A Reinforcement Learning Approach to Demand Balancing and Tariff Optimization in Blockchain-Based Smart Water Networks

Filippo Ansalone, Ioannis Chatzigiannakis
Sapienza University of Rome Rome, Italy
CNIT, Parma, Italy
ansalone.1950936@studenti.uniroma1.it,
ichatz@diag.uniroma1.it

Vincenzo Taormina, Domenico Garlisi
University of Palermo, Palermo, Italy
CNIT, Parma, Italy
INDAM, Roma, Italy
vincenzo.taormina@unipa.it, domenico.garlisi@unipa.it

***Abstract*—The increasing complexity of Water Distribution Networks (WDN), combined with the need for energy efficiency, requires innovative solutions to optimize resource management and reduce operational inefficiencies. This paper introduces a Reinforcement Learning (RL) based framework designed to optimize water consumption patterns and tariff selection for users in smart WDNs. The proposed system leverages smart contracts to dynamically apply tariffs set by suppliers, while users are equipped with water storage tanks that can be filled using consumption patterns aligned with lower tariffs. By distributing diverse consumption patterns among users, driven by dynamic tariffs, the framework mitigates demand peaks at specific times, ensuring balanced network demand throughout the day. The RL algorithm functions using the EPANET simulation, which is fed by real data obtained from the Digital Twin (DT) states, showcasing the actual status of meters interconnected via the Internet of Things (IoT) network. This work demonstrates the potential of combining RL, smart contracts, DT, IoT and distributed consumption patterns to create a resilient and efficient smart WDNs, addressing both user-centric and operator-centric challenges in modern water management systems.**

***Index Terms*—Reinforcement Learning (RL), Smart Contracts, Internet of Things (IoT), Digital Twin (DT), Water Distribution Network (WDN).**

I. Introduction

Water scarcity, exacerbated by climate change, has become a critical challenge for many regions worldwide. Many territories, like the Sicilian island in Italy, try to address the arid climate and prolonged dry spells by relying on small cisterns and water tanks installed in residential buildings to support the management of the water needs of the users. Such solutions have been discussed in the literature, e.g., in [1], [2], and various approaches have been proposed that also consider how to optimize incurred energy costs, showing that the existing extensive infrastructure can be used to improve energy savings and water resource management [3]. Yet, providing sustainable solutions remains a complex task also due to the intricate structure of Water Distribution Network (WDN) and integrating such solutions in an optimal way remains an open problem. As a result, current infrastructures are still struggling to meet the demands of both residents and the continuously increasing touristic flows[1][2][3].

In this paper, we propose a complete system that enables suppliers to optimize water allocation, balance demand, and support informed decision-making that may lead to optimal resource management.

We rely on the following building blocks. First, the large-scale deployment of smart metering across the WDN offering real-time monitoring on water usage [4]–[7]. Second, the transformation of traditional water tanks into intelligent storage units that provide detailed information on available capacity [8]. Third, a WDN Digital Twin (DT) that integrates the real-time data arriving from the smart meters and smart water tanks with simulated data to identify optimal usage patterns for each individual user. These optimal patterns are then provided as feedback to the users to decide independently on how to adjust their consumption pattern individually. In other words, we propose a collaborative selection mechanism of consumption patterns across

This work was partially supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on "Telecommunications of the Future" (PE00000001 - program "RESTART" - CUP E83C22004640001) in the WITS and SPRINT focused project. This work was partially supported by the project SERICS (PE00000014 - CUP D33C22001300002) under the NRRP MUR program funded by the EU-NGEU.
ISBN 978-3-903176-72-0 © 2025 IFIP

[1] finance-commerce.com/2024/07/record-drought-tests-sicilys-water-infrastructure-and-tourism-industry/
[2] euronews.com/green/2024/08/07/sicilys-population-fed-up-of-water-shortages-as-rationing-starts-to-bite
[3] www.theguardian.com/environment/article/2024/aug/19/the-land-is-becoming-desert-drought-pushes-sicilys-farming-heritage-to-the-brink

