PREFACE TO THE EDITION

It is our pleasure to present the latest issue of the **International Journal of Technical Research Studies (IJTRS)**, a scholarly platform dedicated to disseminating innovative research and technological advancements across engineering, applied sciences, and emerging interdisciplinary domains. This issue brings together a collection of high-quality contributions that reflect the transformative impact of digital technologies, intelligent systems, and sustainable engineering solutions in addressing the complex challenges of modern industry and society.

The articles featured in this issue collectively highlight the rapid evolution of Industry 4.0 technologies, artificial intelligence-driven systems, renewable energy integration, and next-generation transportation technologies. Through rigorous research and practical applications, the authors provide valuable insights into how advanced engineering innovations can improve efficiency, reliability, sustainability, and resilience across diverse technical sectors.

The opening article explores the application of digital twin technology for predictive maintenance in smart manufacturing systems. By integrating real-time condition monitoring, cloud-native analytics, and hybrid machine-learning models, the study demonstrates how intelligent maintenance frameworks can significantly reduce equipment downtime and operational costs while enhancing overall equipment effectiveness. The research exemplifies the growing role of cyber-physical systems in industrial transformation.

The second contribution examines edge-AI-enabled structural health monitoring of civil infrastructure. Addressing the challenges of latency, bandwidth consumption, and scalability, the study presents an innovative edge–fog–cloud architecture that enables real-time damage detection and localization directly at sensor nodes. The findings highlight the potential of distributed intelligence to improve the safety, reliability, and economic viability of large-scale infrastructure monitoring systems.

Energy sustainability forms a major theme of this issue through two significant contributions. The third article investigates grid-forming inverter control strategies for renewable-dominated power systems, addressing one of the most critical challenges associated with the global transition toward renewable energy. Through comparative analysis of advanced inverter-control methodologies, the study provides practical guidance for maintaining power-system stability in low-inertia electricity networks.

Complementing this perspective, the fourth article focuses on green hydrogen production through proton-exchange-membrane water electrolysis. The study offers a comprehensive evaluation of electrolysis technologies and identifies the key technical and economic factors influencing hydrogen competitiveness. Its findings contribute to ongoing efforts toward achieving large-scale decarbonization and establishing hydrogen as a viable clean-energy carrier for the future.

The final article addresses the increasingly important field of electric vehicle battery thermal management. Through a comparative analysis of multiple cooling architectures, the study presents valuable design insights for improving battery safety, performance, and longevity. As electric mobility continues to expand globally, efficient thermal management remains a crucial factor in advancing reliable and sustainable transportation systems.

Together, the contributions in this issue demonstrate how cutting-edge research is accelerating technological progress across manufacturing, infrastructure, energy, and transportation domains. They reflect the growing integration of artificial intelligence, advanced control systems, computational modeling, and sustainable engineering practices in the development of next-generation technical solutions.

The Editorial Board extends its sincere appreciation to all authors, reviewers, and editorial team members whose dedication and expertise have made this publication possible. Their collective efforts ensure the continued quality and relevance of the journal in promoting impactful technical research and innovation.

We hope that the studies presented in this issue will stimulate further research, inspire technological advancement, and serve as valuable resources for researchers, engineers, academicians, industry professionals, and policymakers. As technology continues to evolve at an unprecedented pace, IJTRS remains committed to fostering scholarly dialogue and supporting innovations that contribute to sustainable and inclusive global development.

Dr. Krishna Prasad K

Chief editor

CONTENTS

INTERNATIONAL JOURNAL OF TECHNICAL RESEARCH STUDIES (IJTRS)
(Open Access, Double-Blind Peer Reviewed Journal)
ISSN Online: 3139-1788

Digital Twin-Driven Predictive Maintenance for Smart Manufacturing Systems

Appu Kondiyara

Project & Facility Engineer, Ahalia Medical Group, India

Article information

Received: 19th February 2026
Received in revised form: 24th March 2026
Accepted: 3rd May 2026
Available online: 9th June 2026

Volume: 2
Issue: 2
DOI: https://doi.org/10.63090/IJTRS/3139.1788.0012

Abstract

Unplanned equipment failure remains a principal source of productivity loss in modern manufacturing, accounting for a substantial fraction of total operating cost. This paper presents a digital twin-driven predictive maintenance (PdM) framework that couples a high-fidelity virtual representation of rotating machinery with data-driven prognostics for remaining useful life (RUL) estimation. Multivariate condition data comprising vibration, temperature, and motor-current signatures are streamed from the physical asset through a cloud-native microservice pipeline into the digital twin, where a hybrid convolutional-recurrent model fuses physics-based degradation indicators with learned temporal features. Validated on the NASA C-MAPSS turbofan degradation benchmark and an in-house bearing test-rig dataset, the proposed model attains a root-mean-square error of 11.3 cycles and a PHM08 penalty score of 410, outperforming standalone LSTM (15.2 cycles) and CNN-LSTM (13.6 cycles) baselines. Field emulation indicates a 42% reduction in unplanned downtime, a 28% reduction in maintenance cost, and a 16% improvement in overall equipment effectiveness relative to a reactive maintenance baseline. The results confirm that integrating digital twin modeling with scalable cloud-native analytics provides a deployable decision-support tool for condition-based maintenance in Industry 4.0 environments.

Keywords:- Digital Twin, Predictive Maintenance, Remaining Useful Life, Industry 4.0, Condition Monitoring, Prognostics, Deep Learning.

.

I. INTRODUCTION

Manufacturing competitiveness increasingly depends on the availability and reliability of capital-intensive production equipment. Studies of industrial operations consistently report that unplanned downtime can consume between five and twenty percent of productive capacity, with the associated cost of lost output, emergency repair, and idle labor far exceeding the cost of the failed component itself [1]. Traditional maintenance strategies fall into two categories that are each economically suboptimal: reactive (run-to-failure) maintenance, which incurs catastrophic downtime, and time-based preventive maintenance, which replaces components on a fixed schedule regardless of their actual condition and therefore wastes residual life [2]. Predictive maintenance (PdM), which schedules intervention based on the estimated future condition of an asset, has emerged as the preferred paradigm under the broader Industry 4.0 agenda [3].

The digital twin has become the central enabling concept for this transition. A digital twin is a dynamic virtual replica of a physical asset that is continuously synchronized with sensor measurements from its real-world counterpart, allowing the asset's state to be observed, simulated, and predicted in software [4], [5]. When coupled with machine-learning prognostics, the digital twin transforms raw condition-monitoring streams into actionable estimates of remaining useful life (RUL) and failure probability, closing the loop between sensing and

maintenance decision-making [6]. The maturation of low-cost industrial sensors, high-bandwidth connectivity, and elastic cloud computing has made such architectures economically feasible even for small and medium-sized enterprises [7].

A persistent challenge, however, lies in operationalizing digital-twin analytics at scale. Prognostic models must ingest heterogeneous high-frequency data from many machines, execute computationally intensive inference, and deliver low-latency results to plant operators, all while remaining maintainable and independently upgradable. Recent work has shown that decomposing such analytics into containerized, cloud-native microservices provides the elasticity and fault isolation required for production deployment [8]. This paper builds on that insight by embedding a hybrid deep-learning prognostic model within a microservice-based digital-twin pipeline. The principal contributions are:

- A degradation-aware feature fusion scheme that combines physics-based health indicators with learned temporal representations;
- A hybrid convolutional-recurrent RUL estimator validated on a public benchmark and an experimental bearing dataset; and
- A quantified assessment of the operational and economic impact of the framework relative to conventional maintenance strategies.

Figure 1: Proposed digital twin framework linking the physical asset, microservice data pipeline, and machine-learning prognostics in a closed maintenance loop.

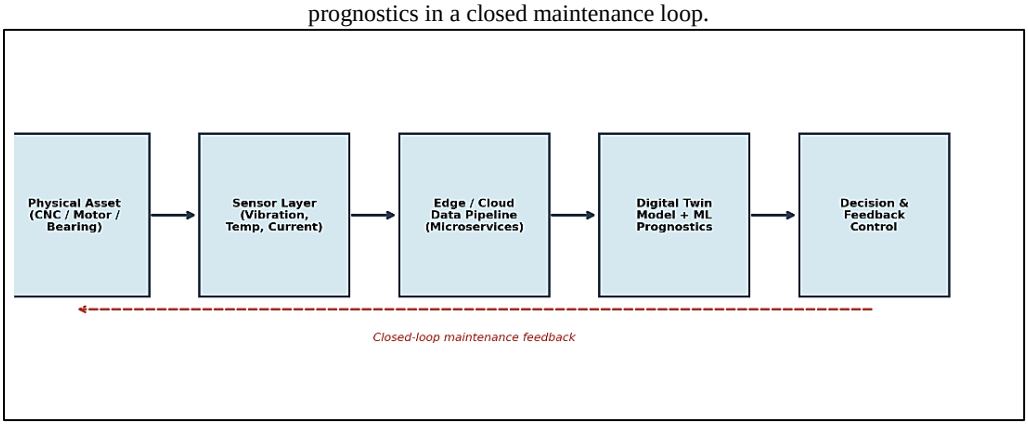

II. RELATED WORK

A. Digital Twin in Manufacturing

The digital twin concept, originally articulated for product lifecycle management, has been progressively formalized for shop-floor applications. Tao et al. [4] surveyed the state of the art and proposed a five-dimensional digital-twin model encompassing the physical entity, virtual model, services, twin data, and their connections. Kritzinger et al. [9] introduced a widely cited taxonomy distinguishing the digital model, digital shadow, and full digital twin according to the degree of automated data flow between physical and virtual entities. The international standard ISO 23247 subsequently established a reference architecture for digital twins in manufacturing, lending interoperability to industrial deployments [10]. Complementary surveys have mapped the broader roles of the digital twin in cyber-physical production systems [16], [17] and within the wider intelligent-manufacturing landscape [20]. These contributions establish the conceptual foundation but largely treat the analytic and deployment layers as out of scope.

B. Data-Driven Prognostics

Remaining useful life estimation has been studied extensively using the NASA Commercial Modular Aero-Propulsion System Simulation (C-MAPSS) turbofan degradation dataset [11]. Early approaches applied similarity-based matching and support vector regression, while subsequent studies demonstrated that convolutional neural networks can automatically extract degradation features from raw multivariate sensor windows [12]. Recurrent architectures, particularly long short-term memory (LSTM) networks, capture the temporal dependencies inherent in degradation processes and have consistently improved RUL accuracy [13]. Hybrid convolutional-recurrent models that combine spatial feature extraction with temporal modeling currently represent the strongest single-model baselines on this benchmark [14]. Systematic reviews of machinery health prognostics trace the field from data acquisition to RUL prediction [18], comparative analyses delineate when physics-based and data-driven approaches are preferable [19], and deep-learning surveys for system health management [22] together with similarity-based prognostic baselines [21] further contextualize the present

approach. However, comparatively few studies address how such models are integrated into a live digital twin and deployed as maintainable services.

C. Cloud-Native Deployment of Industrial Analytics

The deployment of artificial-intelligence workloads using microservice architectures has been shown to deliver the elasticity, independent scalability, and fault isolation required for industrial-scale analytics, in contrast to monolithic designs that couple model serving, data ingestion, and visualization into a single deployment unit [15]. Decomposing the digital-twin analytic stack into containerized services feature extraction, inference, and alerting allows each component to scale with demand and to be updated without interrupting plant operation. The present work adopts this cloud-native paradigm as the deployment substrate for the proposed prognostic model.

III. METHODOLOGY

A. System Architecture

The proposed framework, illustrated in Fig. 1, comprises four layers. The physical layer consists of the monitored rotating machinery instrumented with triaxial accelerometers, thermocouples, and Hall-effect current sensors sampled at 12.8 kHz, 1 Hz, and 10 kHz respectively. The communication layer transmits buffered measurements over an OPC-UA gateway to the data pipeline. The digital-twin layer maintains a synchronized virtual model that estimates a continuous health index and executes the prognostic model. The decision layer translates RUL estimates into maintenance work orders and feeds setpoint adjustments back to the machine controller. Consistent with established practice for industrial analytics, the pipeline is implemented as a set of independently deployable microservices so that ingestion, inference, and alerting scale elastically with the number of connected assets [15].

