Texas A & M University and U.S. Bureau of Reclamation
Hydrologic Modeling Inventory
Model Description Form

August 2007

Name of Model

RiverWare®

Model Type:

Simulation, optimization, water accounting and water rights modeling of multi-objective reservoir and river systems

Modeling Objective(s):

RiverWare® is a general river and reservoir modeling tool for operational scheduling and forecasting, planning, policy evaluation, and other operational analysis and decision processes. RiverWare has the capability to model

- Hydrology and hydrologic processes of reservoirs, river reaches, diversions, distribution canals, consumptive uses, shallow groundwater interaction and conjunctive use;
- Hydropower production and energy uses; and
- Water rights, water ownership, and water accounting transactions.

Agency and Office:

Developed by the University of Colorado Center for Advanced Decision Support for Water and Environmental Systems (CU-CADSWES). Primary sponsors are the Tennessee Valley Authority, the Bureau of Reclamation and the U.S. Army Corps of Engineers. RiverWare is available to the public through the CU Office of Technology Transfer and CADSWES provides training and user support.

Contact:

CU-CADSWES
UCB 421
University of Colorado
Boulder, Colorado 80309-0421
(303) 492-3972
rwinfo@colorado.edu

Technical Contact and Address:

Dr. Edie Zagona, CU-CADSWES
303/492-2189
zagona@colorado.edu

Model Structure or Mathematical Basis:

Overview

RiverWare is a general object-oriented reservoir and river operations modeling tool. It can be used to model any system of rivers, reservoirs and other features such as diversions, canals, stream gages and groundwater interaction.

It takes as input the hydrologic runoff from the watershed at specified nodes. These inflows to the system, along with operations such as reservoir releases and diversions, drive the solution. Alternatively, river gage data can be input and system inflows can be calculated. RiverWare's object-oriented, data-centered approach enables the modeler to represent site-specific conditions by creating a network of simulation objects, populating each with data, and selecting physical process algorithms on each object that are appropriate to the purposes of the object and its representation in the overall model. Data-driven simulations solve the hydrologic variables by propagating solutions of flows and stage through the network. For multi-objective operational analysis and decision-making, RiverWare provides an interface for expression of operational policies as well as both descriptive and prescription solution algorithms driven by these policies.

RiverWare is a hydrologic model in that the primary processes modeled are mass balance of reservoirs, hydrologic routing in rivers, evaporation and other losses, diversions, return flows, and groundwater interaction. Hydropower and economics of hydropower are also a primary objective of the model. Any size system can be modeled from a single reservoir or river reach to an entire basin. Either single events or continuous operations can be modeled.

RiverWare Objects

The basic building blocks of a RiverWare model are objects that represent the features of the physical system. These are represented on the graphical workspace by icons. The user constructs a model of a specific river system by selecting objects from a palette, dragging them onto the workspace, naming the objects, and linking them together to form the river system topology. Each type of object contains the physical process algorithms for the associated river system feature. RiverWare provides the following objects and their primary physical processes:

- Storage Reservoir – models mass balance, evaporation, bank storage, release and various spill types.
- Level Power Reservoir – same as Storage Reservoir plus hydropower, energy, tailwater, operating head, energy in storage.
- Sloped Power Reservoir – same as Level Power Reservoir plus wedge storage for very long reservoirs and timelag routing through the reservoir.
- Pumped Storage Reservoir – same as Level Power Reservoir plus pumped inflow from another reservoir; either pumping or generating occurs during each timestep.
- River Reach – models various methods of routing, seepage, losses, diversions and return flows.
- Aggregate Reach – many River Reaches are aggregated to save space on the workspace.
- Confluence – brings together two Inflows to a single Outflow as in a river confluence.
- Bifurcation – a controlled splitting of the flow from a single Inflow to two Outflows.
- Canal – bi-directional flow in a canal between two reservoirs.
- Diversion – models a diversion structure with gravity or pumped diversion from a River Reach or reservoir.
- Water User – represents an agricultural or M&I user; models depletion (consumption) and return flow, supplemental groundwater pumping, and various methods for generating diversion request from input data.
- Aggregate Water User – multiple Water Users supplied by a diversion from a River Reach or Reservoir; several aggregation/linking structures such as lumped and sequential.

