

Water Pump Operation Optimization under Dynamic Market and Consumer Behaviour

Carmen Cheh
Illinois Advanced Research Center at
Singapore Ltd
Singapore, Singapore
carmen@iarcs-create.edu.sg

Justin Albrethsen
Singapore University of Technology
and Design
Singapore, Singapore
justin_albrethsen@sutd.edu.sg

Zhen Wei Ng
Illinois Advanced Research Center at
Singapore Ltd
Singapore, Singapore
zhenwei.ng@iarcs-create.edu.sg

Binbin Chen
Singapore University of Technology
and Design
Singapore, Singapore
binbin_chen@sutd.edu.sg

Xin Lou
Singapore Institute of Technology
Singapore, Singapore
Illinois Advanced Research Center at
Singapore Ltd
Singapore, Singapore
lou.xin@singaporetech.edu.sg

Zaki Masood
Singapore University of Technology
and Design
Singapore, Singapore
zaki_masood@sutd.edu.sg

David K.Y. Yau
Singapore University of Technology
and Design
Singapore, Singapore
david_yau@sutd.edu.sg

ABSTRACT

In the face of growing energy and water consumption, the pumping costs of water supply systems in high-rise buildings are on the rise. The state of practice uses statically configured water level thresholds or time-based triggers to activate water pumps, while state-of-the-art research works propose to minimize pumping costs by dynamically adjusting the pump schedules. However, the implications of volatile energy price, dynamic consumer water demands, and other important factors - in particular, the impact on water pump health and the disturbance to residents by activating pumps during the night - have not been thoroughly considered in those research works. There is also a lack of thorough evaluation of their performance using real-world data over a prolonged period. Our work addresses those gaps by introducing a model predictive control optimization framework that incorporates machine learning predictions to handle water demand and energy price uncertainty. It combines multiple factors including pump health and resident satisfaction level to find an optimal solution. We used real-world data over prolonged periods of time that exhibit significant pattern changes to evaluate the performance of our dynamic scheduling solution. While significant gain is achieved over state-of-the-art and state-of-the-practice solutions, we also observed

considerable amount of fluctuation in performance gains of such dynamic schemes, especially under varying prediction accuracy of water demand and energy price forecasting, which calls for more future research.

CCS CONCEPTS

• **Computing methodologies** → *Modeling and simulation*; • **Computer systems organization** → **Embedded and cyber-physical systems**.

KEYWORDS

water supply system, optimization, forecasting

ACM Reference Format:

Carmen Cheh, Justin Albrethsen, Zhen Wei Ng, Binbin Chen, Xin Lou, Zaki Masood, and David K.Y. Yau. 2024. Water Pump Operation Optimization under Dynamic Market and Consumer Behaviour. In *The 15th ACM International Conference on Future and Sustainable Energy Systems (E-Energy '24)*, June 04–07, 2024, Singapore, Singapore. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3632775.3661969>

1 INTRODUCTION

In many buildings, the water supply system consists of a ground-level or basement water storage tank that transfers water to a rooftop tank where the water is held and distributed to the residents on demand. In recent years, dynamic energy pricing is becoming more prevalent, in part due to the integration of renewable energy sources [23, 34]. However, the current state of practice uses statically configured water level thresholds or time-based triggers to activate the water pumps. Combined with the dynamic behavior of water usage patterns that varies with seasons, weather conditions, amongst other factors, existing statically-configured scheduling

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

E-Energy '24, June 04–07, 2024, Singapore, Singapore

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0480-2/24/06

<https://doi.org/10.1145/3632775.3661969>

of water pumps may lead to substantially and unnecessarily high energy bills associated with pumping water during periods of high energy prices. This cost burden will eventually be borne by the end users and this sub-optimal practice also affects the effective use of renewable energy. In the case of an average high-rise building with eighteen floors and over a hundred apartment units, our case study estimates the pumping energy cost to be around USD\$1,000 per month. Many works [22, 38] have proposed optimization approaches to reduce that energy consumption cost. If an optimization approach can achieve, say 20% cost savings, for a city of 10,000 such buildings, that would mean a cost savings of around 2 million USD per month.

Given the increased deployment of smart water consumption meters which allow water consumption to be collected on a regular basis [8] and the availability of dynamic energy pricing data that is updated hourly, we can leverage existing machine learning and deep learning techniques to forecast water demand and energy prices. However, the usage of such predictions in dynamic pump scheduling is far from trivial since prediction errors can lead to sub-optimal decisions that cause high energy cost, increased wear and tear in the pump, disruption to water supply, and/or increased resident dissatisfaction due to undesirable pump activation during quiet hours.

Some recent research efforts propose to minimize pumping costs by dynamically adjusting the pump schedules. Their designs, however, have the following limitations that may restrict their applicability to real-world systems: (1) the simplistic assumption of fixed daily energy price profiles [22, 28, 38], (2) the simplistic assumption of the distribution of prediction errors (in terms of energy price or water demand), and (3) the lack of holistic optimization targets, in particular, the lack of modeling the operation impact on residents. While the primary focus of existing optimization works is to minimize pumping costs and ensure the health of transfer pumps, they often overlook the crucial aspect of resident satisfaction. In residential buildings, the activation of transfer pumps during the night can lead to disruptive noise, causing discontent among residents [1, 2]. As such, it is important to optimize the pump schedule such that the pumps are activated as little as possible during the night.

In addition, assuming a fixed energy price profile is not realistic in today's world because of (1) the increasing shift to a dynamic energy market where consumers pay different prices based on the real electricity price and (2) the integration of renewable energy solutions that can offset the energy price depending on the amount of energy generated during different times of the day. Finally, there is also a lack of thorough evaluation of the performance of optimization solutions under real-world conditions where the energy price profile and water demand is unknown.

To overcome the above limitations, this paper proposes a comprehensive solution that integrates machine learning (ML)/deep learning (DL) prediction models into a Model Predictive Control (MPC) optimization approach. The MPC optimization approach uses a dynamic model to forecast system behavior and optimizes the forecast to produce the best control solution at the current time [29]. It is designed to balance long-term planning with reducing forecasting uncertainty. To address the lack of holistic optimization targets, our objective function incorporates three key elements: (1) pumping cost, (2) the pump health, and (3) resident satisfaction.

By incorporating these elements into our objective function, we aim to achieve a more comprehensive and balanced approach to water supply system optimization. We design the objective function such that the water pumps can be turned on for a variable amount of time instead of mandating that the pumps must be kept on until the high level water threshold is met. By doing so, we can avoid pumping too long during periods of high energy price while maintaining the stability of the system. In this manner, we consider a larger solution space for pump schedules, thus allowing for a more adaptive and precise water supply system management. The potential pitfall for such a solution is the increased amount of pumping activations, however, this can be balanced in our optimization formulation with the proper weighting of the pump health against the pumping cost. Subsequently, we use ML/DL models to forecast energy prices and water usage demands using the previous day's historical data. To address the uncertainties introduced by forecasting errors, we develop a flexible optimization approach that brings forward the pumping schedule when the water level in the tank falls dangerously low.