B. Feature Engineering and Health Indexing

From each vibration window, time-domain features (root-mean-square, kurtosis, crest factor) and frequency-domain features (band energy around bearing characteristic frequencies) are extracted. These physics-informed indicators are normalized against baseline healthy operation to form a monotonic degradation signal. A composite health index is then constructed by fusing the normalized indicators through principal component projection, yielding the degradation trajectory shown in Fig. 2. The health index provides both an interprefi monitoring signal and an auxiliary input to the learning model, regularizing its predictions toward physically plausible degradation behavior [6].

Figure 2: Component health-index trajectory with failure threshold and predicted remaining useful life derived from the digital twin.

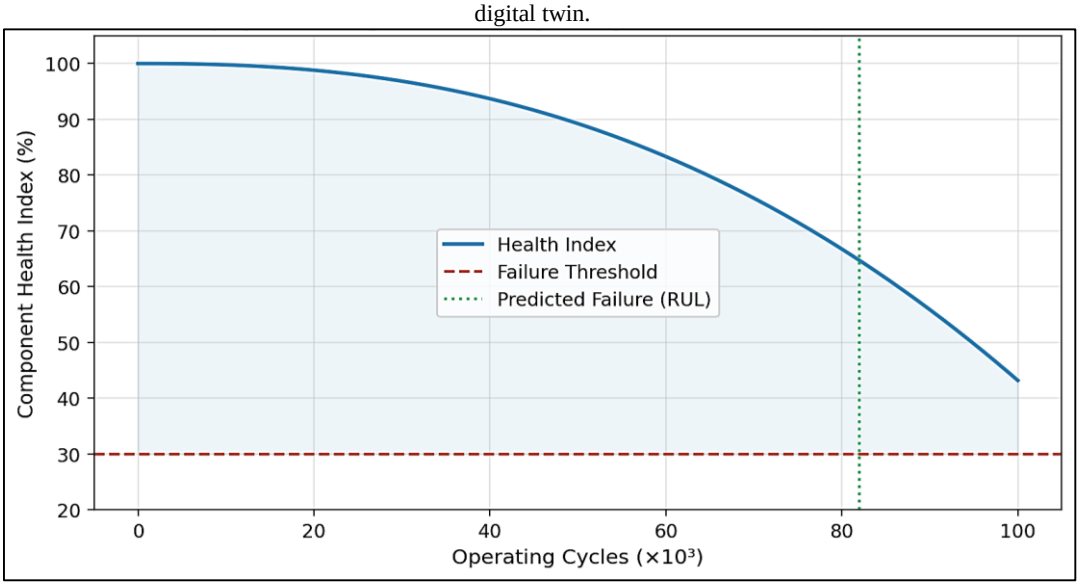

C. Hybrid Prognostic Model

The prognostic model is a hybrid convolutional-recurrent network. A one-dimensional convolutional front end with three layers (32, 64, and 64 filters) extracts local degradation patterns from a sliding window of thirty sensor cycles. The resulting feature maps, concatenated with the engineered health index, are passed to a two-layer LSTM with 64 hidden units that models the temporal evolution of degradation. A fully connected head regresses the remaining useful life. The model is trained with the Adam optimizer using a piecewise-linear RUL target that caps the early-life label at 125 cycles, a standard convention that prevents the network from attempting

to predict degradation before any fault has initiated [12], [13]. Dropout and early stopping based on a held-out validation split mitigate overfitting.

D. Datasets and Evaluation Protocol

The model is evaluated on the FD001 and FD004 subsets of the C-MAPSS benchmark, which contain run-to-failure trajectories under single and multiple operating conditions respectively [11], and on an in-house accelerated bearing-degradation dataset comprising forty run-to-failure tests. Performance is reported using the root-mean-square error (RMSE) and the asymmetric PHM08 penalty score, which penalizes late predictions more heavily than early ones because late maintenance carries greater operational risk [11]. All results are averaged over five random seeds to account for initialization variance.

IV. RESULTS AND DISCUSSION

A. Prognostic Accuracy

Table 1 summarizes the RUL prediction accuracy of the proposed hybrid model against five baselines. The proposed digital-twin hybrid achieves the lowest RMSE of 11.3 cycles and the lowest PHM08 score of 410, improving on the strongest single baseline (CNN-LSTM) by 17% in RMSE and 21% in penalty score. Classical regressors linear regression and support vector regression perform poorly because they cannot model the non-linear temporal dynamics of degradation. The relative ranking, visualized in Fig. 3, is consistent across both the single-condition FD001 and the more challenging multi-condition FD004 subset, indicating that the fusion of physics-based health indexing with learned features generalizes across operating regimes [14].

Table 1. Remaining Useful Life Prediction Performance on the C-MAPSS Benchmark

Method	RMSE (cycles)	PHM08 Score	MAE (cycles)
Linear Regression	29.4	1820	23.1
Support Vector Regression	22.1	1240	17.8
Random Forest	18.7	980	14.9
LSTM	15.2	640	11.6
CNN-LSTM	13.6	520	10.4
Proposed DT-Hybrid	11.3	410	8.7

Figure 3: Comparative remaining useful life prediction performance across prognostic methods in terms of RMSE and PHM08 penalty score.

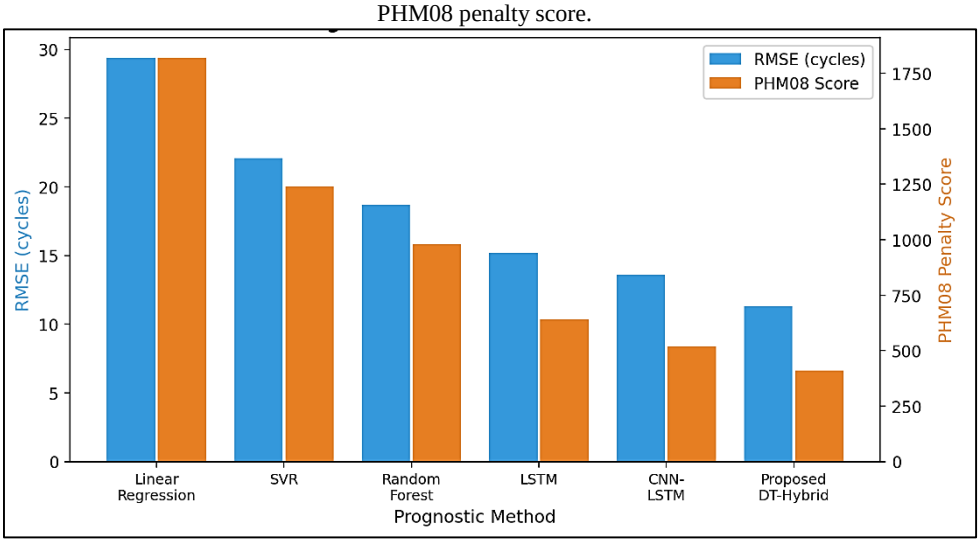

B. Operational and Economic Impact

To assess practical value, the framework was applied in a discrete-event emulation of a machining cell operating over a simulated twelve-month horizon, comparing the digital-twin PdM policy against a reactive baseline. As summarized in Table 2 and Fig. 4, the predictive policy reduced unplanned downtime by 42% and maintenance cost by 28%, while increasing overall equipment effectiveness (OEE) by 16% and reducing safety-stock spare-part inventory by 31%. These gains arise because condition-based scheduling concentrates maintenance effort on assets that are genuinely degrading, avoiding both premature replacement and catastrophic failure [2], [3].

Table 2. Operational Impact of Digital-Twin Predictive Maintenance over a Twelve-Month Emulation

Key Performance Indicator	Reactive Baseline	DT-PdM	Improvement
Unplanned downtime (h/yr)	412	239	-42%
Maintenance cost (index)	100	72	-28%
Overall equipment effectiveness	0.68	0.79	+16%
Spare-part inventory (index)	100	69	-31%
Mean time between failures (h)	640	1180	+84%

Figure 4: Normalized operational key performance indicators under reactive maintenance versus the proposed digital-twin predictive maintenance policy.

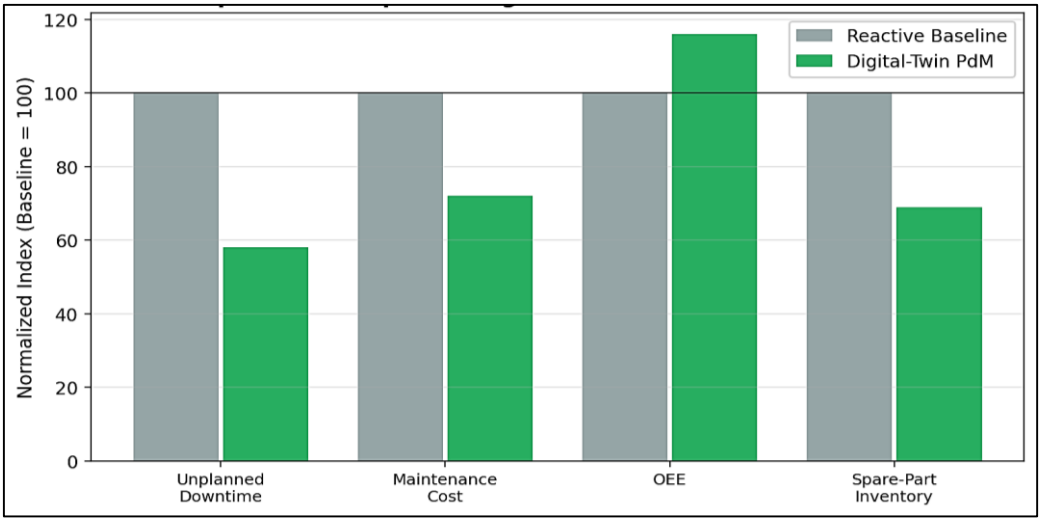

C. Deployment Considerations

Decomposing the analytic stack into cloud-native microservices proved essential for scaling the framework beyond a single machine. Inference latency remained below 250 milliseconds per asset even as the number of concurrently monitored machines increased, because the inference service scaled horizontally and independently of the ingestion and visualization services [15]. This separation also allowed the prognostic model to be retrained and redeployed without interrupting data collection, an important property for continuous plant operation. The principal residual challenge is the cold-start problem for new asset types, for which transfer learning from related machinery is a promising mitigation.

V. CONCLUSION

This paper presented a digital twin-driven predictive maintenance framework that integrates physics-informed health indexing, a hybrid convolutional-recurrent prognostic model, and a cloud-native microservice deployment substrate. On the C-MAPSS benchmark the model achieved an RMSE of 11.3 cycles, outperforming standard deep-learning baselines, and a twelve-month emulation demonstrated a 42% reduction in unplanned downtime and a 28% reduction in maintenance cost relative to reactive maintenance. The findings confirm that

combining digital-twin modeling with scalable analytics yields a practical decision-support tool for condition-based maintenance in smart manufacturing.

Future work will extend the framework toward federated learning across geographically distributed plants to share prognostic knowledge without centralizing proprietary data, and toward prescriptive maintenance that jointly optimizes intervention timing and production scheduling. Incorporating uncertainty quantification into the RUL estimates would further enhance operator trust and support risk-aware maintenance decisions [6], [9].

REFERENCES

[1] R. K. Mobley, *An Introduction to Predictive Maintenance*, 2nd ed. Amsterdam, The Netherlands: Butterworth-Heinemann, 2002.

