

Modeling the Roanoke River Basin Operations with OASIS

Addendum to the
User Manual for OASIS with OCL™

December 2013

Prepared for the
North Carolina Division of Water Resources

DRAFT



Table of Contents

<u>Section</u>	Page
1 – Introduction	3
2 – Model Components	
2.1 Schematic	5
2.2 Model Input	7
2.3 Run Configurations	15
2.4 Model Output	16

Appendices

- A - Model Static Input Data and Run Code
- B – Finalized Inflow Data Development
- C – Provisional Inflow Data Development
- D – Model Weighting Description

Section 1. Introduction

This report describes how OASIS is used to model the operations of the Roanoke River Basin in Virginia and North Carolina. This application of OASIS, known as the Roanoke River Basin Reservoir Operations Model (RRBROM), extends geographically from the headwaters of the Staunton and Dan Rivers in Virginia to the tidal portion of the Roanoke River near Hamilton. This report is not intended to replace the User Manual for OASIS, which is available from the Help menu of the model. Rather, it is intended to document the data used in this application as well as the current operations of the basin. Information about the OASIS platform is included only to the extent necessary to provide context for the application-specific data.

The model is available for registered users on the Division of Water Resources (DWR) server. The model can be used in two modes: (1) a simulation mode to evaluate system performance for a given set of demands, operating policies, and facilities over the historic inflow record; and (2) a position analysis mode for real-time management. The simulation mode contains two default runs, one for conditions today and one for projected 2060 conditions. In the position analysis mode, the model uses multiple ensemble forecasts to provide a probabilistic assessment of conditions up to one year in the future. Although it can be used for other purposes as well, this feature is particularly useful for drought management.

The model uses an inflow data set that extends from January 1, 1930 through December 30, 2011. This data set was developed using a comprehensive approach that (1) relies on over 30 streamflow gages in the basin; (2) accounts explicitly for upstream alterations, or impairments, from reservoir regulation and net water consumption; and (3) uses statistical techniques to complete missing records for these gages.

Real-time drought management requires forecasts of inflow and, as noted below, the forecasts are generated based on inflows through the present day. Updating the inflows requires the collection of impairment data, which can be time intensive. It is envisioned that these data will be collected every five years. In the interim (e.g., through 2016), the inflow data starting January 2012 will be based on a provisional inflow technique so that real-time updates can be made quickly and easily without the need to collect all the impairment data.

The remainder of this document summarizes the components of the model and the major operations in the basin. Appendix A lists the static input data and run code used in the basecase simulation run that is based on today's facilities, operations, and demands. Appendix B describes the approach used to establish the finalized inflow data set. Appendix C describes the approach for generating provisional inflows. Appendix D describes the weighting assigned to the various nodes and arcs so that the model reflects the general priorities for water allocation in the basin.

It is important to note how the OASIS model should and should not be used. OASIS is a generalized type of mass balance model used mainly in evaluating planning and

management alternatives. It is not intended for use in hydraulic routing nor flood management, although it can be linked to other models for those purposes.

In addition, since modeling results are sensitive to inflows, the user must be cautioned about accuracy of the inflows. HydroLogics spent considerable effort in developing the inflow data. The methodology ensures that the monthly naturalized flows at the gage locations match, which assumes that any measurement error is embedded in the impairments and not the streamflow data. DWR agreed to this method, which, although imperfect, is the most reasonable given the available data. Further, it is important to note that we are not trying to replicate history in computing the OASIS inflows; rather, we are trying to build a data set of daily flows whose variation is *representative* of history while preserving monthly gaged flows as “ground truth”.

Due partly to the inaccuracy of some of the impairment data and to time of travel, negative inflows may occur. These can lead to potential model infeasibility. The model code filters out negative inflows, particularly large ones, but preserves the total inflow volume over a short period by debiting those negative inflows from subsequent positive inflows. For example, if a rainstorm hits the upstream part of the reach but not the downstream part, the gaged flow data may indicate a large negative inflow (gain) between the upstream and downstream ends. When the flow attenuates upstream and peaks downstream, the inflow becomes positive, and the negative gains from the days before are debited from the positive inflows in the days after to ensure that the average inflow over that period is preserved. The occurrence of negative inflows is reduced in the main-stem of the Dan and Roanoke Rivers by incorporating time-of-travel equations recommended by DWR. These equations are provided in Appendix B.

