The energy-efficiency benefits of pump-scheduling optimization for potable water supplies

S. M. Bunn
L. Reynolds

Water utilities across the developed world have been installing and operating telemetry and SCADA (supervisory control and data acquisition) facilities for at least three decades. They have amassed substantial quantities of historical operational data held in databases called data warehouses. Significant interest exists for extracting the underlying value from this data to support decision making and performance improvement. The water industry uses approximately 3% of total electricity production in developed countries such as the United Kingdom and the United States. Up to 90% of this electrical energy is consumed by pumps. Small improvements in pump efficiency will yield significant reductions in energy consumption with consequential reductions of carbon emissions to the atmosphere. Technologies to initially select a pump to match the expected performance requirements and then maintain optimal performance through periodic refurbishment are well established. Dynamically optimizing the scheduling of pump operation to improve efficiency under changing diurnal and seasonal water demand patterns is far more complex. This paper summarizes 20 years of progress in the development of pump efficiency improvement techniques and focuses on real-time dynamic optimization technologies and data-mining techniques to improve energy efficiency. Case study results from an automated real-time commercial pump-scheduling system, Aquadapt™, are presented.

Introduction

Electricity consumption by water and wastewater utilities has typically accounted for about 3% of all electrical energy consumption in the United States and United Kingdom [1]. This represents a net annual consumption of 75 billion kWh in the United States, which corresponds to approximately $4 billion of consumption. Potable water alone uses 4.4 billion kWh in the United Kingdom, corresponding to £200 million in 2007. Between 90% and 95% of the electricity purchased is used for pumping [2]. Attempts to increase energy efficiency can reduce electrical consumption for pumping by as much as 5%–25%, and sometimes even greater reductions are realized. Energy efficiency is also associated with other benefits such as reducing the greenhouse gas (GHG) footprint of the utilities. While significant progress has been made in analyzing individual pumps and then matching pump characteristics to specific requirements for the duties of the pumps, these endeavors have relied on assuming that each pump runs at a single pressure and flow operating point. Pump operating requirements are subject to widely variable seasonal and diurnal fluctuations, making it very difficult to consider all operating contingencies. In addition, pumping facilities and water distribution networks have become highly interconnected and complex systems. Pumps do not operate in isolation; for example, because of the incompressibility of water, it is typical that any change in the operating duty of one pump may affect the suction or discharge pressure of other pumps connected to the same pipe system. These other pumps include those operating within the pumping station and, to a lesser extent, pumps elsewhere within the distribution network. These kinds of interdependencies lead to a highly dynamic system.

0018-8646/09/$5.00 © 2009 IBM

This paper focuses on pumping and water distribution networks that supply potable water to customers from treatment plants. Since the early 1980s, these networks have had an increased dependency on telemetry and SCADA (supervisory control and data acquisition) facilities to provide the interface between operators and the equipment to be controlled. Many of these systems read, store, and provide time-stamped values of flow, pressure, and level in the water distribution network, with time intervals ranging from a few seconds to tens of minutes. Many thousands of values are stored at these frequencies, and very large historic databases have been created over the years. Unfortunately, uses for these data have not proceeded apace, and water utilities now recognize that the data could be a valuable resource for improving future performance.

Pump energy efficiency

As mentioned, one aspect of the operation of water networks that is receiving particular attention is energy efficiency and the associated GHG implications of low efficiency. It is reasonable to ask whether the historic data can be used to improve future operation. Given that energy use by water utilities is nationally significant, and more than 90% of this use is for pumping, one might assume that reasonable effort is made to ensure pumps are operating efficiently. In practice, this is often not the case. Energy has been relatively inexpensive, and cost-saving efforts have been focused elsewhere. A major study of pumps [3] commissioned by the European Commission in 2001 determined that very little effort was being made in this area. Three strategies were recommended. First, better efficiency information at procurement time is needed when matching pump characteristics to duty. Second, efforts are needed to recondition pumps in order to counter efficiency deterioration. Third, pumps should be operated at their best efficiency point.

This study also indicated that the efficiency curves of pump manufacturers typically include significant tolerance allowances, which made it difficult for a designer to assess the actual operating efficiency. Further, it was noted that the purchaser's key criteria for pumps were initial purchase price and delivery, reliability, and efficiency. In fact, efficiency was sometimes not considered. This ordering of priorities may be counterintuitive because energy consumed by a pump over its lifetime of 30 or more years results in 90% or more of the life-cycle costs, with initial purchase cost being as little as 2% and the remainder consisting of maintenance costs.

The reconditioning of existing pumps has been shown to be of significant benefit. Pumps wear over their operating life, and although potable water generally wears a pump significantly less than wastewater, performance for a potable water pump will degrade by approximately 1% per year. The European Commission study examined surface roughness of various parts of the pump and potential gains that might be realized from polishing or coating these parts. Efficiency improvements ranged from 5% to 18%.

This gradual performance degradation in potable water pumps over the course of time may be a factor in the lack of urgency exhibited by operations or maintenance staff to institute a proactive preventive maintenance program to measure and improve pump performance. Where water utilities have addressed the needs of pump reconditioning, the results have led to paybacks measured in months rather than years. A coordinated approach of matching pump to duty, along with the refurbishing of pumps, can readily lead to improvements of 20% or more.

A United Kingdom-based company, AEMS Ltd., specializes in pump energy management. This company presented the results for a water pump that has been in operation since 1963 [4], during which time it has run for approximately 100,000 hours without a major overhaul. Running costs were about $160,000 per year. The pump had an average efficiency of 70%. To refurbish the pump so that it would achieve an efficiency of approximately 82% would cost $20,000 and had the potential to save more than $26,000 per year in electricity costs. This leads to a payback period of only 9 months. A second option was to install a new high-efficiency pump and motor at a cost of $60,000, with expected electricity cost savings of $32,000 per year. This leads to a payback period of 1.6 years and allows for the new pump to match the average duty requirement.

In the European Commission study, pump refurbishment and pump duty matching (i.e., making sure the pump had the appropriate size) were two of the three key items for cumulative potential. The third item involves the scheduling of pumps so that they operate closer to the best efficiency point.

Existing pump-scheduling solutions

A number of methods have been proposed or developed to establish dynamic pump-scheduling processes in order to optimize the scheduling of pump run times. Few if any of these methods have been implemented. These methods are intended primarily to leverage time-of-use (TOU) tariff structures in order to minimize energy costs [5–13].

TOU tariffs reflect the real-time supply-and-demand market in energy pricing. When demand is high, the tariff is also high. This typically occurs when heating or air-conditioning loads are high. At night-time, when demand for electricity is low, the tariff is also low. Water utilities are in an advantageous position to maximize pumping to storage during off-peak hours when energy rates are lowest, and then the utilities draw on water in the storage tanks when energy costs are high.