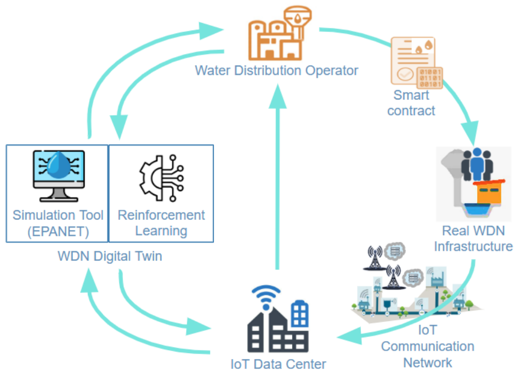

Fig. 1. System infrastructure for smart WDN. DT processes the data related to the real WDN network with the EPANET simulation environment and the RL algorithm.

users that relies on the tank infrastructure to mitigate peak demand issues, thus preventing excessive pressure fluctuations that can potentially lead to pipe leakages.

As depicted in Fig. 1 this process proceeds iteratively, allowing the network to continuously evaluate and dynamically adapt the proposed water consumption patterns to the always varying demand patterns, weather events, or infrastructure updates. We implement the proposed system as open source in a publicly available repository[4] for the scientific community to test various methods on different WDNs. We validate our method on a realistic WDN, demonstrating its potential in reducing peaks and cutting costs.

The remainder of the paper is structured as follows: Sec. II presents the design of the proposed system along with a presentation of the main components. Then in Sec. IV we present the findings of the performance evaluation conducted based on a real-world inspired WDN. In Sec. V we present essential background information and the relevant state-of-the-art. The paper concludes with Sec. VI that summarizes the main findings and provides future work directions.

II. System Design

The starting point is the WDN, visible on the right side of Fig. 1, that is composed of pipelines, storage tanks and reservoirs, pumps and junctions. IoT sensors and smart meters monitor water flow, pressure, and consumption in real time. The data collected from the IoT network feeds into the DT of the WDN that mirrors the real infrastructure in a virtual environment.

The DT integrates a simulation framework with the WNTR/EPANET platform, a prevalent tool for modeling hydraulic behaviors in water distribution networks [9]. WNTR, an open-source software derived from EPANET (developed by the US Environmental Protection Agency), is designed to simulate the hydraulic

[4]https://github.com/WITS-Restart/WDN-tariff-optimizaiton

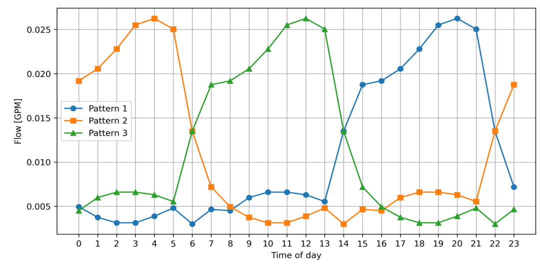

Fig. 2. Daily water consumption patterns used for the evaluation of the proposed DQN model.

and quality dynamics within a WDN [10]. It evaluates network geometry, along with initial parameters such as pipe roughness and diameter and operational rules, to compute various metrics, including flow rates, pressure levels, and water quality indicators.

A Reinforcement Learning (RL) model is used to continuously learn and identify the optimal water consumption strategies based on cost efficiency and network stability to significantly reduce pressure spikes and possible consequent leakages. The identified strategies are then communicated to the water distribution operator for final approval. We integrate blockchain-enabled smart contracts to guarantee on one hand the transparency of the identified optimal water consumption patterns and on the other hand securely record the actual decision of the users to follow these suggestions. The smart contract can also be used by the water suppliers to offer additional services [11]. The decision of the users to potentially adapt their usage patterns based on the recommendation becomes the input of the next iteration, ensuring that the simulation accurately reflects the current state of the physical system. As a result, the DT adapts to real-time data, learning from historical patterns and dynamically adjusting water allocation strategies to minimize costs and balance pressure/demands picks in WDN infrastructure.

A. Usage Patterns

We assume three representative consumption patterns that capture typical variations in water demand throughout the day in an urban environment and can be alternatively selected from the proposed system [12]. These patterns, illustrated in Fig. 2, are characterized by a peak consumption period in one-third of the 24-hour cycle, while the remaining two-thirds exhibit lower demand levels. The three patterns are differentiated by the timing of their peak: morning, afternoon, or evening. These consumption patterns serve as the basis for our optimization framework, where each user could be assigned a specific consumption behavior.

