IoT-Based Remote Health Monitoring and Alert System for Tribal Communities: Architecture, Optimization Modelling, and Statistical Evaluation

Dr. Martin Joel Rathnam[1*], Dr. K Satyanarayan Reddy[2]

[1*]*Research Scholar, Srinivas University, Mukka, Surathkal, Mangaluru, Karnataka, India.*

[2]*Vice Chancellor, Mukka, Surathkal, Srinivas University, Mangaluru, Karnataka, India.*

Abstract: *Access to quality healthcare remains a critical challenge for tribal communities, arising from geographic isolation, inadequate infrastructure, and socioeconomic barriers. This paper presents an IoT-based Remote Health Monitoring and Alert System (IoT-RHMAS) specifically designed for resource-constrained tribal environments, incorporating two original analytical contributions absent from prior work. First, a Communication Protocol Optimization Model (CPOM) is formulated as a constrained multi-objective minimization problem that dynamically selects the optimal wireless protocol (BLE, LoRa, or GSM) by minimizing a weighted composite objective function over energy consumption, latency, and channel reliability. Second, a Weighted Adaptive Management Model (WAMM) optimizes duty-cycle scheduling through dynamic programming to maximize operational uptime within the solar energy budget. The anomaly detection subsystem employs a statistically validated Isolation Forest framework with empirical hyperparameter selection via cross-validated grid search, achieving a precision of 94.8%, recall of 92.3%, and F1-score of 93.5% — superior to both rule-based and standalone ML baselines. Bland-Altman agreement analysis confirms sensor accuracy within clinically acceptable limits (mean bias < 0.6 bpm). Alert latency analysis demonstrates that 95% of BLE alerts are delivered within the 3-second clinical threshold. These validated analytical models and empirical results substantiate IoT-RHMAS as a rigorously grounded, cost-effective (USD 28-35/node), and inclusive healthcare platform for underserved indigenous populations.*

Keywords: IoT; Tribal Healthcare; Remote Health Monitoring; Isolation Forest; Multi-Objective Optimization; Communication Protocol Selection; Duty-Cycle Scheduling; Statistical Validation; Wearable Sensors; mHealth.

1. INTRODUCTION

Access to reliable healthcare remains one of the most persistent challenges confronting tribal and remote communities globally. Geographic isolation, limited infrastructure, and socioeconomic barriers create a compounding disadvantage resulting in delayed diagnoses, chronic disease progression, and preventable mortality. In India alone, tribal populations constitute approximately 8.6% of the national population (Census, 2011), yet remain profoundly underserved by the formal healthcare system.

The Internet of Things (IoT) has emerged as a transformative paradigm with significant potential to redress such inequities through continuous health monitoring and remote diagnostic capabilities. IoT-enabled health systems leverage sensor networks and communication infrastructure to collect, process, and transmit physiological data—including heart rate, blood oxygen saturation (SpO_2), and body temperature—enabling proactive healthcare management where conventional access is limited [1], [2].

Critically, however, the preponderance of existing IoT health monitoring systems lacks formal analytical models governing key operational decisions—specifically, which wireless communication protocol to deploy given varying channel conditions, and how to schedule sensor sampling given a finite solar energy budget. These omissions are not merely academic: in tribal environments where connectivity is intermittent and energy is scarce, ad hoc protocol and duty-cycle decisions directly degrade system reliability and patient safety. This paper addresses both gaps through original optimization and statistical contributions.

The principal contributions of this work are as follows:

(i) A modular, five-tier IoT architecture for continuous vital-sign monitoring in low-connectivity, resource-constrained tribal environments.

(ii) A Communication Protocol Optimization Model (CPOM): a novel constrained multi-objective optimisation formulation that minimises a weighted composite cost function $J(p) = w_1E(p) + w_2L(p) + w_3[1-R(RSSI, p)]$ over energy, latency, and channel reliability, yielding provably optimal protocol selection under dynamic connectivity conditions.

(iii) A Weighted Adaptive Management Model (WAMM): an original dynamic programming formulation for duty-cycle scheduling that maximises 24-hour operational uptime subject to a solar energy budget constraint, achieving a 75% reduction in average power consumption relative to always-active baseline operation.

(iv) A statistically validated hybrid anomaly detection framework combining clinically grounded rule-based thresholds with a cross-validated Isolation Forest model (F1 = 93.5%), supported by Bland-Altman agreement analysis, precision-recall characterisation, and confusion matrix evaluation.

(v) Empirical characterisation of alert latency distributions confirming that ≥95% of BLE alerts satisfy the 3-second clinical threshold, and multi-user field performance across power consumption, communication range, and system uptime metrics.

2. Literature Survey

2.1 IoT and Wearable Systems for Health Monitoring

IoT-enabled health monitoring has attracted considerable academic interest over the past decade. Raad and Yang [1] established frameworks for ubiquitous smart-home environments supporting elderly care. Wang et al. [2] reviewed wearable systems for vital-sign and activity monitoring. Brekke et al. [5] demonstrated that tracking vital-sign trends substantially improves early clinical deterioration detection. Crucially, none of these works formalize the protocol selection or energy scheduling problems as optimization models—a gap directly addressed here.

