2nd International Joint Conference on Water Distribution Systems Analysis & Computing and Control in the Water Industry
Valencia (Spain), 18-22 July 2022
doi: https://doi.org/10.4995/WDSA-CCWI2022.2022.14815

A NONLINEAR MODEL PREDICTIVE CONTROL FRAMEWORK FOR DYNAMIC WATER NETWORK OPTIMIZATION

Ernesto Arandia[1] and James Uber[2]

[1]Xylem Digital Solutions, Viareggio, LU, Italy
[2]Xylem Digital Solutions, Cincinnati, OH, USA

[1]ernesto.arandia@xylem.com, [2]james.uber@xylem.com

Abstract

This paper proposes the application of nonlinear model predictive control (NMPC) strategies within an integrating framework for planning and scheduling real-time water distribution system operations. Model predictive control (MPC) refers to a class of algorithms that make explicit use of a process model to optimize the future predicted behavior of a system. Originally developed to address the control needs of power plants and petroleum refineries, over the past 30 years it has successfully been used in a very wide range of applications in industry. In addition to the predictive model, MPC consists of a performance metric reflecting the control actions, an optimization algorithm that computes a sequence of future control signals that minimizes the performance index subject to a given set of constraints, and a moving horizon strategy, according to which only the first time point of the optimal control sequence is applied online. The predictive models are generally intended to represent the behavior of complex nonlinear dynamical systems and often consist of linear models that circumvent the stability and performance challenges associated with their nonlinear counterparts. Recent developments on methods to model and solve large-scale nonlinear optimization problems lead to a reconsideration of these formulations, in particular, through development of efficient large-scale barrier methods for nonlinear programming (NLP). As a result, it is now realistic to solve NLPs on the order of a million variables. These developments are leveraged here to develop an NMPC dynamic real-time optimization strategy that combines a nonlinear dynamical model of the water network with a large-scale NLP optimization algorithm, and a moving horizon strategy to determine how to exert control on the system at the pump facility and network levels. We consider a recent water utility case study to illustrate these concepts.

Keywords
Dynamic water network optimization, nonlinear model predictive control, operation optimization.

1 INTRODUCTION

Dynamic optimization problems involving real-world water networks present several challenges, including the large scale of the networks, the nonlinearity of the relationships among physical quantities, the discrete nature of operation decisions such as pump selection, and the fact that state-of-the-art optimization solvers do not directly handle integrals or differential equations.

The approach proposed in this paper consists of a combined nonlinear model predictive control (NMPC) strategy designed to address these challenges. The scale problem is addressed through a combination of net- work skeletonization and the use of a large-scale NLP solver that can handle millions of variables efficiently. The nonlinearities are explicitly expressed through hydraulic relationships for all the types of network components involved. The discrete decisions are handled by decomposing the optimization problem into a primary or master problem that solves the network first, and a secondary problem that solves the pump facilities using results from the primary problem. The optimization dynamics is captured through a differential-algebraic formulation in the model constraints.

Previous studies dealing with water network operation optimization consist mostly of academic analyses and very few industrial-scale undertakings. Their methodologies are widely based on the application of heuristic methods that do not address the challenges outlined above. One exception is [2], that describes the development and implementation of a large-scale NLP based pump scheduling system for a city-wide water distribution system. Other examples of NMPC studies designed for large-scale systems are found in the chemical engineering domain, e.g., [1] present a large-scale optimization framework for NLP applied to control chemical processes in real time.

The goal of this paper is to describe the proposed NMPC methodology and the proof-of-concept (PoC) that demonstrates the feasibility and optimization capabilities of the proposed solution. Since increasing energy cost efficiency is often one of the key optimization goals in operations and is also straightforward to use when contrasting optimized scenarios with the status quo, the PoC and this paper focus on energy and pressure optimization first. A near future development will explicitly incorporate water quality in addition to energy and pressure optimization.

The paper is organized as follows. Section 2 describes the NMPC technology from the perspectives of the supply network and the pump facility. Section 3 discusses the proof-of-concept (PoC) established to demonstrate the approach. Section 4 contrasts the actual and optimized operations with focus on energy optimization. Finally, Section 5 presents a summary and discusses future work.

2 METHODOLOGY

Under the proposed approach, the water supply network is a dynamic system that is controlled through optimization over a number of discrete time periods. The water network optimization problem consists thus of (1) a number of dynamic stages, (2) state, control, and algebraic variables, and (3) optimization constraints in the form of differential and algebraic equations (DAEs). The solution of this dynamic optimization problem is formally the time evolution of the network state over a moving (rolling) horizon. This solution is imposed on the facility optimization problem which deals with the selection of pumps to run and their relative speeds, in case they have variable frequency drives (VFDs), over the specified optimization horizon. The approach is general in that any number of pumps are considered and that any of them may have VFDs. The only restriction is, however, that the pumps are arranged in parallel.

Below, we briefly discuss the pump facility abstraction that is introduced to allow formulating the network optimization problem as an NLP. Then we formulate and describe the primary problem in more detail. A discussion of the pump facility problem is beyond the scope of this paper.

2.1 Pump facility abstraction

As illustrated in Fig. 1, under the introduced abstraction the pump facility consists of an upstream node (reservoir), one or more pumps arranged in parallel, one or more throttle control valves, and one or more downstream nodes. The control valves may already exist as shown in the diagram of Fig. 1 that represents a real pump facility in the system described in Section 3. Otherwise, the control valves are added to the model while ensuring the network configuration downstream of the facility is not altered. The purpose of these valves is to remove the concerns about the pump switching behavior by keeping the pumps running continuously while adjusting the control valve settings during optimization.

For consistency, the effective head gains through the pump facility are measured between the valves' downstream nodes (C and D in Fig. 1) and the upstream reservoir. The facility discharge flows are computed as well for each valve using the valves' head loss relationships (see Eq. 5 and 6). Moreover, the electrical power consumption of the facility is calculated per valve using the head gains and flows under the assumption that there is a constant global efficiency for the pump facility. This global efficiency is estimated from the calibrated head and efficiency curves for the

pumps. The power consumption for the whole facility is computed by summing through the control valve powers (see Eq. 11).

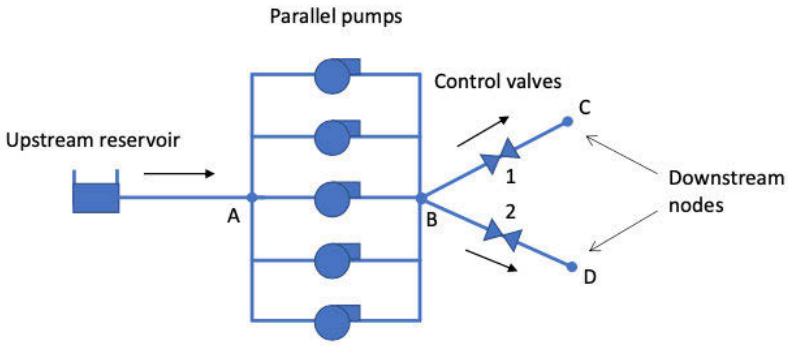

Figure 1. Diagram of the pump facility abstraction.