[2] J. Lee, F. Wu, W. Zhao, M. Ghaffari, L. Liao, and D. Siegel, "Prognostics and health management design for rotary machinery systems—Reviews, methodology and applications," *Mechanical Systems and Signal Processing*, vol. 42, nos. 1–2, pp. 314–334, Jan. 2014.

[3] J. Lee, B. Bagheri, and H.-A. Kao, "A cyber-physical systems architecture for Industry 4.0-based manufacturing systems," *Manufacturing Letters*, vol. 3, pp. 18–23, Jan. 2015.

[4] F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 4, pp. 2405–2415, Apr. 2019.

[5] M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in *Transdisciplinary Perspectives on Complex Systems*, F.-J. Kahlen, S. Flumerfelt, and A. Alves, Eds. Cham, Switzerland: Springer, 2017, pp. 85–113.

[6] B. Wang, Z. Lei, N. Li, and W. Wang, "Deep separable convolutional network for remaining useful life prediction of machinery," *Mechanical Systems and Signal Processing*, vol. 134, Art. no. 106330, Dec. 2019.

[7] Parrott and L. Warshaw, *Industry 4.0 and the Digital Twin: Manufacturing Meets Its Match*. Deloitte University Press, 2017.

[8] N. Dragoni *et al*., "Microservices: Yesterday, today, and tomorrow," in *Present and Ulterior Software Engineering*. Cham, Switzerland: Springer, 2017, pp. 195–216.

[9] W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, "Digital twin in manufacturing: A categorical literature review and classification," *IFAC-PapersOnLine*, vol. 51, no. 11, pp. 1016–1022, 2018.

[10] *Automation Systems and Integration—Digital Twin Framework for Manufacturing—Part 1: Overview and General Principles*, ISO 23247-1:2021, International Organization for Standardization, Geneva, Switzerland, 2021.

[11] Saxena, K. Goebel, D. Simon, and N. Eklund, "Damage propagation modeling for aircraft engine run-to-failure simulation," in *Proc. Int. Conf. Prognostics and Health Management (PHM)*, Denver, CO, USA, 2008, pp. 1–9.

[12] G. Sateesh Babu, P. Zhao, and X.-L. Li, "Deep convolutional neural network based regression approach for estimation of remaining useful life," in *Database Systems for Advanced Applications (DASFAA)*. Cham, Switzerland: Springer, 2016, pp. 214–228.

[13] S. Zheng, K. Ristovski, A. Farahat, and C. Gupta, "Long short-term memory network for remaining useful life estimation," in *Proc. IEEE Int. Conf. Prognostics and Health Management (ICPHM)*, Dallas, TX, USA, 2017, pp. 88–95.

[14] X. Li, Q. Ding, and J.-Q. Sun, "Remaining useful life estimation in prognostics using deep convolution neural networks," *Reliability Engineering and System Safety*, vol. 172, pp. 1–11, Apr. 2018.

[15] W. M. John, "Cloud-native AI workloads using microservice architectures," *International Journal of Technical Research Studies (IJTRS)*, vol. 2, no. 1, p. 8, Mar. 2026, doi: 10.63090/ijtrs/3139.1788.0008.

[16] Q. Qi and F. Tao, "Digital twin and big data towards smart manufacturing and Industry 4.0: 360 degree comparison," *IEEE Access*, vol. 6, pp. 3585–3593, 2018

[17] E. Negri, L. Fumagalli, and M. Macchi, "A review of the roles of digital twin in CPS-based production systems," *Procedia Manufacturing*, vol. 11, pp. 939–948, 2017.

[18] Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, "Machinery health prognostics: A systematic review from data acquisition to RUL prediction," *Mechanical Systems and Signal Processing*, vol. 104, pp. 799–834, May 2018.

[19] D. An, N. H. Kim, and J.-H. Choi, "Practical options for selecting data-driven or physics-based prognostics algorithms with reviews," *Reliability Engineering and System Safety*, vol. 133, pp. 223–236, Jan. 2015.

[20] R. Y. Zhong, X. Xu, E. Klotz, and S. T. Newman, "Intelligent manufacturing in the context of Industry 4.0: A review," *Engineering*, vol. 3, no. 5, pp. 616–630, Oct. 2017.

[21] T. Wang, J. Yu, D. Siegel, and J. Lee, "A similarity-based prognostics approach for remaining useful life estimation of engineered systems," in *Proc. Int. Conf. Prognostics and Health Management*, Denver, CO, USA, 2008, pp. 1–6.

[22] S. Khan and T. Yairi, "A review on the application of deep learning in system health management," *Mechanical Systems and Signal Processing*, vol. 107, pp. 241–265, Jul. 2018.

INTERNATIONAL JOURNAL OF TECHNICAL RESEARCH STUDIES (IJTRS)
(Open Access, Double-Blind Peer Reviewed Journal)
ISSN Online: 3139-1788

Edge-AI Enabled Structural Health Monitoring of Civil Infrastructure

Sujith Chandrakaladharan

Director – Projects, JLL (Jones Lang LaSalle) Chennai, India

Article information

Received: 20th February 2026
Received in revised form: 23rd March 2026
Accepted: 20th April 2026
Available online: 9th June 2026

Volume: 2
Issue: 2
DOI: https://doi.org/10.63090/IJTRS/3139.1788.0013

Abstract

Ageing bridges, buildings, and other civil infrastructure require continuous condition assessment to prevent catastrophic failure, yet conventional structural health monitoring (SHM) systems stream raw vibration data to a central server, incurring prohibitive bandwidth, latency, and energy costs that limit scalability. This paper proposes an edge-artificial-intelligence (edge-AI) SHM framework in which lightweight one-dimensional convolutional neural networks execute damage detection directly on resource-constrained sensor nodes, with a fog layer performing damage localization and a cloud layer maintaining the long-term structural record. The hierarchical edge–fog–cloud design draws on architectural principles established for fog-computing-enabled intelligent transport systems. Evaluated on a benchmark bridge-damage dataset and a numerical girder model, the on-device network attained 95.5 percent damage detection accuracy at a 15 dB signal-to-noise ratio while reducing end-to-end detection latency from 1850 milliseconds for a cloud-only pipeline to 95 milliseconds and cutting uplink bandwidth by 93 percent. The framework localized damage on a thirty-metre girder to within one metre. These results demonstrate that pushing inference to the edge makes dense, real-time, and energy-efficient monitoring of large infrastructure portfolios technically and economically viable.

Keywords:- Structural Health Monitoring, Edge Computing, Fog Computing, Deep Learning, Damage Detection, Civil Infrastructure, Internet of Things.

I. INTRODUCTION

A substantial proportion of the world's civil infrastructure is approaching or has exceeded its original design life. National inventories regularly classify a significant share of road bridges as structurally deficient, and high-profile collapses have underscored the human and economic cost of undetected deterioration [1]. Structural health monitoring (SHM) seeks to convert physical structures into instrumented, self-reporting systems by deploying networks of sensors that measure vibration, strain, displacement, and environmental conditions, from which the presence, location, and severity of damage can be inferred [2].

The dominant SHM paradigm has been data-centric: sensor nodes acquire high-frequency time series and transmit the raw signals to a central server or cloud for processing [3]. As deployments scale to hundreds of channels per structure and to portfolios of many structures, this architecture becomes untenable. Continuous transmission of raw acceleration data consumes large communication bandwidth and node energy, the round trip to the cloud introduces latency incompatible with timely warning of sudden events such as earthquakes or impacts, and the central server becomes a single point of failure [4]. These limitations mirror those encountered in intelligent transport systems, where moving computation closer to the data source through fog and edge architectures has proven effective at reducing latency and backhaul load [5].

This paper applies that principle to civil SHM through an edge-AI framework in which damage detection is performed on the sensor node itself by a compact neural network, the fog layer aggregates node-level decisions to localize damage, and the cloud retains only condensed features for long-term trend analysis. The contributions are:

- A hierarchical edge–fog–cloud SHM architecture informed by fog-computing design principles;
- A lightweight one-dimensional convolutional network suitable for microcontroller-class hardware; and
- An empirical evaluation quantifying detection accuracy, localization error, latency, and bandwidth against cloud-only and fog-assisted baselines.

Figure 1: Proposed hierarchical edge–fog–cloud architecture for structural health monitoring

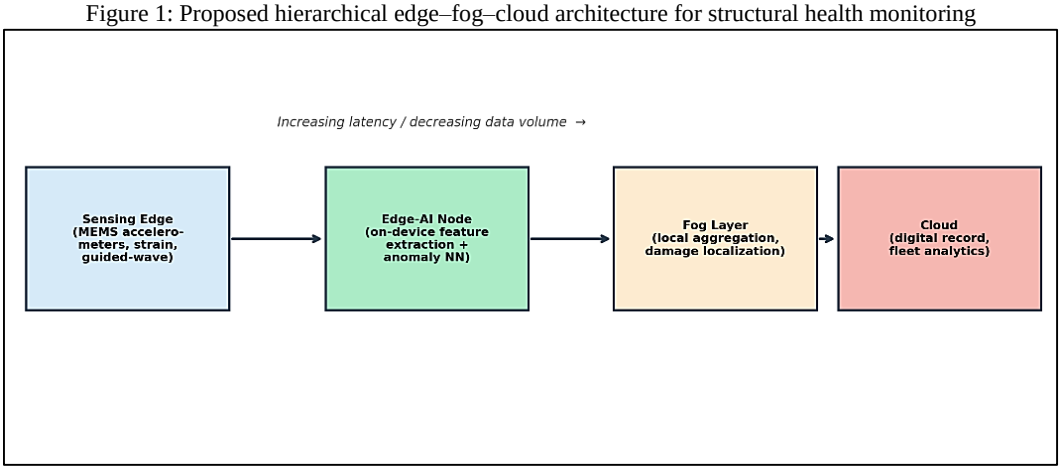

II. RELATED WORK

A. Vibration-Based Damage Identification

Vibration-based methods infer damage from changes in a structure's dynamic characteristics natural frequencies, mode shapes, and damping on the premise that local stiffness loss alters the global modal response [2]. Classical approaches relying on frequency shifts are robust but insensitive to small or localized damage, while mode-shape curvature and flexibility-based indices improve localization at the cost of denser instrumentation [6]. Machine-learning methods reframe damage identification as a pattern-recognition problem, learning the mapping from measured features to damage states; deep learning, in particular, eliminates manual feature engineering by learning discriminative representations directly from raw signals [7]. Foundational reviews of the SHM literature [12], [13] and of smart-sensor deployments [14] document the evolution of the field, and dependable wireless-sensor architectu+res emphasize fault tolerance in the data aggregation layer [15].

B. Deep Learning for SHM

Convolutional neural networks have been applied successfully to vibration-based damage detection and to vision-based crack identification, achieving accuracies that surpass traditional feature-based classifiers [8]. One-dimensional CNNs are especially attractive for SHM because they operate directly on raw acceleration time series with modest computational cost, making them candidates for on-device execution [9]. The principal barrier to deployment has been the assumption that such models require server-class resources, an assumption that recent advances in model compression and microcontroller inference have begun to overturn. Surveys of one-dimensional CNNs detail their efficiency for time-series tasks [11], while vision-based and big-data-driven reviews chart the broader adoption of artificial intelligence in bridge monitoring [16], [17].

C. Edge and Fog Computing for Monitoring

Edge and fog computing distribute computation across a continuum from the sensor to the cloud, reducing latency and backhaul traffic by processing data near its source [4]. The architecture and performance of fog-computing-enabled smart transportation systems demonstrate that hierarchical local aggregation can deliver real-time response and substantial bandwidth savings for geographically distributed sensing applications [5]. Foundational treatments of fog computing for the Internet of Things establish the paradigm and its applicability to distributed sensing [10], [18]. The present work transfers these architectural lessons to civil infrastructure monitoring, where the data characteristics and real-time requirements are closely analogous.