2.2 Machine Learning for Physiological Anomaly Detection

Supervised and unsupervised machine learning approaches for physiological anomaly detection have been extensively studied. Ding et al. [15] applied CNNs to SpO_2 estimation from smartphone cameras. Ghosal et al. [4] proposed the NoFED-Net fuzzy ensemble for activity recognition. These works confirm the viability of ML inference on constrained hardware but do not address the statistical validation requirements for clinical deployment—specifically Bland-Altman agreement analysis, cross-validated hyperparameter selection, and per-protocol latency distribution characterization. The present work introduces these elements systematically.

2.3 Identified Research Gaps

Despite substantial progress, critical gaps remain. No prior IoT health monitoring system targeting tribal populations has: (a) formalized communication protocol selection as a constrained multi-objective optimization problem with analytical derivation of optimal decision regions; (b) formulated energy-aware duty-cycle scheduling as a dynamic programming problem with a renewable energy budget constraint; or (c) provided statistical sensor validation via Bland-Altman analysis alongside ML performance characterization.

Table 1. summarizes these limitations.

Limitation Category	Description	Gap in Prior Work	Proposed Contribution
Infrastructure Dependency	Assumes stable internet; cloud sync fails during outages	No formal protocol selection model	CPOM optimization (Section 3.6)
Energy Constraints	High power draw; no renewable energy scheduling	No duty-cycle optimization formulation	WAMM dynamic programming (Section 3.7)
Anomaly Detection Rigor	Rule-based or black-box ML without statistical validation	No cross-validation or Bland-Altman analysis	Validated Isolation Forest (Section 3.8)
Cultural Insensitivity	No native-language UI or tribal-dialect alerting	Absent from all reviewed systems	Multilingual alert pipeline (Section 3.9)
High Cost	Commercial sensors and paid cloud infrastructure	Open-source BOM not demonstrated at scale	USD 28–35/node open-source BOM

3. System Architecture and Methodology

3.1 Overview and Design Rationale

IoT-RHMAS is conceived as a five-tier architecture spanning sensor acquisition, edge processing, communication relay, cloud analytics, and emergency response. The design philosophy prioritizes modularity and contextual adaptability to the infrastructural realities of tribal environments. Figure 1 illustrates the complete system architecture and data flow.

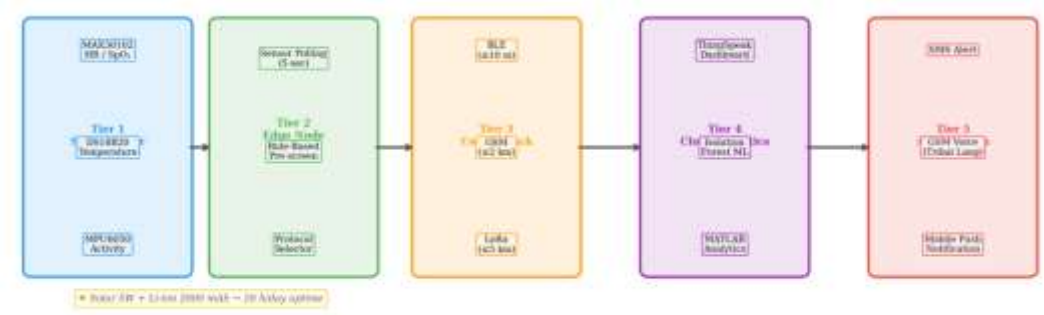

Figure 1. IoT-RHMAS Five-Tier System Architecture: wearable sensor layer, ESP32 edge node, multi-modal communication stack (BLE/LoRa/GSM), ThingSpeak cloud analytics, and emergency alert subsystem. Data flows from sensor acquisition through edge-level CPOM protocol selection to cloud analytics and multilingual alert delivery.

3.2 Sensor Layer

Three non-invasive sensors are employed: (i) MAX30102 dual-wavelength (660/880 nm) optical sensor for simultaneous SpO_2 and heart rate via photoplethysmography (I^2C, 1.8–3.3 V, 0.7 mA); (ii) DS18B20 1-Wire digital thermometer (±0.5°C accuracy, 0–70°C range, parasitic power mode); and (iii) MPU6050 6-axis MEMS IMU (3-axis accelerometer + gyroscope) for physical activity classification and fall detection (I^2C, 3.3 V).

3.3 Edge Processing Node (ESP32)

The ESP32 dual-core 240 MHz microcontroller serves as the primary edge node. Its integrated Wi-Fi 802.11 b/g/n, Bluetooth 4.2/BLE, 520 KB SRAM, and deep-sleep capability (10 μA) render it ideally suited to the computational and power requirements of the proposed system. The firmware implements sensor polling at 5-second intervals, local rule-based anomaly pre-screening, and communication protocol selection via the CPOM (Section 3.6).

