A Scalable MQTT-Based Edge IoT Architecture for Real-Time Distributed Solar PV Panel Monitoring

Ahmed Mohammed
Department of Smart Computing and Cyber Resilience, Faculty of Engineering and Technology, Sunway University, Selangor, Malaysia | Department of Computer Science, Faculty of Science, Gombe State University, Gombe, Nigeria
ahmed.m26@imail.sunway.edu.my

Ranjit Singh Sarban Singh
Research Centre for Human-Machine Collaboration (HUMAC), School of Engineering, Faculty of Engineering and Technology, Sunway University, Selangor, Malaysia
ranjits@sunway.edu.my (corresponding author)

Saad Aslam
Department of Smart Computing and Cyber Resilience, Faculty of Engineering and Technology, Sunway University, Selangor, Malaysia
saada@sunway.edu.my

Yan Chiew Wong
Fakulti Teknologi dan Kejuruteraan Elektronik dan Komputer, Universiti Teknikal Malaysia Melaka (UTeM), Durian Tunggal Melaka, Malaysia
ycwong@utem.edu.my

Received: 14 December 2025 | Revised: 16 January 2026 | Accepted: 2 February 2026

| DOI: https://doi.org/10.48084/etasr.16945

ABSTRACT

Distributed Photovoltaic (PV) installations require panel-level high-frequency monitoring to detect transient electrical and thermal anomalies that are not observable through aggregated inverter-level measurements. Conventional Supervisory Control and Data Acquisition (SCADA) systems rely on centralized architectures, polling-based industrial protocols, and aggressive data aggregation, making them unsuitable for distributed and rooftop PV deployments operating under intermittent network connectivity. Although Internet of Things (IoT)-based monitoring solutions offer low-cost and flexible alternatives, most existing implementations remain cloud-centric and rely solely on Message Queuing Telemetry Transport (MQTT) broker-level Quality of Service (QoS), which does not guarantee durable persistence at the application database layer. This study presents a scalable MQTT-based edge IoT architecture for real-time distributed solar PV monitoring that guarantees true end-to-end data reliability. The proposed system integrates panel-level electrical and environmental sensing with disk-backed edge persistence and an application-level acknowledgment mechanism (MQTT-ACK) that issues delivery confirmation only after successful fog-side database insertion. By coupling acknowledgment with durable storage rather than broker reception, the proposed architecture explicitly decouples sensing reliability from network availability. Experimental evaluation using real-world PV data demonstrates complete data recovery during network disruptions, bounded synchronization delay upon reconnection, and low computational and communication overhead on both edge and fog nodes. The results confirm that combining lightweight MQTT communication with edge-level persistence and application-level acknowledgment provides a practical, low-cost, and resilient alternative to centralized SCADA platforms and cloud-centric IoT PV-monitoring systems.

Keywords-PV monitoring; edge IoT architecture; MQTT; distributed systems; real-time sensing

I. INTRODUCTION

The expansion of Photovoltaic (PV) installations has been driven by declining costs, supportive energy policies, and the transition toward low-carbon energy systems. While utility-scale PV plants account for a substantial share of installed capacity, distributed and rooftop PV deployments are growing rapidly across urban, peri-urban, and rural environments. Despite their advantages, distributed PV systems are particularly vulnerable to performance degradation caused by partial shading, dust accumulation, thermal hotspots, wiring faults, and environmental stress. Empirical studies indicate that such factors can result in annual energy losses ranging from 10% to 20% when faults remain undetected for extended periods [1, 2].

Early detection of these degradation mechanisms requires panel-level, high-frequency monitoring capable of capturing short-duration electrical and thermal variations. However, many existing PV monitoring infrastructures are not designed to operate at this granularity. Instead, they rely primarily on inverter-level measurements and periodic aggregation, which obscure transient anomalies and delay maintenance interventions. As a result, early indicators of degradation are often masked until significant performance losses have already occurred.

Supervisory Control and Data Acquisition (SCADA) systems remain the dominant monitoring solution for large utility-scale PV plants. These systems are traditionally built around centralized architectures, polling-based industrial communication protocols such as Modbus and DNP3, and hierarchical control models inherited from conventional power systems [3]. Although SCADA platforms provide robustness in controlled industrial environments, they are poorly suited for distributed PV deployments. Centralized polling introduces latency, aggressive aggregation reduces temporal resolution, and continuous connectivity to a central server is implicitly assumed. Empirical evaluations have shown that SCADA systems frequently compress hundreds of thousands of raw measurements into coarse hourly summaries, leading to substantial loss of information related to transient faults and short-term performance variations [3].

To overcome these limitations, Internet of Things (IoT)-based PV monitoring systems have gained increasing attention. Advances in low-cost embedded platforms and sensor technologies have enabled compact edge devices to perform panel-level measurements of electrical and environmental parameters. Sensors such as the INA226 enable accurate measurement of voltage, current, and power directly at the PV panel, while digital temperature and humidity sensors provide essential contextual information for interpreting PV performance under varying environmental conditions [4, 5]. These capabilities allow significantly higher spatial and temporal resolution than traditional SCADA-based monitoring.

Communication protocols are crucial for enabling scalable IoT-based monitoring. Unlike SCADA protocols that rely on synchronous polling, IoT architectures typically adopt asynchronous publish-subscribe communication models. Among these, Message Queuing Telemetry Transport (MQTT) has emerged as a preferred protocol due to its lightweight header structure, low bandwidth consumption, and suitability for resource-constrained edge devices [6, 7]. Comparative studies demonstrate that MQTT offers lower latency and improved energy efficiency compared to HTTP-based approaches, while providing better scalability for distributed sensing environments.

Despite these advantages, most existing IoT-based PV monitoring solutions remain cloud-centric and architecturally fragile. In some reported implementations, edge nodes stream sensor data directly to a cloud server or MQTT broker without structured local persistence. Although MQTT provides multiple Quality of Service (QoS) levels, these guarantees apply only to message delivery between the publisher and the broker and do not ensure durable persistence at the subscriber or database layer. Prior evaluations of IoT protocols in real-time systems have similarly highlighted that MQTT does not provide deterministic end-to-end delivery guarantees and remains vulnerable to data loss under unstable network conditions, particularly when persistence is required beyond broker-level confirmation [8]. Consequently, data loss may still occur during broker outages, backend failures, database write errors, or prolonged network disruptions, conditions that are common in rooftop and rural PV deployments [9, 10].