Due to the complexity of most water distribution networks and the multiple production requirements and operational constraints, it is generally considered virtually impossible to explicitly solve the scheduling problem using mathematical techniques. Projecting system energy use from a given pump schedule is, however, straightforward and is already a feature of most off-the-shelf hydraulic modeling packages. Even so, a "brute-force" approach, modeling all possible combinations to arrive at the most preferable solution, is not practical. A system with N pumps that can only have the discrete states of on or off, requiring hourly schedules for a single day, has $(2^N)^{24}$ possible combinations to consider. A system with only 11 pumps has a solution space of 2.96×10^{79} possibilities, which, coincidentally, is almost the same value as recent estimates of the number of atoms in the entire universe. Most sizable water utilities have hundreds or even thousands of pumps. Some of these pumps may operate over a range of speeds so that the assumption of only discrete states is not valid, making the solution space much larger and virtually infinite. Because brute-force solutions are impractical, sophisticated search and optimization techniques have been developed to assist in solving this problem. In 1994, Ormsbee and Lansey [5] evaluated the state of the art to that date in optimization techniques for pumping systems. These researchers used relatively well-known linear and nonlinear programming and dynamic programming tools.

Since around 1994, the most common solution-finding techniques have employed evolutionary algorithms (EAs) to search the solution space. EAs involve a branch of computer science [14] that makes use of techniques borrowed from evolutionary theory to employ random simulations and trials to select "fit parent" solutions from which hopefully even better "child" solutions can be derived. These systems carry out many thousands of guided hydraulic simulation trials to converge on a solution that is a near-optimal pump schedule from the large solution space for relatively small systems. A large body of research papers exists on the application of the most popular form of EA, the genetic algorithm, to the pump-scheduling problem since genetic algorithms are well suited to binary problems in which the pumps are either on or off. A good starting point for the interested reader is a paper published by Savic et al. of Exeter University in 1995 [6]. Related techniques such as swarm optimization [7] and simulated annealing [8] have also been successfully applied to theoretical problems. Several commercial applications of these technologies exist, mainly focusing on the design optimization of pumping systems rather than operations [9]. Most of these applications involve a single objective to reduce pumping costs with respect to TOU electricity tariffs with only small numbers of pumps. Exploratory work has been carried out in multi-objective optimization where more complex factors are evaluated such as reducing the number of pump starts in order to reduce pressure surges [10] or minimize peak electrical demand charges [11]. A useful comparative analysis of various genetic algorithm techniques as applied to pump scheduling was published by Sotelo et al. in 2002 [12].

More recently, the European-based Potable Water Distribution Management (POWADIMA) project [13] has considered speeding the solution process through the use of artificial neural nets (ANNs) to replace the classic hydraulic simulation, and researchers have demonstrated a 1,000-fold improvement in speed. In this study, using standard personal-computer hardware and modeling a modest-sized system that includes 17 pumps for the city of Valencia, optimal solutions were obtained in about 10 minutes. Here, the solution finding benefited substantially from the use of a known feasible solution as a starting point. No time was given for finding the initial feasible solution. This system is yet to be fully implemented but does indicate that the potential exists to use genetic algorithm techniques in real systems, as long as pump numbers are kept quite small.

Attempts to find solutions have been made in order to optimize pump scheduling with respect to TOU tariffs and that also integrate energy efficiency indirectly through pressure control. Here, the assumption is that pumping at lower pressure will reduce energy requirements and have the side benefit of reducing water losses in leaks. Research in prior published papers has generally not made use of real-time calculations of energy efficiency in the optimization function. In this respect, the functioning of the commercial software package Aquadapt** [15, 16] from Derceto Ltd. is unique. This software avoids genetic algorithm techniques and takes advantage of the operating speed of linear programming and mixed integer linear programming, combined with heuristic and nonlinear formulations. This software also explicitly models pump efficiency on the basis of live (i.e., real-time) performance data. It primarily reduces electricity costs for a utility by exploiting differential prices in the energy tariffs while also reducing peak electricity demand charges. Aquadapt specifically addresses flow and pressure constraints, storage-tank-level constraints, pump minimum run times, and cool-down requirements. It can solve problems involving integer solutions, for example, for pumps being only in on or off states, as well as provide continuous solutions such as those associated with variable-speed pumps and flow-control valve setpoints. The software also achieves cost reductions through utilization of lowest-cost water

Table 1 Relative size and complexity of four U.S. water utilities that have implemented Aquadapt.

Customer system	*Population served*	*Storage tanks*	*Pressure zones*	*Pump stations*	*Pumps*	*Control valves*	*Demand (million liters/day)*
East Bay MUD, California	660 thousand	28	26	20	66	4	160–480
Washington Suburban, Maryland	1.6 million	57	15	18	81	25	640–900
WaterOne, Kansas	570 thousand	25	3	26	84	11	190–400
Eastern Municipal Water District, California	630 thousand	68	38	59	149	9	180–450

sources (taking into account that production costs at different treatment plants are not identical) and directing water to the customers over the least expensive possible path in the distribution network.

One case example use of Aquadapt involves the Eastern Municipal Water District (EMWD) from Perris Valley, California. This district used the Aquadapt software package in 2006. The software reduced energy costs by more than 13% and improved efficiency by 8.4%. EMWD received the California/Nevada American Water Works Association (AWWA) Outstanding Energy Management Award in 2007 for this project. East Bay Municipal Utility District (EBMUD) in Oakland, California, has achieved almost identical percentages for cost savings and efficiency improvements. California has made a commitment to GHG reduction at the state level in Assembly Bill 32. With 19% of the energy use of a state [17] involved with the transportation of both raw and potable water, this has led to a focus on energy-efficiency improvements in the water sector. Water District No. 1 of Johnson County, Kansas (WaterOne), operates in a flat electricity tariff market yet achieved cost savings of 20% through peak-demand charge reductions and energy-efficiency improvements.

The Aquadapt algorithm is identical for all four utilities in **Table 1** and requires only three iterations to achieve near-optimal schedules for all water sources, control valves, and pumps, including variable-speed pumps. For even the largest of these utilities, the algorithm requires only a minute or two to find a solution. The optimization algorithm uses the client's own SCADA facility to provide accurate live and historical performance data. However, significant obstacles must be overcome in order to use this data as part of a reliable solution. Some of these obstacles are mentioned in following sections of this paper.

