

IoT-Based Biomedical Architecture with Real-Time Streaming Analytics for Preventive Health Monitoring

Isaac Zablah¹; Edwin Hernandez²; Fiamma Garcia³; Antonieta Zuniga⁴; Antonio Garcia Loureiro⁵

^{1,3}Faculty of Medical Sciences, National Autonomous University of Honduras, Honduras,

jose.zablah@unah.edu.hn, fiamma.garcia@unah.edu.hn

²EGLA Corp., USA, edwinhm@eglacorp.com

⁴Dirección de Medicina Forense, Ministerio Público, Honduras, antonietazuniga2311@yahoo.com

⁵Department of Electronics and Computing, University Santiago de Compostela, Spain, antonio.garcia.loureiro@usc.es

Abstract— *Traditional health monitoring systems rely on periodic sampling, which introduces significant delays in detecting critical physiological events. This paper presents an IoT-based biomedical architecture integrating real-time streaming analytics to enable early detection of adverse health conditions. The proposed system combines wearable biomedical sensors, MQTT communication protocols, and Apache Kafka-based streaming pipelines for continuous physiological data processing. We evaluated the architecture through controlled simulations comparing streaming analytics against conventional periodic sampling approaches as a theoretical approach. Results demonstrate a substantial reduction in event detection latency, with the streaming system achieving mean detection delays of 47.0 seconds compared to 300.0 seconds for periodic sampling ($p < 0.001$). The MQTT-based communication layer exhibited mean latency of 30.2 ms with 95th percentile at 71.2 ms, significantly outperforming HTTP alternatives (mean: 149.9 ms, 95th percentile: 317.4 ms). System scalability testing revealed linear throughput scaling, supporting up to 100 concurrent sensors at 88,000 messages per second. These findings validate the efficacy of streaming analytics in biomedical IoT systems for preventive healthcare applications, with relevance for continuous monitoring of cardiac, respiratory, and metabolic parameters.*

Keywords— *Internet of Things, Biomedical Sensors, Streaming Analytics, Real-Time Processing, Health Monitoring.*

I. INTRODUCTION

The global burden of cardiovascular diseases, diabetes, and respiratory disorders continues to escalate, with the World Health Organization reporting over 17.9 million deaths annually from cardiovascular conditions alone [1]. Traditional healthcare models rely predominantly on episodic patient encounters, limiting continuous monitoring capabilities and often resulting in delayed intervention during critical events [2]. This reactive approach contrasts sharply with the potential of preventive medicine enabled by continuous physiological monitoring.

This paradigm has generated a major interest in technology-driven healthcare systems that prioritize early detection, risk assessment, and rapid clinical decision-making. Wearable sensors, mobile health platforms, and integrated medical devices have all advanced. They now provide the ongoing acquisition of physiological signals in real-world environments, providing a more comprehensive and prolonged

perspective on a patient's health [3]. Such strategies can transform care delivery from hospital-centric, event-driven systems to proactive, individualized interventions, particularly in the management of chronic diseases. Continuous monitoring systems provide the observation and analysis of critical parameters in near real-time. This facilitates the early detection of pathological aberrations, mitigates the risk of acute decompensation, enhances therapeutic results, and alleviates the burden on already overextended healthcare systems [4].

The convergence of Internet of Things (IoT) technologies with biomedical sensing has created unprecedented opportunities for continuous health monitoring [5]. Modern wearable devices equipped with multi-parameter sensors can capture heart rate, oxygen saturation, body temperature, and electrocardiographic signals with clinical-grade accuracy [6]. However, the architectural frameworks supporting these devices frequently employ batch processing or periodic data sampling, introducing latency that may compromise timely detection of acute medical events [7].

Streaming analytics represents a paradigm shift in data processing, enabling real-time analysis of continuous data streams with sub-second latency [8]. Technologies such as Apache Kafka, Apache Flink, and Apache Spark Streaming have demonstrated remarkable performance in financial trading, network monitoring, and industrial IoT applications [9]. Yet, their application in biomedical contexts remains relatively underexplored, particularly regarding architecture design patterns, performance optimization, and clinical validation.