2.2 Network optimization problem

The main characteristics of the NMPC water network optimization model are listed below:

- The network state variables consist of hydraulic heads at non-fixed head nodes. More specifically, the state variables are partitioned into algebraic variables, i.e., hydraulic heads at junctions, and differential variables, i.e., the hydraulic heads at tanks.
- The dynamics of tank water volume is obtained through differential equations that impose flow balance constraints at the tank nodes.
- Control at the pump facilities and throughout the network is exerted through control valves where the decision variables consist of the valves' opening fraction. The pump facility control valves allow for removing the binary decisions associated with pump selection, at the network solution stage.
- Check valves are handled through a simple non-negativity constraint.
- Head loss and gain expressions at network edges consist of three types of continuous functions, namely, the Hazen-Williams equations for pipes, a power-law formulation for pumps, and the head loss equation for throttle control valves (TCVs).
- Mass balance is imposed at all junctions following the so-called node method formulation in the form of algebraic constraints.
- Appropriate upper and lower bounds are imposed to the control, state and other (auxiliary) variables.
- The objective function consists of the combined energy use and energy demand costs over the rolling operation horizon.

The following sets, parameters, and decision variables are introduced to formulate the water network optimization problem.

Sets:

T : the set of time periods

E_P : the set of pipes

E_U : the set of pumps

E_V : the set of valves

E^*_V : the set of pump facility valves

V_J : the set of junctions

V_R : the set of reservoirs

V_T : the set of tanks

F : the set of pump facilities subject to demand charges

E : the set of network links

T_P : the set of peak-hour time periods

T_O : the set of off-peak-hour time periods

Parameters:

: specific weight of water

Δ_T : length of each time period

$c_{j,t}$: electricity cost for facility valve j in time period t

η_j : efficiency associated with facility valve j, assumed constant

$D_{i,t}$: demand of junction i in period t

H_{Ri} : hydraulic head of reservoir I, assumed constant

I_{ij} : incidence of node i and link j

R_{pj} : resistance coefficient of pipe j

R_{vj} : resistance coefficient of valve j

a_{lj} : $l \in \{1, 2, 3\}$, head curve coefficient l of pump j

α : Hazen-Williams exponent

ε : small value set to 10−8

e_i : elevation of node i

$L_{i,min}$: minimum water level in tank i

$L_{i,max}$: maximum water level in tank i

$p_{i,min}$: minimum pressure head of junction i

$p_{i,max}$: maximum pressure head of junction i

A_i : section area of tank I, assumed cylindrical

s_k : demand charge on peak hours for pump facility k

w_k : demand charge on off-peak hours for pump facility k

$\bar{P}_k$: maximum to-date power peak during peak hours for pump facility k

$\bar{W}_k$: maximum to-date power peak during off-peak hours for pump facility k

Decision variables:

$H_{Ji,t}$: hydraulic head of junction i in time period t,

$H_{Ti,t}$: hydraulic head of tank i in time period t,

$h_{j,t}$: hydraulic head loss through link j in time period t,

$x_{j,t}$: fractional opening of control valve j in time period t,

S_k : highest 30-min interval power to date at pump facility k during peak hours,

W_k : highest 30-min interval power to date at pump facility k during off-peak hours.

The primary optimization problem is formulated below:

$$\min \quad \sum_{t\in T}\sum_{j\in E_V^*}\frac{\gamma\Delta_T c_{j,t}}{\eta_j}q_{j,t}h_{j,t}+\sum_{k\in F}\left(s_kS_k+w_kW_k\right), \tag{1}$$

$$\text{s.t.} \quad \sum_j I_{ij}q_{j,t}-D_{i,t}=0 \quad \forall i\in V_J,\forall t\in T, \tag{2}$$

$$q_{j,t}=-R_{p_j}^{-\frac{1}{\alpha+1}}h_{j,t}(h_{j,t}^2+\epsilon)^{-\frac{\alpha}{2(\alpha+1)}} \quad \forall j\in E_P,\forall t\in T, \tag{3}$$

$$q_{j,t}=a_{2j}^{-\frac{1}{a_{3j}}}\left(a_{1j}-h_{j,t}\right)^{\frac{1}{a_{3j}}} \quad \forall j\in E_U,\forall t\in T, \tag{4}$$

$$q_{j,t}=-R_{v_j}^{-\frac{1}{2}}x_{j,t}h_{j,t}(h_{j,t}^2+\epsilon)^{-\frac{1}{4}} \quad \forall j\in E_V,\forall t\in T, \tag{5}$$

$$h_{j,t}=\sum_{i\in V_J}I_{ij}H_{J_{i,t}}+\sum_{i\in V_T}I_{ij}H_{T_{i,t}}+\sum_{i\in V_R}I_{ij}H_{R_i} \quad \forall j\in E,\forall t\in T, \tag{6}$$

$$h_{j,t}\geq 0 \quad \forall j\in E_U,\forall t\in T, \tag{7}$$

$$H'_{T_{i,t}}-\frac{1}{A_i}\sum_j I_{ij}q_{j,t}=0 \quad \forall i\in V_T,\forall t\in T, \tag{8}$$

$$e_i+L_{i,\min}\leq H_{T_{i,t}}\leq e_i+L_{i,\max} \quad \forall i\in V_T,\forall t\in T, \tag{9}$$

$$e_i+p_{i,\min}\leq H_{J_{i,t}}\leq e_i+p_{i,\max} \quad \forall i\in V_J,\forall t\in T, \tag{10}$$

$$P_{k,t}=\sum_{j\in E_{V_k}^*}\frac{\gamma\Delta_T}{\eta_k}q_{j,t}h_{j,t} \quad \forall k\in F,\forall t\in T, \tag{11}$$

$$S_k\geq P_{k,t} \quad \forall k\in F,\forall t\in T_P, \tag{12}$$

$$S_k\geq \overline{P}_k \quad \forall k\in F, \tag{13}$$

$$W_k\geq P_{k,t} \quad \forall k\in F,\forall t\in T_O, \tag{14}$$

$$W_k\geq \overline{W}_k \quad \forall k\in F. \tag{15}$$

The objective function (1) minimizes the energy cost including energy use and demand charges. Since pump facilities may have multiple downstream control valves, the energy use cost consists of the sum of electricity costs through all facility control valves and all time periods, where the same cost $c_{j,t}$ and efficiency η_j are assigned to the valves that belong to the same facility. The energy demand costs consist of electricity charges associated with power peaks across the facilities that are subject to electricity demand charges.