Data mining from large real-time databases

The term *telemetry data* refers to input information received from remote instruments, and this data is stored in a centralized database. Inputs can be digital signals such as the running status of a pump or analog signals such as the amount of flow through a pipe. These signals are generated by instruments, switches, contacts, and other devices in the field. The signals are then captured and converted by devices such as a programmable logic controller (PLC) or remote terminal unit (RTU) located near the instruments. The PLC or RTU then communicates with a remote central server using means such as radios, dial-up telephone lines, cellular phones, or fiberoptic cables. The type of communication path dictates both the quantity and the quality of the data that can be sent. While digital signals can be transmitted relatively efficiently, analog signals usually require digitization, preferably without significant loss of accuracy. A typical out-station, which is another name for an RTU, will use 12-bit analog-to-digital converters, limiting the accuracy of the signal to at best 0.25% of the full range of the instrument. This is acceptable, as the instrument itself is unlikely to have better inherent accuracy. Most analog instruments are programmed to "fail low" under adverse conditions. For example, a 4-mA to 20-mA signal may go to 0 mA to indicate a fault. The RTU may not be able to process this "out-of-range" low signal and may then pass incorrect data to the central station. Many radio systems carry out "report-by-exception" methods in which the system sends new data only when the signal has changed from its previous value by more than some defined threshold. If the thresholds are set too far apart—for example, for thresholds relating to storage-tank levels—then new values may be sent very infrequently during those periods when the levels are barely moving at night.

By far the largest cause of signal error is dropout, that is, a break in the communications path from the outstation to the central server. This means that for periods ranging from minutes to hours, no updated data are received by the SCADA facility from the outstation. Often, the SCADA facility displays the last-known good value when this condition occurs, and it is up to an observant operator to notice that a problem exists.

We have analyzed historical databases from nearly 100 water utilities since 2000. Not only do errors exist in every database, but it is uncommon to find a single day of data without any errors over a span of many years. Removing all "bad days" (i.e., days with data errors), therefore, is not a useful option, as it leaves insufficient valid data to allow useful analysis. However, careful re-creation of missing data can make the entire dataset useful. Full re-creation of data requires using known hydraulic characteristics of the water distribution network together with the use of trusted information. For example, if a pump run status is *on* and the reservoir it supplies is filling, then it can be reasonably assumed that a flow meter reporting zero flow is incorrect. Re-creation and replacement of incorrect data is extremely labor intensive because of the need to address many forms of possible data corruptions.

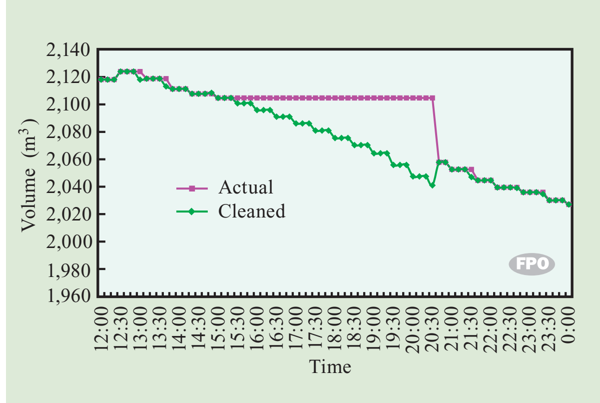

Figure 1

Data dropout (actual data) and correction ("cleaned" data).

Building a new dataset

Rather than store inaccurate data, it is better to detect and correct errors in real time, and then store corrected data only. In practice, this is difficult to achieve. The benefit, however, is a clean (i.e., correct) dataset that can then be used for multiple decision-support and regulatory-reporting purposes. When correcting in real time, it is not possible to use actual future conditions to confirm values or make certain kinds of interpolations or extrapolations. Decisions on accepting, rejecting, or correcting values must be made on the basis of current knowledge of the state of the system.

The methodology used in Aquadapt involves the running of a hydraulic model every 10 minutes, while also receiving updated values from SCADA. The approach also makes use of scheduled future pump and valve settings to predict system behavior. If actual telemetry data diverges from expected behavior, there is a mechanism for detecting errors and also a means of replacing erroneous values with corrected data from the model predictions. These replaced values are tagged as such in the database so they can be identified in later analysis if required. Care must be taken not to reject valid, but unpredicted, values such as low flows created by a pump failing to start.

After seven years of development, the Aquadapt software module called the "data cleaner" is able to achieve reasonable and consistent results. One example of this is shown in **Figure 1**, which shows that a level-signal outage occurred for five hours. The corrected and stored values indicated by the green line are a much better representation of the storage tank level than the live data on the purple line. Resynchronization with live data can be seen when new values start being received again from the SCADA facility. As shown in Figure 1, though better than the raw values, the corrected data are still an approximation, and hence, errors increase the longer the instrument is offline.

The ability to create accurate calibrated hydraulic models of water distribution networks is well established and relatively straightforward. These models are used by planners and engineers. The hydraulic model can be used to accurately predict the future behavior of the hydraulic system for any given starting conditions and known operating parameters for relatively short time horizons such as the next 12 hours. Longer periods of simulation become less reliable because of the uncertainty inherent in estimating future water demand. Armed with these models, large quantities of missing data for at least the current time period can be recreated in real time. Where possible, corresponding values from alternative instrumentation are compared with the hydraulic model forecasts, and if instrument data appear incorrect, the model values are used in place of the missing or erroneous data from the instrument.

Mixing real and assessed data

The availability of databases without missing data and that contain corrected data allows for uses beyond pump-scheduling optimization. Some examples include statutory reporting, fault analysis, and asset management. These databases are large; for example, one such database for a water utility serving about 400,000 customers grows by about 8 GB/yr with several thousand data points added every 10 minutes. Accessing and analyzing these data can be very complex; for example, even the 65,000 rows of the Excel** spreadsheet application are a serious limitation when a user wants to view graphs and trends, and therefore, data visualization tools are essential.

The concept of a *business intelligence dashboard* that shows various key performance indicators (KPIs) is becoming common in a wide variety of industries to monitor performance [18]. These systems generally involve HTML pages on an internal corporate Intranet site. Although not required for optimization, the Aquadapt Dashboard was created to follow this concept and present a graphical view into the large, live, and historical databases.

For a water utility, this data covers production and supply KPIs such as volume of water delivered, cost and utilization of assets, and customer demand for water by geographic area. While this coverage is extremely worthwhile, there is significantly more financial value that can be extracted from the database, especially in the areas of improving asset performance and, specifically, pump performance.

Recall that forward (i.e., future) prediction with a hydraulic model is trivial to accomplish using standard PC hardware. Most hydraulic modeling packages allow the user to calculate expected pressures, flows, pump energy consumption, and energy efficiency. Even on very large systems, these calculations can be accomplished in seconds. If the data is validated through careful calibration of the model and perhaps a few days of on-site manual measuring of unmonitored values such as pressure, then the resultant calibrated model can be used to reasonably estimate actual performance of the pumps. If this approach is combined with real-time (live) data, then useful *virtual-instrumented* values can be derived, where no physical instrument exists, and stored alongside physical values in the database. A good example is using the model to estimate individual pump flows for a multi-pump site even though only a single flow meter exists for the entire pump station.

