

The Resilmesh Architecture: Situation Aware Enabled Cyber Resilience for Dispersed, Heterogenous Cyber Systems

Jorge Bernal Bernabé*, Brian Lee†, Martin Husák‡, Lukáš Sadlek‡, Branka Stojanović§, Michael Somma§, Jorgeley Inacio de Oliveira¶, Ekam Puri Nieto*, Pablo Fernandez Saura*, Antonio Skarmeta*, Vinh Hoa La||

* University of Murcia, Spain, Email: {jorgebernal, ekam.purin, pablofs, skarmeta}@um.es

†TUS, Technnical University of Shannon, Email: {brian.lee}@tus.id

‡Institute of Computer Science, Masaryk University, Emails: husakm@ics.muni.cz, sadlek@mail.muni.cz

§JOANNEUM RESEARCH, Austria, Email: {branka.stojanovic, michael.somma}@joanneum.at

¶Silentpush, Ireland, Email: jorgeley@silentpush.com

||Montimage, France, Email: vinhhoa.la@montimage.eu

Abstract—Cyber systems (CyS) are becoming more and more complex as they are comprised of several infrastructure layers, heterogeneous technologies and dispersed deployments over wide geographical areas (cloud/edge/endpoint) that facilitates multiple attack entry points (vectors). At the same time, CyS attacks are constantly evolving and have become more complex and sophisticated. To address these issues, the ResilMesh architecture aims to provide critical infrastructure security teams with a greater cyber resilience capability by improving cyber resilience using Cyber Situational Awareness (CSA) based security orchestration and analytics framework. The framework enables organizations to achieve real-time defense, reducing attack surface impact by developing tools to combat complexity, disperse infrastructure, delivering flexible placement of security controls across the CyS infrastructure. The architecture combats Advanced Persistent Threat (APT) sophistication by leveraging advanced AI algorithms and tools for early and ongoing attack detection and prediction and improved situation. This paper presents the Resilmesh architecture, a first PoC implementation, as well as an evaluation of the Resilmesh capabilities to detect and mitigate APTs.

Index Terms—Cybersecurity, AI, Resilience, Cyber systems

I. INTRODUCTION

Cyber systems (CyS) are becoming increasingly complex due to the integration of multiple infrastructure layers, heterogeneous technologies, and dispersed deployments across cloud, edge, and endpoint environments. This complexity exposes organizations to a wide range of attack vectors and sophisticated cyber threats. Despite advancements in cybersecurity, existing solutions face several challenges, including limited real-time awareness, as traditional security tools lack comprehensive cyber situational awareness, making it difficult to detect and respond to threats dynamically. Additionally, scalability and adaptability issues arise as conventional SIEMs and security platforms struggle to operate efficiently in heterogeneous and distributed infrastructures. Another key limitation is the lack of explainability in AI models, making it challenging for security teams to validate alerts and responses, ultimately reducing trust in automated detections. Lastly, most current mitigation strategies rely on static response mechanisms, using predefined

rules that fail to adapt to evolving attack patterns, leaving systems vulnerable to sophisticated and persistent cyber threats.

To fill this gap, Resilmesh is designing and developing a novel Cyber Situational Awareness (CSA) based security orchestration and analytics toolset to enable organisations achieve real time defence of essential business functions. ResilMesh will help: reduce CyS attack surface impact by developing tools to combat complexity (better visibility of assets and services and their dependencies, heterogeneity interoperability and extensibility); disperse infrastructure, flexible placement of security controls across the CyS infrastructure; combat Advanced Persistent Threat (APT) sophistication by developing artificial intelligence (AI) algorithms and tools for early and ongoing attack detection and prediction and improved situation and risk awareness.

The overall Resilmesh architecture consists of several planes to aggregate data and events from multiple heterogeneous sources, process and analyse in real time data to manage anomaly detection, event correlation and alerting, and attack detection and prediction. The architecture is endowed with a threat awareness plane for cyber situational assessment capability that in turn supports the operational plane for deciding and enforcing mitigation actions to be taken to respond to a detected incident, and orchestrate these actions as Course of Action (CoA) playbooks.

