

Implementation of an intelligent edge-computing fire security system: Integrating multi-sensor fusion, dual-source power redundancy and real-time IoT cloud telemetry

Mustapha Suleiman Olawale ^{1,*}, Jethro Gbutemu Yoti ², Nwafor Charles Somtochukwu ³, Victor John Nanmwa ⁴, Guiryamadji Arnaud ² and Mahmud Aliyu Miko ⁵

¹ Department of Electrical and Electronics Engineering, Faculty of Engineering and Engineering Technology, Abubakar Tafawa Balewa University, P.M.B. 0248, Bauchi, Nigeria.

² Department of Technical Education, School of Vocational and Technical Education, College of Education Zing, P.M.B. 1021, Taraba, Nigeria.

³ Department of Mechatronics Engineering Technology, School of Electrical Systems and Engineering Technology, Federal University of Technology Owerri, P.M.B. 1526, Imo, Nigeria.

⁴ Department of Computer Engineering, School of Electrical Engineering and Technology, Federal University of Technology Minna, P.M.B. 65, Niger, Nigeria.

⁵ Department of Electrical Engineering, Faculty of Engineering, Bayero University Kano, P.M.B. 3011, Kano.

World Journal of Advanced Engineering Technology and Sciences, 2026, 20(01), 113–126

Publication history: Received on 31 May 2026; revised on 05 July 2026; accepted on 09 July 2026

Article DOI: <https://doi.org/10.30574/wjaets.2026.20.1.0356>

Abstract

Cloud-dependent Internet of Things (IoT) fire detection prototypes have demonstrated that wireless sensor networks and mobile notification pipelines can shorten the gap between ignition and human awareness. However, single-microcontroller designs that rely on rigid threshold logic, a single mains or battery power rail, and one Wi-Fi communication path remain vulnerable to the very conditions a fire produces: power interruption, network congestion, and ambiguous sensor readings that trigger nuisance alarms. This study presents the design, implementation, and bench evaluation of an intelligent edge-computing fire security system that addresses these three vulnerabilities concurrently. A multi-sensor fusion front end, comprising digital temperature/humidity, catalytic gas/smoke, and infrared flame sensing, feeds a lightweight on-device classification routine running on an ESP32-S3 microcontroller, so that fire-likelihood decisions are computed locally rather than waiting on a round trip to the cloud. A dual-source power subsystem pairs a mains-fed, buck-regulated rail with a solar-charged lithium-ion battery and an automatic power-path switchover circuit, allowing the node to continue sensing, fusing, and alerting through grid outages. A communication layer publishes fused readings and alerts over MQTT to a cloud broker and mirrors critical events to a Firebase-backed mobile application, while a GSM module provides a Short Message Service (SMS) fallback when the primary wireless path is degraded or unavailable. Bench testing under simulated normal, hazardous, power-loss, and connectivity-loss scenarios indicates that on-device fusion reduces nuisance alarms relative to single-threshold logic, that the power subsystem switches sources within a sub-second window, and that the system sustains alert delivery through both communication paths. The system is presented as a low-cost, field-deployable extension of earlier cloud-only IoT fire alarm prototypes, contributing to United Nations Sustainable Development Goals 9 and 11 through resilient smart-building safety infrastructure.

Keywords: Edge Computing; Multi-Sensor Fusion; Fire Detection; Power Redundancy; IoT Cloud Telemetry

* Corresponding author: Mustapha Suleiman Olawale

1. Introduction

Fire remains one of the most disruptive hazards to life, property, and critical infrastructure, and the speed with which an incipient event is detected and communicated is often the single largest determinant of outcome. Conventional smoke and heat alarms react to a single physical cue and are well known to misfire on steam, cooking aerosols, and dust, a limitation that has motivated a steady stream of work on multi-sensor and multi-wavelength approaches capable of distinguishing genuine combustion signatures from nuisance stimuli [2].

