DEMONSTRATION OF RIVERWARE

Edie Zagona, Director, Center for Advanced Decision Support for Water and Environmental Systems, University of Colorado, Boulder, CO, (CU-CADSWES) zagona@cadswes.colorado.edu;
Tim Magee, Research Associate, CU-CADSWES, magee@colorado.edu;
Steve Setzer, Professional Research Assistant, CU-CADSWES, setzer@colorado.edu;
David Neumann, Professional Research Assistant, CU-CADSWES, david.neumann@colorado.edu

Abstract: RiverWare™ is a general purpose river and reservoir modeling application. Researched and developed by the University of Colorado Center for Advanced Decision Support for Water and Environmental Systems (CU-CADSWES) in collaboration with the Tennessee Valley Authority (TVA), the U.S. Bureau of Reclamation (USBR), and the U.S. Army Corp of Engineers (USACE), RiverWare can be run either as a stand-alone application, or in concert with other models and databases. RiverWare simulates the hydrologic response of the river system, given unregulated inflows and management decisions such as reservoir releases and diversions. Models are built using a palette of objects that represent features in the basin and user selectable methods that represent physical processes. There are three selectable solution methods in RiverWare: simulation, rulebased simulation, and optimization. In simulation, the user enters all of the data necessary for the objects to solve. Rulebased simulation builds on the simulation framework but the user is able to define a flexible ruleset of operating policies that sets data in the model using IF-THEN-ELSE type logic. By prioritizing the logic, the user is able to simulate complex operating objective. In addition, there is an optimization solution method that strives to meet the objectives using a linear pre-emptive goal programming solution. In addition to the solution controllers, it is possible to track water ownership and/or water type in a water accounting network. RiverWare provides a graphical tool to view and construct the accounting network. Recently, functionality was added to RiverWare to simulate the USACE flood control algorithms in rulebased simulation. The rule function and associated methods attempt to empty the flood pool, prevent downstream flooding and balance reservoirs according to the selected method. RiverWare provied tools to help users analyze model output such as the output manager, plotting, simulation control table, and rules analysis tool. Often, water management agencies develop RiverWare models that are of interest to stakeholders and other interested parties. Two features allow better interaction and communication between the water management agency and stakeholders: the scenario manager and the RiverWare viewer. This paper describes RiverWare, its solution approaches, and these features.

INTRODUCTION

Today's complex water resources management issues require flexible, comprehensive decision support tools that display timely information to water managers and assist in managing river systems for multiple objectives. River systems operate under laws, compacts, treaties and court decrees, while meeting increasing demands that compete for limited fresh water supplies. These demands include agricultural and M&I water use, flood control, hydropower generation, navigation, cooling water for thermal and nuclear power plants, fish and riparian habitat, and numerous types of recreational interests.

RiverWare is a general river and reservoir modeling tool researched and developed at the University of Colorado Center for Advanced Decision Support for Water and Environmental Systems (CU-CADSWES), and also collaboration with the Tennessee Valley Authority (TVA), several offices of the U.S. Bureau of Reclamation (USBR), and the U.S. Army Corp of Engineers (USACE). RiverWare is available to the public through the University of Colorado Office of Technology Transfer. Example applications of RiverWare include the Colorado River basin (Fulp and Harkins, 2001), the Truckee-Carson basin (Coors 2006), and the Upper Rio Grande basin (Sidlow et al., 2006).

RiverWare is a fully supported, documented, and maintained software product. It runs on both Microsoft Windows and UNIX operating platforms. See Zagona et al. (2001 and 2005) or visit http://cadswes.colorado.edu/riverware for more information.

This paper describes RiverWare from a user's perspective. First outlined is building a model. Next, the three solution controllers are described: simulation, rulebased simulation and optimization. Next the water accounting system is described. Finally, tools to help analysis and model sharing are presented.