We note that in a real-world scenario, arbitrarily selecting a pattern is not realistic. In fact, several constraints will affect the actual applicability of a pattern. In the definition of the proposed model, we consider user-specific constraints, where some users may only be eligible for certain patterns based on physical limitations. For example, depending on the capacity of the water storage tank during certain hours, some patterns may or may not be applicable. But there may also be temporal constraints, where some patterns may not be suitable at specific times of the day due to operational restrictions, and constraints based on the physical network, such as clusters of high-demand users, and hence relatively correlated constraints between multiple users.

Future research will enhance user constraints by utilizing evidence from pertinent studies [12] to consider inter-user interactions and network-wide relationships. This integration will enable the WDN DT to simulate water usage more accurately.

B. Tariffs

Operating the pump connected to the water tank installed at each building incurs certain energy costs. In this work, we consider a time-of-use pricing structure that resembles the pricing schemes commonly used in utility pricing. Pricing functions are defined as square wave patterns, where higher rates apply during peak hours and lower rates encourage consumption during off-peak periods. Specifically, we define two pricing windows: one with higher prices from 3:00 to 15:00 and another from 18:00 to 24:00.

In addition to consumption patterns, the optimization phase tries to align storage tank filling and consumption patterns with lower tariffs, thus reducing the costs of the users while promoting a more balanced distribution of demand throughout the day.

III. Pattern selection scheme

The process of choosing the target pattern for each user is using a Deep Q-Learning (DQN) based RL algorithm. The goal is to determine the optimal consumption pattern given the current state of the WDN. Through diversified consumption patterns, driven by the tariffs, it tries to balance network demand throughout the day while maximizing cost savings for the users. As illustrated in Fig. 3 the reinforcement process operates at two levels. The first level consists of selecting an individual user pattern and simulating the effect it brings on the resource usage of the WDN. This allows the system to evaluate different consumption patterns while considering constraints on pattern assignment. This is embedded in a larger reinforcement loop that incorporates tariff definition and, finally, finds optimal price-dependent patterns.

The DQN is a model-free reinforcement learning algorithm that uses a neural network to approximate the

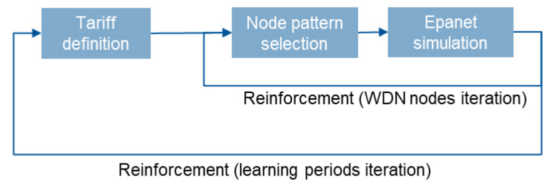

Fig. 3. Schematic representation of the DQN framework. The process consists of two nested reinforcement loops: the inner loop selects consumption patterns and simulates their effects within the EPANET, while the outer loop integrates tariff definition.

Q-function instead of maintaining a traditional lookup table for Q-values as follows:

$$Q(S, a) = \mathit{NeuralNetwork}(S) \qquad (1)$$

where S is the current state of the WDN and a is the selected usage pattern. Therefore, each user is mapped to an agent, an autonomous decision-making entity that interacts with the WDN. At each step, the agent selects a usage pattern based on its current knowledge and explores the impact of this choice on the WDN. To guide the learning process, the reward is based on the response of the WDN to the selected usage pattern.

The Q-table of the agent is a two-dimensional matrix, where $Q(S, a) = Q(s_l, s_p, a)$ represents the Q-value in the Q-table, s_l represents the current tank levels at each node, s_p represents the pressure at each node, and a represents the action selected for the current user. The agent selects action $a \in A$, causing the assignment of a pattern to the current user, with A the set of possible selectable actions.