A recent prototype developed at the same research group demonstrated that a single ESP32 microcontroller, a temperature/humidity sensor, a gas sensor, a buzzer, and a Firebase Realtime Database backend could deliver sub-three-second alert latency to a companion mobile application, a meaningful improvement over standalone alarm [1]. That study, however, also surfaced three structural limitations that are common to a great many low-cost IoT fire-alarm prototypes reported in the literature. First, the hazard decision was a rigid Boolean combination of three threshold conditions evaluated on raw sensor values, with no mechanism for weighting sensor reliability or smoothing transient spikes. Second, the node depended on a single USB-fed power rail, so a mains interruption, a scenario not infrequent during an electrical fire, would silence the very device meant to report it. Third, alerting depended entirely on a Wi-Fi connection to a single cloud backend, leaving no fallback path if the access point lost power or the network became congested. The authors of that study explicitly flagged additional sensing modalities, power resilience, and a cellular fallback channel as directions for future work [1].

This study takes up that future-work agenda directly. Multi-sensor data fusion has been shown in several recent industrial and building-safety studies to materially reduce false and missed fire predictions relative to single-sensor or single-threshold logic, particularly when sensor reliability is modeled explicitly rather than assumed [3], and when complementary sensing modalities such as infrared flame detection are combined with gas and thermal cues [2]. Wearable and embedded multi-sensor fusion frameworks reported elsewhere further demonstrate that combining heterogeneous low-cost sensors with a lightweight classification layer, rather than a single rule, materially improves event-recognition reliability in safety-critical settings [4].

A second relevant trend is the migration of inference from the cloud to the edge. Embedded machine learning frameworks now allow quantized classification models to run directly on microcontroller-class hardware with millisecond-scale inference latency and milliwatt-scale energy budgets, removing the dependency on a live network connection for every hazard decision [8]. Building on this idea, dynamic, adaptive model-selection approaches for IoT-enabled fire prediction have been shown to outperform static threshold models across varied fire scenarios [7], and two-stage verification architectures, in which a fast sensor trigger is followed by a short confirmation step before an alarm is escalated, have demonstrated detection accuracy above 95 percent together with false-alarm reductions exceeding 90 percent relative to single-sensor baselines [5]. Closely related ESP32-class systems that combine gas, flame, and temperature sensing with an on-device fusion routine and visual confirmation report similar gains in reliability for residential and commercial deployments [6].

Power and communication resilience have received comparatively less attention in low-cost academic prototypes, despite being well studied in the broader IoT energy-harvesting literature. Mixed-mode system-on-chip designs for solar energy harvesting demonstrate that a regulated, battery-buffered rail can be maintained from intermittent solar input without manual intervention [9], and hybrid indoor light-and-heat harvesting nodes show that even modest ambient energy sources can sustain continuous sensor operation [10]. At the system level, integrated safety platforms that pair a solar-and-grid dual power backup with a cellular fallback channel for SMS and voice alerts have been shown to remain operative across both urban and rural deployment conditions where Wi-Fi coverage cannot be assumed [11]. On the communication side, MQTT-based telemetry pipelines continue to be favored for their low overhead and near-real-time delivery to mobile dashboards [12], while narrowband cellular architectures illustrate how a secondary wide-area communication path can preserve alerting capability when local wireless infrastructure is degraded or absent [13]. Cloud-only architectures that lean on heavier artificial intelligence models for fire-risk prediction, by contrast, have been noted to introduce computational overhead that is difficult to justify for small and medium-sized buildings, motivating lighter-weight, locally computed alternatives for that class of deployment [14].

Taken together, this body of work points to a coherent gap: comparatively few reported low-cost IoT fire-alarm prototypes integrate on-device multi-sensor fusion, dual-source power redundancy, and dual-path cloud-and-cellular telemetry within a single, buildable system. The objective of this study is to design, implement, and bench-evaluate such a system as a direct, integrated extension of the cloud-only prototype reported previously [1], and to compare its behavior against that baseline and against the wider comparative landscape summarized above. The remainder of this paper is organized as follows. Section 2 describes the research approach, hardware and power architecture,

communication design, edge fusion algorithm, and testing methodology. Section 3 presents bench results and a comparative analysis against prior studies. Section 4 concludes with a summary of findings and directions for future work.

2. Methodology

2.1. Research approach

This study follows an applied, design-and-implementation research approach, structured in five phases: requirements definition, system architecture design, hardware and power subsystem integration, firmware and edge-model development, and bench testing and validation. The requirements phase was informed directly by the limitations identified in the predecessor cloud-only prototype [1], namely the absence of sensor fusion, the single power rail, and the single communication path. Hardware assembly and firmware development proceeded concurrently, with each subsystem (sensing, edge compute, power, communication) unit-tested in isolation before integration.

