Proceedings of the 2020 Winter Simulation Conference
K.-H. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, eds.

MULTI-THREADED SIMULATION OPTIMIZATION PLATFORM FOR REDUCING ENERGY USE IN LARGE-SCALE WATER DISTRIBUTION NETWORKS WITH HIGH DIMENSIONS

David Linz
Behzad Ahmadi
Robert Good

Center for Water-Energy Efficiency
University of California, Davis
215 Sage St., Suite 200
Davis CA, USA

Erin Musabandesu
Frank Loge

Civil and Environmental Engineering
University of California, Davis
215 Sage St., Suite 200
Davis CA, USA

ABSTRACT

Hydraulic simulation models are used to improve the energy use of pumps at water distribution networks through simulation optimizations by selecting operating policies which reduce energy usage while meeting customer water demand. Typical simulated optimizations of complex hydraulic systems have high dimensional decision spaces and require significant time to evaluate. This study presents the design of a new multi-threaded simulation optimization software platform to determine pump operations for water distribution networks. The platform explores rule-based controls for pumps using derivative-free simulation optimization methods as independent, parallelized computational tasks. Decision spaces are reduced through domain division which produces smaller subproblems to be sequentially optimized. The platform is applied to a real urban water distribution system case study to determine energy efficient pump operating policies. The performance of several optimization techniques are compared, indicating that domain division approaches may improve consistency of optimization but are not necessarily beneficial for all optimization techniques.

1 INTRODUCTION

The optimal management of controls within large-scale water distribution systems is a long-standing problem inside the field of hydraulic engineering (Sterling et al. 1975). Optimally controlling pumps has been demonstrated to have large impacts on overall costs, energy use, and environmental impact of a water system (Makaremi et al. 2017). However, given variable water demands, interdependent system elements, and differing energy costs, it can be difficult it determine optimized pump controls to improve system performance (Mala-Jetmarova et al. 2017). This paper presents the design of a new, efficient multi-threaded simulation optimization software platform for determining pump operating policies to improve the energy usage of a water distribution network (WDN).

WDN operators develop pump operating policies to match observed water demand while maintaining water storage levels required for system resiliency (Klise et al. 2015). This decision requires integrating implicit or explicit forecasts of customer water demand with operator intuition to design pump and valve management strategies. Operators determine pump operating policies, usually in the form of conditional rules, which turn pumps “on” or “off” or assign pumps a high or low setting based on water storage levels or the time-of-day. To model scenarios, engineers may use hydraulic simulation software to observe the impact of a given pump operation policy on the system’s energy usage or hydraulic performance.

978-1-7281-9499-8/20/$31.00 ©2020 IEEE

This paper expands on previous research into pump schedule optimization within WDNs (Kang 2014; Bonvin et al. 2019; Fantozzi et al. 2014). Typically, optimization models have focused on finding optimal pump settings defined across specific time periods to minimize the operational cost or energy consumption of a WDN (Mala-Jetmarova et al. 2017). Methods have included linear programming (Jowitt and Germanopoulos 1992), non-linear and dynamic programming (Ormsbee and Reddy 1995), and simulated annealing (Goldman and Mays 2005), among others. Recent contributions to WDN optimization have also expanded the available methodology to Bayesian optimization and random forests with decision trees to model correct pump settings (Candelieri et al. 2018).

Some research has leveraged optimization approaches to select optimal rule-based pump controls dependent on system parameters such as flow, pressure and tank levels, to develop pump control strategies more resilient to uncertain demands (Mala-Jetmarova et al. 2017; Van Zyl et al. 2004). Researchers have used genetic algorithms to determine pump controls at WDNs based on fixed tank level triggers (Paschke et al. 2001), as well as variable tank level triggers where the control condition changes over time (Quintiliani and Creaco 2019; Van Zyl et al. 2004). Later work has explored using a genetic algorithm for optimizing multiple layers of rules which incorporate both system conditions and time-based conditions (Marchi et al. 2017; Blinco et al. 2016). Typically, these studies have limited their optimization methods to genetic algorithms. Additionally, optimization approaches may struggle with high-dimensional decision spaces in large-scale WDNs that contain many interacting pumps.

To address these computational challenges which arise from determining pump operating policies in large, complex WDNs, this paper introduces the design of a new multi-threaded optimization platform for the exploration of optimal operation policy of pumps. In the platform, different simulation optimization approaches are combined with a domain division ("divide-and-conquer scheme") to break the larger optimization problem into sub-problems which prioritize optimizing interconnected pumps sequentially. The platform uses several global search methods formatted to a multi-threaded system to runs hydraulic simulations in parallel to maximize computational resources available to a user in a server or single user environment. As a case study this paper explores the application of the platform to a hydraulic model of single region of a real WDN. Results for each global search method are examined and opportunities to improve the discovery of optimal decision points for practical pump operations are discussed.

2 PROBLEM MODEL AND SIMULATION OPTIMIZATION FORMULATION

The simulation optimization platform searches for a set of rule-based pump controls to minimize simulated hydraulic outputs such as energy use. To model the relationship between pump controls and hydraulic outputs, the platform makes uses of a hydraulic simulation toolkit. Based on the outputs of the hydraulic simulation, the platform applies derivative-free optimization methods to determine a set of pump control rules for reducing energy use.