Many prior works assess their optimization solutions by testing on only one (or a few select) days of water demand and energy pricing data [37, 38]. We recognize that short-term evaluations might not sufficiently justify the performance of an optimization solution as water demand and energy pricing fluctuate across weeks. Instead, we make use of available real-world energy price and water demand data to conduct systematic experiments over prolonged periods of time that exhibit significant pattern changes in terms of both amplitude and volatility. By utilizing forecasted water demand and energy prices based on this substantial dataset, we can more accurately assess the effectiveness and robustness of our optimization approach. Additionally, unlike other works that model uncertainty in forecasted water demand by generating perturbations, we utilize real-world forecasting errors to provide a more realistic evaluation of our solution's performance under different forecasting models with varying accuracy.

Our contributions in the paper are summarized as follows:

- We propose an optimization formulation that aims to minimize pumping cost while maintaining pump health and resident satisfaction. Our formulation is designed to produce a finer-grained pump schedule that allows the pump to be turned on for a variable amount of time, thus enabling us to search a larger solution space for a more optimal cost-saving schedule. We account for uncertainties introduced by water demand forecasting errors by bringing forward the pump schedule when the water in the system is dangerously low.
- We apply our solution to a university building pumping system case study with real-world water demand and energy pricing data. Our results show that we can achieve a 37% cost savings as compared to the standard control algorithm currently used in the case study pumping system when we use the best performing energy price and water demand forecasting models. We also achieve a 73% decrease in pumping duration during the night when residents are asleep, which improves resident satisfaction.
- We conduct a comprehensive comparison of our solution against two other optimization approaches proposed within

the same problem domain. Over time, our solution consistently outperforms the competitors in terms of reducing energy costs, pump active duration during the night, and the number of pump activations. As compared to the other optimization approaches, our solution is more robust in the face of water demand and energy price forecasting errors. However, the prediction accuracy of the water demand and energy price plays a huge role in the final pumping costs achieved by the optimization solutions. As such, we evaluate the effect of water demand loads and forecasting errors on our optimization solution. Our results show that in general, the higher the accuracy of the forecasting model, the higher the cost savings achieved with optimization. However, that is not always the case under different water demand loads and volatile energy prices. So the optimization solution is very volatile under a dynamic energy price profile.

The rest of this paper is organized as follows. We describe the related works in Section 2. Section 3 describes the water pump system. Sections 4 and 5 details our optimization formulation and how we integrate ML/DL prediction models into our MPC optimization approach. Section 6 discusses the experimental results of applying our optimization framework on a university building pumping system case study with real-world energy price and water demand data. Finally, we conclude in Section 7.

2 RELATED WORK

Many existing works develop control frameworks to optimize pump schedules with the aim of decreasing pumping costs [20, 26, 28]. For example, Quintiliani and Creaco [28] adjusted the water level thresholds triggering the activation of the pumps in such a way that the pumps are turned on during low energy price time slots. These additional time slots help to drive water levels close to the highest and lowest acceptable level at the beginning and end of the peak tariff period, respectively. The major drawback to these approaches are that they do not consider time-varying water demand and energy prices, which is unrealistic as we will see in the datasets that we used in Section 6.1.

There are many works now that explicitly model the changing water demand as a fundamental part of the optimization formulation [15, 37, 38]. The authors in [37] proposed an economic MPC framework that controls the pumps in real-time for a water distribution system. At each timestamp, they use the forecasted water demand to solve an optimization problem minimizing pumping cost and use the solution for the control policy at the next time step. Similarly, Wanjiru et al. [38] explored the use of an open-loop and closed-loop control model to minimize the urban consumer's energy cost for pumping water from the ground storage tank to the rooftop tank. However, all those approaches assume a fixed daily energy price profile, which is not practical in today's changing energy market. On the other hand, some works such as [19] define energy price classes and compute an optimal pumping schedule for each of those energy price classes. However, the authors do not consider water demand in their optimization formulation.

Another major limitation of those approaches is that they do not explore the volatility of the optimization results when prediction of water demand and/or energy prices are introduced into the

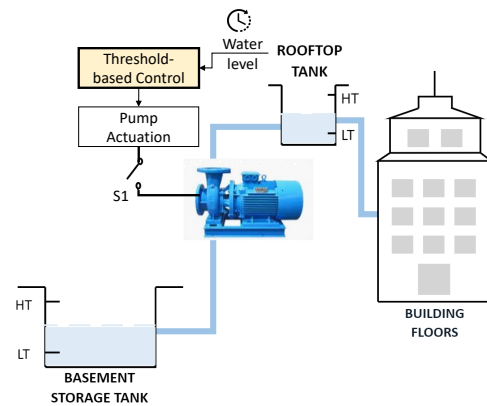


Figure 1: System model of the domestic water supply system.

system. Stuhlmacher and Mathieu [33] offer a solution to tackling prediction uncertainty by explicitly modeling the prediction errors of water demand and power demand. Once the forecast errors are known, the control policies are adjusted to correct for those errors while ensuring that the system constraints are adhered to. The authors' main focus, however, is optimizing the water distribution network's electricity cost by controlling the supply pump flow rate and tank net outflow. Sopsakis et al. [32] models the prediction uncertainty using a scenario-based stochastic approach and uses concrete forecasting models (ARIMA for energy price and SVR for water demand) and thus, their work constitutes a first step towards understanding how actual prediction models may affect optimization results. However, their solution applies to a water distribution system and only involves the evaluation of one week's worth of data. Our paper, on the other hand, is focused on the water supply system in buildings. Furthermore, we experiment with different concrete forecasting models that can be used to obtain real-world predictions instead of relying on statistical models of perturbations. Finally, our approach is the first work that explicitly models resident satisfaction and evaluates optimization solutions based on the time in a day when the pump is activated. In Section 6, we will perform a comprehensive comparison of our approach against the works of [28] and [38] because those two works are designed for similar water pump systems as our case study.

3 WATER PUMP SYSTEM AND CASE STUDY

In this paper, we consider domestic water supply systems that provide water to building residents. Water from the main supplier is stored in a basement tank, from which it is pumped to a rooftop tank for storage under high pressure. Upon demand, the water is then distributed to residents. Most water supply systems utilize a simple control algorithm to transfer water from the basement tank to the rooftop tank. Those systems rely on the deployment of water level sensors in the two tanks as shown in Fig. 1. Those coarse-grained sensors can detect whether the water level in the tank exceeds or falls below a given threshold. The standard control algorithm turns the transfer pumps on when the water level in the rooftop tank reaches the low level threshold and only turns the pumps off when the rooftop tank reaches the high level threshold.

However, such an algorithm is agnostic of the dynamic pricing of today's energy market, i.e., the transfer pumps are often turned on at a time when energy prices are high. We aim to optimize the pumping schedule to achieve energy cost savings while ensuring the stability of the system and maintaining resident satisfaction. Such a goal is achievable due to: (1) the availability of real-time data for energy pricing and water consumption meter readings, and (2) the low level of granularity of water level readings by more fine-grained sensors.