This paper describes the Resilmesh architecture, the first PoC implementation as well as the first performance evaluation conducted to validate the approach. The results achieved highlight the effectiveness of our approach in identifying anomalies, correlating complex attack patterns, and dynamically adapting security responses, reinforcing the practical applicability of CSA-driven cybersecurity solutions for CPS.

The rest of this paper is organized as follows: Section II reviews the state of the art. Section III presents the ResilMesh architecture. Section IV describes the implementation of the first Proof of Concept (PoC). Section V discusses the performance evaluation. Section VI concludes the paper.

II. STATE OF THE ART

Cyber attacks on critical infrastructure have surged due to rapid digitalization, shifting security concerns beyond IT networks to sectors like energy and transportation [1]. Cyber-physical systems (CPS), integrating IoT devices and industrial control systems (ICS), introduce new vulnerabilities [1]. Recent incidents, where attackers exploited outdated access rights to manipulate critical processes, highlight the urgent need for preemptive detection systems [2].

In CPSs, anomaly localization is as crucial as detection, ensuring swift and targeted responses to cyber threats [3]. While intrusion detection research has primarily focused on improving accuracy and reducing detection time [4], localization remains underexplored [1]. Existing approaches either require expert intervention [5] or rely on complex neural architectures like GNNs [6]. Some modular methods attempt to divide networks into zones but fail to align with hardware functionalities, limiting decentralization [7].

Centralized anomaly detection, while effective, demands substantial energy, contributing to a high carbon footprint [8]. Edge computing mitigates these issues by processing data locally, reducing latency and energy consumption. Federated learning methods, such as Deep Autoencoder for Federated Learning (DAEF), show promise in enhancing anomaly detection while significantly lowering environmental impact [9].

III. ARCHITECTURE

The ResilMesh architecture consists of six planes as displayed in Fig. 1 and discussed in this section.

A. The Infrastructure Plane

This plane includes both physical and virtual network elements responsible for tasks such as forwarding the traffic according to rules received by the Network Controllers through the southbound API. In virtualized scenarios, this infrastructure includes the cloud computing technologies (i.e., computing, storage and networking) to deliver virtual Infrastructure-as-a-Service (IaaS). Special controllers are used to control devices in more heterogeneous scenarios including IoT.

Security probes or sensors are deployed in the infrastructure to support monitoring the managed system. For instance, Packetbeat agents is an analyser that gathers and sends data from hosts and containers, FileBeat agents for collecting logs from security devices, cloud, containers, hosts. In addition, special purpose probes for OT can be shipped for gather logs, and data from the OT network protocols and system logs.

B. Aggregation Plane

The aggregation plane is the initial pre-processing step needed to ensure the data coming from the Infrastructure Plane flows properly throughout the pipeline. It collects, aggregates, and normalises data and events from multiple heterogeneous sources including logs, IDS, network sources, AI models, etc. It contains a rich set of data/event (including network traffic) filtering, fusion, logging, storage and forwarding functions.

It is comprised of three functional sub-layers. The *Data Transformation Layer* receives the data from on-device agents and other sensors and filters and aggregates the incoming raw data. It contains the Event Aggregation (EA) functional component, which ingests event logs from multiple sources, transforms them via various operations (routing / logging / fusion / filtering / augmentation / reduction / monitoring), and then outputs the data to one or more destinations via the streaming layer. Outputted events are transferred to the *Data Normalisation Layer* which adapts the aggregated data to a common data format scheme - the Elasticsearch Common Schema (ECS). It transforms data from multiple disparate formats coming from different sources, to a single common format that can then be used for analytics, visualisation, reporting, etc. The *Data Streaming Layer* distributes the normalised data to the upper layers and also performs further processing on the data streams, if required. It contains the Message Broker (MB), Event Stream Processing (ESP) and Event Enrichment (EE) functions. Message brokers can be used to route and deliver messages to the required destinations. EE enriches the events with contextual information from the enrichment API that scans, clusters, scores and enriches the global IPv4 range in a first-party database that outputs Indicators Of Future Attack (IOFA) – domain, IP and URL data that explains the relationship between billions of observable data points across the internet.

