Reprinted from the Proceedings of the First Federal Interagency Hydrologic Modeling Conference, Las Vegas, Nevada, April 19-23, 1998, pp. 5-113-120.

# RIVERWARE:
# A GENERAL RIVER AND RESERVOIR MODELING ENVIRONMENT

**By Edith A. Zagona, Research Associate, Center for Advanced Decision Support for Water and Environmental Systems, University of Colorado, Boulder, CO; Terrance J. Fulp, Operations Research Analyst, U.S. Bureau of Reclamation, Lower Colorado Region, Boulder, CO; H. Morgan Goranflo and Richard M. Shane, Technical Specialists, River System Operations, Tennessee Valley Authority, Knoxville, TN**

**Abstract:** A general river basin modeling environment for operations and planning requires a high degree of software flexibility to allow users to model any river basin, manage data input and output efficiently enough for near real-time operations, and provide a selection of solution algorithms, all through a user-friendly interface. RiverWare is an extensible, maintainable software framework which provides a modeling environment to meet all the modeling needs of managers and operators of river and reservoir systems.

## INTRODUCTION

Water management agencies and utilities face increasingly difficult challenges in managing water resources. Environmental considerations, increasing demands on dwindling water supplies, outspoken recreational interests, the specter of climate change, and the restructuring of the power utility industry all have converged at a time when federal resources for developing modeling tools are minimal. Planning and operational river basin models developed in the previous decades are often not adequate to represent the changing multiple objectives of the projects and cannot be updated without significant expense.

To meet this challenge, the U.S. Bureau of Reclamation (USBR) and the Tennessee Valley Authority (TVA) are investing in a project with the Center for Advanced Decision Support for Water and Environment Systems (CADSWES) at the University of Colorado (CU) in Boulder to develop a general river basin modeling tool which can be used for a wide range of applications. The tool, called RiverWare, has been developed and applied to several basins by the two sponsoring agencies, and continues to be enhanced and improved.

To meet the goal of providing a modeling tool that can be applied to any river basin for both operational and planning applications, the RiverWare software was designed to meet the following general requirements:

- Be flexible enough to use for a variety of applications including daily scheduling, operational forecasting and long-range planning. This requires a range of timestep sizes and appropriate physical process modeling variability to support this range.
- Support various modeling solution methodologies. An organization's decision as to which basic modeling approach to use, simulation or optimization, depends on the specific goals of the model as well as the traditional way of looking at their system. Offering a variety of approaches allows each organization to continue historical practices as well as explore new approaches.
- Allow tailoring of applications by providing many basin features and many alternative methods for modeling these features. Even more important than the range of available features and methods is the extensibility of the software to provide ease in adding new methods. It must be recognized that there will always be applications that require enhancing the software and that most

agencies have some computational methods to which they are wedded for institutional reasons.
- Represent policy as input data. Many older models are obsolete because the operating policies were hard-coded and mixed in with the physical process model where they cannot be easily changed, or in some cases, cannot even be understood. Easy policy evaluation and modification by the user must be seen as a basic requirement for all new modeling tools.
- Provide an easy-to-use interface. A water resources engineer should be able to build, run and analyze model results relatively quickly, easily and without excessive training requirements. An operations scheduler should be able to view selected data in a convenient format, make changes to the operation, rerun the model and analyze results quickly.
- Fit into existing data and model interfaces. Every water management organization has an existing framework of databases, real-time data, supporting models, reporting tools, etc. to which the model must be connected. A general modeling tool must be flexible enough to tailor the application to any existing or changing configuration.
- Be supported by an organization which provides continued maintenance, enhancements, user support and technology transfer.

In the remainder of this paper, we present several of the major features of the RiverWare modeling software.

## MODEL CREATION

**Objects and Slots**: The basic building blocks of a RiverWare river basin model are *objects* which represent the features of the river basin. The objects are represented by icons on the workspace which can be opened to show the list of *slots*, which are the variables associated with the physical process model equations for that feature. For example, all reservoirs have slots, among others, called Inflow, Outflow, Storage and Pool Elevation.