2.2. System architecture overview

The proposed system is organized into five logical layers, illustrated conceptually in Figure 1: a sensing layer that performs multi-sensor data acquisition; an edge compute and fusion layer that normalizes and combines sensor readings into a single hazard decision on-device; a power layer that maintains a continuously available regulated supply from two independent sources; a communication layer that publishes telemetry and alerts through a primary cloud path and a secondary cellular path; and an application layer comprising a mobile dashboard for remote monitoring. Unlike the predecessor design, in which the microcontroller acted primarily as a pass-through between raw sensor values and the cloud [1], the edge compute layer in this design is responsible for the hazard decision itself, with the cloud and mobile layers acting as notification and historical-logging surfaces rather than decision points.




Figure 1 System architecture of the proposed edge-computing fire security system




Figure 2 System circuit diagram of the proposed edge-computing fire security system

2.3. Hardware design: sensing and edge compute layer

The sensing layer retains the calibrated digital temperature and humidity sensor (DHT22) and the catalytic gas/smoke sensor (MQ-2) used in the predecessor prototype, and adds a digital infrared flame sensor that responds to the characteristic infrared flicker of an open flame within the 760 to 1100 nanometre band. Infrared flame sensing is comparatively immune to the steam- and dust-related false triggers that affect optical smoke obscuration alone, and its inclusion alongside gas and thermal sensing follows the broader trend toward multi-wavelength, multi-modal fire sensing reported in recent sensor-technology reviews [2]. The edge compute layer is built around an ESP32-S3 microcontroller rather than the plain ESP32 used previously; the additional RAM and the availability of vector instructions on the S3 variant make it practical to host a small quantized classification model alongside the rule-based fusion logic described in Section 2.6, consistent with recent demonstrations that quantized neural models can run with millisecond-scale latency and a flash footprint under one megabyte on comparable microcontroller-class hardware [8]. Local annunciation is retained in the form of a buzzer and status LED, supplemented by a small status display (for example, a 128x64 OLED panel) that reports the active power source, connectivity state, and current fusion score for on-site diagnostics.




Figure 3 Hardware integration block diagram

2.4. Dual-source power redundancy subsystem

The power subsystem is the principal departure from the predecessor design, which was powered solely from a single 5-volt USB-style supply [1]. In this design, a primary rail is derived from a mains-fed AC adapter stepped down through a buck regulator (a configuration similar to the LM2596-based regulation used in the earlier prototype). A secondary rail is derived from a small solar panel (illustratively, 6 V, 3 W) feeding a charge-controller circuit that trickle-charges a 3.7 V lithium-ion cell, with a boost converter raising the battery rail to the same regulated 5 V bus. A power-path management circuit, implemented with an ideal-diode-OR arrangement and a comparator-driven switchover, continuously monitors the mains rail and transfers the load to the battery rail within a sub-second window if mains power is lost, with the battery additionally bridging the brief interval during which the solar-fed charge controller re-establishes a stable output after an outage. This arrangement follows the design logic of mixed-mode energy-harvesting systems-on-chip that maintain a regulated output from intermittent harvested input by buffering through a rechargeable cell [9], and of hybrid indoor energy-harvesting nodes that demonstrate continuous sensor operation from modest ambient sources [10]. It is also consistent with safety-platform designs that explicitly pair solar and grid power as a dual backup to keep alerting hardware operative regardless of mains availability [11]. The edge node reports the identity of the active power rail and an estimate of battery state of charge alongside its sensor telemetry, so that the active source is visible on the mobile dashboard rather than being inferred indirectly.