To perform the hydraulic simulation of WDNs the platform leverages the EPANET toolkit (Rossman 2000). EPANET simulates WDN operations through a series of demand-driven steady-state calculations to ensure mass and energy balances throughout all network assets (Rossman 2000). Results of the simulation include time-series data of hydraulic head and flow throughout the network, as well as pump energy use and storage levels of local water reservoirs, herein referred to as tanks. The operation of pumps is controlled through sets of logical rules. These rules use the value of a particular system variable as a threshold for assigning a setting. Since the EPANET simulations are driven by customer water demand, EPANET contains hardcoded demand patterns that represent the demand at a given set of nodes for a particular time. A successful policy for pump operation will require a minimum amount of energy while meeting customer water demand and ensuring adequate tank levels.`

This paper defines a policy as a set of values which completely define rule-based pump logic inside a simulation. Within a given hydraulic model, the problem formulation uses a set of pumps (indexed $u \in (1...U)$) along with a set of tanks (indexed $r \in (1...R)$). Each pump is operated based on a set of rules that determine the on/off status as a function of tank levels or the time-of-day of the simulation. Each policy is written as a decision vector $\hat{x}$ which represents the pump status and threshold values across a set of pump rules. Vector $\hat{x}$ is used to write new rules to govern pump operation inside the simulation of a WDN. For example, a basic set of three rules to control operations for pumps written in EPANET would be:

LINK PUMP1 **Open** If TANK1 Below **17.00**
LINK PUMP1 **Closed** If TANK1 Above **22.00**
LINK PUMP1 **Open** AT TIME **10.68** HOURS

and could then be represented with the numeric decision vector:

$$\hat{x} = [x_1 = 1, x_2 = 0, x_3 = 1, x_4 = 17, x_5 = 22, x_6 = 10.68].$$

In this case each variable $(x_1, x_2,\ x_3, x_4, x_5, x_6)$ represents a value that defines the set of rules; x_1 for the pump setting on the first rule, x_2 for the pump setting on the second rule, x_3 for the pump setting on the third rule, x_4 for the threshold level for the first rule, x_5 for the threshold level for the second rule, x_6 for the threshold level for the third rule. Therefore, based on EPANET simulations, $Sim()$, the relevant hydraulic quantities can be modeled as a deterministic function of policy vector $\hat{x}$ and timeframe T. Observable hydraulic outputs α are formulated as:

$$\alpha = \ Sim_{\alpha}(\hat{x}, T).$$

A central requirement of WDNs is to always meet water customer demands. Since EPANET is demand-driven and requires pump operations and tank storage to mass-balance with water demand, demands can be met by imposing minimum storage requirements for tanks that are within hydraulic zones that serve water demand directly. Therefore, optimal rule-based management can be expressed as a simulation optimization problem with an imposed tank-level constraint. The full formulation of the optimization problem can be written as:

$$min_{\hat{x}} Y(\hat{x}) = \ min_{\hat{x}}\left(\sum_{u=1}^{U} Sim_{PumpEnergy_u}(\hat{x}, T)\right)$$

such that:

$$Sim_{TankLevel_r}(\hat{x}, T) > Tank_{min}$$

$$\hat{x} \epsilon X$$

where $Sim_{TankLevel_r}\ (\hat{x}, T)$ and $Sim_{pumpEnergy_u}$ are the simulated pump energy uses and tank levels over time period T for tank r and pump u, respectively. The constraint constant $Tank_{min}$ ensures that the tank maintains minimum water storage for hydraulic stability under shifting customer water demands. The domain X is the policy search space defined as upper and lower bounds for each rule-settings decision value.

Through the artificial penalty value added to the objective function if the tank level drops below $Tank_{limit}$, the constrained optimization problem is transformed into an unconstrained simulation optimization problem inside a set of multi-dimensional box-constraints. Using this formulation of the problem, the platform applies modifications of existing simulation optimization methods to determine pump operating policies that reduce energy usage in a WDN.

3 SIMULATION OPTIMIZATION METHODS

A multi-threaded simulation optimization platform was programmed in C++ to explore pump operating policies while performing EPANET simulations. The platform includes several alternative derivative-free global optimization methods that can be applied to the hydraulic simulation results. The platform also allows the user to specify a "divide-and-conquer" domain division scheme to improve the optimization efficiency for high dimensional decision spaces (X) where decomposing the system into smaller optimization problems may improve the computational efficiency of the optimization. The platform is designed to leverage multi-threading to run simulations and optimization algorithms in parallel to maximize performance with modern hardware on servers or individual computers (see Figure 1).