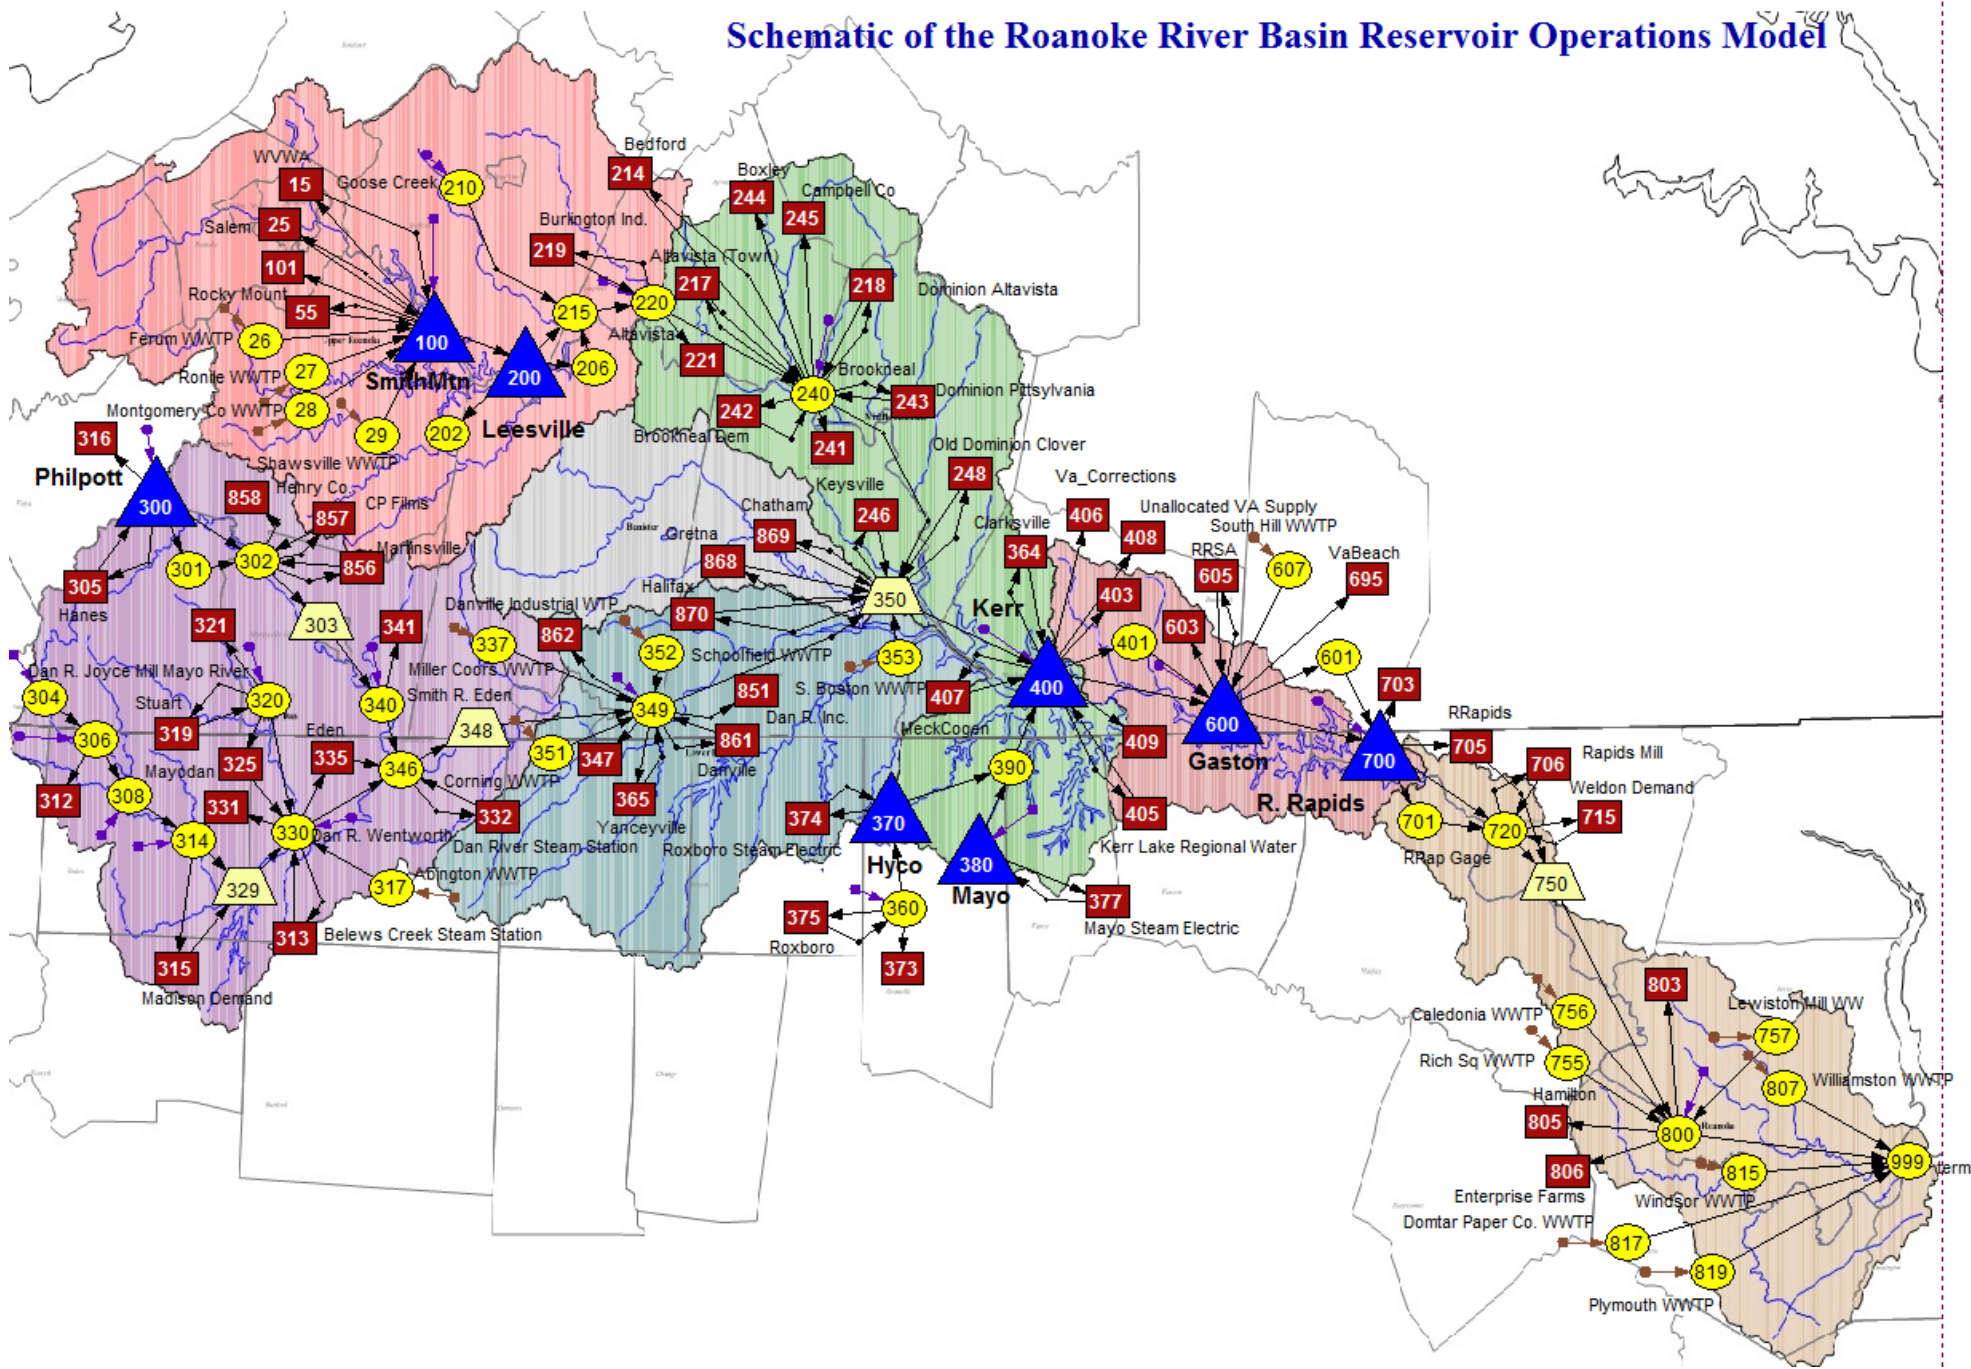
Section 2. Model Components

2.1 Schematic

The model uses a map-based schematic that includes nodes for withdrawals (agricultural, municipal, and industrial), discharges (municipal and industrial), reservoirs, gage locations, and points along the rivers where flows are of interest. Arcs represent means of water conveyance between nodes. The model schematic is shown on the following page. The schematic and associated physical data were developed with input from basin stakeholders at model review meetings.

In total, the model has approximately 115 nodes and 150 connecting arcs. There are 13 reservoir nodes (8 actual reservoirs and 5 artificial storage nodes used for time-of-travel flow routing), over 10 agricultural demand nodes, over 45 municipal and industrial demand nodes, over 15 independent wastewater discharge nodes (i.e., not tied into a water withdrawal node), over 15 natural inflow nodes (including the reservoir nodes), and other miscellaneous nodes to account for minimum flow requirements, interconnections, and instream flow assessment for ecological needs.

Schematic of the Roanoke River Basin Reservoir Operations Model



The user can click on any node or connecting arc on the schematic to access specific information, like reservoir elevation-storage-area data or minimum streamflow requirements. These data are also contained in tables contained on other tabs of the model.

2.2 Model Input

Input data for the model is stored in three forms: static and pattern data, timeseries data, and user-defined data using operations control language (OCL). The timeseries data are stored outside the model run. The other data are embedded in the run and copy over automatically when creating a new run.

Static and pattern data are contained in the GUI and represent data that do not change during the model simulation (such as physical data like reservoir elevation-storage-area relationships) or repeating data that occurs every year in the simulation (like monthly demand patterns or seasonal minimum release patterns). Timeseries data change with each day in the simulation record and typically consist of inflows and reservoir net evaporation. OCL allows the user to define more elaborate operating rules than are permitted from the GUI.

Static and Pattern Data

Tables containing the model's static and pattern data can be found in Appendix A. Reservoir information includes elevation-storage-area relationships, minimum and maximum allowable storage, and any rule curves which dictate the preferred operating elevation throughout the year.

Minimum flows and reservoir releases are defined by minimum flow patterns on arcs.

Water treatment plant and transmission constraints are defined by maximum capacities on arcs.