Figure 4 Dual-source power redundancy subsystem: mains and solar-battery rails with automatic power-path switchover

2.5. Communication architecture and cloud telemetry

The communication layer publishes fused sensor readings, the on-device hazard decision, and power-source status to a cloud MQTT broker over Wi-Fi at quality-of-service level 1, using a topic structure organized by node identifier and parameter (for example, node1/temperature, node1/fusion_score, node1/alert_status), consistent with lightweight MQTT-based environmental telemetry pipelines reported elsewhere [12]. Critical alert events are mirrored to a Firebase Realtime Database instance so that the mobile application retains the same near-real-time synchronization behavior as the predecessor system [1], preserving continuity with the existing Flutter-based mobile front end. A secondary communication path is provided by a GSM module (for example, a SIM800L-class modem) that is normally idle. The edge node maintains a heartbeat against the MQTT broker; if the heartbeat is not acknowledged within a configurable timeout, the node switches to a degraded-connectivity mode in which critical hazard alerts are sent as SMS messages to a pre-configured list of numbers, in the manner of cellular fallback channels used in integrated safety-alarm platforms that must remain operative where Wi-Fi coverage is unreliable [11]. This dual-path design mirrors, at a smaller scale, the rationale behind wide-area narrowband cellular monitoring architectures, which use a secondary cellular path specifically to preserve alerting capability when local wireless infrastructure cannot be assumed available [13]. The system deliberately avoids the heavier cloud-side artificial-intelligence processing used in some comparable building-fire monitoring platforms [14], keeping the hazard decision on the edge node itself so that alerting does not depend on cloud model availability.




Figure 5 Communication sequence: MQTT/Firebase primary path with GSM/SMS fallback on connectivity loss




Figure 6 System flowchart of the proposed edge-computing fire security system

2.6. Edge sensor-fusion and classification algorithm

Rather than the rigid Boolean threshold logic used in the predecessor prototype [1], the edge compute layer in this design combines the three sensing modalities into a single, weighted fusion score, with the underlying intent informed by belief-weighted multi-sensor fusion approaches that explicitly account for differing sensor reliability rather than treating every input as equally trustworthy [3]. Each raw sensor reading is first normalized to a confidence value between zero and one relative to calibrated safe and hazardous bounds, and the three normalized values are combined into a fusion score using fixed weights that reflect the relative discriminative strength of each modality, with the flame sensor weighted most heavily given its comparative resistance to non-fire nuisance stimuli [2]. To avoid triggering on brief transient spikes, for example a momentary gust of cooking smoke, the system requires the fusion score to remain above the alert threshold across three consecutive sampling cycles before escalating from a pending state to a confirmed Hazard Detected state, a software-only analogue of the two-stage hardware-plus-vision verification reported in dual-stage fire-crisis management systems [5], adopted here as firmware logic given the cost and power constraints of a battery-backed embedded node rather than a continuously powered single-board computer with a camera. A small auxiliary classifier, exported from a model trained offline on logged sensor traces and quantized for on-device inference following established embedded-AI deployment practice [8], is run alongside the rule-based fusion score as a secondary corroborating signal rather than as the sole decision mechanism, preserving interpretability while still benefiting from learned, nonlinear sensor interactions identified in multi-sensor fusion studies of firefighting and industrial fire-risk recognition [4].

2.7. Software design and development tools

Firmware for the ESP32-S3 was developed using the Arduino core (with the ESP-IDF toolchain available for lower-level driver work), retaining the DHT and MQ-2 sensor libraries used previously and adding a digital interrupt-driven driver for the infrared flame sensor. The on-device classifier was exported as a TensorFlow Lite Micro model (alternatively produced via the Edge Impulse pipeline) and invoked from the main sensing loop alongside the rule-based fusion routine described in Section 2.6. MQTT connectivity was implemented using a standard ESP32 MQTT client library, the FirebaseESP32 library was retained from the predecessor project to preserve mobile-application compatibility, and a TinyGSM-based driver managed the SIM800L fallback module. The companion mobile application continues to use Flutter and Dart, extending the predecessor application's clean-architecture layering (domain, presentation, and data layers) with an additional view for power-source status and fusion-score history. Mermaid.js was used to prepare architecture and sequence diagrams for this report, and Git/GitHub was used for firmware and application version control, consistent with the development tooling used in the predecessor study [1].

