Interfaces

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Relieving Pressure: Optimizing Water Distribution Pressure Management at Valley of the Moon Water District

Segev Wasserkrug, Alexey Tsitkin, Alexander Zadorojniy

Interfaces
Vol. 44, No. 5, September–October 2014, pp. 509–524
ISSN 0092-2102 (print) | ISSN 1526-551X (online)

http://dx.doi.org/10.1287/inte.2014.0761

Relieving Pressure: Optimizing Water Distribution Pressure Management at Valley of the Moon Water District

Segev Wasserkrug, Alexey Tsitkin, Alexander Zadorojniy
IBM Research, Haifa Lab, Haifa, 31905, Israel {segevw@il.ibm.com, alexey.tsitkin@gmail.com, zalex@il.ibm.com}

Efficiently managing pressure in a water distribution network is an issue for water utilities worldwide. Utilities must maintain a delicate balance between lowering the pressure as much as possible to reduce water loss and electricity usage, while keeping pressure high enough to maintain the required level of service. In this paper, we describe how we created and deployed an advanced decision support solution to help the Valley of the Moon Water District (VOMWD) in Sonoma County, California improve its pressure management. Our solution enables VOMWD to efficiently manage water pressure in its network and better handle the pressure changes resulting from seasonal variations in demand. It provides a comprehensive view of the pressure status in the network and incorporates a novel optimization algorithm and problem formulation, which efficiently solve a nonconvex optimization problem and provide recommendations for demand and input pressure changes in the network.

Following the deployment of our solution, VOMWD reduced the number of leaks and bursts by 16 percent compared to the previous year and by 19 percent compared to the average of the previous three years. Less quantifiable results included a reduction in pressure spikes and improvements in tank water levels and water turnover. Our solution has widespread applicability; therefore, we plan to use it to help water utilities worldwide significantly improve their pressure-management capabilities.

Keywords: water; optimization; management; decision support.

Efficient pressure management in water distribution networks is a significant challenge for many water utilities. The pressure in these networks influences many of their key operational aspects, including water loss and electricity usage. In 2005, the World Bank estimated the total cost of water loss to be $14 billion worldwide (Kingdom et al. 2006). Another study estimated that cutting the water losses in Asia in half could supply an additional 150 million people with water (Asian Development Bank 2010). In the United States, moving and treating water and wastewater requires an estimated four percent of the national energy supply, and electricity expenses account for 80 percent of municipal water-processing and distribution costs (Center for Sustainable Systems 2012).

Improved pressure management helps reduce water loss, because lowering pressure reduces both the amount of water lost through leaks and the number of new leaks and bursts. As a rule of thumb, a 10 percent pressure reduction results in a 10 percent reduction of water loss caused by existing leaks and a 14 percent reduction in the number of new leaks and bursts (Lambert 2013). Therefore, pressure management is one of the four main methods that the International Water Association recommends for reducing water loss (Farley and Trow 2003). Energy demand and costs may also be associated with pressure in many water networks because providing the target pressure may require additional pumping. Therefore, the need to reduce energy costs can be an additional incentive for reducing pressure. However, a sufficiently high pressure level must be maintained for various goals, such as ensuring that water reaches all customers at reasonable flow rates, maintaining required levels of pressure for fire hydrants or automatic sprinkler systems to meet firefighting standards, and storing sufficient water in tanks to handle failures in the network or emergencies such as earthquakes.

Therefore, operators of water networks must maintain a delicate balance between lowering pressure as much as possible and keeping it sufficiently high to provide the required level of service. Multiple types of network components (e.g., valves, pumps, tanks) affect pressure in water networks further complicating a utility's ability to maintain this balance. Moreover, the components in water networks are so highly interconnected that a change in the setting of one component could potentially impact the pressure in large segments (and sometimes in all) of the network.

The benefits of pressure management, coupled with the complexity associated with its efficient use, make it an attractive domain for applying analytics and optimization; however, doing so requires overcoming several significant challenges, including capturing the behavior of an actual physical system and providing solutions that water operators—people with little or no experience in advanced decision support systems—can use easily.

In this paper, we describe how we created a pressure-management solution with advanced optimization techniques and algorithms and successfully applied it to improve the pressure management in the Valley of the Moon Water District (VOMWD), despite the challenges mentioned previously. VOMWD, a water utility in Sonoma County, California, faced the constant challenge of balancing these pressure-level trade-offs. The solution was created by a team from the IBM Research Haifa Lab, working closely with VOMWD's general manager and head of operations. As we discuss in this paper, applying our solution significantly enhanced pressure management in VOMWD's water network.

Challenges of Pressure Management in VOMWD

VOMWD (http://www.vomwd.com) provides water to 23,000 customers over a complex water distribution network that is composed of valves, wells, tanks, pumps, and more than 92 miles of pipes. Its network is divided into 10 pressure zones (a pressure zone is a part of the network in which the pressure can be managed more or less independently of the other pressure zone). The largest zone encompasses 85 percent of the network and serves as the backbone for the remainder of the network. Figure 1 shows a schematic of the water network in this pressure zone.

This main pressure zone is located in a valley, and additional pressure zones are located on the mountains situated on the sides of this valley. The network also

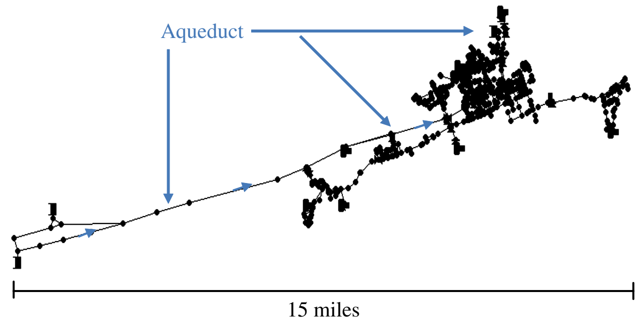

Figure 1: (Color online) This map displays a pipe map of the VOMWD main pressure zone, which is located in a valley, with the aqueduct marked. The arrows on the aqueduct indicate the direction of water flow.

contains several water tanks that are used to ensure that sufficient pressure can be provided to the higher pressure zones and as emergency storage reserves for events such as earthquakes. The Sonoma County Water Agency (SCWA), a regional distributor, provides 80 percent of the water that VOMWD supplies to its customers. SWCA delivers this water through an aqueduct that runs through VOMWD's main pressure zone (see Figure 1). VOMWD receives the remaining 20 percent of its water directly from groundwater wells located within the distribution network.

In managing its water pressure, VOMWD attempts to achieve several, sometimes conflicting, goals:

- Dynamically adapt the pressure in the network to seasonal changes in demand to reduce the number of new leaks and bursts that result from these demand changes.
- Maintain sufficient pressure to provide the required service to customers.
- Maintain the pressure that firefighting regulations require for fire hydrants and automatic sprinkler systems.
- Maintain sufficient pressure to ensure that the required water level in tanks is available for emergency situations (e.g., earthquakes).
- Ensure that pressure is low enough to enable sufficient exchange of water in the tanks to avoid water-quality issues.