BUILDING A MODEL

To develop a model, the user pulls objects off the palette onto the workspace, names them, links them together, selects desired physical process algorithms, and populates the slots with data. The object palette contains simulation objects such as storage reservoirs, power reservoirs, reaches, water users, pipelines, and distribution canals. Figure 1 shows the RiverWare workspace, palette, and object view. Following the object oriented modeling approach, each object contains its own data (called slots) and physical process algorithms in user selectable methods. Data structures include: timeseries slots, table slots, scalar slots, and user defined expression slots, and slots with results of statistical analysis of other slots.

Each object contains algorithms to represent physical processes. The modeler selects the desired algorithm, or "method," based on timestep size, resolution of data, institutional preferences, or other modeling needs. For example, every Reach object has about a dozen routing methods from which the modeler must make a selection. The methods allow each object to be simulated as needed. The user links objects together to form a network. Links are created between the desired slots on objects (e.g., the Outflow slot on a reservoir object may be linked to the Inflow slot on a downstream Reach object). Links are not directional but propagate information in either direction.

The user controls the time range of the run, i.e. the length of the timestep, the start time and the end time of the run. Available timestep lengths are: hourly, 6-hour, 12-hour, daily, weekly, monthly and annual. In addition, the user selects the solver to be used for the runs. The three main solvers are Simulation, Rulebased Simulation and Optimization. These solution controllers are described in the following sections.

SIMULATION

The simulation solution is executed entirely by the objects, each solving and propagating values from one object to another via links. The solution algorithm for each object consists of the equations that describe the basic hydrologic processes for the object, together with the specification for data needed to complete the solution. For example, a Reservoir object has as its primary equation, the conservation of mass: Storage = Previous Storage + (Inflows – Outflows)?t. Assuming that the previous storage is known, the equation can be solved for one of the remaining variables: Inflow, Outflow, or Storage, if the other two are known. Depending on the knowns and unknowns, the Reservoir will automatically select the correct form of the equation.