In this paper, our study is demonstrated on a water supply system in a seven-storey residential university building. Water is supplied to the floors via a two metres-tall rooftop tank with a three metres diameter. When the water level in the rooftop tank falls below 0.4 metres, the pump in the basement is activated to pump water from the basement storage tank to the rooftop tank. The pump is stopped when the water level in the rooftop tank reaches 1.6 metres.

Based on the physical dimensions of the water supply system, we constructed a simulation model using EPANET [6], a widely used water distribution system simulator. Our model is shown in Fig. 2 where each component of the water supply system is modeled as a vertex. Each edge between the vertex represents a pipe connecting the components. Properties of the component (e.g., tank level, pipe length) are input using EPANET's GUI. Finally, the water demand profile (at the Consumers node) and the energy price profile are modeled using EPANET's patterns. We use the EPANET model to (1) provide real-time water level updates and (2) simulate the effect of using different control algorithms on the system. Note that, while our work focuses on this particular building water pump system, it can be effortlessly adapted to other building water pump systems.




Figure 2: EPANET model of the university pumping system.

4 OPTIMIZATION PROBLEM FORMULATION

We propose an optimization problem formulation which calculates the optimal pump schedule that minimizes energy consumption cost and maintenance cost while incorporating the resident satisfaction and adhering to system constraints. The problem formulation takes in as input the real-time water level sensor values, as well as dynamic energy pricing and water demand data.

We perform optimization over a time horizon of T minutes. To improve the scalability of the solution, we divide the time horizon into sampling periods t_s so that the total number of samples is N . Then, we model the pump schedule as a collection of decision variables: $(p(i), y(i))$ where $i \in [0, N]$. The decision variable $p(i)$ represents the number of minutes that the transfer pump is on during the i -th time period whereas the binary decision variable $y(i)$ indicates whether the transfer pump is turned on at any point during the i -th time period. We define $y(i)$ so that we can constrain the number of times that the transfer pump is turned on, which affects the pump's health and resident satisfaction.

The two decision variables representing the pump schedule are related via the following constraints:

$$y(i) \leq 1, \quad i \in [0, N] \quad (1)$$

$$0 \leq y(i) \leq p(i) \leq M \times y(i), \quad i \in [0, N] \quad (2)$$

where M refers to the maximum number of minutes that the pump is allowed to be on (which is the amount of time that it takes to pump water into the rooftop tank such that the water level increases from the low threshold (AL) to the high threshold (AH)). Those two constraints ensure that when the pump is turned on at time period i , i.e., $y(i) = 1$, then the pump must be on for a minute or longer, i.e., $p(i) \geq 1$. Similarly, if the pump is off during time period i , i.e., $y(i) = 0$, then the pump must not be on at all, i.e., $p(i) = 0$.

To maintain the stability of the system such that the rooftop tank does not overflow or drain out, we set constraints to ensure that the water level in the rooftop tank remains within the low and high level thresholds. We assume that the water demand $D(i)$ is known at time period i . When the pump is turned on, the flow rate is a constant $Q \text{ m}^3$ per minute and the total area of the tank is A . Then, we can formulate the following constraint for the water level in the rooftop tank at a given time i as

$$AL \leq h(0) + \frac{1}{A} \sum_{k=0}^i (Qt_s p(k) - D(k)) \leq AH, \quad i \in [0, N] \quad (3)$$

where $h(0)$ is the initial water level of the tank.

Finally, we learned from a discussion with city authorities in Singapore that residents are dissatisfied with water pumps being activated during the night. Thus, to improve resident satisfaction, we do not want the pump to operate during the night when people are asleep. As such, we define G to be a set of time periods where we do not want the pump to operate. Then, we define the following hard constraint:

$$\sum_{i=1}^{|G|} y(g_i) = 0. \quad (4)$$

We also add the hard constraint (5) to ensure that the water tank is full at the end of the time horizon, which will ensure that the pump is not forced to activate in the next time period (which may be a period of high energy prices):

$$h(0) + \frac{1}{A} \sum_{k=0}^N (Qt_s p(k) - D(k)) = AF, \quad (5)$$

where AF is the height of the water tank. The hard constraints (4) and (5) regarding resident satisfaction and water level at the end of the day may result in (1) the optimization problem being intractable or (2) a high energy consumption cost. Thus, we enforce them as soft constraints instead.

Our objective function serves to minimize the energy consumption cost of operating the transfer pumps while improving resident satisfaction and reducing the cost of wear-and-tear from turning the pumps on frequently. We include wear-and-tear cost in our objective function because it is important to consider the long-term effect of increasing the frequency of pumping as it will lead to increased maintenance and repair cost. Those goals are often conflicting with each other (e.g., minimizing the cost may result

in the pumps needing to work in the night, which decreases resident satisfaction) and, thus, need to be weighted according to their importance to the stakeholder. Thus, we formulate the objective function as

$$\begin{aligned} \min_{p,y} \quad & \alpha \sum_{i=0}^N C(i)p(i) + \beta \sum_{i=0}^N y(i) + \gamma \sum_{i=0}^{|G|} y(g_i) \\ & + \delta(|AF - h(0) - \frac{1}{A} \sum_{i=0}^N (Q_{t_s}p(i) - D(i))|) \\ \text{s.t. } \forall i \in [0, N] \quad & y(i) \leq 1, \\ & 0 \leq y(i) \leq p(i) \leq M \times y(i), \text{ and} \\ & AL \leq h(0) + \frac{1}{A} \sum_{k=0}^i (Q_{t_s}p(k) - D(k)) \leq AH, \end{aligned} \quad (6)$$

where $C(i)$ is the energy price at time period i and $\alpha, \beta, \gamma, \delta$ are weights. The weights are assigned by stakeholders so that the optimization solution will prioritize certain components of the objective function over others. The first component of the objective function refers to the energy consumption cost, the second refers to the wear-and-tear cost (or pump start-up cost), the third refers to the night-hours cost, and the fourth refers to the low water-level cost.

5 OPTIMIZATION IN REAL-WORLD SETTING

The optimization problem requires knowledge of the energy prices and the water demand in real-time and needs to be solved in an online manner. To achieve that, we train ML/DL models and use the trained models to predict energy prices and water demand using recent historical data. Then, we use an online MPC optimization framework that incorporates those predictions, as depicted in Fig. 3. Finally, we describe how the optimization solution can be implemented in control centers to decide when the transfer pumps will be activated.

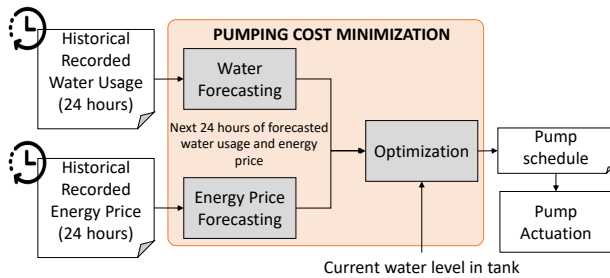


Figure 3: Workflow of our optimization solution.