Communication protocols constitute a critical architectural component in IoT systems. While HTTP remains ubiquitous for web services, its request-response model and connection overhead introduce latency unsuitable for real-time biomedical applications [10]. Message Queue Telemetry Transport (MQTT) has emerged as a lightweight publish-subscribe protocol optimized for constrained devices and unreliable networks, demonstrating superior performance in IoT deployments [11]. However, comparative evaluations of MQTT versus HTTP in biomedical streaming contexts remain limited [12].

This research addresses these gaps by presenting comprehensive IoT-based biomedical architecture integrating streaming analytics for preventive health monitoring. We hypothesize that real-time streaming analytics enable significantly earlier detection of critical physiological events compared to conventional periodic sampling approaches. The architecture employs MQTT for sensor communication, Apache Kafka for stream processing, and machine learning algorithms for anomaly detection. We validate the system through controlled simulations with realistic physiological data, quantifying performance improvements in detection latency, communication overhead, and system scalability.

II. MATERIALS AND METHODS

A. System Architecture

The proposed architecture comprises four primary layers: sensor layer, communication layer, processing layer, and application layer. The sensor layer integrates commercial wearable devices (simulated) capable of measuring heart rate (HR), peripheral oxygen saturation (SpO_2), and body temperature at 1 Hz sampling frequency. These parameters were selected based on their clinical relevance in detecting cardiac arrhythmias, respiratory distress, and metabolic dysregulation [13].

The communication layer implements MQTT protocol with Quality of Service (QoS) level 1, ensuring at-least-once delivery while minimizing overhead. An Eclipse Mosquitto broker manages publish-subscribe operations, with sensors publishing to topic hierarchies organized by patient identifier and physiological parameter. This design enables selective subscription and horizontal scaling of processing nodes [14].

The processing layer leverages Apache Kafka for stream ingestion and Apache Spark Structured Streaming for real-time analytics. Data flows through three processing stages: ingestion, transformation, and anomaly detection. The transformation stage applies moving average filters (window size: 60 seconds) to reduce noise while preserving signal characteristics. Anomaly detection employs threshold-based algorithms validated against clinical guidelines: bradycardia ($\text{HR} < 50$ bpm), tachycardia ($\text{HR} > 120$ bpm), hypoxemia ($\text{SpO}_2 < 90\%$), and fever (temperature $> 38.0^\circ\text{C}$) [15].

B. Simulation Framework

To enable controlled, reproducible evaluation, we developed a comprehensive simulation framework generating synthetic physiological data with realistic temporal dynamics. The simulator models circadian variations, stochastic noise, and clinically relevant pathological events. Heart rate synthesis incorporates baseline circadian rhythm (70 ± 10 bpm) with superimposed Gaussian noise ($\sigma = 2$ bpm) and programmed critical events including bradycardia episodes (45 ± 2 bpm, duration: 5 minutes) and tachycardia episodes (135 ± 3 bpm, duration: 7 minutes).

SpO_2 simulation maintained normal values ($97.5 \pm 0.5\%$) with episodic desaturation events ($88 \pm 1\%$, duration: 6

minutes) representing hypoxemia. Temperature data incorporated circadian variation (0.3°C amplitude) and fever episodes ($38.5 \pm 0.15^\circ\text{C}$, duration: 2 hours). All simulations extended 24 hours with 1 Hz sampling rate, generating 86,400 data points per parameter. This approach ensures statistical robustness while avoiding ethical concerns associated with inducing pathological states in human subjects.

C. Performance Metrics

System performance evaluation employed four primary metrics: detection latency, communication latency, system throughput, and alert precision. Detection latency quantifies the time interval between physiological event onset and algorithmic detection, measured in seconds. Communication latency measures end-to-end message transmission time from sensor to processing node, measured in milliseconds. System throughput evaluates maximum sustainable message processing rate, expressed as messages per second. Alert precision calculates the proportion of generated alerts corresponding to true clinical events, computed as $\text{TP}/(\text{TP}+\text{FP})$ where TP denotes true positives and FP denotes false positives.