Constraints (2) model the flow balance at the network junctions, where the flows $q_{j,t}$ are determined through constraints (3) for pipes, (4) for pumps, and (5) for valves. Note that, flow reversals through pipes and valves are handled by using a hyperbolic approximation of the

absolute value of the link flow [3]. The head loss $h_{j,t}$ through link j in time period t is set by constraints (6) using the incidence relationship I_{ij} and the hydraulic heads of the incident nodes i. Constraints (7) prevent flow reversal through pumps. We introduce the differential constraints (8) to account for the tank flow dynamics; note that the formulation is for cylindrical tanks. Constraints (9) and (10) are used to maintain the tank levels and junction pressure heads within admissible ranges. Constraints (11) are used to determine the power use of facility k in time period t for the facilities subject to demand charges. Finally, constraints (12) to (15) are auxiliary constraints used to determine the power peaks throughout the optimization horizon.

2.3 Architecture of the Water Network Optimizer

This section describes the high-level water network optimization software architecture illustrated in Fig. 2. The water network optimizer (WNO) runs on a rolling-horizon basis following a configurable time interval (e.g., hourly). At the start of the process, real-time digital twin boundary conditions, such as tank levels, are obtained and the demand forecaster is activated to provide demand forecasts over the incoming planning horizon. As a next step, the optimization model is created (parameters, variables, constraints, discretization, optimization objective, etc.) using the boundary and demand data. Finally, the model is solved using IPOPT [4] to generate an optimized operational plan that includes control valve settings and provisional pump facility head-flow operation points. These results are used by the pump facility optimizer to generate pump and valve schedules. The subsequent steps repeat the following sequence: (1) update water demand data (from the digital twin services), (2) update the initial boundary conditions for the incoming operational plan, (3) forecast demands, (4) update the optimization model, and (5) solve the optimization formulation to deliver a new operation plan.

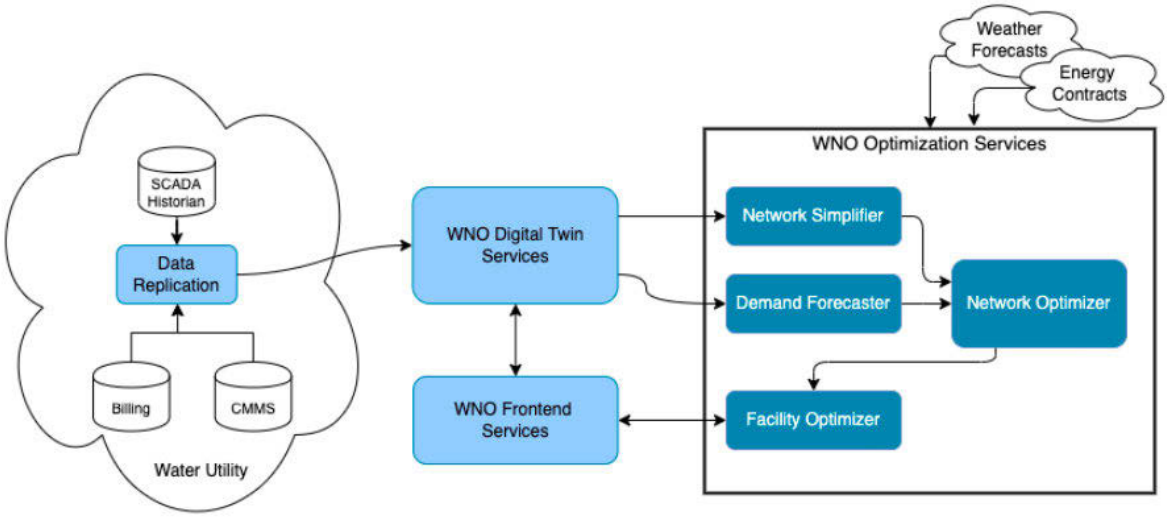

Figure 2. Diagram of the water network optimizer architecture.

3 PROOF OF CONCEPT

The Hillsborough Water Resources Department (HWRD) South-Central system is used as the testing and validation supply system. A modified and anonymized representation of the HWRD network model is illustrated in Fig. 3. The network includes three pump facilities named A, B, and C where all pumps are equipped with VFDs. Facility A accounts for most (64%) of the energy

consumed for pumping operations. Both A and B operate under a time-of-use and demand-charge energy schedule while facility C operates under a fixed-rate schedule. In addition, the HWRD network is instrumented with several pressure sensors denoted in Fig. 3 by the red squares labelled from a through n. Their data was analyzed statistically to determine appropriate minimum and maximum pressure constraints that were used during optimization.