Research has explored improvements in the scalability, security, and ingestion performance of IoT systems through layered architectures, encryption mechanisms, and intrusion detection techniques [11, 12]. While these efforts address important aspects of IoT robustness, the challenge of end-to-end data persistence under intermittent connectivity remains insufficiently addressed in PV monitoring contexts. Many systems implicitly assume continuous cloud availability, leaving a significant reliability gap between edge-level sensing and long-term data storage. In [13], end-to-end acknowledgment mechanisms in MQTT-based systems were explored by introducing End-to-End MQTT (E-MQTT) protocol-level extensions to enable publishers to verify message reception by subscribers through modified QoS semantics. Although effective, such approaches require changes to the MQTT protocol stack and introduce additional message exchange overhead, which may limit compatibility with existing deployments and resource-constrained edge devices.

The current work addresses this gap by proposing a scalable MQTT-based edge IoT architecture that explicitly integrates disk-backed edge persistence with an application-level acknowledgment mechanism. The key novelty lies in the architectural enforcement of true end-to-end data reliability at the application layer, whereby each PV measurement is acknowledged only after successful persistence in the fog database. Unlike conventional MQTT-based solutions that rely solely on broker-level QoS, the proposed approach couples fog-side database commits with edge-side record lifecycle management, ensuring that measurements are neither lost nor prematurely discarded. Records remain cached locally during network outages and are synchronized only after confirmed database insertion. In contrast to protocol-level extensions such as E-MQTT, the proposed architecture operates entirely at the application layer, preserving full MQTT compatibility while guaranteeing durable persistence through disk-backed edge caching and database-level acknowledgment.

Accordingly, this study makes the following contributions: (1) the design and implementation of a low-cost, edge-based PV monitoring node with integrated electrical and environmental sensing, (2) the development of a lightweight application-level acknowledgment mechanism that guarantees durable end-to-end data persistence beyond broker-level QoS semantics, (3) experimental validation using real-world PV measurements, including quantitative evaluation of data recovery rate, synchronization delay, and resource overhead under network disruptions, and (4) a systematic architectural comparison demonstrating how the proposed approach addresses the limitations of centralized SCADA systems and conventional cloud-centric IoT PV monitoring solutions.

II. METHODOLOGY

This section describes the design, implementation, and operational logic of the proposed MQTT-based edge IoT architecture for real-time distributed solar PV monitoring. The methodology is presented in terms of system architecture, sensing and edge processing, communication and synchronization mechanisms, and security considerations to ensure reproducibility.

A. System Architecture

The proposed system adopts a distributed edge-fog architecture composed of PV sensing nodes, an MQTT-based communication layer, and a fog ingestion and storage layer. The proposed monitoring system adopts a distributed edge-fog architecture, in which each PV panel is instrumented with an autonomous edge node responsible for sensing and local persistence. Measurements are transmitted asynchronously using MQTT to a fog node that performs durable storage, acknowledgment generation, and data dissemination to a dashboard application for real-time visualization and analysis. This architecture decouples sensing reliability from cloud availability while enabling scalable multi-node deployment. The high-level interaction between edge nodes, MQTT brokers, and the fog layer is illustrated in Figure 1.

Each PV panel is monitored by an autonomous edge node responsible for data acquisition, local persistence, and controlled synchronization with the fog. The communication layer uses MQTT to decouple publishers and subscribers, while the fog layer performs durable storage, acknowledgment generation, and system-level monitoring.

The architecture is explicitly designed to tolerate intermittent connectivity and backend disruptions by decoupling sensing reliability from network availability. This is achieved through disk-backed persistence at the edge and application-level acknowledgment issued only after a successful fog-side database insertion. The communication configuration and reliability mechanisms adopted in the proposed system are summarized in Table I.

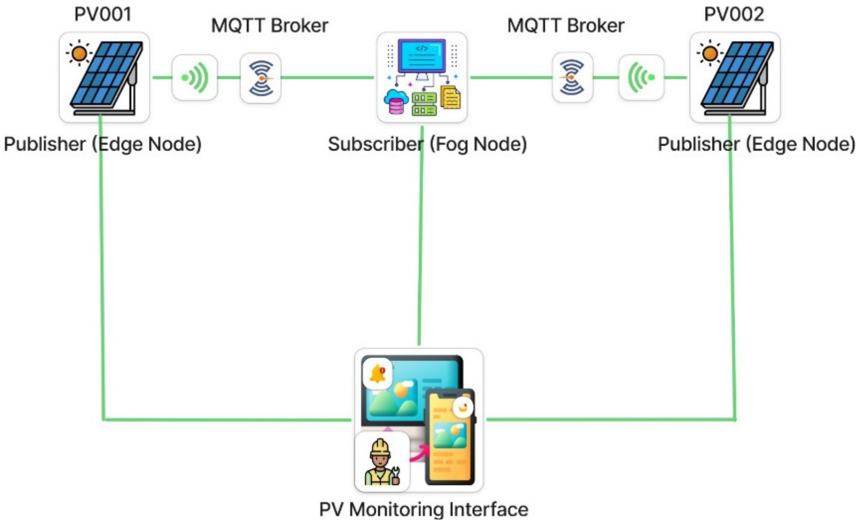

Fig. 1. High-level architecture of the proposed MQTT-based edge-fog system for distributed solar PV monitoring.

TABLE I. MQTT COMMUNICATION CONFIGURATION AND RELIABILITY MECHANISMS

Parameter	Configuration / mechanism	Rationale
Communication protocol	MQTT	Lightweight, low-latency, suitable for resource-constrained devices
MQTT QoS level	QoS 1	Ensures delivery to broker while avoiding QoS 2 handshake overhead
Publish model	Asynchronous publish-subscribe	Decouples sensing from cloud availability
Application-level acknowledgement (ACK)	ACK after database commit	Guarantees end-to-end data persistence
Edge persistence	Local MariaDB database	Prevents data loss during outages
Fog ingestion	Non-blocking queue + UPSERT	Prevents backlog and ensures idempotency
Duplicate handling	Sequence-based record tracking	Safe retransmission under QoS 1

B. Edge Node Design and Operation

1) Data Acquisition and Local Persistence

Each edge node was built around a Raspberry Pi Zero W and interfaced with multiple electrical and environmental sensors using standardized digital buses. The INA226 sensor was used for voltage, current, and power measurements, whereas the DS18B20 and SHT45 sensors provided panel and ambient environmental data. The detailed wiring and GPIO interfacing configuration used to integrate these sensors with the edge node is shown in Figure 2.

