

Toward Edge Implementation of Soft Sensors for Water Distribution Networks

Fotis N. Koumboulis^{#1}, Nikolaos D. Kouvakas^{#2}, Maria P. Tzamtzi^{#3}, Dimitrios G. Fragkoulis^{#4}, Antonios N. Menexis^{#5}, Klimis K. Katsiavrias^{#6}, Konstantinos Katsiavrias^{#7}, John Dimitropoulos^{*8}

[#]*Robotics, Automatic Control and Cyber-Physical Systems Laboratory, Department of Digital Industry Technologies, National and Kapodistrian University of Athens, Euripus Campus, Euboea, Greece*

^{1,2,3,4,5}{fkoumboulis, nkouvak, mtzamtzi, dfragkoulis, amenexis}@dind.uoa.gr,

^{*}AKATT S.A.

251, Agiou Dimitriou str, 173 42, Agios Dimitrios, Athens, Greece

^{6,7,8}{c_katsiavrias, k_katsiavrias, j_dimitropoulos}@akatt.gr

Abstract— This paper presents an edge-oriented implementation of a discrete-time bank of switching observers for chlorine estimation in a benchmark water distribution network (WDN). The approach combines local execution of multiple observers within a Node-RED runtime, deterministic data acquisition from a programmable logic controller (PLC) via Modbus Transmission Control Protocol (TCP), and cloud-supervised switching. The architecture enables parallel observer execution, telemetry arbitration across a single stream, and continuous adaptation of the runtime. The experimental validation is carried out in a hardware-in-the-loop setup in which a hydraulic and quality simulated model provide the measurement data for a physical PLC in real-time, allowing to assess the entire edge–cloud pipeline under realistic conditions of an industrial communication system.

Keywords— Soft Sensors, WDNs, Edge Computing, Switching Observers, Hardware-in-the-Loop Validation

I. INTRODUCTION

Edge computing plays a critical role in water quality monitoring within WDNs, particularly when timely detection and response are required. By executing estimation functions at the edge, continuous operation can be maintained even under intermittent cloud connectivity. Moreover, edge-based processing enables scalable deployment across geographically distributed infrastructures, minimizes communication load, and enhances resilience. Thus, it is particularly suitable for real-time soft-sensing applications in WDNs.

Edge and Internet-of-Things (IoT) architectures are increasingly adopted in WDNs to cope with latency, bandwidth, and scalability constraints that arise in large sensing deployments. An explicit edge–cloud monitoring paradigm for water quality in a WDN is given in [1], illustrating the practical value of distributing data processing between near-field nodes and cloud services. At a broader scale, massive IoT deployments for water distribution systems, and the associated communication, data-management, and evaluation challenges, are analyzed in [2], while a concrete large-scale smart-water IoT framework and its performance evaluation are studied in [3] and [4]. Complementary to platform-level frameworks, data-efficient monitoring and inference mechanisms motivated by energy and communication constraints are addressed in [5]. These works establish the enabling edge/IoT foundation for distributed monitoring in WDNs. However, explicit implementations of model-based soft sensors executed at the edge (and coordinated via cloud-supervised switching), to the best of our knowledge, have not as yet been studied.

This paper investigates the transition of chlorine soft sensors from simulation-oriented designs to deployable cyber-physical implementations by presenting an edge-oriented realization of a discrete-time bank of switching observers for a benchmark water distribution network. The proposed approach executes multiple observer instances locally within a Node-RED runtime, acquires measurements deterministically from a PLC through Modbus TCP, and performs cloud-supervised observer activation and selection. A gating stage enforces single-stream telemetry while preserving parallel observer execution at the edge, thereby enabling uninterrupted runtime adaptation. Experimental validation is performed in a hardware-in-the-loop configuration, where a real-time hydraulic and water-quality simulator generates sensor-like measurements that are injected into a physical programmable logic controller, allowing the full edge–cloud pipeline to be assessed under realistic industrial communication conditions.

