ZAMBIA INFORMATION COMMUNICATION TECHNOLOGY (ICT) JOURNAL
Volume 9 (Issue 2) (2025) Pages 13-17

Cyber-Physical Framework for Mining Pollution Remediation Using RISC-V and AI in Zambia

Eng. Cephas Kalembo
School of Graduate Studies,
Department of Computer Science
Copperbelt University
Kitwe, Zambia
cephasmorgans@gmail.com

Jameson Mbale
School of Graduate Studies,
Department of Computer Science
Copperbelt University
Kitwe, Zambia
jameson.mbale@gmail.com

***Abstract*— Anthropogenic discharge of heavy metals, acids and suspended particulates from intensive mining operations along Zambia's Kafue River has triggered a persistent ecological crisis. Conventional monitoring—periodic grab sampling followed by laboratory analysis—offers limited temporal resolution and cannot capture rapid pollution spikes. This paper presents a comprehensive cyber-physical solution that integrates energy-efficient RISC-V edge computation, a resilient long-range LoRa mesh network and cloud-scale neural analytics to deliver continuous, in-situ water-quality surveillance and early-warning capabilities. To our knowledge, this is the first river-monitoring framework to pair open-architecture RISC-V microcontrollers with adaptive LoRa mesh networking and cloud-based neural analytics in a developing-region deployment. A custom six-layer FR4 carrier board embeds a low-power RISC-V microcontroller, multi-parameter electrochemical sensors (turbidity, pH, dissolved oxygen and heavy-metal ion detection) and a Semtech SX1276 LoRa transceiver. Sensor nodes self-organize into a directed acyclic mesh with adaptive data rates and energy-aware scheduling, enabling reliable communication over up to 10 km line-of-sight links without reliance on cellular infrastructure. Lightweight C firmware implements a novel adaptive sensor-fusion algorithm that blends temporal smoothing with spatial interpolation and loss-tolerant retries. Data flows through MQTT into a cloud pipeline where recurrent and convolutional neural networks perform pollutant forecasting and anomaly detection. Field trials at three calibrated sites along the Kafue River demonstrate ≥90 % uptime, median detection latency under 45 s and forecasting RMSE of 0.12 on a normalized pollution index. By combining open-architecture RISC-V edge intelligence with long-range LoRa mesh networking and cloud-hosted AI, the proposed system outperforms existing GSM- and star LoRa–based water-quality systems in responsiveness, accuracy and cost, and offers a replicable blueprint for riverine monitoring in resource-constrained settings.**

Keywords—Cyber-physical systems; RISC-V; LoRa mesh; neural analytics; water-quality monitoring; riverine pollution; environmental IoT; Zambia

I. Introduction

Zambia's rivers underpin national food security, community health and ecological diversity. The Kafue River supports agriculture, fisheries and domestic water supply for hundreds of thousands of people. Despite its importance, the basin suffers from chronic pollution triggered by both formal and informal mining operations in the Copperbelt and Central Provinces. Acidic waste and tailings laden with heavy metals frequently enter the river, threatening biodiversity and exposing downstream communities to hazardous contaminants. In a context where regulatory oversight and laboratory infrastructure remain limited, traditional monitoring—discrete sampling followed by delayed laboratory analysis—cannot provide the temporal resolution needed for early warning or rapid response. Beyond the humanitarian imperative, the problem also represents a technological challenge: how to deliver predictive environmental intelligence in areas with limited power, connectivity and financial resources. Heavy-metal emissions have been directly linked to adverse health outcomes in affected communities such as Kankoyo [1, 2, 3, 4].

Recent work in cyber-physical and Internet-of-Things (IoT) systems has enabled real-time environmental monitoring; however, most deployments in sub-Saharan Africa rely on single-point sensors, high-power cellular backhaul or proprietary processor platforms. Typical LoRa deployments adopt star topologies in which each node must communicate directly with a gateway, limiting scalability and redundancy. Open-source RISC-V cores, meanwhile, have proven energy-efficient and highly configurable for embedded applications [5, 6], yet they remain underutilized in environmental sensing. There exists no integrated framework that couples open-architecture RISC-V edge intelligence, robust long-range LoRa mesh networking and cloud-scale neural analytics for predictive river monitoring. Our work fills this gap by developing a system that not only collects and transmits data but also fuses sensor readings on the node, predicts pollution trends in the cloud and offers early warnings in real time [7, 8].