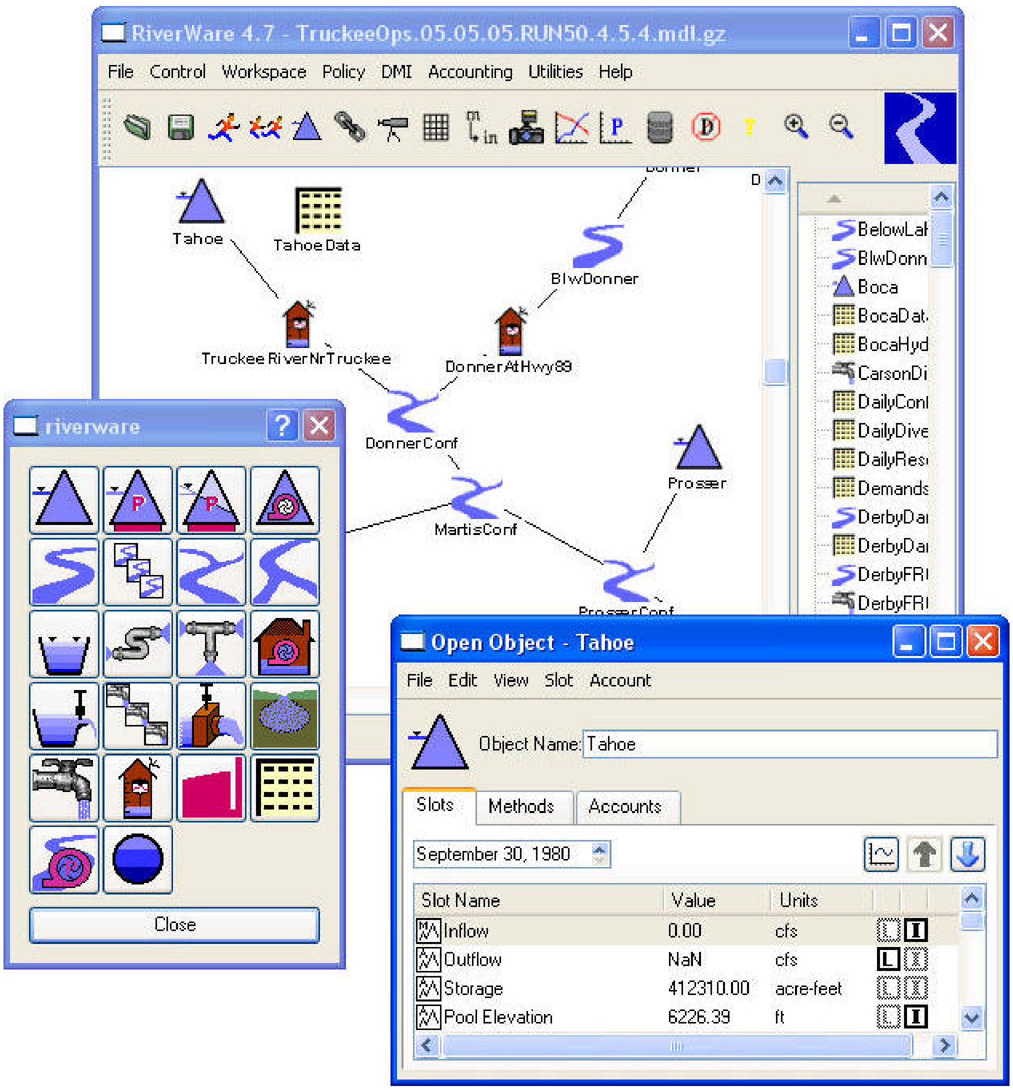

Figure 1 RiverWare Workspace, Palette, and Open Object view

Simulation is an exactly defined solution. The model must have the correct number of knowns and unknowns to solve. Conflicting information results in an error message and stopped run; too little information will result in an object not solving. In cases where multiple slot linkages make the objects mutually dependent on a solution, objects iterate until a solution meets convergence criteria.

RULEBASED SIMULATION

Rulebased simulation builds on the simulation framework but the user is able to define a flexible set of operating policies that sets data in the model thus causing the objects to solve. Using prioritized IF-THEN-ELSE type logic, the user is able to set values based on the state of the system. In general, the operating policies, called *rules*, contain logic for operating the system based on hydrologic conditions, time of year, demands, and other considerations. The operating policy is defined in a "ruleset." The rules are also organized in a prioritized fashion where only higher priority rules can overwrite values set by lower priority rules. By prioritizing the logic, the user is able to simulate complex operating objective. In addition, the user can organize the ruleset into groups of rules and functions that have to do with a geographic area, policy type, or by any other grouping. Individual rules or even groups can be toggled on or off so that policies can be tested.

The user defines the rules in the form of assignments, user defined functions, and pre-defined functions using the Riverware Policy Language (RPL). The RPL is a functional programming language that includes a palette of functions and expressions that guides the user through the construction of rules. Hundreds of predefined, compiled functions are available for the user. In addition, the user is encouraged to develop his or her own functions. Properly named functions allow the user to make the ruleset readable to stakeholders and decision makers.

Rulebased simulation begins with an under-determined simulation model, i.e., input values are insufficient to solve the model. Typically the modeler supplies the hydrology and lets the rules determine the operations. The rulebased simulation works as follows. Control is first given to simulation: any objects that have enough data to solve, based on user inputs only, do so. When no objects can solve, control switches to the rule processor. The processor starts with the lowest priority rule, executing rules until a rule results in an assignment to one or more slots on the workspace. For example, on a certain date, the outflow from a reservoir is set to the monthly minimum required release if specified conditions in the basin are true. After the assignment, control is given back to the simulation, and the objects solve based on the results of the rule assignment. To continue the example, the simulation system forces the reservoir to solve given outflow and the user specified inflow. When all objects having enough data have simulated, control goes back to the rule processor. As the rules and simulation continue to alternate, rule assignments are limited by priority such that lower priority rules cannot overwrite higher priority rules. Rules will re-execute if a value in the rule evaluation has been changed because of an assignment by another rule. The timestep is complete when all the rules have fully executed and all objects having enough data have solved.