II. BENCHMARK WATER DISTRIBUTION NETWORK AND OBSERVER FORM

Maintaining an adequate free chlorine residual throughout a water distribution network is a central operational requirement (see [6]-[13] and the references therein) because chlorine is the last protective barrier against microbial regrowth and contamination events after treatment. However, the residual is not spatially uniform and can change rapidly with hydraulic transients, demand patterns, water age, and reactions in the bulk water and at pipe walls. In practice, utilities cannot instrument every critical pipe section with in-line chlorine analysers, so soft sensors become essential to infer chlorine concentration at unmonitored locations from routinely available telemetry at the network's discrete-time sampling rate. For the derivation of the soft sensor, here a physics-based approach will be used (see [13]-[14]). Toward this goal, a mathematical model of the water distribution network is required to provide the causal link between hydraulics and chlorine transport / decay, to support systematic estimator / observer design with analysable convergence properties, and to avoid relying solely on black-box predictors whose reliability can degrade under varying operating conditions. For later-on observer design purposes, a benchmark water distribution network of the form presented in Figure 1 will be used. The configuration consists of a star-shaped network composed of three reservoirs interconnected through three pipes to a single internal junction while aggregated consumer demand is also introduced at the node.

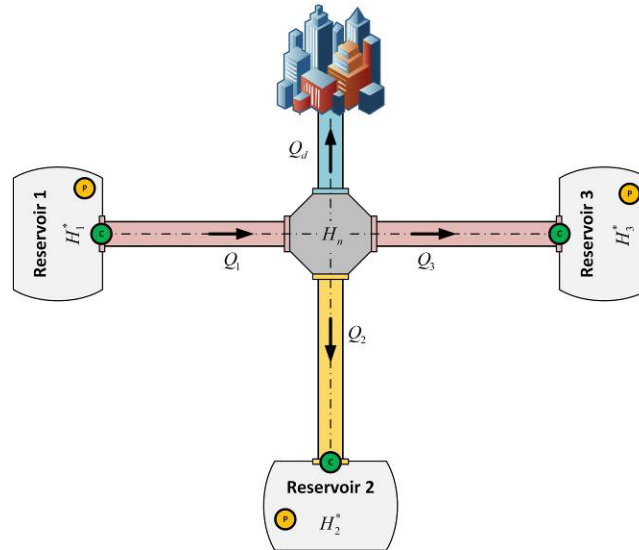


Fig. 1 Benchmark Water Distribution Network Setup

Applying series of computations (see [14]), the dynamic model of the WDN is

$$E\dot{x} = \Gamma(x) + Z(x)u \quad (1)$$

where $x = [x_1 \ \dots \ x_{10}]^T = [Q_1 \ Q_2 \ Q_3 \ H_n \ c_{1,1} \ c_{1,2} \ c_{2,1} \ c_{2,2} \ c_{3,1} \ c_{3,2}]^T$, $u = [u_1 \ \dots \ u_5]^T = [H_1^* \ \dots \ H_2^* \ H_3^* \ Q_d \ c_0]^T$, Q_i ($i \in \{1,2,3\}$), H_i^* ($i \in \{1,2,3\}$) is boundary (actuatable) pressure heads imposed by the reservoir i , Q_i is the volumetric flow rate in conduit i , Q_d is the aggregated consumption flow rate, $c_{i,j}$ is the chlorine concentration at section j of conduit i and c_0 is the influent chlorine concentration to the network from reservoir 1, while E , $\Gamma(x)$ and $Z(x)$ are appropriate system matrices depending upon the geometry of the network and water flow properties. Based on the model (1), a linear approximant of the following form can be produced

$$E\delta\dot{x} = J_s(\bar{x}, \bar{u})\delta x + Z(\bar{x})\delta u \quad (2)$$

where \bar{x} and \bar{u} are the vectors of nominal values of the state and input vectors of the nonlinear model. Clearly, the system matrices J_s and Z depend upon the nominal values of the system. Following a quasi-static approach, i.e. assuming that the water flow model is static and applying zero-order-hold interpolation, the discrete time dynamics of the chlorine concentration in the network take on the form