We compared the proposed streaming analytics system against a baseline periodic sampling architecture representative of conventional telemetry systems. The periodic system samples at 5-minute intervals, aggregates data into batches, and performs analysis during each cycle. This configuration reflects common practice in clinical telemetry units and commercial health monitoring platforms [16].

D. Statistical Analysis

Python 3.9 was employed to conduct statistical analyses, utilizing the NumPy, SciPy, and Pandas libraries for numerical computations, statistical evaluations, and data manipulation, respectively. Before performing inferential analysis, the distributions of the data were examined to evaluate any possible deviations from normality. Due to the non-normal distributions of detection latency measures, the Wilcoxon rank-sum test, a non-parametric method, was employed to compare the groups. This assessment does not presuppose a Gaussian distribution and is applicable to independent samples [17], [18].

A detailed collection of descriptive statistics, encompassing the arithmetic mean, median, standard deviation, and 95th percentile, was utilized to examine communication latency, clearly demonstrating both central tendency and tail behavior pertinent to performance-sensitive applications. Linear regression analysis was employed to evaluate throughput scalability. This allows for an examination of how performance varies with increasing system load, facilitating the identification of potential bottlenecks or saturation effects. All statistical tests were conducted using a two-tailed approach, with the predetermined threshold for significance established at $\alpha = 0.05$. All simulation code, data pretreatment protocols, and

analytic scripts have been made publicly accessible in the supplemental materials to enhance transparency and ensure repeatability.

III. RESULTS

A. Detection Performance

The streaming analytics architecture demonstrated substantial improvements in event detection latency compared to periodic sampling. Table I summarizes detection performance metrics across three simulated critical events. The streaming system achieved mean detection latency of 47.0 seconds, while periodic sampling exhibited mean latency of 300.0 seconds, representing an 84.3% reduction in detection time.

TABLE I
DETECTION PERFORMANCE COMPARISON

Metric	Streaming Analytics	Periodic Sampling
Mean Detection Latency (s)	47.0	300.0
Median Detection Latency (s)	47.0	300.0
Events Detected	3/3 (100%)	3/3 (100%)
Latency Reduction	84.3% (253.0 seconds)	

Figure 1 illustrates box plots that compare the distributions of detection latency between the two topologies. The streaming system shows negligible variation due to continuous real-time processing, whereas periodic sampling reveals intrinsic variability influenced by the timing of events in relation to sampling intervals. The Wilcoxon rank-sum test demonstrated statistically significant differences ($p < 0.001$), robustly validating our claim concerning the superiority of streaming analytics in event detection.




Fig. 1 Detection latency comparison between streaming analytics (proposed) and periodic sampling (traditional) approaches. Box plots display median (center line), interquartile range (box), and overall range (whiskers). Red diamonds indicate mean values.

B. Communication Performance

Communication layer evaluation revealed substantial differences between MQTT and HTTP protocols in terms of message latency and reliability. MQTT exhibited mean end-to-end latency of 30.1 ms (SD = 13.2 ms) with 95th percentile at 56.8 ms. In contrast, HTTP demonstrated mean latency of 149.0 ms (SD = 75.2 ms) with 95th percentile at 289.5 ms, representing a 79.8% increase in average latency compared to MQTT.

Figure 2 depicts latency distributions for both protocols using histograms. The MQTT distribution demonstrates significant positive skew with a concentration around the lower limit, signifying reliable low-latency performance. The distribution of HTTP latency exhibits increased variability and pronounced tails, indicating a heightened vulnerability to network congestion and overhead in connection formation. These attributes render MQTT far more appropriate for time-sensitive biological applications necessitating reliable, low-latency communication.

C. System Scalability

Scalability testing evaluated system throughput under varying sensor loads ranging from 1 to 100 concurrent devices. Table II presents detailed throughput measurements for both MQTT-based and HTTP-based architectures across six load conditions. The MQTT architecture demonstrated near-linear scaling, achieving 950 messages per second with a single sensor and 88,000 messages per second with 100 sensors (scaling factor: 92.6x). HTTP architecture exhibited sublinear scaling, reaching 450 messages per second at single-sensor load and 28,000 messages per second at 100-sensor load (scaling factor: 62.2x).