The DQN architecture consists of two sequential neural networks, each composed of linear layers interleaved with ReLU activations. The first network processes, for each user, 24 hourly pressure values. Note that pressure values are influenced by the water consumption of every node of the WDN that depends on both the assigned pattern and its own tank level. These are transformed through layers with 256, 128, and 128 neurons, producing a 32-dimensional output vector. Additionally, since every node can have a different amount of daily water consumption, we added an embedding layer that, starting from the node ID, outputs a feature vector of 16 elements. These outputs are then concatenated with 24 values representing the hourly selected tariff, forming a 72-dimensional $(32+16+24)$ feature vector that serves as input to the second neural network. The second network includes two hidden layers of 128 neurons each and an output layer that assigns the optimal consumption pattern for the current node. All feature inputs are normalized between 0 and 1 before being processed. In particular, the pressure features are normalized using the max-min method, where the minimum value is set to zero and the maximum corresponds to the highest simulated pressure in the network, as computed by EPANET.

The reward function used here has two objectives. The first objective is to ensure the supply at the district by regulating the level in the water tank. The second is to reduce the operating costs by optimizing the satisfied water in all the users by maintaining the pressure level at each user in a safety range. We want to avoid that the pressure goes upper or lower than a certain range. In order to obtain an optimal control policy which leads the WDN to fulfill the desired control objectives, a cost (reward) function is defined as

$$C(S,a) = \beta \cdot \text{WaterCost}(S, s_f, a) + \delta \cdot \text{Penalization}(S,a) \tag{2}$$

it represents the immediate cost for the current simulation after taking action a in state S with the applied tariff s_f. We have considered a cost function composed of two components: WaterCost for the user calculated on the water consumption for the applied tariff, and a Penalization component due to unsatisfied water demands. These components are weighted by the coefficients β and δ, respectively. The reward function incorporates an incentive for the agent to balance economic efficiency with service reliability. In fact, these coefficients determine the relative importance of different aspects of the optimization problem, such as cost minimization and demand balancing. In this work, we chose to assign a higher weight to demand satisfaction compared to cost to ensure the stability and reliability of WDN.

The algorithm learns to balance exploration (trying new actions) and exploitation (using known actions) to identify the most cost effective tariffs. As a result, the process of learning the optimal strategy using Deep Q-Learning aligns with the concept of identifying the most efficient distribution of the consumption pattern.

In the training process, the Q-values are updated based on the Bellman [13] equation:

$$\begin{aligned} Q(S_t, a) &= Q(S_t, a) \\ &\quad + \alpha \left[R_t + \gamma \max_{a'} Q(S_{t+1}, a') - Q(S_t, a) \right] \end{aligned} \tag{3}$$

where α is the learning rate, R_t is the reward received after taking action a in state S_t. γ: discount factor and $\max_{a'} Q(S_{t+1}, a')$ represent the maximum Q-value for the next state. The γ parameter is set to 0.95 and represents the discount factor applied to future rewards. Determines the importance of future rewards relative to immediate decisions. A higher γ value encourages the algorithm to prioritize long-term rewards.

The agent stores past experiences in a *memory buffer* which stores the most recent 2000 steps executed during training. The memory allows to sample past experiences to train the neural network and learn from a diverse set of states and actions, rather than relying solely on the most recent interactions.

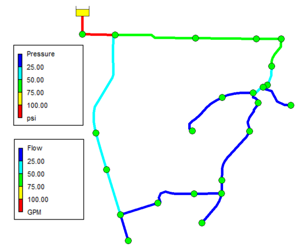

Fig. 4. Schematic representation of the WDN topology used, comprising 23 nodes (junctions), 1 reservoir (top left), and 24 pipes.

Since Q-values start at zero, the agent has no preference for actions initially. To balance exploration (trying random actions) and exploitation (choosing the best-known action), an epsilon-greedy (ϵ-greedy) strategy is used by following the formula:

$$\text{action} = \begin{cases} \arg\max_a Q(S,a) & \text{if } r > \epsilon \quad \text{(exploitat.)}, \\ \text{random action} & \text{otherwise} \quad \text{(explorat.)}. \end{cases} \tag{4}$$

At each step, the agent checks if a random number $r \in [0,1]$ exceeds ϵ - a parameter which controls the exploration-exploitation trade-off, has a range $[1, 0.05]$. Exploration rate (ϵ) starts at 1 (100% random exploration) and decays over episodes, reducing exploration as the agent learns (up to 0.05). This ensures that the agent explores early but gradually shifts to exploiting learned Q-values. As ϵ decreases, the algorithm increasingly relies on the trained neural network to select actions, shifting from exploration to exploitation.