$$\delta q(k+1) = A_d(\bar{u})\delta q(k) + B_d(\bar{u})\delta u(k) \quad (3)$$

The discrete time linear approximant system matrices in (3) will be used for observer / chlorine soft sensor design purposes. Following the procedure in [14], the soft sensor comprises of a predefined bank of observers of the form

$$\mathfrak{S}:\delta\hat{q}(k+1) = (A_d(\bar{u}) - L(\bar{u})C_d)\delta\hat{q}(k) + B_d(\bar{u})\delta u(k) + L(\bar{u})\delta y(k) ; \delta\hat{q}(0-) = \delta\hat{q}_0 \quad (4)$$

where $L(\bar{u})$ is an observer matrix to be determined by the designer. Clearly, each observer depends upon the nominal values of the system. Hence, as the system diverges from the operating point, the estimation error is bound to increase and hence switching between the observers of the soft sensor is necessary. In what follows, a system architecture for edge implementation of the soft sensor system (indicatively consisting of four individual observers) will be presented.

III. SYSTEM ARCHITECTURE - EDGE EXECUTION WITH CLOUD-SUPERVISED SWITCHING

The proposed platform implements a complete edge-to-cloud pipeline for soft sensing in water distribution networks, with emphasis on runtime operation (see Figure 2). The architecture is centred on a Node-RED edge runtime that a) acquires industrial measurements from a programmable logic controller, b) executes a bank of discrete-time observers in parallel, and c) forwards to the cloud the state variables estimate, produced by the bank, as well as measurements of the system output to be used to trigger a supervisory switching signal.

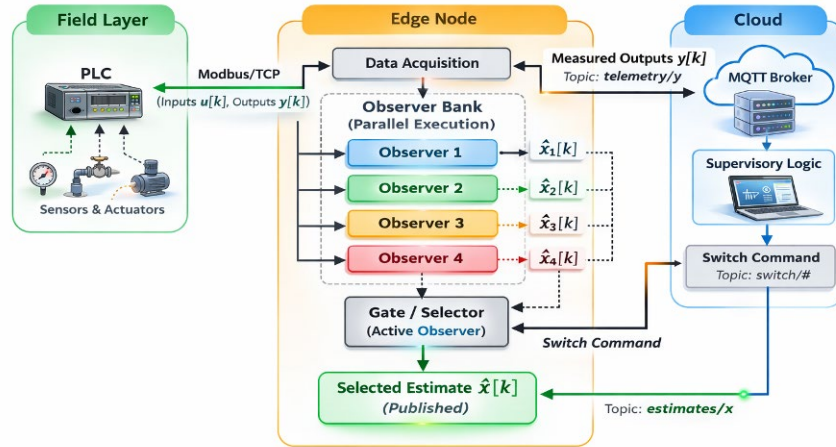


Fig. 2 System Architecture

This structure (see Figure 3) allows the computationally critical estimation tasks to be performed locally at the edge, while retaining a cloud layer for supervision, monitoring, and remote adaptation of the estimation mode during operation.

At the measurement interface, the edge runtime connects to the field layer through a Modbus TCP link to the PLC. A Modbus client node cyclically reads a block of seven holding registers, which arrive in Node-RED as a sequential array (see Figure 4). A parsing stage then unpacks this array into individual scalar signals, assembling a five-dimensional control-input vector and a two-dimensional measured-output vector that drive the observer-bank execution. The use of a single block read reduces protocol overhead, yields a predictable sampling pattern, and simplifies time alignment of signals at the edge.