The throughput differential between MQTT and HTTP widened progressively with increasing sensor load. At single-sensor load, MQTT processed approximately $2.1\times$ more messages per second than HTTP; this ratio increased to $3.1\times$ at the 100-sensor condition. This widening gap reflects the cumulative advantage of MQTT's reduced per-connection overhead and persistent connection model under concurrent load. At maximum tested load, the system sustained 88,000 messages per second across 100 sensors, corresponding to 880 readings per sensor per second—a throughput substantially exceeding the 1–10 Hz sampling rates typical of vital sign monitoring applications. The linear scaling profile further suggests that the architecture could accommodate substantially larger sensor networks through horizontal infrastructure scaling, without requiring fundamental architectural modifications. This characteristic makes the design suitable for institutional deployments requiring simultaneous monitoring of large patient populations.

Figure 3 visualizes throughput scaling characteristics through line plots comparing MQTT and HTTP architecture. Linear regression analysis of MQTT throughput yielded $R^2 = 0.9987$, indicating excellent linear fit. HTTP architecture demonstrated $R^2 = 0.9912$, suggesting slight deviation from linearity at higher loads. The superior scalability of MQTT-

based architecture stems from reduced connection overhead, persistent connections, and efficient binary message encoding compared to HTTP's text-based, stateless protocol design.

TABLE II
SYSTEM THROUGHPUT UNDER VARIABLE SENSOR LOADS

Concurrent Sensors	MQTT Throughput (msg/s)	HTTP Throughput (msg/s)
1	950	450
5	4,800	2,100
10	9,500	4,000
20	18,500	7,500
50	45,000	16,000
100	88,000	28,000




Fig. 2 Communication latency distributions for MQTT (left) and HTTP (right) protocols. Histograms display frequency distributions with overlay showing mean (red dashed) and 95th percentile (orange dashed) values. $n = 10,000$ message transmissions per protocol.

IV. DISCUSSION

This research demonstrates that IoT-based architectures incorporating streaming analytics can substantially reduce detection latency for critical physiological events compared to conventional periodic sampling systems. The 84.3% reduction in mean detection time achieved by our streaming approach translates to clinically meaningful improvements in intervention timeliness. For cardiac events such as ventricular

tachycardia or severe bradycardia, every minute matters; earlier detection enables prompt medical intervention, potentially preventing progression to cardiac arrest or hemodynamic collapse [19].




Fig. 3 System throughput comparison between MQTT-based (green) and HTTP-based (red) architectures under increasing sensor loads. Lines connect measured throughput values across six load conditions (1-100 concurrent sensors). Error markers indicate measurement endpoints.

The superior performance of MQTT over HTTP aligns with theoretical expectations based on protocol design characteristics. MQTT's publish-subscribe model eliminates polling overhead and enables event-driven architecture, while persistent TCP connections reduce handshake latency. HTTP's request-response paradigm requires explicit polling or complex workarounds like long polling or server-sent events, introducing both latency and implementation complexity [20]. Our measured 79.8% latency reduction with MQTT validates its suitability for time-critical biomedical applications.

System scalability results indicate the proposed architecture can support large-scale deployments without performance degradation. The near-linear scaling observed with MQTT ($R^2 = 0.9987$) suggests the architecture could accommodate hundreds or thousands of concurrent sensors with appropriate infrastructure provisioning. This scalability characteristic is essential for institutional deployments spanning hospitals, nursing homes, or population-health programs monitoring multiple patients simultaneously.

Several important considerations warrant attention regarding clinical translation. First, the anomaly detection component employed clinically validated, fixed-threshold algorithms. This choice was deliberate: threshold-based rules offer interpretability, direct grounding in established clinical guidelines [13], and deterministic behavior—properties that are valuable for initial validation of a streaming architecture. However, fixed thresholds have well-documented limitations in real-world monitoring: they do not account for inter-individual physiological variability, circadian patterns, or medication-induced baseline shifts, and are susceptible to elevated false-positive rates in the presence of motion artifacts or sensor noise [21]. Machine learning approaches offer a principled path to overcoming these limitations. Recurrent neural networks, particularly LSTM-based architectures, have