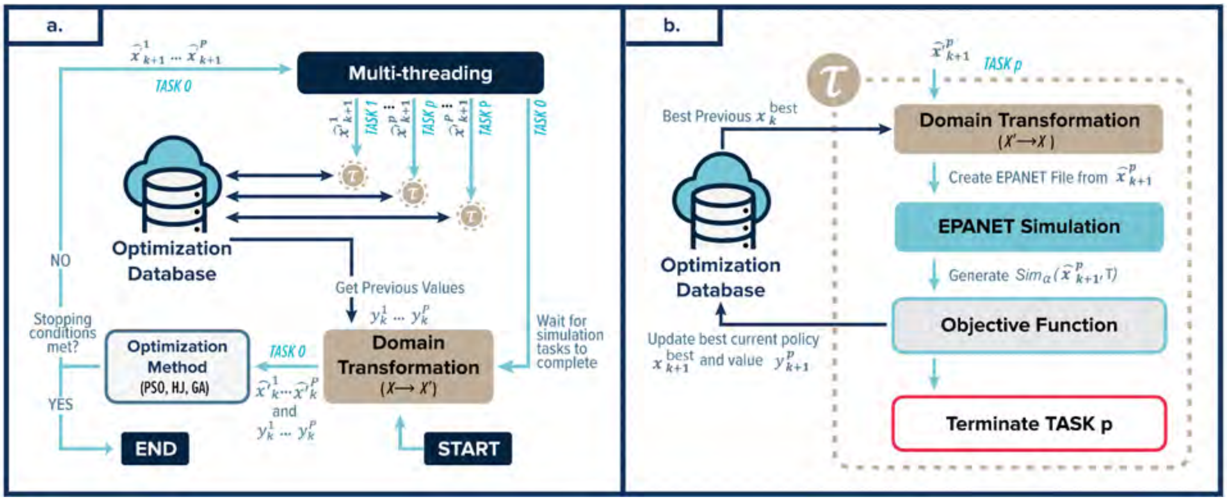

Figure 1: The optimization platform and its interaction with the EPANET simulation framework with multiple threads. Figure 1.a shows the platform first transforming the domain (domain division), point sampling via the optimization methods, and the beginning of the multi-threaded component of the platform. Figure 1.b shows the activity inside a single thread, first performing a domain transformation by combining the sampled point with the best previously sampled point, running the simulation, and then computing the objective value before terminating the thread.

Throughout the optimization, the platform maintains a central data repository, the "optimization database", that houses policy vectors, hydraulic simulation results and other intermediate results required to perform the optimization and is safely accessible to all parallelized tasks simultaneously utilizing mutexes for each data collection within the database. At a given iteration k the optimization process decomposes the problem to search on a smaller domain X'. Using the previous policies $({x'_k}^1 \ldots {x'_k}^P)$ and previous objective values $({y'_k}^1 \ldots {y'_k}^P)$, a selected optimization algorithm generates a batch of P new policy inputs, ${x'_{k+1}}^1 \ldots {x'_{k+1}}^P$. For each policy p, a task (τ_p) is submitted into a multi-threaded queue to perform the hydraulic simulation. Then on each thread, a domain transformation $X'\ to\ X$ is performed by combining the partial policy $({x'_{k+1}}^P)$ with the best current policy to set a form a full policy $({x_{k+1}}^P)$. Each full policy is used to write the controls for an EPANET simulation which runs on the thread τ_p to generate a set of hydraulic values $(Sim_\alpha(\,{x_{k+1}}^P, T))$. Finally, the objective values are computed from the hydraulic values, and then the objective values $({y'_{k+1}}^1 \ldots {y'_{k+1}}^P)$ are updated to the optimization database for the next iteration and the threaded task τ_p terminates. Optimization iterations continue until the maximum number of function evaluations are exceeded or the performance does not improve over a given number of iterations.

Three common optimization algorithms were implemented in the platform which perform the domain transformations, batching and policy computations including Particle Swarm Optimization (PSO) (Kennedy and Eberhart 1995), Hookes-Jeeves Pattern Search (HJ) (Hooke and Jeeves 1961), and a simple

Genetic Algorithm (GA) (Eiben et al. 1994). The format of each method is briefly discussed along with modifications used inside the simulation optimization platform to make the algorithms compatible with the multi-threaded design.

PSO is a popular population-based search algorithm that searches a domain based on moving a set of sampled policies (particles P). Each particle's movement is characterized as a proportion (ω) of a random velocity and bias (φ_1) towards the best policy observed by a single particle and a bias (φ_2) towards the best policy observed by the entire system of particles. To enable batching of the simulation optimization tasks, at the start of every iteration all potential moves are determined simultaneously with updated values and used to generate the numerous simulation optimization tasks.

The HJ optimization method is a common pattern-search algorithm where new solutions are located by moving a starting policy in a random direction within a given domain with a fixed step size (δ). After a direction no longer improves the objective function observed, the algorithm selects a new random direction and moves with a step size reduced by a ratio (ρ) until the objective function no longer improves. The method was formatted to increase its efficiency through the inclusion of multiple starting points (P). Additionally, each non-improving point will only check one random direction before attempting to move. The algorithm batches one new policy for each point, which represents either an improving direction or a test direction. Particle locations are updated at the beginning of each iteration and individually batched as tasks for computation.

Lastly, the basic GA focuses on an evolutionary approach to exploring new potential policies based on the recombination and variation of already discovered policies with promising objective values. Starting with a given population of sampled policies (P), a designated number of well-performing "elites" are copied over to the next batch (c_{elite}), the best points from the current population are then recombined randomly to create a proportion (ρ recombine) of the new potential polices for the next iteration. Finally, the remaining percentage of poorly performing policies are "mutated" by re-selecting them uniformly across the bounds of the domain. After recombination and mutation, updated policies are batched and run on multiple threads.