Fixed parameters such as represented by pump curves and certain pipe data need only be updated every few years, and live operating data can be transferred directly into the hydraulic model from the SCADA facility. Typically, the only extra information the model requires in order to predict future values is a future schedule for pumps, valves, and settings for flow control devices. Given this mix of actual and predicted data, very complex relationships can be evaluated.

Extracting pump performance data

To assist in understanding the next section of this paper, a short primer on pump theory is useful. Water is generally pumped using centrifugal pumps. A modern centrifugal pump transfers water from its suction side, typically near the shaft of the pump and accelerates it by centrifugal force applied by a spinning impellor, reminiscent of an airplane propeller. The high kinetic energy water is then discharged into a chamber near the outer radius of the pump with the kinetic energy converted to pressure. The energy required to accomplish this is typically supplied by an electric motor, gas engine, or diesel engine. Newton's first and second laws of motion are applicable, and the energy imparted into the water can easily be measured by examining a mass of water that a pump may deliver to an increased discharge height in meters or feet. The time taken to deliver the water, which is related to the rate of flow, gives a measure of the input power required.

As the flow from the pump increases, frictional losses in the pipes also increase in proportion to the square of the flow rate. Hence, doubling the flow rate quadruples the frictional losses, making the pump work harder. The flow rate, lift, and losses must be taken into account by an engineer when selecting a pump. There are additional complications to consider since the discharge pressure of the pump is determined by many other factors such as water demand and elasticity of this demand, which relates to the degree to which demand is pressure dependent. For example, the higher the pressure at a tap or a showerhead, the larger the volume of water used in a fixed time period. The change in discharge head with flow is represented as a *system curve*, and it will be different between day and night, depending on water demand. The term *head* is often defined as any resistance to the flow of a pump. When pump manufacturers indicate a head pressure, this refers to a vertical discharge pressure that may be thought of as the vertical lift in height (e.g., measured in feet of water) at which a pump can no longer exert sufficient pressure to move water. The greater the head pressure of a pump, the more powerful the pump.

A designer typically selects a point on the system curve to correspond with an average flow demand sufficient to easily fulfill storage and supply demand within a reasonable time, for example, 8 hours. This means that the pump will satisfy demand on typical days, keeping another 16 hours of pumping time in reserve for peak demand days. A pump is then selected so that it runs at its best efficiency point (BEP) on the expected system curve at the desired head and flow. In **Figure 2**, a typical set of pump curves is shown with the head curve, the efficiency curve (which indicates the variation in efficiency over the pump flow range), and the power curve all depicted on the same graph. The BEP occurs at the apex of the efficiency curve. Here, the term "power" refers to the shaft input power, and we note that this pump could be driven by diesel, gas, or even a water-driven turbine.

The problem is that with demand changing diurnally, from high demand in the mornings and evenings and low demand overnight, it is very difficult to select a pump that will operate optimally over the entire expected operating range. A designer will also take into account expected growth in demand in an area over the approximately 40 years of anticipated pump operating life and tend to specify a much larger pump than is required at the installation time to cope with future projected growth.

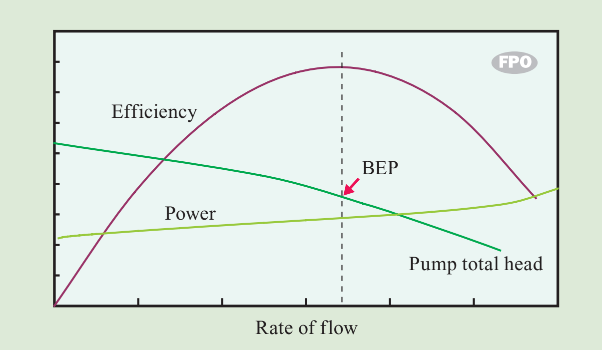

Figure 2

Generalized pump curve example showing best efficiency point (BEP). (Axis units are arbitrary for these schematic curves.)

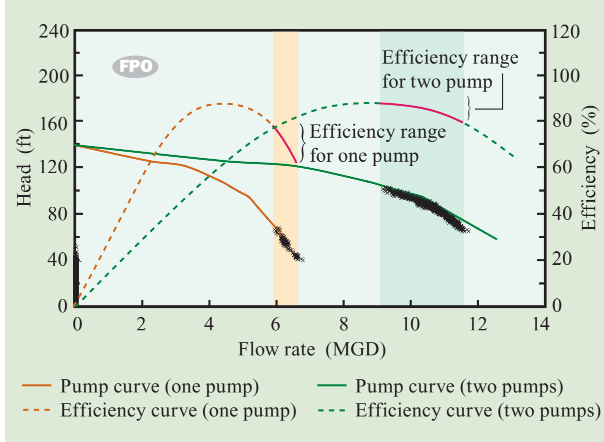

Figure 3

Telemetry data overlaid on pump curves. The efficiency ranges shown as purple curve segments are simply those portions of the efficiency curves that correspond to the data points on the pump curves.

This means that pumps rarely operate close to their BEP even when new. A third issue is that the pump will wear over its lifetime, changing its characteristics from the manufacturer's curves.

What is required is a process that calculates actual pump curves and system curves in real time and then matches the best available pump or selection of pumps to the current instantaneous duty. If a pump only needs to run eight hours in a day to satisfy demand and replenish storage, then this introduces further options with respect to which hours to use during the day. There is no universal rule to aid in pump selection such as always choosing the newest pump. Consider that a badly worn pump may run much more efficiently at night, while a new pump may run more efficiently during the day at one pump station and exactly the opposite situation may be true at another pump station in the same network. In **Figure 3**, the pump curves for a single pump and generated curves for two of these pumps operating in parallel are shown. Here, the generated curve is calculated from data from the two individual pumps. The corresponding efficiency curves are also given. Actual telemetry data, after validation, has been overlaid as ticks on the graph. Note that two pumps working together do not give twice the flow of a single pump; for example, note that about six million gallons per day (MGD) flow for a single pump and only approximately 10 MGD for two pumps. Recall that pressure is dependent on the square of flow; thus, the use of two pumps creates higher back-pressure in the discharge pipeline. In this example, one pump operates between 40 and 60 feet of head, but two pumps generate 65 to 100 feet of head, this extra back-pressure is primarily a result of the increase in flow rate. Note also that the ticks indicating telemetry values cover a range of values on the pump curves, because a changing diurnal demand exists. Efficiency may be determined by reading the efficiency curve located above the operating point on the *y*-axis. Also note that in this example, two pumps operate between 80% and 90% efficiency, while one pump on its own only operates between 60% and 76%; thus, it is always better to run two pumps together.