- Aggregate Delivery Canal – generates demands for a Diversion Object to which it is linked, and models supplies to off-line water users linked to the canal segments; routing flow through multiple linked canal segments, each of which can model seepage and losses.
- Inline Power Plant – generates power via head and flow, without storage.
- Inline Pumping Plant – provides head to water.
- Pipeline – simulates pipe flow.
- Groundwater Storage Object – receives water from Water User return flows for temporary storage and from seepage from River Reaches and Delivery Canals; releases water according to aquifer elevation or storage back to the river or to deeper groundwater, or supplies pumping demands from Water Users.
- River Gage – specified flows imposed at a river node.
- Thermal Object – this object represents the a thermal power system; its primary purpose is to represent the economics of the thermal system which play a role in determining the value of hydropower produced by the system's Power Reservoirs.
- Data Object – user-specified data: data for policy statements and user-tailored output combinations from various objects.

Slots and User-Selectable Methods

Objects can be "opened" to view the list of *slots*, i.e., variables and parameters associated with the physical process models. Also, each object has User-Selectable Methods. For example, the user selects a routing method on River Reach objects and a method for hydropower calculation on Power Reservoirs. Each object has general slots, e.g., all reservoirs have inflow, outflow and storage. Method-specific slots are associated with user-selectable methods, e.g., Muskingum coefficients are associated with the Muskingum Routing method.

Model Computation Techniques

RiverWare provides three basic solution approaches for the river/reservoir systems: pure simulation, rulebased simulation and optimization. Simulation and rulebased simulation can be coupled with Water Quality computations or Water Ownership (Accounting) modeling. The first step in model development is building a pure simulation model for calibration and verification. The simulation can be used for scenario-based runs. Alternatively, rules representing the system operational policies can be developed to drive the simulation. Optimization runs are always followed by simulation to predict the full non-linear effects of the optimal solution.

Pure Simulation

The simulation solution is executed entirely by the objects, solving in turn and propagating values from one object to another via the links. The solution method 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 equation of conservation of mass:

$$S_t = S_{t-1} + (\text{Inflows} - \text{Outflows})\, \Delta t$$

Assuming that the previous storage, $S_{t-1,}$ is known, the equation can be solved for one of the variables: Inflow, Outflow, or Storage, if the other two are known. The Reservoir knows the three possible forms of the equation and the data requirements (known and unknown variables) of each. During the simulation, the object receives inputs or values propagated across a link from another object that has solved. The Reservoir waits for new data, and when received, it checks the data that it has against the data requirements for the equations it can solve. When the Reservoir has the right combination of known and unknown variables for one of its equations, it solves the equation and writes the value of the solved variable into the slot. If the slot is linked, the value propagates to another object.

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

Pure simulation is useful for calibration and *if-then* scenarios to investigate the results of particular inputs. Some operators use RiverWare in simulation, inputting data rather than rules to come up with a solution. In this case, the operating policy is implicit in the inputs.

Water Quality

Water quality can be calculated after the physical water solution has been completed at each timestep or after the entire run. 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.

Rulebased Simulation

Rulebased simulation provides a means for simulation based on logical policy statements rather than explicitly specified input values for operations such as reservoir releases, storages, diversions, etc. In general, the operating policies, called *rules*, contain logic for operating the system based on hydrologic conditions, time of year, demands, and numerous other considerations.