IV. Performance Evaluation

We evaluated our approach using the network topology available from the Open Water Analytics community repository [14]. The selected WDN benchmark, shown in Fig. 4, contains 23 nodes (junctions), 1 reservoir, and 24 pipes. The users of the WDN are indicated by circles (nodes), and the pipes by lines. The reservoir is located at the top left. The pipes and nodes are colored on the basis of their current flow and pressure at a certain time step, aligned with the colormap shown in the figure. Our objective is to show that the proposed selection scheme surpasses the reference baseline and optimally adjusts the chosen pattern for each WDN node to minimize user costs while meeting water needs. To achieve this, we apply the proposed approach to the chosen WDN, taking into account the consumption pattern and tariffs detailed in Section II-A. Additionally, we assign values of 100 and 1000 to the coefficients β and δ in formula (2).

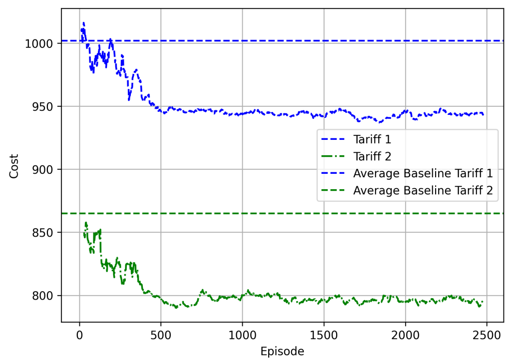

Fig. 5. Evolution of the cost function over training episodes on average and for the two tariffs.

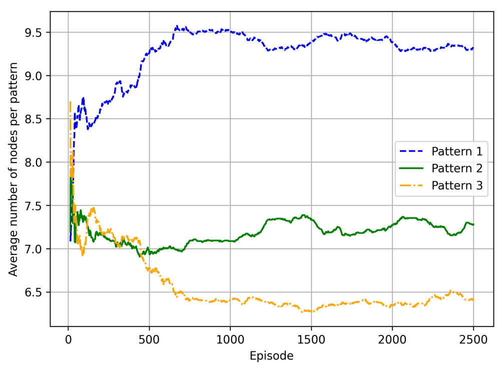

Fig. 6. For a specific tariff, the evolution of the average number of nodes assigned to each consumption pattern across training episodes. Highlighting how the optimal tariff is preferred.

The results are highlighted in Fig. 5 and in Fig. 6. 5 illustrates the evolution of the cost function between training episodes for the two chosen tariffs. The plot displays the average cost computed as the rolling average with a window set to 20 samples. Initially, the RL agent starts from random solutions, leading to higher cost values. However, as the training progresses, the algorithm optimizes the assignment of consumption patterns, gradually reducing costs. The convergence stabilizes after approximately 500 episodes, maintaining a relatively steady behavior. The RL algorithm achieves minimum costs slightly above 800 for tariff 2 and just below 950 for tariff 1. A key factor influencing training performance is the computational time per episode, which is predominantly dictated by the EPANET simulation. In our computational environment (Intel Xeon w7-3455, 2.5 GHz, 24 Core, 187GB RAM), each episode takes an average of 39 seconds.

To assess the quality of these results, and given the infeasibility of exhaustively exploring all possible pattern combinations across the nodes, we define baseline configurations for comparison. The baselines include uniform assignments, mixed assignments using two patterns at a time and random assignments. The baselines are constructed while respecting the node constraints outlined above to serve as a comparison with the RL results. Along the three, the lowest cost baselines, shown in Fig. 5, are derived from random pattern assignments while respecting these constraints, yielding costs of 1007 and 861 for tariff 1 and tariff 2, respectively. Conversely, the most expensive baselines arise from applying one pattern to all nodes, considering constraints. For figure clarity, Fig. 5 excludes this baseline, along with the uniform baseline.

