Human Interaction and Emerging Technologies (IHIET 2026), Vol. 236, 2026, 191–200
https://doi.org/10.54941/ahfe1008095

A Modular Edge-to-Cloud Architecture for Remote Monitoring and Condition-Based Maintenance in Scalp Cooling Cyber-Physical Systems

Tauqeer Ali Khan[1,2], Omar Huerta[1], Dipo Olaosun[2], Jonathan Binder[2], Ertu Unver[2], and Peter Culmer[1]

[1]School of Mechanical Engineering, Leeds, LS2 9JT, UK
[2]Paxman Coolers, Huddersfield HD8 0LE, UK

ABSTRACT

Scalp cooling (SC) systems constitute refrigeration-based therapeutic cyber-physical systems in which embedded control, closed-loop thermal regulation, and multi-sensor acquisition must satisfy strict reliability and medical software compliance requirements. Existing embedded architectures for SC devices function as isolated controllers providing local temperature regulation but lack system-level telemetry, condition-based health monitoring, and traceable compliance verification across treatment cycles. This paper presents a modular edge-to-cloud architecture for connected medical refrigeration cyber-physical systems that preserves deterministic embedded control while enabling device observability and fleet-level analytics. The proposed framework integrates heterogeneous sensor acquisition with differentiated sampling cadences: 3-second scheduled paired monitoring of coolant tank temperature and flow rate, event-driven treatment stage capture, and low-frequency coolant pH monitoring. Fault-tolerant local persistence is achieved through an embedded SQLite datastore with periodic batch synchronisation to AWS cloud infrastructure. The architecture also enables automated compliance verification against Instructions for Use (IFU) treatment-stage constraints and condition-based maintenance analytics derived from thermal performance, coolant quality, and device utilisation metrics. The architecture was evaluated across fifteen controlled deployment sessions at two independent sites under ambient conditions ranging from 14.7°C to 26.0°C. All fifteen sessions successfully achieved the target operational temperature of −4°C, confirming that the monitoring layer does not interfere with the original device thermal control. The proposed framework provides a scalable retrofit pathway for legacy refrigeration-based medical devices to transition into connected, observable, and compliance-verifiable cyber-physical system fleets suitable for large-scale clinical deployment.

Keywords: Scalp cooling, Cyber-physical systems, Edge computing, Scheduled monitoring, Condition-based maintenance, Medical IoT

INTRODUCTION

Scalp cooling (SC) reduces chemotherapy-induced alopecia through four distinct controlled mechanisms: vasoconstriction, reduced drug uptake, reduced hair follicle cell division, and reduced metabolic activity. Randomised

Received June 11, 2026; Revised July 5, 2026; Accepted July 29, 2026; Available online August 25, 2026

clinical evidence has established that maintaining defined thermal thresholds yields significant hair preservation outcomes (Nangia et al., 2017; Rugo et al., 2017). From a systems perspective, refrigeration-based SC devices constitute cyber-physical systems (CPS) characterised by tight coupling between sensing, actuation, embedded computation, and thermodynamic heat exchange (Lee, 2008; Rajkumar et al., 2010).

IoT connectivity has demonstrated substantial clinical value across medical device domains. Philip et al. (2021) demonstrated that IoT-enabled in-home health monitoring systems can deliver continuous vital sign acquisition, remote clinician access, and alerting through edge-cloud architectures. Pradhan et al. (2021) established that sensor-cloud integration improves both operational visibility and patient safety outcomes across embedded healthcare devices. Pace et al. (2019) further demonstrated that edge-based IoT architectures can support latency-sensitive healthcare applications in clinical environments while satisfying connectivity and data governance constraints of regulated deployments.

Despite this broader progress, SC devices continue to operate as isolated refrigeration controllers without system-level telemetry, fleet observability, or automated compliance verification. The architectural modelling and system-level analysis of connected therapeutic refrigeration CPS therefore remain largely absent from the embedded systems literature, representing the primary gap this work addresses.