C. The Collaboration Mesh Plane

This is the Collaboration Mesh connectivity underlay as described previously in the Architecture Foundations section. Note that the SM is an optional component and is deployed only for specific scenarios. Also while a micro-service based architecture (with container RO) is the preferred application software architecture, other approaches such a bare-metal may be used when required.

D. The Threat Awareness Plane

The Threat Awareness layer includes an **AI-based Anomaly Detection (AID)** module to detect anomalies for any type of IT or OT application and/or network protocol. Models can be placed at the endpoint/edge or in the cloud as required. While ResilMesh supports the use of any ML or AI techniques to develop models, the focus is on the use of deep learning techniques. It evaluates different multi-view deep learning approaches, such as multi-view fusion-based methods and multi-view alignment-based methods, to deal with the intrinsic heterogeneity of mixed technology domains. This layer also includes a federated learning platform - the Privacy FL Framework (PPFL) - to support federated training of the different types of deep learning algorithms mentioned above. The platform supports both horizontal and vertical federated learning, as well as a combination of both. The final output of these detectors will be events sent to the AIC and/or SIEM correlation function, or as features to be passed on to other models. **The AI Correlator (AIC)** provides AI methods to correlate security events includes with application to i)

classify security events for event detection, event grouping, and event pattern extraction, ii) Intrusion detection which deals with multi-stage and targeted attacks or anomaly detection to notify the security administrator about misuses and deviations from normal behaviour, respectively and iii) Intrusion/attack projection based on incoming events, which allows early detection of intruder targets. AI correlation may be embedded within other components e.g. such as NSE or NDR or may exist as a stand alone component. The **SIEM** is a central component in ResilMesh and provides a number of capabilities including event logging (through OpenSearch), event correlation as well as a range of XDR capabilities. The **THF (TTP-based Hunting and Forensics)** TTP-based threat hunting focuses on identifying adversaries by their tactics, techniques, and procedures—the “how” of their operations rather than the “what” or “when.” This approach leverages the ATT&CK framework developed by MITRE. THF supports the use of TTP-based hunting techniques for cyber attack investigation. THF includes a UI where the user can carry out the mentioned hunting and analysis features.

E. The Situation Assessment Plane

The Situation Assessment plane contains a set of functional components that collectively provide cyber situational assessment capability to ResilMesh - and which, together with the Threat Awareness plane implements CSA in ResilMesh. These components are:

1) *Infrastructure and Service Information Model (ISIM)*: Captures and represents all the entities of interest in the environment including devices, networks applications (services) users and data. The information model interconnects the pieces of information on the assets in the environment. When deploying the system in a new environment there is a need to fill the database with data from external sources. For each category of assets, their enumeration will be collected from existing databases, repositories, or service or custom tools.

2) *Cyber Asset Attack Surface Management (CASM)*: Ingests data from the ISIM. The principal role of the CASM is to monitor the organisation’s internal and external attack surface and security posture. It provides an interactive query capability to allow operators to determine cyber security posture based on the relationship between assets – based on ISIM. It also carries out domain (and sub-domain) enumeration to discover so-called ‘shadow-IT’ and to identify and manages threats discovered in internet-facing assets using independent scans of the organisation attack surface. CASM can detect changes of asset status and trigger required actions or alerts to the Mitigation Manager (MM) as required.

3) *Network Detection and Response (NDR)*: Continuously monitors and analyzes raw enterprise network traffic using a hybrid detection approach that combines signature-based and anomaly-based techniques. The signature-based approach detects known threats by matching traffic patterns against predefined attack signatures and threat intelligence feeds, enabling rapid identification of previously observed malicious activities. Meanwhile, the anomaly-based approach builds a dynamic baseline of normal network behavior, identifying

deviations that may indicate emerging or unknown threats. NDR ingests data from multiple sources, including ISIM (Information System Inventory Management) for asset information and SIEM (Security Information and Event Management) for security events and alerts. When a potential threat is detected—either through signature matching or behavioral anomaly detection—NDR generates alerts for security teams and can invoke automated mitigation actions via the Mitigation Manager, ensuring swift and proactive threat response. This dual-layered detection strategy enhances the adaptability, accuracy, and effectiveness of the ResilMesh security framework.