Operational policy is expressed in the RiverWare Policy Language (RPL), an interpreted language developed for, and exclusive to, RiverWare. RPL is a functional language in which assignments (to slots) are made only at the highest level of the rules. The language is rich enough to express the highly complex logic needed for many basins. The main elements of RPL are:

- Assignments (a slot is set to the results of a logical expression);
- Object/slot values from the workspace;
- Mathematical and logical operators, including looping mechanisms;
- Date/time expressions;
- Built-in functions; and
- User-defined functions.

A rule is constructed in a syntax-directed editor that accesses a palette containing these elements.

In addition to convenience, the use of functions and slot names provides a means of simplifying the rules and making them comprehensible when read.

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. For the mathematical solver, RiverWare accesses the CPLEX mathematical programming subroutine library. The goals are input by the user through a graphical Constraint Editor. Each goal can be either a simple objective or a set of constraints that is transformed to an objective to minimize the deviations from the constraints. The goals/constraints specify limitations on the values of various slots (variables) on the objects. The objects reformulate the goals/constraints to be linear expressions of the basic decision variables. The advanced user can select alternative linearization techniques and parameters that result in more accurate linearizations. The objects automatically generate the physical constraints of the system that reflect the mass balance, continuity, and upper and lower bounds of the variables. 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 object may involve maximizing the avoided thermal cost from hydropower generation. After an optimization run, the user can select to have a post-optimization simulation run automatically set up in which the optimal reservoir releases are inputs to the simulation run that solves for storages, elevations, hydropower, etc. The post-optimization simulation run predicts the effects of the optimal release schedule without inaccuracies due to linear approximation.

Water Accounting

The user creates a separate network of water accounts on the objects. The accounts are linked together to form the possible transfers of the "paper" water. Storage Accounts, Flow Accounts and Diversion Accounts all represent the legal accounts in a water rights system. Accrual, carry-over, exchanges and other water accounting mechanisms can be represented. Pass-through accounts are created to keep track of the ownership of water in transit through the system. The accounts calculate their balances as water is transferred from one account to another through the supply links. The account network solution behaves like a spreadsheet in that it immediately updates the balances as data is entered. Rules accessed through the rulebased simulation solver can inspect values and set values on the accounts in addition to setting values in the physical water system. In the current implementation, the user controls the reconciliation between the physical and accounting systems through input or through the rules driving the solution.

Priority Water Rights Allocation Solver

A RPL function solves the global water rights problem through an iterative procedure that allocates flow and checks for violation of senior rights. The function will solve the entire allocation of the natural flow and return the solution to the rule set.

Representation of Uncertainty

Optional uncertainty calculations using first-order, second-moment propagation can be executed with the simulation solution. The simulation input values for the basic reservoir variables (inflow, outflow, storage, pool elevation) represent the expected values. For each input value, the standard deviation or 95% confidence value is input as a percentage of the mean value. The uncertainties of the output values are calculated and displayed in the uncertainty slots. The uncertainties can be propagated across links. River reaches can propagate inflow uncertainty through the Muskingum routing method along with uncertainties of the Muskingum routing parameters.

Model Execution Options

RiverWare executes in single-run mode with the following solution options:

- Pure Simulation with option for Water Quality after Water Quantity at end of each timestep or after all timesteps.
- Rulebased Simulation with option for Water Quality at end of each timestep.
- Pure Simulation with Water Accounting at end of each timestep or after all timesteps.
- Rulebased Simulation with Water Accounting at end of each timestep.

- Optimization with optional automatic set-up of post-optimization simulation run.

All above runs can be executed through the GUI or in batch mode using command line arguments to load model files and run DMIs, etc.