5.1 Water Demand and Energy Price Forecasting

Water demand forecasting is widely studied, with many approaches utilizing traditional time series models like the Auto-Regressive Integrated Moving Average (ARIMA) [18, 30, 39], dynamic deep neural networks (DNNs) [21], or Gated Recurrent Unit (GRU) [31], among other models. A number of the water demand forecasting approaches also take into account additional contextual data such as sunny hours, population, seasonal trends, temperature, and precipitation [13, 18, 27, 30]. Similar to water demand forecasting,

traditional statistical methods [34–36, 40] and deep learning models (LSTM, GRU, and CNN) [16, 24] can also be applied to predict energy prices.

In this paper, we implement various forecasting approaches to predict the water demand and energy price for the next 24-hour time period. For the DL models that we implement, we input a sliding window of the previous 24 hours of historical data (water consumption or energy prices), day of the week (1 to 7, where 1 denotes Sunday), and hour number (0 to 23, where 0 denotes the time between 12:00am and 1:00am). The day of the week and hour number features help to capture information about the expected building occupancy (e.g., we expect higher water consumption on a Sunday night). We evaluate the impact of the accuracy of the various forecasting approaches on the optimization results in Section 6.

5.2 Online MPC Optimization

Since our forecasting algorithms predict the water demand and energy price up to the next 24 hours, we restrict the time horizon of our optimization problem to a day (or 1440 minutes). The straightforward way to integrate the forecasted data into the optimization process is to directly input the 24-hours prediction. However, that approach does not take into account real-time updates of water consumption and energy prices, which can improve the accuracy of predictions for later hours of the day.

We adopt a MPC approach [11] where optimization is performed over a long time horizon (24 hours) but only the first optimization result (i.e., the optimal schedule for the first time period) is used to control the system (Fig. 4). Then, in the next time period, the forecasting algorithm is repeated with the knowledge of the latest water consumption and energy price data. The predictions are fed into the optimization algorithm to determine the optimal schedule.




Figure 4: MPC optimization with a time horizon of 24 hours. Each box represents an hour in a day. After one round of optimization, the first hour (box 1) is colored yellow to indicate the pumping schedule is implemented for that hour. The “Prediction” window then slides by one hour to take in the historical input data and predict the following 24 hours of water demand and energy price data.

That optimal schedule informs the system that the pump needs to be active for a certain duration t during a time period. Since the sampling time period t_s is longer than t , we need to determine the exact time when the pump needs to be turned on during that time period. There are two considerations for determining when to turn the pump on: (1) ensuring minimal energy cost and (2) adhering to the anticipated pumping duration of t minutes. Since we define $C(i)$ to be the energy price at time period i , that means the energy cost is constant throughout the sampling time period. Then, we only

```

1 Timer.schedule.Start := t_s - t
2
3 IF water_level <= AL + 0.01 THEN
4   Pump.On := TRUE;
5   schedule.Start := t_{now};
6 END_IF;
7
8 IF NOT schedule.Stop AND schedule.Elapsed THEN
9   Pump.On := TRUE;
10 END_IF;
11
12 IF schedule.Elapsed THEN
13   Pump.On := FALSE;
14 END_IF;

```

Figure 5: Pseudocode of control center logic.

need to determine when to turn the pump on so that the pump can remain active for t minutes without violating system constraints.

If the pump is turned on at the start of the time period, the water level may reach the high water level threshold AH before the end of t minutes. That means the pump is only active for less than t minutes. On the other hand, if we pump at the end of the time period, the water level will be at its lowest and, thus, the pump can achieve the maximum active pumping duration. So we choose to activate the pump t minutes before the end of the time period.

However, the water demand forecasting may predict a lower water demand that results in the water level falling below the low water level threshold before the pump is active. Thus, we need to implement a fail-safe measure that initiates the pumping earlier than the scheduled time if the water level falls dangerously low, i.e., the scheduled pumping at the end of the time period will be brought forward. To achieve that, we can implement the following pseudocode (shown in Fig. 5) as a script in the control center. The control center will run the optimization every hour to obtain the optimal duration that the pump should be turned on at the end of the sampling time period. When the control center receives sensor values in real-time, it checks whether the water level reading is near the low level water threshold AL . If so, then it sends a start command to the pump. Otherwise, it waits until the end of the sampling time period ($t_s - t$) to send the start command. When t time has elapsed after the start command, a stop command is sent to the pump. In case of forecasting errors or the event that the network between the control center and the pump is disconnected, the local controller at the pump still executes the default threshold-triggered control, hence making sure that the system constraints are not violated. Thus, our solution is easy to integrate into water supply systems with no change required to the pump controller.

6 EXPERIMENTS AND PERFORMANCE EVALUATION

In this section, we evaluated our optimization solution using the university building pumping system case study introduced in Section 3 and compared the performance of our solution against various benchmarks. To solve our proposed optimization problem, we used the CVXPY library [10, 17] with the GLPK_MI solver [7]. We also used the Water Network Tool for Resilience (WNTR) [9], a Python library that interfaces with EPANET, to simulate the effects of using different pump control algorithms. We modified parts of WNTR's source code in order to calculate the energy consumption cost for the pumps for the duration of the simulation.

6.1 Datasets for Forecasting Models

To model the water consumption for our case study, we used the hourly water consumption dataset (in gallons) collected by Bejarano et al. [12] from different university buildings August 1st, 2018 to December 8th, 2018. Since the university buildings encompass facilities such as laboratories and classrooms, we specifically chose the buildings that are purely residential so that it follows a similar demand pattern as our case study. We used the models proposed in [12] (LSTM, GCRF, ARIMA, and LR) and [31] (GRU), with several updates to the model structure and hyperparameters to enhance the accuracy on our dataset. All the models are trained on water demand data of each specific building, with the LSTM, GRU, and GCRF models being trained on water demand data that was augmented with temporal features (day number and hour). The first 75% of the data was used to train the models while the remaining 25% was used as the test set to evaluate the model and select the optimal hyperparameters (where needed).

In our experiments, we considered two types of energy pricing scenarios: (1) fixed energy price profile with the price during the day (7am-6pm) being USD\$90/MWh and the price during the night (7pm-6am) being fixed at USD\$301/MWh, and (2) dynamic energy pricing that changes every half-hour interval. We considered the fixed energy price profile for fair comparison against other optimization solutions. The fixed energy price profile was obtained from the city authority in Singapore. For the dynamic energy pricing scenario, we obtained energy pricing data over two years (Jan 2022 to Dec 2023) from the Energy Market Company (EMC) in Singapore [5]. The dataset comprises information on the Wholesale Electricity Price (WEP) of every half-hour interval. The WEP is the net purchase price paid by retailers, which includes administrative costs and uplift charges in addition to the energy price. We used the WEP as the ground truth energy consumption cost in our simulation experiments. We used the stock price forecasting models in [25] (CNN-LSTM) and both the top performing energy price forecasting models (LSTM, GRU) and typical machine learning models (RF) in [24], with several updates to the model structure and hyperparameters to enhance accuracy on our dataset. All models are trained with energy price data that was augmented with temporal features. We trained our models using data from 2022 and validated its performance with data from 2023.