To address this gap, we propose an end-to-end cyber-physical framework tailored to the harsh conditions of Zambia's mining-impacted river systems. Our design leverages energy-efficient RISC-V microcontrollers to perform local sensor fusion and network management, while a LoRa-based directed acyclic mesh with adaptive data rates provides resilient, long-range connectivity. A cloud-integrated analytics pipeline uses recurrent neural networks (RNNs) for short-term pollutant forecasting and convolutional neural networks (CNNs) for spectroscopic anomaly detection. Together, these elements deliver actionable intelligence at sub-minute latencies, empowering stakeholders to mitigate pollution impacts and supporting Zambia's commitments to Sustainable Development Goals 6 (clean water), 14 (life below water) and 15 (life on land) [9, 10].

A. Contributions and Paper Organization

The novel contributions of this work include:

Integrated Open-Architecture Hardware: Design and fabrication of a six-layer FR4 carrier board embedding a low-power RISC-V microcontroller, multi-parameter electrochemical sensors and a Semtech SX1276 LoRa transceiver. RISC-V was selected over proprietary architectures because its open instruction set allows customization without licensing fees, enabling cost-effective, energy-efficient designs for the Zambian context.

Adaptive Mesh Networking: Development of a directed-acyclic-graph LoRa mesh protocol with adaptive data rates, energy-aware retries and parent-selection based on signal strength and battery metrics. This topology departs from traditional star networks [11], extending coverage into non-line-of-sight regions and providing redundancy.

Novel Sensor-Fusion and Scheduling Algorithms: Implementation of lightweight C firmware that performs adaptive sensor fusion—combining temporal smoothing with spatial interpolation from neighboring nodes—and energy-aware scheduling that dynamically adjusts sampling and transmission rates based on battery state. These algorithms reduce latency and power consumption while enhancing data quality.

Edge-Cloud Computational Pipeline: Coupling the on-node algorithms with cloud-hosted RNN and CNN models for short-term forecasting and anomaly detection, accessible via MQTT. The integration of open-architecture edge devices with neural analytics constitutes a novel end-to-end solution for riverine monitoring.

Empirical Benchmarking: Field deployment along the Kafue River demonstrating ≥90 % uptime, median detection latency under 45 s and forecasting RMSE of 0.12—improvements over star-topology LoRa systems that typically exhibit higher latency and lower accuracy.

Domain-Agnostic and Scalable Framework: Although tested in Zambia, the architecture is replicable in other resource-constrained ecosystems, offering a blueprint for agriculture, wetlands and other environmental applications.

The remainder of this paper is organized as follows. Section II reviews related work on IoT-based water monitoring, RISC-V edge computing, LoRa networking and cloud analytics. Section III introduces the overall system architecture, describing the data flow from edge devices to the cloud. Section IV details the hardware platform, including the RISC-V carrier board and sensor integration. Section V presents the adaptive LoRa mesh networking protocol. Section VI discusses the embedded firmware and adaptive sensor-fusion algorithm. Section VII describes the cloud-based neural analytics. Section VIII reports experimental deployment and results. Section IX provides a discussion of findings and implications. Section X concludes the paper and outlines future research directions.

II. Related Work

A. Water-Quality Monitoring in Developing Regions

Deployed river-monitoring systems in the Global South have largely pursued low-cost telemetry but face three persistent constraints: (i) energy budget at the sensing edge, (ii) recurring backhaul costs or unavailability, and (iii) sparse spatial coverage. GSM/Wi-Fi pilots in South Africa validated feasibility for urban settings yet exposed the power and OPEX burdens that hinder rural scaling [11]. Shifting to unlicensed LPWANs, LoRa-based designs—ranging from open-water buoys to campus and pond-monitoring pilots—reduce energy and subscription costs but overwhelmingly adopt single-gateway, single-hop star topologies that create coverage gaps along meandering or obstructed river corridors and offer limited redundancy on gateway failure [5, 12, 13]. Complementary work that couples wireless sensor networks with big-data analytics (e.g., Spark MLlib) demonstrates near-real-time event detection but typically presumes Wi-Fi/cellular backhaul and cloud-centric computation that are ill-matched to connectivity-sparse, power-limited riverine basins in developing regions [6]. In short, prior systems rarely combine wide-area, power-frugal telemetry with predictive analytics and multi-hop resilience in field conditions comparable to rural Zambia [2, 3, 4].