WDN operators and engineers seeking to improve pump operations themselves will typically explore pump or pump station controls one at a time or based on geographic or network grouping. Improvements are often iterative, starting from an initial, trusted set of pump controls and changing certain thresholds to measure the overall impact on the system performance. Since pumps at different locations may have independent operation, optimizing geographically distinct or hydraulically distant pump groups may more efficiently determine promising pump control policies. Incorporating the concept of this approach, the platform allows users to specify a "divide-and-conquer" scheme dividing the optimization problem into sub-problems. In each sub-problem, the optimization focuses on limited pump variables while keeping the remaining variables set at the best currently known values. The division creates a new search space ($X' \subset X$) as subset of the total domain and can potentially reform an intractable high-dimension optimization problem into a series of easier to solve lower-dimension sub-problems.

Altogether, the simulation optimization platform provides a variety of options for exploring optimal pump settings. First, the platform uses three provided simulation optimization methods. Second, the platform offers the option of a "divide-and-conquer" scheme to separate a larger pump control problem in a series of smaller pump control problems. Based on each one of these selections, a variety of different optimization approaches are available to a user. To illustrate the available methods for pump policy determination, the paper explores a case study and compares the relative efficiency of the various methods for determining energy-efficient pump policy recommendations.

4 CASE STUDY AND NUMERICAL RESULTS

To test the performance and features of the simulation optimization platform with a case study, a hydraulic simulation model was optimized using all three optimization methods, PSO, HJ and GA, for a non-divided domain and two distinct "divide-and-conquer" schemes.

The case study explores the optimization of the pump operations in a single pressure zone, or hydraulically independent region, of the potable water distribution system managed by the Moulton Niguel

Water District (MNWD) in southern California. The hydraulic model used in the case study is a component of a larger hydraulic model built and calibrated by the Center for Water-Energy Efficiency at the University of California, Davis for the MNWD in 2018-2019 for the purpose of operational planning and system analysis (see Figure 2 for a simplified diagram). The simulation length is 168 hours with a resolution of 60 seconds and represents typical operation in a single hot-weather week. All locational and identifiable data have been stripped from the hydraulic model.

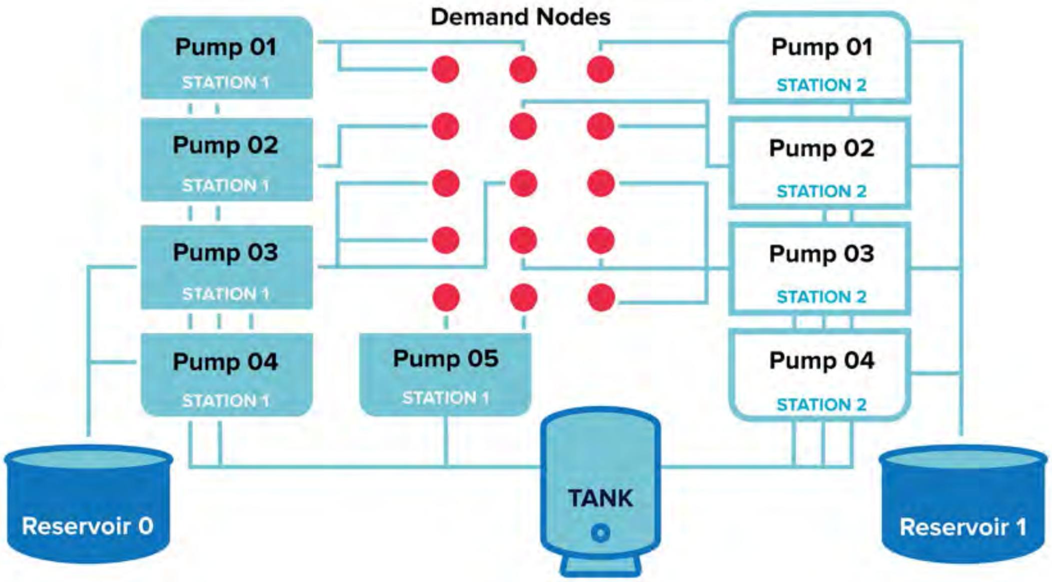

Figure 2: Presented is a simplified illustration of the case study hydraulic model. Here two sets of operating pumps lift water from lower reservoirs to a central water tank and water customer demand nodes.

The case study model includes over 700 demand nodes and contains two pump stations, one with four pumps and the other with five pumps. Both pump stations draw water from connected pressure zones and supply an above ground storage tank with a capacity of 3.9 million gallons. The connected pressure zones are modeled as water sources for the simulation with variable hydraulic head and infinite volume. Four pumps in this zone have static operations, one pump continually operates and three pumps are turned off, and were therefore not included in the optimization procedure. The remaining four pumps, two at each pump station, were selected as the subject of this optimization problem. Each of these pumps are controlled by two sets of rules: one that turns the pump on when the tank reaches a low level and one that turns the pump off when the tank is full. Using original controls, the total pump energy use is 29.9 megawatt-hours for the simulation period.