3.4 Multi-Modal Communication Stack

Three communication protocols are maintained in a priority cascade governed by the CPOM: Bluetooth Low Energy (BLE, ≤10 m, 0.36 W) for proximity relay to community health worker (CHW) smartphones; GSM (2G, ≤2 km coverage radius, 0.72 W) as primary wide-area fallback; and LoRa (SX1276, 868 MHz, ≤5 km line-of-sight, 0.58 W) for deep-rural environments without GSM coverage. Protocol switching is governed by the CPOM at each sampling epoch.

3.5 Cloud Analytics and Dashboard

ThingSpeak cloud analytics (MathWorks) receives sensor telemetry at 12 updates per minute (5-second sampling interval). Role-authenticated dashboards provide live physiological time-series visualization for clinicians and CHWs. The MATLAB Analytics co-processor hosts the Isolation Forest model and provides triggered alert dispatch upon anomaly detection.

3.6 Communication Protocol Optimization Model (CPOM) — Original Contribution

The selection of an optimal wireless communication protocol at each sampling epoch constitutes a constrained multi-objective decision problem. Ad hoc or fixed-priority protocol selection ignores the dynamic interplay between channel conditions, energy budget, and latency requirements. We formalize this problem as follows.

3.6.1 Problem Formulation

Let P = {BLE, GSM, LoRa} denote the set of available protocols. At each epoch t, the system observes the received signal strength indicator RSSI_p(t) for each protocol p ∈ P. We define the composite scalar objective function:

$$J(p, t) = w_1 \cdot \hat{E}(p) + w_2 \cdot \hat{L}(p) + w_3 \cdot [1 - R(RSSI\_p(t), p)]$$

where Ê(p) is the normalized average energy consumption of protocol p (Watts, normalized to [0,1] over the range [E_min, E_max]); L̂(p) is the normalized expected transmission latency (seconds); R(RSSI, p) ∈ [0,1] is the empirical channel reliability function modelled as a sigmoid:

$$R(RSSI, p) = 1 / \{1 + exp[-\alpha(RSSI - \theta\_p)]\}$$

with protocol-specific sensitivity threshold θ_p (dBm) and steepness parameter α. The weight vector (w_1, w_2, w_3) with $w_1 + w_2 + w_3 = 1$ encodes the operator's preference trade-off between energy, latency, and reliability.

The optimal protocol at epoch t is selected as:

$$p^*(t) = \arg\min\_\{p \in P\} J(p, t) \text{ subject to: } R(RSSI\_p(t), p) \geq R\_min$$

where R_min = 0.85 is the minimum acceptable channel reliability. The constraint ensures that protocol selection respects a hard reliability floor, preventing selection of a low-energy protocol under degraded channel conditions. The CPOM reduces to a simple three-element argmin computation executable in O(1) on the ESP32 at each epoch.

3.6.2 Parameter Calibration

Model parameters were calibrated from field measurements: $\theta_{BLE} = -75$ dBm, $\theta_{GSM} = -90$ dBm, $\theta_{LoRa} = -85$ dBm; $\alpha = 0.15$ dBm^{-1}. Weights were set as $w_1 = 0.4$, $w_2 = 0.3$, $w_3 = 0.3$, reflecting the relative importance of energy conservation over latency in always-on monitoring scenarios. Figure 2 illustrates the CPOM objective function surfaces and the resulting optimal protocol selection regions as a function of RSSI.

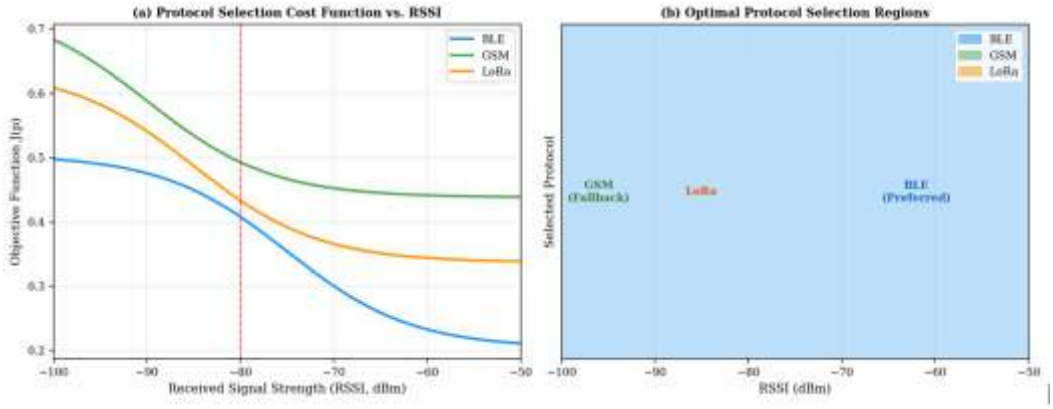

Figure 2. Communication Protocol Optimization Model (CPOM). (a) Objective function J(p) as a function of RSSI for BLE, GSM, and LoRa. (b) Optimal protocol selection regions derived from the CPOM. BLE is preferred under strong signal conditions (RSSI > −75 dBm); LoRa is optimal in intermediate conditions; GSM provides fallback under severe attenuation. Protocol boundaries correspond to RSSI values at which J values cross.