Using a hydraulic model to determine head, flow, and energy for each individual pump and all pump combinations is relatively simple for a single pump station, as long as water demand is predicted reasonably accurately, for example, for the remainder of the day. The use of variable speed pumps, where the motor driving the pump can operate over a range of speeds, is becoming more common by water utilities. These variable speed pumps provide increased operational flexibility; however, they also vastly increase the complexity of the calculations. Actions at any one pump station may affect the operating point of many other pump stations in the network, requiring all calculations to be repeated. As a result of these interactions, scheduling pumps for optimal efficiency needs to be solved globally. This is a much more difficult problem to solve than pump scheduling for tariff concerns alone, and a review of the available literature suggests that this has only been addressed in the Aquadapt solution.

The solver in Aquadapt pre-calculates expected flow, energy, and efficiency for all pumps in the distribution network for a generalized solution with separate values for each half-hour timeslot of the day. It then uses these *a*

Table 2 Greenhouse gas reductions for four U.S. case study systems. [The Emissions and Generation Resource Integrated Database (eGRID) is a comprehensive inventory of environmental attributes of electric power systems. EPA: U.S. Environmental Protection Agency.]

Customer system	*Average MWh per year*	*Average efficiency gain using Aquadapt*	*EPA eGRID 2004 CO_2 emissions (tons/MWh)*	*Extrapolated CO_2 reduction per year (tons)*
East Bay MUD, California	26,000	6.1%	0.502	800
Washington Suburban, Maryland	7,000	8.4%	0.515	300
WaterOne, Kansas	99,000	8.3%	0.547	4,500
Eastern Municipal Water District, California	94,000	6.0%	0.845	4,800

priori data in the first iteration of its solution to produce a pump schedule. The solver then uses this schedule in the hydraulic model in combination with actual telemetry data, flows, and pressures to generate more accurate pump flow and energy values. This creates the second iteration, which is again evaluated by the hydraulic model to generate new predicted flows and pressures. By the third iteration, new values for expected flow, pressure, and energy are usually sufficiently close to the second iteration to no longer generate changes in the pump schedule, and hence, the Aquadapt solver has converged on a near-optimal solution. If successive iterations are run beyond this point, they tend to oscillate between two similar pump schedules with almost identical expected costs. A simple example of this kind of oscillation involves an eight-hour block of time with a fixed tariff, such as during the off-peak night period in which only four hours of pump time is required to fill a tank. Any four hours could be selected based only on tariff to give the same cost; however, when combined with efficiency data, Aquadapt is able to select the most efficient four hours. There may be very little cost difference between two similar schedules, hence leading to the tendency to oscillate.

Using both energy and efficiency in the objective function has generated some interesting results in the Aquadapt implementations. If the optimization algorithm only targeted tariff cost reduction, the algorithm would lead to attempts to move energy use from peak times to off-peak times. If the differential in tariff between peak and off-peak is substantial, this is indeed what the Aquadapt algorithm will do. Where tariffs are very flat, for example, 5 cents/kWh peak and 4.5 cents/kWh off-peak, then with only 0.5 cents between tariff bands, the Aquadapt solver may determine numerically that the efficiency seeking process leads to lower cost since a single kilowatt-hour saved due to delivering water more efficiently is worth at least 4.5 cents, nine times that of moving kilowatt-hour use from peak to off-peak. This also exposes the weakness of not taking efficiency into consideration in the genetic algorithm-based pump-scheduling solutions proposed so far, which only involve time-of-use tariffs.

Results from four case histories

Four significant implementations of Aquadapt have been running for up to four years in the United States. They have achieved audited energy efficiency improvements of between 6% and nearly 9%, and energy cost savings of 12%–25% per annum (Table 1 and **Table 2**).

The cleaned, corrected data from multiple sensors are combined with calculated data from "virtual instruments" (which make use of model-based estimates) and stored in the very large historical databases that are then "data-mined" and displayed using a dashboard visualization tool. The databases range in size from 3 GB to 22 GB to date and are growing at a rate of between 1 GB and 8 GB per year depending on instrument numbers and update frequency. As mentioned, the dashboard is implemented as a series of Intranet pages on the client's network.

Figure 4 illustrates an example screenshot of data extracted from a large dataset and shows water demand data as a bar graph for the week starting on October 29, 2007, with actual values in yellow plotted next to average expected values in blue. As can be seen, demand for this week was higher than the seasonal average every day. What may not be so clear is that expected values are categorized by day type, and weekend demand was expected to be slightly lower than weekday. The blue bars are slightly shorter on the weekend days. The pie graph breaks down these total data into the client-specified pressure zones or customer areas. Total water-demand figures are very commonly used by water utilities, mainly for mandatory reporting for abstraction licenses, which are regulatory licenses to divert either surface water or ground water for a designated purpose. Typically, this reporting was achieved manually from a limited number of data loggers or dial counters located at each water source, and then large numbers of calculations used to