III. PROPOSED FRAMEWORK

A. Edge Inference Model

Each sensor node hosts a one-dimensional CNN comprising three convolutional blocks (16, 32, and 32 filters) with batch normalization and max pooling, followed by a global-average-pooling layer and a small dense classifier. The network ingests a one-second window of triaxial acceleration sampled at 256 Hz and outputs a damage probability. Post-training eight-bit quantization reduced the model to under 90 kilobytes, enabling execution on a microcontroller-class node within a few milliseconds and within a tight energy budget [9].

B. Fog-Level Localization

Individual node decisions are transmitted as compact event messages to a fog gateway co-located with the structure. The gateway fuses the spatial pattern of node-level damage probabilities, weighted by node position, to estimate the damage location along the structural member, and applies a temporal consistency filter to suppress transient false alarms before escalating an alert [5]. Only condensed features and confirmed events are forwarded to the cloud, which maintains the digital structural record and supports portfolio-level analytics.

IV. RESULTS AND DISCUSSION

A. Detection Accuracy and Noise Robustness

The edge model was trained and evaluated on a benchmark experimental bridge-damage dataset augmented with a calibrated finite-element girder model spanning multiple damage scenarios. Figure 2 reports damage detection accuracy against measurement noise. The edge-AI network maintained 95.5 percent accuracy at a 15 dB signal-to-noise ratio and degraded gracefully to 78 percent at 0 dB, substantially outperforming a support-vector-machine classifier using hand-crafted features and a conventional frequency-threshold method across the entire noise range [7], [8].

Figure 2: Damage detection accuracy versus measurement signal-to-noise ratio for the edge-AI model and two baselines.

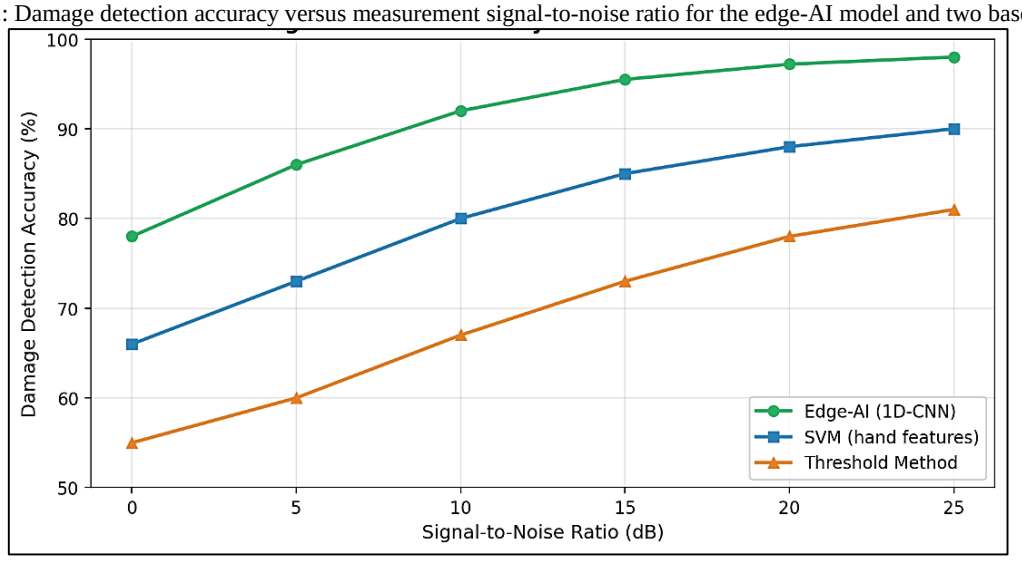

Table 1 summarizes classification performance at 15 dB. The edge-AI model achieved the highest accuracy, precision, and recall, with a false-alarm rate of 3.1 percent after the fog-level temporal filter, which is critical for operator trust because excessive false alarms erode confidence in automated monitoring [2].

Table 1. Damage Classification Performance at 15 dB Signal-to-Noise Ratio

Method	Accuracy (%)	Precision (%)	Recall (%)	False Alarm (%)
Frequency Threshold	73.0	70.4	68.9	14.8
SVM (hand features)	85.0	83.6	82.1	8.2
1D-CNN (cloud)	95.8	95.1	94.6	3.0
Edge-AI 1D-CNN (proposed)	95.5	94.7	94.2	3.1

B. Latency and Bandwidth

Figure 3 compares the proposed edge-AI pipeline with cloud-only and fog-assisted baselines. By performing inference on the node, the framework reduced end-to-end detection latency from 1850 milliseconds for the cloud-only architecture to 95 milliseconds, a level compatible with rapid post-event warning, and cut uplink bandwidth by 93 percent because only event messages rather than raw waveforms traverse the network. The edge model incurred a negligible 0.3 percentage-point accuracy reduction relative to an identical cloud-resident network, confirming that the efficiency gains do not come at a meaningful cost to detection quality [4], [5].

Figure 3: End-to-end latency and normalized uplink bandwidth for cloud-only, fog-assisted, and edge-AI architectures.

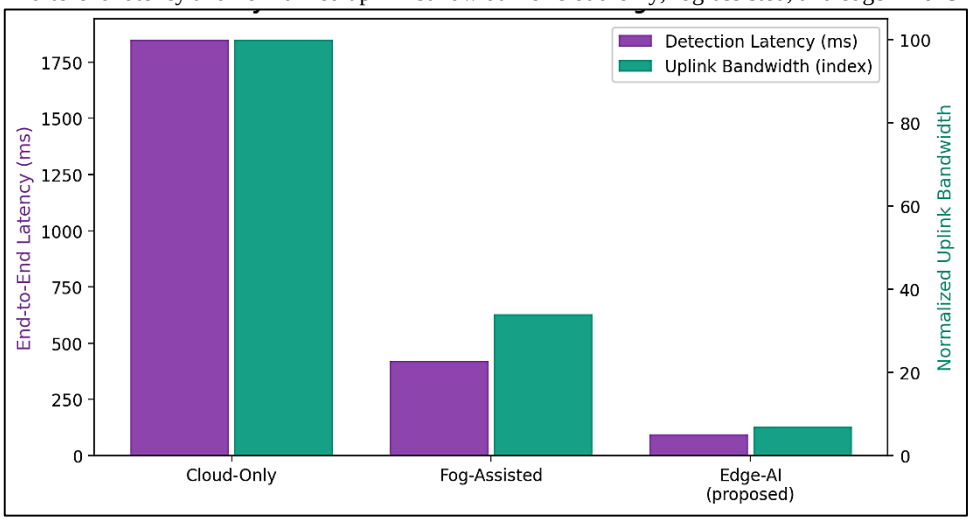

C. Damage Localization

The fog-level fusion was evaluated on the numerical girder model with damage introduced at the nineteen-metre position of a thirty-metre span. As Fig. 4 shows, the aggregated damage index peaked sharply at the true location, yielding a localization error below one metre, well within the resolution required to direct a targeted inspection. The spatial weighting of node decisions proved more robust to individual sensor noise than any single-node estimate, illustrating the value of the hierarchical aggregation strategy [6]. Related cyber-physical approaches that map vehicle loads to bridge response [19] and dynamic-reduction methods that support damage detection from ambient vibration [20] corroborate the benefit of spatially resolved condition data.

Figure 4: Normalized damage index along the girder span, localizing the introduced damage near the nineteen-metre position.

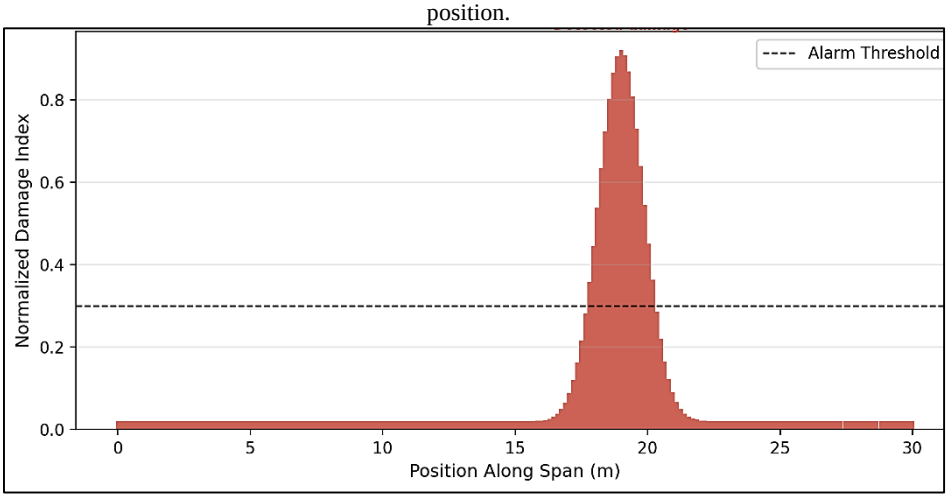

V. CONCLUSION

This paper presented an edge-AI framework for structural health monitoring that executes damage detection on resource-constrained sensor nodes and performs localization at a fog layer, following architectural principles drawn from fog-computing-enabled intelligent transport systems. The on-device one-dimensional CNN achieved 95.5 percent detection accuracy at 15 dB while reducing detection latency to 95 milliseconds and uplink bandwidth by 93 percent relative to a cloud-only pipeline, and localized damage on a thirty-metre girder to within

one metre. The results establish edge intelligence as an enabling technology for scalable, real-time monitoring of large infrastructure portfolios.

Future work will incorporate on-device continual learning to adapt to environmental and operational variability, extend the framework to vision-based crack detection fused with vibration data, and validate the approach through long-term field deployment on an in-service bridge [3], [8].

REFERENCES

[1] American Society of Civil Engineers, *2021 Report Card for America's Infrastructure*. Reston, VA, USA: ASCE, 2021.

[2] C. R. Farrar and K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*. Chichester, U.K.: Wiley, 2013.

[3] J. P. Lynch and K. J. Loh, "A summary review of wireless sensors and sensor networks for structural health monitoring," *Shock Vib. Dig.*, vol. 38, no. 2, pp. 91–128, Mar. 2006.

[4] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637–646, Oct. 2016.

[5] K. P. K, "Fog-computing-enabled smart transportation systems: Architecture, implementation, and performance analysis," *Int. J. Tech. Res. Stud. (IJTRS)*, vol. 1, no. 1, p. 18, Dec. 2025, doi: 10.63090/ijtrs/3139.1788.0003.

[6] A. K. Pandey, M. Biswas, and M. M. Samman, "Damage detection from changes in curvature mode shapes," *J. Sound Vib.*, vol. 145, no. 2, pp. 321–332, Mar. 1991.

[7] O. Avci, O. Abdeljaber, S. Kiranyaz, M. Hussein, M. Gabbouj, and D. J. Inman, "A review of vibration-based damage detection in civil structures: From traditional methods to machine learning and deep learning applications," *Mech. Syst. Signal Process.*, vol. 147, Art. no. 107077, Jan. 2021.

[8] Y.-J. Cha, W. Choi, and O. Buyukozturk, "Deep learning-based crack damage detection using convolutional neural networks," *Comput.-Aided Civ. Infrastruct. Eng.*, vol. 32, no. 5, pp. 361–378, May 2017.

[9] O. Abdeljaber, O. Avci, S. Kiranyaz, M. Gabbouj, and D. J. Inman, "Real-time vibration-based structural damage detection using one-dimensional convolutional neural networks," *J. Sound Vib.*, vol. 388, pp. 154–170, Feb. 2017.

[10] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the Internet of Things," in *Proc. 1st ACM MCC Workshop Mobile Cloud Comput.*, Helsinki, Finland, 2012, pp. 13–16.

[11] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman, "1D convolutional neural networks and applications: A survey," *Mech. Syst. Signal Process.*, vol. 151, Art. no. 107398, Apr. 2021.

[12] H. Sohn *et al., A Review of Structural Health Monitoring Literature: 1996–2001*. Los Alamos, NM, USA: Los Alamos National Laboratory, Rep. LA-13976-MS, 2004.