The user constructs the model on the graphical workspace by selecting objects from a palette, dragging the objects with the mouse onto the workspace, naming the objects, and linking them together. Objects are linked together to form the topology of the river basin using the graphical link editor. Specifically, a slot on one object is linked to a slot on another object. During the simulation, the solution process involves propagation of the information among objects via the links. Currently, the RiverWare palette contains the following objects and the main water quantity physical processes which they model in a river basin:

- *Storage Reservoir* - mass balance, evaporation, bank storage, spill;
- *Level Power Reservoir*- Storage Reservoir plus hydropower, energy, tailwater, operating head;
- *Sloped Power Reservoir* - Level Power Reservoir plus wedge storage for very long reservoirs;
- *Pumped Storage Reservoir* - Level Power Reservoir plus pumped inflow from another reservoir;
- *Reach* - routing in a river reach, diversion 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* - depletion and return flow from a user of water;
- *Aggregate Water User*- multiple Water Users supplied by a diversion from a Reach or Reservoir;
- *Aggregate Delivery Canal* - generates demands and models supplies to off-line water users;

- *Groundwater Storage Object* - stores water from 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.

Data required by the model is entered into the slots, of which there are three basic kinds. *Time Series Slots* contain data at specified times. The slot manages the time keeping and generates the time series for the data. The default time series inherits the start time, end time and timestep size from the Run Control dialog unless the user configures the time series differently on the slot. *Table Slots* contain functional relationship data such as area-elevation-volume tables, or simple parameter data required by the model equations. *Scalar slots* are single values. Data can be entered into slots manually by typing the number into the slot dialog, or by file import which can import an entire timeseries or table of values at once. Values can also be imported through the Data Management Interface (DMI) utility, which is described in a subsequent section.

**Units:** All internal representation of slot values and computations is done in the default SI units. However, the user may enter and display values in any selected unit of a similar type. For example, the internal RiverWare unit for all slots with unit type of FLOW is cubic meters per second (CMS). The user units are specified in the GUI and can be set to any other FLOW unit, e.g., cfs, acre-feet per day, etc.

**Methods:** Each object has a list of User-Selectable Method Categories. For each one, a method must be selected for the detailed modeling equations used in the physical process model. (Methods are described in detail in the Engineering Methods section below.)

**Run Control:** The intended model run is set up on the Run Control dialog. The start time and end time and the timestep size of the run are specified. Timestep size ranges from hourly to yearly. There is no limit on the range of the runs. RiverWare currently supports dates from 1800 to 2300 A.D. This time range could be extended easily without additional memory usage or performance deterioration. The user also selects the solution type or “Controller” on this dialog. The list of slots and User Methods which appear are dependent on the controller selection.

**Saving Models:** Models are saved as text files. All data including the objects and their names, their topologic arrangement on the workspace, all input data in the slots, user units, method selections, run control selections, and all GUI settings are saved in the model file. The results of a model run are saved optionally.

## ENGINEERING METHODS

The objects on the workspace represent features of the river basin. Methods on each object contain the physical process model for the feature. The methods are flexible in handling a variety of input/output combinations of the basic data. In addition, the specific equations and physical representation of the processes are variable to accommodate a wide range of timestep sizes, data availability and resolution requirements, and modeling preferences.

To accomplish this flexibility, each object has two basic types of methods. *Dispatch Methods* map the input/output configuration specified by the user to the correct solution algorithm. *User Select-*

*able Methods* are alternative model representations which are selected by the user through the graphical user interface (GUI). For example, all reservoir objects have many dispatch methods for solving the mass balance equations. If Inflow and PoolElevation are known, the dispatch method for solving for Outflow and Storage is invoked. In addition, the user may select from a number of Evaporation methods, each of which calculates the evaporation loss in the reservoir as part of the mass balance calculation.

