linearization method assumes a fixed operating head (the initial operating head is assumed) and defines P in terms of the decision variable, QT. Based on the PC and QT tables, a P vs. QT relationship is developed for each power reservoir as shown in Fig. 1. Power is approximated by a piecewise linear function with segments indexed by *i*, as

$$P = \sum_i a_i QT_i \qquad \forall i \qquad \text{(Eq. 5)}$$

$$QT = \sum_i QT_i \qquad \forall i \qquad \text{(Eq. 6)}$$

$$0 \le QT_i \le u_i \qquad \forall i \qquad \text{(Eq. 7)}$$

where u_i is the maximum length of the segment. The power linearization can be characterized by any number of piecewise segments, although two (as shown in Fig. 1) or three are most common.

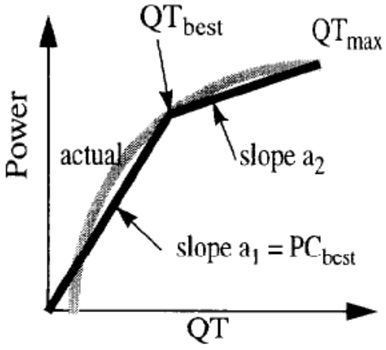

Figure 1. Power Linearizations as function of turbine flow

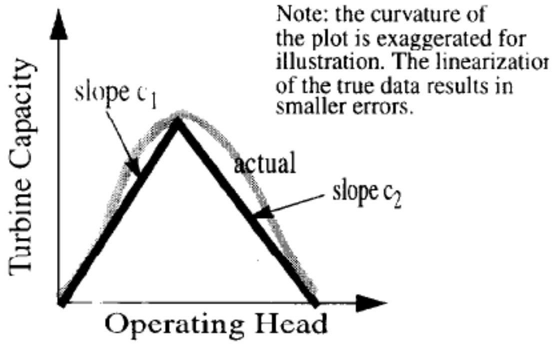

Figure 2. Turbine Capacity Linearizations as function of operating head.

Turbine capacity constraints are applied automatically in RiverWare as part of the physical constraints on the system. Typically these constraints become tight either as a result of no-spill policy constraints or an economic objective. QT is a convex nonlinear function of operating head, approximated by two or three linear segments as shown in Fig. 2, and represented by equations similar to Eqs 5 - 7. OH is translated to a function of headwater elevation minus a (constant) tailwater elevation. Headwater elevation is in turn translated to storage, a decision variable.

Estimated Operating Head

If the user can estimate the operating head at each timestep, the power approximation in Eqs. 5 - 7 can be a time-varying approximation with less error.

Convex Combination of Operating Points

The Lambda Power Linearization method was developed to provide improved approximation at a computational cost for reservoirs where head fluctuation is an important control variable for optimizing power generation. Possible operating points are generated from combinations of outflow, tailwater base value, and headwater elevation. Input data typically includes estimated minimum, maximum and median values of the headwater elevation for the run, and approximately three similar values for tailwater base value. (Tailwater base value is the downstream elevation to which a function of outflow is added to define the tailwater elevation.) RiverWare automatically adds three outflow values: zero, best efficiency turbine flow, and turbine capacity. More outflow points can be added by the user, for example, if spill is anticipated. The operating points based on these values are calculated using the PC and QT data tables. Each operating point consists of an operating head, outflow, tailwater, turbine release, spill, and power value. The optimization solution finds a convex combination of these possible points, by finding optimal λ values, where λ_i is the fraction of each operating point used in the solution. The generated constraints at each timestep *t* for *j* operating points have the form

$$P_t = \sum_j \overline{P_j} \lambda_{tj} \quad \text{(Eq. 8)}$$

where $\overline{P_j}$ is the power for operating point j, calculated prior to the optimization. Similar constraints are added for QT, headwater elevation, tailwater elevation and storage. The final constraint requires the operating point fractions to sum to one.

$$\sum_j \lambda_{tj} = 1 . \quad \text{(Eq. 9)}$$

Hydropower Economic Objective

RiverWare's Thermal object provides a policy variable, Net Avoided Cost, to formulate an economic objective to maximize the thermal replacement value of the hydropower and other alternate power sources. Typically, this objective,

$$\text{Maximize} \sum_t \text{NetAvoidedCost}_t \quad \text{(Eq. 10)}$$

identifies the best operations when the system of reservoirs is operating in a "normal" range, i.e., it satisfies most system level and flow constraints to meet water supply, flood control, environmental, navigation, etc., and still has some operational flexibility. The economic value represents a trade-off of the value of immediate hydropower generation against future expected value of water in storage. The objective is interpreted in either of two ways, based on which Thermal Replacement Value method is selected.