RESEARCH GAPS AND CONTRIBUTIONS

This work addresses three structural gaps. First, refrigeration-based SC devices lack a formalised cyber-physical systems architecture that integrates deterministic embedded control with system-level telemetry and fleet observability. Second, no existing embedded architecture for therapeutic SC systems incorporates an automated compliance verification engine capable of validating treatment sessions against IFU constraints. Third, condition-based maintenance (CBM) frameworks adapted for medical refrigeration devices remain absent from the literature.

The primary contribution of this paper is a monitoring and compliance architecture. The original device retains full responsibility for closed-loop thermal regulation; this architecture observes, records, verifies, and reports. The paper presents: (i) a modular three-tier edge-to-cloud architecture for non-invasive retrofit of existing SC devices; (ii) an embedded IFU compliance verification engine with binary session-level pass/fail reporting and (iii) a condition-based maintenance framework comprising three composite health indicators. The architecture is evaluated across fifteen real-world deployment sessions at two independent clinical sites.

METHODOLOGY

The proposed system implements a three-tier edge-to-cloud architecture designed to retrofit existing scalp cooling devices with scheduled monitoring, IFU compliance verification, and condition-based maintenance capabilities without modifying the underlying closed-loop thermal control system.

The architecture follows an edge-first operational model: all safety-critical data acquisition and compliance logic execute locally on the embedded edge controller, ensuring uninterrupted therapeutic operation independent of network availability. Cloud connectivity is treated as an asynchronous analytics channel rather than a control dependency.

The three tiers are: (1) the Device Tier, handling deterministic sensor acquisition, treatment stage event capture, and local fault-tolerant storage; (2) the Edge-Cloud Synchronisation Tier, managing batch extraction, payload formatting, and secure transmission to cloud infrastructure and (3) the Cloud Analytics Tier, providing relational storage, role-based access control, fleet-level monitoring, and condition-based maintenance dashboards. The overall architecture is illustrated in Figure 1.

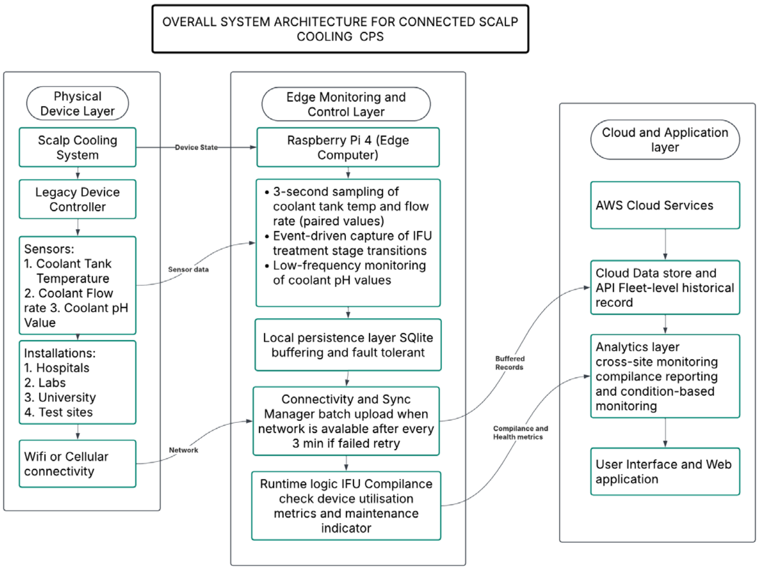

Figure 1: SC System architecture, illustrating the three-tier edge-to-cloud design.

Each SC unit is retrofitted with a Raspberry Pi single-board computer acting as the embedded edge controller. The host platform, Paxman Scalp Cooling System, the global market leader in scalp cooling and an established Class II medical device deployed across more than 60 countries already holds regulatory clearance, so the Raspberry Pi is integrated non-invasively, it interfaces directly with the existing device, intercepting sensor signals and capturing treatment-stage state transitions without replacing or modifying the original control firmware. This approach preserves the device's regulatory approval status and minimises hardware redesign complexity. The choice of Raspberry Pi reflects the research and evaluation context of this deployment, production implementations targeting mass deployment would warrant evaluation of medical-grade embedded alternatives such as dedicated microcontroller units (MCUs) with appropriate certification profiles.