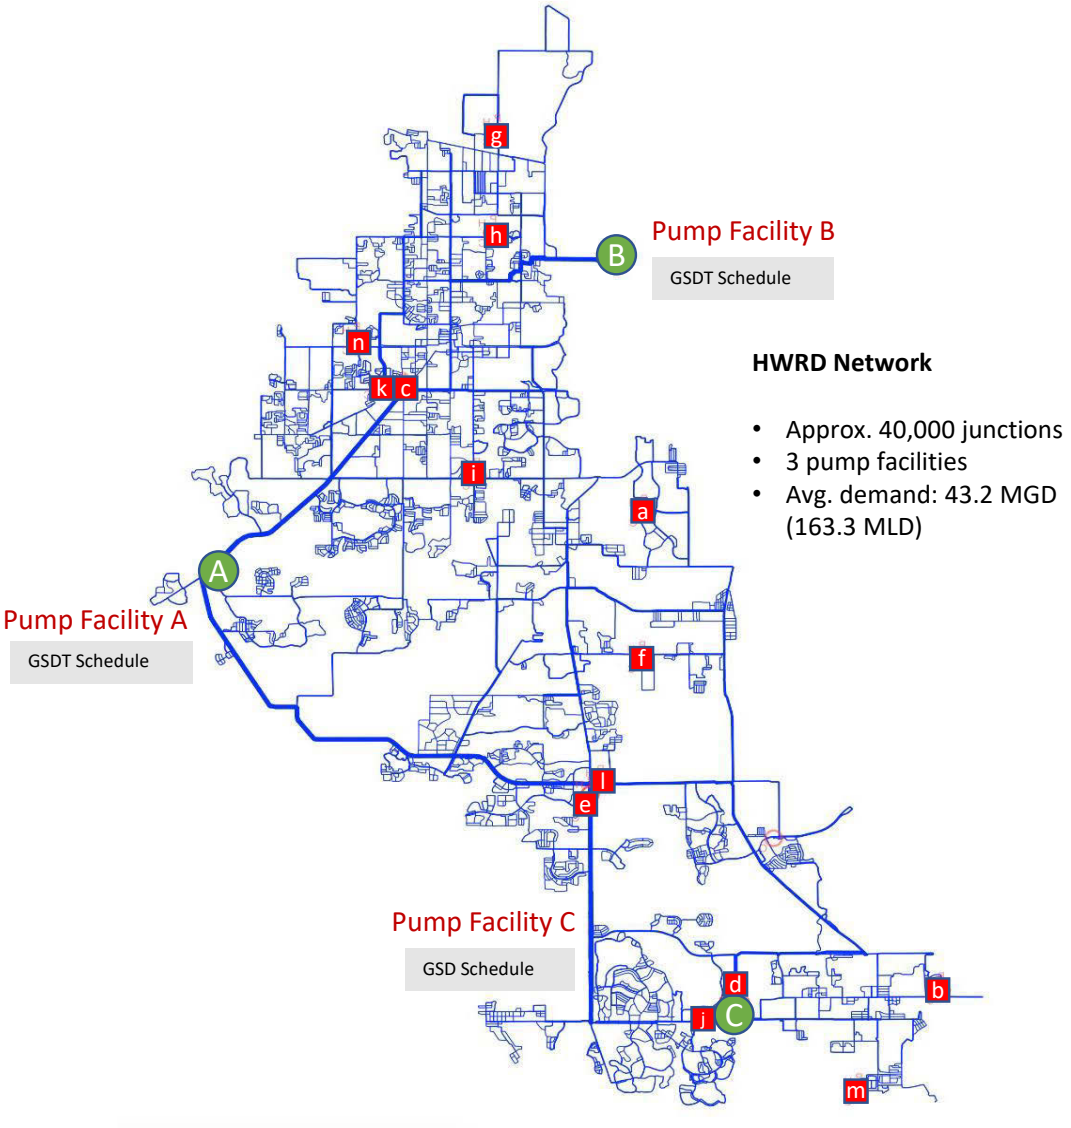

Figure 3. The HWRD network model used in the PoC. Here, the real network has been modified geometrically and anonymized. The time frame of October 2020 is chosen for the PoC due to its interesting characteristics, such as water demand variability and the peak/off-peak demand charge tariff differences (for analysis of energy use and demand costs).

The time frame of October 2020 is chosen for the PoC due to its interesting characteristics, such as water demand variability and the peak/off-peak demand charge tariff differences (for analysis of energy use and demand costs).

A view of the actual network operation output in terms of energy demand and energy use cost is provided in Fig. 4. The energy demand or power input (Fig. 4a) was computed from SCADA measurements of pump facility discharge flow, suction pressure, and discharge pressure. The

energy use cost (Fig. 4b) was computed from the energy demand time series by applying the appropriate energy schedule to each facility according to the energy contracts.

One noticeable feature of the time series of Fig. 4a is the presence of two power peaks during weekdays that are caused mainly by irrigation demands. Another visible feature is that power and energy use are lower during Sundays (Oct 4, 11, 18, and 25), in response to different consumption patterns during weekends with respect to weekdays.

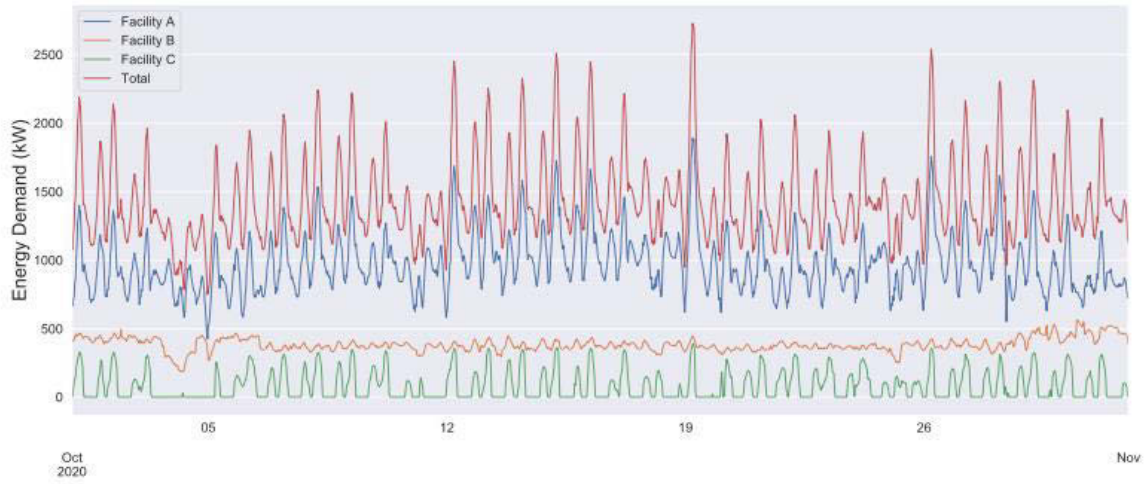

(a)

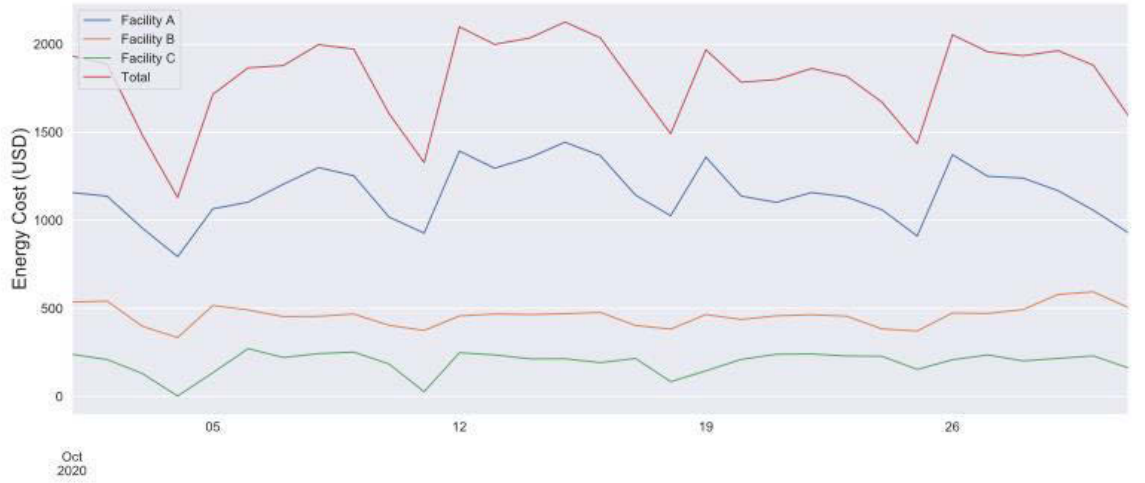

(b)

Figure 4. Actual energy demand and energy use costs for the HWRD system in the month of October 2020. (a) Instantaneous energy demand in KW per pump facility and total calculated from pressure and flow data; (b) Daily energy use cost (USD) per pump facility and total. Facilities A, B, and C account for 64, 25, and 11 percent of the total monthly energy use cost, respectively.