Multiple Run Management

For some modeling purposes, it is necessary to make many runs and use the aggregated results from all the runs. RiverWare includes a utility called Multiple Run Management (MRM) that sets up and executes multiple runs automatically and sends the results to output files that can be analyzed by post-processing programs. There are three basic approaches to multiple run management as described in the following sections

Using the *Concurrent* mode of MRM, the user can make many runs over a planning horizon, using many traces of stochastically generated hydrologic inputs. MRM exports the results of the runs to one or more files in RiverWare Data Format (rdf). Then, post-processing analysis programs can import the rdf files and generate probabilistic information about the occurrence of certain events or the effectiveness of proposed operating policies. The hydrologic traces can be generated externally by a statistical package, or alternatively, a single set of hydrologic traces (usually the period-of-record hydrology) can be input directly into the model file. The MRM has an option that permutes these data according to the index sequential method, creating a new combination of the historical data for each of the multiple runs. (This method is commonly used by the BOR and other agencies.)

With the use of the Graphical Policy Analysis Tool (GPAT), an Excel-based tool developed jointly by CADSWES and Reclamation, the output of stochastic MRM can be used to compare two or more proposed operating policies in terms of their probabilistic effects on specified basin measurement criteria. In NEPA environmental impact studies, the measurement criteria may include a stream flow or lake elevation that is expected to comply with biological recommendations. To use GPAT for policy comparison, the multi-trace runs are performed for each policy alternative, and the results imported into GPAT. GPAT can provide statistical information in various ways over time, hydrologic trace, and policies.

Iterative MRM uses logic in the form of RPL expressions to control iterative runs of a single timeframe. The logic can look at values from within a run, and set values for the next run, using some specified criteria to end the iterations. This is used, for example, for yield studies.

Model Parameters:

In simulation, the basic parameters reflect the mass balance solution. Reservoirs are characterized by elevation-storage-area relationships, and maximum release and spill tables. Other processes such as hydropower, tailwater elevation, evaporation, river reach routing, consumption and return flow on Water Users, and water quality are represented according to specific methods selected by the user. The available user-selectable method categories and their menus of methods, along with the input, output and parameter requirements are described in the RiverWare technical documentation.

Rulebased Simulation does not require any other specific parameters except as needed by the policies. Optimization requires upper and lower bounds for optimization variables and linearization parameters.

Spatial Scale Employed in the Model:

RiverWare's spatial representation is based on discrete basin features such as reservoirs, diversions, water users, etc. Aggregation objects such as the Aggregate Reach, Aggregate Water User and Distribution Canal allow economic workspace representation. Spatial scale of River Reach objects should be consistent with the computational timestep and the routing method selected. For example, a "No Routing" reach with a monthly timestep could be hundreds of miles long, whereas a reach modeled with Muskingum-Cunge at an hourly timestep must have a length consistent with the routing parameters. Some routing methods such as Kinematic Wave and Storage Routing automatically discretize the reach into spatial elements that are consistent with the timestep for computational purposes.

The Sub-basin Editor allows the user to group objects in the model into Sub-basins that have specific names. The sub-basin designation does not affect the model solution. It is a convenience for referring to groups of objects in the rule and constraint languages.

Temporal Scale Employed in the Model:

The following computational timesteps are available: 1 hour, 6 hours, 12 hours, day, week, month, year.

RiverWare uses an absolute, i.e., historical time designation (rather than indexing) for time management. In specifying the run duration, specific start and end date/times are required and all input and output data are associated with specific date/times. Timeseries slots generate their timeseries automatically for time-varying data.

Input Data Requirements:

Required and Optional Input Values

The following is a partial list of required and optional inputs to a RiverWare model:

- Hydrologic Inflows into the rivers and reservoirs at particular points, or river gage data and reservoir storages that can be used to calculate the hydrologic inflows.
- Volume-elevation-area tables for level reservoirs, or headwater/storage/flow tables for sloped reservoirs.
- Data for optional reservoir methods include: hydropower characteristic curves, energy in storage tables, release and spill curves, evaporation and bank storage parameters, tailwater curves, pump/generator capacities, seepage coefficients and sediment accumulation parameters.
- River reach and canal data can (depending on methods selected) include routing parameters, loss and seepage coefficients, stage-discharge coefficients, diversion options, and bank storage parameters.
- Water Users required inputs (depending on methods selected) include: diversion requests or depletion requests or data such as land use or population for specific methods which can result in diversion requests; pumping capacities for supplemental groundwater usage; and data to support various return flow method calculations.
- Thermal power system data is required on the Thermal Object for certain hydropower economic calculations in simulation and optimization.