In addition to the direct PLC measurements, the architecture supports ingestion of auxiliary quantities through Message Queuing Telemetry Transport (MQTT). A topic hierarchy of the form Measurements/# (see Figure 5) is used to deliver values needed by the edge logic (i.e. operating-point values, configuration and other parameters that are updated less frequently than the main sampling loop). Each topic is isolated through a dedicated filter, and the most recent value is stored in Node-RED's flow context. MQTT updates modify cached context values whenever they arrive, whereas observer execution reads a snapshot of those cached values at the time a new measurement tick occurs.

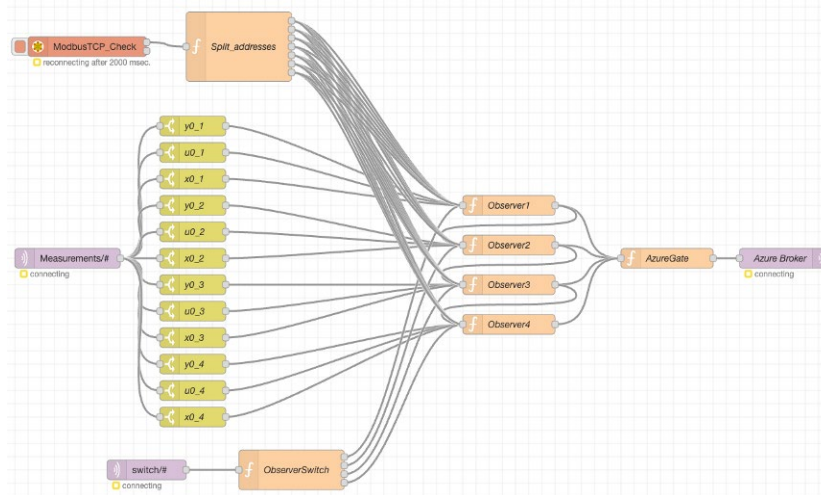


Fig. 3 Node-Red flows Overview

```

1 var u1 = {payload:msg.payload[0]};
2 var u2 = { payload: msg.payload[1] };
3 var u3 = { payload: msg.payload[2] };
4 var u4 = { payload: msg.payload[3] };
5 var u5 = { payload: msg.payload[4] };
6 var y1 = {payload:msg.payload[5]};
7 var y2 = { payload: msg.payload[6] };
8 return [u1, u2, u3, u4, u5, y1, y2];

```

Fig. 4 Modbus/TCP Splitting Registers Function

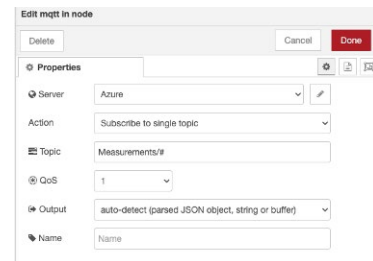


Fig. 5 Measurements MQTT Topic

The estimation layer is implemented as a bank of four observers that run in parallel within the same edge runtime. The observers share the same execution template and message interface, but use different internal parameterizations, corresponding to distinct locally valid models (e.g., different operating regions). From a system point of view, the key architectural feature is that the observer bank can be executed at the edge.

A key element of the architecture is the explicit separation between estimation and selection. Switching is realized through a dedicated supervisory path. Node-RED subscribes to a switch/# topic, interprets received commands as an active-observer index, and stores the selected index in flow context. Importantly, reception of a switch command does not interrupt the acquisition pipeline or the observer-bank execution. Instead, it updates only the selection state that will be used downstream to determine which estimate stream is published. This design supports remote adaptation of the soft sensor during runtime while avoiding disruptions of the edge process, which is critical when the edge node must remain operational under intermittent connectivity or during frequent operating changes.

To enforce single-stream publication, a gating mechanism is placed immediately before the cloud-facing output. Each observer tags its outgoing message with an identifier, and a gate compares this identifier against the stored active-observer index.

On the cloud interface, the selected estimate stream is published to a Microsoft Azure MQTT broker. Architecturally, this yields three distinct communication paths: a) the field-to-edge data path via Modbus TCP, b) the cloud-to-edge supervisory path via MQTT (both Measurements/# and switch/#), and c) the edge-to-cloud telemetry path via gated MQTT publication. The separation of these paths improves modularity and operational resilience since loss or delay of supervisory messages does not halt estimation at the edge, and intermittent cloud connectivity does not interfere with acquisition from the programmable logic controller.