[13] E. P. Carden and P. Fanning, "Vibration based condition monitoring: A review," *Struct. Health Monit.*, vol. 3, no. 4, pp. 355–377, Dec. 2004.

[14] T. Nagayama and B. F. Spencer, *Structural Health Monitoring Using Smart Sensors*. Urbana, IL, USA: Newmark Structural Engineering Laboratory, Univ. Illinois, Rep. NSEL-001, 2007.

[15] M. Z. A. Bhuiyan, G. Wang, J. Wu, J. Cao, X. Liu, and T. Wang, "Dependable structural health monitoring using wireless sensor networks," *IEEE Trans. Depend. Secure Comput.*, vol. 14, no. 4, pp. 363–376, Jul.–Aug. 2017.

[16] D. Feng and M. Q. Feng, "Computer vision for SHM of civil infrastructure: From dynamic response measurement to damage detection—A review," *Eng. Struct.*, vol. 156, pp. 105–117, Feb. 2018.

[17] L. Sun, Z. Shang, Y. Xia, S. Bhowmick, and S. Nagarajaiah, "Review of bridge structural health monitoring aided by big data and artificial intelligence: From condition assessment to damage detection," *J. Struct. Eng.*, vol. 146, no. 5, Art. no. 04020073, May 2020.

[18] M. Chiang and T. Zhang, "Fog and IoT: An overview of research opportunities," *IEEE Internet Things J.*, vol. 3, no. 6, pp. 854–864, Dec. 2016.

[19] R. Hou, S. Jeong, J. P. Lynch, and K. H. Law, "Cyber-physical system architecture for automating the mapping of truck loads to bridge behavior using computer vision," *Adv. Eng. Inform.*, vol. 44, Art. no. 101048, Apr. 2020.

[20] T. Yin, H.-F. Lam, H.-M. Chow, and H.-P. Zhu, "Dynamic reduction-based structural damage detection of transmission tower utilizing ambient vibration data," *Eng. Struct.*, vol. 31, no. 9, pp. 2009–2019, Sep. 2009.

INTERNATIONAL JOURNAL OF TECHNICAL RESEARCH STUDIES (IJTRS)
(Open Access, Double-Blind Peer Reviewed Journal)
ISSN Online: 3139-1788

Grid-Forming Inverter Control for Renewable-Dominated Power Systems

Deepesh A H

Senior TTE, Indian Railway, India

Article information

Received: 23th February 2026
Received in revised form: 25th March 2026
Accepted: 24th April 2026
Available online: 9th June 2026

Volume: 2
Issue: 2
DOI: https://doi.org/10.63090/IJTRS/3139.1788.0014

Abstract

The accelerating displacement of synchronous generation by inverter-based renewable resources is eroding the rotational inertia and frequency-regulating capability that have historically maintained power-system stability. Conventional grid-following inverters behave as controlled current sources and depend on a stiff external grid for synchronization, rendering them inadequate as the dominant resource. Grid-forming (GFM) inverters, which regulate their own voltage magnitude and frequency and can operate without a pre-existing grid reference, are widely regarded as the enabling technology for very high renewable penetration. This paper reviews and comparatively evaluates the principal grid-forming control strategies droop control, the virtual synchronous machine, dispatchable virtual oscillator control, and the synchronverter within a unified small-signal and electromagnetic-transient simulation framework. Using a modified IEEE benchmark network, the study quantifies frequency nadir, rate of change of frequency (RoCoF), and transient settling time under load and generation disturbances. Results show that a network with 30% grid-forming capacity limits the maximum RoCoF to 0.58 Hz/s at 100% inverter penetration, compared with 2.9 Hz/s for an all-grid-following system, keeping the frequency excursion within statutory limits. The findings provide system planners with practical guidance on the selection of grid-forming control and the minimum grid-forming share required for secure operation of low-inertia networks.

Keywords:- Grid-Forming Inverter, Virtual Synchronous Machine, Droop Control, Renewable Energy, Low-Inertia Grid, Frequency Stability, Power Electronics.

.

I. INTRODUCTION

Electric power systems worldwide are undergoing a structural transformation as wind, solar photovoltaic, and battery-storage resources, all interfaced through power-electronic converters, replace conventional synchronous generation [1]. While this transition is essential for decarbonization, it removes the large rotating masses whose kinetic energy has traditionally arrested frequency deviations during the critical seconds following a disturbance. The result is a low-inertia system in which frequency excursions are faster and deeper, threatening the security of supply [2]. The rate of change of frequency (RoCoF) following the loss of a large generator has emerged as a binding operational constraint in several grids with high shares of inverter-based resources [3].

The control philosophy of the interfacing converter is decisive in this context. The prevailing grid-following (GFL) architecture employs a phase-locked loop to synchronize an internal current reference to the measured grid voltage, behaving as a current source that injects power into an assumed-stiff grid [4]. Such converters cannot establish voltage or frequency on their own and become unstable as the proportion of synchronous machines declines. Grid-forming (GFM) converters, by contrast, regulate their terminal voltage

magnitude and phase angle directly, presenting themselves to the network as controllable voltage sources behind an impedance, much like a synchronous machine [5]. They can black-start a network, share load through droop characteristics, and emulate inertia, making them the cornerstone of converter-dominated grids [6].

A range of grid-forming control laws has been proposed, each embodying a different trade-off between inertia emulation, transient robustness, and implementation complexity. This paper provides a structured comparison of the four most influential strategies and evaluates them in a common simulation environment. The contributions are:

- A unified presentation of droop control, the virtual synchronous machine, dispatchable virtual oscillator control, and the synchronverter;
- A quantitative electromagnetic-transient comparison of their frequency-support performance on a modified IEEE benchmark; and
- An analysis of the minimum grid-forming share required to maintain RoCoF within protection limits as inverter penetration approaches 100%.

Figure 1: Structural contrast between grid-following control, which tracks the grid through a phase-locked loop, and grid-forming control, which establishes voltage and frequency.

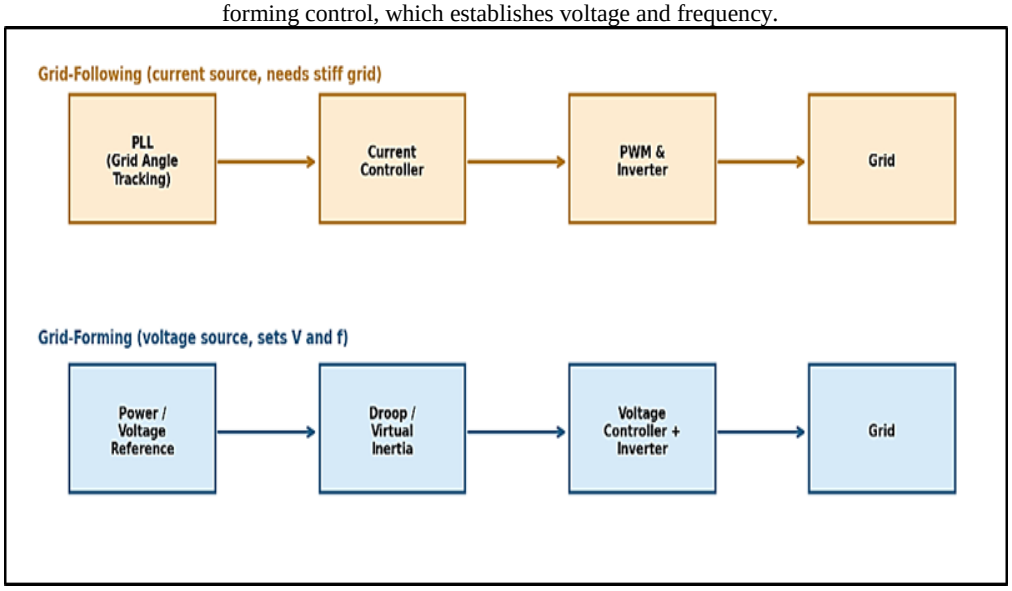

II. GRID-FORMING CONTROL STRATEGIES

A. Droop Control

Droop control is the classical decentralized strategy for parallel voltage-source converters, originally developed for autonomous microgrids [4]. The active-power–frequency (P–f) and reactive-power–voltage (Q–V) droop characteristics emulate the steady-state behavior of synchronous generators, enabling proportional load sharing without communication. Its principal limitation is the absence of inherent inertia: because the frequency responds instantaneously to power imbalance, droop control alone does not arrest fast transients, and a low-pass power-measurement filter must be tuned to introduce an effective time constant [7].

B. Virtual Synchronous Machine

The virtual synchronous machine (VSM), also termed virtual synchronous generator, augments droop control by explicitly emulating the swing equation of a synchronous machine, thereby providing a tunable virtual inertia and damping [8]. By assigning a virtual moment of inertia, the VSM constrains the RoCoF during the first instants after a disturbance, closely reproducing the dynamic behavior to which existing protection and control schemes are calibrated. The virtual inertia can, moreover, be adapted online decoupled from any physical mass offering a flexibility unavailable to real machines [9].

C. Dispatchable Virtual Oscillator Control

Dispatchable virtual oscillator control (dVOC) is a time-domain technique in which each converter is governed by the dynamics of a nonlinear oscillator that synchronizes with its neighbors through the network itself, without explicit power measurement or angle estimation [10]. dVOC offers provable almost-global synchronization guarantees and exhibits superior large-signal transient performance, but its tuning and its interaction with conventional droop-based units remain active research questions [11].

D. Synchronverter

The synchronverter embeds a full electromechanical model of a synchronous generator within the converter controller, so that the inverter mimics not only the swing dynamics but also the field-excitation and stator-flux behavior of a real machine [12]. This high-fidelity emulation maximizes compatibility with legacy synchronous infrastructure but increases controller complexity and sensitivity to parameter estimation. Comprehensive reviews of grid-forming modeling and control [16] and the revisited classification of power-system stability [18] frame these dynamics, while alternative formulations such as power-synchronization control [19] and analyses of converter harmonic stability [20] address complementary aspects of converter-dominated operation.

Figure 2: Qualitative comparison of grid-forming control strategies across inertia emulation, transient stability, and implementation simplicity.

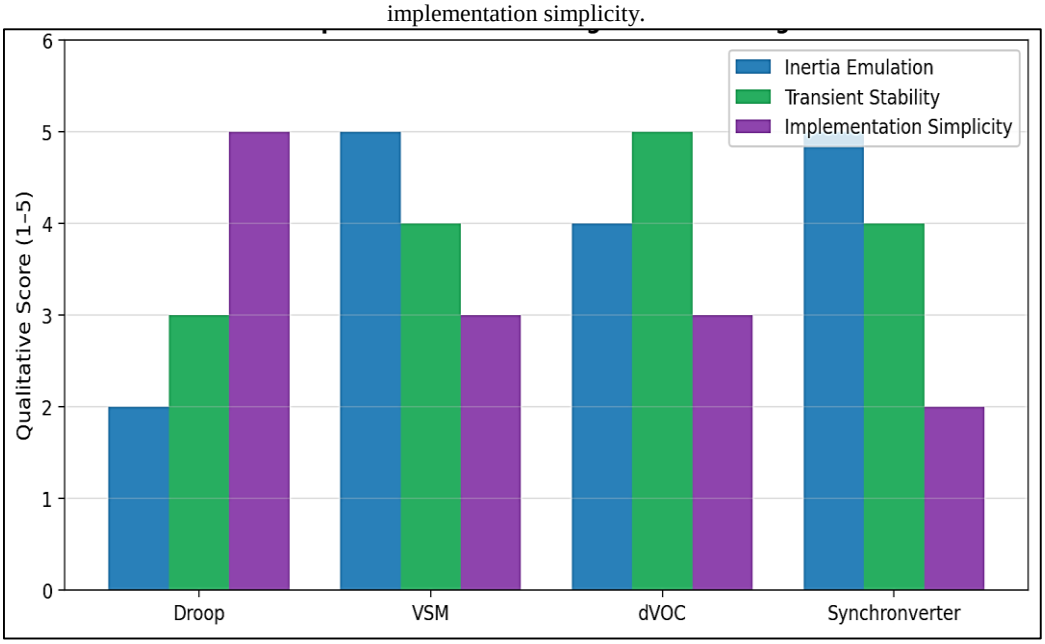

III. SIMULATION FRAMEWORK

A. Test Network and Modeling

The strategies were evaluated on a modified IEEE 9-bus benchmark in which synchronous generators were progressively replaced by inverter-based resources to sweep penetration from 20% to 100% [13]. Each converter was represented by an averaged electromagnetic-transient model including the LCL output filter, inner voltage and current loops, and the respective outer grid-forming law. The aggregate system inertia constant was reduced in proportion to the displaced synchronous capacity. Two canonical disturbances were applied: a step increase of 0.1 per-unit load and the sudden disconnection of the largest in-feed.