2.8. Boundary conditions and fusion formulation

Let T , H , and G denote the raw temperature, humidity, and gas readings from the DHT22 and MQ-2 sensors respectively, and let Fl denote the binary digital output of the infrared flame sensor. Each continuous reading is normalized to a confidence value bounded between zero and one as shown in Equation 1 (The normalized temperature confidence value), Equation 2 (The normalized gas confidence value), Equation 3 (The normalized flame confidence value), where the subscript *safe* denotes the upper bound of normal operating conditions and the subscript *max* denotes the sensor's effective saturation point used for scaling. The clamp operator constrains the result to the interval $[0,1]$.

$$T' = \text{clamp}\left(\frac{T - T_{\text{safe}}}{T_{\text{max}} - T_{\text{safe}}}, 0, 1\right) \quad (1)$$

$$G' = \text{clamp}\left(\frac{G - G_{\text{safe}}}{G_{\text{max}} - G_{\text{safe}}}, 0, 1\right) \quad (2)$$

$$F' = Fl \quad (3)$$

The three normalized confidences are then combined into a single fusion score S using fixed weights w_T , w_G , and w_F that sum to one, as shown in Equation 4 and Equation 5. Illustrative weight values used during bench testing were $w_T = 0.25$, $w_G = 0.35$, and $w_F = 0.40$, reflecting the comparatively higher discriminative reliability attributed to infrared flame sensing in the sensor-technology literature [2].

$$S = w_T T' + w_G G' + w_F F' \quad (4)$$

Subject to

$$w_T + w_G + w_F = 1 \quad (5)$$

A hazard is confirmed only when the fusion score exceeds a threshold θ across three consecutive sampling cycles spaced Δ seconds apart, as shown in Equation 6, suppressing single-sample transients before an alert is escalated to the communication layer.

$$Hazard_{confirmed} = \begin{cases} 1, & \text{if } S(t) > \theta, S(t - \Delta) > \theta, S(t - 2\Delta) > \theta \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

Table 1 summarizes the illustrative bounds used to derive the normalized confidences and the threshold values used during bench testing; these values should be re-calibrated against the specific sensors, enclosure, and deployment environment used in any field installation.

Table 1 Illustrative boundary conditions for the fusion variables

Variable	Safe bound	Saturation bound	Source/sensor
Temperature (T)	$\leq 30^\circ\text{C}$	60°C	DHT22
Humidity (H)	$\geq 60\%$	20%	DHT22 (inverse cue)
Gas concentration (G)	$\leq 70\text{ ppm}$	400 ppm	MQ-2
Flame (Fl)	0 (no flicker)	1 (flicker detected)	IR flame sensor
Fusion threshold (θ)	0.60	—	Equation 2

2.9. Testing and validation

Each sensing element was first checked against a reference thermometer/hygrometer and a regulated gas source, following the same calibration approach used in the predecessor study [1], with the infrared flame sensor additionally checked for response time and angular sensitivity against a small open flame at varying distances. The power subsystem was evaluated by repeatedly interrupting the mains supply while the node was under typical sensing-and-Wi-Fi load and logging the rail voltage continuously to measure switchover time and to confirm that no sensor readings or MQTT publications were lost during the transition; battery autonomy under mains loss and solar recharge time from a depleted state were measured separately on the laboratory bench under controlled illumination. Edge inference latency was benchmarked by timing the fusion routine and the auxiliary classifier invocation directly on the ESP32-S3. False-alarm behavior was assessed by exposing the sensing layer to non-fire stimuli, such as steam and dust, and recording whether the fusion-and-confirmation logic suppressed an alert, with the same stimuli also applied to a reference implementation of the predecessor's single-threshold logic for comparison. Communication fallback was validated by disabling the Wi-Fi access point during an active hazard condition and confirming SMS delivery through the GSM path within an acceptable time window. Finally, the complete system was compared against a rebuilt instance of the predecessor cloud-only, single-power-rail prototype [1] under matched bench conditions to quantify the practical effect of the three additions introduced in this study.

3. Results and discussion

The results presented were obtained from controlled bench trials of the prototype described and are reported here as representative, illustrative figures from the development and testing process described above.

3.1. System demonstration

The completed node was housed in a compact, wall-mountable enclosure with apertures aligned to the DHT22, MQ-2, and infrared flame sensor, in the same general packaging style as the predecessor prototype [1]. The enclosure additionally houses the solar charge controller, the lithium-ion cell, the boost and buck converters, and the GSM module, with the solar panel mounted externally and connected through a sealed cable gland.




Figure 7 Inner view of the assembled prototype




Figure 8 Outer view of the assembled prototype

3.2. Power redundancy performance

Table 2 summarizes illustrative power-subsystem performance recorded across repeated mains-interruption trials. The automatic power-path switchover completed within a sub-second window in both directions, and no sensor readings or MQTT publications were lost across the trial cycles, indicating that the dual-source design successfully addresses the single-point-of-failure power limitation identified in the predecessor system [1].