As part of the PoC optimization process, the distribution system was skeletonized by the network simplifier (see Fig. 2) resulting in a hydraulically equivalent network of approximately 400 nodes. Since the overall optimization process was configured to generate a 24-hour operation plan detailed at 30-min time intervals, the simplified network yielded over 21,000 decision variables after the initial optimization model construction performed by the optimizer. Most of these variables (400 x 48 = 19,200) correspond to the hydraulic heads at nodes. Moreover, the optimizer

was configured to run as it would in production over a period of one month starting 2020-10-01 at 00:00 and ending 2020-11-01 at 00:00. The optimizer was setup to update the 24-hour operation plan every hour using current demand forecasts and initial conditions, such as tank levels. Other aspects of the configuration included the service-level constraints such as minimum pressures and tank levels, and the energy use and demand charge schedules. The results consist of 24 x 31 = 744 sets of optimized hydraulic states, energy costs and pump facility head-flow operation points. These results are summarized and described in the next section.

4 RESULTS

In this section we briefly discuss the computational performance and compare the optimization results with the actual operation by contrasting the energy use and costs. We also present a comparison of the levels of service before and after optimization focusing on network pressure and storage.

In terms of computational performance, Table 1 summarizes the CPU time employed by the IPOPT solver for the first and subsequent rolling-horizon iterations. While the solver took approximately 257 CPU seconds (4.28 min) for the first iteration, it took an average 10 seconds for the subsequent iterations. The total processing time including the solver CPU time, WNO service calls, etc., was 281 s (4.68 min) for the first iteration and 29.2 s on average for the subsequent iterations.

Table 1. Network Optimization Computational Performance Summary

Optimization Process Attribute	Value
Optimization output	
Number of variables	21,216
Number of equality constraints	20,971
Number of inequality constraints	1,568
Initial optimization number of iterations	3,315
Subsequent optimizations number of iterations	127
Processing time (s)	
IPOPT CPU time – first iteration	257.54
IPOPT CPU time – subsequent iterations (average)	10.08
Total processing time – first iteration	281.00
Total processing time – subsequent iterations (average)	29.20

Since the optimization objective is to minimize the sum of energy use costs and energy demand charges, it is important to describe how the network optimizer manages power use and power peaks throughout the optimization process. Fig. 5 and Table 2 present a comparison of the main features of energy demand (kW) and energy consumption (kWh) before and after optimization for pump facility A, the largest in the network. Note that we use the terms energy demand and input power as synonyms.

The energy contract for facilities A and B specifies demand charges based on the input power peak over the billing month. The power peaks are computed by averaging the facility input power time series at 30-min intervals and then selecting the maximum values both for peak and off-peak hours. The peak hours for the PoC are 12:00 Noon to 9:00 PM as specified in the utility's energy contract. Since the monthly power peaks cannot be known ahead of time, the optimizer starts from an initial guess of the power peak values and updates these estimates as the optimization

progresses throughout the month. The updating process can be observed in Fig. 5 where power peaks are increasing progressively over time. By the end of the month, after optimization, the peak input power and cost for facility A were reduced by about 10 and 26% during peak and off-peak hours, respectively (see Table 2).

The energy consumption costs are related to the time-of-day charges. The optimizer aims at reducing these costs by combining the use of storage and pumps while respecting hydraulic and level of service constraints, in such a way that more pumping is performed during off-peak hours and more storage is used during peak hours. As shown in Table 2, the on-peak energy consumption and their associated costs were reduced by about 35% after optimization. On the other hand, the off-peak energy consumption and costs were reduced by approximately 23% by the optimizer.

In combination, through power peak reduction and improved time-of-use energy management, the network optimizer was able to reduce the total energy costs by approximately 23% with respect to the actual operation, for pump facility A.

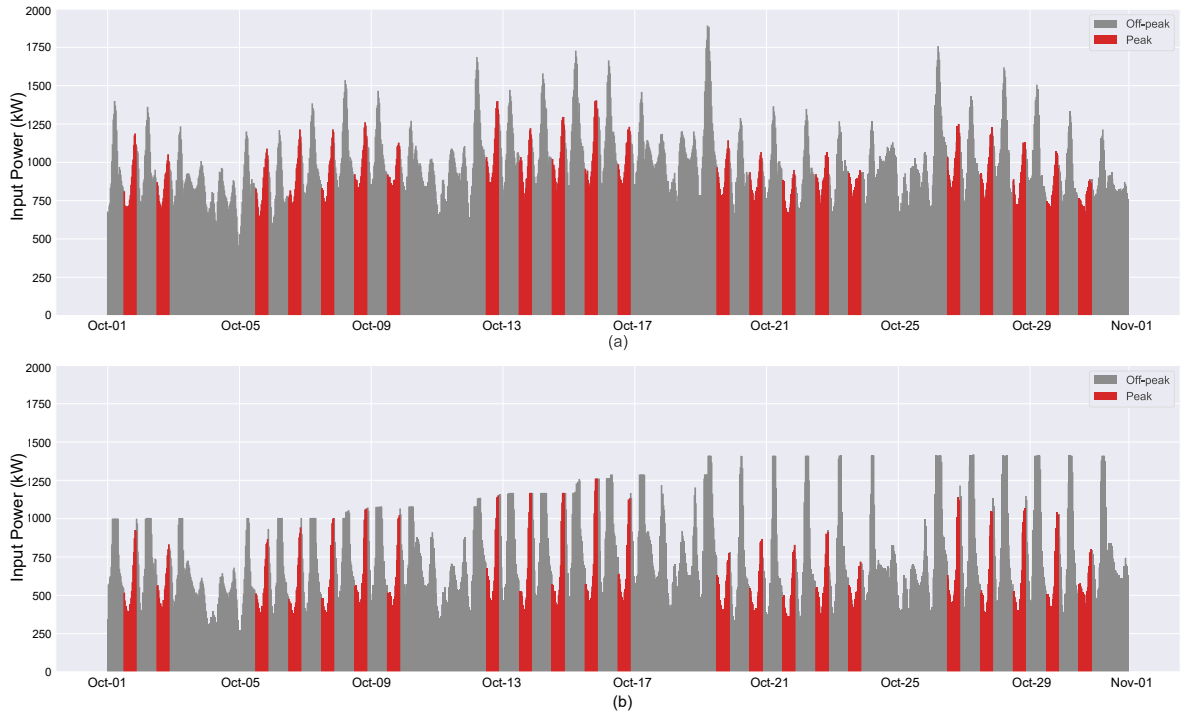

Figure 5. Comparison of input power at pump facility A: (a) Actual input power; (b) Optimized input power. The input power is expressed as a 30-min average in kW. The input power during peak and off-peak consumption hours are represented in red and gray, respectively.