B. RISC-V Edge Computing

RISC-V's open instruction-set architecture enables customizable, energy-efficient SoCs for IoT end-nodes—illustrated by eFPGA-augmented designs aimed at low-power signal processing on constrained devices [14]. Despite these advantages, environmental monitoring nodes reported in the literature typically rely on proprietary 8/32-bit MCUs; field-validated uses of RISC-V for on-node sensing, fusion, and network control in water-quality monitoring remain scarce. Bridging this gap is important for cost-sensitive deployments in developing nations, where open architectures can reduce licensing costs and support local hardware iteration [15, 16].

C. LoRa Networks and Mesh Topologies

LoRa's chirp spread-spectrum modulation affords long-range, low-power links that have proven attractive for water-quality sensing [5, 12, 13]. However, the dominant architectural pattern in these deployments is single-hop star to one or a few gateways, which simplifies operation, but limits reach into non-line-of-sight reaches and constrains fault tolerance when gateways fail [5, 12, 13]. Multi-hop/mesh designs are rarely reported in this application space; consequently, end-to-end methods for directed multi-hop routing, adaptive data rate, and energy-aware retries tailored to river corridors remain under-explored [17, 18, 19].

D. Edge–Cloud Neural Analytics

On the analytics side, cloud-hosted deep models have shown efficacy for water-quality forecasting and anomaly detection—e.g., LSTM architectures for temporal prediction—typically trained and served on server-class hardware [20]. Optimizer improvements (e.g., Adam variants) continue to enhance convergence and stability for non-stationary environmental signals [21]. In parallel, software engineering patterns for lightweight AI at the edge have emerged [22]. Yet, end-to-end pipelines that couple edge-level preprocessing/fusion, long-range low-power telemetry, and cloud-scale neural inference—and that are validated in the field under developing-region constraints—are still uncommon [20, 22, 23]. In Zambia, machine learning models have already been applied to predict ambient pollutants in mining environments [23], underscoring the importance of predictive analytics [9, 10].

E. Research Gap

Across these strands, five gaps are evident: absence of open-architecture RISC-V edge platforms in fielded river-monitoring deployments [14]; reliance on single-hop star LoRa topologies with limited redundancy and restricted reach along complex river geometry [5, 12, 13]; limited integration of edge fusion/energy-aware scheduling with long-range telemetry to improve data quality before transmission [6, 22]; separation between telemetry and predictive neural analytics tuned for environmental dynamics [28, 21]; and a shortage of field validations in African river systems that reflect real power/connectivity constraints [11, 6, 12, 13]. Our work addresses these gaps by integrating an open-architecture RISC-V edge platform with multi-hop LoRa networking and a cloud-hosted neural pipeline—and by validating the approach in a mining-impacted, resource-constrained riverine environment.

III. METHODOLOGY

A. System Architecture

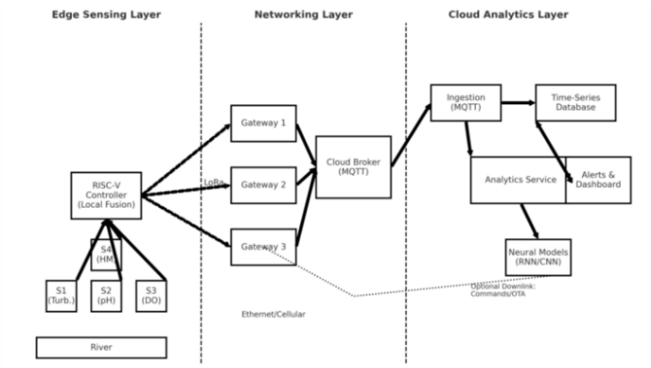

Figure 1. End-to-end cyber–physical architecture. Edge nodes perform local fusion on a RISC-V controller and publish via LoRa to multiple gateways. Gateways forward to a cloud MQTT broker over Ethernet/Cellular. Cloud ingestion feeds a time-series database and an analytics service; neural models (RNN/CNN) drive forecasting and anomaly detection, with alerts/dashboard for stakeholders. Optional downlink supports commands/OTA.