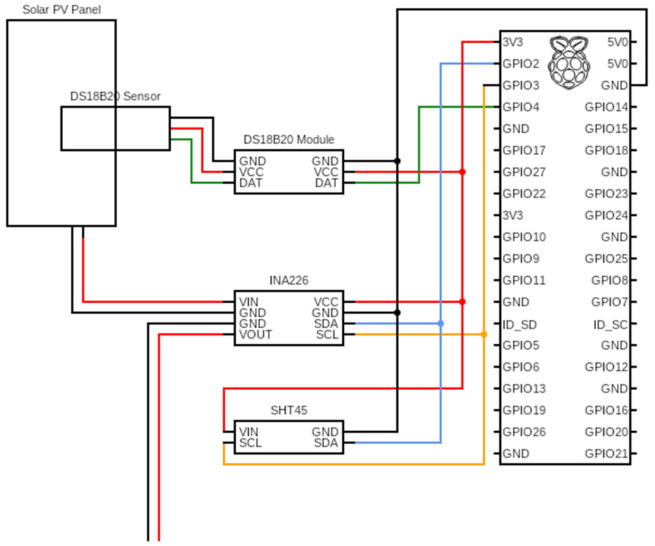

Fig. 2. Electrical and environmental sensor interfacing with the Raspberry Pi Zero W edge node.

Each record was assigned a monotonically increasing sequence number and a timestamp generated at the edge. The records were initially marked as unsynchronized and remained stored locally, regardless of the network state. This disk-backed persistence ensures that no measurements are lost during network outages, broker failures, or cloud-side unavailability.

2) Edge Metrics and Resource Monitoring

To quantify the system overhead, the edge node continuously monitors its own resource utilization. CPU usage, memory consumption, and estimated energy usage were measured and embedded within each transmitted payload. These metrics enable quantitative evaluation of the computational and energy overhead introduced by the proposed reliability mechanism and support the claims regarding suitability for resource-constrained edge devices.

C. MQTT Communication and Synchronization Strategy

1) MQTT-ACK Configuration

MQTT is employed as the primary communication protocol using a publish-subscribe model. The edge nodes publish measurement payloads using QoS level 1 to ensure reliable delivery to the broker while tolerating occasional duplicates. QoS level 2 is intentionally avoided due to its multi-step handshake and increased processing overhead on constrained devices.

Distinct topics are used for data transmission, "solar/panel/data", and acknowledgments, "solar/panel/ACK", allowing independent scaling and simplified message handling. Edge nodes retransmit unsynchronized records upon reconnection without a fixed retry timeout, relying on acknowledgment reception to determine completion.

2) Application-Level Acknowledgment (MQTT-ACK)

To guarantee end-to-end data reliability, the system employs an application-level acknowledgment mechanism that explicitly links acknowledgment issuance to successful database persistence at the fog node. Measurements were logged locally, transmitted via MQTT, and the status was updated to synced at the edge node storage only after acknowledgment following database insertion. The operational workflow of sensing, buffering, transmission, acknowledgment, and looping behavior is depicted in Figure 3.

Unlike standard QoS guarantees, which terminate at the broker, the proposed MQTT-ACK mechanism explicitly couples acknowledgment with durable cloud-side persistence.

The synchronization workflow operates as follows:

- The edge node publishes a measurement payload containing a unique sequence number and timestamp.
- The fog subscriber receives the payload and places it into a non-blocking processing queue.
- The payload is validated and inserted into the fog database using idempotent UPSERT semantics.
- Only after a successful database commit does the fog node publish an acknowledgment message containing the corresponding sequence number.
- Upon receiving the acknowledgment, the edge node marks the local record as being synchronized.

Records remain cached on the local edge node until the system receives a valid acknowledgment confirming that the data have been successfully stored in the fog node, not just received by the broker.

D. Fog Ingestion and Storage Layer

The fog node is responsible for ingesting MQTT messages, persisting data, and issuing acknowledgments. To prevent blocking under high-throughput conditions, incoming messages are placed into a bounded, non-blocking queue prior to database insertion.

Database operations employ UPSERT semantics, ensuring idempotent handling of retransmitted records and preventing duplication caused by MQTT QoS 1 retransmissions. This design allows the system to tolerate duplicate deliveries without compromising data integrity.

In addition to data persistence, the fog node records the processing latency, CPU utilization, and memory usage, enabling quantitative evaluation of the backend overhead and synchronization behavior.

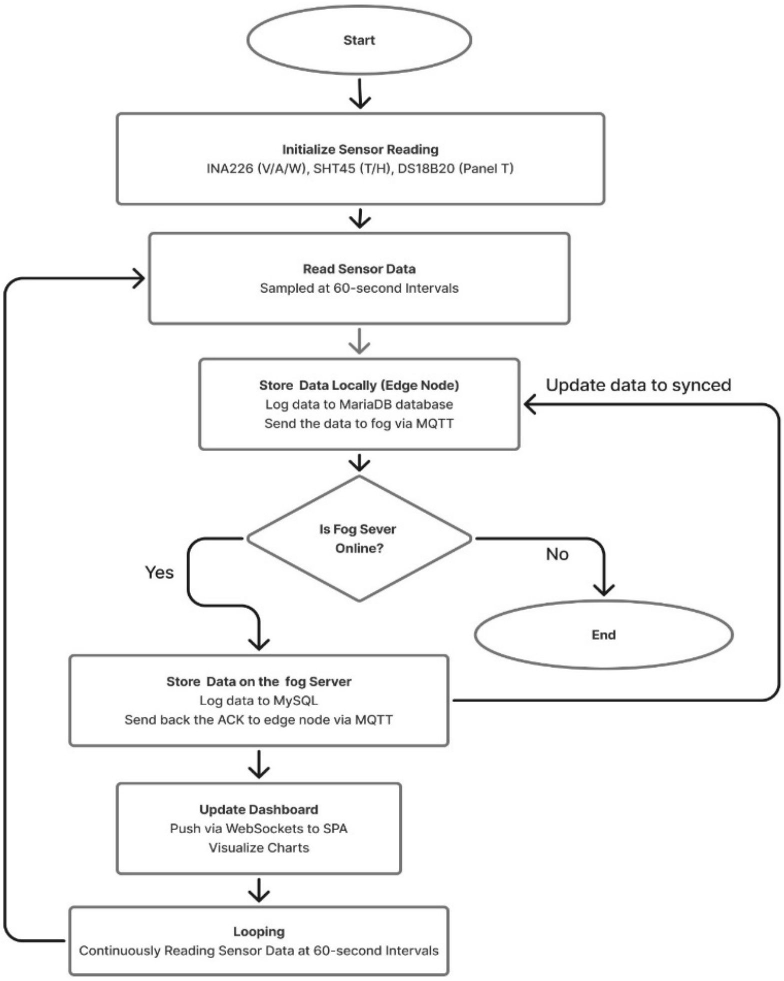

Fig. 3. Application-level acknowledgment (MQTT-ACK) workflow for reliable data synchronization.

E. Failure Handling and Autonomous Operation

The system was explicitly designed for autonomous operation under failure conditions. Edge nodes continue to sense and store data during network outages without blocking or data loss. The fog layer includes automatic reconnection logic for database and broker failures, ensuring continued operation without manual intervention.