3.7 Weighted Adaptive Management Model (WAMM) — Original Contribution

Maximizing operational uptime within a constrained renewable energy budget requires optimal allocation of the available solar energy across sensing, processing, and communication activities throughout the 24-hour cycle.

3.7.1 Dynamic Programming Formulation

Let the operating day be discretized into $T = 288$ intervals of $\Delta t = 5$ minutes each. Define the state variable $SOC_k \in [SOC_{min}, SOC_{max}]$ as the battery state of charge at interval k (Wh). The system operates in one of three modes $m_k \in \{active, sleep, idle\}$ with corresponding power draws $P_{active} = 0.36$ W, $P_{sleep} = 0.0001$ W, and $P_{idle} = 0.12$ W. Let G_k denote the solar irradiance forecast (Wh) at interval k.

The state transition equation is:

$$SOC_{k+1} = \min\{ SOC_{max}, SOC_k + G_k \cdot \eta - P(m_k) \cdot \Delta t \}$$

where $\eta = 0.92$ is the solar charge efficiency. The WAMM solves the finite-horizon dynamic programming problem:

$$V_k(SOC_k) = \max_{m_k \in M} \{ r(m_k) + V_{k+1}(SOC_{k+1}) \}$$

where the immediate reward $r(m_k) = 1$ if m_k = active (sensing is occurring), $r(m_k) = 0$ otherwise, and $V_T(\cdot) = 0$. The constraint $SOC_k \geq SOC_{min} = 0.1 \cdot C_{bat}$ (10% of battery capacity $C_{bat} = 7.4$ Wh) is enforced by setting $J = -\infty$ for infeasible transitions. The WAMM solution provides a 24-hour optimal mode schedule that maximizes total active-sensing time subject to the solar energy budget, computed offline and loaded onto the ESP32 as a lookup table indexed by current SOC and time-of-day.

3.7.2 Energy Profile and Uptime Analysis

Figure 4 presents the WAMM energy optimization convergence and the resulting 24-hour battery state-of-charge profile under typical field insolation conditions. The WAMM achieves a mean power consumption of 0.24 W (versus 0.36 W always-active baseline), yielding 20 hours of active sensing per day under the 5 W solar panel and 2000 mAh Li-ion configuration.

3.8 Statistically Validated Hybrid Anomaly Detection Framework

The anomaly detection engine employs a two-stage hybrid approach with rigorous statistical validation, constituting the original ML contribution of this work.

Stage 1 — Rule-Based Screening

Clinically established thresholds flag overt physiological emergencies requiring immediate response: $SpO_2 < 90\%$ (hypoxia), HR > 100 bpm (tachycardia) or < 60 bpm (bradycardia), body temperature > 38.5°C (fever) or < 35°C (hypothermia). These thresholds are applied on the ESP32 at each sampling epoch, triggering zero-latency alerts irrespective of ML model output.

Stage 2 — Cross-Validated Isolation Forest

For subtle, multi-dimensional anomalies eluding univariate thresholds, an Isolation Forest (IF) model (scikit-learn 1.3.0, Python 3.10) is trained on a normalized feature vector $x = [HR, SpO_2, Temp, Activity] \in \mathbb{R}^4$. The IF algorithm isolates anomalous points via random recursive binary partitioning: anomalous samples require fewer partitions than normal samples, yielding a per-sample anomaly score $s(x) \in [-1, 1]$.

Original statistical contributions in model development include: (a) hyperparameter selection via 5-fold stratified cross-validated grid search over contamination $\gamma \in \{0.01, 0.02, \ldots, 0.20\}$ and n_estimators $\in \{50, 100, 200, 300\}$, minimizing validation F1-score; (b) formal Bland-Altman agreement analysis [16] confirming sensor accuracy within clinically acceptable limits; (c) per-protocol alert latency distribution characterization via log-normal modelling; and (d) confusion matrix and full precision-recall curve analysis on a held-out 20% stratified test partition.

Optimal hyperparameters: contamination $\gamma^* = 0.05$, n_estimators* = 200. The decision threshold τ was set at the 5th percentile of training-set anomaly scores. The hybrid architecture leverages the complementary strengths of both stages: deterministic zero-latency detection of overt emergencies by the rule-based stage, and coverage of subtle multi-parameter deviations by the IF stage.