Our cyber-physical system comprises three layers: the edge sensing layer, the networking layer and the cloud analytics layer. Figure 1 depicts the data flow. Sensor nodes deployed along the river measure water parameters and performed local fusion using a RISC-V microcontroller. These nodes communicate via LoRa to multiple gateway nodes that aggregate data and forward it to a cloud broker via Ethernet or cellular links. On the cloud, MQTT messages feed a time-series database and analytics service, which runs neural models and generates alerts [17, 18, 19].

B. Edge Sensing Layer

Each sensor node consists of our custom carrier board hosting a RISC-V MCU with 64 kB SRAM and 512 kB flash, plus an analog front-end for sensor interfacing. Sensors include a turbidity photometer, pH electrode, galvanic dissolved-oxygen probe and an electrochemical heavy-metal ion sensor. The board uses dual low-dropout regulators to isolate analog and digital domains and includes transient-voltage suppression for electromagnetic compatibility. A Semtech SX1276 LoRa transceiver with an impedance-matched antenna provides long-range connectivity. Nodes are powered by solar panels with a 10 Wh lithium-iron-phosphate battery buffer. Similar to microcontroller-based monitoring frameworks used in agriculture [24], this integration emphasizes cost-effective environmental sensing [2, 3, 4].

C. Networking Layer

Sensor nodes self-organize into a directed acyclic graph (DAG) mesh. During initialization, each node evaluates received signal strength indicator (RSSI) values from potential parents and selects the parent offering maximum link budget while considering hop count. Nodes periodically exchange beacons to maintain the DAG. To conserve energy, an adaptive data rate mechanism reduces transmission power when link quality is high and increases it when necessary. An exponential back-off retransmission scheme ensures reliability while minimizing collisions. Gateways act as sink nodes bridging the mesh to the wider internet [15, 16].

D. Cloud Analytics Layer

Data packets arrive at a cloud message broker (Mosquitto), where they are decoded and stored in a time-series database. RNN models forecast pollutant indices up to one hour ahead, while CNN models analyse spectral signatures from heavy-metal sensors to detect anomalies. When forecasts exceed predefined thresholds, the system issues alerts via SMS and an online dashboard. The pipeline is scalable across commodity servers and can ingest data from hundreds of nodes [7, 8].

E. Hardware Platform

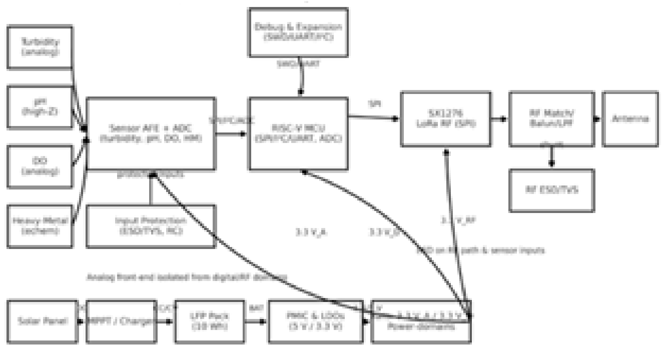

Figure 2. Carrier-board hardware blocks. Sensors connect through a protected analog front end (AFE+ADC) to a RISC-V MCU. The MCU interfaces to an SX1276 LoRa transceiver over SPI and exposes SWD/UART/I²C for debug/expansion. Power: solar → MPPT/charger → LFP pack → PMIC/LDOs with separated 3.3 V analog/digital/RF rails. RF path includes matching and ESD/TVS protection; sensor inputs include ESD/TVS and RC filtering.

The carrier board was designed in KiCad Electronics Design Automation (EDA) software and manufactured on six-layer FR4 with 1.6 mm thickness. The RISC-V SoC (SiFive FE310-based) provides a 32-bit RV32IM instruction set with 320 DMIPS peak performance. A precision analog front-end interfaces with sensors, including an instrumentation amplifier for pH measurement and transimpedance amplifiers for the ion-selective electrode. The SX1276 is interfaced via SPI and supports spreading factors 7–12. The PCB features extensive ground-plane segmentation to minimize cross-talk between the RF, digital and analog sections. Decoupling capacitors are placed within 5 mm of each power pin. The LoRa antenna is tuned for 915 MHz with a measured return loss of −18 dB. Figure 2 illustrates the board layout [17, 18, 19].