Acknowledgments are issued strictly after successful database insertion, guaranteeing that data are never prematurely updated to synced from edge storage.

F. Security Considerations

Security mechanisms are incorporated into multiple layers of the architecture. MQTT connections support TLS encryption to protect data in transit. Broker authentication prevents unauthorized publishers and subscribers. Database access is restricted using credential-based authentication.

Acknowledgment messages include record identifiers to prevent arbitrary ACK injection. Although comprehensive intrusion detection and replay attack mitigation are outside the scope of this study, these aspects are identified as important directions for future work.

III. RESULTS AND DISCUSSION

This section presents and discusses the experimental results obtained from the deployed MQTT-ACK edge IoT PV monitoring system. The evaluation focuses on real-time sensing performance, data reliability under network disruptions, synchronization behavior, and interpretation of observed PV electrical and thermal variations.

A. Experimental Deployment and Data Collection

The proposed architecture was experimentally validated using a PV panel instrumented with electrical and environmental sensors and connected to a Raspberry Pi Zero W edge node. The monitoring system was deployed under semi-outdoor conditions using a compact enclosure to protect the edge electronics while maintaining sensor exposure. This setup reflects the practical installation constraints in rooftop and distributed PV environments. The deployed edge node used for the experimental evaluation is displayed in Figure 4. During system development and validation, a prototype setup was used to verify sensor integration, data acquisition, and communication logic under controlled conditions. The prototype configuration, including the PV module, sensors, and edge processing unit, is illustrated in Figure 5.

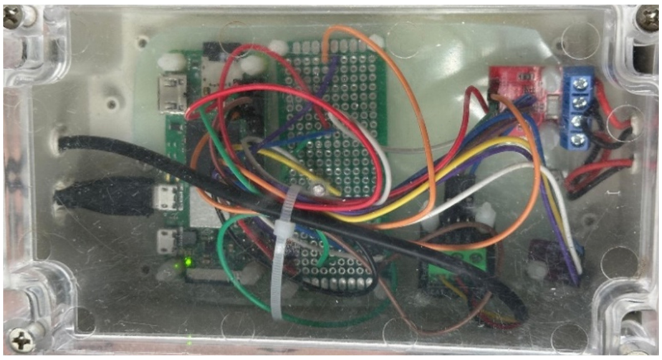

Fig. 4. Deployed edge monitoring node enclosed for semi-outdoor operation.

Electrical parameters (voltage, current, and power) and environmental parameters (panel temperature, ambient temperature, and humidity) were sampled at 60-s intervals and stored locally at the edge before fog synchronization.

The dataset consists exclusively of real-world measurements collected during normal PV operation and during controlled network disruption scenarios. No synthetic data were used for the evaluation. Local storage logs confirmed continuous data acquisition, regardless of network availability, demonstrating autonomous edge operation.

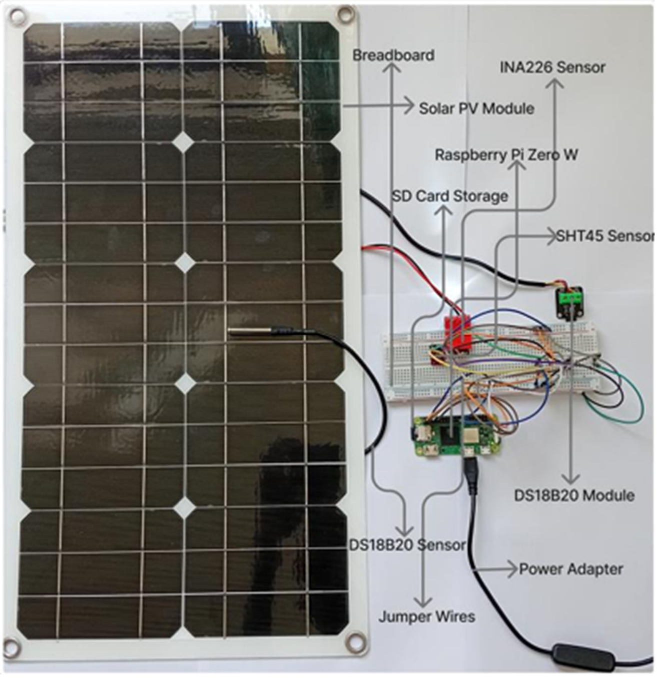

Fig. 5. Prototype PV monitoring setup used during system development and testing.