6.2 Experimental Setup

For our optimization experiments, we used a duration of a month (or 30 days) to compare the performance of our solution against other approaches. We ran our optimization framework in tandem with the EPANET simulator to simulate the effect of using our optimized control algorithm on our case study system.

We selected the last month of water demand data from one of the university residential building as the testing set, shown in Fig. 6. The data for other residential buildings are similar. For comparison, we chose two different months worth of energy pricing data in 2023 (shown in Fig. 7) that represent the typical price patterns — one where the prices are lower and more stable (Fig. 7a) and one where the prices are volatile and have high fluctuations (Fig. 7b).

City authorities typically recommend that the water tank be sized so that it accommodates a day's worth of water consumption, so the

pump will only activate once in a day [14]. However, the average daily water consumption can increase due to hot weather seasons or global events such as pandemics [3, 4]. To model the effect of the optimization on systems with different water demand loads and tank capacity, we experimented with different scaling of the water demand data so that the average number of pump activations in a day varies between 0.8–3 times. Each pump activation takes 6.5 minutes to pump from the low water threshold (0.4m) to the high water threshold (1.6m). We assigned the initial water level, W_0 , to 1.5m which is near the high water level threshold (1.6m) of our system’s rooftop tank (which has a total volume of 14m^3).

For our objective function, we assigned the weights as follows: the energy consumption cost weight ($\alpha = 1$), the pump start-up cost weight ($\beta = 1$), the night-hours cost weight ($\gamma = 1$), and the low water-level cost weight ($\sigma = 0.2$). We ran our algorithm and simulations in a container on a Linux machine with a 3 GHz CPU core and 32 GB of RAM.

6.3 Performance Evaluation

The performance of the optimization algorithms are measured by three metrics: (1) pumping energy cost, (2) the number of pump activations, and (3) the duration that the pump is active during the night (which we define as the period from 11pm to 6am). In this section, we (1) demonstrate the necessity of different components of our objective formulation, (2) explore the system settings and forecasting approaches under which the optimization solution yields the best cost savings, and (3) compare the solution achieved by our approach against other state-of-the-art solutions. For all our experiments, we compare the performance with that of the standard control algorithm, which is the existing solution deployed in the building system that activates the pumps upon reaching the low-level water threshold (AL).






Figure 6: A month’s worth of water demand data in a residential building on campus.

6.3.1 Ablation Study. First, we performed an ablation study to determine the necessity of each cost component in the objective function. We ran our MPC optimization with ground truth knowledge about energy price and water demand. Table 1 shows that although the energy cost decreases by up to 3% when certain components of the objective function are removed, we see that the number of pump activations increases by 8% and 26% when the low-water cost and startup cost is removed respectively and the duration that the pump is active during the night increases by 73% when the night-hours cost is removed. The rise in pump activations when the startup or low-water cost is removed will adversely affect the pump health whereas the increase in pumping during the night when the night-hours cost is removed will result in decreased resident satisfaction.



(a) August energy pricing. Note that there is a sharp rise in energy price around the 600-th period which goes up to \$450 / MWh.



(b) May energy pricing.

Figure 7: A month’s worth of energy pricing data in 2023.

Table 1: Evidence of the necessity of the objective function formulation.

	Cost (\$)	# Pump Activation	Duration Active in Night (min)	
No optimization	15.63	50	25	
Optimization with full objective function	8.69	60	0	
Optimization without	Startup cost	8.21	73	6
	Low-water cost	8.67	64	0
	Night-hours cost	8.49	57	74

6.3.2 Performance of Optimization under Various System Settings and Forecasting Models. In this section, we consider the performance of our optimization solution when (1) the average number of pump activations in a day varies between 0.8-3 times and (2) the forecasting errors vary due to the selection of forecasting models for water demand and energy price.

Fixed energy price profile: Under the fixed energy price profile, we note that the amount of cost savings that can be achieved using our solution increases when the pump is activated fewer times in a day, as shown in Fig. 8a. When the water demand is lower (i.e., the pump is activated fewer times), then the ratio of pumping in the night (when the price is higher) to the day is higher. Thus, the potential for cost savings is higher. However, there is a slight dip in the cost savings when the pump is activated very rarely. That is because the water demand is so low that the amount of pumping during the night also drops until the ratio of pumping in the night to the day is lower. So the amount of cost savings dips.

When forecasting is introduced, the amount of cost savings drops due to the forecasting errors. From Fig. 8a, we note there is one exception. Upon further inspection, we found that the optimization solution under forecasting achieves 6% lower cost than the solution under ground truth but the pump is activated 12% more than the solution under ground truth. Thus, the balance between pump health and energy cost caused the optimization solution under forecasting to achieve a lower cost while sacrificing pump health.

As expected, the more frequently the pump is activated in a day, the longer the duration the pump is active in the night. Fig. 8b shows that with our optimized pump schedule, we are consistently

able to reduce this duration to less than 10 minutes with knowledge of the water demand. When we introduce forecasting, this duration increases slightly but remains relatively close to the amount achieved by the best-performing pump schedule.




Figure 8: Performance metrics (percentage of cost savings in (a) and pump activation duration during the night in (b)) of our optimization solution for different water load demands.




Figure 9: Cost savings achieved with different water demand forecasting models and different water load demands.

We also experiment with different forecasting models to evaluate how the accuracy of forecasting water demand affects the optimization solution in terms of the amount of cost savings. Assuming that the energy cost of the standard control algorithm (i.e., the baseline) is C_{sim} , energy cost achieved with optimization using ground truth is C_{GT} , and energy cost achieved with optimization using forecasting is C_F , then we measure the impact of the forecasting model by taking $100 \times \frac{(C_{GT} - C_F)}{C_{Sim} - C_{GT}}$, which is the percentage of cost difference achieved between the two optimization solutions. In Fig. 9, we see that the optimization solution with GRU model achieves the closest cost savings to the best-performing pumping schedule (on average, 18.5% difference with the cost savings achieved with ground truth). The performance is similar to that achieved with the LSTM forecasting model (on average, 21.6% difference with cost savings achieved with ground truth) due to the similarities of the models. In general, the accuracy (measured by the root mean square error (RMSE)) of the water demand forecasting models (shown in Table 2) closely reflect the optimization solution performance. The lower the RMSE, the higher the amount of cost savings achieved.

Moreover, we notice that the forecasting models ARIMA and LR perform worse when the water demand is very low. That is because those two models predict a much higher water consumption during the early morning hours (i.e., the night). So not all of the pumping in the night can be moved to the day due to the forecasting errors. When the water demand load is very low, the optimization solution




Figure 10: Percentage of cost savings achieved with different energy price distribution.

can either (1) move most of the pumping from the night to the day to save cost but sacrifice pump health by pumping a little during the night, and then, again in the day or (2) keep a larger fraction of the pumping in the night and not pump during the day. Based on our weighting of the startup cost, our optimization opts for the second option, thus resulting in a larger energy cost.