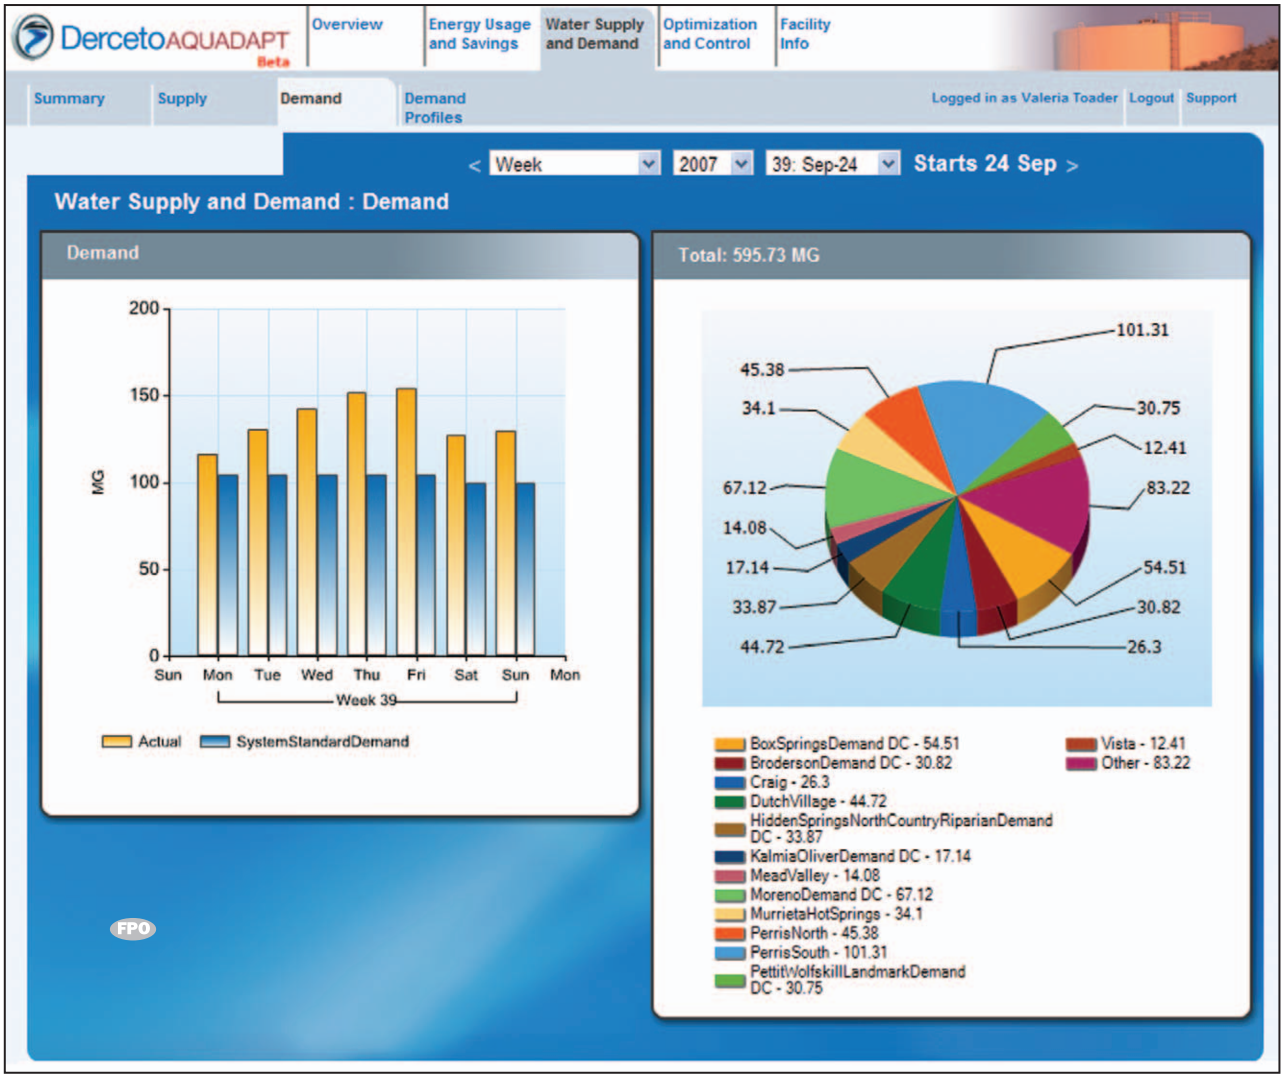

Figure 4

Dashboard view of values extracted by data mining.

determine where this water went to in the distribution network. A number of water utilities in the United Kingdom reported that this was a task requiring many man-months to complete for annual statutory Ptime and store the results. An audit trail is maintained by tagging values that are from virtual instruments, and storing the validation algorithm used.

Some of the most interesting data to be extracted from the database is the pump station energy and efficiency data. An overall analysis of energy use for all installations has shown that Aquadapt is running pump stations more efficiently than before the software was installed. **Figure 5** is a plot of energy use, in units of kilowatt-hours, per million gallons (MG) delivered for the largest pump at one Californian utility pump station. Efficiently delivering a fixed volume of water with a pump means delivering fewer kilowatt-hours per million gallons than compared to the case of running the same pump inefficiently. Values at the left side of the x-axis of the graph are therefore more efficient than those to the right. Note that Aquadapt is operating at an 11.5% improvement in efficiency, using an average of 605 kWh per MG instead of 685 kWh per MG before it was installed. Aquadapt is also doing this consistently, as indicated by the reduced variance of the operating band.

Recall that pumps have an operating curve of flow versus pressure. The actual operating point of a pump is a function of the pump curve and its interaction with the

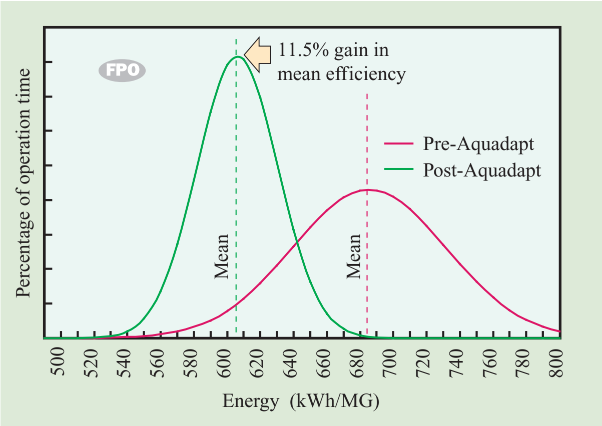

Figure 5

Pump station uses less electrical energy to deliver water after installation of the Aquadapt optimizer.

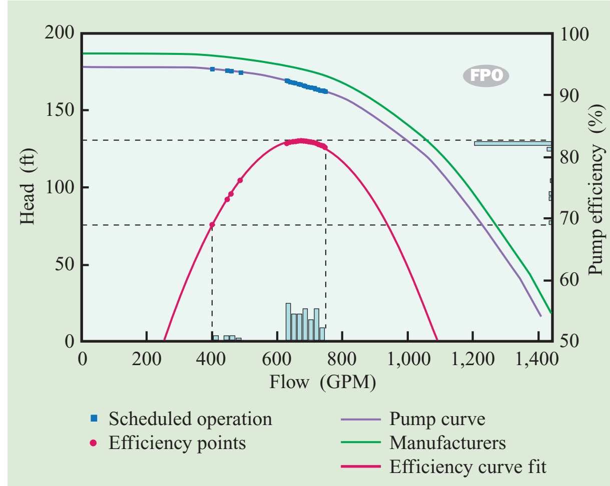

Figure 6

Pump station curve generation in real-time (screenshot from graphical dashboard).

system curve, which reflects the pipe network and the water demand in that network. If we plot half-hourly samples of flow versus pressure across the pump, we obtain a representation of how the pump is actually behaving throughout the day. While this would be very useful, the problem with this approach is that instrumentation, especially flow meters, are expensive to buy and install. In practice, it is rare to find one flow meter per pump and sometime not even one serving an entire pump station. As mentioned, Aquadapt is able to provide "virtual instruments" that indicate estimated flow for each individual pump in a multi-unit pump station. Mining and displaying this virtual data alongside real instrument data facilitates an understanding of pump station operation.

In **Figure 6**, the manufacturer's pump curve, shown in green, and the current measured curve (which we sometimes refer to as a calibrated curve), shown in purple, are almost identical. This is typical of a new pump. Blue squares on the purple pump curve indicate where each half-hour flow/pressure sample lies, with a histogram on the flow axis (x-axis) showing the frequency of operation at each flow range. These graphs can be extracted from the current live operation or retrieved from the historical database. The visualization tool is programmed to show operation on the pump curve where individual flow meters do not exist for each pump. The calibrated manufacturer's "efficiency" curve in pink also makes use of the same operating points, with a histogram to the right showing the frequency of operation at each point on the efficiency axis (y-axis).