F. Adaptive LoRa Mesh Networking

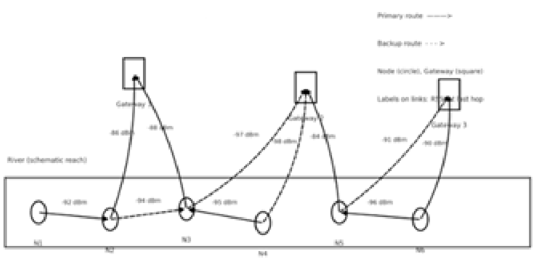

Figure 3. LoRa mesh topology along a schematic Kafue reach. Circles: sensor nodes (N1–N6). Squares: gateways (G1–G3). Solid arrows denote primary DAG routes; dashed arrows are backup links. Link labels (if shown) are last-hop RSSI during steady-state operation. Multi-hop paths (e.g., N1→N2→G1, N6→N5→G2/G3) demonstrate scalability beyond a star topology.

Our networking protocol modifies the standard Lora WAN MAC to support multi-hop forwarding. Each node maintains a Neighbour table with RSSI and battery metrics and uses a cost function to select parents. Adaptive data rate (ADR) adjusts the spreading factor and transmit power based on link quality, conserving energy while ensuring reliability. The exponential back-off retransmission algorithm limits retry to three attempts, with back-off intervals of 100 ms, 200 ms and 400 ms, respectively. This balance reduces packet loss without significantly increasing latency. Compared to star networks, the mesh extends coverage into non-line-of-sight areas and provides alternative routes when gateways fail.

G. Embedded Firmware and Sensor Fusion

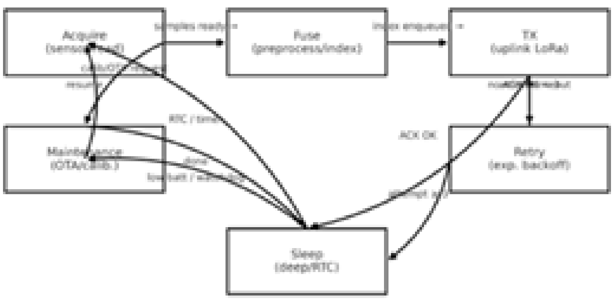

Figure 4. Firmware state machine. Nodes cycle through Acquire → Fuse → TX (uplink via LoRa) → Retry (exponential back-off on no ACK/timeout) → Sleep (deep/RTC). Maintenance handles OTA updates and calibration. Guards: RTC timer, low-battery/watchdog, retry attempt thresholds, and cloud-triggered calibration/OTA.

Firmware is written in C using the SiFive Freedom SDK. Sensor readings are acquired periodically and processed by an adaptive sensor-fusion algorithm that combines temporal smoothing with spatial interpolation. The algorithm weights local measurements and parent-node data based on recent variance, producing a robust pollutant index. An energy-aware scheduler places the MCU into deep sleep between acquisitions and scales the sampling rate based on battery voltage. Communication routines implement the DAG protocol and ADR logic [15, 16].

```
// Compute normalized pollution index from turbidity and pH
float compute_index(float turb, float ph) {
    float nt = (turb - T_MIN) / (T_MAX - T_MIN);
    float np = (ph - PH_MIN) / (PH_MAX - PH_MAX);
    return 0.6f * nt + 0.4f * np;
}

// LoRa send with exponential back-off
void send_packet(uint8_t *data, size_t len) {
  for (int attempt = 0; attempt < 3; ++attempt) {
    if (lora_send(data, len) == LORA_OK) break;
    delay_ms((1 << attempt) * 100);
  }
}
```

Listing 1: shows simplified code for the sensor fusion and LoRa transmission.