Table 2. *Summary of energy use/cost, input power and demand charges at pump facility A*

Description	After Optimization	Before Optimization	Reductio After Optimization (%)
On-peak peak input power (kW)	1,253.85	1,403.09	10.64
Off-peak peak input power (kW)	1,399.71	1,892.45	26.04

On-peak demand charge ($)	10,206.34	11,421.18	10.64
Off-peak demand charge ($)	6,564.63	8,875.58	26.04
On-peak consumption (kWh)	118,134.06	181,993.83	35.09
Off-peak consumption (kWh)	417,835.20	545,354.73	23.38
On-peak consumption cost ($)	7,697.62	11,858.72	35.09
Off-peak consumption cost ($)	18,342.97	23,941.07	23.38
Total consumption cost ($)	26,040.58	35,799.79	27.26
Total demand charge ($)	16,770.96	20,296.75	17.37
Total charges ($)	42,811.54	56,096.54	23.68

A view of the total costs for the network both before and after optimization is presented in Fig. 6. These costs were computed from the hourly optimized and actual costs. The reduction of energy use costs is clearly noticeable by comparing the actual (purple-dashed line) and the optimized (red line) time series of costs. The actual demand charges were calculated by dividing the monthly demand charge by the 31 days of the analysis month. These costs are constant unlike the optimized demand charges, which are progressively increasing because of the updates that take place during optimization. A small reduction in demand charges was obtained through optimization as shown by the actual (orange dashed line) and optimized energy demand (last value of green-dash-dot line) time series.

To summarize the total daily energy cost savings after optimization, Fig. 7 presents a day-by-day sequence of percent savings obtained after optimization. These values range from 20 to 35%.

Another important aspect of the optimization process is to verify the levels of service in terms of network pressures and storage. The optimized pressures are compared with the pressures before optimization in Fig. 8. Overall, a considerable reduction of pressure was obtained throughout the network and the predetermined minimum pressure head constraint (29 m) was respected in all cases. This reduction of service pressures is responsible for much of the energy use cost savings.

Finally, the differences in water storage in the tanks are presented in Fig. 9. For visualization purposes, the optimized storage plot of Fig. 9a is showing the complete set of rolling-horizon tank levels (light colored lines). The green solid line is constructed by selecting the first time point in the sequence of optimized tank levels over the rolling operation horizon results. This line shows that the network optimizer is filling the tanks when needed but keeping the overall storage relatively low, on average, within the established minimum and maximum head constraints (14.5 and 24 m). On the other hand, the actual tank levels before optimization are shown in (Fig. 9b), where tanks are filled and drained cyclically and filled over the weekend, when energy is cheap, in preparation for the Monday demand peak in early hours. In addition, Fig. 9c shows the demands of the network to allow visualizing the opposite trends of demand versus storage, especially in the case of actual storage management.

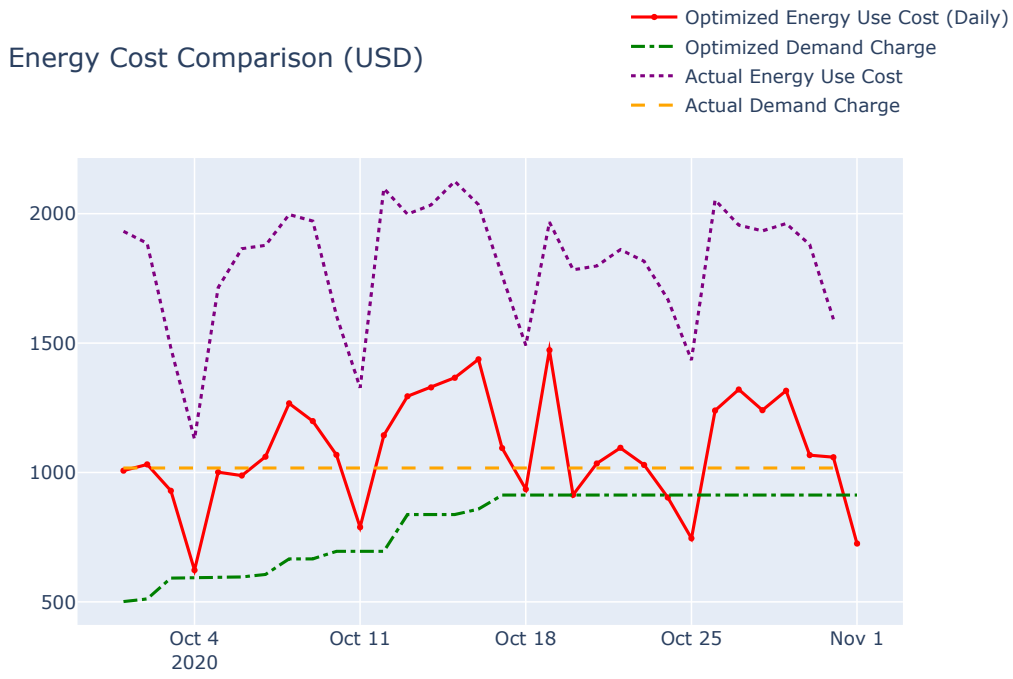

Figure 6. Comparison of total daily energy cost (USD) including use and demand charges. The red solid line corresponds to the daily energy use cost after optimization. The purple-dashed line represents the actual computed daily energy use cost. The green-dash-dot line represents the progressively increasing demand charge (apportioned by day) as computed during optimization; the last value in this series is the optimized daily demand cost. The orange-dashed line corresponds to the actual demand charge (apportioned by day).

5 SUMMARY AND CONCLUSIONS

We proposed a combined strategy designed to address challenges associated with optimizing a water network, namely, the network scale, the nonlinearity of hydraulic relationships, and the discrete nature of operation decisions. The strategy decomposes the problem in two subproblems, one that optimizes the network first, and the other that optimizes the pump facilities using results from the first subproblem. As part of the scope, we presented the NMPC methodology that represents the water supply network as an optimization problem where the state dynamics are modeled as a system of DAEs. The water network optimizer architecture includes the necessary tools that translate the optimization model from a high-level to a low-level and leverages IPOPT, a powerful large-scale NLP solver.

The proposed method was demonstrated as a rolling-horizon optimizer on a PoC where one-month worth of water supply network data was processed to obtain optimized operation policies for control valves and pump facility hydraulic head and flow operation points. The results demonstrate that it is feasible in terms of processing time (under 30 seconds) to run the NMPC optimizations to support the operation decisions of a real water network. The results also show that, for the network evaluated, considerable energy savings can be obtained on most operational days, especially when the supply system is not hydraulically stressed due to water consumption solicitation. In addition, we showed that how the network optimizer respects the level of service constraints including pressure and tank level limits.

Future work includes incorporating water quality in the optimization process, enhancing the scalability and robustness of the optimization process, and further developing and tuning the

pump facility optimizer that works in coordination with the water network optimizer described here.