Fig. 6 illustrates, for a specific tariff, the evolution of the average number of nodes assigned to each consumption pattern over training episodes. As the RL algorithm learns, it converges towards a distribution where one pattern is preferred over the others. This makes perfect sense since it is the consumption pattern that uses most of the water during the hours of lower price. We note that, despite achieving optimal system cost through uniform node distribution across the three patterns, the proposed selection method identifies which nodes should be assigned a specific pattern. It is essential to determine the pattern assignment for nodes while taking into account constraints related to tank levels and tariffs.

V. Related Works

a) User consumption optimization in WDN: DQN, introduced by Mnih et al. [15], marked a significant advance in RL by combining Q-learning with deep neural networks. One of the key innovations is the inclusion of experience replay, which breaks the correlation between consecutive samples. This approach demonstrated human-level performance in complex tasks, establishing deep RL as a powerful framework for decision-making in high-dimensional spaces. A limited number of studies in the relevant literature have investigated the optimization of WDNs using DQN in controlling pump speed [16]–[18]. These works highlight the potential of DQN to improve efficiency by dynamically adjusting pump operations to reduce energy consumption and maintain optimal pressure levels. To the best of our knowledge, this is the first study that investigates how DQN can facilitate the operation of a number of water storage tanks installed throughout the WDN to obtain optimal pricing-dependent patterns.

b) Smart contract in WDN: Smart contracts can improve water resource management by involving citizens in policy-making, voting, and implementation, ensuring transparency and mutual accountability [11], [19]. This approach advocates for sustainable management and empowers citizens by integrating them into decision-making, promoting communal goals. Smart contracts have the potential to facilitate various applications in the retail water market, including contract signing and secure data management of water usage for billing purposes [20]. Smart contracts have also been considered for the conservation of residential water via a resource sharing system among consumers, highlighting a consumer-centric model that allows households to manage surplus water through a decentralized Ethereum-based platform, facilitating peer-to-peer transactions [21]. Mohanta *et al.* [22] describe a smart home simulation using the Ethereum blockchain to store temperature and humidity data and support smart contracts for efficient data management. This study examines the use of smart contracts to manage the interaction between water suppliers and the consumer, facilitating dynamic cost determination and validating tank fill levels. This ensures that suppliers have accurate status information, preventing fraudulent attempts by users to manipulate patterns.

VI. Conclusions and Future Work

In this paper, we proposed a transformative approach to water management by integrating real-time IoT data monitoring the real infrastructure with a detailed digital twin. The flow rates, pressure levels, and consumption patterns continuously collected by the smart meters and sensors deployed across the network are transmitted to the digital twin, which mirrors the physical system and allows dynamic updates to the simulation model. A collaborative selection mechanism of consumption patterns between users has been designed to mitigate peak demand issues, thus preventing excessive pressure fluctuations that can potentially lead to pipe leakage. The proposed approach considers DQN as a reinforcement learning strategy to optimize water demand management. The approach significantly outperforms several baseline strategies—including uniform, random, and mixed pattern assignments. The resulting system represents a paradigm shift in water management. Not only may it address the immediate challenges of water scarcity, but it also lays the foundation for a sustainable and adaptive system capable of withstanding future climate uncertainties.

References

[1] A. Campisano and C. Modica, "Optimal sizing of storage tanks for domestic rainwater harvesting in sicily," *Resources, Conservation and Recycling*, vol. 63, pp. 9–16, 2012.

[2] "Potential for peak flow reduction by rainwater harvesting tanks," *Procedia Engineering*, vol. 89, 2014, 16th Water Distribution System Analysis Conference, WDSA2014.

[3] Y. Chang, G. Choi, J. Kim, and S. Byeon, "Energy cost optimization for water distribution networks using demand pattern and storage facilities," *Sustainability*, vol. 10, 2018.

[4] D. Amaxilatis, I. Chatzigiannakis, C. Tselios, N. Tsironis, N. Niakas, and S. Papadogeorgos, "A smart water metering deployment based on the fog computing paradigm," *Applied Sciences*, vol. 10, 2020.

[5] M. Zecchini, A. A. Griesi, I. Chatzigiannakis, D. Amaxilatis, and O. Akrivopoulos, "Identifying water consumption patterns in education buildings before, during and after covid-19 lockdown periods," in *2021 IEEE International Conference on Smart Computing (SMARTCOMP)*.