OPTIMIZATION

RiverWare's optimization solution is a linear, pre-emptive goal programming algorithm that optimizes reservoir outflows for a prioritized set of user-specified objectives. See Eschenbach et al. (2001) and Biddle (2006) for more information on optimization and its application to the Tennessee Valley Authority system. It solves simultaneously over multiple time steps. Because RiverWare optimizes simultaneously over multiple time steps, some simulation processes cannot be modeled and some nonlinear processes must be approximated by linear functions.

Preemptive linear goal programming is an extension of linear programming, maximizing or minimizing a single linear objective with respect to a set of simultaneous linear constraints. A preemptive goal program replaces the single objective in linear programming with a series of prioritized objective functions. Each objective is optimized in order and then constrained to its optimal value. Thus, a high priority objective is never sacrificed for a lower priority objective, just as in rulebased simulation. Objectives may include explicit objective functions, such as maximizing the revenue from hydropower or minimizing spill. In addition, a set of constraints can be converted to an objective: specifically, minimizing the violation of the constraints.

Once formulated, RiverWare uses the commercially available optimizer, CPLEX, to solve the individual linear programs as part of the goal programming optimization. The optimal reservoir outflows from the goal program are returned to the RiverWare workspace as inputs and used to drive a simulation. This post-optimization simulation removes any approximation errors in the optimization and simulates processes that are not modeled in the optimization.

USACE FLOOD CONTROL

A policy commonly used by the USACE for flood management in a basin that has multiple reservoirs is to release the flood pool as quickly as possible, avoid flooding at downstream channel control points, and balance storages in the reservoirs. The algorithm finds a simultaneous solution for releases from all basin reservoirs, considering all control points, calculated over a specified forecast period with forecasted inflows. In RiverWare, this policy can be modeled along with other basin objectives in a ruleset.

There are two flood control methods implemented in RiverWare. The *Operating Level Balancing method* calculates flood-control releases using an algorithm that was developed by the Southwest Division of USACE and previously implemented in their SUPER model. This method balances storages in the reservoirs using numbered operating levels that are associated with specific storage or elevation levels on each reservoir. For example, all reservoirs associate the top and bottom of the flood pool with the same numbered operating levels. Additional levels are defined within the flood pool. Two reservoirs are considered to be balanced when their storages are at the same operating level number. The other USACE flood-control method, the *Phase Balancing method*, uses an algorithm that was developed by the Kansas City District. This method balances the storages in the reservoirs using "phases" that are associated with storage levels in the reservoirs.

The rule function *FloodControl* executes the selected method on the objects in the subbasin, then returns to the rule the list of reservoirs and computed flood-control release values. The rule assigns the values slots on the reservoirs. In the subsequent simulation, each reservoir solves its mass balance equation using the flood control release value as outflow, along with any other releases such as surcharge releases and hydropower releases. See Daylor et al. (2006) for more information on the algorithm and the implementation in RiverWare.

WATER ACCOUNTING

In many basins, it is necessary to track water type and water ownership, not just the physical volume of water. In RiverWare, this is accomplished by creating a water accounting network.

In RiverWare "physical" water is modeled in the simulation object network and "paper" water (water type and ownership) is modeled in the accounting network. For example, a reservoir's storage indicates the quantity of physical water in the reservoir at each timestep. The storage account on a reservoir shows the volume of water in the reservoir owned by a separate entity or water of a different legal type. In RiverWare, water accounts are created on the simulation objects and linked together to form a network of accounts. These links are called "supplies" and represent supply and demand relationship between two accounts. Both supplies and accounts in general have attributes. Accounts can have a "water type" and "water owner." Supplies can have a "supply type" and a "release type." These attributes are useful to describe the system and to facilitate rules accessing accounts. The accounting network is controlled by user inputs or by rules; i.e. the user sets the outflow from an accounts. Predefined and user defined methods control how water is introduced into and removed from the accounting network, typically using information from the physical simulation. For example, the physical inflow to a reservoir may be shared between two accounts but the reservoir evaporation may be charged to just one of the accounts. Currently there are three types of accounts: storage, pass-through and diversion.