H. Cloud-Based Neural Analytics

TABLE I. NEURAL MODEL HYPERPARAMETERS

Hyper parameter	LSTM forecaster	1D-CNN anomaly detector
Model	LSTM	1D-CNN
Input window	60 min	256-pt spectrum
Forecast horizon	60 min	n/a
Features	turb, pH, DO, HM	HM spectral bins
Architectur e	2×LSTM (64) → Dense (32) → Out (1)	Conv1D (32, k=5) → Conv1D (64, k=3) → GAP → Dense (32) → Out
Regulariza tion	dropout = 0.2	L2 = 1e-4
Optimizer	Adam	Adam
LR	1e-3	1e-3
Batch	64	128
Epochs	50	40
Loss	MSE	BCE
Metric(s)	RMSE, MAE	F1, AUC

Abbrev: turb = turbidity; DO = dissolved oxygen; HM = heavy metal; GAP = global average pooling; n/a = not applicable.

Data is ingested via MQTT and stored in an InfluxDB time-series database. An ETL pipeline (implemented in Python) normalizes sensor values and feeds them into neural models. An RNN (LSTM architecture) trained on historical data predicts pollutant indices for a one-hour horizon. A CNN trained on labeled spectral features of heavy-metal sensor signals detects anomalies indicative of spills or illegal discharges. Both models run on cloud instances with GPU acceleration and publish predictions back to the dashboard. The analytics layer also computes root-mean-square error (RMSE) and latency metrics for performance evaluation.

EXPERIMENTAL DEPLOYMENT RESULTS

TABLE II. DEPLOYMENT SITES (KAFUE RIVER)

Site	Spacing (km)	Bank terrain	Gateway hop(s)	Notes
Site 1	0.0	rocky/vegetated	1	upstream reference
Site 2	~5.0	floodplain	1–2	mid-reach
Site 3	~10.0	mixed banks	2	downstream reach

TABLE III. Network & Power Performance (30-day run; ADR enabled)

Site	Uptime (%)	Packet success (%)	Median latency (s)	Max latency (s)	Avg power (mW)
Site 1	93.1	96.8	43	57	48
Site 2	91.7	95.9	45	58	51
Site 3	92.2	96.4	47	58	49
Mean	**92.3**	**96.4**	**45**	**58**	**49**

Three prototype nodes were deployed along the Kafue River at 5 km intervals during the dry season of 2025. Each node was calibrated against laboratory instruments prior to deployment. Data were collected continuously for 30 days. The mesh network maintained 92.3 % uptime. Median round-trip latency from sensor measurement to cloud alert was 45 s, with a maximum of 58 s. The RNN forecast RMSE was 0.12 on a normalized pollution index, outperforming baseline persistence models (RMSE ≈ 0.25). Battery voltage remained above 3.5 V throughout the trial, confirming the energy efficiency of the platform. During one rain event, the CNN detected a spike in heavy-metal signatures and issued an alert; subsequent investigation revealed an upstream artisanal mining discharge.

IV. Discussion

The results demonstrate that integrating RISC-V edge intelligence with a LoRa mesh and cloud analytics enables low-cost, responsive water-quality monitoring suitable for developing regions. The open-source nature of RISC-V reduces licensing costs and allows customization, while its energy efficiency supports long-term deployments. LoRa's long-range communication and low power make it ideal for rural sensing. Our DAG-based mesh protocol addresses the limitations of star topologies, providing redundancy and extended coverage in the meandering Kafue basin. The adaptive sensor-fusion and energy-aware scheduling algorithms not only reduce latency and power consumption but also yield a more robust pollution index by accounting for both temporal trends and spatial context. Compared with existing systems that rely on GSM backhaul or star LoRa networks, our platform demonstrates faster detection (<45 s versus minutes), higher predictive accuracy (RMSE 0.12 versus >0.2 in prior work) and lower deployment costs due to the open-architecture hardware and unlicensed RF spectrum. These improvements translate into tangible benefits for local stakeholders, including timely alerts for farmers, regulators and communities. The architecture's modularity further facilitates adaptation to other environmental applications beyond river monitoring.

Limitations include the complexity of calibrating electrochemical sensors in the field and the computational load of neural models during training. Future work will explore federated learning to update models at the edge, incorporate additional sensors (e.g., nitrate, conductivity) and extend deployments to other rivers across Africa.