[6] M. Zecchini, A. A. Griesi, I. Chatzigiannakis, I. Mavrommati, D. Amaxilatis, and O. Akrivopoulos, "Using iot data-driven analysis of water consumption to support design for sustainable behaviour during the covid-19 pandemic," in *2021 6th SEEDA-CECNSM*. IEEE, 2021.

[7] P. Pisani and R. J. L. Taloma, "L'intelligenza artificiale per le reti idriche," Servizi a Rete, Tech. Rep., august 2024.

[8] M. A. Wister, E. Leon, A. Alejandro-Carrillo, P. Pancardo, and J. A. Hernandez-Nolasco, "Using iot for cistern and water tank level monitoring," *Applied System Innovation*, no. 6, 2024. [Online]. Available: https://www.mdpi.com/2571-5577/7/6/112

[9] K. A. Klise *et al.*, "A software framework for assessing the resilience of drinking water systems to disasters with an example earthquake case study," *Environmental Modelling & Software*, vol. 95, 2017.

[10] J. Liu and Y. Guoping, "Iterative methodology of pressure-dependent demand based on epanet for pressure-deficient water distribution analysis," *Journal of Water Resources Planning and Management*, vol. 139, pp. 34–44, 01 2013.

[11] A. Bracciali, I. Chatzigiannakis, A. Vitaletti, and M. Zecchini, "Citizens vote to act: Smart contracts for the management of water resources in smart cities," in *2019 First International Conference on Societal Automation (SA)*. IEEE, 2019.

[12] R. J. L. Taloma, F. Cuomo, D. Comminiello, and P. Pisani, "Machine learning for smart water distribution systems: exploring applications, challenges and future perspectives," *Artificial Intelligence Review*, vol. 58, no. 4, p. 120, Jan 2025.

[13] F. L. Lewis and K. G. Vamvoudakis, "Reinforcement learning for partially observable dynamic processes: Adaptive dynamic programming using measured output data," *IEEE Transactions on Systems, Man, and Cybernetics*, 2011.

[14] "OpenWaterAnalytics," accessed: 2024-04-15. [Online]. Available: https://raw.githubusercontent.com/OpenWaterAnalytics/epanet-example-networks/master/epanet-tests/large/NW_Model.inp

[15] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, 2015.

[16] G. Hajgató, G. Paál, and B. Gyires-Tóth, "Deep reinforcement learning for real-time optimization of pumps in water distribution systems," *Journal of Water Resources Planning and Management*, vol. 146, no. 11, p. 04020079, 2020.

[17] S. Hu, J. Gao, D. Zhong, R. Wu, and L. Liu, "Real-time scheduling of pumps in water distribution systems based on exploration-enhanced deep reinforcement learning," *Systems*, vol. 11, no. 2, p. 56, 2023.

[18] J. Xu, H. Wang, J. Rao, and J. Wang, "Zone scheduling optimization of pumps in water distribution networks with deep reinforcement learning and knowledge-assisted learning," *Soft Computing*, vol. 25, pp. 14 757–14 767, 2021.

[19] M. Zecchini, A. Bracciali, I. Chatzigiannakis, and A. Vitaletti, "On refining design patterns for smart contracts," in *European Conference on Parallel Processing*. Springer, 2019.

[20] T. Sarantakos, D. Amaxilatis, A. Pagano, D. Garlisi, R. Laceda Taloma, T. Cattai, I. Chatzigiannakis, V. Vythoulka, and C. Zaroliagis, "A tool to facilitate the design of smart contracts in smart water distribution networks," in *IFIP*, 2024.

[21] S. Tiwari, J. Gautam, V. Gupta, and N. Malsa, "Smart contract for decentralized water management system using blockchain technology," *Int J Innov Technol Explor Eng*, vol. 9, no. 5, pp. 2046–2050, 2020.

[22] B. K. Mohanta, D. Jena, S. Ramasubbareddy, M. Daneshmand, and A. H. Gandomi, "Addressing security and privacy issues of iot using blockchain technology," *IEEE Internet of Things Journal*, 2020.