demonstrated capacity to model temporal dependencies in physiological time series and detect subtle deviations from patient-specific baselines [22]. Autoencoders trained on normal physiological patterns can flag anomalies as high reconstruction error without requiring labeled pathological events, which is advantageous given the scarcity of annotated clinical datasets [23]. Isolation Forest and One-Class SVM methods offer computationally efficient alternatives suitable for edge deployment. Integration of such models within the existing Kafka-Spark pipeline is architecturally feasible, as Spark MLlib and TensorFlow Serving are compatible with the structured streaming framework used here. The primary challenge lies not in architectural integration but in obtaining sufficient real patient data for model training and validation—reinforcing the urgency of the clinical validation phase outlined above. Incorporating adaptive, ML-based detection represents the highest-priority enhancement for future iterations of this architecture.

Second, although data security and patient privacy were not implemented or evaluated in the current simulation-based study, they constitute non-negotiable requirements for any clinical deployment of this architecture. The present work intentionally focused on characterizing architectural performance under controlled conditions; security mechanisms were excluded to isolate latency and throughput as measurement targets without confounding overhead. In a production deployment, the architecture would require end-to-end TLS encryption for MQTT transmissions, certificate-based device authentication at the broker level, role-based access control for the Kafka topic hierarchy, and data anonymization pipelines compliant with applicable regulatory frameworks—including HIPAA in the United States, GDPR in the European Union, and equivalent instruments in other jurisdictions [24]. The streaming pipeline introduces specific security considerations not present in batch systems, including the continuous nature of data exposure and the challenge of enforcing access control at high ingestion rates. Quantifying the latency overhead introduced by TLS and authentication mechanisms under realistic sensor loads is an important open question that should be addressed in the next evaluation phase.

Third, energy efficiency represents a practical constraint that warrants careful consideration in wearable IoT deployments. Although our simulation focused exclusively on architectural latency and throughput, the communication protocol selected at the sensor node can substantially affect power consumption. MQTT is generally associated with lower communication overhead than HTTP because it uses a lightweight messaging structure and more efficient exchange patterns for constrained IoT environments. Experimental evidence from embedded IoT platforms has shown lower average power consumption and longer battery life for MQTT than for HTTP under real-time monitoring conditions [25]. For a wearable device transmitting vital signs at 1 Hz during continuous operation, these per-message differences can accumulate over time and materially affect battery longevity, which is a critical factor in ambulatory monitoring scenarios.

Future work should therefore include direct measurement of current draw and energy per transmitted message for both protocols on representative hardware platforms, such as ESP32 or Nordic nRF52 devices, together with evaluation of duty-cycling strategies, adaptive sampling rates, and energy-harvesting mechanisms to improve operational autonomy [26].

Fourth, system reliability and fault tolerance are critical requirements for any architecture intended for clinical use, and merit explicit discussion even though they were not empirically evaluated in this study. The proposed architecture incorporates several mechanisms that inherently support fault resilience. MQTT's Quality of Service level 1 (QoS 1) guarantees at-least-once message delivery, protecting against packet loss on unreliable networks. Apache Kafka's distributed log architecture provides built-in data replication across broker nodes, ensuring message persistence and availability even during individual broker failures. Apache Spark Structured Streaming supports checkpointing, which allows pipeline recovery from processing failures without data loss. However, our simulation did not subject these mechanisms to deliberate fault injection scenarios—such as broker crashes, network partitions, or sensor disconnection events—and therefore cannot quantify recovery times or data loss rates under failure conditions. Single points of failure remain in the current design, including the Eclipse Mosquitto broker and the Spark driver node. Future iterations should address these through broker clustering, multi-region Kafka replication, and watchdog services. Empirical fault-injection testing, consistent with prior IoT dependability studies that deliberately inject faults into message brokers to exercise self-healing and fault-tolerance mechanisms, is recommended as part of a rigorous pre-deployment validation protocol [27].

