Editorial

Innovative Systems-Based Decision Support: Tales for the Real World

Tirusew Asefa, Ph.D., P.E., D.WRE, M.ASCE

Dept. of Modeling and Systems Decision Support, Tampa Bay Water, Clearwater, FL 33763. E-mail: tasefa@tampabaywater.org

DOI: 10.1061/(ASCE)WR.1943-5452.0000565

In this paper, the author shares decades of collective experience in the development, implementation, and application of systems-based decision support tools at Tampa Bay Water—one of the largest wholesale water providers in the Southeast United States. The emphasis here is on the use of the systems-based modeling rather than on searching for "optimal" decisions concerning individual system components that may or may not result in overall system optimality. To this end, the paper highlights some state-of-the-practice models, approaches, and tools in the hope that researchers and academic professionals will appreciate the complex decision-making environment that practicing engineers face. For reasons that will become clear in in the text that follows, the use of systems-based decision support tools, although attractive in many ways, poses important challenges. From a utilities perspective, these tools require a commitment to in-house expertise and computing resources. In terms of the research community, the proposed innovations must avoid academic simplifications of the challenges real systems face and must actually make sense in complex decision-making environments.

Utility

Tampa Bay Water is a wholesale water provider for an area of 2.3 million customers in west central Florida. It provides high-quality water to six member governments consisting of three counties and three cities whose elected officials constitute the board that oversees the agency's function. Each member government is billed a uniform rate per thousand gallons to cover the cost of water delivery now and in the future (the agency has unequivocal agreement to meet the region's demand, and a budget is set every year based on projected member government demand). Tampa Bay Water delivers wholesale water at 19 connection points to meet specific levels of service (water quality, pressure, and flow).

Since 1998, when 100% of the region's water supply came from groundwater, there has been a significant shift to alternative sources because of the environmental impacts resulting from heavy groundwater pumpage. Today the mix of supply sources varies year to year, but approximately 7% comes from desalinated seawater; 40%, from surface water; and 53%, from groundwater. As a result, the region has witnessed a remarkable recovery in wetlands and lake water levels. Daily production ranges 120–260 million gallons. Table 1 shows facility changes since 1998. The overall value of the current infrastructure is estimated at $1.2 billion or more.

The following sections summarize the systems-based decision support tools that are used to inform operation of this infrastructure across various timescales.

Table 1. Tampa Bay Water Facility Overview

Parameter	1998	2013
Supply sources	Groundwater	Groundwater, surface water, desalinated seawater
Number of delivery points	9	19
Production wells	207	171
Water plants	2	12
Pumping/booster stations	4	14
Transmission lines (miles)	160	240
Horsepower	28,000	106,025
Pumps	228	259
Chemical feed systems	14	37
Ground storage tanks	5	10
Reservoir storage (billion gallon)	0	15.5
Monitoring/control nodes	92	428
SCADA input/output (I/O) points	1,000	22,500
Communication lines	20	65
Major permits	6	21

Weekly Decision Support Tools

The optimized regional operation plan (OROP) is an integer optimization model, based on a mathematical programming language (*AMPL*) (Fourer et al. 2003), that runs every week with a four-week optimization horizon. OROP produces a weekly system operation schedule based on the upcoming week's demand, supply availability, operational constraints, and recent hydrological conditions.

OROP submodels include delivery point demand forecast, operational models translating steamflow amounts to permitted withdrawals at specific locations, integrated hydrological output to provide a unit response relating groundwater-level changes to unit pumpage, and initial groundwater-level projection at control points.

Use

Given seasonal source allocation decisions (see the section "Seasonal Decision Support Tools"), the weekly decision support tools meet upcoming forecasted demand with a mix of supply sources (direct surface water from two rivers, reservoir, groundwater, and/or desalinated seawater). Although operators have flexibility in responding to member governments' daily demands at a delivery point, a weekly production schedule is generated by OROP every Friday that constitutes the schedule for the upcoming week starting on Saturday. Staff from the Operations, Regulatory Compliance, and Decision Support departments meet every Thursday to set the constraints and expectations for the OROP run. At the end of the run, an electronic report is published detailing how much water to use from which supply source to meet projected demand. The report also lists the priorities of standby groundwater wells in case demand exceeds forecast.

OROP's objective, given alternate sources of supply, is to distribute groundwater production among wellfields and the wells (175) in them based on recent hydrologic conditions. The model provides production schedules by maximizing surficial water levels at key locations through weights tied to a given locality's wetlands recovery target. If, for example, there has been higher rainfall and more favorable conditions in one part of the system than in others, OROP shifts wellfield production accordingly. Not only does this require a well-connected system that satisfies demand at a location in more than one way; the acquisition and use of recent data is also critical. The agency maintains a network of monitoring locations that are available in near real time. Environmental conditions are continuously monitored and assessed to determine the impacts of groundwater pumpage. If there is an impact in one area, there are protocols in place for detailed site assessment and eventually to make that location a control point over which the optimization model operates.