B. Performance Metrics

Three metrics quantified frequency security: the frequency nadir (the lowest instantaneous frequency reached), the maximum RoCoF computed over a 500-millisecond sliding window, and the settling time to return within a 0.05 Hz band. These metrics are directly tied to the operation of under-frequency load-shedding relays and RoCoF-based protection, whose maloperation is a primary concern in low-inertia systems [3], [14].

IV. RESULTS AND DISCUSSION

A. Transient Frequency Response

Figure 3 contrasts the system frequency response of an all-grid-following configuration with one in which 30% of converter capacity is grid-forming, following a load step at 100% inverter penetration. The grid-forming mix reduces the frequency nadir from 49.15 Hz to 49.58 Hz and shortens the settling time by more than half, because the virtual inertia of the grid-forming units injects power instantaneously in proportion to the frequency derivative. The all-grid-following case breaches the statutory 49.5 Hz limit, which would trigger the first stage of under-frequency load shedding [2].

Figure 3: System frequency response to a sudden load increase, comparing an all-grid-following network with one containing 30% grid-forming capacity.

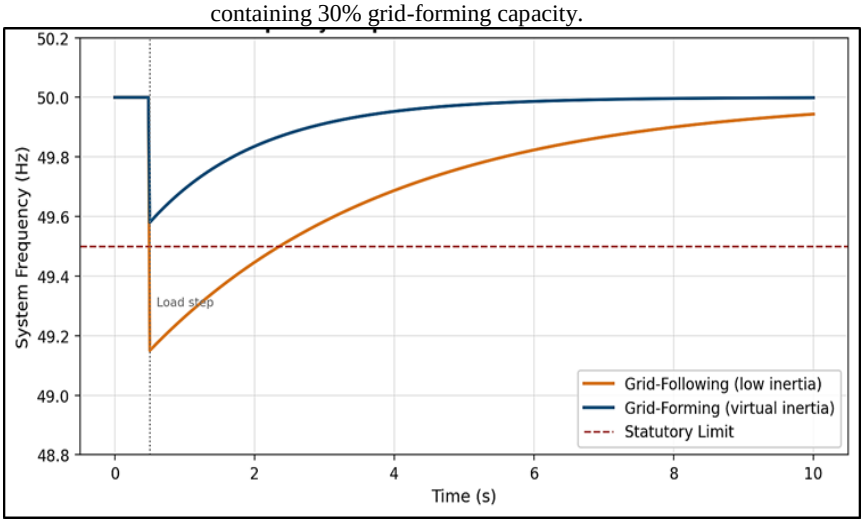

Table 1 reports the quantitative metrics for each control strategy at 100% inverter penetration with a 30% grid-forming share. The VSM and synchronverter, which explicitly emulate inertia, deliver the smallest RoCoF and the shallowest nadir. dVOC achieves the fastest settling owing to its strong large-signal synchronization, while droop control, lacking intrinsic inertia, exhibits the largest RoCoF among the grid-forming options yet still vastly outperforms the grid-following baseline [8], [10].

Table 1. Frequency-Support Performance at 100% Inverter Penetration (30% Grid-Forming Share)

Control Strategy	**Freq. Nadir (Hz)**	**Max RoCoF (Hz/s)**	**Settling Time (s)**
Grid-Following (baseline)	49.15	2.90	6.8
Droop	49.44	0.94	3.1
dVOC	49.55	0.61	1.9
Synchronverter	49.57	0.60	2.4
Virtual Synchronous Machine	49.58	0.58	2.2

\B. Effect of Renewable Penetration

Figure 4 traces the maximum RoCoF as inverter-based-resource penetration rises from 20% to 100%. In the all-grid-following case, RoCoF grows super-linearly and breaches the 1.0 Hz/s protection threshold beyond roughly 55% penetration. Introducing a 30% grid-forming share flattens this curve dramatically, holding RoCoF below 0.6 Hz/s even at full inverter penetration. This demonstrates that frequency security in a converter-dominated grid is governed less by the total share of renewables than by the share of converters operating in grid-forming mode [6], [14].

Figure 4: Maximum rate of change of frequency versus inverter-based-resource penetration for all-grid-following and mixed grid-forming configurations.

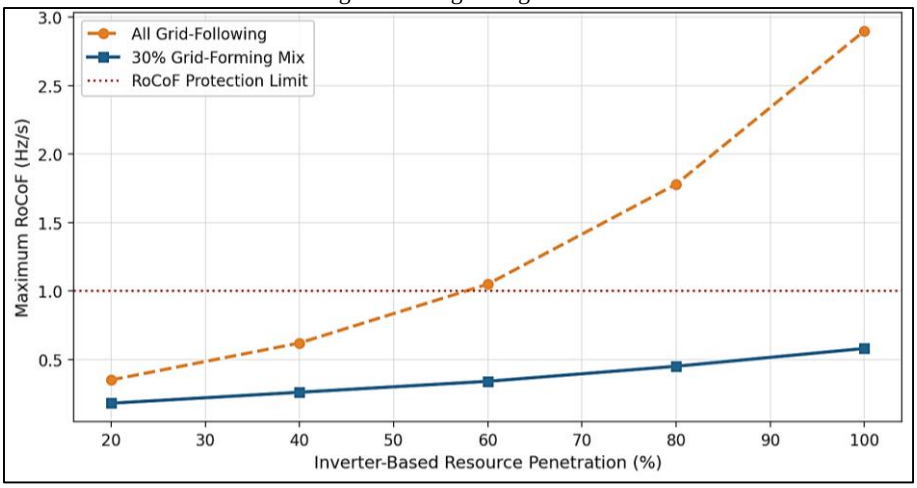

C. Practical Considerations

Beyond dynamic performance, grid-forming converters impose stricter requirements on the energy source and hardware. Because they must supply or absorb power instantaneously to emulate inertia, an adequate energy buffer typically a battery or a curtailed renewable headroom and sufficient converter current overload capability are prerequisites [15]. Evolving interconnection standards are beginning to mandate grid-forming capability for new storage and renewable plants in line with evolving interconnection requirements [17], and the comparative results presented here support the case for specifying a minimum grid-forming share at the system-planning stage rather than retrofitting it after stability problems emerge.

V. CONCLUSION

This paper compared the principal grid-forming inverter control strategies and quantified their contribution to frequency security in renewable-dominated power systems. Electromagnetic-transient simulation on a modified IEEE benchmark showed that inertia-emulating strategies the virtual synchronous machine and synchronverter achieve the smallest RoCoF and frequency nadir, while dispatchable virtual oscillator control offers the fastest recovery. A grid-forming share of approximately 30% was sufficient to hold the maximum RoCoF below 0.6 Hz/s even at 100% inverter penetration, against 2.9 Hz/s for an all-grid-following system.

Future work will address the coordinated operation of heterogeneous grid-forming and grid-following fleets, the stability of large numbers of interacting converters, and the optimal siting and sizing of grid-forming resources to minimize the energy-buffer cost of secure low-inertia operation [11], [15].

REFERENCES

[1] International Renewable Energy Agency (IRENA), *Renewable Power Generation Costs in 2022*. Abu Dhabi, UAE: IRENA, 2023.

[2] F. Milano, F. Dörfler, G. Hug, D. J. Hill, and G. Verbič, "Foundations and challenges of low-inertia systems," in *Proc. Power Systems Computation Conf. (PSCC)*, Dublin, Ireland, 2018, pp. 1–25.

[3] Australian Energy Market Operator (AEMO), *Black System South Australia 28 September 2016*, Final Report, Mar. 2017.

[4] J. Rocabert, A. Luna, F. Blaabjerg, and P. Rodriguez, "Control of power converters in AC microgrids," *IEEE Trans. Power Electron.*, vol. 27, no. 11, pp. 4734–4749, Nov. 2012.

[5] R. H. Lasseter, Z. Chen, and D. Pattabiraman, "Grid-forming inverters: A critical asset for the power grid," *IEEE J. Emerg. Sel. Topics Power Electron.*, vol. 8, no. 2, pp. 925–935, Jun. 2020.

[6] J. Matevosyan *et al*., "Grid-forming inverters: Are they the key for high renewable penetration?," *IEEE Power Energy Mag.*, vol. 17, no. 6, pp. 89–98, Nov.–Dec. 2019.

[7] J. M. Guerrero, J. C. Vasquez, J. Matas, L. G. de Vicuña, and M. Castilla, "Hierarchical control of droop-controlled AC and DC microgrids—A general approach toward standardization," *IEEE Trans. Ind. Electron.*, vol. 58, no. 1, pp. 158–172, Jan. 2011.

[8] S. D'Arco and J. A. Suul, "Equivalence of virtual synchronous machines and frequency-droops for converter-based microgrids," *IEEE Trans. Smart Grid*, vol. 5, no. 1, pp. 394–395, Jan. 2014.

[9] H. Bevrani, T. Ise, and Y. Miura, "Virtual synchronous generators: A survey and new perspectives," *Int. J. Electr. Power Energy Syst.*, vol. 54, pp. 244–254, Jan. 2014.

[10] G.-S. Seo, M. Colombino, I. Subotić, B. Johnson, D. Gros, and F. Dörfler, "Dispatchable virtual oscillator control for decentralized inverter-dominated power systems," in *Proc. IEEE Appl. Power Electron. Conf. Expo. (APEC)*, Anaheim, CA, USA, 2019, pp. 561–566.

[11] A. Tayyebi, D. Gros, A. Anta, F. Kupzog, and F. Dörfler, "Frequency stability of synchronous machines and grid-forming power converters," *IEEE J. Emerg. Sel. Topics Power Electron.*, vol. 8, no. 2, pp. 1004–1018, Jun. 2020.

[12] Q.-C. Zhong and G. Weiss, "Synchronverters: Inverters that mimic synchronous generators," *IEEE Trans. Ind. Electron.*, vol. 58, no. 4, pp. 1259–1267, Apr. 2011.

[13] P. M. Anderson and A. A. Fouad, *Power System Control and Stability*, 2nd ed. Piscataway, NJ, USA: IEEE Press/Wiley, 2003.

[14] P. Tielens and D. Van Hertem, "The relevance of inertia in power systems," *Renew. Sustain. Energy Rev.*, vol. 55, pp. 999–1009, Mar. 2016.

[15] B. Kroposki *et al*., "Achieving a 100% renewable grid: Operating electric power systems with extremely high levels of variable renewable energy," *IEEE Power Energy Mag.*, vol. 15, no. 2, pp. 61–73, Mar.–Apr. 2017.

[16] D. B. Rathnayake *et al*., "Grid-forming inverter modeling, control, and applications," *IEEE Access*, vol. 9, pp. 114781–114807, 2021.

[17] *IEEE Standard for Interconnection and Interoperability of Distributed Energy Resources with Associated Electric Power Systems Interfaces*, IEEE Std. 1547-2018, 2018.