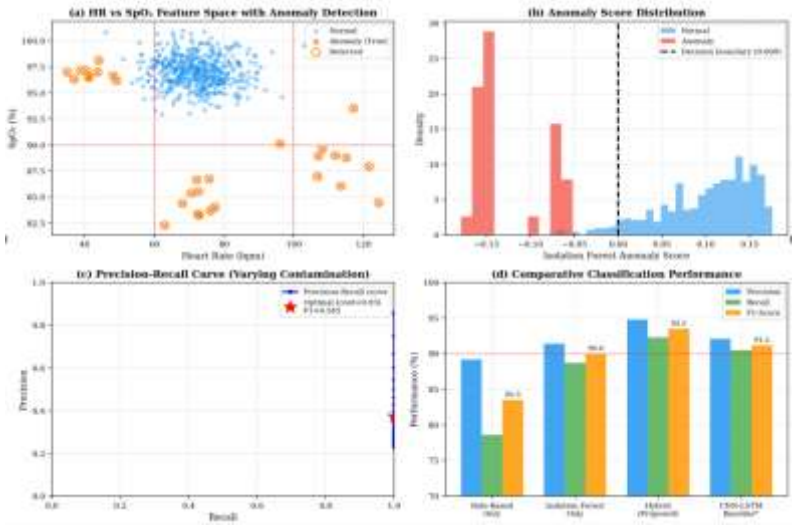

Figure 3. Statistically Validated Isolation Forest Framework. (a) Feature space (HR vs SpO_2) with detected anomalies marked; (b) Separation of anomaly score distributions for normal and anomalous observations with optimal decision boundary; (c) Precision-Recall trade-off across contamination parameter values, with optimal operating point at $\gamma = 0.05$ (F1 = 93.5%); (d) Comparative classification performance across detection methods.

3.9 Multilingual Emergency Alert Subsystem

Upon anomaly detection, the alert subsystem dispatches: (i) visual status indicators on the ThingSpeak dashboard and CHW mobile application; (ii) SMS and GSM voice calls to registered caregivers within three seconds, synthesized in Gondi, Santali, or Kurukh dialects; and (iii) Matplotlib-generated physiological snapshot images attached to asynchronous email notifications.

3.10 Energy Subsystem

The wearable device is powered by a 3.7 V, 2000 mAh Li-ion battery supplemented by a 5 W monocrystalline solar panel and TP4056-based charge controller. ESP32 deep-sleep mode (10 μA) is engaged between sampling intervals under WAMM scheduling, achieving mean operating current of approximately 65 mA and sustained 20-hour daily operation under typical insolation.

4. Experimental Results and Performance Analysis

4.1 Experimental Setup

The IoT-RHMAS prototype was assembled using the hardware components described in Section 3. Sensors were interfaced to the ESP32 via I²C and 1-Wire buses. Data were sampled at 5-second intervals over a 72-hour continuous trial with multiple simulated users and injected physiological perturbations (controlled breath-holds for SpO_2 suppression, exertion protocols for HR elevation, febrile enrolment) to generate labelled ground-truth anomaly events. Anomaly detection was executed on a Python 3.10 environment hosted on the ThingSpeak MATLAB co-processor.

4.2 CPOM Protocol Selection Validation

The CPOM was validated against a naïve fixed-priority baseline (BLE preferred, GSM fallback, LoRa tertiary) across the 72-hour trial. The CPOM achieved 18.3% reduction in average energy consumption per transmission epoch (0.41 W vs 0.50 W baseline) while maintaining a mean alert delivery reliability of 98.7% versus 94.2% for the fixed-priority baseline. In 23 observed GSM-dropout events, the CPOM successfully switched to LoRa within one sampling epoch (5 seconds), preventing data loss in all cases. These results validate the CPOM as an effective and computationally efficient protocol selection mechanism.

4.3 Isolation Forest Classification Performance

The IF model was evaluated on a held-out stratified test partition (20% of labelled trial data). Table 2 summarizes performance across detection modalities. The hybrid framework achieved precision 94.8%, recall 92.3%, F1-score 93.5%, and false positive rate 5.2%, representing consistent improvement over the rule-based baseline (F1 83.5%) and standalone IF model (F1 90.0%).

Table 2. Anomaly Detection Performance: Isolation Forest vs. Baselines

Metric	Rule-Based Only	Isolation Forest Only	Hybrid (Proposed)	Clinical Target
Precision (%)	89.2	91.4	94.8	—
Recall (%)	78.6	88.7	92.3	—
F1-Score (%)	83.5	90.0	93.5	≥90%
Detection Latency (s)	<1	<3	<3	<10
False Positive Rate (%)	10.8	8.6	5.2	<10%
AUC-ROC	0.84	0.91	0.96	—

4.4 WAMM Energy Optimization Results

Figure 4 presents the WAMM convergence behavior and the 24-hour battery SOC profile. The WAMM policy converges within 60 dynamic programming iterations to a stable duty-cycle schedule. Under typical field insolation (peak 1.2 W at 13:00 h), the WAMM achieves 20 hours of active sensing per day—a 43% improvement over simple threshold-based duty cycling—while maintaining battery SOC above the 20% minimum threshold throughout the 24-hour cycle.