Once the weekly operational outlook is published and implemented, an electronic report is automatically generated detailing the weekly demand and supply forecast versus observations (usually within a week or two after the forecast). This report provides a snapshot of model performance. It includes a comparison of forecast and actual values for (1) each demand delivery point—that is, what the decision support tool expected a member government to ask for versus what it actually asked for; (2) forecast versus actual supply available; (3) scheduled versus actual wellfield production; (4) scheduled versus actual surface-water production; and (5) scheduled versus actual surface-water source allocation. Mismatches between forecast and observed data are used to assess such factors as operational constraints not yet captured and/or submodel performance. Short-term demand and supply forecasts are highly dependent on near-term weather conditions, and the agency is continually improving its submodels using state-of-the-practice operationally available forecast products such as the Climate Prediction Center's seven-day quantitative precipitation forecast (QPF) and the global ensemble forecast system (GEFS). QPFs are currently automatically downloaded and integrated into the agency's database. The utility of GEFS to improve weekly demand forecasts has only recently been finalized.

Seasonal Decision Support Tools

The seasonal source allocation tool updates at a monthly time step and has an outlook of up to three months ahead. Information from this tool is fed to the weekly decision support tools. Given current hydrological conditions, reservoir storage, operational constraints, and seasonal climate outlook, the seasonal source allocation tool provides a source allocation schedule to meet expected regional demand for the upcoming month, staying within the projected annual budget (source allocation).

Submodels include seasonal rainfall models, seasonal rainfall-runoff models, seasonal demand outlook models, and water shortage mitigation models.

Use

At the beginning of a fiscal year (coinciding with the water year), the source allocation models provide the projected demand of each member government and the projected source allocation to meet that demand. The three criteria that drive the decision tool are environmental stewardship (captured through wellfield permits; see Table 1), cost (each source has a different cost), and reliability (i.e., delivering quality water to member governments). Once a water year starts, based on year-to-date allocations, current hydrological conditions, and the upcoming weather outlook [e.g., the onset of an El Ni ño southern oscillation (ENSO) wet or dry state], the source allocation model provides an updated allocation schedule for the next one to three months.

At the beginning of every month, staff from departments across the agency coordinate to set operational expectations and identify any activities, such as scheduled maintenance, that may affect seasonal operations. The outlook for seasonal streamflow is conditioned on seasonal ENSO outlooks, provided by the Climate Prediction Center and the International Research Institute, which drive the seasonal rainfall simulation. Especially during fall and winter season these general circulation models have skill that is significant over much of Florida affecting supply availability over much of Florida. Intrayear demand changes are mainly driven by weather, not by socioeconomic factors.

Demand-side management decision making is performed in parallel and is supported by a water shortage mitigation plan (WSMP) tool that assesses the state of the system using current and recent past hydrological and system variables, such as regional rainfall cumulative deficits, streamflow deficits, and reservoir levels. It is used as an early warning system. Tampa Bay Water is different from many utilities that often use a drought management plan instead of water shortage mitigation plan. The WSMP provides a system-level response to prevailing hydrological conditions, looking not only at weather but also at initial conditions, system storage, and operational constraints. This tool classifies a current condition from Phase I (water shortage warning) to Phase IV (shortage crisis). The agency advises its member governments on any change in phase so they can impose water restrictions to suit the situation at hand. This information is fed back into the seasonal allocation models. With the information provided by the WSMP tool, member governments can decide whether to aggressively implement and enforce water restrictions or to resort to an expensive alternative such as desalinated seawater (if there is still untapped capacity). Currently, the feedback between the WSMP and the seasonal allocation models is loosely coupled.

Annual to Decadal Decision Support Tools

The systemwide performance evaluation model runs at a daily timescale over years to decades to characterize system behavior based on changes in demand projections, supply variability/availability, and system operational constraints. Demand projection updates are performed every year to capture any shift in key regional socioeconomic indicators.

Submodels include rainfall-runoff, system operation, long-term demand forecasting, and facility availability (in progress)

Use

The main purpose of the systemwide performance evaluation model is to assess the performance of the entire system for a wide array of hydrological conditions and future demand scenarios. This assessment becomes part of a regional water master plan update that must be completed every five years. Because of environmental requirements that dictate how much water can be withdrawn from surface-water sources, daily simulation time steps are used in the submodels. The typical simulation length is a few decades, and the results are used to compare system performance under current versus future conditions—for example, increased demand on the system. The scenarios created are used for quantifying the timing and quantity of new water supply sources as demand for water consumption grows. One of the key characteristics of the submodels is that, whereas both demand and supply are assumed to be driven by random variables (weather and socioeconomic factors), they are correlated and their joint probability distribution is modeled to provide probabilistic future supply needs (the probabilistic gap between supply and demand accounts for the uncertainties in both).