Thermal Unit Method Objective

If the Thermal Unit method is selected, Eq. 10 is interpreted by RiverWare's optimization software as

$$\text{Maximize} \sum_t \begin{pmatrix} \text{ThermalCost}_t \\ \text{without} \\ \text{Hydropower} \\ \text{Energy} \end{pmatrix} - \begin{matrix} \text{Cost}_t \\ \text{ofThermal} \\ \text{Units} \\ \text{Replaced} \end{matrix} - \begin{matrix} \text{Future} \\ \text{Valueof} \\ \text{Water} \\ \text{Used} \end{matrix} \Bigg) \quad \text{(Eq. 11)}$$

subject to, for all t,

$$\begin{matrix} \text{Usageof} \\ \text{UnitsReplaced}_t \end{matrix} + \begin{matrix} \text{Hydropower} \\ \text{EnergyProduced}_t \end{matrix} - \begin{matrix} \text{Pumping} \\ \text{Energy}_t \end{matrix} + \begin{matrix} \text{Allocated} \\ \text{Energy}_t \end{matrix} = \begin{matrix} \text{System} \\ \text{Load}_t \end{matrix}$$

The objective, Eq. 11, can replace thermal units with conventional and pumped storage power and allocated energy as indicated by the modeler. Thermal Cost without Hydropower Energy, is a constant. The Cost of Thermal Units Replaced is the unit cost per MW times the unit usage.

$$\begin{matrix}\text{Cost}_t\\ \text{ofThermal}\\ \text{Units}\\ \text{Replaced}\end{matrix} = \text{CostCoeff} \times \text{UnitUsage} \qquad \text{(Eq. 12)}$$

The future value of water used and released is

$$\begin{matrix}\text{FutureValueof}\\ \text{WaterUsed}\\ \text{andSpilled}\end{matrix} = \sum_{rt}[(\text{Spilled}_{rt} + \text{QT}_{rt}) \times \begin{matrix}\text{FutureValue}\\ \text{ofWater}\end{matrix}_r] \qquad \text{(Eq. 13)}$$

Unit usages are the capacity of the *i* units replaced by hydropower; they are decision variables, constrained by

$$0 \le \begin{matrix}\text{Unit}\\ \text{Usage}_{i,t}\end{matrix} \le \begin{matrix}\text{Unit}\\ \text{Capacity}_{i,t}\end{matrix} \qquad \text{For all} \quad i, t \qquad \text{(Eq. 14)}$$

Allocated energy is constrained to meeting the user-specified availability of the (e.g., purchased) power. For example, daily values of allocated energy may be specified for an hourly or 6-hourly model. RiverWare optimizes the distribution of that energy among the timesteps in the day.

The optimizer does not order the alternate power sources in the solution, as in simulation, but the user can specify which alternate power sources are used to replace thermal units.

Hydro Block Method Objective

If the Thermal Unit method is selected, Eq. 10 is interpreted by RiverWare's optimization software as

$$\text{Maximize} \sum_t \left[\begin{matrix}\text{HydroBlock}\\ \text{Usage}\end{matrix}_t \times \begin{matrix}\text{Replacement}\\ \text{Value}\\ \text{ofBlocksUsed}_t\end{matrix} - \begin{matrix}\text{Future}\\ \text{Valueof}\\ \text{Water}\\ \text{Used}\end{matrix}_t \right] \qquad \text{(Eq. 15)}$$

subject to

$$\begin{matrix}\text{Hydropower}\\ \text{EnergyProduced}_t\end{matrix} - \begin{matrix}\text{Pumping}\\ \text{Energy}\end{matrix}_t + \begin{matrix}\text{Allocated}\\ \text{Energy}\end{matrix}_t = \begin{matrix}\text{Capacityof}\\ \text{HydroBlocks}\\ \text{Used}\end{matrix}_t \qquad \forall t \quad \text{(Eq. 16)}$$

with the same optional treatments of allocated energy as for the thermal unit method.

Summary

RiverWare can be used for planning and operating multiobjective river and reservoir systems using simulation or optimization. Hydropower objectives can be modeled in various ways. In simulation, RiverWare models hydropower using several methods appropriate at various timesteps and ranging from general empirical relations to detailed modeling of each generator. Economic evaluation methods report the thermal replacement value of the hydropower, and other alternate power sources, used. The optimizer includes several linearization methods for accurate representation of hydropower in the solution, and meets and economic objective to maximize the thermal replacement value of the hydropower produced.