VOMWD's main challenge in pressure management was determining how to best adapt the pressure in its main pressure zone to seasonal changes in demand. A change in season between summer and winter causes a significant drop in demand (see Figure 2(a)) and

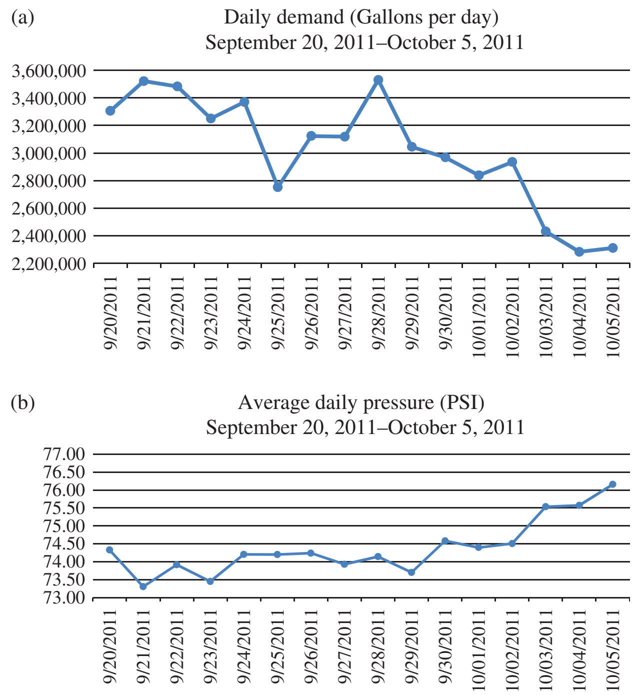

Figure 2: (Color online) These graphs show a drop in demand in the VOMWD network in October 2011 as a result of a transition to winter (2(a)), and a corresponding increase in pressure (2(b)).

a corresponding increase in pressure (see Figure 2(b)). These uncontrolled pressure changes resulted in pressure spikes (i.e., abrupt and significant increases in pressure), as Figure 3 shows; in turn, these spikes caused an increase in the number of bursts in the network's pipes.

VOMWD's primary means of pressure control in its main pressure zone are 10 pressure-reducing valves (PRVs) installed at 10 outlets of the SCWA aqueduct. PRVs facilitate pressure control by enabling a reduction in the pressure of the water flowing through them. However, controlling the pressure through these valves is challenging. All segments of the primary pressure zone are interconnected, and a change in the setting of any single valve could impact the pressure in the entire network. Compounding this difficulty are three factors that influence the pressure in this zone: tanks and wells in the network; seasonal and diurnal changes in demand; and water pressure in the SCWA aqueduct, which in turn depends on factors that VOMWD does not control, such as the water demands of additional cities that depend on the aqueduct.

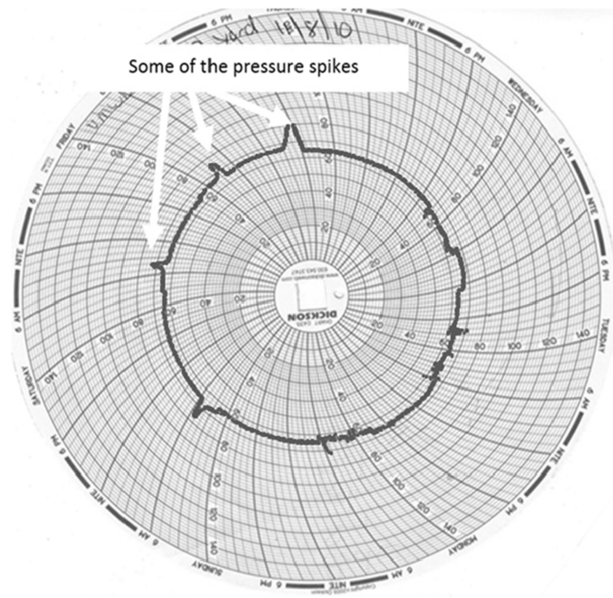

Figure 3: The pressure-reading chart shows pressure spikes during seasonal transitions. Pressure readings are marked by the thick line.

Previous Methodology of Pressure Management in VOMWD

Prior to our involvement, VOMWD's method of managing pressure was reactive, tedious, and based primarily on rules of thumb and operator experience. VOMWD based its pressure-management decisions on the following information:

- Information available through a supervisory control and data acquisition (SCADA) system: SCADA systems are information management systems commonly installed in water networks to gather and display data about the status of the network in near real time. Although these systems sometimes enable remote management of equipment such as pumps or valves, VOMWD's SCADA system did not provide such control for the PRVs. Instead, operators had to manually change the valve settings by physically accessing the valves (see Figure 4). The relevant information that the SCADA system provided about water pressure was the level of the water in the water tanks and several pressure-measurement points:
- An analog pressure logger at the center of the network: This logger recorded the pressure on paper charts, as the example in Figure 3 shows.

Figure 4: This photo shows a VOMWD employee manually changing the PRV setting.

• Monthly usage information on the outlets of the aqueduct: This information was provided by meters that are installed on each outlet of the aqueduct to measure the water flowing through that outlet. Although these meters reported the flow in real time to SCWA's SCADA system, the information was summarized and provided to VOMWD only as monthly totals for billing VOMWD.

Based on this information, VOMWD managed pressure by reacting to three types of events: anomalies in tank-level heights, as indicated by the SCADA system; pressure spikes appearing in the chart logger (see Figure 3); and customer complaints regarding low pressure. If a VOMWD staff member felt that an event warranted a change in valve pressure settings, an operator followed the procedure outlined as follows:

1. Select a single valve that is the most likely to influence the affected area. For example, in the case of pressure spikes at the chart logger, the operator would select the valve for which the most flow came through, based on the monthly billing information received from SCWA.

2. Slightly modify the setting on that single valve, by either raising or lowering the setting by approximately three pounds per square inch (psi).

3. Monitor the situation and repeat the aforementioned process as necessary.

This procedure has two shortcomings. First, as a result of the complexity of the network and the possible interdependencies of the valves, VOMWD operators could not change the settings of more than one valve at a time. Despite their experience and familiarity with the network and its behavior, they were also unable to predict the effects of these changes. This resulted in a rule-of-thumb-based, severely suboptimal pressure-management procedure. Second, VOMWD had no method to detect potential issues in advance, and operators could only react to undesirable events after they occurred. These shortcomings resulted in inefficiencies, such as pressure spikes and large numbers of leaks and bursts during seasonal transitions.

VOMWD partnered with IBM to address these challenges and improve the pressure management in its network, with the goal of reducing the water loss and improving service. This joint effort focused on the operational management of the network pressure. VOMWD needed a tool that its operators could use to (1) understand the current status of pressure in the network, and (2) obtain near real-time recommendations on how to change the settings of the valves to optimize pressure. Providing such a tool required data consolidation, data analytics, and advanced decision support algorithms, for which we used mathematical optimization. In the remainder of this paper, we describe the tool we developed and its application.

Solution Details

As a significant part of our work on the overall solution, we created an algorithm that could provide recommendations on how to optimally set the valves, while considering both the particulars of VOMWD's network and various objectives. Such an optimization algorithm was required to consider both the dependencies of the valve settings and the network-wide effects. We detail our optimization algorithm in the *Optimization Algorithm* section, and we describe our overall system in the *System Overview* section.

Optimization Algorithm

The decision variables in our optimization problem are the settings of the 10 PRVs on the aqueduct. For each valve, these variables specify the maximum pressure of water passing through the valve. Optimizing water networks, including optimizing the settings of equipment such as valves, has been a popular research topic for many years. One of the main challenges of solving optimization problems in this space is that the physical equations, which describe the hydraulic behavior of the water in the water networks, result in nonlinear and nonconvex equality constraints, and therefore in a nonconvex optimization formulation. Specifically, the equation that describes the energy loss on a pipe between two points is a nonlinear equality constraint in the flow, which is in turn influenced by the valve settings—the main decision variables in our problem. Moreover, the form of this equation depends on the direction of flow between these two points, requiring explicit modeling and consideration of the flow direction. In Appendix B, we describe this modeling of energy conservation.