For this case study, the platform is used to optimize the controls in the hydraulic model by modifying the values of the conditional thresholds based on the level of the zone's water storage tank. The settings for the pump statuses were not modified in the decision vector. As such, this optimization-simulation problem optimizes eight continuous variables that describe the threshold tank levels for turning the pump settings on or off.

To maintain system resiliency, the constraint on the level of the zone's single water storage tank ($Tank_{min}$) was set to 10 % above the minimum tank level specified in the hydraulic model. The upper and lower bounds (box constraints), which determine the search domain of the problem (X) are defined as 5% above the minimum tank level and below the maximum tank level respectively. Additionally, the original control settings were used as the default feasible value for pump status triggers to assist the optimization in finding feasible results.

For each of the optimization methods described in Section 3, the simulation optimizations used a standard set of method parameters. The PSO method uses $P = 5$ particles with an initial velocity coefficient $\omega = 0.5$ and with the local and global velocity components set at $(\varphi_1 = 0.5)$ and $(\varphi_2 = 0.5)$, respectively. The HJ algorithm uses $P = 5$ particles, a starting step of $(\delta = 0.3)$ multiplied by the length of the respective box dimension, and a step-size reduction ratio of $\rho = 0.5$. Finally, the GA uses a population of $p = 5$, with a single elite $(c_{elite} = 1)$, and recombining 60% of the policy points with the remainder of the policy points being randomly mutated.

For the divide-and-conquer scheme, two methodologies were examined for dividing the optimization problem into smaller sub-problems. The first domain division scheme divides the variables associated with each of the four pumps by location ("divide-and-conquer scheme 1"), with the pumps from group 1 being optimized over the first half of an algorithm's iterations and the remaining iterations dedicated to optimizing the threshold setting variables of the other two pumps in the second pump grouping. For comparison, this paper examines a second domain division scheme which groups one pump from each geographic region together, dedicating the first half of the method's iterations to optimizing the first group of variables and the second half of the method's iterations to the second (see Figure 3 for result distributions).

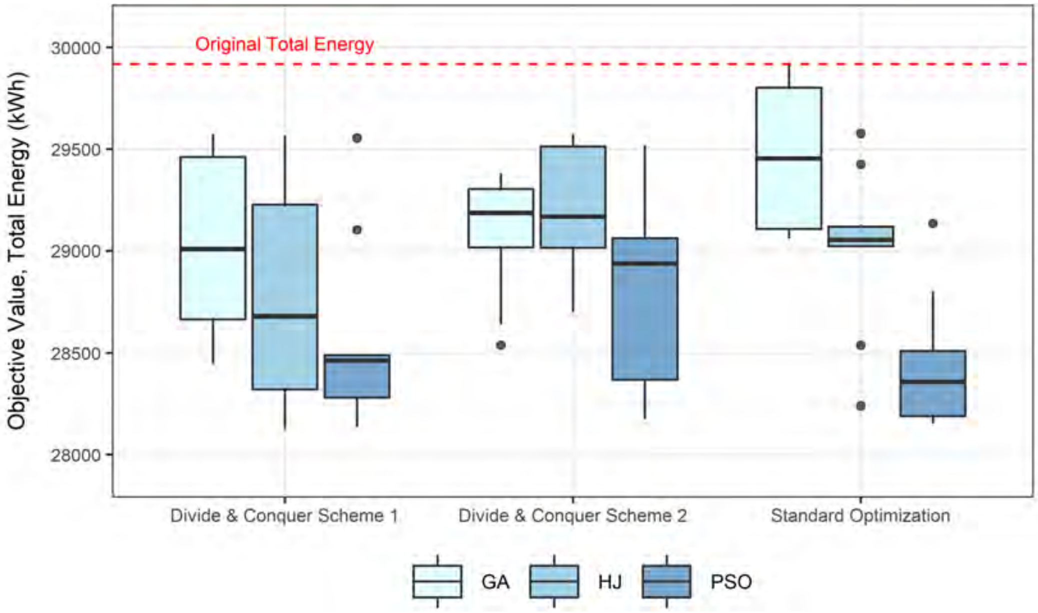

Figure 3: The box plots present a comparison of the energy usage (kWh) of policies found by each method applied to the case study optimization problem. Each box plot represents the distribution of final objective values for the best determined pump policy across each of the ten replications.

To limit the optimization's computational budget, the maximum number of function evaluations was set to 50 for each experiment. Using each optimization method paired with a divide-and-conquer scheme, including using no division of the domain, required nine optimization experiments with ten replications of each experiment. Each of the nine experiments were characterized by a range of objective values, or total energy consumption in kilowatt-hours (kWh), through the repeated discovery of improved operating policies.

Each of the optimization methods generates a policy that improves on the initial default policy, where each pump is active for fewer hours while still meeting demand and maintaining the tank's lower limit. The relative effectiveness for each optimization approach used for this case study is compared based on the observed medians across each of the replications for each experimental method (see Table 1). All methods made improvements from the baseline energy demand of 29,919.39 kWh by approximately 2-5%.