Municipal and industrial demand nodes use an annual average demand subject to a monthly pattern, and an associated wastewater discharge based on a fraction of the monthly demand. Wastewater discharges not associated with demand nodes are modeled using an annual average return subject to a monthly pattern.

The model allows the user to systematically adjust all municipal and industrial demands in the basin by invoking the demand multiplier option on the Setup tab. This is useful when doing sensitivity analyses on the impact of demand growth in the basin. Note that agricultural demands and independent wastewater returns are not adjusted using this multiplier. The agricultural demand can be adjusted as described below, and the independent wastewater returns can be adjusted manually in the pattern tables.

Timeseries Data

The timeseries data are stored in a basedata timeseries file (*basedata.dss*), which contains all the inflow and net evaporation (evaporation less precipitation) data. The sources for these data are provided in Appendix B along with a more detailed description of how the inflows were developed. As noted, updating the timeseries data can be done in two ways: (1) using the comprehensive approach described in Appendix B; or (2) using the provisional approach for facilitating real-time drought management described in Appendix C. The provisional approach relies on data from select gaging and precipitation stations throughout the basin.

The provisional updates can be done directly from the interface by selecting the Update Record tab. First the user presses the Download Data Button; once the data has downloaded the user needs to check for any blanks or erroneous values. After verifying the data, the inflows can be updated by pressing the Update Record button. The update record algorithm will calculate the inflows to all the OASIS inflow nodes and net evaporation for all reservoir nodes and write them to the *basedata.dss* file automatically.

Agricultural water use is modeled as a timeseries over the historic hydrologic record. It is broken out by county and depends on livestock count, crop usage, livestock and crop water consumption, and rainfall. Evapotranspiration equations for each crop are used in conjunction with the timeseries precipitation record so that crops are only irrigated when necessary. The water use can be easily adjusted from the model interface by opening the Edit Agricultural Data dialog box on the Setup tab. The model automatically converts the input data on crop acreage and livestock count into water use values. The agricultural demand nodes are a summation of the agricultural water usage in a particular reach of interest.