A number of techniques have previously been used to solve this nonconvex optimization problem. The most straightforward approach uses evolutionary or metaheuristic algorithms. Such an approach uses hydraulic simulators, such as EPANET (United States Environmental Protection Agency 2014), which are tools that enable the solving of the hydraulic equations (i.e., deriving the pressure at nodes and the flows through pipes) for a fixed setting of the network equipment (e.g., pumps, valves). In this approach, the evolutionary algorithms are used to generate solutions, which are then tested using the hydraulic simulator. Kapelan et al. (2004) discuss an example of applying this approach in water networks. A major drawback of this approach is that it usually offers no guarantee or bound on the optimality of this solution; moreover, the algorithm's running time is generally too long, precluding its use as an operational tool. Another approach uses mixed-integer nonlinear programming (MINLP)-based optimization that directly captures the nonlinear constraints. Klempous et al. (1997), Bragalli et al. (2012), Gleixner et al. (2012), Burgschweiger et al. (2009), and Eck and Mevissen (2013) provide examples. The drawbacks of the MINLP approach are that existing algorithms cannot usually guarantee a global optimum, and the solvers that can solve this type of formulation are less efficient than mixed-integer linear programming (MILP) solvers.

Because efficient solvers exist for MILP problems, several researchers have used a linearization approach. The most straightforward approach (and also the approach we attempted first) approximates the nonlinear constraints by piecewise linear approximations; however, this approach introduces a large number of binary variables, making it extremely inefficient. Other linearization approaches include iterative approaches based on the first-order Taylor series approximation of the nonlinear equations (Sterling and Bargiela 1984) and coupling linear relaxations, including first-order Taylor approximations, with branch-and-bound techniques (Sherali and Smith 1997). Similarly, our approach uses a first-order Taylor approximation to efficiently solve the problem using a MILP solver; however, it differs in that it integrates a hydraulic simulator and a local-search heuristic in a way that eliminates the need for multiple iterations.

In addition to efficiently solving a nonconvex optimization problem, our algorithm must meet the following criteria:

- Because the valve settings can only be changed manually by personnel who physically access the valves, changes can be made at most every few days, and the preferred frequency is at most once each month. Therefore, the recommendations generated must be robust with regard to variations in customer demand between setting changes and variations in input pressure to the valves (i.e., the pressure on the aqueduct controlled by SCWA, which external factors such as water usage by additional cities influence).
- The objective function in the optimization algorithm must correctly capture the multiple types of pressure-management objectives. Moreover, this objective function must be based on pressure goal specifications provided by a water network operator, typically a person who has little experience in advanced decision support tools.
- Each valve may shut off completely, depending on the pressure in the remainder of the network. This results in discontinuous behavior that the optimization formulation must consider.

We addressed the previous criteria with an innovative problem formulation and optimization algorithm.

Traditionally, two types of objective functions are used in pressure-optimization formulations; the first minimizes overall pressure, while ensuring that minimum pressure levels are met at critical points in the network (Eck and Mevissen 2013); the second minimizes energy usage (Skworcow et al. 2010). In our case, neither option was acceptable, given the varying pressure goals within the network. For example, pressure may need to be lowered to reduce pressure spikes in one part of the network,whereas pressure in another part may need to be increased to drive up the water level in some tanks. Therefore, the objective function had to be able to consider these types of goals. As a result, our objective function formulation is similar to the formulation in Sterling and Bargiela (1984), who defined a set of points such that a pressure target is specified for each point. The optimization then seeks to minimize the sum of absolute differences between the predicted pressure and the desired pressure at these points. Our objective function is a generalization of this objective function, as we describe next.

In evaluating the robustness requirements described previously, consider the following. One physical law that must be maintained in water distribution networks (and, therefore, in the optimized solution) is conservation of mass. According to this law, the difference between the amount of water flowing into and out of a node must equal the demand in this node; however, when demand varies, how to best formulate this constraint is unclear.

To define an appropriate objective function and handle the robustness requirement, we used the following process. First, we defined a scenario-based approach. Our optimization formulation considers a set of scenarios; in each scenario, we define a specific demand and input pressure. The scenarios are representative of the variations in demand and the network's input pressure, both over the course of a day and within a particular month. In conjunction with VOMWD, we defined the demand scenarios to use. To take into account the variation in demand in the different months, our solution dynamically updated the scenarios using historical information for the monthly demand.

Based on these scenarios, we defined the objective function as follows:

• VOMWD defined a set of eight pressure critical points in the network that capture the status of pressure in the various parts of the network.

• For each critical point, VOMWD's operator defined three targets: minimum, maximum, and average desired pressure over all scenarios.

• Because determining the precise pressure targets (i.e., the minimum, maximum, and average targets) for each point is difficult, we then defined the optimization objective as minimizing some distance function between the optimized pressure values at the critical points and the desired targets. This allowed for a degree of fuzziness in the definition of the individual pressure goals.

Another unique feature of our formulation is that it allows the user to influence and change the definition of the distance function used in the objective, and thus achieve a better pressure-management outcome. One way the user can influence this function is by defining the purpose of the pressure target for each critical point. For example, the user can specify that the goal of the pressure setting at a particular critical point is to increase the tank level. In such a case, the distance function is automatically modified so that an optimized average value below the average target will be penalized by the distance function more than an optimized average value higher than the targeted average. Another way the user can influence the distance function is by defining the goals at some critical points as having higher priorities than others (see Figure 5). Appendix A shows the formulation of the optimization problem.

Another issue addressed by our formulation is the need to model the discontinuous behavior of physical equipment. For example, consider a PRV between two network nodes, 1 and 2. At any given time, two possibilities exist:

• Water flows through the valve from node 1 to node 2. In this case, the valve ensures that the pressure at node 2 is not higher than the setting.

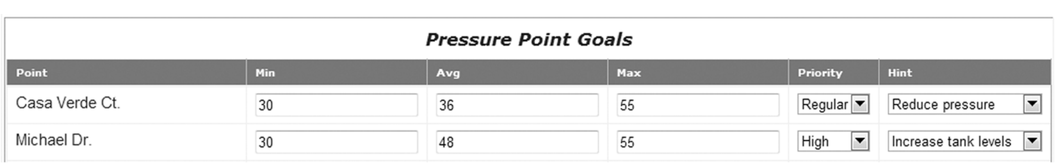

Pressure Point Goals

Point	Min	Avg	Max	Priority	Hint
Casa Verde Ct.	30	36	55	Regular	Reduce pressure
Michael Dr.	30	48	55	High	Increase tank levels

Figure 5: Our approach allows an operator to specify pressure goals in a Web-based user interface.

• Water flows to node 2 from another source in the network. In this case, the pressure at node 2 may be higher than the setting, causing the valve to close. Whenever the valve is closed, no direct connection between nodes 1 and 2 effectively exists in the network.

Our formulation models such discontinuous behavior using a set of integer variables that define a set of if-then rules to capture the desired behavior. To the best of our knowledge, no previous optimization work has addressed this need for detailed valve modeling. On the contrary, previous work assumed that the pressure at the outlet of the valve was less than or equal to the setting (Eck and Mevissen 2013). Our application required detailed modeling; the structure of VOMWD's network is such that without such modeling, the optimization model became infeasible for some scenarios.

The final point our model addressed was the efficient solution of the nonconvex optimization problem. Similar to Sterling and Bargiela (1984) and Sherali and Smith (1997), we linearized the nonlinear constraints using first-order Taylor approximations. However, unlike previous work, we then used this formulation in an algorithm that we created, which combines hydraulic simulation with MILP optimization to efficiently solve the optimization problem without requiring multiple iterations of the MILP solver.

As we describe previously, the main cause of nonlinearities in the optimization model is that the energy conservation equation related to pressure loss on a pipe is nonlinear in the flow within the pipe. To efficiently solve the problem, we linearized these equations using their first-order Taylor approximations. We based our decision to linearize in this manner on the following observations:

• Observation 1: In practice, for a given flow value, the difference between the actual energy loss on a pipe and the linearized version of the energy-loss equation is very small, even for large deviations of the flow (see Figure 6).

• Observation 2: The energy loss is extremely robust relative to flow changes; that is, significant changes in flow within a pipe lead to relatively small energy losses as a result of friction (see Figure 6).