Table 1: The median values generated by each of the various optimization methods paired with different domain division schemes.

Optimization Method	Divide-and-Conquer Scheme 1 (kWh)	Divide-and-Conquer Scheme 2 (kWh)	Standard Optimization (kWh)
GA	29,007.63	29,185.46	29,452.66
HJ	28,677.19	29,167.42	29,051.59
PSO	28,462.77	28,936.60	28,356.68

The descent profiles of the replications from each experimental method (see Figure 4) show the relative effect of function evaluations on optimizer performance in terms of objective value. Generally, the policies with the lowest, or most optimum, objective values are achieved by the PSO method. The lowest median objective value is generated by PSO without division of the domain, followed by the PSO with the geographic division strategy “Divide-and-Conquer Scheme 1”. The GA generated the least optimum median objective values out of the various optimization methods except under the ordering division strategy “Divide-and-Conquer Scheme 2” where it performs similarly to HJ.

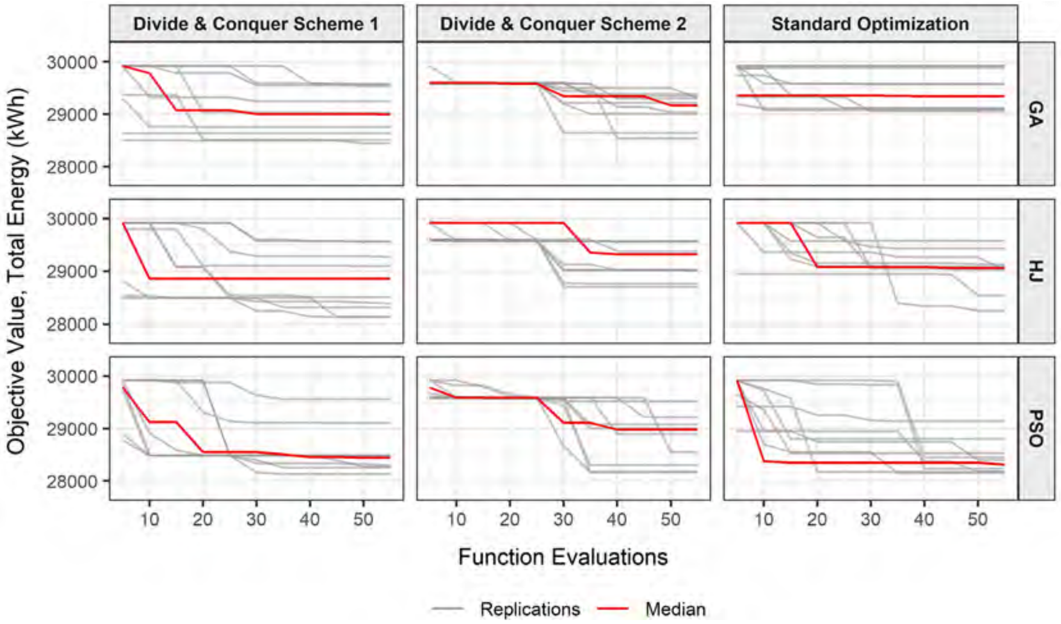

Figure 4: Descent profiles for each of the nine experiments show the objective value (total energy use in kilowatt-hours (kWh)) versus the number of function evaluations. Each replication’s descent profiles are illustrated in grey. The descent profile of the replication with the median final value is highlighted in red.

The “Divide-and-Conquer Scheme 1” improved the performance of the optimizations under HJ and GA, while the “Divide-and-Conquer Scheme 2” generally reduced optimization performance in terms of median objective value with the exception of GA.

A general tendency of the GA and HJ descent profiles indicate shallower objective value decreases, with more consistent decreases in objective value from the PSO approach. Furthermore, the full domain optimization and the first divide-and-conquer scheme show much more variance across the slope of descent for each of the optimization methods. The results suggest that simulation optimizations using the GA and HJ algorithm or using the “Divide-and-Conquer Scheme 2” are more affected by the selection of an initial starting policy for each of the particles. Moreover, the descent profiles suggest that poorly performing optimization methods are not sufficiently exploring the domain and are unable to escape local optima. Further changes in parameter selection may allow the evaluated methods to overcome this hurdle and spend more of the computational budget exploring the domain. While all optimization methods reduce the initial pump policy's energy usage within the constraints of the problem, the PSO is overall the most effective optimizer for this case study, with the most effective approach being the application of the PSO solver without any division of the domain. While showing benefits when paired with some optimization methods, the domain division schemes do not uniformly improve the efficiency of the optimization and, as in the case of the "Divide-and-Conquer Scheme 2", may decrease the efficiency of the optimization. These results may be the consequence of the limited number of decision variables, and that both sets of pumps connect to a common storage tank within a single hydraulic zone.

5 CONCLUSIONS AND FUTURE RESEARCH

This paper presents the design of a new, multi-threaded simulation optimization platform for the efficient optimization of water distribution network operations. The platform searches for pump operating policy improvements using any of a series of derivative-free optimization methods. Domain division of pumps into sub-problems for sequential optimization has been implemented in the platform to improve performance of optimizations for large, complex water distribution networks. The platform design is generalized to support custom division schemes, optimization methods, and simulation strategies while enabling performance scalability on individual computers or server environments.