Operating Rules

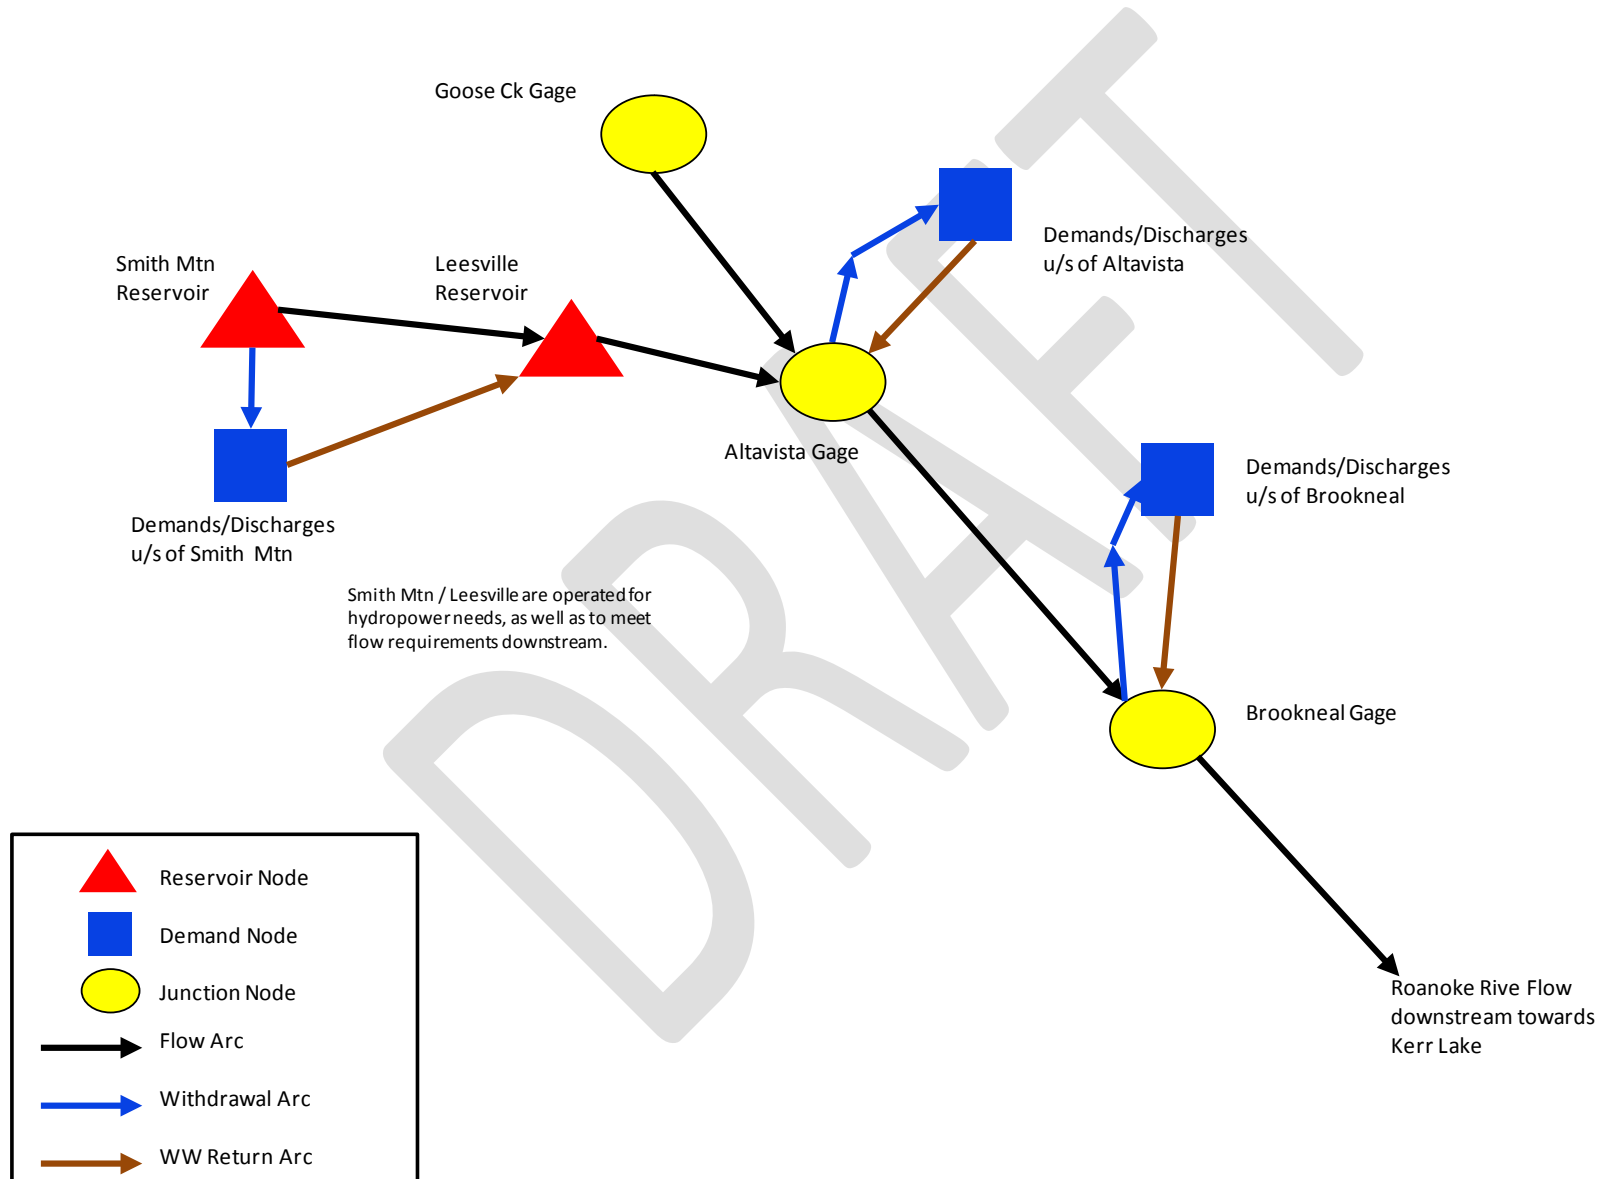
As described in more detail in Appendix D, most of the water allocation priorities are set by the user in the GUI by applying weights to nodes and arcs. The most common operating rules are for storing water in reservoirs versus releasing the water to meet local demands or minimum releases, and these are reflected by the weighting scheme. Simply stated, if a minimum flow in a river is more important than meeting the local water supply demand, a higher weight on the minimum flow means water supply deliveries will be scaled back if necessary in a drought to meet the minimum flow.