IV. EXPERIMENTAL RESULTS

The experimental validation will be conducted under a hardware-in-the-loop configuration in order to emulate realistic field conditions while preserving full control over the underlying plant dynamics. In this setup, a computer-based simulation environment implements the water distribution network model and generates, in real time, the measurable outputs that would normally be produced by physical sensors. The

execution from that moment forward is completely real. As already mentioned, the system implements a bank consisting of four observers. These observers are designed on the following input operating points: a) $\bar{u}_1 = 20[\text{m}]$, $\bar{u}_2 = 15[\text{m}]$, $\bar{u}_3 = 10[\text{m}]$, $\bar{u}_4 = 0.006[\text{m}^3/\text{s}]$, $\bar{u}_5 = 2[\text{mg/l}]$, b) $\bar{u}_1 = 26[\text{m}]$, $\bar{u}_2 = 13[\text{m}]$, $\bar{u}_3 = 8.4[\text{m}]$, $\bar{u}_4 = 0.005[\text{m}^3/\text{s}]$, $\bar{u}_5 = 1.9[\text{mg/l}]$, c) $\bar{u}_1 = 35[\text{m}]$, $\bar{u}_2 = 10[\text{m}]$, $\bar{u}_3 = 6[\text{m}]$, $\bar{u}_4 = 0.0035[\text{m}^3/\text{s}]$, $\bar{u}_5 = 1.75[\text{mg/l}]$, d) $\bar{u}_1 = 44[\text{m}]$, $\bar{u}_2 = 7[\text{m}]$, $\bar{u}_3 = 3.6[\text{m}]$, $\bar{u}_4 = 0.002[\text{m}^3/\text{s}]$, $\bar{u}_5 = 1.6[\text{mg/l}]$.

It can readily be verified that the accuracy regions corresponding to the above operating points satisfy the dense web principle. Following the procedure in [14], the corresponding observer matrices are derived to be

$$L(\bar{u}) = \begin{bmatrix} -0.0044 & -0.002 & 0.092 & 0.5448 & 0.0028 & 0.0021 \\ 0.0021 & -0.0032 & 0.0178 & 0.009 & 0.0055 & 0.0926 \end{bmatrix}^T$$

$$L(\bar{u}) = \begin{bmatrix} -0.0017 & -0.0013 & 0.0155 & 0.2205 & -0.0009 & 0.0019 \\ 0.0021 & -0.0033 & 0.0043 & 0.0054 & 0.0022 & 0.0044 \end{bmatrix}^T$$

$$L(\bar{u}) = \begin{bmatrix} 0.0005 & 0.0003 & 0.0035 & 0.0877 & 0.0023 & -0.0004 \\ 0.0021 & -0.0023 & -0.0023 & 0.0025 & 0.0009 & 0.0176 \end{bmatrix}^T$$

$$L(\bar{u}) = \begin{bmatrix} 0.0013 & 0.0002 & 0.002 & 0.0434 & 0.0017 & -0.0006 \\ 0.0025 & -0.0002 & -0.0019 & 0.0008 & -0.0001 & 0.0083 \end{bmatrix}^T$$

Regarding the sampling period it will be assumed that it is equal to 1 minute, being typical for online chlorine concentration measurements in WDNs, while for the network simulation, the data presented in [14] will be used. Initializing the observer with zero estimation error, in Figures 6 to 9 chlorine concentration plots for the non-measurable are produced for transitions between five points. It can readily be verified that “measurements” and soft sensor estimations during the experiment are almost identical.