• Observation 3: In some segments of the network, the flow is not influenced by pressure, as Figure 7 illustrates.

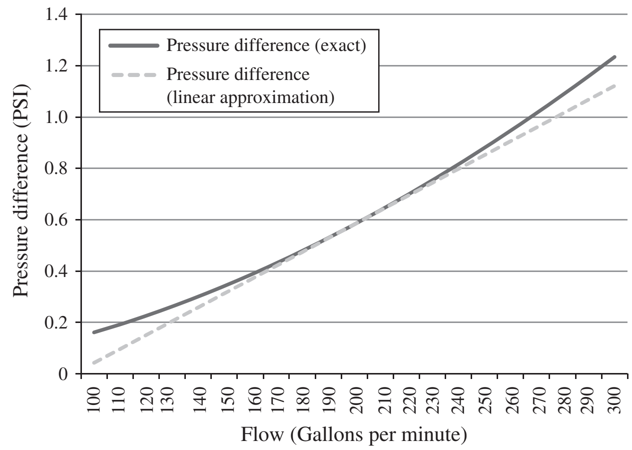

Figure 6: This graph displays variations in pressure reductions (equal to energy loss) plotted according to the actual friction-loss formula (curved line) and the linearized friction-loss formula (straight line). The differences displayed are for flows between 100 gallons per minute (GPM) and 300 GPM on a typical pipe in VOMWD's network. The linearization is obtained by a first-order Taylor approximation at flow value 200 GPM (a typical flow on such pipes).

Observations 1 and 2 imply that given an estimate of the flow in the network, which is sufficiently close to the actual flow in the optimal solution, the linearized model of pressure loss should be a good relaxation of the actual optimization problem. Therefore, we require a good estimate of the flow in the network.

In our algorithm, to obtain such a flow estimate, we use the knowledge that we are optimizing an existing network in which settings for the valve already exist and have been used to operate the network prior to the required change. Based on these settings and using a hydraulic model of the network, we use EPANET (United States Environmental Protection Agency 2014) to calculate initial flow values. We then carry out the Taylor relaxation around these flow values. Note that

Figure 7: In some segments of the network, flow is determined solely by demand. In this example, this occurs because only a single path exists to each node.

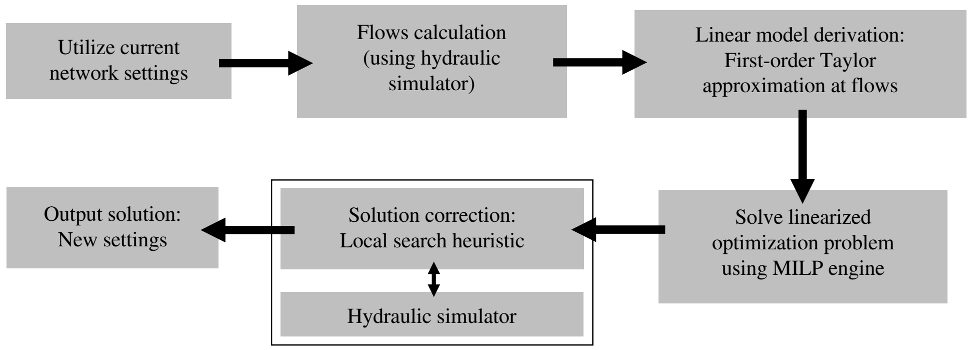

Figure 8: This flowchart is a schematic representation of our optimization algorithm that provides detailed recommendations for valve settings in the network. The algorithm combines hydraulic simulation with mathematical optimization.

we separately calculate such initial flow values for each scenario. Observation 3 then states that because the flow in some parts of the network is determined solely by demand, we have some reason to believe that, at least in parts of the network, this initial flow will correspond to the flow in the optimal solution of the original optimization problem. Therefore, in these parts of the network, the Taylor approximation of friction loss will correspond to the actual friction loss.

After computing the initial flow values using EPANET and deriving the linearized Taylor approximations, we solve the linearized optimization problem using a MILP solver. As a final step, to compensate somewhat for the fact that our problem was solved by a linearized version of the original optimization problem, we implement a local search heuristic that uses the EPANET hydraulic simulator as a black box to attempt to find the local optimum of the actual optimization problem around the solution provided by the linearized version. The heuristic evaluates and compares possible solutions by applying them to the hydraulic simulator and then computes the objective function based on the simulated results. This heuristic is basically a variation of performing an exhaustive local search on the possible valve settings while maintaining computational feasibility. We implement this by combining the following techniques (depending on the actual problem size): making slight changes in valve settings, discretizing the intervals of possible settings, and filtering out some of the valves that are less influential to the objective function. The heuristic then iteratively decreases the search space until no better solution is found locally. We can summarize the overall approach as follows: a heuristic approach attempts to use the linearized optimization problem to obtain a solution in the neighborhood of the global optimal solution for the original problem, and then uses the local search to improve the solution and ensure feasibility of the solution for the original nonconvex optimization problem.

We summarize the steps of our algorithm here and in Figure 8, as follows:

1. Input the current valve settings into the hydraulic simulation model (with other initial conditions, such as tank levels).
2. Use the hydraulic simulator to calculate the flows through the network with the current valve settings in the network.
3. Create a linearized version of the optimization problem by using a first-order Taylor approximation around the flows calculated in step 2 (Appendix B provides more details about the exact form of this linearization for pipes).
4. Using a MILP solver (e.g., CPLEX), solve the optimization problem with the linearized pipe constraints.
5. Perform a local search around the solution provided by the MILP solver using a heuristic that employs the hydraulic simulation model, which contains the nonlinear energy-conservation constraints, as a black box.
6. Output the best solution obtained by the local search.

To validate our algorithm in a laboratory setting, we compared our algorithm to the MINLP implementation

Network	No. of pipes	No. of valves	No. of hours	Computation time (seconds)		Objective value	
				MINLP	Our algorithm	MINLP	Our algorithm
SB25 (benchmark)	36	3	1	1 s	1 s	576	568
SB25 (benchmark)	36	3	24	78 s	15 s	13,991	13,991
Customer network #1	128	3	1	1 s	1 s	2,273	2,272
ExNet (benchmark)	2,465	3	1	200 s	8 s	63,907	63,104
VOMWD	789	10	1	303 s	54 s	41,579	41,684
VOMWD	789	10	24	—	1,170	—	945,874

Table 1: A performance comparison between our algorithm and the MINLP algorithm shows that our algorithm performs significantly faster than the MINLP with equivalent (and usually slightly better) objective function values.

described in Eck and Mevissen (2013) using a simplified version of the objective function. Table 1 shows the results of this comparison. In terms of objective function values achieved, the results are very close and the running times of our algorithm are significantly shorter for larger problem sizes. In one instance, the MINLP implementation could not find a feasible solution, because it did not capture the discontinuous behavior of the valves.

The MINLP algorithm did not serve only as a basis for comparison; we also used it to find optimality bounds for some networks based on semidefinite programming (SDP) relaxations. In all cases in which we used such relaxations, the optimality gaps were less than seven percent; in many cases, they were less than 0.3 percent, thus serving as additional evidence of the quality of solutions provided by our algorithm.

System Overview

Although the optimization algorithm is a major part of our solution, providing a useful operational system to VOMWD required additional components, as we describe next.