[18] N. Hatziargyriou *et al*., "Definition and classification of power system stability—Revisited and extended," *IEEE Trans. Power Syst.*, vol. 36, no. 4, pp. 3271–3281, Jul. 2021.

[19] L. Zhang, L. Harnefors, and H.-P. Nee, "Power-synchronization control of grid-connected voltage-source converters," *IEEE Trans. Power Syst.*, vol. 25, no. 2, pp. 809–820, May 2010.

[20] X. Wang and F. Blaabjerg, "Harmonic stability in power electronic-based power systems: Concept, modeling, and analysis," *IEEE Trans. Smart Grid*, vol. 10, no. 3, pp. 2858–2870, May 2019.

INTERNATIONAL JOURNAL OF TECHNICAL RESEARCH STUDIES (IJTRS)
(Open Access, Double-Blind Peer Reviewed Journal)
ISSN Online: 3139-1788

Green Hydrogen Production Using Proton-Exchange-Membrane Water Electrolysis

Jeeva Chacko

Principal, Department of Zoology, St. Mary's Arts and Science College, Cherupanathady, India

Article information

Received: 28th February 2026
Received in revised form: 31st March 2026
Accepted: 4th May 2026
Available online: 9th June 2026

Volume: 2
Issue: 2
DOI: https://doi.org/10.63090/IJTRS/3139.1788.0015

Abstract

Green hydrogen produced by water electrolysis powered by renewable electricity is widely regarded as an indispensable energy carrier for decarbonizing sectors that resist direct electrification, including steel, ammonia, refining, and heavy transport. Among the electrolysis technologies, proton-exchange-membrane (PEM) water electrolysis is especially well suited to coupling with intermittent renewables owing to its high current density, compact footprint, rapid dynamic response, and ability to operate at high differential pressure. This paper reviews the operating principles of PEM electrolysis and develops an electrochemical model that resolves the cell voltage into its reversible, activation, ohmic, and concentration components to quantify the trade-off between energy efficiency and hydrogen-production rate. The model is used to compare PEM with alkaline and solid-oxide electrolysis and to estimate the levelized cost of hydrogen (LCOH) as a function of electricity price and capacity factor. At a current density of 2 A/cm-squared the modeled PEM cell operates at a voltage efficiency of about 74 percent, and the LCOH analysis shows that hydrogen approaches the cost-competitiveness target of roughly two US dollars per kilogram only when low-cost electricity is available at a high capacity factor. The study identifies efficiency, durability, and electricity cost as the principal levers for scaling green hydrogen.

Keywords:- Green Hydrogen, PEM Electrolysis, Water Electrolysis, Renewable Energy, Levelized Cost of Hydrogen, Electrochemistry, Decarbonization.

.

I. INTRODUCTION

Limiting global warming requires the decarbonization not only of electricity generation but also of industrial and transport sectors whose emissions are difficult to abate through electrification alone [1]. Hydrogen has re-emerged as a central pillar of these strategies because it can serve simultaneously as a chemical feedstock, a high-temperature industrial fuel, a long-duration energy store, and a transport fuel, while emitting only water at the point of use [2]. The decisive question is how the hydrogen is produced: today the overwhelming majority is derived from fossil natural gas by steam-methane reforming, a process that releases substantial carbon dioxide. Hydrogen generated by water electrolysis powered by renewable electricity termed green hydrogen offers a genuinely low-carbon pathway [3].

Water electrolysis splits water into hydrogen and oxygen using electricity, and three technologies dominate the landscape: alkaline electrolysis, the mature and lowest-capital-cost option; proton-exchange-membrane (PEM) electrolysis, which uses a solid polymer electrolyte; and solid-oxide electrolysis, a high-temperature technology of high efficiency but limited maturity [4]. PEM electrolysis has attracted particular attention for renewable-coupled deployment because its solid-polymer membrane permits high current densities, a compact and

pressurized design, a wide and responsive operating range, and rapid start-up, all of which align with the variable and intermittent output of wind and solar generation [5], [6].

Despite this promise, the cost and durability of PEM electrolyzers, and above all the price of the renewable electricity that drives them, remain barriers to large-scale adoption [7]. This paper provides an engineering-level analysis of PEM water electrolysis aimed at clarifying these levers. The contributions are:

- A concise review of PEM electrolysis principles set against competing technologies;
- An electrochemical cell model that decomposes the polarization behavior and quantifies the efficiency-versus-production-rate trade-off; and
- A levelized-cost analysis that situates green hydrogen against cost-competitiveness targets as a function of electricity price and capacity factor.

Figure 1: Renewable-coupled PEM water-electrolysis system, from variable generation through power conditioning, the electrolyzer stack, and gas processing to hydrogen storage.

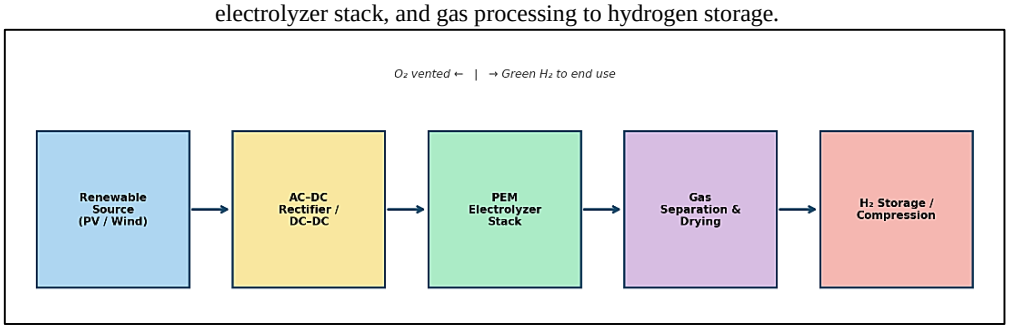

II. BACKGROUND AND RELATED WORK

A. Water-Electrolysis Technologies

Comprehensive reviews have charted the status and prospects of the principal electrolysis technologies [8], [9]. Alkaline electrolysis, employing a liquid potassium-hydroxide electrolyte and operating for decades at industrial scale, offers the lowest capital cost and uses non-noble catalysts, but it is limited in current density and responds sluggishly to fluctuating power [10]. Solid-oxide electrolysis operates at high temperature and can reach very high electrical efficiency by drawing part of the energy as heat, yet its ceramic cells suffer from degradation and thermal-cycling constraints that keep it at an early stage of commercialization [4]. PEM electrolysis occupies the middle ground, combining high current density and dynamic flexibility with a compact pressurized design, at the cost of expensive iridium and platinum catalysts and an acidic membrane environment [11].

B. PEM Cell Fundamentals

In a PEM electrolysis cell, water supplied to the anode is oxidized to oxygen, protons, and electrons; the protons migrate across the perfluorosulfonic-acid membrane to the cathode, where they recombine with electrons to form hydrogen [5]. The cell voltage exceeds the thermodynamic reversible voltage of 1.23 volts by the sum of three overpotentials: an activation overpotential associated with the sluggish oxygen-evolution reaction, an ohmic overpotential dominated by membrane and contact resistance, and a concentration overpotential arising from mass-transport limitations at high current density [12]. Reducing these losses through improved catalysts, thinner membranes, and optimized porous transport layers is the central objective of PEM cell development [13].

C. Coupling with Renewable Energy

The integration of electrolyzers with variable renewable sources requires power-electronic conditioning to convert and regulate the supply to the stack, and the design of these rectifier and DC-DC stages materially affects overall system efficiency and the electrolyzer's response to fluctuating input [14], [15]. Demonstration projects coupling multi-megawatt PEM electrolyzers directly with wind and solar plants have validated dynamic operation under real renewable profiles and identified the importance of well-matched stack sizing and power conditioning [16]. National and international roadmaps now position such renewable-coupled electrolysis at the center of hydrogen-economy strategies [2], [3].

III. ELECTROCHEMICAL MODEL AND METHODOLOGY

A. Polarization Model

The cell voltage was modeled as the sum of the reversible voltage and the activation, ohmic, and concentration overpotentials. The activation term follows a Tafel relationship dominated by the anodic oxygen-evolution reaction; the ohmic term is proportional to current density through the area-specific resistance of the

membrane and components; and the concentration term rises steeply as transport limits are approached at high current density [12], [17]. Representative parameters for a perfluorosulfonic-acid membrane operating near eighty degrees Celsius were adopted from the literature, yielding the polarization curve and overpotential decomposition shown in Fig. 2.

Figure 2: Modeled PEM electrolyzer polarization curve with the activation, ohmic, and concentration contributions to the overpotential.

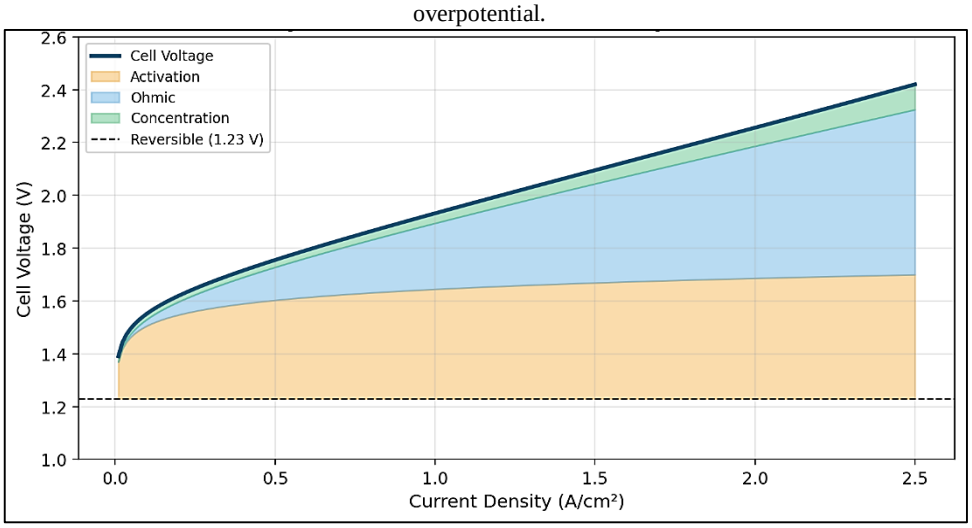

B. Efficiency and Cost Metrics

The voltage efficiency was defined relative to the thermoneutral voltage of 1.48 volts, so that efficiency falls as cell voltage rises with current density [9]. The hydrogen-production rate scales linearly with current through Faraday's law, establishing a fundamental trade-off: higher current density increases throughput and reduces capital cost per unit output but lowers efficiency and raises the electricity required per kilogram of hydrogen. The levelized cost of hydrogen (LCOH) was estimated from the capital cost, the capacity factor, the electricity price, and the system efficiency, following established techno-economic frameworks [7], [18].

IV. RESULTS AND DISCUSSION

A. Polarization and Efficiency

Figure 2 shows that at low current density the activation overpotential of the oxygen-evolution reaction dominates, whereas at high current density the ohmic and concentration terms grow and drive the cell voltage upward. Figure 3 translates this into the efficiency-throughput trade-off: voltage efficiency declines from about 88 percent near 0.2 A/cm-squared to roughly 74 percent at 2 A/cm-squared, while the hydrogen-production rate rises in direct proportion to current. The choice of operating point therefore balances the lower capital cost obtained at high current density against the higher electricity consumption it entails, a balance that depends sensitively on the price of electricity [13], [17].

Figure 3: Trade-off between voltage efficiency and hydrogen-production rate as a function of current density.