The clinical utility of continuous monitoring extends beyond acute event detection to longitudinal trend analysis, medication titration monitoring, and early detection of gradual physiological deterioration. Streaming analytics platforms can readily accommodate these applications by maintaining historical context and computing temporal features such as heart rate variability, SpO₂ trend slopes, or temperature patterns. Integration with electronic health records (EHR) enables correlation of physiological data with clinical notes, laboratory results, and medication administration records, supporting comprehensive clinical decision-making [28].

Several limitations of this study should be acknowledged. First and most notably, the evaluation relies entirely on synthetic physiological data generated through a controlled simulation framework rather than recordings from real patients. This methodological choice was deliberate: simulation enables precise control over event timing, signal characteristics, and pathological patterns, ensuring reproducible and unconfounded performance benchmarking at the architectural design stage. This approach is consistent with established practice in biomedical engineering, where simulation-based evaluation precedes clinical deployment to identify architectural bottlenecks and optimize system parameters before exposing patients to experimental protocols

[29], [30]. Nonetheless, real-world physiological signals exhibit inter-individual variability, motion artifacts, sensor drift, and clinical workflow complexities that synthetic data cannot fully replicate. Prospective validation using real sensor data from consenting participants, conducted under institutional review board approval, is therefore an essential next step before clinical translation. Second, the current evaluation was conducted in a single simulated environment without geographic distribution, network heterogeneity, or multi-institutional infrastructure variability. Third, no formal usability assessment was performed with healthcare providers, which is necessary to confirm clinical acceptability and workflow integration.

Future research directions include integration of advanced machine learning models for personalized risk prediction, expansion to additional physiological parameters such as electrocardiography and continuous glucose monitoring, development of edge computing capabilities for on-device analytics, and comprehensive clinical trials evaluating patient outcomes, healthcare utilization, and cost-effectiveness. The foundational architecture presented here provides a robust platform for these extensions.

V. CONCLUSION

This paper presents a comprehensive IoT-based biomedical architecture that integrates streaming analytics for continuous health monitoring, substantially reducing event detection latency. The proposed strategy reduces detection delay by 84.3% relative to conventional periodic sampling methods. MQTT exhibited a 79.8% reduction in mean latency relative to HTTP. The assessment of system scalability indicates that throughput increases roughly linearly, suggesting its applicability across various healthcare settings, from individual patients to extensive populations.

This research offers technological advancements by presenting architectural design patterns for biomedical streaming systems, assessing performance across multiple dimensions, and corroborating the findings through controlled simulations with authentic physiological dynamics. These findings provide evidence-based guidance for the implementation of real-time health monitoring systems, particularly relevant to preventive care, chronic disease management, and post-discharge monitoring.

Numerous technical challenges remain to be addressed regarding security, privacy, power management, and clinical validation. Nonetheless, the enhancements in performance indicate that streaming analytics is an effective and beneficial approach for next-generation biomedical IoT systems. As wearable sensor technology advances and healthcare systems increasingly prioritize preventative treatment, frameworks facilitating real-time physiological monitoring will become integral to personalized, proactive therapy.

Future initiatives should prioritize upcoming clinical trials that validate safety and efficacy in real-world settings, the integration of advanced machine learning algorithms for

personalized risk assessment, and the development of strong security measures to protect patient information. The integration of IoT technology, streaming analytics, and artificial intelligence has the capacity to significantly transform healthcare delivery, shifting the emphasis from reactive treatment to predictive prevention.

ACKNOWLEDGMENT

This research was supported by the Institute for Research in Medical Sciences and Right to Health (ICIMEDES), the Faculty of Medical Sciences (FCM), and the Directorate of Scientific, Humanistic, and Technological Research (DICIHT) of the National Autonomous University of Honduras (UNAH).