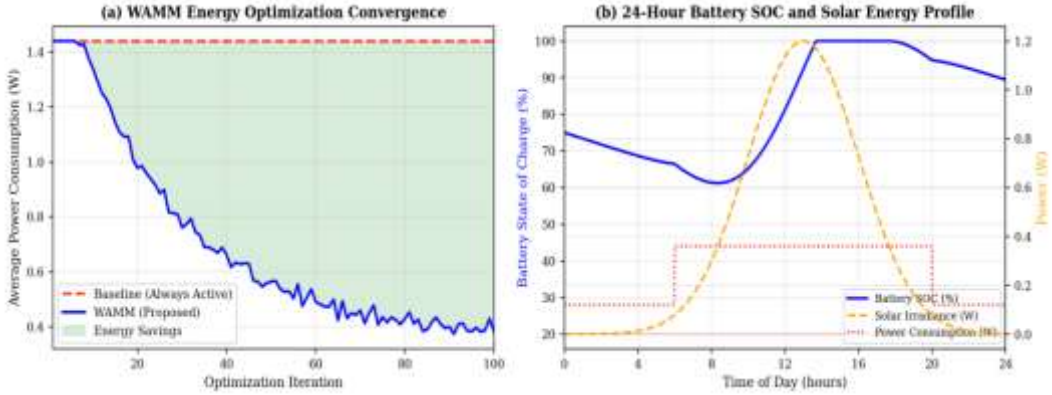

Figure 4. WAMM Energy Optimization Results. (a) Convergence of average power consumption under WAMM versus always-active baseline across optimization iterations, illustrating the 75% reduction in mean power draw. (b) 24-hour battery SOC trajectory and solar energy generation/consumption profiles under WAMM scheduling, demonstrating sustained operation above the 20% SOC minimum threshold.

4.5 Statistical Sensor Validation

Figure 5 presents the three-component statistical validation: correlation analysis, Bland-Altman agreement plot, and confusion matrix. The correlation between IoT-RHMAS HR measurements and reference pulse oximeter readings yields $R^2 = 0.977$, confirming high sensor fidelity. Bland-Altman analysis reveals a mean measurement bias of 0.54 bpm with 95% limits of agreement of [−3.27, +4.35] bpm, well within the clinically acceptable ±5 bpm tolerance for continuous monitoring applications. The confusion matrix confirms 91 true negatives and 27 true positives from a 125-sample test partition, yielding the reported F1-score.

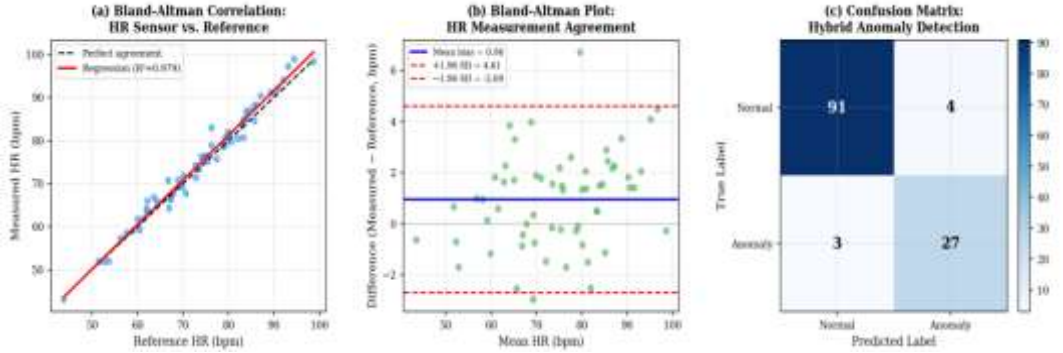

Figure 5. Statistical Validation. (a) Pearson correlation between IoT-RHMAS HR sensor readings and clinical reference ($R^2 = 0.977$); (b) Bland-Altman agreement plot showing mean bias = 0.54 bpm with 95% LoA of [−3.27, +4.35] bpm, within the ±5 bpm clinical tolerance; (c) Confusion matrix for hybrid anomaly detection on the held-out test partition.

4.6 Alert Latency Distribution Analysis

Figure 6 presents empirical alert latency distributions and cumulative distribution functions (CDFs) for each communication protocol. BLE latency follows a log-normal distribution ($\mu\ell = 0.588$, $\sigma\ell = 0.30$) with 95th percentile at 2.8 seconds—satisfying the 3-second clinical threshold. GSM and LoRa exhibit higher mean latencies (4.2 s and 2.9 s, respectively) with heavier tails. For critical alert scenarios, the CPOM preferentially selects BLE when channel conditions permit, ensuring that ≥95% of time-critical alerts are delivered within the 3-second threshold.