This object-oriented modeling approach mirrors the object-oriented software implementation and both benefit from this technology. From a software perspective, the benefit is extensibility: new methods can be added and integrated quickly and easily. From a modeling perspective, the benefit is the flexibility gained by selecting the physical process modeling methods individually on each object. Since the methods are easy to add, it is possible to have a large selection which includes some methods which may be quite particular to how one agency models one site, but necessary to that organization for institutional reasons. Table 1 contains a few examples of RiverWare's objects and User Methods.

**Table 1: Selected User Methods in RiverWare**

| Object Type | User Method Category | User Methods | |
|---|---|---|---|
| Reservoirs | Evaporation & Precipitation | No Evaporation Pan and Ice Evaporation<br>Daily Evaporation | Input Evaporation<br>CRSS Evaporation |
| | Spill | Unregulated Spill<br>Regulated Spill<br>Unregulated Plus Regulated | Regulated Plus Bypass<br>Unregulated Plus Regulated Plus Bypass |
| Power Reservoirs | Power | Plant Power<br>Unit Generator Power | Peak Base Power<br>LCR Power |
| | Tailwater | Tailwater Base Value Only<br>Tailwater Base Value Plus Lookup Table | Tailwater Stage Flow Lookup Table<br>Tailwater Compare<br>Hoover Tailwater |
| Reaches | Routing | No Routing<br>Time Lag Routing<br>Variable Time Lag Routing<br>SSARR | Muskingum<br>Kinematic Wave<br>Muskingum-Cunge<br>MacCormack |
| Water User (on AggDiversion) | Return Flow | Fraction Return Flow<br>Proportional Shortage | Variable Efficiency |

In addition to water quantity modeling, RiverWare provides several options for water quality calculations. The user may select to model dissolved solids only, temperature only, or combinations of these and dissolved oxygen. If modeling total dissolved solids only, a simple, well-mixed model is available. Temperature and DO models use a 2-layer reservoir model and discretized reaches in which the water quality equations are coupled with hydraulic routing, either with or without dispersion.

A special object on the palette is the Thermal Object. This object evaluates of the avoided costs from replacement of thermal power by hydropower.

## MULTIPLE SOLUTIONS AND CONTROLLERS

Alternative approaches to modeling multi-objective river basins have been developed, discussed and debated by water management agencies and academicians over the years. RiverWare endeavors to provide both prescriptive and descriptive techniques which are easy to formulate, analyze and apply to real planning and operations problems. Three fundamental solution methods are provided in RiverWare: simple simulation, rulebased simulation, and optimization. The first allows straight-forward scenario runs in which user-supplied inputs drive the solution. In the other two solution techniques, operational policies drive the solution. All operational policies are part of the input data set to permit easy modification and evaluation. In addition, the user may track water ownership by creating a network of water accounts in parallel with the river basin topology, and solve the accounting network independently of the simulation. The specific details of these solution methods have been designed and implemented to assure ease of use in solving a broad range of modeling application problems.

**Simulation:** Pure simulation solves a uniquely and completely specified problem. Each object must have enough information to invoke and solve a Dispatch Method, but may not have too much information. The solution is based on an object-oriented modeling paradigm: each object waits until it has enough information to solve, then it executes its Dispatch Method. 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. In the 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.

Although the simulation clock advances forward in time, the objects may solve for any timestep whenever they receive new information at that timestep. This allows some flexibility in specifying models where the solution is not propagating 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 may be specified, where a future target storage is met by adjusting the reservoir’s outflow over a specified timeframe.

**Rulebased Simulation:** Whereas in pure simulation the model is exactly specified, in rulebased simulation there is not enough information on the objects to solve the system. The additional information is added by prioritized policy statements (rules) which are specified by the user, interpreted by the rule processor and which set slot values on the objects based on the state of the system. The rules themselves are basically if-then constructs which examine the state of system (functions of values of slots on the objects) in the antecedent (if) clause and then set slot values depending on that state. The rule set is global in that each rule has a unique priority even though it may pertain to only one or a few objects.