REFERENCES

- [1] World Health Organization, "Cardiovascular diseases (CVDs)," WHO Fact Sheet, Jul. 31, 2025. [Online]. Available: [https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-\(cvds\)](https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds)).
- [2] P. A. Heidenreich et al., "Forecasting the impact of heart failure in the United States: A policy statement from the American Heart Association", *Circ Heart Fail*, vol. 6, no. 3, pp. 606-619, 2013. doi: 10.1161/HHF.0b013e318291329a
- [3] E. J. Topol, "High-performance medicine: the convergence of human and artificial intelligence", *Nat Med*, vol. 25, no. 1, pp. 44–56, Jan. 2019, doi: 10.1038/s41591-018-0300-7.
- [4] S. M. Steinhubl, E. D. Muse, and E. J. Topol, "Can mobile health technologies transform health care?", *JAMA*, vol. 310, no. 22, pp. 2395–2396, Dec. 2013, doi: 10.1001/jama.2013.281078
- [5] S. M. R. Islam, D. Kwak, M. H. Kabir, M. Hossain, and K. S. Kwak, "The Internet of Things for health care: A comprehensive survey", *IEEE Access*, vol. 3, pp. 678-708, 2015. doi: 10.1109/ACCESS.2015.2437951
- [6] A. Pantelopoulos and N. G. Bourbakis, "A survey on wearable sensor-based systems for health monitoring and prognosis", *IEEE Trans Syst Man Cybern C Appl Rev*, vol. 40, no. 1, pp. 1-12, Jan. 2010. DOI: 10.1109/TSMCC.2009.2032660
- [7] M. A. Habib, M. Mohktar, S. B. Kamaruzzaman, K. S. Lim, T. M. Pin, and F. Ibrahim, "Smartphone-based solutions for fall detection and prevention: Challenges and open issues", *Sensors*, vol. 14, no. 4, pp. 7181-7208, 2014. DOI: 10.3390/s140407181
- [8] T. Akidau et al., "The dataflow model: A practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing", *Proc VLDB Endow*, vol. 8, no. 12, pp. 1792-1803, 2015. doi: 10.14778/2824032.2824076
- [9] J. Kreps, N. Narkhede, and J. Rao, "Kafka: A distributed messaging system for log processing", in *Proc 6th Int Workshop Netw Meets Databases (NetDB)*, Athens, Greece, 2011, pp. 1-7.
- [10] R. Fielding et al., "Hypertext transfer protocol -- HTTP/1.1", *RFC 2616, Internet Engineering Task Force*, 1999. doi: 10.17487/RFC2616
- [11] U. Hunkeler, H. L. Truong, and A. Stanford-Clark, "MQTT-S—A publish/subscribe protocol for wireless sensor networks", in *Proc 3rd Int Conf Commun Syst Softw Middleware Workshops (COMSWARE)*, Bangalore, India, 2008, pp. 791-798. doi: 10.1109/COMSWA.2008.4554519
- [12] V. J. Aski, V. S. Dhaka, A. Parashar, S. Kumar, and I. Rida, "Internet of Things in healthcare: A survey on protocol standards, enabling technologies, WBAN architectures and open issues", *Phys. Commun.*, vol. 60, Art. no. 102103, 2023, doi: 10.1016/j.phycom.2023.102103.
- [13] A. J. Camm et al., "Guidelines for the management of atrial fibrillation: The Task Force for the Management of Atrial Fibrillation of the European Society of Cardiology (ESC)", *Eur Heart J*, vol. 31, no. 19, pp. 2369-2429, 2010. doi: 10.1093/eurheartj/ehq278
- [14] S. Khan, E. Maimaitijiang, I. Kures, and Y. Mamatjan, "A Novel Modular Framework for Secure and Scalable Remote Health Monitoring: RHMS", *Appl. Sci.*, vol. 15, no. 23, Art. no. 12623, 2025, doi: 10.3390/app152312623