4) *Critical Service Awareness (CSA)*: Provides hierarchical risk assessment to aggregate infrastructure risk into a risk for the critical service or business mission (CS/M). It implements tools to assess such risks. To achieve this, the component uses data stored in ISIM (information model) and forward the outputs to the NSE (Network Situation Evaluation).

5) *Network Situation Evaluation (NSE)*: Provides a risk assessment of the overall network based on input from other functions and can also project the attack intensity for the network. It provides visualisation of both current and future network risk status. It uses inputs from a number of tools to do this including the ISIM, SIEM, CSA and NDR.

F. The Security Operations Plane

The Security Operations Plane contains a set of functional components aiming to decide the most suitable mitigation actions against a detected incident, orchestrate these actions as CoA playbooks and analyse collected information to identify adversaries by tactics/techniques carried out during the incident. These components are:

1) *Mitigation Manager (MM)*: Responsible for deciding which mitigation actions, if any, are taken to deal with an incident. MM relies on CSS and NSA to get information about mission, risk, and network status - factors in the mitigation decision process. MM utilizes a dynamic algorithm to select the best mitigation actions to enforce. Once mitigation actions have been decided, MM dispatches the appropriate mitigation through the Playbooks Tool component.

2) *Playbooks Tool (PT)*: Triggers the orchestration of selected mitigation playbook(s). Based on tools such as Shuffle, PT launches CoA playbooks that execute workflows which contain a set of several actions to mitigate incidents, such as OpenC2 commands. Some mitigation actions include network filtering, quarantining hosts or updating outdated software. Workflow actions also can enrich SIEM adding new rules (like yara or sigma rules) or execute other tools like Security Orchestrator that can enforce security or privacy policies.

3) *Workflow Orchestrator (WO)*: Supports PT by orchestrating and automating complex actions. WO performs playbook actions which are difficult to implement in PT when requested by PT within a CoA playbook execution. WO then uses a subset of tools from Target Actuators to carry out the requested actions. PT can request an action to the WO via a REST/gRPC API, and then WO can provide information about the progress

and outcome of the performed actions, establishing a constant feedback loop.




Fig. 1. Resilmesh functional architecture.

IV. IMPLEMENTATION

A. Framework

ResilMesh is implemented using a microservices architecture, ensuring modularity, scalability, and efficient processing of security events. The framework is structured into three key foundational layers, each playing a distinct role in event collection, normalization, enrichment, and correlation. Vector serves as the observability pipeline tool, responsible for collecting and normalizing logs from diverse sources such as Syslog, Windows Event Logs, and other network telemetry streams. NATS functions as the message broker, efficiently aggregating and distributing events based on their processing state (e.g., normalized, enriched, or correlated). Each component in the framework interacts with NATS to retrieve the relevant security events, perform specific processing tasks, and subsequently push the enriched or analyzed data back into the pipeline in a normalized format. Wazuh acts as the SIEM solution within ResilMesh, ingesting security events from the Infrastructure Plane after they have passed through the pipeline. A key feature of Wazuh's integration in ResilMesh is its ability to correlate security events with threat intelligence feeds from MISP. Moreover, Wazuh's seamless integration with OpenSearch provides advanced capabilities for log visualization, anomaly

detection, and forensic investigation, empowering analysts with deep insights into evolving cyber threats. Enrichment Module (EM) provided by Silent Push, is deployed as a containerized service within the framework. This module augments security events with critical contextual information like scores, reputations, PADNS data, etc, enhancing their relevance for analysts and improving threat detection accuracy. The enrichment process involves multi-dimensional data augmentation, depending on the nature of the event (IP, domain, subnet, ASN, etc.) The inclusion of this enrichment module significantly enhances situational awareness, enabling threat analysts to draw actionable insights directly from the enriched event data.