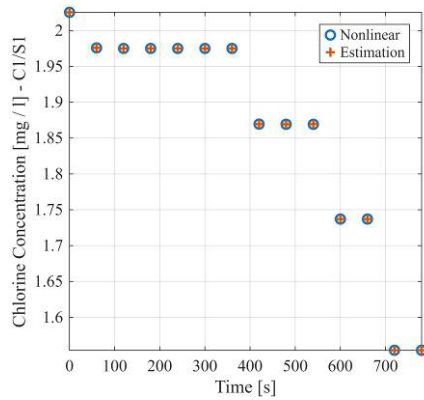


Fig. 6 Chlorine concentration Vs estimation in the section 1 of pipe 1

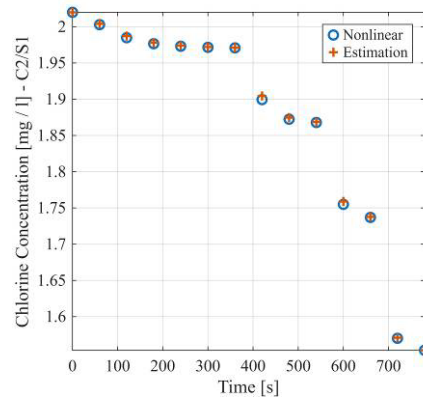


Fig. 8 Chlorine concentration Vs estimation in section 1 of pipe 2

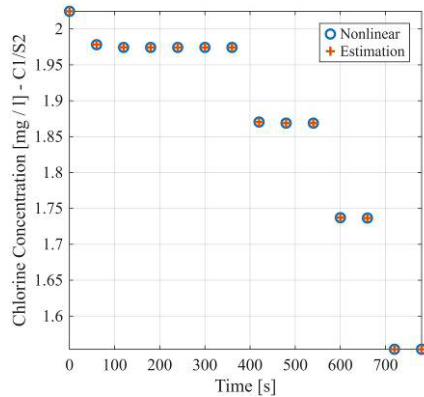


Fig. 7 Chlorine concentration Vs estimation in section 2 of pipe 1

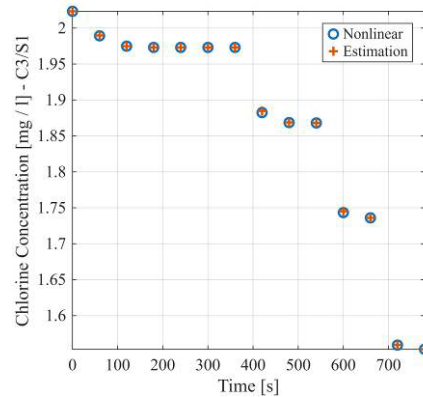


Fig. 9 Chlorine concentration Vs estimation in the section 1 of pipe 3

V. CONCLUSIONS

An edge-oriented implementation of a physics-based chlorine soft sensor for water distribution networks, based on a discrete-time bank of switching observers designed around multiple operating points, has been introduced. The edge–cloud architecture that has been proposed executes all observers continuously within a Node-RED runtime, acquires measurements from a programmable logic controller via Modbus TCP, and uses cloud-supervised switching through MQTT. A gating mechanism ensures that only the cloud-selected observer publishes estimates, enabling seamless runtime adaptation without interrupting acquisition or estimation. Validation has been conducted in a hardware-in-the-loop setting, confirming accurate chlorine reconstruction during operating-point transitions and demonstrating the feasibility of deploying switching soft sensors at the network edge. A general edge implementation framework for a parametric number of step wise safe observers is currently under investigation.

ACKNOWLEDGMENT

This work was conducted within “SUB1.1 Clusters of Research Excellence (CREs)” Action of the NRRP “Greece 2.0”, co-funded by Greece and EU – NextGenerationEU, under the project YII3TA-0560428-DeSw.