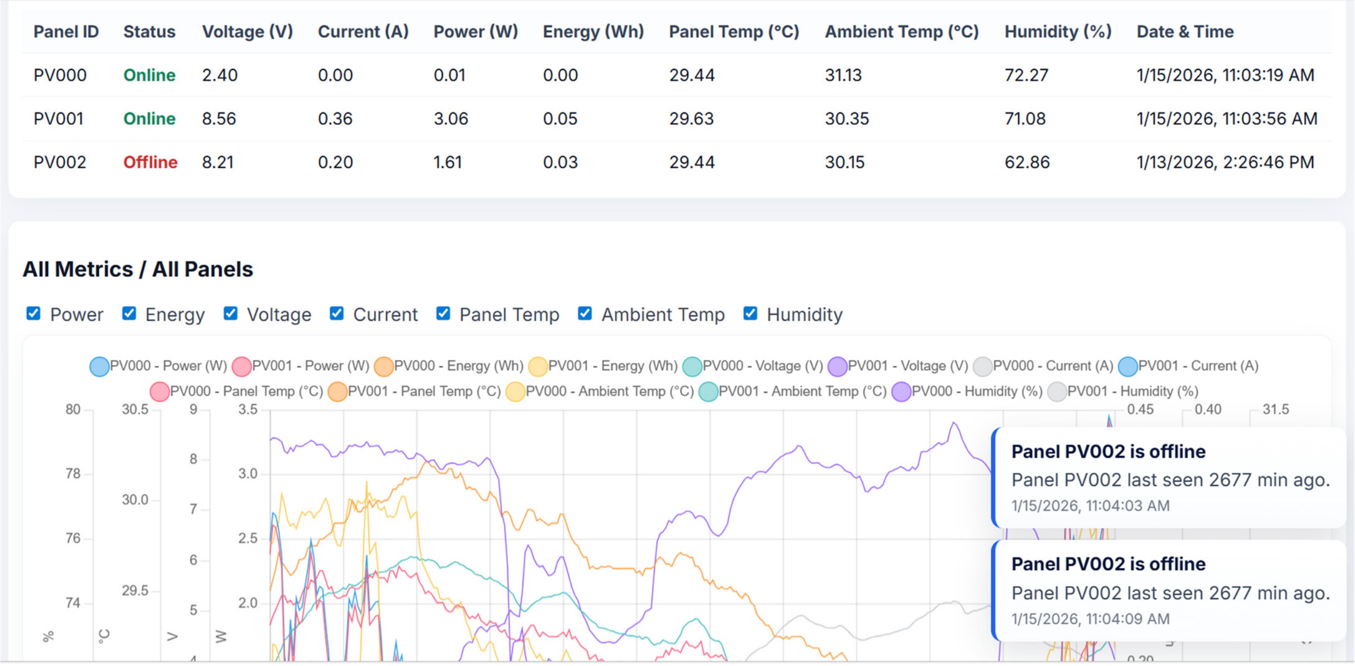

Panel ID	Status	Voltage (V)	Current (A)	Power (W)	Energy (Wh)	Panel Temp (°C)	Ambient Temp (°C)	Humidity (%)	Date & Time
PV000	Online	2.40	0.00	0.01	0.00	29.44	31.13	72.27	1/15/2026, 11:03:19 AM
PV001	Online	8.56	0.36	3.06	0.05	29.63	30.35	71.08	1/15/2026, 11:03:56 AM
PV002	Offline	8.21	0.20	1.61	0.03	29.44	30.15	62.86	1/13/2026, 2:26:46 PM

All Metrics / All Panels

☑ Power ☑ Energy ☑ Voltage ☑ Current ☑ Panel Temp ☑ Ambient Temp ☑ Humidity

Panel PV002 is offline
Panel PV002 last seen 2677 min ago.
1/15/2026, 11:04:03 AM

Panel PV002 is offline
Panel PV002 last seen 2677 min ago.
1/15/2026, 11:04:09 AM

Fig. 6. Real-time PV monitoring dashboard showing multi-panel measurements, node availability status, and offline detection alerts.

B. *Real-Time Monitoring Performance*

During normal connectivity, the system successfully transmitted the sensor measurements to the server in near real time. The data were visualized through a web-based dashboard using WebSocket updates, enabling continuous monitoring of the PV electrical output and environmental conditions. The dashboard aggregates measurements from multiple PV panels, displays the current electrical and environmental parameters, and reports node availability status. Offline detection and alert notifications are triggered when an edge node fails to synchronize within a predefined interval of at least 5 min. An example of the real-time monitoring interface during system operation is shown in Figure 6. The end-to-end data flow from sensing to visualization exhibited stable operation without observable packet loss or dashboard inconsistencies.

The low communication overhead of MQTT allowed the edge node to maintain consistent performance despite limited processing and memory resources. This confirms the suitability of lightweight publish-subscribe communication for continuous PV telemetry in constrained environments.

C. *Reliability under Network Disruptions*

To evaluate reliability under intermittent connectivity, controlled network outages were introduced by disconnecting the edge nodes from the network for 30 min during a 24 h run. Outage durations were chosen to reflect realistic short- to medium-term connectivity loss in rooftop deployments. During this outage, sensor data continued to be acquired and stored locally in the MariaDB database. During network disruptions, the edge node continues to operate autonomously by persisting all measurements locally until connectivity is restored. Once the network becomes available, the buffered records are synchronized in batches and acknowledged only after successful fog-side persistence.

Inspection of the local database during outage confirmed that all measurements were retained with a synchronization flag indicating unsynced status. Upon network restoration, the edge nodes initiated synchronization by publishing buffered records to the server using the MQTT-ACK mechanism.

Analysis of synchronization logs showed that all buffered records were successfully transmitted and persisted in the fog database, resulting in a 100% data recovery rate across the tested outage durations. No measurements were lost or duplicated at the fog database layer, confirming that the proposed architecture ensures reliable data continuity under network instability. To quantify the impact of edge caching and application-level acknowledgment on data reliability, a comparative analysis was conducted under identical network disruption conditions, as outlined in Table II.

In the absence of edge-side persistence and acknowledgment, records generated during network outages are permanently lost once transmission fails, resulting in partial data delivery. In contrast, the proposed architecture retains unsent records locally at the edge and synchronizes them only after successful cloud-side database persistence is confirmed. As shown in Table II, the proposed approach achieved complete data recovery with zero data loss under identical experimental conditions, clearly demonstrating the effectiveness of disk-backed caching combined with application-level acknowledgment.

TABLE II. COMPARISON OF DATA DELIVERY OUTCOMES WITH AND WITHOUT EDGE CACHING AND APPLICATION-LEVEL ACKNOWLEDGMENT

Test scenario/case	**Records generated**	**Records cached at edge node**	**Records successfully delivered to fog node**	**Records lost**	**Data recovery rate (%)**
Without edge cache and ACK	2618	0	2558	60	97.7%
With edge cache + MQTT-ACK (proposed)	2618	60	2618	0	100%

D. *Synchronization Delay and Latency Analysis*

The synchronization delay is defined as the time difference between record creation at the edge and acknowledgment receipt after a successful database insertion at the fog node. This metric reflects the buffering behavior, queueing delay, and database processing time. The distribution of synchronization delay observed after network recovery is exhibited in Figure 7.

Communication latency was evaluated by measuring the time between message publication at the edge and reception at the fog node. This metric captures the network transmission delay and broker-level processing time. The observed latency distribution for all transmitted records is shown in Figure 8.

This behavior is expected, as buffered data are transmitted in batches rather than continuously. The synchronization delay did not affect data integrity, and all records eventually persisted without manual intervention.

Latency measurements include network transmission time, queueing delay, and database processing time. The bounded delay distribution confirms that the MQTT-ACK mechanism does not introduce excessive overhead and scales predictably with workload size.

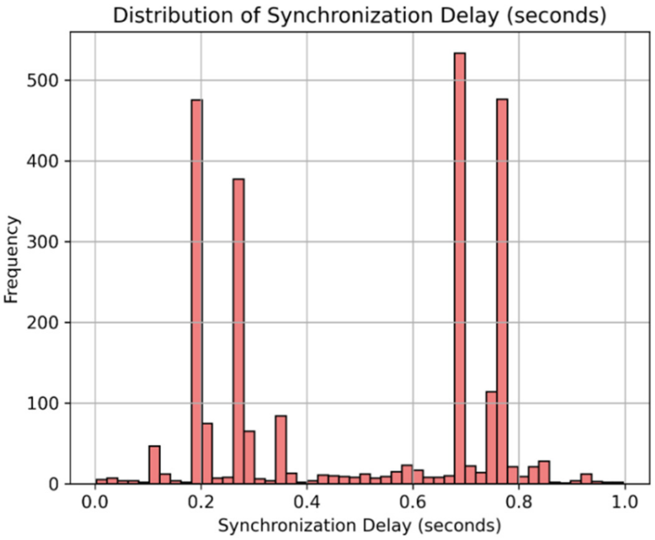

Fig. 7. Distribution of synchronization delay following acknowledgment-based recovery, measured as the time between record creation at the edge and database-commit acknowledgment at the fog node.