Projecting long-range water supply demand is notoriously difficult because of uncertainties in underlying socioeconomic factors and their great variability from year to year. To respond to this difficulty, the agency accounts explicitly for uncertainties and plans water supply needs through an adaptive management framework that aims to produce incremental supply as needed. One of the most challenging aspects remains the apparent mismatch between the timescales of demand projection uncertainties and the timescale required to bring additional supply online. For example, five to ten years is needed to complete a water supply project depending on complexity, public acceptance, and financial, and regulatory constraints; however, decisions must be made under uncertainty without under- or overbuilding facilities before regional demands are met.

In planning future water supply projects, Tampa Bay Water is incorporating a level-of-service concept to capture member governments' willingness to accept a given level of risk, to plan for it, and to adopt a management strategy to handle extreme events rather than build a facility to accommodate them. This management strategy is at the core of the water shortage mitigation plan (WSMP) model, discussed earlier, which is used as an early warning system to execute demand-side management options. Combining a level-of-service criterion and quantifying the reliability of meeting that criterion through a full probabilistic description of demand and supply, given environmental and operational constraints, enables the agency to meet a multiobjective goal (cost, environmental sustainability, and reliability of supply).

As part of optimizing system operational efficiency, a facility availability model is currently being developed. When complete, it will provide information on the sustainable capacity of facilities (not to be confused with rated or daily max capacities). Currently, a facility's sustainable capacity is based solely on expert opinion and is deterministic in nature. The new model will provide a sustainable capacity that is probabilistic and embeds both expert opinion and data-informed rules. Developing such a model for each major treatment or pumpage facility entails estimating its mechanical reliability through fault tree mapping and failure probability estimation. Failure probabilities are estimated using laboratory-based data (manufacturer supplied) or actual data collected through a computerized maintenance management system (CMMS). The goal is to incorporate both planned and unplanned (based on mechanical reliability) outage and to generate facility availability realizations that will be linked to system models, replacing deterministic operational constraints.

Because of the nature of the problem, such models have high computational demands, which are met through a distributed computing environment where simulations are performed on clusters of computers. In this way simulations that would take months to complete on a single machine are completed in matter of days. One of the advantages of this approach is that, once the simulations are done, a variety of postprocessing questions (what-if scenarios) can be answered because the full probability distribution of demand and supply are numerically estimated.

Summary of Value and Broader Challenges

Tailor-Made versus Commercially Available Software

Utilities face a unique challenge in the complexity of the system they use and the specific questions they need to answer: whether to use commercially available systems models, develop their own, or combine options. For example, Tampa Bay's weekly decision support tool is a tailor-made optimization model with a simulation engine (by way of a unit response matrix) driven by an integrated hydrological model based on *MODFLOW* and *HSPF*. This model is required to capture central Florida's unique hydrological characteristics (significant interaction between open-water and shallow-groundwater level that requires physical modeling). The driving force behind it is environmental: it is part of regional strategy for the recovery of wetlands and lakes from heavy groundwater pumpage. The agency's systems models at seasonal and decadal scales are mainly tailor-made and based on the *MATLAB* computing environment. This environment provides unparalleled flexibility in model development, visualization, and distributed computing capabilities.

Gap between Applied Research and Real-World Applications

A number of utilities across the nation now recognize the need to partner with universities and local research institutions to solve complex water resources management issues. One example is the Water Utilities Climate Alliance (www.wucaonline.org). However, the challenge is moving a partnership beyond one project at a time toward a framework that regularly transfers knowledge and methods that can be used in the real world. If these models/methods are implemented and vetted in field applications, they provide important feedback to the scientific community on their limitations and opportunities for improvement. From a utility perspective, these projects can be seen as continuous improvement tools—for example, to optimize an operational and/or planning activity. A key challenge that the applied research community needs to address is how to avoid the use of simplifying assumptions that may limit the usefulness of models/methods in a practical setting.

Utilities' Commitment to Maintain In-House Expertise and Computing Resources

Transferring applied research into real-world applications is a two-way street. Although the research community advances the use of innovative decision support tools in the real world, utilities must be in a position to take advantage of such tools to guide short- and long-term operational and planning activities. Unfortunately, decision makers may not be equipped to sufficiently evaluate the benefits of innovative decision support tools. In addition, the challenge of climate change means that utilities cannot assume that the future will be consistent with the past. The benefits of collaborative effort are significant: a resilient infrastructure able to withstand uncertainties, savings to rate payers through systems optimization, protection of the environment, and encouragement of sustainable water resource use now and the future.

Acknowledgments

The author thanks Dr. Patrick Reed of Cornell University for his valuable input.

References

Fourer, R., Gray, D. M., and Kernigham, B. W. (2003). *AMPL: A modeling language for mathematical programming*, 2nd Ed., Duxbury Press, Canada.