Multi-parameter sensor acquisition is implemented across three distinct modalities with differentiated sampling cadences, as summarised in Table 2. Coolant temperature and flow rate are acquired as synchronous paired measurements at a fixed interval of 3 seconds constituting scheduled periodic sampling, not hard real-time control in the CPS sense. Coolant pH is acquired once per week, reflecting the electrochemical dynamics of coolant degradation. Treatment stage events are captured as asynchronous triggers generated by operator interaction with the existing device UI.

Table 1: Sensor sampling parameters.

Parameter	Sensor Type	Sample Rate	Storage Target	Rationale
Coolant Temp	NTC Thermistor	3 seconds	SQLite (edge)	Primary therapeutic control variable; pull-down and steady-state monitoring
Coolant Flow Rate	flow sensor	3 seconds	SQLite (edge)	Detects blockage, pump degradation, and flow efficiency trends
Coolant pH	Electrochemical pH probe	Once per week	SQLite + AWS	Coolant quality indicator; acceptable range 8.0–9.5
Treatment Stage Events	UI button trigger	Event-driven	SQLite + AWS	IFU stage state machine: pre-infusion, infusion, post-infusion

SQLite Edge Database

All sensor records and event logs are written synchronously to a local SQLite database hosted on the Raspberry Pi edge controller. SQLite was selected for three properties: zero-configuration operation requiring no separate server process, ACID-compliant transaction semantics guaranteeing write integrity under power interruption, and a sufficiently compact footprint for constrained storage capacity. The local persistence layer serves as both a fault-tolerant buffer guaranteeing data capture continuity during network disruptions, and as the authoritative local audit trail satisfying IEC 62304 traceability requirements independent of cloud availability.

Edge-to-Cloud Synchronisation

A background synchronisation service implemented as a cron-scheduled process on the Raspberry Pi executes batch extraction and transmission at a fixed interval of 3 minutes. At each execution, the service queries the local SQLite database for records accumulated since the previous synchronisation timestamp, formats them into structured JSON payloads, and transmits them to the AWS cloud backend over either LTE-M or Wi-Fi connectivity. The 3-minute interval balances cloud ingestion load against monitoring latency, providing near-real-time fleet visibility while avoiding per-record transmission overhead.

Cloud Layer

The cloud layer is implemented on AWS infrastructure, comprising a relational database for persistent record storage, a backend application server managing authentication, device registration, API request validation, and role-based access control, and a web-based monitoring dashboard. Three access roles are defined: administrator access for device fleet management, service engineer access for maintenance analytics and alert management, and R&D access for full historical time-series data and compliance report export. Inbound data is authenticated at the API gateway layer using device-specific credentials and validated against a defined schema before database insertion (see Figure 2).

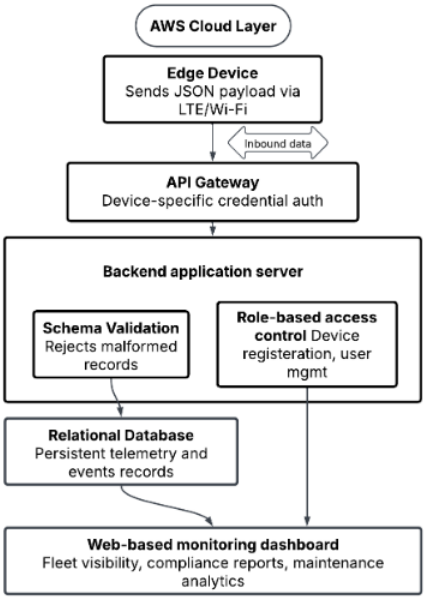

Figure 2: SC system architecture, illustrating the three-tier edge-to-cloud design.