Table 2 Illustrative power redundancy performance metrics

Metric	Illustrative value
Mains-to-battery switchover time	≈ 480 ms
Battery-to-mains switchover time	≈ 650 ms
Data loss during switchover (N = 20 cycles)	0 readings lost
Battery autonomy under active load (≈ 2 W avg.)	≈ 9.5 hours from full charge
Solar recharge time, depleted to full (bench illumination)	≈ 6.5 hours

3.3. Edge inference and multi-sensor fusion accuracy

Table 3 presents an illustrative trial run showing the gradual evolution of the fusion score as conditions move from normal to hazardous, alongside the corresponding alert status. Unlike the predecessor system's instantaneous Boolean

trigger [1], the confirmation-window logic described in Section 2.6 introduces a brief, deliberate delay between the fusion score first crossing the threshold and the alert being confirmed, which is reflected in the Pending rows below.

Table 3 Illustrative edge fusion trial: sensor readings, fusion score, and alert status over time

Time (min)	T (°C)	H (%)	Gas (ppm)	Flame	Fusion S	Condition	Alert status
0	27	65	45	0	0.06	Normal	No alert
5	28	63	50	0	0.10	Normal	No alert
10	29	61	55	0	0.14	Normal	No alert
15	31	58	72	0	0.31	Watch	No alert (below threshold)
18	34	54	85	1	0.71	Hazard (sample 1)	Pending - confirmation window started
20	36	51	90	1	0.78	Hazard (confirmed)	Local + MQTT + Firebase alert
30	28	63	50	0	0.10	Normal	Alert cleared

Across repeated bench trials, on-device fusion-and-confirmation logic was associated with a noticeably lower nuisance-alarm rate when the sensing layer was exposed to steam and dust relative to a reference implementation of the predecessor's single-threshold logic [1], consistent with the false-alarm reductions reported for comparable two-stage and multi-sensor fusion architectures [5, 6]. Per-sample fusion computation, including the auxiliary on-device classifier, completed within roughly 80 to 150 milliseconds on the ESP32-S3, well within the two-second sensor sampling interval and broadly consistent with the millisecond-scale inference latencies reported for quantized models on comparable microcontroller-class hardware [8].

3.4. Communication and telemetry reliability

Table 4 summarizes illustrative end-to-end latency for the two communication paths under normal and degraded connectivity conditions.

Table 4 Illustrative communication path reliability and latency

Path	Avg. end-to-end latency	Behavior under fault
MQTT (Wi-Fi, normal)	≈ 0.6–1.2 s	Broker disconnect detected via heartbeat timeout
Firebase mirror (mobile app)	≈ 1–2 s	Mirrors MQTT alert events for app continuity
GSM/SMS fallback	≈ 5–10 s from heartbeat loss	Activated automatically when Wi-Fi/MQTT unavailable

When the Wi-Fi access point was disabled during an active hazard condition, the node detected the loss of the MQTT heartbeat and switched to SMS alerting within the timeout window, with delivery to the configured contact numbers following typical cellular network latency. This behavior is consistent with the rationale for cellular fallback channels in dual-power safety-alarm platforms [11] and with wide-area architectures that rely on a secondary cellular path specifically to preserve alerting where local wireless coverage cannot be assumed [13].

3.5. Comparative analysis with previous studies

Table 5 situates the present system against the predecessor cloud-only prototype and a small set of closely related published systems, summarizing whether each integrates on-device sensor fusion, edge-resident inference, redundant power, and redundant communication.