Table 2: RMSE of water demand and energy price forecasting models.

Model	Water Demand					Energy Price				
	GRU	GCRF	LSTM	ARIMA	LR	Aug	RF	LSTM	CNN-LSTM	GRU
RMSE	130	141	155	315	378	May	204	448	744	1054
							607	719	1103	1173




Figure 11: Cost savings achieved with different energy price forecasting models and different water load demands. Forecasting models with (a) May energy prices and (b) August energy prices.

Dynamic energy price: For both the May and August energy prices, we determine the amount of cost savings achieved with ground truth knowledge of both water demand and energy prices. As shown in Fig. 10, the cost savings achieved for the month of May is significantly higher than that in August. Since May's energy prices includes multiple high fluctuations, the standard control algorithm results in many pump activations scheduled during those periods. With knowledge of the water demand and energy prices, we can achieve a higher percentage of cost savings by scheduling those activations during periods of lower energy prices.

When energy price forecasting is introduced (and knowledge of water demand is assumed), we note that unlike the results we obtained for the fixed energy price profile, the energy cost savings achieved under different models and different water load demands is extremely volatile. Despite this volatility, we can still see (from

Table 2) that on average, the lower the RMSE of the forecasting model for a particular month, the higher amount of cost savings achieved. The RMSE of all models in May is higher than in August due to the volatility of the energy price fluctuations. Thus, the cost savings is much lower in May than in August, as shown in Fig. 11. The exception is the CNN-LSTM model when the water demand is low. That peak in Fig. 11b for the CNN-LSTM model occurs because pumping is scheduled during one of the highest energy price peaks due to the forecasting errors. In general, we observe that the RF forecasting model performs the best.

6.3.3 *Performance Comparison of Optimization Solution.* We compared our optimization solution against the following benchmarks:

- Quintiliani's [28] optimization approach that changes the water level threshold at certain fixed timeslots. They assume a fixed energy price profile where the energy prices are lower than average during a certain off-peak period. The intuition is to choose a high water level threshold just before the start of the peak pricing period so that the tank will remain near peak capacity and will not require the pump to operate as often during the peak pricing period, as shown in Fig. 12. Then, a low water level threshold is set just before the start of the off-peak pricing period so that the tank remains close to empty.
- Wanjiru's MPC optimization approach [38] that optimizes for energy cost and pump health but does not consider resident satisfaction. Unlike our formulation which distinguishes between the sampling time period and the scheduled pump activation duration, Wanjiru's approach assigns the sampling time period as 10 minutes and assumes the pump operates during the whole duration until the water tank level hits the high threshold. Additionally, the optimization problem seeks to maintain a low water tank level at the end of 24 hours.

We chose the other two optimization approaches because those solutions are also designed for water transfer pump systems in high-rise buildings and they have shown good pumping cost savings in their evaluation. To ensure that the water in the rooftop tank does not overflow or drain out, we always enforce the standard control algorithm on top of these benchmarks.




Figure 12: The optimization approach adopted by Quintiliani. The decision variables are the new water level thresholds (w_l and w_h) and the duration before the price change where the new thresholds would take effect (t_l and t_h).

Table 3 shows the performance of the different algorithms after a month. For the dynamic energy pricing, we compared our approach against Wanjiru's solution [38]. Since Quintiliani's [28] approach assumes a fixed energy price profile, we did not include it

	Cost (\$)	# Pump Activation	Duration Active in Night (min)
Standard Algorithm	15.63	50	26
Our MPC	Ground truth	60	0
	Forecasting	58	7
Wanjiru [38]	Ground truth	70	68
	Forecasting	57	22

(a) Results for August dynamic energy pricing.

	Cost (\$)	# Pump Activation	Duration Active in Night (min)
Standard Algorithm	10.51	50	26
Our MPC	Ground truth	62	0
	Forecasting	68	0
Wanjiru [38]	Ground truth	68	5
	Forecasting	64	11
Quintiliani [28]	Ground truth	106	13
	Forecasting	101	13

(b) Results for fixed energy price profile.

Table 3: Results of different optimization solutions with and without forecasting.

in the comparison for dynamic energy pricing. For the forecasting models, we chose GRU for water demand and RF for energy price forecasting. We also chose a moderate water demand load with the pump activating 1.7 times in a day on average.

Fixed energy price profile: The results shown in Table 3b indicate that our optimization solution, Wanjiru's solution, and Quintiliani's solution achieves a cost savings of 33%, 35%, and 38% respectively compared to the standard control algorithm. When energy prices and water demand are unknown, and thus, requires forecasting, our optimization solution achieves a cost savings of 39% whereas Wanjiru's solution and Quintiliani's solution only achieves a cost savings of 23% and 38% respectively.

In terms of the number of pump activations, all optimization solutions require the pump to be activated more frequently. That is because the pump needs to be active for a minimum duration of, say t minutes, to address the given water demand. In order to achieve the minimum number of pump activations, the pump should be active for the longest duration when it is activated, i.e., it should be activated when the water level is at its lowest. That can be achieved with the standard control algorithm. Any attempt to optimize for energy cost would require the pump to be activated when the water level is higher, thus requiring the pump to be activated more often to address the water demand. Our solution causes a 36% increase in pump activations which is comparable to Wanjiru's solutions (28%) and much better than Quintiliani's (102%). Finally, our approach results in a reduction of 100% of the duration that the pump is active in the night whereas Quintiliani and Wanjiru's solutions result in a 50% and 23% decrease in pumping duration respectively.

In Fig. 13, we observe that comparing the different solutions over the course of a few days does not demonstrate the effectiveness of one solution over another. Under the assumption that water demand is known, our approach results in a lower energy cost than the standard control algorithm but performs slightly worse than Quintiliani's and Wanjiru's approach (shown in Fig. 13a). When forecasting is introduced, then our approach consistently performs better than Wanjiru's solution from the beginning (shown in Fig. 13b). Again, our approach performs worse than Quintiliani's solution in