The rules are expressed through the graphical Structure Editor which helps the user formulate syntactically correct statements. The rule language permits the creation of functions which may perform complicated calculations to support the decisions made by the rules. The rule statements are parsed and interpreted, and the instructions are then executed by the Rule Processor.

The interaction between the simulation and the rules at each timestep is as follows. The model simulates until all objects have executed all the Dispatch Methods they can, given the user inputs. Then the Rule Processor fires the highest priority rule. The rule may fail if some of the slot values it needs are not yet set, or if it tries to overwrite values set by a higher priority rule. If the rule fails, the next highest priority rule is fired, and so on until a rule is successfully executed and new slot values are set on the model. After a rule fires, it is taken off the list of currently active rules until any of its dependencies (slots it accesses in the antecedent clause) change. After the rule fires, the simulation continues until it has solved everything it can, then the Rule Processor is invoked again. This continues until there are no more rules which can fire at that timestep, then the clock advances and the next timestep is executed. Just as in pure simulation, there is no guarantee of a solution. RiverWare provides diagnostic tools, however, to aid the analyst in understanding which rules successfully fired, as well as which objects dispatched which methods.

**Optimization:** RiverWare's optimization utilizes pre-emptive goal programming, using linear programming (LP) as an engine to optimize each of the prioritized goals input by the user. The optimal solution of a higher priority goal is not sacrificed in order to optimize a lower priority goal. The goals are input by the user through the graphical Constraint Editor tool. Each goal can be either a simple objective, or a set of constraints which is turned into an objective to minimize the deviations from the constraints. RiverWare accesses the CPLEX mathematical programming subroutine library for the solving engine.

One of the challenges of optimizing river basin operations using LP is representing the nonlinear processes. RiverWare provides automatic linearizing of nonlinear variables. The user may formulate goals or constraints on a wide range of model variables (slots). The underlying optimization software reformulates the objective as a linear expression in the basic decision variables. The nonlinear relationships are represented by table data entered by the user. The advanced user can select alternative linearization techniques and parameters which result in more accurate linearizations. RiverWare's optimization software also takes advantage of the basic model data available in the objects and links on the workspace to automatically generate the physical constraints of the system which reflect the 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 an appropriate goal programming formulation.

In addition to policies governing flows, elevations, spill and other variables in the physically-based model, power economic objectives may be brought into the analysis through methods developed on the Thermal Object. The user may specify an objective which involves maximizing the avoided thermal cost from hydropower generation. The economic value represents a trade-off of the value of immediate hydropower generation against future expected value of water in storage. The current value of hydropower is defined as the avoided costs of thermal power resulting from the addition of hydropower to the overall power mix. As the most expensive thermal units are replaced by hydropower, the marginal cost of power generation decreases.

When the optimization run is made, the physical constraints are generated and sent to the solver as the highest priority objective. Then each user-specified goal is interpreted, linearized, and solved. For each goal, the solutions of the higher priority objectives are maintained as constraints. The optimization solution, the values of the decision variables, are returned from CPLEX and entered into the slots on the object. After an optimization run, the user can set up a post-optimization simulation run which automatically enters the optimal reservoir release schedule as inputs in the simulation and solves for storages, elevations, hydropower, etc. The simulation can be used to refine the optimization output.

**Water Accounting:** Water Accounts are created by the user on the objects. The sources of water to fill the accounts are supply links from other accounts and slots such as those which provide Hydrologic Inflow to the objects. Storage Accounts, Flow Accounts and Diversion Accounts all represent legal accounts in a water rights system. Pass-through accounts are created to keep track of the ownership of water in transit in the system.

In the initial prototype implementation of the water accounting system, the accounts calculate their balances as water is transferred from one account to another through the supply links. The account network solution behaves much like a spreadsheet in that it immediately updates the balances as data is entered. This allows for the accounting of water in an "after-the-fact" model. Future development will allow for prescriptive solutions of the network based on operations and water priorities.

**Controllers:** RiverWare's "controllers" are the software mechanisms for directing the model solution. The software architecture was designed to support any number of controllers. The controllers parallel the solution methods, but can also manage a combination of solutions. Currently the following controllers can be selected by the user on the graphical Run Control dialog where the user also specifies the begin and end dates of the run, and the timestep size.