- Integration of various data sources: A VOMWD operator needs to know the current pressure status in the network on an ongoing basis. Prior to our work, VOMWD did not have the pressure sensors in place to provide such knowledge; no pressure measurements were available for the majority of the pressure critical points. Therefore, VOMWD installed pressure sensors at all defined critical points as part of implementing our solution. We had to capture the information provided by these loggers and integrate it with a variety of additional information, including (1) hourly information from the SCWA about the amounts of water input into the network (previously only available to VOMWD as a monthly billing summary), (2) VOMWD SCADA information about tank levels and pressure at the pre-existing measurement points, and (3) geo-locations of existing pressure reducing valves in the network and their current status and settings.
- Alerting and trending capabilities: In conjunction with VOMWD, we used the consolidated data to define key performance indicators (KPIs) that were constantly monitored. Examples of such KPIs include the actual minimum, maximum, and average pressure for the current hour for each pressure critical point and the height of water in the tanks. Our system enabled VOMWD to see the current status of these KPIs on an ongoing basis. In addition, we defined several threshold levels for each KPI. Alerts are automatically generated whenever these thresholds are crossed. Moreover, the system enables the operator to view historical and trending information for each KPI. Thus, the user can know the current network state in terms of pressure and become aware of problematic situations significantly earlier (see Figure 9).
- Optimization: The optimization algorithm detailed in the *Optimization Algorithm* section is integrated into the larger system. The operator can enter the goals for each critical point, run the optimization, and obtain detailed operational recommendations on how to set the valves to optimize the pressure in the network based on the goals entered. Moreover, the operator can change the settings at some valves and see the predicted pressures at each critical point. This feature is useful for performing what-if analyses.

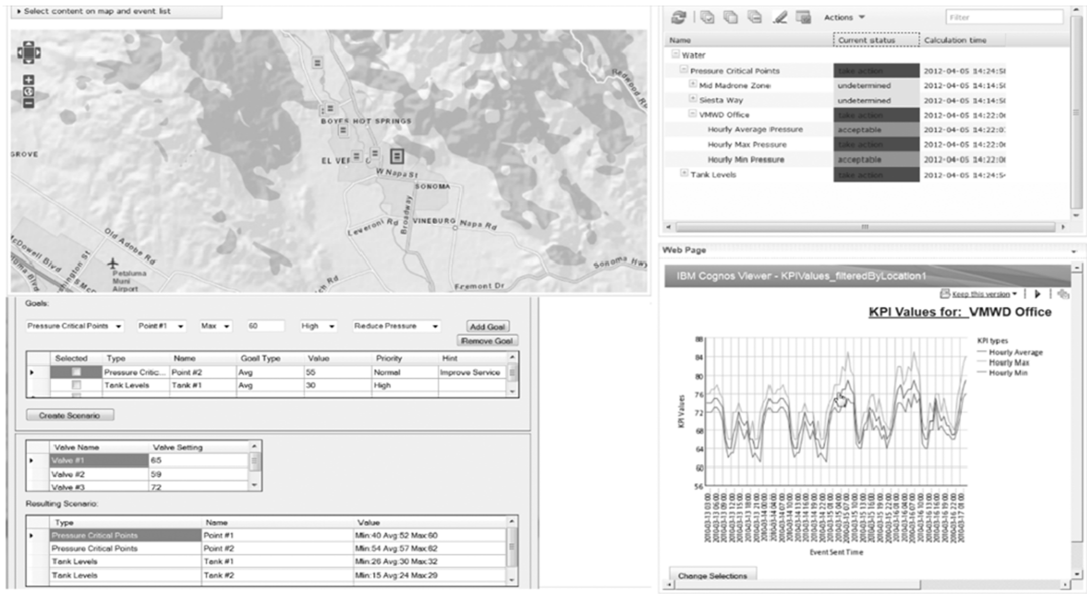

Figure 9: Operators can access our pressure-management system via a Web-based interface, allowing them to change settings and examine predicted pressure values from any location.

Implementation

The implementation of the pressure-management system at VOMWD required significantly more effort than merely developing and deploying the solution described in the *Solution Details* section.

First, digital pressure loggers were installed at the six critical pressure points at which no pressure measurements capabilities were previously in place, and their data were integrated into the system. Calibrating the simulation model is another important step in our work. As we describe in the *Optimization Algorithm* section, our solution requires a hydraulic model of the network in a format that can be solved using a hydraulic simulator. When we began our work, we found that VOMWD had such a model; however, it had not been updated for several years. Moreover, its main purpose was to help VOMWD make network design decisions, whereas we required a model suitable for operational decisions. Therefore, one of our major initial efforts was to update the model and attempt to calibrate it for operational decision making. Our simulation model was much more accurate than the previous one; however, despite our efforts, this simulation model did not completely capture the network's behavior. For example, as the demand changed during the day and week, it did not fully capture the pressure variations at pressure critical points. To address this discrepancy between the simulation model and the actual network, we defined a special usage process based on the following concept: Instead of treating the pressure goals as absolute pressure values, we should treat them as goals relative to the current state. For example, if an operator's goal is to increase the average pressure at some critical point by 2 psi, that operator would carry out the following process (see Figure 10):

1. Run the hydraulic simulation with the current setting of the valves to calculate the current average pressure at this critical point based on the simulation.
2. Increase the value obtained by the required amount (i.e., 2 psi).
3. Set this updated value as the average goal for this point for the optimization.

In summary, our solution's deployment comprises the following main steps:

1. Define pressure critical points.
2. Deploy pressure loggers at critical points at which no pressure loggers existed previously.
3. Update and calibrate the simulation model.
4. Define the new usage process.

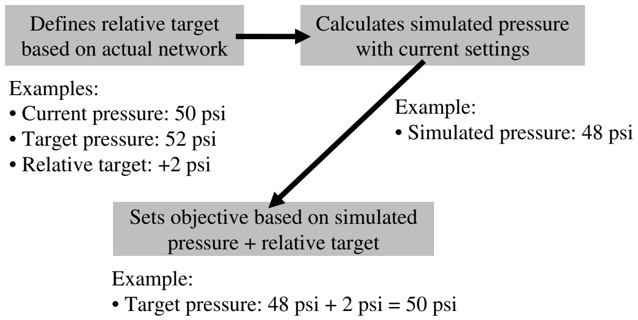

Figure 10: In this example, our optimization process addresses the discrepancy between the simulation model and the actual network.

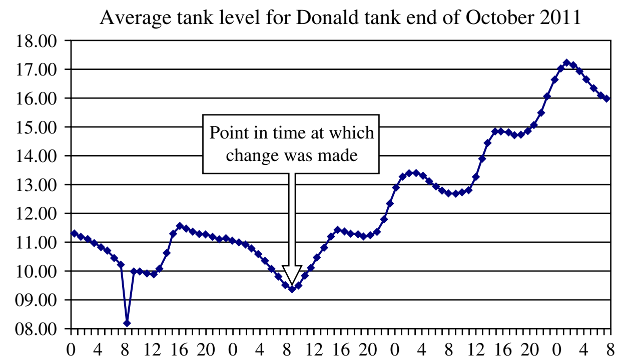

Figure 11: (Color online) We addressed a decrease in tank water level following a well failure by applying our optimization algorithm. The implementation of the recommended changes in valve settings led to an almost immediate increase in the tank's water level.

5. Formulate and implement the optimization algorithm.
6. Integrate the various solution components (e.g., data integration, alerting and trending, optimization) into a single Web-based solution.
7. Implement and deploy the solution to generate PRV-setting recommendations.

Initial Implementation

We began our work in June 2011. Our first target was to help VOMWD with the seasonal transition from summer to winter in October 2011. At the beginning of October, using the consolidated data created as part of our solution, we noted the demand reductions as a result of the seasonal transition (see Figure 2(a)) and the corresponding pressure increases (see Figure 2(b)). Based on these data, in conjunction with VOMWD, we decided to change the valve settings during the last week of October.

Based on our optimization algorithm, we established recommendations for new valve settings and were planning to present them to VOMWD. However, when we arrived at VOMWD, one of its wells was malfunctioning, resulting in a drop in water level at the Donald water tank. By adapting our objective function definition, we were able to help VOMWD address this failure. After rerunning our optimization with the adapted objective function definition, we came up with a recommendation that involved simultaneous updates to four of the valves and raised the pressure near the Donald tank, while attempting to maintain a lower pressure at other critical points. VOMWD personnel implemented these changes—marking the first time that multiple valves were updated simultaneously. These changes achieved the new primary goal of compensating for the well failure and significantly raised the water level in the Donald tank (see Figure 11).