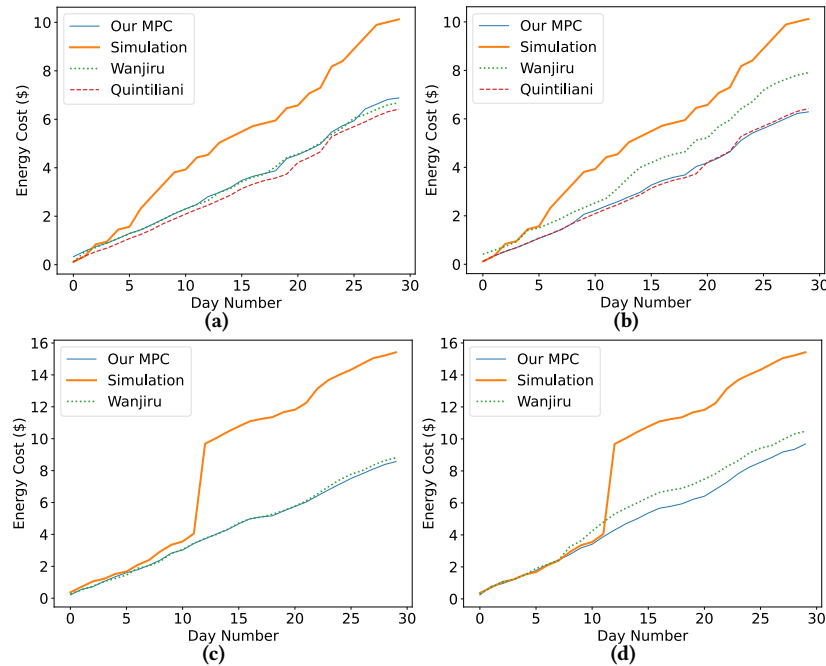


Figure 13: Accumulative energy cost over a month under different optimization solutions. (a) and (b) show the costs under a fixed energy price profile, (c) and (d) show the costs under a dynamic energy price profile. Water demand and energy prices are known in (a) and (c). Water demand is forecasted in (b). Both water demand and energy price are forecasted in (d).

the middle of the month, but by the end of the month, our solution has managed to outperform Quintiliani’s solution.

Dynamic energy pricing: The results shown in Table 3a indicate that our optimization solution and Wanjiru’s solution achieves a cost savings of 44% and 42% respectively compared to the standard control algorithm. When energy prices and water demand are unknown, and thus, requires forecasting, our optimization solution achieves a cost savings of 37% whereas Wanjiru’s solution only achieves a cost savings of 31%. The MPC approaches (i.e., our solution and Wanjiru’s solution) perform better than the standard control algorithm because in the presence of a changing energy price, the MPC approach has more opportunities to carefully select the pumping schedule to minimize energy cost.

Our solution and Wanjiru’s solution causes a 16% and 14% increase in the pump activations. Finally, our approach results in a 73% reduction of the duration that the pump is active in the night. In comparison, Wanjiru’s solution results in a 15% reduction. That is because their optimization formulation does not take into account the duration that the pump is active in the night. On the other hand, we explicitly model that metric as a soft constraint.

Our approach performs worse than the standard control algorithm for the first 5-8 days (shown in Figs. 13c and 13d). That is because our approach seeks a pump schedule that not only achieves lower energy cost but results in a higher water level at the end of the day. This trade-off results in the energy cost being slightly higher. Also, our solution achieves similar cost savings to Wanjiru’s when ground truth of water demand and energy price is provided. However, Fig. 13d shows that as time progresses, our algorithm achieves a much higher cost savings than any other approach due to its ability to adapt to changing water demand and energy cost. That

is because Wanjiru’s optimization formulation represents the pump schedule as a coarse-grained binary decision variable whereas we define a integer decision variable that allows the pump to be turned on for a varying amount of time. Thus, Wanjiru’s optimization solution generates a strict pump schedule, and is not as flexible as our approach in dealing with forecasting errors. Therefore, our results show that the short timeframe evaluation conducted by the related works in this domain is insufficient to prove that their proposed approach achieves better energy cost savings. In particular, we observe that our MPC formulation performs best over an extended time period, which is more reflective of real-world settings.

Summary: Unlike existing literature that only evaluates the amount of energy cost savings when the energy price profile and water demand is known, we demonstrate in this section that the amount of energy cost savings is highly dependent on the forecasting techniques and the water load demand. In general, our solution achieves the best energy cost savings overall. Our solution is also better able to handle forecasting errors, resulting in a much lower energy cost than both the standard control algorithm and the related works. Under a fixed energy price profile, when the water demand load is lower, we are able to achieve a higher percentage of cost savings. It is prudent to adopt a water forecasting model with a low RMSE so that the optimized solution will be as close to the best-performing solution as possible. Our solution is also able to reduce the pump activation duration during the night down to less than ten minutes (for the experiments that we conducted). Under a dynamic energy price profile, when water demand and energy price is known, we achieve higher cost savings when the energy prices fluctuate greatly. However, when energy price is not known, the optimization solution results are volatile and unpredictable. While

more experiments are needed to determine the amount of volatility introduced by forecasting, our current results already show that the energy savings fluctuate greatly across different forecasting techniques and water load demand. In general, a forecasting model with low RMSE will achieve better energy cost savings but the results depend highly on the water load demand.

7 CONCLUSION

In this paper, we proposed to use a MPC optimization framework that integrates ML/DL predictions of water demand and energy prices to minimize the pump operation cost, maintain pump health, and increase resident satisfaction in buildings. The ML/DL models forecast the next 24 hours of water demand and energy pricing, utilizing historical data from the previous day. Those predictions, along with the current water level in the tank, are passed into the MPC optimization framework. The MPC approach strategically balances long-term planning and uncertainty reduction in the optimization process. We conducted a case study by simulating a university building pumping system with real-world data for water demand and energy prices spanning a prolonged period of time. Our results showed that our optimization solution achieves 37% cost savings with our water demand and energy price forecasting models. It also reduces the duration that the pump is active in the night by 73-100%, thus improving resident satisfaction. Additionally, we show that the accuracy of water demand forecasting models directly affect the amount of cost savings achieved with optimization. On the other hand, the amount of cost savings achieved under a dynamic price profile is very volatile. While the accuracy of energy price forecasting models does have some correlation to the cost savings, the actual amount of cost savings depends highly on the water load demand and water consumption pattern. Thus, our work presents a first step towards understanding how forecasting models affect the optimization results in real-world systems and the development of more robust optimization solutions under such prediction errors.

ACKNOWLEDGMENTS

This research/project is supported by the National Research Foundation, Singapore, and Ministry of National Development, Singapore under its Cities of Tomorrow R&D Programme (CoT Award COT-V2-2020-1). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore and Ministry of National Development, Singapore.