B. Context awareness

Context awareness, in the context of the Resilmesh system, covers the situation assessment plane. Its implementation builds upon the combination of existing tools and their improvements and novel additions. Namely, Resilmesh uses and enhances the components of CRUSOE, a toolset aiming data collection, analysis, and visualization enabling CSA [10].

The **ISIM** component use Neo4j graph database and GraphQL API. The database content is structured according to the data model inferred from the CRUSOE data model [10] with several differences. First, ISIM aims to represent all asset types as defined by the Cyber Defense Matrix¹, i.e., not only devices, applications, and user, but also networks and data. Second, ISIM adds new layers representing hardware and physical location, namely to support industrial robotics and other OT use cases. Finally, ISIM adds a labeling attributes to distinguish types of network nodes based on MITRE D3FEND's Digital Artifact Ontology (DAO)² and custom labels for internal purposes, e.g., which sensor provided which piece of information.

The **CASM** aims at being extensible set of data collection tools. In the initial implementation, we employed an existing attack surface management toolset, EasyEASM³ to conduct the tasks of active scanning of network hosts and services and enumeration of subdomains. All scans produce CPE identifiers or discovered hosts and services, which are then matched with vulnerabilities enlisted in NVD. The corresponding CVEs and their CVSS scores (in version 4) are then fetched to ISIM. Running the CASM is orchestrated via Temporal⁴, which enables detailed scheduling of scans and control over the workflow, including the repetition of failed tasks. Scans are executed by workers that can be deployed inside of outside the protected network, thus performing internal or external attack surface monitoring. In the future work, we aim at incorporating and orchestrating other data collection tools that would either collect additional (e.g., IoT or OT scanners) or more precise data (e.g., dedicated vulnerability scanners), depending on the needs of deployment environment.

¹<https://cyberdefensematrix.com/>

²<https://d3fend.mitre.org/dao/>

³<https://github.com/goldencybersec/EasyEASM>

⁴<https://temporal.io/>

The CSA component currently implement only the data management, which follows the approach proposed in CRUSOE toolset [10]. The critical services or business missions can be represented as a tree structure, in which nodes representing cyber assets (devices or services) or children of nodes representing the missions or its logical component. In addition, AND/OR nodes can be placed between missions and assets to model situations, in which the mission relies on multiple assets simultaneously (AND) or one of interchangeable assets (OR). Tree structure can be provided as JSON files, stored in ISIM, and visualized in the dashboard.

C. AI-based incident detection

ResilMesh utilizes both centralized and edge-based anomaly detection to balance computational efficiency with responsiveness. The centralized approach processes large volumes of data in the cloud, leveraging aggregated insights from multiple sources to detect sophisticated attack patterns. In contrast, the edge-based approach operates on local devices, reducing latency and enabling real-time anomaly detection without excessive data transmission. This distributed detection framework enhances resilience and adaptability in OT and ICS environments, and addresses privacy preserving needs.

One of key features of ResilMesh's AID module is its multi-view learning capability, which integrates diverse data types and sources. These include IT network traffic, OT sensor readings, and system log files. Different multi-view methods are being considered in the project. E.g. multi-view fusion methods extract high-level representations by combining different data modalities, while multi-view alignment techniques ensure consistency across heterogeneous inputs.

The evaluation focused on anomaly detection performance using a decentralized, edge-computing-based approach. The experiment leveraged autoencoder-based deep learning models trained on segmented data corresponding to different edge areas within an ICS network. The testbed included multiple industrial devices, with models deployed both centrally and on edge nodes, and utilized OT sensor readings.

The centralized model benefitted from the availability of aggregated data, which allowed for a broader view of system behaviour. However, it required higher computational resources and incurred latency due to data transmission requirements. Despite these challenges, the centralized approach remains a valuable component of a hybrid detection strategy.

The proof-of-concept results underscore the potential of distributed AI-based anomaly detection in OT and ICS environments. Future iterations will refine model architectures and explore hierarchical deep learning and multi-view techniques to further enhance detection accuracy and operational efficiency, ensuring robust cybersecurity for industrial systems.