Compliance Verification and Condition Monitoring

The efficacy of scalp cooling in preventing chemotherapy-induced alopecia is contingent upon strict clinician adherence to a three-phase treatment protocol, where each phase serves a distinct physiological purpose and deviation at any stage risks negating the cumulative protective effect. During pre-infusion, the scalp must reach its target temperature to induce vasoconstriction in the follicular microvasculature, reducing both drug perfusion and cellular metabolic activity before cytotoxic agents enter systemic circulation (Grevelman & Breed, 2005; Dunnill et al., 2018). Continuous cooling throughout the infusion phase then sustains this protection during the period of peak plasma drug concentration (Nangia et al., 2017), while the mandatory post-infusion period maintains follicular protection across the pharmacokinetic clearance window against residual cytotoxic exposure—a duration that is regimen-dependent, ranging from 20 minutes for docetaxel to 60 minutes for weekly paclitaxel and up to 90 minutes for other alopecia-inducing regimens per the device instructions for use (IFU) (Komen et al., 2013). Given the high potential for human error in time-critical clinical environments, strict IFU compliance is essential not only for patient safety but for ensuring therapeutic efficacy remains consistent and reproducible across treatment sessions. To address this, the present study implements an embedded verification module on the Raspberry Pi–based SC device that autonomously evaluates each session against key IFU parameters, including attainment of the minimum 30-minute pre-infusion period (extendable to 45 minutes for thick or type 3–4 hair), correct sequential phase progression, and completion of the regimen-specific post-infusion cooling duration. Stage-annotated sensor data enable automatic reconstruction of the full treatment timeline, providing a continuous, automated safeguard against protocol non-compliance across deployed devices.

In parallel, condition monitoring assesses device health using thermal performance, coolant quality, and utilisation indicators. These metrics analyse cooling performance relative to ambient temperature, monitor long-term pH trends, and track device runtime and treatment counts to support condition-based maintenance and early fault detection. Further validation of maintenance effectiveness and false alarm rates will require larger-scale fleet deployments.

EXPERIMENTAL RESULTS

Experimental validation was carried out through 15 controlled deployment sessions conducted across two sites: Paxman Headquarters (Paxman HQ), Huddersfield, and the University of Leeds. Sessions were performed under heterogeneous ambient conditions ranging from 14.7°C to 26.0°C to assess monitoring architecture behaviour and variability in pull-down time under representative real-world operating environments. Each device was instrumented with a Raspberry Pi controller connected to the device sensor interfaces. The cooling device was operated using a small cooling cap connected to a glass head model emulating the thermal load of the human scalp. All measurements were stored locally within the embedded SQLite datastore and synchronised periodically to AWS cloud infrastructure.

Table 2: Experimental deployment dataset (n = 15).

Exp.	Location	Ambient Temp (°C)	Pull-Down Time (min)	Target −4°C Achieved
1	Paxman HQ	15.7	8	Yes
2	Paxman HQ	16.4	8	Yes
3	Paxman HQ	14.7	9	Yes
4	Paxman HQ	19.8	9	Yes
5	University of Leeds	23.0	17	Yes
6	Paxman HQ	17.0	9	Yes
7	University of Leeds	22.6	16	Yes
8	University of Leeds	18.0	15	Yes
9	University of Leeds	18.0	15	Yes
10	Paxman HQ	26.0	13	Yes
11	University of Leeds	22.0	17	Yes
12	Paxman HQ	16.6	8	Yes
13	University of Leeds	18.2	14	Yes
14	Paxman HQ	16.0	10	Yes
15	University of Leeds	21.0	17	Yes

Cross Site Deployment

The higher average ambient temperatures observed at the University of Leeds (20.7°C versus 17.8°C at Paxman HQ) directly explain the longer pull-down times and slower cooling rates recorded there. Both devices achieved the −4°C target in every session, confirming the robustness of the underlying closed-loop thermal control system and the non-interference property of the monitoring architecture.

Table 3: Comparison.