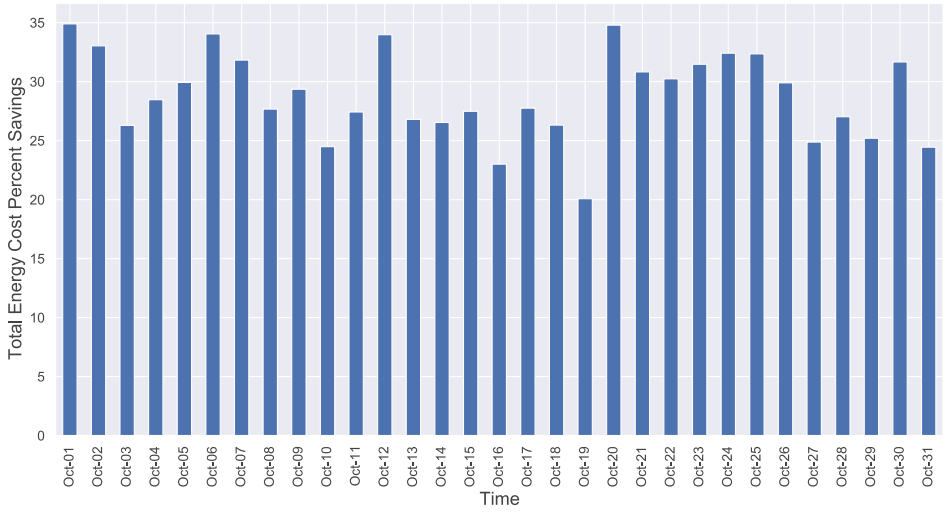

Figure 7. Daily savings of energy use cost calculated from the optimized operation schedules. The savings range from 20 to 35%.

6 ACKNOWLEDGEMENTS

We thank Dr. John McCary, Planning Section Manager and the technical staff at HWRD for providing all of the data necessary to develop the PoC presented in this paper.

7 REFERENCES

[1] Lorenz T Biegler and David M Thierry. Large-scale optimization formulations and strategies for nonlinear model predictive control. IFAC-PapersOnLine, 51(20):1–15, 2018.

[2] Jacek Blaszczyk, Krzysztof Malinowski, Alnoor Allidina, N Callaos, T Gill, and B S´anchez. Optimal pump scheduling by non-linear programming for large scale water transmission system. Proc. CCISE, pages 7–12, 2013.

[3] Ajit Gopalakrishnan and Lorenz T Biegler. Economic nonlinear model predictive control for periodic optimal operation of gas pipeline networks. Computers & Chemical Engineering, 52:90–99, 2013.

[4] Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Mathematical programming, 106(1):25–57, 2006.