A simulation run requires exactly enough data to solve the mass balance and inflow/outflow equations of the objects. Many combinations of inputs are possible, allowing the objects to solve upstream or downstream, or forward or backwards in time. Each object has a number of possible solution algorithms (called "dispatch methods") based on the combinations of inputs for which a solution is possible. For example, given the previous storage, a reservoir can solve if the user inputs the pool elevation and the inflow is either input or propagated across the link from the upstream object's outflow. In this case, the reservoir solves for release and storage. Other input variable options are: outflow, storage, energy, a “max” flag on the outflow slot or a “best efficiency” or “max capacity” flag on the energy slot. If inflow, outflow and storage are all specified, the reservoir can solve for local inflow. Reaches can solve downstream or upstream (for some simple routing methods). Similarly all the objects can solve for various combinations of input data. Often the input data options are dependent on user-selected methods on the objects. A description of the input options for all objects is in the RiverWare technical documentation.

Rulebased simulation requires a set of policies to drive the simulation, and any special data needed by those policies. Input values are also accepted and are treated as the highest priority values, i.e., cannot be overwritten by any of the rules. Optimization requires as input a set of goals (objectives and constraints) and the associated data. Specific user input values are possible in the optimization solution also. These values are given the highest priority in the goal programming solution. For example, a particular reservoir pool elevation can be specified by the user to satisfy a recreational requirement. This elevation will hold while the system optimizes for other objectives such as hydropower economics.

Units

Each input and output value in RiverWare can be expressed in units selected by the user. The values are stored internally, and all computations are performed in SI units. The display units can be changed at any time. For units which do not have constant conversions, such as per-month or per-year rates, the conversion uses the actual number of seconds in the timestep, e.g., acre-ft per month for February 2000 has a different conversion to cms than it would for March 2000 or February 2001, accounting for the variation in seconds in these months. The Water Accounting data can be displayed either as flow or volume units; the conversion is automatic and uses the model’s computational timestep.

Input Mechanisms

Input values can be loaded into RiverWare by any of three mechanisms:

- Manually through the GUI. Values for each timestep (or table data) can be typed directly into the variable’s data display window. The System Control Table, a spreadsheet-like formatting of the data, can also accept direct input by the user.
- Loading data files. For each data variable in the model, the GUI provides an option to load the data structure with data from a flat file on the file on the user’s file system. The user specified the units of the values.
- Data Management Interface (DMI). The DMI provides a way to automate retrieving and loading large amounts of data into the model by making a single selection from RiverWare’s GUI that invokes an external program. The external program retrieves data from databases, files, or any other source, and formats the data in flat files, which can be imported into RiverWare’s data structures. A control file maps the file names to specific data structures in the model. DMIs are used to initialize models with current values, input forecasts and demand schedules, or to load historical data.
- Direct data connection with DSS or HDB. Through RiverWare's GUI, slots can be mapped directly to records in a specified DSS file (the Corp's Data Storage System) or in the Oracle Hydrologic Database used by Reclamation.

Model Output:

Output Values

The output values are the timeseries of values that are solved for on the various objects; the list depends on the specific model configuration and user methods selected (see RiverWare Technical Documentation for exact output for each object and method). In general, outputs are reservoir and reach outflows and/or reservoir storages, elevations, and energy. Outputs can also include water quality and water accounting information.