D. Mitigation

For its main decision-making algorithm, the Mitigation Manager uses a graph-based approach where an attack is divided into steps, each edge representing an attacker's action as progress towards a goal such as exfiltrating data or causing

damage to critical infrastructure. Incoming alerts are mapped to potential attack graphs - these attack graphs help the Mitigation Manager infer additional information, such as the likely objectives of an attacker, but are mainly used to track the steps an attacker has taken and predict their next steps.

The Mitigation Manager dispatches mitigation strategies from both the PT and the WO components. These strategies have associated metadata with detailed information on their cost of execution, effects on the system/network they're applied to, and any pre-requirements for the strategy to be applicable.

The mitigation process consists of the following steps: parsing the alert, locating matching attack graphs, evaluating graph conditions and advancing graph progress, updating risk values, selecting nodes likely to be triggered, locating applicable mitigation strategies, and dispatching them.

V. PERFORMANCE EVALUATION

A. Context awareness

We evaluated the context awareness plane in the campus network of Masaryk University. Using the common network scanners, we were able to populate ISIM database with information on a network covering /16 IPv4 address range serving 15,000 active devices at peak times and 40,000 users. The periodically scheduled scan of the whole network took from 17 hours to 3 days (depends on the setting of scanning all ports or the most popular ones and rate limiting applied to avoid congestion) and produced over 3.5 GB of data in Neo4j database, including the details on approximately 31,000 unique hosts and 98,000 network services linked via 650 CPE fingerprints to 2,900 unique CVEs, together posing as 640,000 nodes and 12,000,000 edges. Even with such amounts of data, the graph database responded promptly to queries and did not cause significant delays.

B. AI-based incident detection

The evaluation of the AI-based incident detection is performed using a publicly dataset (BATADAL), consisting of different sensor readings over time. The dataset is split into three segments, each mapped to a specific edge area based on network topology, ensuring data alignment with the physical layout and PLC positions. Results from the evaluation demonstrated that while the edge-based models exhibited a minor 0.4% reduction in detection accuracy compared to a centralized model, they significantly improved response time, reducing the mean time-to-detection by 1.5% [11]. This trade-off is beneficial in OT environments, where rapid response to anomalies is critical for system integrity.

To analyze energy consumption and carbon emissions, we used the CodeCarbon framework for training and centralized inference, and Jetson Stat for edge inference. These tools tracked energy usage across RAM, CPU, and GPU, with emissions estimated using CodeCarbon's conversion rates. A benchmark model was first trained to measure energy consumption, followed by optimized centralized and edge-based Autoencoders. By aggregating results from training and

inference, we assessed the impact of edge computing in reducing energy use and emissions compared to centralized detection. Localized processing at the edge significantly reduced network bandwidth consumption and energy usage, with an estimated 75% reduction in emissions (Fig. 2).




Fig. 2. Energy consumption and produced emissions for training different edge and centralized models.

C. Mitigation

Component performance was tested standalone and deployed with the other Security Operations Plane components. The primary objective is to quantify end-to-end latency. Testing was performed using synthetic alerts, together with an existing corpus of mitigation playbooks and the ISIM database collected from Masaryk University. Results are shown in left-to-right order in Fig. 3.

The first evaluation analysed the average time (0.23s avg.) spent by the Mitigation Manager between receiving an alert and dispatching a mitigation. Testing showed that the Mitigation Manager spent 72% of the time (0.17s) executing the playbooks, which included expensive computation such as iterating through all active attack graphs and querying external components.

The second test then measures the complete end-to-end latency (3.2s avg.) between the suspicious activity and the playbook's effect. Measurements showed that the Mitigation Manager occupied around 8% of the total time (0.26s), with detection and mitigation execution taking 26% (0.82s) and 66% (2.12s) respectively.