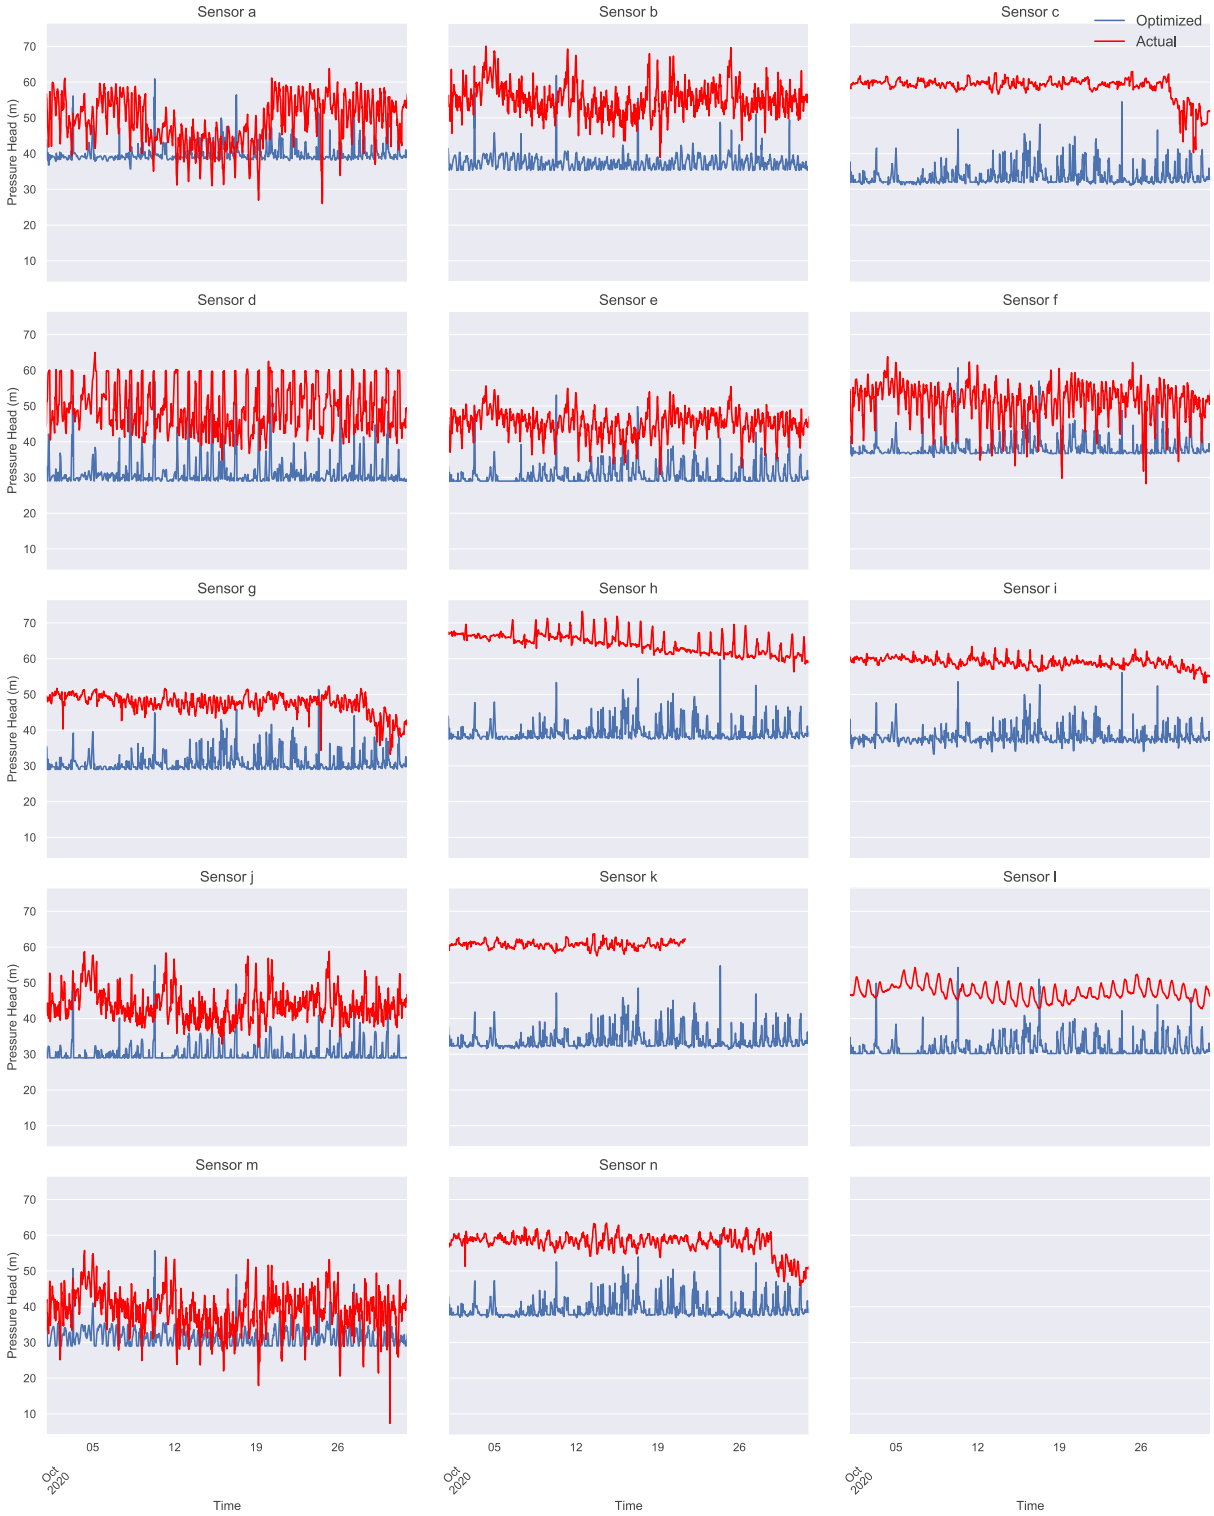

Figure 8. Comparison of optimized and actual pressure heads (m) at network junctions where pressure sensors are available. The optimization constraints at these junctions specify that pressures must be above 29 m (41 psi). The purpose of this plot is to illustrate how the levels of service in the form of pressure constraints are being applied and respected.

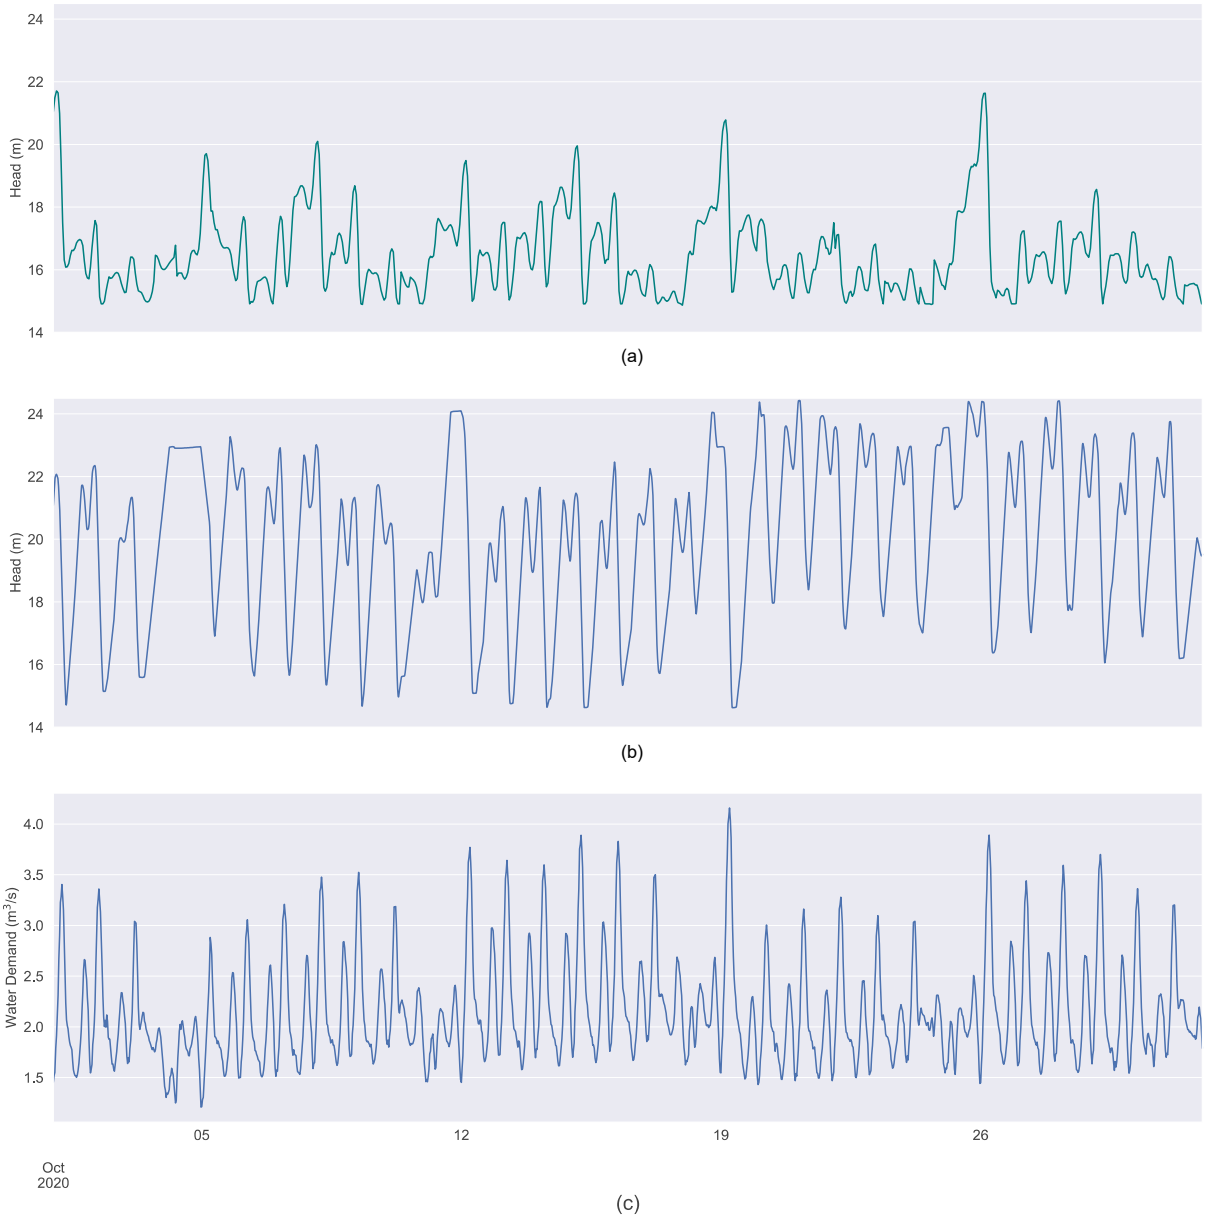

Figure 9. A comparison of actual and optimized storage management. The actual operation consists of filling the tanks during the weekend when energy is cheap in preparation for the Monday demand peak in early hours. On the other hand, the network optimizer is filling the tanks when needed but keeping the overall storage lower within limits.