The results demonstrate that the MQTT-ACK mechanism effectively balances reliability and communication efficiency without overwhelming constrained edge devices.

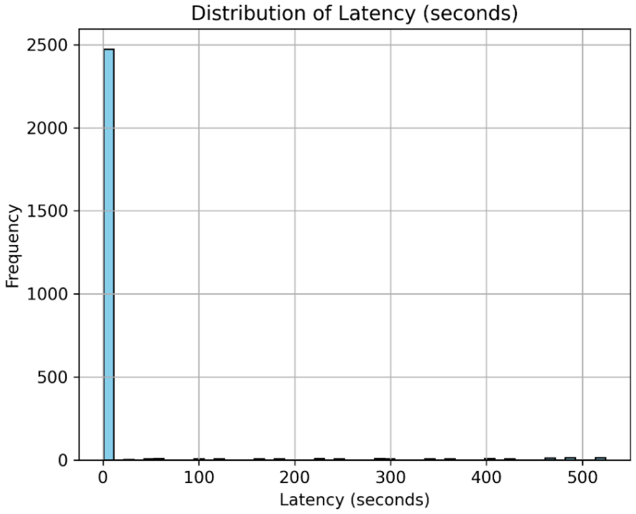

Fig. 8. Distribution of end-to-end communication latency between edge publication and fog-side message reception during normal connectivity.

E. *Edge and Fog Resource Overhead*

Edge resource monitoring shows that CPU utilization, memory usage, and energy consumption remained low and stable throughout system operation. The additional overhead introduced by local persistence and acknowledgment handling was minimal, confirming suitability for low-power edge devices. To evaluate the computational and energy overhead introduced by local persistence and application-level acknowledgment, system resource utilization was continuously monitored at the edge node during normal operation and post-outage synchronization. Metrics include CPU utilization, memory consumption, and estimated energy expenditure, providing a holistic view of edge-side overhead. The statistical distributions of these resource metrics are presented in Figure 9.

Similarly, fog resource monitoring indicates stable CPU and memory usage, even during post-outage synchronization bursts. The use of non-blocking queues and idempotent database writes prevented backlog accumulation and ensured consistent ingestion performance.

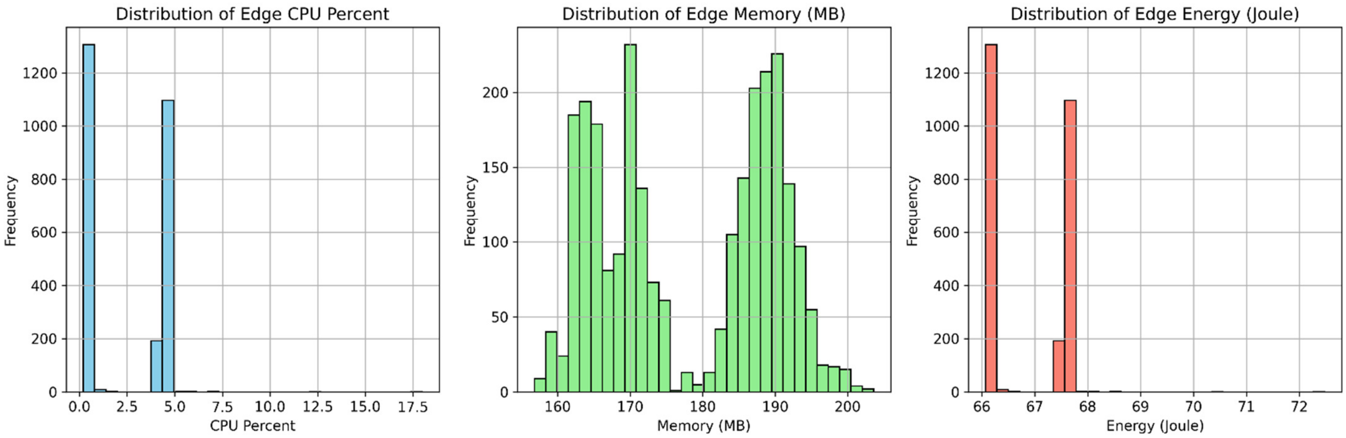

Fig. 9. Distribution of edge node resource utilization: CPU usage, memory consumption, and estimated energy expenditure during system operation.

F. *Interpretation of PV Electrical and Environmental Behavior*

Panel-level monitoring enables the observation of short-term electrical and thermal dynamics that are typically obscured by inverter-level aggregation. Time-series analysis revealed the temporal relationship between voltage, current, and environmental conditions, as well as the influence of temperature on the PV output. Representative electrical and environmental trends captured during the experimental deployment are displayed in Figure 10.

These observations confirm that panel-level high-frequency monitoring provides valuable insights into short-term performance variations that are typically obscured by inverter-level aggregation in conventional SCADA systems. By capturing synchronized electrical and environmental data, the proposed system supports more accurate fault diagnosis and performance assessment.

G. *Comparison with SCADA and Conventional IoT Monitoring*

Compared to SCADA-based monitoring systems, the proposed architecture provides significantly higher temporal and spatial resolution, enabling the detection of transient anomalies that are otherwise lost through aggregation. Unlike conventional IoT-based PV monitoring solutions that depend on continuous cloud connectivity, the proposed system introduces structured edge persistence and controlled synchronization, ensuring reliable operation under real-world network conditions.

The experimental results confirm that the proposed MQTT-ACK edge architecture successfully bridges the reliability gap between centralized SCADA platforms and cloud-dependent IoT solutions, offering a practical approach for distributed and rooftop PV monitoring. A comparison between SCADA-based systems, conventional IoT monitoring, and the proposed architecture is provided in Table III.