Encouraged by these results, VOMWD operators decided to make additional changes to improve the exchange of water in another water tank—by reducing the pressure in the vicinity of this tank—and to reduce pressure in other areas of the network, without impacting the water level in the Donald tank. They achieved the new goals, resulting in overall improved pressure.

Overall Implementation Process

Since October 2011, when VOMWD first used the optimization algorithm to derive the valve setting changes, the company has used our system's recommendations to make all changes to valve settings. In June 2012, we helped it use our system to adapt to the transition from winter to summer. In October 2012, VOMWD operators made the change to network pressure without any support from us, thus attesting to the system's intuitive nature and its suitability for water network operators.

Solution Benefits

VOMWD's primary goal in implementing the pressure-management system was to reduce the number of leaks (some of which may have been large bursts) resulting from seasonal changes. Therefore, the main parameter we used to measure the results was the number of leaks and bursts as compared to previous years.

	Nov	Dec	Jan	Feb	Mar	Apr	May	June	July	Aug	Sep	Oct	Total
2010–2011	5	3	8	1	7	4	3	6	6	8	5	11	67
2011–2012	2	2	6	4	1	10	3	3	8	8	4	5	56

Table 2: One year after deployment of our optimization algorithm, the number of leaks decreased significantly.

To ensure that the results were reliable, we measured the number of leaks and bursts for a full year after deployment—from the end of October 2011 until the end of October 2012. When we compared the number of leaks and bursts for the year after deployment (November 2011 to October 2012) to the number in the same period during the previous year (see Table 2), we found that the leaks decreased by more than 16 percent—56 leaks compared to 67 previously. In addition, they decreased by more than 19 percent compared to the previous three-year average of 69.4 leaks.

This reduction in leaks translates to cost savings both as a result of less wasted water and not having to fix 16 percent more leaks. If this reduction in leaks were to correspond to the expected 19 percent reduction in water loss, the result would be 17,639,104 gallons of water saved per year and a potential two percent reduction in overall annual operating expenditures. For VOMWD, this is approximately $60,000; for larger water utilities, the potential savings are significantly greater, because the amount of money represented by two percent of the overall operating expenditure is much higher.

Doing a month-by-month comparison of the number of leaks and bursts (see Table 2), seven of the 12 months show a reduction in the number of leaks; in the other five months, the number of leaks is the same for two months. The only months after the optimization for which the number of leaks is greater are February, which shows only one leak in the previous year, July, with eight leaks compared to six, and April, with 10 leaks compared to four in the previous year. We were unable to find any reason for the unusually large number of leaks in April; however, the large number of leaks in July may indicate that when VOMWD operators made the transition from winter to summer, they might have overcompensated for the increase in demand by specifying pressure goals that resulted in an exceedingly high pressure increase. Therefore, this seasonal transition should be a focus of improvement for VOMWD in the coming years. The month-to-month variation in improvement also emphasizes the importance of measuring the improvement over a sufficiently long period.

For VOMWD, our work should result in additional savings beyond those resulting from reducing the number of leaks and bursts. Such savings include the reduction of water lost through existing leaks and reductions in property damage costs caused by large bursts. According to VOMWD staff, one such burst in August 2009 caused $20,000 in damages.

Our work also resulted in several benefits that are more difficult to quantify. The first benefit is a reduction in the number of pressure spikes; in Figure 12, the pressure line is flat with no sharp spikes. This is in contrast to Figure 3, which shows the pressure values for November 2010 (prior to our work), in which the large pressure spikes are readily apparent. Additional benefits attested to by VOMWD are improvements in

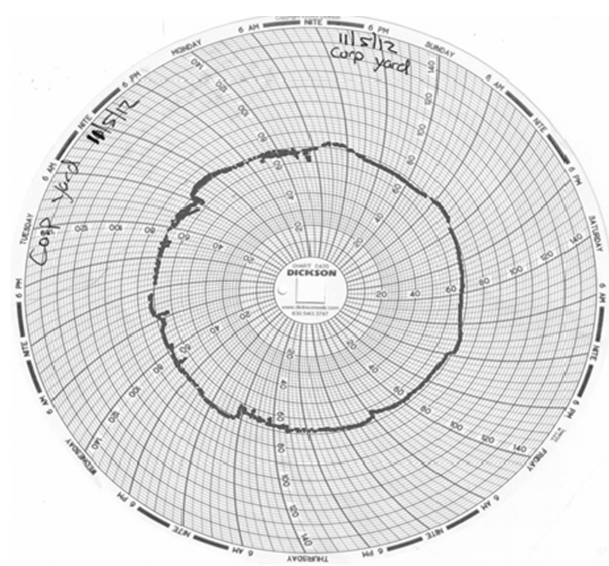

Figure 12: After the deployment of our solution, this pressure chart logger for November 2012 showed steadier pressure with no pressure spikes.

levels of service resulting from improvements in tank water levels and tank turnover (i.e., reduction in the amount of time that water remains in the tank—lengthy residence times may lead to stagnation), and the ability to change from a fully reactive to a proactive process.

We also observed that the system gave added confidence to VOMWD staff members; they had the ability and confidence to make changes much more frequently than they had in the past. VOMWD previously tried to limit changes to twice a year—in the fall and spring transitions between summer and winter. Company management now feels comfortable that operators can make changes more frequently—monthly or even once every few weeks. This ability to change pressure settings more often enables VOMWD to better align the network pressure with the actual demand for water.

Finally, our solution resulted in very high satisfaction and a new appreciation on the part of VOMWD staff and management for the capabilities of operations research (OR). The primary user of our system, VOMWD's head of operations, told us he is extremely happy with the system's recommendations. Despite his 30 years of experience, he said he would never have thought of making the changes the system recommended. He also explained that although he was initially skeptical about applying analytics and OR tools to address the problem, he is now a very strong supporter of such techniques.

Lessons Learned

As a result of our work, VOMWD has a new and efficient process and tool for managing water pressure. Moreover, its staff can monitor the status of pressure in the network and understands the benefits of efficient pressure management; OR, specifically optimization, was vital to this success. The application of an optimization algorithm that enabled the efficient solution of a complex, nonconvex optimization problem achieved a significant reduction in leaks and bursts.

Although the optimization algorithm was a vital part of the project's success, the application would not have been successful without the creation of an end-to-end solution that included deploying additional sensors, data consolidation, a practical usage methodology, and an objective function formulation and user interface tailored to the needs of the end user. Another important lesson we learned is that when applying OR to physical systems, many challenges must be overcome. These include dealing with limited precision, frequently erroneous sensors, and the need to address discrepancies between the physical systems being optimized and the models available for their optimization. We believe that addressing these discrepancies and dealing with data that originate from physical sensors are interesting topics for future academic research.

To efficiently solve the resulting optimization problem, we created and implemented an optimization algorithm that combines hydraulic simulation, a MILP optimization algorithm, and a local-search heuristic in a novel way. This may be a specific example of a much more general approach, based on combining a linearization around an existing operating point, a method for efficiently estimating the effect of a specific operating point on the physical system, and a local heuristic. Therefore, possible topics for future research include evaluating whether this approach can be generalized and formalized, and whether it can be applied to significantly different domains, such as an electrical power grid, in which an existing system must be optimized and for which suboptimal settings are already in place.