Metric	Paxman HQ	University of Leeds
Number of Sessions	8	7
Ambient Temp Range	14.7–26.0°C	18.0–23.0°C
Average Ambient Temp	17.8°C	20.7°C
Pull-Down Time Range	8–13 min	14–17 min
Average Pull-Down Time	9.25 min	15.9 min
Average Cooling Rate	2.17°C/min	1.50°C/min
All Sessions Reached −4°C	Yes	Yes

Treatment Stage Observability

Treatment stage observability and compliance verification form the key monitoring objective of the proposed architecture, enabling automated assessment of treatment execution in accordance with the IFU protocol. Each treatment session comprises three stages: pre-infusion cooling, infusion cooling, and post-infusion cooling. Across all fully instrumented sessions conducted at both sites, the recorded stage durations consistently satisfied IFU-defined requirements, with pre-infusion durations of approximately 30 minutes, post-infusion durations of at least 20 minutes, and correct sequential stage progression, resulting in full compliance across all evaluated sessions (see Figure 3).

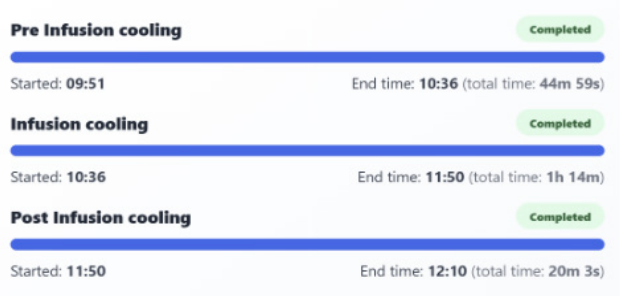

Figure 3: Treatment stage timeline illustrating pre-infusion, infusion, and post-infusion.

CONCLUSION

This paper presented and validated a modular edge-to-cloud monitoring architecture for connected medical refrigeration cyber-physical systems through fifteen deployment sessions across two independent sites. The proposed framework combines embedded IFU compliance verification with remote monitoring, enabling automated assessment of treatment adherence, device performance, coolant quality, and utilisation trends while maintaining complete separation from the device's closed-loop thermal control functions. By incorporating ambient-normalised thermal performance analysis, the architecture provides a practical approach for condition-based maintenance that distinguishes environmental influences from potential mechanical degradation, addressing limitations of traditional calendar-based maintenance strategies. All deployment sessions successfully achieved the target operational temperature of −4°C while the monitoring system remained active, confirming non-interference with therapeutic operation. Beyond its technical contributions, the framework supports improved healthcare delivery by relieving clinicians of manual compliance checks, safeguarding the protocol adherence on which therapeutic efficacy depends, and helping ensure device readiness so that patients receive consistent, uninterrupted treatment together creating the conditions for the highest achievable patient outcomes. Overall, the results demonstrate the potential of the proposed framework as a scalable monitoring, compliance, and maintenance solution for future large-scale clinical deployments of connected therapeutic refrigeration devices.

REFERENCES

Beghi, A., Cecchinato, L., Cosi, G. and Rampazzo, M. (2011). Fault detection and diagnosis for HVAC systems. Applied Thermal Engineering, 31(17–18), pp. 3612–3622.

Bonomi, F., Milito, R., Zhu, J. and Addepalli, S. (2012). Fog computing and its role in the Internet of Things. Proceedings of the MCC Workshop on Mobile Cloud Computing, pp. 13–16.

Chen, Y., Zheng, Y., Li, W. and Huang, X. (2023). Edge intelligence for IoMT: A survey on resource-constrained inference and on-device learning. IEEE Internet of Things Journal, 10(5), pp. 3878–3897.

Dunnill, C. J., Al-Tameemi, W., Collett, A., Haslam, I. S., & Georgopoulos, N. T. (2018). A clinical and biological guide for understanding chemotherapy-induced alopecia and its prevention. The Oncologist, 23(1), 84–96

Grevelman, E. G., & Breed, W. P. M. (2005). Prevention of chemotherapy-induced hair loss by scalp cooling. Annals of Oncology, 16(3), 352–358.