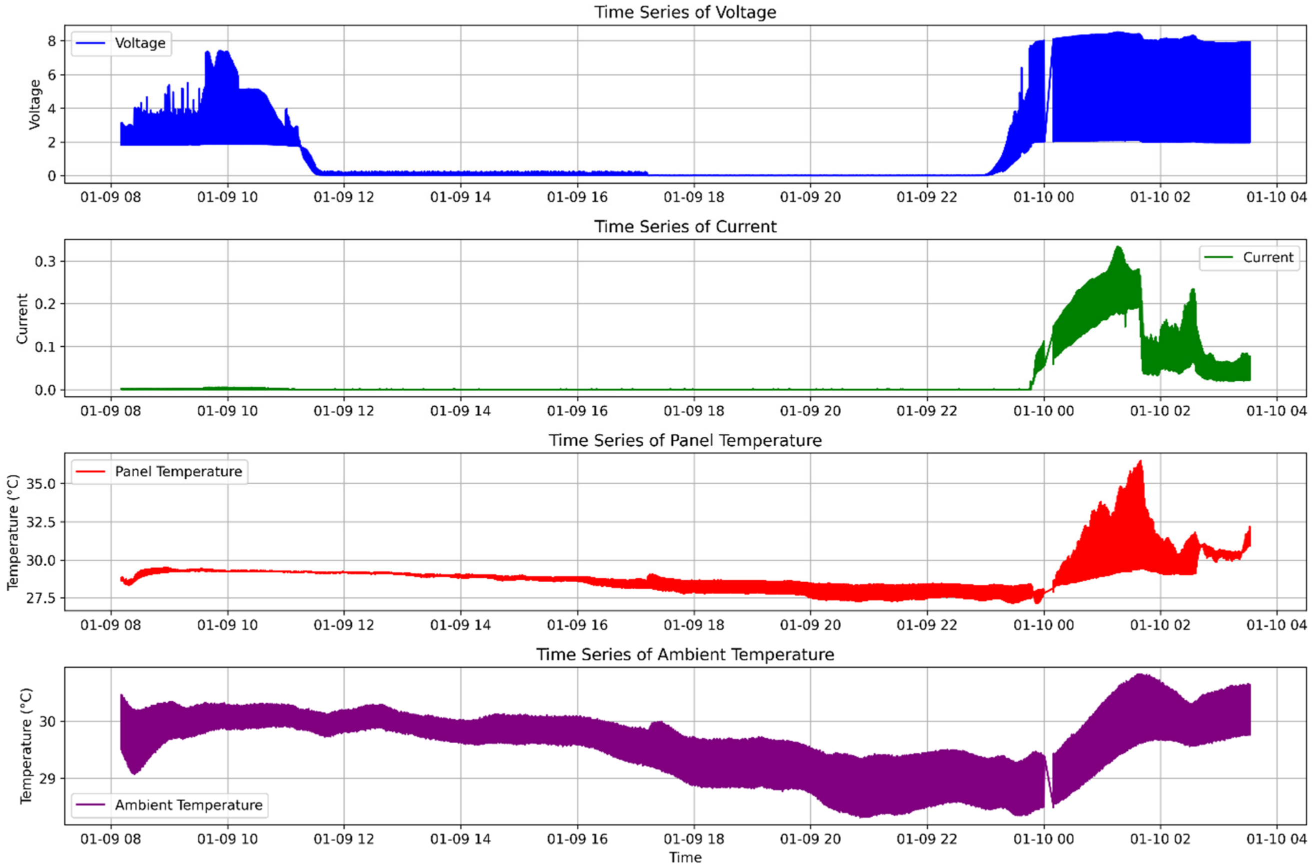

Fig. 10. Time-series variation of PV parameters during experimental deployment: voltage, current, panel temperature, and ambient temperature.

TABLE III. ARCHITECTURAL COMPARISON OF PV MONITORING APPROACHES

Feature	**SCADA-based monitoring**	**Conventional IoT monitoring**	**Proposed architecture**
Measurement granularity	Inverter-level	Panel-level	Panel-level
Sampling frequency	Low (min–h)	Moderate	High (60 s)
Communication model	Centralized polling	Cloud-centric publish-subscribe	Distributed edge-fog publish-subscribe
Edge autonomy	Limited	Limited	Full autonomous operation
Local data persistence	No	Usually absent	Disk-backed persistence
Reliability guarantee	Central server dependent	Broker-level QoS only	End-to-end (ACK after DB commit)
Tolerance to network outages	Low	Moderate	High
Scalability	Limited by polling	Broker-dependent	Horizontally scalable
Suitability for rooftop PV	Limited	Moderate	High

H. Scalability Considerations

Although the experimental validation was conducted using two edge nodes, the proposed architecture is inherently scalable. Each edge node operates independently with its own local database and synchronization logic, while MQTT topics isolate communication streams. The broker remains stateless with respect to data persistence, and backend databases can be horizontally scaled.

As the number of deployed nodes increases, synchronization events occur asynchronously based on individual connectivity conditions, reducing the likelihood of synchronized traffic bursts. This design makes the architecture suitable for large-scale geographically distributed PV deployments.

IV. LIMITATIONS AND FUTURE WORK

Although the proposed MQTT-ACK edge IoT architecture demonstrates reliable operation and complete data recovery under intermittent network connectivity, several limitations remain and warrant further study.

First, the experimental validation was conducted using only two PV-monitoring nodes. Although the architecture is inherently scalable and analytically supports multi-node deployment, large-scale experimental validation involving tens or hundreds of edge nodes was not performed in this study.

Future work will focus on evaluating the system performance under dense deployments, including broker throughput, database write performance, and synchronization behavior under simultaneous reconnection events.

Second, the edge persistence mechanism relies on SD card-based storage, which may present endurance limitations during long-term high-frequency data logging. Although SD card wear was not observed during the experimental period, extended deployments may require alternative storage solutions or wear-leveling strategies. Future implementations may incorporate external flash storage or periodic data compaction to improve long-term reliability.

Third, while the proposed MQTT-ACK mechanism guarantees end-to-end data persistence, it introduces a bounded synchronization delay after prolonged network outages. Although this delay does not affect data integrity, it may impact real-time analytics immediately following reconnection. Future work will investigate adaptive synchronization strategies and prioritization mechanisms to minimize post-outage latency.

Fourth, the security considerations in this study were limited to transport-layer encryption, authentication, and basic acknowledgment validation. Advanced security mechanisms such as intrusion detection, replay attack mitigation, and fine-grained access control were not implemented. These aspects will be addressed in future studies, particularly for deployments in adversarial or public network environments.

Finally, the current study focuses on reliable data acquisition and persistence rather than advanced analytics. Future extensions will integrate machine learning-based fault diagnosis, anomaly detection, and predictive maintenance models using a reliable data stream enabled by the proposed architecture.

V. CONCLUSIONS

This study presents a scalable Message Queuing Telemetry Transport Acknowledgement (MQTT-ACK) edge Internet of Things (IoT) architecture for real-time distributed Photovoltaic (PV) monitoring, explicitly designed to guarantee end-to-end data reliability under intermittent network connectivity. Unlike conventional Supervisory Control and Data Acquisition (SCADA)-based and cloud-centric IoT monitoring solutions, the proposed system integrates disk-backed edge persistence with an application-level acknowledgment mechanism that issues delivery confirmation only after a successful cloud-side database insertion.

Experimental evaluation using real-world PV data demonstrated that the proposed architecture achieved complete data recovery during network disruptions with bounded synchronization delay upon reconnection. By coupling acknowledgment with durable database persistence and employing idempotent cloud-side processing, the system ensures that measurements are neither lost nor duplicated, even under retransmissions and temporary backend unavailability. These results confirm that reliability is enforced at the application layer rather than relying solely on broker-level Quality of Service (QoS) guarantees.