REFERENCES

- [1] 2008. Annoyed over water pump noise. <https://condosingapore.com/showthread.php/2669-Annoyed-over-water-pump-noise>
- [2] 2018. New water pump causes sleepless nights for tenant. <https://ottawasun.com/life/homes/rental-guide-apartment-buildings-new-water-pump-causes-sleepless-nights-for-tenant>
- [3] 2022. Households used more water in 2020, companies less. <https://www.cbs.nl/en-gb/news/2022/12/households-used-more-water-in-2020-companies-less>
- [4] 2022. Singapore's daily water consumption rose in 2021 for second straight year. <https://www.straitstimes.com/singapore/politics/budget-debate-daily-consumption-of-water-rose-again-in-2021-to-reach-158-litres-per-capita>
- [5] 2023. EMC Price Information. <https://www.emcsg.com/MarketData/PriceInformation>
- [6] 2023. EPANET - Application for Modeling Drinking Water Distribution Systems. <https://www.epa.gov/water-research/epanet>
- [7] 2023. GLPK - GNU Linear Programming Kit. <https://www.gnu.org/software/glpk/>
- [8] 2023. Smart Water Meter - Singapore. <https://www.smartnation.gov.sg/initiatives/urban-living/smartwatermeter/>
- [9] 2023. USEPA/WNTR - An EPANET compatible python package. <https://github.com/USEPA/WNTR>
- [10] Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. 2018. A rewriting system for convex optimization problems. *Journal of Control and Decision* 5, 1 (2018), 42–60.
- [11] David Angeli and Matthias A. Müller. 2019. *Economic Model Predictive Control: Some Design Tools and Analysis Techniques*. Springer International Publishing, Cham, 145–167.
- [12] G. Bejarano, A. Kulkarni, R. Raushan, A. Seetharam, and A. Ramesh. 2019. SWaP: Probabilistic Graphical and Deep Learning Models for Water Consumption Prediction. In *Proc. of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation* (New York, USA). ACM, 233–242.
- [13] R. Bruce Billings and Donald E. Agthe. 1998. State-Space versus Multiple Regression for Forecasting Urban Water Demand. 124, 2 (1998), 113–117.
- [14] Public Utilities Board. 2022. *Handbook on Application for Water Supply*.
- [15] Antonio Candelieri, Raffaele Perego, and Francesco Archetti. 2018. Bayesian optimization of pump operations in water distribution systems. *Journal of Global Optimization* 71, 1 (2018), 213–235.
- [16] Zihan Chang, Yang Zhang, and Wenbo Chen. 2018. Effective Adam-Optimized LSTM Neural Network for Electricity Price Forecasting. In *2018 IEEE 9th International Conference on Software Engineering and Service Science (ICSESS)*, 245–248.
- [17] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17, 83 (2016), 1–5.
- [18] Mohammad Ebrahim Banihabib and Pezhman Mousavi-Mirkalaei. 2019. Extended linear and non-linear auto-regressive models for forecasting the urban water consumption of a fast-growing city in an arid region. 48 (2019), 101585.
- [19] Bradley Eck, Sean McKenna, Albert Akhiev, Akihiro Kishimoto, Paulito Palmes, Nicole Taheri, and Susara van den Heever. 2014. *Pump Scheduling for Uncertain Electricity Prices*. 426–434.
- [20] Bissan Ghaddar, Joe Naoum-Sawaya, Akihiro Kishimoto, Nicole Taheri, and Bradley Eck. 2015. A Lagrangian decomposition approach for the pump scheduling problem in water networks. *European Journal of Operational Research* 241, 2 (2015), 490–501.
- [21] M. Ghiassi, David K. Zimbra, and H. Saidane. 2008. Urban Water Demand Forecasting with a Dynamic Artificial Neural Network Model. 134, 2 (2008), 138–146.
- [22] H.S. Kang, H. Kim, J.K. Lee, I. Lee, B.-Y. Kwak, and H.J. Im. 2014. Optimization of pumping schedule based on water demand forecasting using a combined model of autoregressive integrated moving average and exponential smoothing. *Water Supply* 15, 1 (10 2014), 188–195.
- [23] Rüdiger Kiesel and Michael Kusterman. 2016. Structural models for coupled electricity markets. *Journal of Commodity Markets* 3, 1 (2016), 16–38.
- [24] Jesus Lago, Fjo De Ridder, and Bart De Schutter. 2018. Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. 221 (2018), 386–405.
- [25] Wenjie Lu, Jiazheng Li, Yifan Li, Aijun Sun, Jingyang Wang, and Abd E. I.-Baset Hassanien. 2020. A CNN-LSTM-Based Model to Forecast Stock Prices. *Complexity* 2020 (1 2020), 10 pages.
- [26] Tiago Luna, João Ribau, David Figueiredo, and Rita Alves. 2019. Improving energy efficiency in water supply systems with pump scheduling optimization. *Journal of cleaner production* 213 (2019), 342–356.
- [27] David R. Maidment and Shaw-Pin Miaou. 1986. Daily Water Use in Nine Cities. *Water Resources Research* 22, 6 (1986), 845–851.
- [28] Claudia Quintiliani and Enrico Creaco. 2019. Using additional time slots for improving pump control optimization based on trigger levels. *Water Resources Management* 33, 9 (2019), 3175–3186.
- [29] J. Rawlings, D.Q. Mayne, and Moritz Diehl. 2017. *Model Predictive Control: Theory, Computation, and Design*.
- [30] Danielle C. M. Ristow, Elisa Henning, Andreza Kalbusch, and Cesar E. Petersen. 2021. Models for forecasting water demand using time series analysis: a case study in Southern Brazil. 11, 2 (2021), 231–240.
- [31] Tony Salloom, Okyay Kaynak, and Wei He. 2021. A novel deep neural network architecture for real-time water demand forecasting. 599 (2021), 126353.
- [32] P. Sotasakis, A.K. Sampathirao, A. Bemporad, and P. Patrinos. 2018. Uncertainty-aware demand management of water distribution networks in deregulated energy markets. *Environmental Modelling & Software* 101 (2018), 10–22.
- [33] Anna Stuhlmacher and Johanna L. Mathieu. 2020. Chance-Constrained Water Pumping to Manage Water and Power Demand Uncertainty in Distribution Networks. *Proc. IEEE* 108, 9 (2020), 1640–1655.
- [34] Bartosz Uniejewski, Grzegorz Marczasz, and Rafał Weron. 2019. Understanding intraday electricity markets: Variable selection and very short-term price forecasting using LASSO. 35, 4 (2019), 1533–1547.

- [35] Bartosz Uniejewski, Jakub Nowotarski, and Rafał Weron. 2016. Automated Variable Selection and Shrinkage for Day-Ahead Electricity Price Forecasting. 9, 8 (2016), 621.
- [36] Bartosz Uniejewski and Rafał Weron. 2018. Efficient Forecasting of Electricity Spot Prices with Expert and LASSO Models. 11, 8 (2018), 2039.
- [37] Ye Wang, Kevin Too Yok, Wenyan Wu, Angus R. Simpson, Erik Weyer, and Chris Manzie. 2021. Minimizing Pumping Energy Cost in Real-Time Operations of Water Distribution Systems Using Economic Model Predictive Control. *Journal of Water Resources Planning and Management* 147, 7 (2021), 04021042.
- [38] Evan M Wanjiru, Lijun Zhang, and Xiaohua Xia. 2016. Model predictive control strategy of energy-water management in urban households. *Applied Energy* 179 (2016), 821–831.
- [39] S. L. Zhou, T. A. McMahon, A. Walton, and J. Lewis. 2000. Forecasting daily urban water demand: a case study of Melbourne. 236, 3 (2000), 153–164.
- [40] Florian Ziel. 2016. Forecasting Electricity Spot Prices Using Lasso: On Capturing the Autoregressive Intraday Structure. 31, 6 (2016), 4977–4987.