The Operations Control Language (OCL) allows the user to model more complex operating rules, such as hydropower production or drought management protocols that tie demand reductions to reservoir levels or river flows. These files are accessible from the model interface. The OCL files associated with the basecase simulation run that represents current conditions is included in Appendix A. The key OCL files include *main.ocl*, which initializes the run and refers to all the other OCL files (as well as filtering the inflows for any negative gains in the provisional record); *WW_returns.ocl*

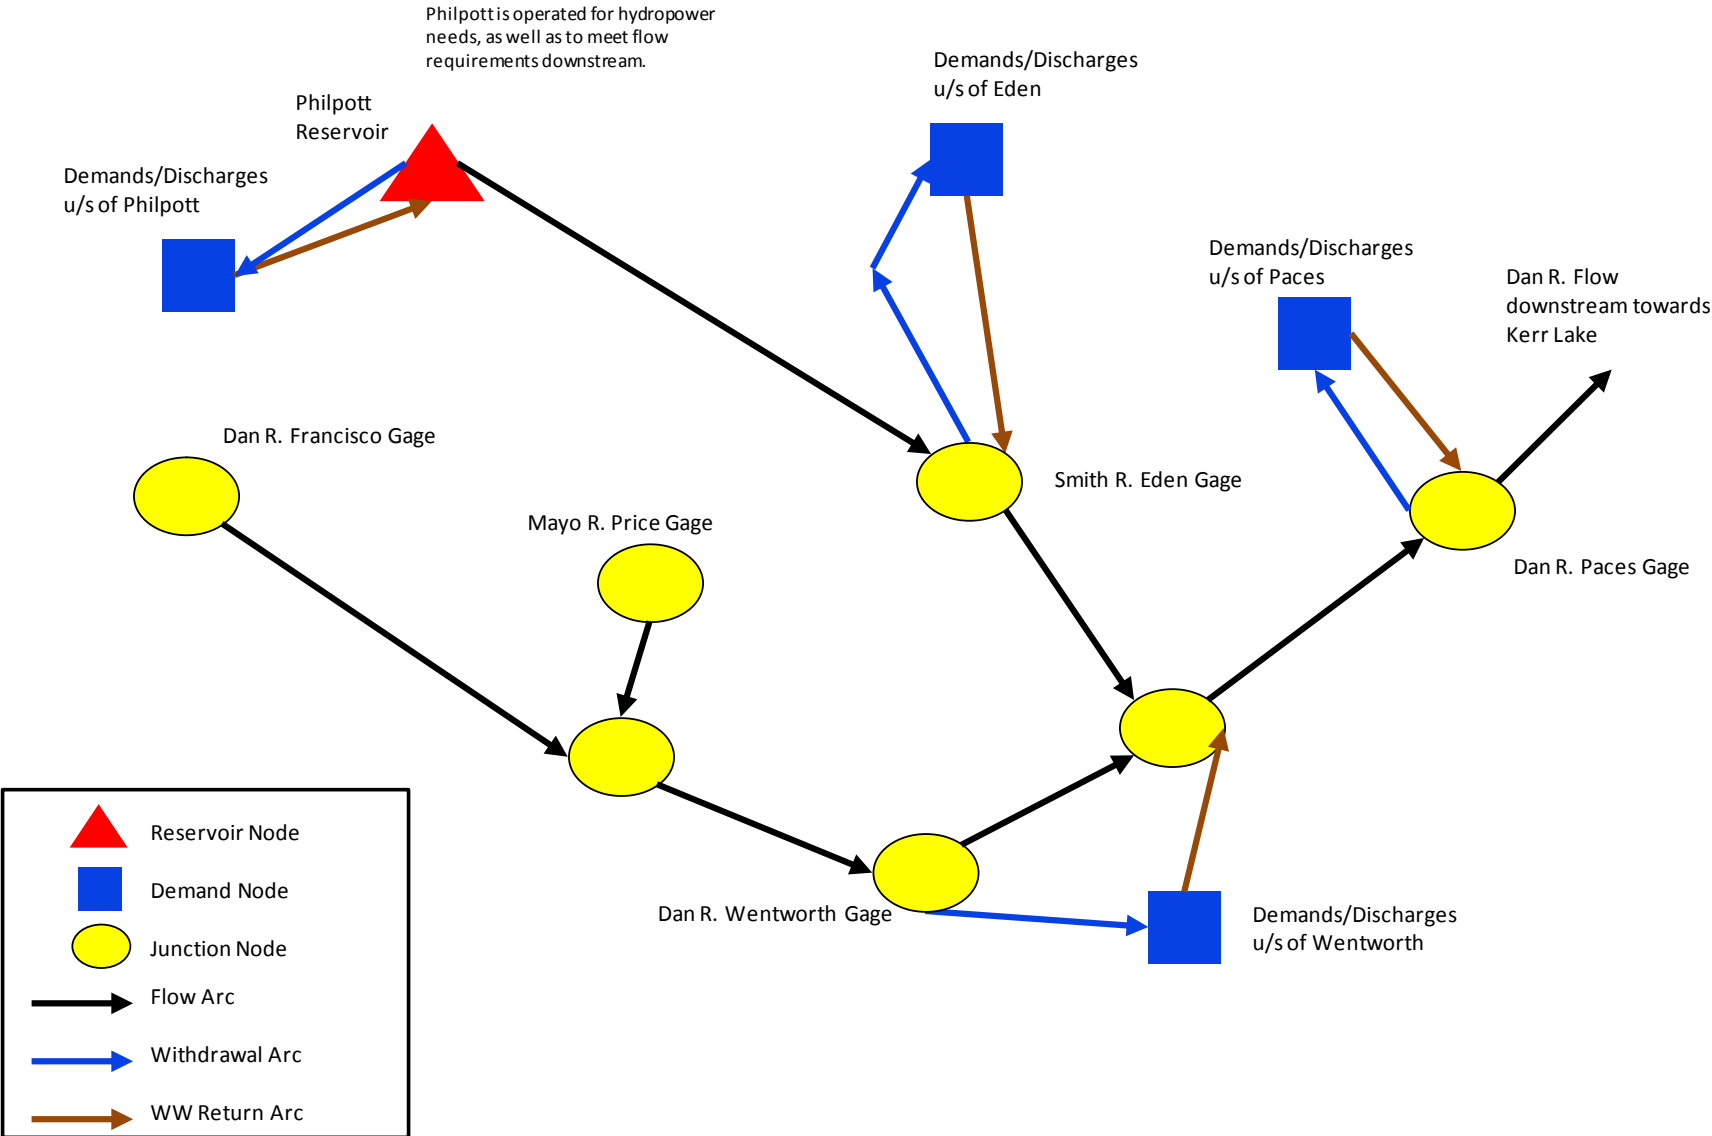
which sets the wastewater returns; *routing.ocl*, which routes the flows through the basin; a number of files which dictate the operating policies for the Smith Mountain / Leesville system, and for Kerr Lake down to Roanoke Rapids Lake; and *drought_plans.ocl* which code the North Carolina Water Shortage Response Plans submitted by utilities to DWR. A number of other OCL files dictate the operating policies for other systems, and can be found in Appendix A. Appendix D details the weighting which also controls operations in the basin.