- Simulation - can be run with or without Water Quality. Water Quality can be run In-line (at the end of each timestep) or post-process (at the end of the run).
- Rulebased Simulation - manages the interaction of the simulation with the Rule Processor. Water Quality can be coupled with rulebased simulation only in post processing mode.
- Optimization - pre-emptive linear goal programming solution; automatically sets up the post-optimization simulation if the controller is switched to "Simulation" after an optimization run.
- Accounting - executes accounting methods and solves accounts; currently this controller should be selected as a post-processor after a simulation run has been completed;
- Rulebased Simulation / Accounting InLine - solves each timestep for rulebased simulation then accounting, allowing the operational rules to access the previous timestep's accounting values.

## DATA MANAGEMENT INTERFACE

The successful application of any model for operational scheduling requires the model's data to be updated quickly and easily to reflect current data such as real-time measurements, inflow forecasts, scheduled hydroplant operations, and special operations. For a general modeling tool, the data communications must be possible regardless of the specific database application or the sources of other data. Similarly, planning studies often require a large number of runs with varied data which may come from other sources such as historical databases. Automatic importing and exporting of data can be achieved through RiverWare's Data Management Interface utility.

The Data Management Interface (DMI) in RiverWare provides the means of using external programs to automatically load data into RiverWare. These routines are written by the user or the user's organization in any programming language and invoked through the RiverWare GUI.

Scenario-based DMI's execute a suite of individual routines to bring many data sources together into the model. The DMI can also be used to advance the run control start and end times.

## MODEL ANALYSIS AND SCENARIO MANAGEMENT

The RiverWare modeling environment is designed to give the user the tools needed to build and run simulations to meet the needs of scheduling and planning activities. This is achieved through graphical user interface tools including the following:

- Simulation Control Table: a spreadsheet-like display of the data from a RiverWare model. The user constructs and configures an SCT, and can construct as many as needed to display various combinations of data. These can be iconified and brought up as different views on the model are needed. The SCT is interactive in that the user can specify inputs, run the simulation, and view the output from this interface.
- Diagnostics: information to assist the user in analyzing modeling problems during runs goes to an output dialog that saves the messages and offers searching and browsing features. The user may filter the information according to object, slot, time, etc. Warnings and errors are always displayed.
- Dispatch Info Table: a graphical tool showing what methods each object executed at each timestep and which slots were known and unknown. This is helpful to analyze over/under determination problems in simulation.
- Expression Slots: slots on the Data Object which hold user-defined expressions which are algebraic combinations of other slots, e.g., sum of all hydropower generation in the basin.
- Plotting: individual slots can be plotted quickly by selecting one button. Graphs with plot of many slots and/or many runs can be configured by the user. Plot set-ups can be saved.
- Snapshot Manager: the values of selected slots can be saved in this special data object for successive runs. If plotting the values, the new traces automatically are added to the plots. These data can also be accessed by the expression slots, the SCT and other output forms.
- Multiple Run Management: Many runs can be set up and executed in advance, automatically changing slot data or policies between runs. The runs can all have the same start/end times or be sequential. An Index Sequential setting on selected slots automatically permutes a series of historic data set as specified by the user (useful in using historic inflow data for planning studies).
- Excel Writer: a utility that takes outputs from RiverWare runs and creates an Excel input file for user-specific post-processing analysis.

## TECHNOLOGY TRANSFER

RiverWare is under continued development at CADSWES, with new versions released several times a year. User support is provided to sponsors, and a web-based bug-tracking facility allows users to log problems directly. It is the intention of CADSWES and the sponsors to maintain the software in the future and provide enhancements as requested. In addition, formal training sessions are given at CADSWES. A three-day introductory training session is available, and new training courses for Rulebased Simulation and Optimization are anticipated.

RiverWare is a c++ application which runs on a Sun workstations under the solaris 2.3 (or greater) operating system.