Our solution is not specific to VOMWD. We believe it can provide significant benefits to many water networks worldwide. For many utilities, the expected monetary benefits would be significantly larger than for VOMWD. This is not only because of the higher expenditures of larger utilities, but also because of the cost of water loss; for example, in one city on the West Coast of the United States, the estimated amount of water loss is approximately seven percent, and the cost of one percent of water lost is almost \$3 million. For such a city, a 19 percent reduction in water loss represents a savings of almost \$4.2 million a year, without considering benefits related to reductions in pipe repairs or property damage. In addition, improvements in pressure management would reduce the risk of large bursts in busy metropolitan areas. One such burst could result in dramatic disruptions and huge monetary costs. Finally, in many cities in which pumps provide water to consumers, improved pressure management could result in reductions in energy use and its related costs. We have had ongoing discussions with several utilities in the United States, Europe, Africa, and Asia, and hope to use our system and approach to improve the water-pressure management for many of these utilities.

Acknowledgments
The authors thank the dedicated staff of Valley of the Moon Water District (VOMWD) and the Sonoma County Water Agency (SCWA). We would specifically like to thank Krishna Kumar, Paul Gradolph, Shari Walk, and Matthew Fullner from VOMWD, and Jay Jasperse, Don Seymour, and Todd Schram from SCWA for their outstanding collaboration throughout this project. This project could not have been successful without their significant efforts, openness, continuous feedback, and their willingness to innovate.

Appendix A. Optimization Model
In this appendix, we describe a somewhat simplified version of our formulation of the optimization problem. We use the following notation:

- $G=(V,E)$ is a graph representing the water network. V and E are the set of nodes and pipes respectively.
- S is a set of scenarios that represent the variations in input pressure and demand that the optimal solution must consider.
- $r_{(i,j)}$ is a pipe-specific resistance coefficient (see Appendix B).
- e_i is the elevation of node i (i.e., the height above some reference elevation).
- $Q'_{(i,j),s}$ is the initial estimate of flow on pipe (i,j) in scenario $s \in S$.

Decision variables:

- $v_{(i,j)} \in \{\mathbb{R}^+ \cup 0\}$ are the settings of the valves ($v_{(i,j)}$ is the setting of valve v between node i and node j).
- $closed_{(i,j),s}, open_{(i,j),s} \in \{0,1\}$ are binary variables indicating whether valve (i,j) is closed or open, respectively, in scenario $s \in S$.
- $Q_{(i,j),s}$ is the flow on pipe i,j in scenario $s \in S$.
- $p_{i,s} \in \{\mathbb{R}^+ \cup 0\}$ is the pressure in node i in scenario s.
- $H_{i,s} \in \{\mathbb{R}^+ \cup 0\}$ is the total energy (also known as head). $H_{i,s}=(1/\gamma)p_{j,s}+e_i$ where γ is a unit-specific conversion constant between pressure and elevation.
- $H_s^{pu} \in \{\mathbb{R}^+ \cup 0\}$ is the head (energy) added by pump pu in scenario s.

Constraints:

- Conservation of mass:

$$\forall i, \forall s, \quad \sum_j Q_{(j,i),s} - \sum_k Q_{(i,k),s} = D_{i,s}. \tag{A1}$$

- Linear approximation of conservation of energy in pipes (see Appendix B):

$$\begin{aligned}&\forall \text{pipes}(i,j), \forall s \in S,\\ &\left(\frac{1}{\gamma}p_{i,s}+e_i\right)-\left(\frac{1}{\gamma}p_{j,s}+e_j\right)\\ &\quad = r_{(i_1,i_2)}[\text{sign}(Q'_{(i_1,i_2),s})(|Q'_{(i_1,i_2),s}|)^{1.852}\\ &\qquad +1.852(|Q'_{(i_1,i_2),s}|)^{0.852}(Q_{(i_1,i_2),s}-Q'_{(i_1,i_2),s})].\end{aligned} \tag{A2}$$

- Pressure-reducing valve (PRV) modeling: These constraints capture the possibly discontinuous behavior of the valves.

$$\forall \text{valves } (i,j), \forall s, \quad open_{v,s}+closed_{v,s}=1. \tag{A3}$$

$$\forall \text{valves } (i,j), \forall s, \quad closed_{(i,j),s} \to (Q_{(i,j),s}=0). \tag{A4}$$

$$\begin{gathered}\forall \text{valves } (i,j), \forall s \in S,\\ (p_{j,s} > v_{i,j}) \to (closed_{(i,j),s}=1).\end{gathered} \tag{A5}$$

$$\begin{gathered}\forall \text{valves } (i,j), \forall s \in S,\\ ((p_{j,s} \le v_{i,j}) \wedge (H_{i,s} \ge H_{j,s})) \to (open_{(i,j)}=1).\end{gathered} \tag{A6}$$

$$\begin{gathered}\forall \text{valves } (i,j), \forall s,\\ \left((open_{(i,j),s}=1)\wedge\left(H_{i,s} \ge e_j+\frac{1}{\gamma}v_{(i,j)}\right)\right)\\ \to (p_{j,s}=v_{(i,j)}).\end{gathered} \tag{A7}$$

$$\begin{gathered}\forall \text{valves } (i,j), \forall s,\\ \left((open_{(i,j),s}=1)\wedge\left(H_{i,s} \le e_j+\frac{1}{\gamma}v_{(i,j)}\right)\right)\\ \to (H_{j,s}=H_{i,s}).\end{gathered} \tag{A8}$$

- Pumps: A linearization (similar to valves) of the energy a pump adds to the network (sh, c, and e are pump-specific constants):

$$\begin{aligned}\forall pu_{(i,j)}, \forall s \in S, \quad H_s^{pu_{(i,j)}} &= sh - c[(Q'_{(i,j),s})^e + e(Q'_{(i,j),s})^{e-1}\\ &\quad \cdot (Q^{pu}_{(i,j),s}-Q'_{(i,j),s})].\end{aligned} \tag{A9}$$

Our objective function is based on the deviation of the pressure value at each critical point. It uses the following additional inputs:

- A set of critical points C: This is the set of points whose pressure represents the pressure in the entire zone.
- tmn_c, tmx_c, tav_c: These are, for each critical point $c \in C$, the target minimum, maximum, and average values across the scenarios $s \in S$.
- $w_s \in \mathbb{R}^+$: This is the relative importance of scenario s with regard to the pressure targets.
- $wdevmx_c^+$, $wdevmn_c^+$, $wdevav_c^+$: This is the relative importance of each type of deviation above the respective pressure target.
- $wdevmx_c^-$, $wdevmn_c^-$, $wdevav_c^-$: This is the relative importance of each type of deviation below the respective pressure target.

To calculate the objective function, we define the following additional decision variables:

- mn_c, mx_c, $av_c \in \mathbb{R}$: The minimum, maximum, and average values of pressure over all the scenarios in S (i.e., $mn_c = \min_{s\in S} p_{c,s}$, $mx_c=\max_{s\in S} p_{c,s}$, $av_c = (\sum_{s\in S} w_s p_{c,s})/(\sum_{s \in S} w_s)$).
- $devmn_c^+$, $devmx_c^+$, $devav_c^+ \in \{\mathbb{R}^+ \cup 0\}$: The respective deviations of the critical point pressure values above the targets.

- $devmn_c^-$, $devmx_c^-$, $devav_c^- \in \mathbb{R}^+ \cup \{0\}$: The respective deviations of the critical point pressure values below the targets.

Based on this, we have the following:

$$\forall c \in C, \quad mn_c - tmn_c + devmn_c^+ - devmn_c^- = 0, \tag{A10}$$

$$\forall c \in C, \quad mx_c - tmx_c + devmx_c^+ - devmx_c^- = 0, \tag{A11}$$

$$\forall c \in C, \quad mav_c - tav_c + devav_c^+ - devav_c^- = 0, \tag{A12}$$

and the objective is then to minimize the maximum weighted deviation from any target:

$$\begin{aligned} &\text{minimize } \max_{c\in C} \max(wdevmn_c^+ \cdot devmn_c^+, \\ &wdevmx_c^+ \cdot devmx_c^+, wdevav_c^+ \cdot devav_c^+, \\ &wdevmn_c^- \cdot devmn_c^-, wdevmx_c^- \cdot devmx_c^-, \\ &wdevav_c^- \cdot devav_c^-). \end{aligned} \tag{A13}$$

The respective deviation weights are derived from the user preferences regarding the priority and target of the goal at each critical point.