In the past, users found that a complex accounting network with many accounts and supplies was difficult to visualize. Users were required to construct and maintain a diagram or flowchart of the network. Recently, functionality was added to RiverWare to present a graphical view of the accounting networking. This tool is extremely useful to help model developers communicate with decision makers and stakeholders in the basin. The accounting network is presented in a separate view from the standard schematic depiction of the physical network or workspace view. The accounting view provides a representation of the physical objects as well as the accounts on those objects. Drawing properties like line color or shading can be changed based on attributes of the accounts and supplies like water type or supply type.

ANALYSIS TOOLS

RiverWare provides tools to help modelers, decision makers, and stakeholders analyze a model run. This section lists and briefly describes the available tools.

Output Manager The Output Manager allows users to customize output devices like plots or data files for export.

Plotting RiverWare contains a fully functional tool to plot timeseries slots, table slots, and periodic slots. The user is able to customize and save the plotting configuration.

Snapshots The Snapshot facility allows selected data to be saved from sequential runs. Plots can be updated with the latest run results, allowing easy internal comparison of runs.

Model Run Analysis RiverWare provides a tool to show how each object solved. In rulebased simulation, it also shows the priority used to set the value that caused the object to solve. The display is configurable to color code timesteps and objects that were controlled by a certain rule.

Simulation Control Table (SCT) The user is able to define an SCT that provides a spreadsheet view of the model. In addition, the user can make runs and input or exported data directly from this table.

Rules Analysis RiverWare provides a tree-view of the RPL ruleset. This is useful to see the functions that call other functions and vice versa. In addition, RPL performance information is calculated and displayed.

Data Management Interface (DMI) RiverWare provides a Data Management Interface to connect a RiverWare model to any external database or data files. Input DMI's bring data into a model while output DMIs take data from a model run and write it to a database. A control file maps data variables in RiverWare to flat files. DMI's use an external executable developed by the user to get data from an external source and write to the data file, or to process data that is exported from RiverWare to the data files. DMI's can be executed manually or from a rule. Recent development will allow internal mapping of data directly to external data sources and sinks. Currently supported are direct data links to the Corps of Engineers Data Storage System (DSS) files and to Excel spreadsheets.

Diagnostics Comprehensive and detailed diagnostic messaging is possible and can be controlled by the user by filtering by objects, timesteps, etc. This powerful tools assists in debugging or understanding the details of the model computations.

Multiple Run Management (MRM) allows the automation of setup and execution of many runs that differ only in input data, or ruleset. The output of MRM runs can be analyzed statistically using the Graphical Policy Analysis Tool (GPAT), an Excel-based tool develop for RiverWare.

TOOLS TO ALLOW INTERACTION AMONG STAKEHOLDERS

Often, water management agencies use RiverWare models that are of interest to stakeholders and other interested parties. Two features allow better communication between the water management agency and stakeholders: the scenario manager and the RiverWare viewer.

Scenario Manager The scenario manager provides non-expert modelers (i.e. those not involved in the model development) a framework to enter sensible and feasible alternative input values, make model runs and compare the results of these alternative runs without impacting the underlying base model. Scenario analysis consists of three parts: a set of scenario slots (data), a baseline model, and a set of alternative input data describing the scenario. The scenario slot list is a set of slots to be used in the scenario. These slots contain the input data that the user wants to modify and test. Typically the model developer selects this list of slots and is able to specify minimum and/or maximum values for this data which helps to prevent unrealistic input data. The modeler then saves the model specifying that it is a "baseline model." Once a model has been saved as a baseline model, topology, methods, and non-scenario slots become uneditable and slots cannot be added or deleted from the scenario slot list. By preventing editing of the baseline model and the scenario slot list, the integrity of the baseline model is preserved and the scope of the scenarios is limited.