Options for Output Value Representation

Output values can be viewed and/or exported in the following ways:

- Each individual variable's data structure can be opened to view the input and output values at each timestep.
- The System Control Table (SCT), a spreadsheet-formatted interface for the model, displays input and output values for selected variables in the model.
- Plots of variables can be configured to show a single variable or several variables plotted over time.
- "Snapshots" are data from the same variables saved over several runs. These can be viewed or plotted to show how results vary with different input sets.
- Data Management Interface (DMI) routines can be invoked to export the data to databases or other destinations (see explanation in "Input Mechanisms" above.
- Selected data can be exported in RiverWare output format. Post-process tools are available to format this output to an Excel binary file or a standard comma-delimited spreadsheet file format.

Additional Visual Output for Run Analysis

A graphical analysis tool can be displayed for simulation runs, rulebased simulation runs and optimization runs. The simulation analysis tool displays a matrix of objects and timesteps and shows which objects solved and which did not solve during each timestep. Additional details indicate what values were known and unknown. This information allows the user to identify over- or under-determination problems. The rulebased simulation analysis tool includes the same information, and in addition, it shows which operational policies (indicated by priority number) affected the solution of each object at each timestep. The visual analysis tool thus shows how policies are applied at various times and how the effects of the policies propagate through the river/reservoir system. The optimization analysis tool similarly indicates which goal (by priority) was responsible for the final values in each object at each timestep.

Diagnostics and Error Messages

A user-configurable diagnostic facility outputs various categories of diagnostics during model runs to assist in model analysis and debugging. All error messages are also sent to the diagnostic output window. The diagnostics are color-coded by category for ease of reading, can be searched for key words, and can be sent to a file and saved.

Parameter Estimation / Model Calibration:

RiverWare does not provide automatic parameter estimation routines.

Model Testing and Verification:

Each model that is constructed using the RiverWare modeling tool must be calibrated tested and verified by the model developer.

Model Sensitivity:

RiverWare's computations set values to within a specified "convergence" tolerance. An improved value will not be set in a slot if it is within the slot's convergence tolerance of

the old value. A default convergence criterion of .01 percent is set on all slots. The user must verify that the convergence criterion is small enough to result in accurate computations, and change the convergence criterion if needed. For example, a relatively small change of elevation in a reservoir may result in a significant change in storage. However, if the convergence criterion for elevation is too small, the change will not be made to the elevation hence will not be reflected in the storage.
The modeler must select the physical processes to model that are significant to the results and meaningful within the computational timestep selected.

Model Reliability:

RiverWare models have been developed to replace previous models on a number of river/reservoir systems including the Colorado River and the TVA system. These models have been shown to produce identical results to the previous models. RiverWare models have also been developed and calibrated to historical data on other basins.

Model Application / Case Studies:

RiverWare has been used to develop planning and operational models of the TVA reservoir and river system, the Colorado River, the San Juan River, the Pecos River, the Yakima and Umatilla basins, the Upper Rio Grande, the Truckee-Carson River, and several other reservoir/river systems in the U.S. and abroad.

Computer Requirements:

RiverWare currently runs on Windows 2000, XP or Vista and Solaris Unix 8.0 or higher. Memory and disk requirements depend on model size, which has no limits. Recommended minimum of 256 MB of memory, using the maximum memory for large models.

Documentation and Training:

General and technical documentation is available on the CADSWES web site http://cadswes.colorado.edu. The technical documentation requires Acrobat Reader. The following training courses are offered at CADSWES: Introduction to RiverWare Simulation Modeling, Introduction to Rulebased Simulation, and Introduction to RiverWare Optimization. Water Accounting training class will be offered for the first time in 2007. Each training course provides a detailed training manual with guided tutorials and reference materials.

Other Comments:

RiverWare continues to be enhanced at CADSWES under the sponsorship of the TVA, BOR, USACE and other agencies. RiverWare’s software is designed to be easily extensible, permitting the addition and modification of methods to meet evolving needs as they are identified.

CADSWES provides user support with each license. License revenues contribute to maintenance and small enhancements of the software.