Appendix B. Linearization of the Energy-Conservation Equations in Hydraulic Networks

The general form of the energy loss caused by friction on a pipe p between two nodes i, j is given by Equation (B1) (Burgschweiger et al. 2009), where H_{i_1}, H_{i_2} is the total energy (also called head) at points i_1, i_2, respectively, $Q_{(i_1,i_2)}$ is the flow on the pipe between point i_2 and i_1, and $r_{(i_1,i_2)}(Q_{(i_1,i_2)})$ is a pipe-specific, flow-dependent resistance coefficient.

$$H_{i_2} - H_{i_1} = r_{(i_1,i_2)}(Q_{(i_1,i_2)})Q_{(i_1,i_2)}|Q_{(i_1,i_2)}|. \tag{B1}$$

Several simplified empirical forms of this formula exist. One such form, which is in common use in the United States, is the Hazen Williams equation (Haested Methods et al. 2003). The Hazen Williams equation can be written as:

$$\left(\frac{1}{\gamma}p_{i_1} + e_{i_1}\right) - \left(\frac{1}{\gamma}p_{i_2} + e_{i_2}\right) = r_{(i_1,i_2)}Q^{1.852}_{(i_1,i_2)}. \tag{B2}$$

The differences between Equation (B2) and Equation (B1) are the following:

- $r_{(i_1,i_2)}$ is a pipe-specific constant (no longer dependent on flow).
- We assume that the flow is now only positive, and flows from i_1 (point with higher energy) to i_2 (point with lower energy).

Because we use Equation (B2) in our model, we have to address two issues: the nonlinearity of the constraint and the directionality of $Q_{(i_1,i_2)}$. To address the nonlinearity of the constraint, we assume (as described in the *Optimization Algorithm* section) that we have a good approximation $Q'_{(i_1,i_2)}$ of the flow of the pipe. To consider the directionality of flow, we then further allow the flow to be negative when the flow is from i_2 to i_1, and modify the first-order Taylor approximation to obtain the following linearized constraint:

$$\begin{aligned} &\left(\frac{1}{\gamma}p_{i_1} + e_{i_1}\right) - \left(\frac{1}{\gamma}p_{i_2} + e_{i_2}\right) \\ &= r_{(i_1,i_2)}[\text{sign}(Q'_{(i_1,i_2)})(|Q'_{(i_1,i_2)}|)^{1.852} \\ &\qquad + 1.852(|Q'_{(i_1,i_2)}|)^{0.852}(Q_{(i_1,i_2)} - Q'_{(i_1,i_2)})]. \end{aligned} \tag{B3}$$

References

Asian Development Bank (2010) The issues and challenges of reducing non-revenue water. Accessed March 1, 2013, http://www.miya-water.com/user_files/Data_and_Research/miyas_experts_articles/2_NRW/92.pdf.

Bragalli C, D'Ambrosio C, Lee J, Lodi A, Toth P (2012) On the optimal design of water distribution networks: A practical MILNP approach. *Optim. Engrg.* 13(2):219–246.

Burgschweiger J, Gnädig B, Steinbach M (2009) Optimization models for operative planning in drinking water networks. *Optim. Engrg.* 10(1):43–73.

Center for Sustainable Systems (2012) U.S. Water Supply and Distribution Fact Sheet. Accessed March 12, 2014, http://css.snre.umich.edu/css_doc/CSS05-17.pdf.

Eck B, Mevissen M (2013) Fast non-linear optimization for design problems on water networks. *Proc. 2013 World Environ. Water Resources Congress* (American Society of Civil Engineers, Reston, VA).

Farley M, Trow S (2003) *Losses in Water Distribution Networks: A Practitioner's Guide to Assessment, Monitoring and Control* (IWA Publishing, London).

Gleixner AM, Held H, Huang W, Vigerske S (2012) Towards globally optimal operation of water supply networks. *Numer. Algebra Control Optim.* 4(2):695–711.

Haested Methods, Walski TM, Chase DV, Savic DA, Grayman WM, Beckwith S, Koelle E (2003) *Advanced Water Distribution Modeling and Management* (Haested Press, Waterbury, CT).

Kapelan Z, Savic D, Walters G (2004) Incorporation of prior information on parameters in inverse transient analysis for leak detection and roughness calibration. *Urban Water J.* 2(1):129–143.

Kingdom B, Liemberger R, Marin P (2006) The challenge of reducing non-revenue water (NRW) in developing countries. How the private sector can help: A look at performance-based service contracting. Report, The World Bank, Washington, DC.

Klempous R, Kotowski J, Nikodem J, Ułasiewicz J (1997) Optimization algorithms of operative control in water distribution systems. *J. Comput. Appl. Math.* 84(1):81–99.

Lambert A (2013) Leakage reductions: The fundamental role of pressure management. *Water World Internat.* 28(1):17–21.

Sherali H, Smith E (1997) Global optimization approach to a water distribution network design problem. *J. Global Optim.* 2(11):107–132.

Skworcow P, Ulanicki B, AbdelMeguid H, Paluszczyszyn D (2010) Model predictive control for energy and leakage management in water distribution systems. Accessed March 12, http://watersoftware.dmu.ac.uk/uploads/publications/paperv6.pdf.

Sterling M, Bargiela A (1984) Leakage reduction by optimised control of valves in water networks. *Trans. Inst. Measurement Control* 6(6):293–298.

United States Environmental Protection Agency (2014) EPANET. Accessed February 28, http://www.epa.gov/nrmrl/wswrd/dw/epanet.html.

Segev Wasserkrug is the technical leader of the Business Optimization Group at the IBM Haifa research lab. Segev has over 12 years practical experience in applying advanced optimization and analytical techniques to customer problems in a variety of domains including workforce management and logistics. Over the past several years, Segev has led and participated in work on applying advanced analytical and optimization techniques to improving the operations and management of water distribution networks in Europe, Asia and the U.S. This work has included working closely with water operators and managers in defining optimization solution requirements, algorithm development, and solution implementation in the customer's environment. At present, he is leading work on optimizing the operations of wastewater treatment plants. Segev received his PhD in information systems management, and his MSc and BA in computer science from the Technion—Israel Institute of Technology. Segev has numerous academic publications and patents.

Alexey Tsitkin is a software architect with over 10 years of experience in the software industry. Alexey earned his BA with honors, and studied for an MSc at the computer science faculty, in the Technion—Israel Institute of Technology, where he also served as a teaching assistant in several courses. Alexey led the software architecture and development of the aforementioned water optimization solutions while working at the IBM Haifa Research Lab. Alexey currently works at the Kenshoo Company, focusing on optimization platforms for online advertising.

Alexander Zadorojniy is a research scientist in the business optimization group at IBM Research—Haifa. Alexander holds a PhD from the Tel-Aviv University and a BSc and MSc from the Technion—Israel Institute of Technology, all in electrical engineering. His research involved development of efficient optimization algorithms for discrete time stochastic control problems, and solving real-life design and management problems. Alexander's awards for scientific publications include INFORMS Wagner finalist, GYSS finalist, CSER first runner-up theoretical award, INFORMS Nicholson finalist, and ORSIS winner. Previous to IBM, Alexander spent 10 years at Intel IDC working on hardware design, where, in addition to leading unit designs and designing complex clock networks, he specialized in verification and validation of leading edge Intel processors. At present he is working on development of optimization platforms and algorithms for water network management and the optimization of waste water treatment plants operations.