The optimization platform is leveraged in a case study exploring pump policy solutions for the hydraulic model of a single hydraulic zone from the potable water network of the Moulton Niguel Water District. Among the simulation optimization methods, PSO was the most efficient approach to discovering optimal pump policies. The performance of domain division in improving optimization was dependent on the selected simulation optimization method and precise domain division strategy. In the context of optimizing the policies of pumps serving a single hydraulic zone, the domain division strategy did not uniformly improve optimization performance.

Generally, the multi-threaded simulation optimization platform would be useful for water distribution systems and engineers seeking to evaluate opportunities for improved pump control policies to reduce the energy consumption within a distribution system. The availability of different optimization methods that are integrated into the multi-threaded design will enable the determination of optimized pump operations on a limited computational budget where decision support requires rapid response on modest hardware. Commercial and academic research communities in hydraulic modeling would benefit from the platform through the ability to generate customizable optimizations for user-driven hydraulic objectives that can be performed utilizing background threads without requiring users to pause other hydraulic modeling activities.

This research was limited by both the scope and complexity of the case study and the computational budget. In larger case studies, there is an opportunity to observe improved application of the domain division schemes by leveraging partial independence in the operation of pumps within different hydraulic zones. An increased number of function evaluations may also lead to improved solutions across all domain division schemes. Further research will apply the platform to optimize the operations of hydraulic networks with several, complex hydraulic zones and explore the impact of additional function evaluations. Additionally, the independence of simulation and optimization tasks presents the opportunity to seamlessly combine optimization approaches at different iterations or within different sub-problems through a robust optimization management strategy. Future research should investigate the benefit of using exploration-focused algorithms at early iterations paired with methods that exploit local solutions at later iterations. Future research should also expand to include stochastic customer demands to improve simulation accuracy. Expansions to the algorithms to support stochastic simulation optimizations will be necessary for the platform to accommodate these stochastic demands.

ACKNOWLEDGMENTS

Special thanks to the California Energy Commission for financing this study with funds from the Electric Program Investment Charge (EPIC), the ratepayer surcharge authorized by the California Public Utilities Commission, through agreement number EPC-16-062. Special thanks to the electricity ratepayers in California for funding this and other research studies. Special thanks to the Moulton Niguel Water District (MNWD) and its staff for their excellence in supporting this and other research into the water-energy nexus and energy efficiency in the water sector. In particular the authors appreciate the support and collaboration of Joone Lopez, General Manager at the MNWD; Drew Atwater, Director of Finance and Water Resources at the MNWD; Ronin Goodall, Superintendent Of Operations at the MNWD; and many other operations, engineering, and technical support staff at the MNWD.

This document was prepared as a result of work sponsored by the California Energy Commission. It does not necessarily represent the views of the Energy Commission, its employees, or the State of California. The Energy Commission, the State of California, its employees, contractors, and subcontractors make no warranty, express or implied, and assume no legal liability for the information in this document; nor does any party represent that the use of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the Energy Commission nor has the Energy Commission passed upon the accuracy of the information in this report.

REFERENCES

Blinco, L. J., A. R. Simpson, M. F. Lambert, and A. Marchi. 2016. “Comparison of Pumping Regimes for Water Distribution Systems to Minimize Cost and Greenhouse Gases”. *Journal of Water Resources Planning and Management* 142(6):04016010.

Bonvin, G., S. Demassey, and A. Lodi. 2019. “Pump Scheduling in Drinking Water Distribution Networks with an LP/NLP-Based Branch and Bound”. The Center for Direct Scientific Communication. https://hal-mines-paristech.archives-ouvertes.fr/hal-02158535, accessed 1st April 2020.

Candelieri, A., R. Perego, and F. Archetti. 2018. “Bayesian Optimization of Pump Operations in Water Distribution Systems” . *Journal of Global Optimization* 71(1):213–235.

Eiben, A. E., P. E. Raue, and Z. Ruttkay. 1994. “Genetic Algorithms with Multi-Parent Recombination”. In *Parallel Problem Solving from Nature,* edited by Y. Davidor, H.P. Schwefel, R. Männer, 78–87. Berlin, Heidelberg: Springer.

Fantozzi, M., I. Popescu, T. Farnham, F. Archetti, P. Mogre, E. Tsouchnika, C. Chiesa, A .Tsertou, M. Castro Gama, and M. Bimpas. 2014. “ICT for Efficient Water Resources Management: The ICeWater Energy Management and Control Approach”. *Procedia Engineering* 70:633–640.

Goldman, F. E., and L. W. Mays. 2005. “Water Distribution System Operation: Application of Simulated Annealing”. *Water Resouces Systems Management Tools*. New York: McGraw Hill.

Hooke, R., and T. A. Jeeves. 1961. “'Direct Search' Solution of Numerical and Statistical Problems”. *Journal of the Association for Computing Machinery* 8 (2):212–229.

Jowitt, P. W., and G. Germanopoulos. 1992. “Optimal Pump Scheduling in Water-Supply Networks”. *Journal of Water Resources Planning and Management* 118(4):406–422.