A series of stylized flowcharts are provided below summarizing the overall operations of the basin as captured in the model.

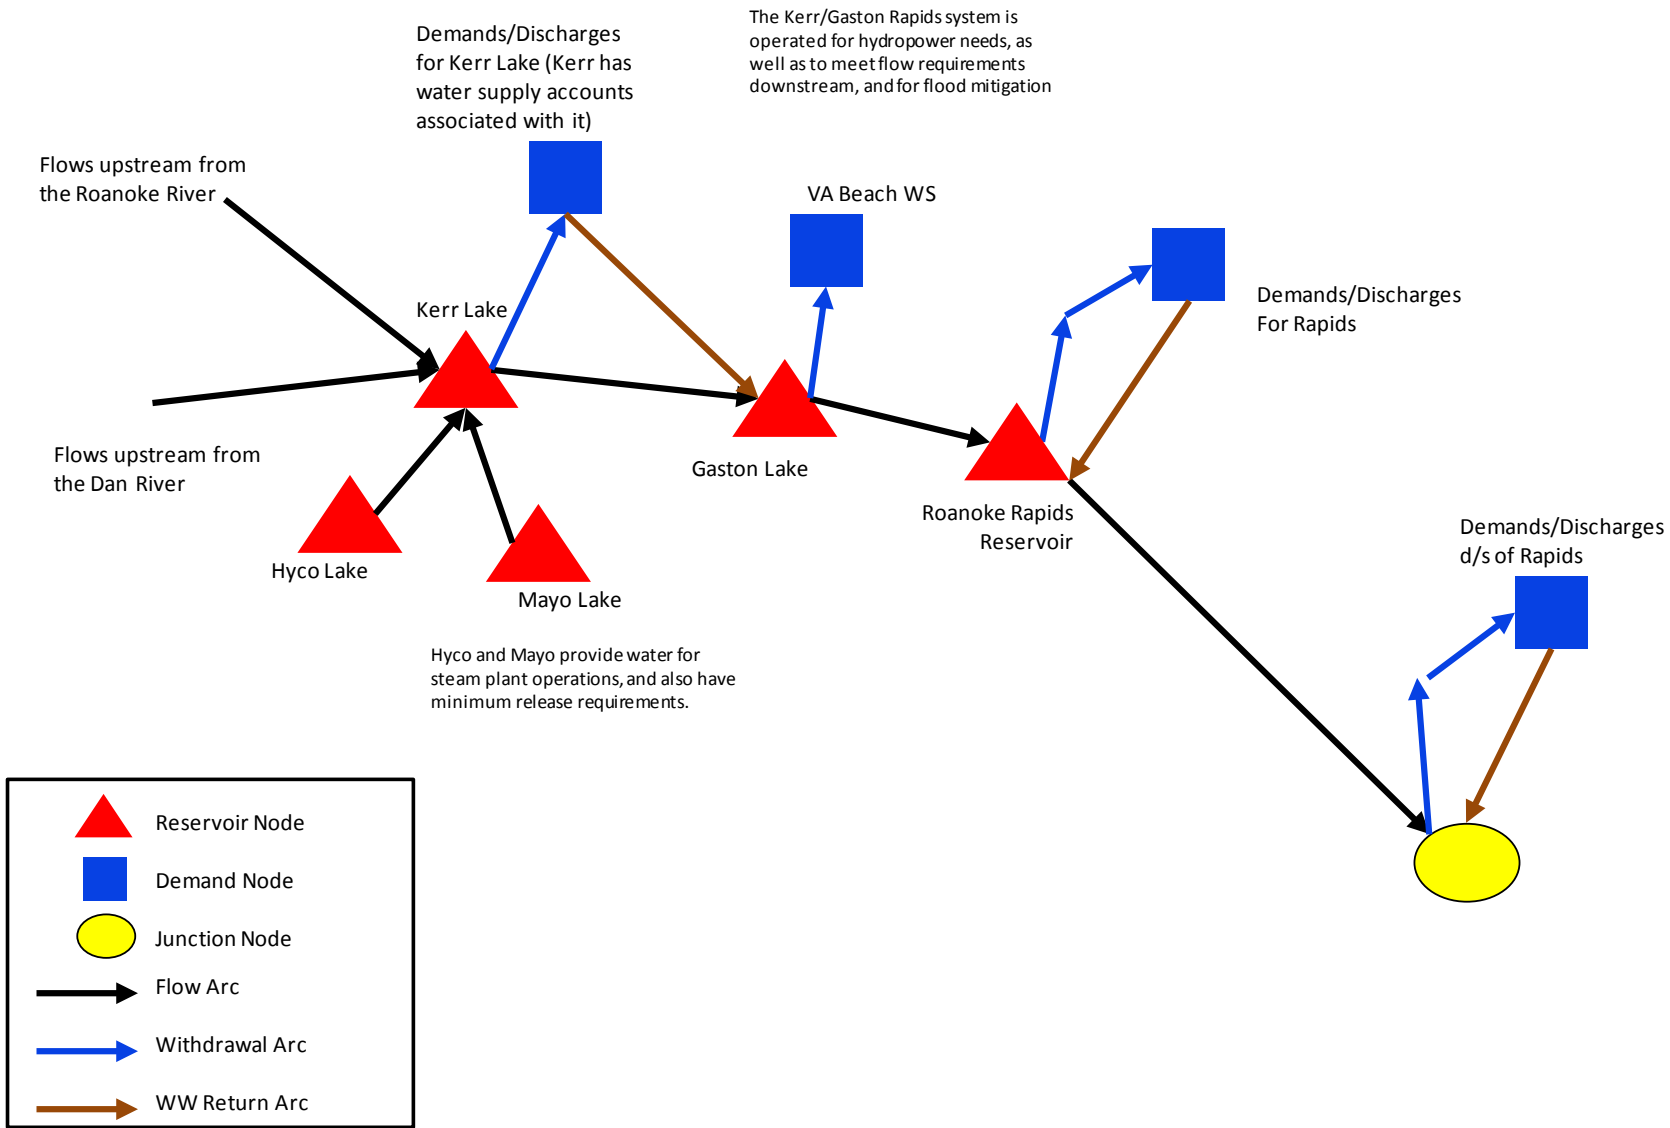
Flow Chart of Major Nodes in the Upper Roanoke Basin