Resource utilization measurements further showed that the additional overhead introduced by local persistence, acknowledgment handling, and synchronization remained low in terms of CPU usage, memory consumption, and energy expenditure. This validates the suitability of the proposed approach for deployment in resource-constrained edge devices in large-scale distributed PV installations.

Compared to centralized SCADA platforms, the proposed architecture provides significantly higher temporal resolution and autonomous operation, enabling the detection of short-duration electrical and thermal anomalies that are typically obscured by inverter-level aggregation. In contrast to conventional IoT-based PV monitoring systems that assume continuous cloud availability, the proposed design explicitly decouples sensing reliability from network connectivity, making it well-suited for rooftop and geographically distributed PV deployments.

Finally, the results demonstrate that combining lightweight MQTT communication with disk-backed edge persistence and application-level ACK offers a practical, low-cost, and resilient monitoring solution for modern distributed solar PV systems. This architecture provides a reliable foundation for advanced analytics, fault detection, and predictive maintenance in real-world PV deployments.

DECLARATION OF COMPETING INTERESTS

The authors declare no conflicts of interest.

ACKNOWLEDGMENT

This work was supported under the research grant GRTIN-RAG-DCIS-04-2024, titled “Real-Time IoT-Based Solar Photovoltaic Panel Parameters Analysis and Fault Notification Hardware System Device Development.” The support is gratefully acknowledged. The authors would like to acknowledge Sunway University for providing research grant support for conducting and necessary financial support to publish this research work.

DATA AVAILABILITY

The dataset used in this study consists of real-world measurements collected from an experimental solar photovoltaic deployment. The data are available upon request from the corresponding author.

AI USE AND DECLARATION OF GENERATIVE AI USE

During the preparation of this manuscript, generative AI tools (ChatGPT) were used only for language refinement and editing. All experimental work, technical content, analyses, methodologies, and conclusions were conceived, developed, and validated entirely by the authors. The authors reviewed and edited all AI-generated suggestions and take full responsibility for the content of this publication.

REFERENCES

[1] S. Al-Ali, A. G. Olabi, and M. Mahmoud, "A review of solar photovoltaic technologies: developments, challenges, and future

perspectives," *Energy Conversion and Management: X*, vol. 27, July 2025, Art. no. 101057, https://doi.org/10.1016/j.ecmx.2025.101057.

[2] I. Lillo-Bravo, P. González-Martínez, M. Larrañeta, and J. Guasumba-Codena, "Impact of Energy Losses Due to Failures on Photovoltaic Plant Energy Balance," *Energies*, vol. 11, no. 2, Feb. 2018, Art. no. 363, https://doi.org/10.3390/en11020363.

[3] M. Bouzguenda, S. Chtourou, M. Alarfaj, R. M. Sumsudeen, and M. Shwehdi, "Arduino Uno Wi-Fi DeMilitarized Zone-based monitoring of solar photovoltaic systems," *Measurement and Control*, vol. 55, no. 3–4, pp. 136–145, Mar. 2022, https://doi.org/10.1177/00202940221090553.

[4] S. Ikemba, K. Song-hyun, T. O. Scott, D. R. E. Ewim, S. M. Abolarin, and A. A. Fawole, "Analysis of solar energy potentials of five selected south-east cities in nigeria using deep learning algorithms," *Sustainable Energy Research*, vol. 11, Jan. 2024, Art. no. 2, https://doi.org/10.1186/s40807-023-00096-7.

[5] S. Ansari, A. Ayob, M. S. H. Lipu, M. H. M. Saad, and A. Hussain, "A Review of Monitoring Technologies for Solar PV Systems Using Data Processing Modules and Transmission Protocols: Progress, Challenges and Prospects," *Sustainability*, vol. 13, no. 15, July 2021, Art. no. 8120, https://doi.org/10.3390/su13158120.

[6] M. Mahmood, P. Chowdhury, R. Yeassin, M. Hasan, T. Ahmad, and N.-U.-R. Chowdhury, "Impacts of digitalization on smart grids, renewable energy, and demand response: An updated review of current applications," *Energy Conversion and Management: X*, vol. 24, Oct. 2024, Art. no. 100790, https://doi.org/10.1016/j.ecmx.2024.100790.

[7] S. Haddad, R. Chaker, M. A. Moussa, L. Saci, and A. Khezzar, "A Scalable Real-Time Monitoring System for East-West PV Systems Leveraging the Internet of Things," in *2024 International Conference on Advances in Electrical and Communication Technologies (ICAECOT)*, Setif, Algeria, Oct. 2024, pp. 1–5, https://doi.org/10.1109/ICAECOT62402.2024.10828963.

[8] A. Gavrilov, M. Bergaliyev, S. Tinyakov, K. Krinkin, and P. Popov, "Using IoT Protocols in Real-Time Systems: Protocol Analysis and Evaluation of Data Transmission Characteristics," *Journal of Computer Networks and Communications*, vol. 2022, Aug. 2022, Art. no. 7368691, https://doi.org/10.1155/2022/7368691.

[9] K. Kull, B. Asad, M. A. Khan, M. U. Naseer, A. Kallaste, and T. Vaimann, "Faults, Failures, Reliability, and Predictive Maintenance of Grid-Connected Solar Systems: A Comprehensive Review," *Applied Sciences*, vol. 15, no. 21, Oct. 2025, Art. no. 11461, https://doi.org/10.3390/app152111461.

[10] P.-H. Pham Le, Q. N. Do, T. Q. Dinh, H.-T.-N. Pham, and L. V. Nguyen, "A comparative security analysis of MQTT brokers against DoS attacks," *Journal on Information Security*, vol. 2026, Jan. 2026, Art. no. 5, https://doi.org/10.1186/s13635-026-00224-y.

[11] B. Aydin, H. Aydin, and S. Gormus, "Intrusion detection systems in IoT: A detailed review of threat categories, detection strategies, and future technologies," *Journal of Information Security and Applications*, vol. 95, Dec. 2025, Art. no. 104291, https://doi.org/10.1016/j.jisa.2025.104291.

[12] S. Kaganurmath, N. G. Cholli, and M. R. Anala, "DLKS-MQTT: A Lightweight Key Sharing Protocol for Secure IoT Communications," *Engineering, Technology & Applied Science Research*, vol. 15, no. 2, pp. 21532–21538, Apr. 2025, https://doi.org/10.48084/etasr.10216.

[13] Y. Im and M. Lim, "E-MQTT: End-to-End Synchronous and Asynchronous Communication Mechanisms in MQTT Protocol," *Applied Sciences*, vol. 13, no. 22, Nov. 2023, Art. no. 12419, https://doi.org/10.3390/app132212419.