The user, typically the stakeholder, creates a scenario and enters data on the scenario slots. The scenario data must be in the range defined in the scenario slot. This scenario data is saved separately from the model file making it easily shared between users and model developers. The user loads the scenario into the baseline model and activates the scenario to make a model run.

Once a run has been made, output in the model file represents the scenario. The baseline model cannot be overwritten with this new data, but the results can be re-saved as a non-baseline run.

RiverWare Viewer In order to make RiverWare models accessible to interested parties who do not wish to invest in a license or the expertise needed to set up and execute runs, CADSWES supports a free "read-only" version of RiverWare called the RiverWare Viewer (referred to as the "Viewer"). The Viewer can load a saved model; all features of RiverWare that have to do with viewing or exporting data are functional, including creating and viewing SCTs; exporting data through DMIs; loading and examining Rulesets, and creating plots. However, data cannot be edited and runs cannot be made with the Viewer. In order to allow saving DMIs, plot configurations and other changes, Viewer users may re-save the model file. To track subsequent saving and re-saving, RiverWare's model comment utility records the user who last saved the model and the time it was saved, keeping this information in a non-editable chronology.

SUMMARY

RiverWare is a general tool used to simulate river basins. Included are three solution mechanisms: simulation, rulebased simulation, and optimization. These mechanisms provide flexibility to modelers to represent their basin and policy in various applications including operations, after-the-fact water accounting, short term forecasting, long term planning, and impact studies. Included also in RiverWare is an accounting system that tracks water ownership and water type. Finally, there are tools to analyze data and communicate the results with stakeholders and other interested parties. This provides a comprehensive yet flexible tool that helps water managers make informed decisions about the water operations.

REFERENCES

Biddle, S. and Magee T., (2006). "Managing TVA's Hydropower System Using RiverWare," Proceedings of the Third Federal Interagency Hydrologic Modeling Conference, Reno, NV.

Coors S. (2006). "Truckee-Carson Basin RiverWare Operations Model," Proceedings of the Third Federal Interagency Hydrologic Modeling Conference, Reno, NV.

Daylor, J., Cotter, J., Zagona, E.A., (2006). "Development and Use of USACE-SWD Flood Control and Hydropower Algorithms in RiverWare," Proceedings of the Third Federal Interagency Hydrologic Modeling Conference, Reno, NV.

Eschenbach, E.A., Magee, T., Zagona, E., Goranflo, H.M., and Shane, R. (2001). "Goal Programming Decision Support System for Multiobjective Operation of Reservoir Systems,"ASCE Journal of Water Resources Planning and Management, 127(2).

Fulp, T., and J. Harkins (2001). "Policy Analysis Using Riverware: Colorado River Interim Surplus Guidelines," Proc ASCE World Water & Environmental Res Congress, Orlando, FL.

Sidlow, M., Roark, M., Sanders, A., and Zagona, E.A., (2006). "Stakeholder participation in Target Flow Modeling on the Middle Rio Grande," Proceedings of the Third Federal Interagency Hydrologic Modeling Conference, Reno, NV.

Zagona, E.A., Fulp T.J., Shane R., Magee, T., and Goranflo, H.M. (2001). "RiverWare: A Generalized Tool for Complex Reservoir Systems Modeling," Journal of the American Water Resources Association, AWRA, 37(4), pp 913-929.

Zagona, E.A., Magee T., Goranflo, H.M., Fulp T., Frevert, D., and Cotter J. (2006). "RiverWare," Chp 21 of Watershed Models. Ed. Singh, V. and Frevert, D. CRC Press.