VI. CONCLUSIONS

In this paper, we introduced ResilMesh's novel AI-driven cybersecurity framework designed to enhance CSA and improve threat detection and mitigation in CPS. By integrating real-time




Fig. 3. Time distribution for the example mitigation scenarios.

network monitoring, threat intelligence correlation, and AI-powered anomaly detection, ResilMesh provides a scalable and adaptive security solution for protecting critical infrastructures. The first PoC implementation demonstrated the feasibility of the proposed architecture, manage cyber situational awareness and detect and mitigate cyberattacks.

Moving forward, future work will focus on expanding the scalability of ResilMesh, integrating advanced machine learning techniques for attack prediction, and refining automated mitigation strategies to enhance proactive cyber defense. Further experimentation in real-world industrial environments will also be conducted to validate the robustness of the framework in heterogeneous and distributed infrastructures. By bridging the gap between traditional SIEM systems and AI-driven cybersecurity orchestration, ResilMesh paves the way for a more intelligent, responsive, and resilient security architecture for critical infrastructure protection.

ACKNOWLEDGMENT

This project has been funded by the European Union as part of the Horizon Europe Framework Program (HORIZON), under the grant agreement 101119681, Resilmesh.

REFERENCES

- [1] N. Tuptuk, P. Hazell, J. Watson, and S. Hailes, "A Systematic Review of the State of Cyber-Security in Water Systems," *Water*, vol. 13, no. 1, p. 81, 2021.
- [2] S. Ikeda, "Two More Cyber Attacks on US Water Supply Highlights Concerns About Vulnerabilities, but Sensational Headlines Sometimes Overstate the Threat to Public Safety," <https://www.cpomagazine.com/cyber-security/two-more-cyber-attacks-on-us-water-supply-highlights-concerns-about-vulnerabilities-but-sensational-headlines-sometimes-overstate-the-threat-to-public-safety/>, Jun. 2021.
- [3] D. Li, N. Gebraeel, K. Paynabar, and A. P. S. Meliopoulos, "An Online Approach to Cyberattack Detection and Localization in Smart Grid," *IEEE Transactions on Power Systems*, vol. 38, no. 1, pp. 267–277, 2023.
- [4] B. Stojanović, H. Neuschmied, M. Winter, and U. Kleb, "Enhanced anomaly detection for cyber-attack detection in smart water distribution systems," in *Proceedings of the 17th International Conference on Availability, Reliability and Security*, 2022, pp. 1–7.
- [5] R. Taormina and S. Galelli, "Deep-Learning Approach to the Detection and Localization of Cyber-Physical Attacks on Water Distribution Systems," *Journal of Water Resources Planning and Management*, vol. 144, no. 10, p. 04018065, 2018.
- [6] L. Tsiami and C. Makropoulos, "Cyber-Physical Attack Detection in Water Distribution Systems with Temporal Graph Convolutional Neural Networks," *Water*, vol. 13, no. 9, p. 1247, 2021.
- [7] M. Housh, N. Kadosh, and J. Haddad, "Detecting and Localizing Cyber-Physical Attacks in Water Distribution Systems without Records of Labeled Attacks," *Sensors (Basel, Switzerland)*, vol. 22, no. 16, p. 6035, 2022.
- [8] U. Gupta, Y. G. Kim, S. Lee, J. Tse, H.-H. S. Lee, G.-Y. Wei, D. Brooks, and C.-J. Wu, "Chasing Carbon: The Elusive Environmental Footprint of Computing," in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 2021, pp. 854–867.
- [9] D. Novoa-Paradela, O. Fontenla-Romero, and B. Guijarro-Berdiñas, "Fast deep autoencoder for federated learning," *Pattern Recognition*, vol. 143, p. 109805, 2023.
- [10] M. Husák, L. Sadlek, S. Špaček, M. Laštovička, M. Javorník, and J. Komárková, "CRUSOE: A toolset for cyber situational awareness and decision support in incident handling," *Computers & Security*, vol. 115, p. 102609, 2022.
- [11] M. Somma, A. Flatscher, and B. Stojanović, "Edge-based anomaly detection: Enhancing performance and sustainability of cyber-attack detection in smart water distribution systems," in *2024 32nd Telecommunications Forum (TELFOR)*. IEEE, 2024, pp. 1–4.