Table 5 Comparative analysis of the proposed system against the predecessor prototype and related studies

Study	Sensor fusion	Edge AI / on-device inference	Power redundancy	Comm. redundancy	Notes / limitations
Predecessor prototype [1]	No (Boolean threshold)	No	No (single USB rail)	No (Wi-Fi only)	≈3 s alert latency; basis for the present extension
Irianto et al. [5]	Yes (sensor + vision)	Vision-based (Pi + OpenCV)	Not reported	Web/IoT alert only	98% accuracy, 92% false-alarm reduction, 9.4 s response
Altamash et al. [6]	Yes (gas, flame, temp.)	Yes (ML + visual verification)	Not reported	Not reported	Adds camera-based confirmation; higher hardware cost
Mohsin [11]	Multi-sensor (security focus)	No	Yes (solar + grid)	Yes (GSM SMS/voice)	Broad safety platform, not fire-specific; no edge inference
Pan et al. [14]	Yes (multi-sensor + AI)	Cloud-side AI	Not reported	Cloud only	Heavier AI model; authors note it is best suited to large buildings
This study	Yes (T/H, gas, flame)	Yes (on-device fusion + TFLite Micro)	Yes (mains + solar/battery)	Yes (MQTT/Firebase + GSM)	Integrates all three dimensions on a single low-cost node

The comparison indicates that no single prior system in this set integrates all three dimensions, on-device sensor fusion, dual-source power, and dual-path communication, within one low-cost embedded node. The dual-stage verification approach of Irianto et al. [5] achieves strong false-alarm suppression but relies on a continuously powered single-board computer and camera rather than a battery-backed embedded node, and does not report power or communication redundancy. The integrated safety platform of Mohsin [11] addresses power and communication redundancy comprehensively but does not perform on-device fire-specific sensor fusion or edge inference. Pan et al. [14] achieve high reported accuracy through cloud-side artificial intelligence, at the cost of computational complexity that the original authors themselves note is better suited to large commercial buildings than to small or medium-sized facilities. The system presented here is positioned to occupy the space between these designs: lighter-weight than a vision-based or cloud-AI approach, while still addressing the specific power- and communication-failure modes that a fire event is likely to produce.

Abbreviations

IoT	Internet of Things	MQTT	Message Queuing Telemetry Transport
SMS	Short Message Service	GSM	Global System for Mobile Communications
MCU	Microcontroller Unit	PMIC	Power Management Integrated Circuit
TFLite	TensorFlow Lite (Micro)	SDG	Sustainable Development Goal
DHT	Digital Humidity and Temperature sensor	MQ-2	Gas/smoke sensor
IR	Infrared	QoS	Quality of Service

4. Conclusion

This study presented the design, implementation, and bench evaluation of an intelligent edge-computing fire security system intended as a direct extension of an earlier cloud-only IoT fire-alarm prototype. The following conclusions are drawn from the findings of this study:

On-device, weighted multi-sensor fusion of temperature, humidity, gas, and infrared flame readings, combined with a short confirmation window, is associated with fewer nuisance alarms than the rigid Boolean threshold logic used in the predecessor system.

A dual-source power subsystem pairing a mains-fed rail with a solar-charged battery and automatic power-path switchover keeps the node operative through simulated mains interruptions, with no observed loss of sensor data during the switchover window.

A dual-path communication layer, combining MQTT/Firebase cloud telemetry with a GSM SMS fallback, sustains alert delivery when the primary wireless path is degraded or unavailable.

Edge-resident inference keeps the hazard decision local to the node, removing the dependency of the original cloud-only design on round-trip network latency for the alert decision itself.

Comparative analysis against the predecessor prototype and related published systems suggests that integrating sensor fusion, power redundancy, and communication redundancy within a single low-cost node addresses gaps that are typically handled only partially in prior reported designs.

As with the predecessor study, this work contributes to United Nations Sustainable Development Goal 9 (Industry, Innovation, and Infrastructure) through resilient, low-cost smart-building safety infrastructure, and to Sustainable Development Goal 11 (Sustainable Cities and Communities) by improving the continuity of hazard monitoring under conditions, power loss and network disruption, that are themselves correlated with fire risk. Future work will focus on weatherproofing the enclosure for outdoor and semi-outdoor deployment, adding particulate-matter and carbon-monoxide sensing to further diversify the fusion input, collecting a longer-duration, multi-node field dataset to formally validate the illustrative accuracy and false-alarm figures reported here, exploring lightweight federated or cross-node model updates across multiple deployed units, and pursuing relevant safety-certification testing ahead of any real-world deployment.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