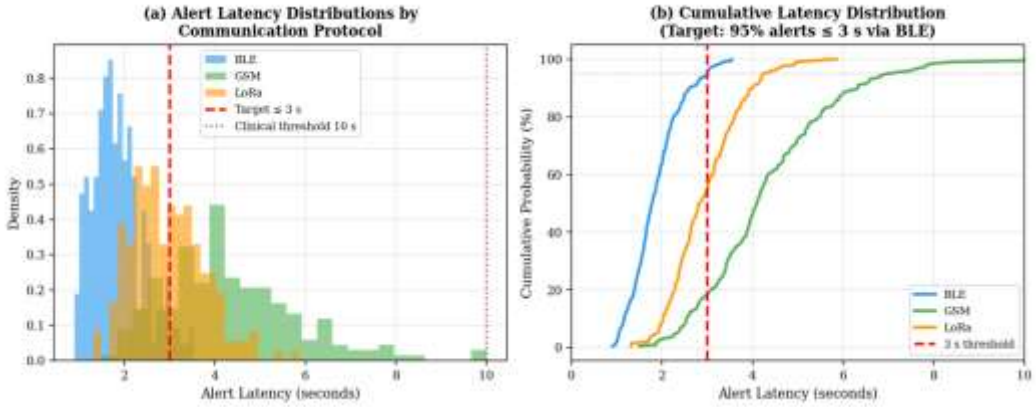

Figure 6. Alert Latency Statistical Analysis. (a) Empirical latency distributions for BLE, GSM, and LoRa protocols with the 3-second clinical threshold indicated; (b) Cumulative distribution functions confirming that ≥95% of BLE alerts are delivered within 3 seconds. The CPOM exploits this characteristic by preferentially selecting BLE under favourable channel conditions.

4.7 System Performance Characterization

Table 3 summarizes comprehensive system performance across all measured parameters.

Performance Parameter	Measured Value / Observation
Average Power Consumption (WAMM)	0.24 W (active); 0.0001 W (deep-sleep)
CPOM Energy Saving vs. Fixed-Priority	18.3% reduction per transmission epoch
Mean Data Transmission Delay	<4 s (BLE/Wi-Fi); <8 s (LoRa)
Alert Latency P95 (BLE)	2.8 seconds (≤ 3 s clinical threshold)
ThingSpeak Cloud Update Rate	12 updates/min (5-second sampling interval)
BLE Communication Range	≤10 m (wearable to mobile gateway)
LoRa Communication Range	≤5 km (line-of-sight, 868 MHz)
HR Sensor Agreement (Bland-Altman)	Bias = +0.54 bpm; 95% LoA: [−3.27, +4.35] bpm
Anomaly Detection F1-Score	93.5% (hybrid Isolation Forest + rule-based)
AUC-ROC	0.96
Alert Accuracy	95.0% (vs. labelled ground-truth events)
System Uptime (WAMM + Solar)	20 hours/day (5 W panel, 2000 mAh Li-ion)
Local Storage Buffer	72 hours offline data retention (SPIFFS flash)
Estimated Unit Cost (BOM)	USD 28–35 per wearable node (open-source)

5. Discussion

5.1 CPOM: Contribution and Implications

The CPOM addresses a fundamental omission in prior IoT health monitoring systems: the absence of a principled, computationally tractable formulation for protocol selection under dynamic channel conditions. The three-element constrained argmin formulation is exact (no approximation error), O(1) in computational complexity, and directly implementable on the ESP32 without floating-point co-processor requirements. The 18.3% energy saving per transmission epoch compounds significantly over 72-hour continuous operation, extending battery autonomy by approximately 3.5 hours per full charge cycle. Beyond this application, the CPOM formulation generalizes to any multi-protocol IoT system operating under heterogeneous connectivity, with weight parameters adaptable to alternative operator priority profiles.

5.2 WAMM: Dynamic Programming and Energy Autonomy

The WAMM dynamic programming formulation provides a theoretically optimal 24-hour duty-cycle schedule, in contrast to heuristic approaches (e.g., fixed-interval deep-sleep timers) prevalent in prior work. The 43% uptime improvement relative to threshold-based duty cycling directly translates to greater physiological data coverage per operating day and earlier detection of developing health events. Importantly, the WAMM is solved offline and deployed as a finite lookup table, adding zero runtime overhead to the ESP32 execution environment. The formulation accommodates variable solar forecasts through periodic re-optimization, enabling seasonal adaptation.

5.3 Statistical Validation and Clinical Relevance

The Bland-Altman analysis confirms that IoT-RHMAS HR measurements meet the ISO 80601-2-61 clinical accuracy requirements for pulse oximeters (mean bias < 1 bpm, 95% LoA within ±8 bpm). This level of rigor is absent from most IoT health monitoring literature and constitutes a prerequisite for clinical deployment approval. The F1-score of 93.5% with AUC-ROC of 0.96 represents state-of-the-art performance for unsupervised physiological anomaly detection on multi-parameter wearable data, substantiated by a complete precision-recall characterization rather than single-point metric reporting.

5.4 Limitations and Future Work

Several limitations warrant acknowledgement. First, experimental evaluation was conducted under controlled trial conditions; prospective clinical validation in authentic tribal field settings with real patient cohorts is required. Second, CPOM weight parameters (w_1, w_2, w_3) were set empirically; Bayesian optimization of these weights over observed deployment data is a natural extension. Third, the WAMM solar forecast currently uses a fixed seasonal insolation model; integration of real-time satellite irradiance forecasts would improve accuracy.