Flow Chart of Major Nodes in the Dan/Smith River Basin



Flow Chart of Major Nodes in the Lower Roanoke River Basin



Because of the complexity of operations in this basin, which are driven by hydropower production, additional operational details are provided below.

Smith Mountain/Leesville (captured in *leesville* and *user_def_ops* ocl files in Appendix A)

This pump-storage facility, owned by Appalachian Power Company, a division of American Electric Power, has minimum release requirements at Leesville, and, during low flows, additional releases are made to meet flow targets at the Brookneal gage downstream. The low flow protocol that it follows uses probability-based triggers that are coded in the model. These triggers are invoked when there is a certain likelihood of reaching certain lake elevations in the near future and these stipulate what the resulting release targets should be.

Philpott (see *interchange.ocl* file in Appendix A)

Philpott coordinates power operations with Kerr Reservoir downstream, both Corps of Engineers projects. Like Kerr, it operates to a guide curve and produces a minimum firm energy. As part of the operations interchange, when lake levels are low, some of the power generation is shifted to Kerr to help conserve storage in Philpott (and vice versa if Kerr is low).

Kerr-Gaston-Roanoke Rapids (Rapids) (see *kerr_declaration*, *coe_rules*, *set_firm_energy*, *rapids_releases*, *spawning*, and *betterment* ocl files as well as the energy run OCL files described below, all provided in Appendix A)

The complex power operations of these three reservoirs are driven by a weekly energy declaration that establishes the target volume of water that will be released from Kerr from Saturday, start of week, to Friday, end of week. The weekly release must meet the firm energy requirements. If there is enough storage in Kerr (relative to its guide curve) to generate at least 1200 MWh of secondary energy during the week, the minimum release is increased to that amount. In addition, the declaration is adjusted depending on the relative storage levels in Philpott and Kerr based on their interchange rules. The weekly declaration is the sum of the firm, secondary, and interchange energy levels. The Corps has the option of adjusting the declaration as the week progresses if inflow projections change. The model will adjust the declaration (called a redeclaration) through mid-week (Wednesday) if the new declaration is at least 10 percent higher than the original declaration set on Saturday.

The releases from Kerr are coordinated with operations of the smaller lakes downstream (Lake Gaston and Rapids). Dominion Virginia Power, which owns these two reservoirs, dispatches the weekly declaration for Kerr from the three projects based on hourly power prices. The model does the same, although since this is a simulation model, a representative weekly pricing pattern for each month is used when simulating the historic inflow record from 1930 to present. The pricing pattern is based on average prices from 2009 to 2012. There is a different pattern applied to each month since prices are sensitive to time of year.

RRBROM is a daily model. It essentially has two components: a water run which determines that water use over the course of the week (daily calculations), and an energy run which slots the water into the highest priced hours during the week to maximize hydropower revenue (hourly calculations), which reflects how the system is dispatched. These are done in one pass when running the model. OCL files for the energy run (included in Appendix A) are prefixed with the word “*power*”. There are other objectives that may override, or even constraints (which must be met), and these are described below. The user has the option of running the hourly energy component in two modes: one where the model allocates the daily volume of water using an average daily power price for each day (called hourly/daily using a value for the constant DoPowerOptim in the Constants table of 2), the other using hourly power prices for each day (called hourly using a constant value of 1). The difference in energy runs is significant in run time. Using the hourly/daily optimization, the model runs in about 10 minutes for the period of record (from 1930 to present) compared to about two hours more for the hourly optimization. The differences in results, including dispatch volume and revenue, are insignificant overall, so the user is encouraged to use the hourly/daily mode.

Flood control: normal operations limit the daily release from Rapids to 8,000 cfs. When Kerr reaches 300 feet or the upper end of the spawning bubble (maximum of 302 feet), the flood “ratchet” is increased to 20,000 cfs, which is approximately equal to the maximum turbine capacity at Rapids. At 312 feet, the ratchet increases to 25,000 cfs, meaning spill at Rapids occurs. At 315 feet, the ratchet is 35,000 cfs; at 320 feet, the ratchet is 35,000 cfs or 85% of the inflow to Kerr, whichever is higher; and at 321, the ratchet is 100% of the inflow to Kerr.