- [15]P. Kligfield et al., “Recommendations for the standardization and interpretation of the electrocardiogram: Part I: The electrocardiogram and its technology”, *Circulation*, vol. 115, no. 10, pp. 1306–1324, 2007, doi: 10.1161/CIRCULATIONAHA.106.180200
- [16]L. Tarassenko, D. A. Clifton, M. R. Pinsky, M. T. Hravnak, J. R. Woods, and P. J. Watkinson, “Centile-based early warning scores derived from statistical distributions of vital signs”, *Resuscitation*, vol. 82, no. 8, pp. 1013–1018, 2011, doi: 10.1016/j.resuscitation.2011.03.006
- [17]C. R. Harris, K. J. Millman, S. J. van der Walt et al., “Array programming with NumPy”, *Nature*, vol. 585, pp. 357–362, Sept. 2020, doi: 10.1038/s41586-020-2649-2
- [18]F. Wilcoxon, “Individual comparisons by ranking methods”, *Biometrics Bull.*, vol. 1, pp. 80–83, Dec. 1945, doi: 10.2307/3001968
- [19]P. Kudenchuk et al., “Part 10: Electrical therapies: Automated external defibrillators, defibrillation, cardioversion, and pacing”, *Circulation*, vol. 132, no. 18 suppl 2, pp. S447–S463, 2015. doi: 10.1161/CIR.0000000000000262
- [20]C. Bormann, A. P. Castellani, and Z. Shelby, “CoAP: An application protocol for billions of tiny internet nodes”, *IEEE Internet Comput.*, vol. 16, no. 2, pp. 62–67, 2012. doi: 10.1109/MIC.2012.29
- [21]J. Allen, “Photoplethysmography and its application in clinical physiological measurement”, *Physiol. Meas.*, vol. 28, no. 3, pp. R1–R39, 2007. doi: 10.1088/0967-3334/28/3/R01
- [22]S. Boda, M. Mahadevappa, and P. K. Dutta, “An automated patient-specific ECG beat classification using LSTM-based recurrent neural networks”, *Biomed. Signal Process. Control*, vol. 84, Art. no. 104756, Jul. 2023, doi: 10.1016/j.bspc.2023.104756
- [23]M. A. Rassam, “Autoencoder-Based Neural Network Model for Anomaly Detection in Wireless Body Area Networks”, *IoT*, vol. 5, no. 4, pp. 852–870, 2024, doi: 10.3390/iot5040039
- [24]K. N. Vokinger, D. J. Stekhoven, and M. Krauthammer, “Lost in anonymization—A data anonymization reference classification merging legal and technical considerations”, *J. Law Med. Ethics*, vol. 48, no. 1, pp. 228–231, 2020, doi: 10.1177/1073110520917025
- [25]H. J. Jara Ochoa, R. Peña, Y. Ledo Mezquita, E. Gonzalez, and S. Camacho-Leon, “Comparative Analysis of Power Consumption between MQTT and HTTP Protocols in an IoT Platform Designed and Implemented for Remote Real-Time Monitoring of Long-Term Cold Chain Transport Operations”, *Sensors*, vol. 23, no. 10, Art. no. 4896, 2023, doi: 10.3390/s23104896.
- [26]S. Ullah, H. Higgins, B. Braem, B. Latré, C. Blondia, I. Moerman, S. Saleem, Z. Rahman, and K. S. Kwak, “A comprehensive survey of wireless body area networks on PHY, MAC, and network layers solutions”, *J. Med. Syst.*, vol. 36, no. 3, pp. 1065–1094, 2012, doi: 10.1007/s10916-010-9571-3
- [27]M. Duarte, J. P. Dias, H. S. Ferreira, and A. Restivo, “Evaluation of IoT Self-healing Mechanisms using Fault-Injection in Message Brokers”, in *Proc. 4th Int. Workshop on Software Engineering Research and Practices for the Internet of Things (SERP4IoT '22)*, 2022, pp. 9–16, doi: 10.1145/3528227.3528567
- [28]M. Nawaz and J. Ahmed, “Cloud-based healthcare framework for real-time anomaly detection and classification of 1-D ECG signals”, *PLOS ONE*, vol. 17, no. 12, e0279305, 2022, doi: 10.1371/journal.pone.0279305
- [29]S. Meliá, S. Nasabeh, S. Luján-Mora, and C. Cachero, “MoSIoT: Modeling and Simulating IoT Healthcare-Monitoring Systems for People with Disabilities,” *Int. J. Environ. Res. Public Health*, vol. 18, no. 12, Art. no. 6357, 2021, doi: 10.3390/ijerph18126357
- [30]O. Said and A. Tolba, “Design and Evaluation of Large-Scale IoT-Enabled Healthcare Architecture,” *Appl. Sci.*, vol. 11, no. 8, Art. no. 3623, 2021, doi: 10.3390/app11083623.