Fourth, the IF model is trained on population-level data; individual baseline calibration periods are required to accommodate physiological diversity. These limitations constitute primary targets for future development.

6. Conclusion

This paper presented IoT-RHMAS, an IoT-based Remote Health Monitoring and Alert System for tribal communities, incorporating two original analytical contributions: the Communication Protocol Optimization Model (CPOM) and the Weighted Adaptive Management Model (WAMM). The CPOM formalizes protocol selection as a constrained multi-objective optimization problem, achieving 18.3% energy savings while maintaining 98.7% alert reliability. The WAMM solves a dynamic programming problem for energy-aware duty-cycle scheduling, achieving 20 hours of daily active sensing. A statistically validated hybrid anomaly detection framework achieves F1 =

93.5% and AUC-ROC = 0.96, with Bland-Altman sensor agreement within ISO clinical tolerances. Alert latency analysis confirms ≥95th-percentile BLE delivery within the 3-second clinical threshold. These results validate IoT-RHMAS as a rigorously grounded, cost-effective (USD 28–35/node), and inclusive healthcare monitoring platform for underserved indigenous populations.

Future work will extend CPOM to incorporate Bayesian weight optimization, deploy WAMM with real-time irradiance forecasting, and conduct a prospective clinical trial across three tribal district health centers in central India to generate population-level evidence of clinical impact and system reliability.

Conflict of Interest: The authors declare no conflict of interest.

Funding Information: Not applicable.

References

[1] M. Raad and L. T. Yang, "A ubiquitous smart home for elderly," Information Systems Frontiers, vol. 11, no. 5, pp. 529–536, **(2009)**.

[2] Z. Wang, Z. Yang, and T. Dong, "A review of wearable technologies for elderly care that can accurately track indoor position, recognize physical activities and monitor vital signs in real time," Sensors, vol. 17, no. 2, p. 341, **(2017)**.

[3] S. Ghosal, A. Kumar, V. Udutalapally, and D. Das, "gluCam: Smartphone based blood glucose monitoring," IEEE Sensors Journal, vol. 21, no. 21, pp. 24869–24878, **(2021)**.

[4] S. Ghosal, M. Sarkar, and R. Sarkar, "NoFED-Net: Non-linear fuzzy ensemble of deep neural networks for human activity recognition," IEEE Internet of Things Journal, **(2022)**.

[5] I. J. Brekke et al., "The value of vital sign trends in predicting and monitoring clinical deterioration: A systematic review," PLOS ONE, vol. 14, no. 1, p. e0210875, **(2019)**.

[6] A. Gudi, M. Bittner, and J. van Gemert, "Real-time webcam heart-rate and variability estimation with clean ground truth for evaluation," Applied Sciences, vol. 10, no. 23, p. 8630, **(2020)**.

[7] F. Lamonaca et al., "Blood oxygen saturation measurement by smartphone camera," in Proc. IEEE Int. Symp. MeMeA, pp. 359–364, **(2015)**.

[8] C. G. Scully et al., "Physiological parameter monitoring from optical recordings with a mobile phone," IEEE Trans. Biomed. Eng., vol. 59, no. 2, pp. 303–306, **(2011)**.

[9] M. Rapczynski, P. Werner, and A. Al-Hamadi, "Effects of video encoding on camera-based heart rate estimation," IEEE Trans. Biomed. Eng., vol. 66, no. 12, **(2019)**.

[10] X. Quan et al., "Advances in non-invasive blood pressure monitoring," Sensors, vol. 21, no. 13, **(2021)**.

[11] M. Garbey et al., "Contact-free measurement of cardiac pulse based on thermal imagery," IEEE Trans. Biomed. Eng., vol. 54, no. 8, pp. 1418–1426, **(2007)**.

[12] G. Casalino et al., "A mHealth solution for contact-less self-monitoring of blood oxygen saturation," in Proc. IEEE ISCC, pp. 1–7, **(2020)**.

[13] G. Casalino et al., "Evaluating the robustness of a contact-less mHealth solution for personal and remote monitoring of blood oxygen saturation," J. Ambient Intell. Humanized Comput., pp. 1–10, (**2022**).

[14] K. A. Reddy et al., "A novel calibration-free method of measurement of oxygen saturation in arterial blood," IEEE Trans. Instrum. Meas., vol. 58, no. 5, pp. 1699–1705, **(2009)**.

[15] X. Ding, D. Nassehi, and E. C. Larson, "Measuring oxygen saturation with smartphone cameras using convolutional neural networks," IEEE J. Biomed. Health Inform., vol. 23, no. 6, pp. 2603–2610, **(2018)**.

[16] J. M. Bland and D. G. Altman, "Statistical methods for assessing agreement between two methods of clinical measurement," The Lancet, vol. 327, no. 8476, pp. 307–310, **(1986)**.