This kind of data has rarely been presented before. Traditionally, this would have required an expensive pump data-monitoring tool on each pump, which is often impossible to achieve due to restrictions in the degree to which flow meters can be effective when installed close to pumps. By viewing the pump curve in Figure 6, we see that, most of the time, the pump operates close to peak efficiency of approximately 83%, but note that only a small increase in discharge head of as little as 5 feet causes a reduction in flow from approximately 700 gallons per minute (GPM) to 400 GPM and a substantial drop in efficiency to 69%. The increase in head may be due to dynamic water demand changes or a reservoir-level rising, causing a static pressure increase. For example, if operators preferred to keep reservoirs completely full, we might find that this pump is always operating at that inefficient lower flow rate end of the curve, whereas allowing the reservoir to drop by a couple of feet can significantly improve efficiency.

Figure 7 corresponds to a pump station with two pumps. When purchased, these were specified as a nominal 2,000-GPM pump and a smaller nominal 800-GPM pump. As can be seen from the graphs, the smaller pump is not operating at the specified flow and is actually averaging more than 1,000 GPM, possibly because it is operating at lower pressures than the design called for. Here, it is interesting to see that the 2,000-GPM pump is well worn, as indicated by the fact that the current measured pump curve, shown in purple, is more than 25% lower than the original manufacturer's curve, shown in pink. Yet somewhat surprisingly, the pump is operating well, with an average efficiency of 88%. Note also that the second pump, while almost new as indicated by the very

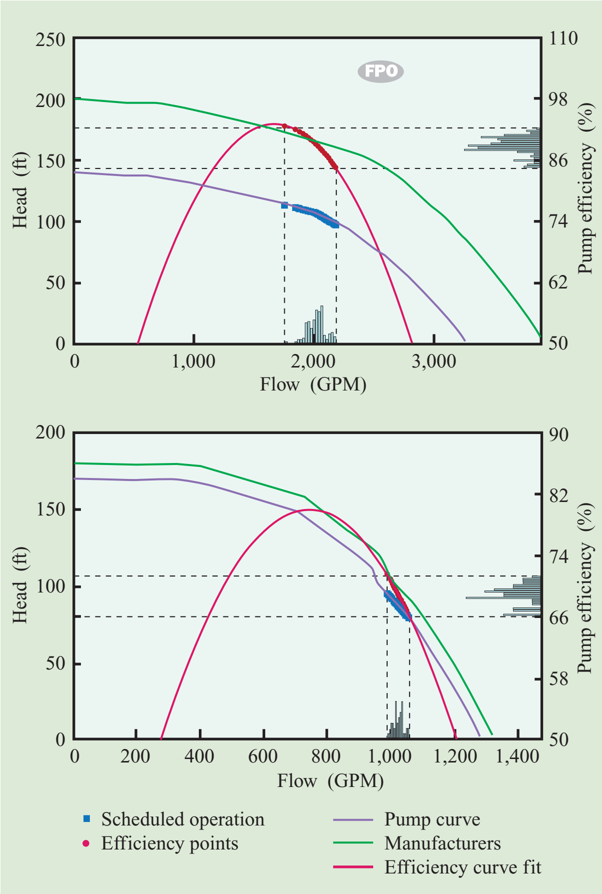

Figure 7

Two fixed-sized pumps at the same station have different characteristics and operate with differing efficiency (screenshot from graphical dashboard).

similar manufacturer's and actual pump curves, is not operating as efficiently. The 800-GPM pump is experiencing only 80–95 feet of total dynamic head (TDH), which is a measure of the differential between the suction pressure and discharge pressure of a pump. This is approximately 25% less than the ideal value of 125 feet of differential pressure it needs to operate at the top of its efficiency curve.

In **Figure 8**, the actual operating points (blue squares) are not on the purple pump curve because real instruments provide data on flow and pressure for this pump; thus, we know exactly where this pump is operating. In particular, the blue squares below the calibrated pump curve (purple line), rather than on it, indicate that either a problem has arisen, such as a stuck valve, or the pump needs to be calibrated again. From this graph, we can deduce that the pump is probably oversized for the required duty because even flows as low

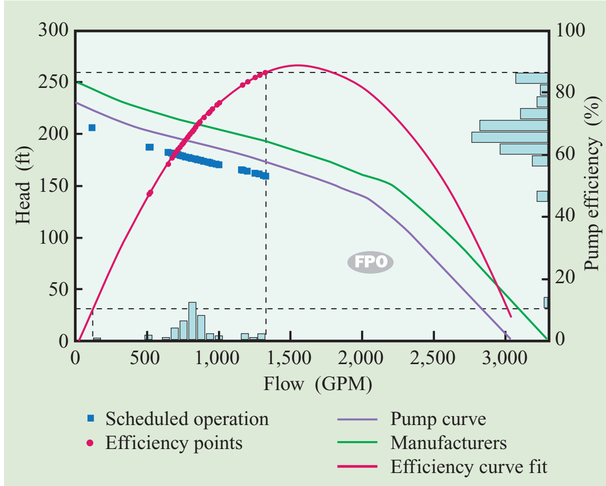

Figure 8

Pump with individual pressure and flow meters allows actual operating points to be shown (screenshot from graphical dashboard).

as 750 GPM are generating high pressures. Since the system curve defines the requirements of the pump, a pump operating most efficiently at 900 GPM and 170 feet of head would be ideal because this matches what we are measuring, although an engineer may prefer a higher-flow pump delivering higher pressures.

The Aquadapt pump-scheduling application is able to make these calculations in real time and for some hours into the future using the hydraulic model. Pumping is scheduled in order to achieve lowest overall cost, simultaneously using both future efficiency information and known time-of-use tariff data in the calculations.

Conclusions

After labor costs, energy costs are typically the highest or second highest controllable cost for water utility operations. Energy use by water utilities is not trivial, with up to 3% of all electrical energy generated in a country going to water or wastewater pumping. Traditional methods for improving pump efficiency have concentrated on the static process of carrying out a pump curve calibration, assuming an operating point, and then either replacing, modifying, machining, polishing, or coating the pump surfaces. While these measures have been shown to be beneficial, they do not take into account the dynamic nature of the actual operating range of a pump.

A small number of tools exist to assist designers of water-supply networks to select appropriate pumps to minimize energy cost, but few, if any, exist to actually automate operations online and in real time. This paper presented one such tool, Aquadapt, which collects, automatically corrects, and stores data from existing instrumentation and telemetry and creates new values from virtual instruments generated within a calibrated hydraulic model. It builds large datasets of mixed real and virtual data that are used in the optimization algorithm but that can also be mined to view operational information that has previously remained hidden from view or was so unreliable in its raw form, that it was considered worthless.