Kang, D. 2014. “Real-Time Optimal Control of Water Distribution Systems”. *Procedia Engineering* 70:917–923.

Kennedy, J., and R. Eberhart. 1995. “Particle Swarm Optimization”. In *Proceedings of International Conference on Neural Networks*, November 27th - December 1st, Perth, Australia, 4:1942–1948.

Klise, K. A., R. Murray, and N.W. La Tonya. 2015. “Systems Measures of Water Distribution System Resilience”. Technical Report No. SAND2015-20746R, Sandia National Lab, Albuquerque, New Mexico.

Makaremi, Y., A. Haghighi, and H. Reza Ghafouri. 2017. “Optimization of Pump Scheduling Program in Water Supply Systems Using a Self-Adaptive NSGA-II; a Review of Theory to Real Application”. *Water Resources Management* 31(4):1283–1304.

Mala-Jetmarova, H. N. Sultanova, and D. Savic. 2017. “Lost in Optimisation of Water Distribution Systems?” Open Research Execter. http://hdl.handle.net/10871/25685, accessed 1st April 2020.

Marchi, A., A. R. Simpson, and M. F. Lambert. 2017. “Pump Operation Optimization Using Rule-Based Controls”. *Procedia Engineering* 186:210–217.

Ormsbee, L. E, and S. L. Reddy. 1995. “Nonlinear Heuristic for Pump Operations”. *Journal of Water Resources Planning and*

Management 121(4):302–309.

Paschke, M., K. Spencer, N. Waniarcha, A. R. Simpson, and T. Widdop. 2001. "Genetic Algorithms for Optimising Pumping Operations". In *19th Federal Convention, Australian Water Association*, April 1st - 5th, Canberra, Australia.

Quintiliani, C., and E. Creaco. 2019. "Using Additional Time Slots for Improving Pump Control Optimization Based on Trigger Levels". *Water Resources Management* 33(9):3175–3186.

Rossman, L. A. 2000. "EPANET 2: Users Manual Rep. No. EPA/600/ R-00/057". Water Supply and Water Resources Division, National Risk Management Research, Cincinnati, Ohio.

Sterling, M. J. H., and B. Coulbeck. 1975. "Technical Note. A Dynamic Programming Solution to Optimization of Pumping Costs". *Proceedings of the Institution of Civil Engineers* 59(4):813–818.

Van Zyl, J. E., D. A. Savic, and G. A. Walters. 2004. "Operational Optimization of Water Distribution Systems Using a Hybrid Genetic Algorithm". *Journal of Water Resources Planning and Management* 130(2):160–170.

AUTHOR BIOGRAPHIES

DAVID D. LINZ is a Research Engineer at the Center for Water-Energy Efficiency at the University of California, Davis specializing in optimization, modeling, and decision-making under uncertainty. David obtained his Bachelor degree in Mechanical Engineering from UC Davis with a subsequent Masters and PhD in Industrial Engineering from the University of Michigan and the University of Washington respectively. David's current research interests concern simulation optimization, dynamic decision-making models, and project management. His email is ddlinz@ucdavis.edu.

ERIN MUSABANDESU is a Graduate Student Researcher at the Center for Water-Energy Efficiency at the University of California, Davis. Her current research focuses on water and wastewater energy optimization and management. Prior to starting her Master's/PhD program at UC Davis, she worked in the private engineering sector for six years and in international development for two years. Erin earned her bachelor of science degree from UC Davis in Civil and Environmental Engineering with an emphasis in environmental engineering. Her email is enmusabandesu@ucdavis.edu.

BEHZAD AHMADI is a Research Engineer for the Center for Water-Energy Efficiency at the University of California, Davis specializing in water resources management, optimization and automation of water distribution system, and developing decision support tools for water distribution management. Behzad obtained his PhD in civil and environmental engineering from Portland State University. His email is behahmadi@ucdavis.edu.

ROBERT GOOD is the Engineering Manager at the Center for Water-Energy Efficiency at the University of California, Davis where he leads and manages the center's research projects. His research interests include optimization and automation of water distribution networks, real-time data visualization, distributed water treatment and reuse, and the development of technologies to improve regional water resource and energy management. Robert received a M.S. in Civil and Environmental Engineering from UC Davis. His email is rtgood@ucdavis.edu.

FRANK LOGE is a Professor in the Department of Civil and Environmental Engineering and Director of the Center for Water-Energy Efficiency at the University of California, Davis. He began his career studying water and wastewater treatment and has since expanded his research to encompass the interconnected relationship between energy, water, and greenhouse gas emissions. He advocates for fundamental research to advance knowledge, as well as the development and deployment of technologies or policies that improve environmental quality and the efficiency of resource consumption. His research interests are: water and energy efficiency in urban and agricultural systems, water/energy demand management, sustainable building design, water reuse, conservation-based water rates, data privacy/security, data analytics, and entrepreneurship. Dr. Loge holds a Ph.D. in Civil and Environmental Engineering from the UC Davis and is a licensed Professional Engineer in the state of California. His email is fjloge@ucdavis.edu.