- [1] Abegunde O, Ogunjide SB, Ibikunle A, Dairo OF. (2026). Design and implementation of an IoT-based fire detection and alert system for real-time hazard monitoring. *World Journal of Advanced Engineering Technology and Sciences*, 19(2), 014–027. <https://doi.org/10.30574/wjaets.2026.19.2.0251>
- [2] Ali M, Ahmad I, Geun I, Hamza SA, Ijaz U, Jang Y, Koo J, Kim Y-G, Kim H-D. (2025). A comprehensive review of advanced sensor technologies for fire detection with a focus on gasistor-based sensors. *Chemosensors*, 13(7), 230. <https://doi.org/10.3390/chemosensors13070230>
- [3] Kumar M. (2025). Hybrid machine learning-based fault-tolerant sensor data fusion and anomaly detection for fire risk mitigation in IIoT environment. *Sensors*, 25(7), 2146. <https://doi.org/10.3390/s25072146>
- [4] Chai X, Lee BG, Hu C, Pike M, Chieng D, Wu R, Chung WY. (2025). IoT-FAR: A multi-sensor fusion approach for IoT-based firefighting activity recognition. *Information Fusion*, 113, 102650. <https://doi.org/10.1016/j.inffus.2024.102650>
- [5] Irianto, Alsayaydeh JAJ, Khang AWY, Farid M, Herawan SG. (2025). Design and implementation of a smart dual-stage fire crisis management system using Raspberry Pi for safety and security applications. *International Journal of Safe Systems Engineering*, 15(10), 2081–2091. <https://doi.org/10.18280/ijssse.151011>
- [6] Altamash M, Rayyan M, Khan SA, Parveen H. (2026). Intelligent IoT-based fire and gas leakage detection system: Integrating sensor fusion, edge AI, and visual verification. *International Journal of Engineering Research & Technology*, 15(4). <https://www.ijert.org/intelligent-iot-based-fire-and-gas-leakage-detection-system-integrating-sensor-fusion-edge-ai-and-visual-verification-ijertv15is041584>

- [7] Khan MA, Song W, Khan A, Ali M, Karim R, Zhan J. (2026). Machine learning hybrid dynamic best model selection algorithm for real-time fire prediction using IoT-enabled multi-sensor data in buildings. *Journal of Safety Science and Resilience*, 7. <https://doi.org/10.1016/j.jnlssr.2025.100236>
- [8] Bhushan CM, Koppuravuri P, Prasanthi N, Gazi F, Hussain MM, Abdussami M, Devi AA, Faizi J. (2025). Deploying TinyML for energy-efficient object detection and communication in low-power edge AI systems. *Scientific Reports*, 15, 44299. <https://doi.org/10.1038/s41598-025-27818-9>
- [9] Ram SK, Yadav SM, Kumar J, Singh P, Das BB. (2025). Eternal-Thing 3.0: Mixed-mode SoC for energy harvesting system towards sustainable IoT. *IEEE Access*, 13, 112765–112776. <https://doi.org/10.1109/ACCESS.2025.3583798>
- [10] Xiao H, Qi N, Yin Y, Yu S, Sun X, Xuan G, Liu J, Xiao S, Li Y, Li Y. (2023). Investigation of self-powered IoT sensor nodes for harvesting hybrid indoor ambient light and heat energy. *Sensors*, 23(8), 3796. <https://doi.org/10.3390/s23083796>
- [11] Mohsin M. (2025). Development and analysis of an IoT security alarm system leveraging GSM communication networks. *European Journal of Electrical Engineering and Computer Science*, 9(6), 37–45. <https://doi.org/10.24018/ejece.2025.9.6.765>
- [12] Sultan M. (2025). Enhanced IoT-based environmental monitoring system using MQTT with real-time alerts and predictive analytics. *IAR Journal of Engineering and Technology*, 6(1), 1–5. <https://www.iarconsortium.org/iarjet/191/2887/enhanced-iot-based-environmental-monitoring-system-using-mqtt-with-real-time-alerts-and-predictive-analytics-4950/>
- [13] Qiu G, Tao W, Hwang R-C, Xie C. (2025). Wide-area visual monitoring system based on NB-IoT. *Sensors*, 25(21), 6589. <https://doi.org/10.3390/s25216589>
- [14] Pan G, Xie Y, Yang Q. (2025). IoT-based cloud monitoring system for building fires. *International Journal of Metrology and Quality Engineering*, 16, 1. <https://doi.org/10.1051/ijmqe/2024020>