International Electrotechnical Commission (2006). IEC 62304: Medical device software — Software life cycle processes. Amendment 1, 2015.

International Organization for Standardization (2016). ISO 13485: Medical devices — Quality management systems.

Islam, M.M., Rahaman, A. and Islam, M.R. (2015). The Internet of Things for health care: A comprehensive survey. IEEE Access, 3, pp. 678–708.

Katipamula, S. and Brambley, M.R. (2005). Methods for fault detection, diagnostics, and prognostics for building systems. HVAC&R Research, 11(1), pp. 3–25.

Komen, M. M. C., Smorenburg, C. H., van den Hurk, C. J. G., & Nortier, J. W. R. (2013). Factors influencing the effectiveness of scalp cooling in the prevention of chemotherapy-induced alopecia. The Oncologist, 18(8), 885–891.

Lee, E.A. (2008). Cyber physical systems: Design challenges. Proceedings of IEEE ISORC, pp. 363–369.

Lee, J., Davari, H., Singh, J. and Pandhare, V. (2018). Industrial artificial intelligence for Industry 4.0-based manufacturing systems. Manufacturing Letters, 18, pp. 20–23.

Medical Device Coordination Group (2019). MDCG 2019-16: Guidance on cybersecurity for medical devices. European Commission.

Mobley, R.K. (2002). An Introduction to Predictive Maintenance. 2nd ed. Oxford: Butterworth-Heinemann.

Nangia, D., Baranwal, N., Bhatt, P. and Singhal, S. (2017). Effect of a scalp cooling device on alopecia in women undergoing chemotherapy for breast cancer: The SCALP randomised clinical trial. JAMA, 317(6), pp. 596–605.

Pace, P., Aloi, G., Gravina, R., Caliciuri, G., Fortino, G. and Liotta, A. (2019). An edge-based architecture to support efficient applications for healthcare industry 4.0. IEEE Transactions on Industrial Informatics, 15(1), pp. 481–489.

Pang, G., Shen, C., Cao, L. and Hengel, A.V.D. (2021). Deep learning for anomaly detection: A review. ACM Computing Surveys, 54(2), pp. 1–38.

Philip, N.Y., Rodrigues, J.J.P.C., Wang, H., Fong, S.J. and Chen, J. (2021). Internet of Things for in-home health monitoring systems: Current advances, challenges and future directions. IEEE Journal on Selected Areas in Communications, 39(2), pp. 300–310.

Pradhan, B., Bhattacharyya, S. and Pal, K. (2021). IoT-based applications in healthcare devices. Journal of Healthcare Engineering, 2021, Art. no. 6632599.

Rahmani, A.M., Gia, T.N., Negash, B., Anzanpour, A., Azimi, I., Jiang, M. and Liljeberg, P. (2018). Internet of Things for healthcare: Applications, challenges, and solutions. IEEE Reviews in Biomedical Engineering, 11, pp. 68–83.

Rajkumar, R., Lee, I., Sha, L. and Stankovic, J. (2010). Cyber-physical systems: The next computing revolution. Proceedings of ACM/IEEE DAC, pp. 731–736.

Rugo, H.S., Klein, P., Melin, S.A., Hurvitz, S.A., Melisko, M.E., Moore, A., Park, G., Mitchel, J. and Somlo, G. (2017). Association between use of a scalp cooling device and alopecia after chemotherapy for breast cancer. JAMA, 317(6), pp. 606–614.

Shi, W., Cao, J., Zhang, Q., Li, Y. and Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), pp. 637–646.

U.S. Food and Drug Administration (2023). Cybersecurity *in medical de*vices: Quality system considerations and content of premarket submissions. FDA Guidance.

Wang, S. and Xiao, F. (2004). AHU sensor fau*lt diagnosis using pri*ncipal component analysis. Energy and Buildings, 36(2), pp. 147–160.

Yin, S., Ding, S.X., Xie, X. and Luo, H. (2015). Data-based techniques focused on modern industry: An overview. IEEE Transactions on Industrial *Electronics, 62*(1), pp. 657–667.