By highlighting pumps that are operating inefficiently, this data is directly beneficial for asset replacement assessment. Merely viewing this data does not in itself allow an operator to select the most efficient pumps in real time. Scheduling pumps to minimize energy cost driven by time-of-use tariff patterns has been attempted by a number of water utilities [19]. However, to achieve cost minimization including efficiency savings, an automatic pump-scheduling process is required that takes into account predicted pressures and flows per pump in real time and uses these predictions in its optimization algorithm to select the most appropriate pumps to operate.

The Aquadapt software tool has been demonstrated to fully automate water distribution scheduling and achieve energy cost savings of 10%–20%, with four major U.S. water utilities currently using the software in a fully automated and online mode (Table 1) Aquadapt has also achieved reductions in energy consumption of 6%–8.4% (Table 2). This has in turn reduced the carbon footprint for these utilities by the same percentages, adding up to thousands of tons of CO_2 per year. If this was applied to the entire U.S. water and wastewater pumping networks, using an average greenhouse gas value of 0.668 short tons of CO_2 per MWh of electricity production [20], between 3 million and 4 million tons of CO_2 emissions could be saved annually.

**Trademark, service mark, or registered trademark of Derceto Limited and Microsoft Corporation in the United States, other countries, or both.

References

1. Electric Power Research Institute, *Energy Audit Manual for Water/Wastewater Facilities*, Electric Power Research Institute, Palo Alto, CA, 1999.
2. D. Denig-Chakroff, National Regulatory Research Institute, "Reducing Electricity Used for Water Production: Questions State Commissions Should Ask Regulated Utilities," June 13, 2008, p. 5; see *http://nrri.org/pubs/water/reducing_electricity_used_for_water_08–06.pdf*.
3. D. Reeves, "Study on Improving the Energy Efficiency of Pumps," European Commission, 2001. AEAT-6559/v, 5.1; see *http://re.jrc.ec.europa.eu/energyefficiency/motorchallenge/pdf/SAVE_PUMPS_Final_Report_June_2003.pdf*.
4. M. Yates, "The Iceberg Effect—Hidden Costs of Pump Operation," *World Pumps* **381**, 54–56 (1998).
5. L. Ormsbee and K. Lansey, "Optimal Control of Water Supply Pumping Systems, *J. Water Resour. Plng. Mgmt*. **120**, No. 2, 237–252 (1994).
6. D. A. Savic, G. A. Walters, and M. Schwab, *Multiobjective Genetic Algorithms for Pump Scheduling in Water Supply*, Springer-Verlag, London, 1997.
7. C. Wegley, M. Eusuff, and K. E. Lansey, "Determining Pump Operations Using Particle Swarm Optimization," *Proceedings of the Joint Conference on Water Resources Engineering and Water Resources Planning and Management,* Minneapolis, MN, 2000.
8. R. S. Powell and G. McCormick, "Derivation of Near-Optimal Pump Schedules for Water Distribution by Simulated Annealing," *J. Operational Res. Soc*. **55**, 728–736 (2004).
9. Z. Wu, P. Boulos, W. Schaetzen, C. Orr, and M. Moore, "Using Genetic Algorithms for Water Distribution System Optimization," *Proceedings of the ASCE Environmental and Water Resources Institute's (EWRI's) World Water & Environmental Resource Congress*, May 20–24, Orlando, FL.
10. K. E. Lansey and K. Awumah, "Optimal Pump Operations Considering Pump Switches," *J. Water Resources Planning Mgmt*. **120**, 17–35 (1994).
11. G. McCormick and R. S. Powell, "Optimal Pump Scheduling in Water Supply Systems with Maximum Demand Charges," *J. Water Resoures Planning Mgmt*. **129**, No. 5, 372–379 (2003).
12. A. Sotelo, C. Lücken, and B. Barán, "Multiobjective Evolutionary Algorithms in Pump Scheduling Optimization," *Proceedings of the Third International Conference on Engineering Computational Technology*, Stirling, Scotland, 2002, pp. 175–176.
13. F. Martinez, V. Hernandez, J. Alonso, Z. Rao, and S. Alvisi, "Optimising the Operation of the Valencia Water-Distribution Network," *J. Hydroinformatics* **9**, No. 1, 65–78 (2007).
14. J. Holland, *Adaptation in Natural and Artificial Systems*, The University of Michigan Press, Ann Arbor, Michigan, 1975.
15. S. Bunn, *Pump Scheduling Optimization in Four US Cities: Case Studies*, American Society of Civil Engineers, Cincinnati, Ohio, 2006.
16. Derceto Ltd., Overview of Aquadapt Software for Water Distribution; see *http://www.derceto.com/derceto-water.html*.
17. G. Klein, "California's Water-Energy Relationship," California Energy Commission, Sacramento, CA, 2005. CEC-700-2005-011-SF.
18. S. E. Ante and J. McGregor, "Giving the Boss the Big Picture," *Business Week* February 13, 2006, pp. 48–50; see *http://www.businessweek.com/magazine/content/06_07/b3971083.htm*.
19. P. Bounds, K. Ulanicka, B. Ulanicki, B. Dacre, and G. Cummings, "Optimal Scheduling of South-Staffordshire Water Supply System Using the FINESSE Package," *Proceedings of the 7th Computing and Control for the Water Industry (CCWI2003) Conference*, Imperial College, London, 2003, pp. 283–292.
20. U.S. Department of Energy, "Updated State-level Greenhouse Gas Emission Coefficients for Electricity Generation," Office of Integrated Analysis and Forecasting, U.S. Department of Energy, 2002.

Received June 2, 2008; accepted for publication June 27, 2008

Simon M. Bunn *Derceto Ltd., L6 Call Plus Centre, Auckland, New Zealand (sbunn@derceto.com).* Mr. Bunn is the Chief Technical Officer of Derceto Ltd. He received his B.E. degree from the University of Canterbury (New Zealand) in 1983. He worked for the Beca Group, an engineering consultancy in New Zealand, from 1990 to 2002, rising to become a partner in the firm in 1999 before starting the spin-off high-technology group, Derceto, producing energy optimization software for water utilities. In 2005, he was named New Zealand Engineering Entrepreneur of the year.

Laurie Reynolds *Mouchel Ltd., West Byfleet, Surrey, UK, KT14 6EZ (laurie.reynolds@mouchel.com).* Mr. Reynolds is a Chartered Control Engineer, and he has spent 30 years with Thames Water, the largest water services company in the United Kingdom. Mr. Reynolds was responsible for the design and planning of the SCADA facilities at the London Water Control Centre, one of the most advanced water control centers. Currently, he is a consultant with Mouchel on advanced control for process and network optimization.