Spawning: fixed releases from Rapids occur during the spawning season from April 1 to June 15. A spawning bubble raises the normal operating guide curve at Kerr during this period (up to 302 feet) so that additional storage can be made available to meet the spawning releases. No energy peaking from Rapids is allowed. Spawning releases are based primarily on Kerr elevations, meaning spawning releases will get reduced the lower the lake drops. If Kerr elevation exceeds the spawning elevation bubble because of net inflows, the model will release additional water up to the flood ratchet to get back to the guide curve.

Betterment: the Corps’ betterment plan is invoked on the receding limb after flood releases drop to a certain level. In this case, the Rapids releases drop by 5000 cfs increments every 4 days starting at 20,000 cfs once the lake has dropped back down to 312 feet.

Drought Protocol: because it is not easily codified, no drought protocol is modeled for power operations for Philpott-Kerr-Gaston-Rapids. In extreme droughts (like 2002), flows from the system have been reduced below the normal minimums (through variances) and purchases of firm energy from the open market to allow lower generation (and thus releases) from the system. Consequently, the computed elevations during certain drought events (like 2002) may be lower than those historically.

Gaston

The lake elevation is tightly controlled and is normally operated within a one foot band. The lake can surcharge a few feet during high inflow events.

Rapids

Releases from Rapids are controlled by a FERC minimum release, energy requirements, spawning requirements, the flood ratchets, and maintaining a 5 foot elevation range. FERC minimum releases are increased at certain times of the year based on the weekly declaration.

2.3 Run Configurations

The model can be used in two modes: (1) a simulation mode to evaluate system performance for a given set of demands, operating policies, and facilities over the historic inflow record; and (2) a position analysis mode for real-time management. General information on creating, modifying, and executing runs is provided in the User Manual for OASIS, which is available from the Help menu of the model.

Simulation:

In simulation mode, on the Setup tab, the user can select from three radio buttons: No Forecasts, Conditional Forecasts, and Non-conditional Forecasts. The latter two enable the user to evaluate forecast-based operating policies (which is used for the Smith Mountain/Leesville low flow protocol in the basecase scenario), with inflow forecasts generated for each week in the historical inflow record. Conditional forecasts account for antecedent flow conditions while non-conditional forecasts are made independent of how wet or dry the basin is. The forecasts for the simulation mode are generated outside the GUI and stored in the basedata folder. The current forecast file is developed from the timeseries *basedata.dss* file that extends through December 2011. The forecast file should only be updated in conjunction with the comprehensive inflow updates (anticipated every five years with the next update in 2017).

To enable all utility drought plans in a run, set the *Drought Plans On* variable in the OCL Constants Table to 1. A value of 0 will turn all drought plans off.

The GUI allows for all municipal and industrial demands in the model to be uniformly increased or decreased by a user-specified fraction. To enable the demand multiplier, check the *Use Multiplier* box on the Setup tab, and enter a number in the *Multiplier Value* box. For example, setting the value to 0.9 will decrease all M&I demands by 10%, and setting it to 1.1 will increase them by 10%.

Position Analysis:

In position analysis mode, the user can select from Conditional or Non-Conditional Forecasts on the Setup tab. By executing a run, the model will produce a forecast (typically of river flows or reservoir elevations) for up to the next 365 days. A forecast can be made on any date in the historic inflow record or no more than one day after the end of the inflow record. Typically it will be used starting the day after the last update of the inflow and net evaporation record. For example, if these records end December 31, 2011, the user can run a forecast for January 1, 2012. If a month has passed, and the user wants to run a forecast for February 1, 2012, the user would update the inflows and net evaporation for January using the Update Record tab and then start the position analysis run on February 1. For a reservoir, or locations affected by the operation of a reservoir upstream, the forecast is dependent on the starting elevation of the reservoir. On the Setup tab, the user simply inputs the starting elevation (or storage), the starting date of the forecast run, and clicks Run.

2.4 Model Output

The model allows the user to customize output files (tables or plots) and save them for routine use. Alternatively, the user can click on any node or arc on the schematic or go to the Setup tab and select Quick View to access and save tabular or plotted output. A number of tables and plots have been provided for points of interest in the basin. The balance sheet can also serve as a useful tool for tracking water through the system.

Included in the model output tables is a file called *xQy_ClimaticYear_Rapids.lv*. This file allows the user to compute instream flow statistics, such as 7Q10, for a specific site, in this case the Roanoke River at Roanoke Rapids gage. To generate statistics for a different site, the user would copy and rename the file, then change the name and associated arc listed in the file. When viewing the generated output, the default layout shows two columns, for 7- and 30-day low flows (these periods can be changed in the .lv file). Scrolling to the bottom of the output file shows Log Pearson percentiles for each column. If the user is interested in the 7Q10 (7-day low flow, 10th percentile) flow, the user would look at the first column, and the row labeled LPrs.100.

In addition, the model is capable of automatically determining the safe yield for a specific demand node, in this case the demand from Philpott Lake. To generate statistics for a different site, the user would copy and rename the file (currently called *SafeYield_Philpott.lv*), then change the name and associated demand node listed in the file. The safe yield can be determined for each year in the historic inflow record (annual safe yield analysis) or for the entire period of record. The user inputs the adjustment criteria by selecting the Run Safe Yield Analysis button on the Setup tab. The safe yield routine works by tracking demand shortages for the chosen demand node, and iteratively works towards the maximum demand that produces no shortages from the supply source (in this example, Philpott Lake). Note under the current output file configuration, the drought plans should be turned off when using the safe yield routine, as the demand reductions resulting from drought restrictions inherently produce a 'shortage' from the

normal demand. The output file configuration can be modified if needed for the specific drought plan of each system.