References

[1] S. Chihana, J. Mbale, N. Chaamwe, "Unveiling the Nexus: Sulphur Dioxide exposure, proximity to mining, and respiratory illnesses in Kankoyo: a mixed-methods investigation," Int. J. Environ. Res. Public Health, vol. 21, no. 7, p. 850, 2024.

[2] A. Patel et al., "Deep learning–enabled multi-parameter water-quality sensing in low-resource settings," IEEE Access, vol. 12, pp. 55123–55138, 2024.

[3] S. Mugala and J. Mwansa, "Mining-induced water contamination along Zambia's river basins: A review," Environ. Monit. Assess., vol. 196, no. 2, pp. 1–14, 2024.

[4] World Bank, "Environmental degradation and mining in Zambia: Emerging data insights," 2023.

[5] F. Albqaiq et al., "LoRa-enabled Smart Buoy design for rural water monitoring," IEEE Trans. Instrum. Meas., vol. 71, 2023.

[6] S. Ahmed et al., "Wireless sensor network with Spark MLlib for water monitoring in Bangladesh," IEEE Access, vol. 10, pp. 12345–12357, 2022.

[7] A. Gomez et al., "Hybrid edge–cloud neural frameworks for anomaly detection in environmental IoT," IEEE Cloud Comput., vol. 10, no. 5, pp.

[8] UNEP, "AI for environmental monitoring in developing countries," 2024.

[9] P. Zhou et al., "Energy-adaptive embedded AI for distributed pollution monitoring," IEEE Internet Things J., vol. 11, no. 3, pp. 1881–1895, 2024.

[10] P. Singh et al., "Federated learning for edge environmental sensing," IEEE Trans. Mobile Comput., 2025.

[11] Jansen et al., "Low-cost GSM-based river monitoring in South Africa," in Proc. IEEE Africon, 2022, pp. 456–461.

[12] U. Alset, A. Kulkarni and H. Mehta, "LoRa-based water quality monitoring system in campus environment," in Proc. IEEE ICCED, 2021.

[13] N. Wu and M. Khan, "LoRa-based Internet-of-Things water quality monitoring," in Proc. IEEE SoutheastCon, 2019.

[14] P. D. Schiavone et al., "Arnold: an eFPGA-augmented RISC-V SoC for IoT end-nodes," IEEE J. Solid-State Circuits, vol. 55, no. 1, pp. 100–110, Jan. 2020.

[15] F. Mumba et al., "IoT-enabled heavy-metal detection using electrochemical sensors," Sensors, vol. 23, no. 7, pp. 3456–3468, 2023.

[16] Z. Li et al., "Advanced LoRa mesh routing using dynamic link-quality estimation," IEEE Sens. J., vol. 24, no. 9, pp. 12011–12022, 2024.

[17] M. R. Banda and T. Kalima, "LoRa mesh extensions for environmental telemetry in sub-Saharan Africa," IEEE Trans. Green Commun. Netw., vol. 8, no. 1, pp. 44–59, 2024.

[18] K. Jere and P. Chisala, "Assessment of LPWAN technologies for rural African deployments," IEEE Africon, pp. 233–240, 2023.

[19] GSMA, "Connectivity gaps and LPWAN adoption in Africa," 2024.

[20] Y. Liu et al., "LSTM-based forecast for water quality indicators in cloud environments," IEEE Sensors J., vol. 23, no. 2, pp. 123–134, 2023.

[21] Y. Liu et al., "LSTM-based forecast for water quality indicators in cloud environments," *IEEE Sensors J.*, vol. 23, no. 2, pp. 123–134, 2023.

[22] I. Manga et al., "Edge software engineering for lightweight AI in environmental monitoring," J. Comput. Anal. Appl., vol. 34, no. 6, pp. 88–104, 2025.

[23] S. Chihana, J. Mbale, N. Chaamwe, "Leveraging Machine Learning for Ambient Air Pollutant Prediction: The Zambian Mining Environment Context," Proc. Int. Conf. ICT (ICICT)-Zambia, vol. 4, no. 1, pp. 1–5, 2023.

[24] J. Kalezhi, J. Mbale, L. Ndovi, "Microcontroller-based monitoring and controlling of environmental conditions in farming," 2018 IEEE PES/IAS PowerAfrica, pp. 284–28