REFERENCES

- [1] E. Q. Shahra, W. Wu, S. Basurra, and A. Aneiba, “Intelligent Edge-Cloud Framework for Water Quality Monitoring in Water Distribution System,” *Water*, vol. 16, no. 2, Art. no. 196, 2024.
- [2] A. Pagano, D. Garlisi, I. Tinnirello, F. Giuliano, G. Garbo, M. Falco, and F. Cuomo, “A survey on massive IoT for water distribution systems: Challenges, simulation tools, and guidelines for large-scale deployment,” *Ad Hoc Networks*, vol. 168, Art. no. 103714, 2025.
- [3] A. Pagano, D. Garlisi, F. Giuliano, T. Cattai, R. J. L. Taloma, and F. Cuomo, “Introducing and evaluating SWI-FEED: A smart water IoT framework designed for large-scale contexts,” *Computer Communications*, vol. 237, Art. no. 108146, 2025.
- [4] A. Pagano, D. Garlisi, F. Giuliano, T. Cattai, and F. Cuomo, “SWI-FEED: Smart Water IoT Framework for Evaluation of Energy and Data in Massive Scenarios,” in *Proc. IFIP Networking Conf.*, 2024, pp. 583–585.
- [5] T. Cattai, S. Colonnese, D. Garlisi, A. Pagano, and F. Cuomo, “GraphSmart: A Method for Green and Accurate IoT Water Monitoring,” *ACM Trans. Sensor Netw.*, vol. 20, no. 6, Art. no. 130, 2024.
- [6] World Health Organization, *Guidelines for Drinking-water Quality: Fourth Edition Incorporating the First Addendum*. Geneva, Switzerland: World Health Organization, 2017.
- [7] World Health Organization, *Principles and Practices of Drinking-Water Chlorination: A Guide to Strengthening Chlorination Practices in Small- to Medium-Sized Water Supplies*. New Delhi, India: World Health Organization, Regional Office for South-East Asia, 2017.
- [8] Centers for Disease Control and Prevention, *About Water Disinfection with Chlorine and Chloramine*. Feb. 14, 2024.
- [9] R. A. Li, J. A. McDonald, A. Sathasivan, and S. J. Khan, “Disinfectant residual stability leading to disinfectant decay and by-product formation in drinking water distribution systems: A systematic review,” *Water Research*, vol. 153, pp. 335–348, Apr. 2019.
- [10] M. E. Bledsoe, A. Goel, M. Adelgren, T. M. LaPara, and R. M. Hozalski, “Effect of chloramine disinfection of community water system on Legionnaires’ disease outbreak, Minnesota, USA, 2024,” *Emerging Infectious Diseases*, vol. 32, no. 1, pp. 74–81, Jan. 2026.
- [11] M. Djerioui, M. Bouamar, M. Ladjal, and A. Zerguine, “Chlorine soft sensor based on extreme learning machine for water quality monitoring,” *Arabian J. Sci. Eng.*, vol. 44, no. 3, pp. 2033–2044, 2019.
- [12] M. A. Ogundero, T. A. Fashanu, F. O. Agunbiade, K. Orolu, A. A. Yinusa, U. A. Daudu, and M. O. H. Amuda, “A soft sensor-based inference engine for water quality assessment and prediction,” *Air, Soil and Water Research*, vol. 18, no. 1, pp. 1–12, 2025.
- [13] N. D. Kouvakas, F. N. Koumboulis, M. P. Tzamtzi, and D. G. Fragkoulis, “Toward the design of chlorine soft sensors via stepwise safe switching observers for a primary water distribution network,” in *Artificial Intelligence Applications in Engineering and Medicine: Proceedings of the 2nd Olympiad in Engineering Science*, D. Pavlou, H. Adeli, and J. Correia, Eds., Structural Integrity, vol. 36. Cham, Switzerland: Springer, 2026.
- [14] N. D. Kouvakas, F. N. Koumboulis, M. P. Tzamtzi, D. G. Fragkoulis, K. Katsiavrias, K. Katsiavrias and J. Dimitropoulos, “On the Discrete Time Stepwise Safe Switching Soft Sensor Design for Water Distribution Networks,” *12th International Conference on Control Engineering & Information Technology*, Istanbul, Turkey, April 27-30, 2026.