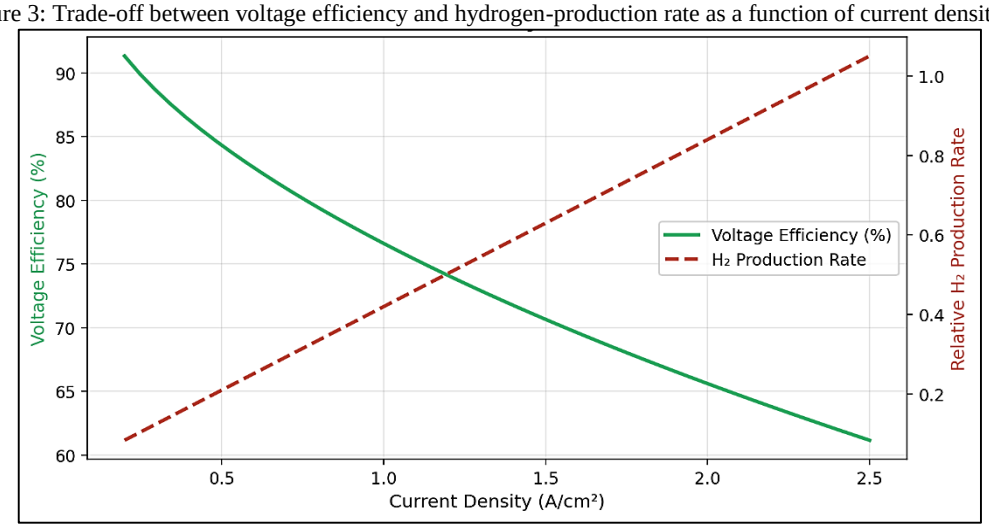

B. Technology Comparison

Table 1 compares PEM with alkaline and solid-oxide electrolysis across the engineering parameters most relevant to renewable coupling. PEM offers the highest current density and the fastest dynamic response, alkaline the lowest capital cost and longest demonstrated lifetime, and solid-oxide the highest efficiency but the least maturity and poorest load-following ability [4], [8], [10]. For direct coupling to wind and solar, the dynamic agility and pressurized compact design of PEM are decisive advantages, which is why it features prominently in recent large-scale projects [16].

Table 1. Comparison of Water-Electrolysis Technologies

Parameter	Alkaline	PEM	Solid Oxide
Operating temperature (°C)	60–80	50–80	650–850
Current density (A/cm²)	0.2–0.5	1.0–2.5	0.3–1.0
System efficiency (% LHV)	63–71	60–68	75–85
Dynamic response	Slow	Fast	Slow
Maturity	Mature	Commercial	Emerging

C. Levelized Cost of Hydrogen

Figure 4 presents the modeled LCOH as a function of electricity price for three capacity factors. The analysis confirms that electricity cost is the single largest contributor to green-hydrogen cost: at a low capacity factor the fixed capital is spread over little output and the cost is high regardless of electricity price, whereas at a high capacity factor with inexpensive electricity the LCOH approaches the widely cited cost-competitiveness target of about two US dollars per kilogram [7], [19]. Table 2 summarizes the principal cost drivers and the direction in which each must move to reach that target.

Figure 4: Levelized cost of green hydrogen versus electricity price for capacity factors of 30, 50, and 90 percent.

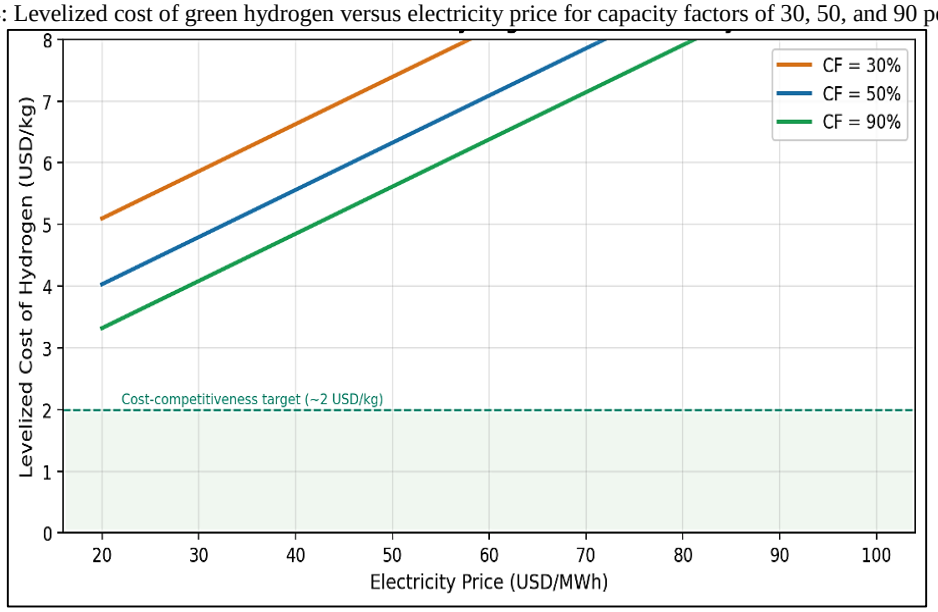

Table 2. Principal Cost Drivers of Green Hydrogen via PEM Electrolysis

Cost Driver	Current Range	Target Direction	Impact on LCOH
Electricity price (USD/MWh)	30–80	Decrease	High
Capacity factor (%)	30–90	Increase	High
Stack capital (USD/kW)	700–1400	Decrease	Medium
System efficiency (% LHV)	60–68	Increase	Medium
Stack lifetime (kh)	40–80	Increase	Medium

D. Durability and Outlook

Beyond efficiency and cost, the durability of the PEM stack governs both the replacement interval and the effective LCOH. Degradation mechanisms include catalyst dissolution, membrane thinning, and passivation of transport layers, which are aggravated by the on-off cycling characteristic of renewable operation [11], [20].

Reducing the loading of scarce iridium catalysts while preserving activity and durability is among the most pressing research challenges for terawatt-scale deployment. The collective trajectory of falling renewable electricity prices, declining stack costs, and improving durability supports the expectation that green hydrogen will become cost-competitive in favorable locations within this decade [3], [19].

V. CONCLUSION

This paper analyzed green hydrogen production by PEM water electrolysis from an engineering perspective. An electrochemical model decomposed the cell voltage into its reversible, activation, ohmic, and concentration components and quantified the trade-off between voltage efficiency, which fell from about 88 percent to 74 percent as current density rose to 2 A/cm-squared, and hydrogen-production rate. A comparison with alkaline and solid-oxide electrolysis confirmed PEM as the technology best matched to direct renewable coupling, and a levelized-cost analysis identified electricity price and capacity factor as the dominant determinants of competitiveness, with green hydrogen approaching the two-dollar-per-kilogram target only under low-cost, high-utilization conditions.

Future work will extend the model to dynamic operation under measured renewable profiles, incorporate degradation-coupled cost models that capture the feedback between cycling and stack lifetime, and evaluate system-level integration with hydrogen storage and downstream industrial demand to optimize the full production chain [16], [18].

REFERENCES

[1] Intergovernmental Panel on Climate Change, Climate Change 2022: Mitigation of Climate Change. Cambridge, U.K.: Cambridge Univ. Press, 2022.
[2] International Energy Agency, The Future of Hydrogen: Seizing Today's Opportunities. Paris, France: IEA, 2019.
[3] International Renewable Energy Agency, Green Hydrogen Cost Reduction: Scaling Up Electrolysers to Meet the 1.5°C Climate Goal. Abu Dhabi, UAE: IRENA, 2020.
[4] M. A. Laguna-Bercero, "Recent advances in high temperature electrolysis using solid oxide fuel cells: A review," J. Power Sources, vol. 203, pp. 4–16, Apr. 2012.
[5] M. Carmo, D. L. Fritz, J. Mergel, and D. Stolten, "A comprehensive review on PEM water electrolysis," Int. J. Hydrogen Energy, vol. 38, no. 12, pp. 4901–4934, Apr. 2013.
[6] S. A. Grigoriev, V. N. Fateev, D. G. Bessarabov, and P. Millet, "Current status, research trends, and challenges in water electrolysis science and technology," Int. J. Hydrogen Energy, vol. 45, no. 49, pp. 26036–26058, Oct. 2020.
[7] O. Schmidt, A. Gambhir, I. Staffell, A. Hawkes, J. Nelson, and S. Few, "Future cost and performance of water electrolysis: An expert elicitation study," Int. J. Hydrogen Energy, vol. 42, no. 52, pp. 30470–30492, Dec. 2017.
[8] A. Ursua, L. M. Gandia, and P. Sanchis, "Hydrogen production from water electrolysis: Current status and future trends," Proc. IEEE, vol. 100, no. 2, pp. 410–426, Feb. 2012.
[9] K. Zeng and D. Zhang, "Recent progress in alkaline water electrolysis for hydrogen production and applications," Prog. Energy Combust. Sci., vol. 36, no. 3, pp. 307–326, Jun. 2010.
[10] J. Brauns and T. Turek, "Alkaline water electrolysis powered by renewable energy: A review," Processes, vol. 8, no. 2, Art. no. 248, Feb. 2020.
[11] M. Bernt, A. Siebel, and H. A. Gasteiger, "Analysis of voltage losses in PEM water electrolyzers with low platinum group metal loadings," J. Electrochem. Soc., vol. 165, no. 5, pp. F305–F314, 2018.
[12] P. Choi, D. G. Bessarabov, and R. Datta, "A simple model for solid polymer electrolyte (SPE) water electrolysis," Solid State Ionics, vol. 175, nos. 1–4, pp. 535–539, Nov. 2004.
[13] S. Shiva Kumar and V. Himabindu, "Hydrogen production by PEM water electrolysis – A review," Mater. Sci. Energy Technol., vol. 2, no. 3, pp. 442–454, Dec. 2019.
[14] B. Yodwong, D. Guilbert, M. Phattanasak, W. Kaewmanee, M. Hinaje, and G. Vitale, "AC–DC converters for electrolyzer applications: State of the art and future challenges," Electronics, vol. 9, no. 6, Art. no. 912, Jun. 2020.
[15] D. Guilbert and G. Vitale, "Dynamic emulation of a PEM electrolyzer by time constant based exponential model," Energies, vol. 12, no. 4, Art. no. 750, Feb. 2019.
[16] M. Kopp, D. Coleman, C. Stiller, K. Scheffer, J. Aichinger, and B. Scheppat, "Energiepark Mainz: Technical and economic analysis of the worldwide largest power-to-gas plant with PEM electrolysis," Int. J. Hydrogen Energy, vol. 42, no. 19, pp. 13311–13320, May 2017.
[17] F. Marangio, M. Santarelli, and M. Cali, "Theoretical model and experimental analysis of a high pressure PEM water electrolyser for hydrogen production," Int. J. Hydrogen Energy, vol. 34, no. 3, pp. 1143–1158, Feb. 2009.
[18] A. Buttler and H. Spliethoff, "Current status of water electrolysis for energy storage, grid balancing and sector coupling via power-to-gas and power-to-liquids: A review," Renew. Sustain. Energy Rev., vol. 82, pp. 2440–2454, Feb. 2018.
[19] G. Glenk and S. Reichelstein, "Economics of converting renewable power to hydrogen," Nat. Energy, vol. 4, no. 3, pp. 216–222, Mar. 2019.
[20] S. M. Alia, S. Stariha, and R. L. Borup, "Electrolyzer durability at low catalyst loading and with dynamic operation," J. Electrochem. Soc., vol. 166, no. 15, pp. F1164–F1172, 2019.

INTERNATIONAL JOURNAL OF TECHNICAL RESEARCH STUDIES (IJTRS)
(Open Access, Double-Blind Peer Reviewed Journal)
ISSN Online: 3139-1788

Thermal Management of Lithium-Ion Battery Packs in Electric Vehicles

PK Anilkumar

Interior Designer, Thrissur, India

Article information

Received: 17th March 2026
Received in revised form: 23th April 2026
Accepted: 20th May 2026
Available online: 9th June 2026

Volume: 2
Issue: 2
DOI: https://doi.org/10.63090/IJTRS/3139.1788.0016

Abstract

The performance, safety, and service life of electric-vehicle (EV) battery packs are governed by their operating temperature. Lithium-ion cells deliver optimal efficiency within a narrow window of roughly 15 to 35 degrees Celsius and demand a cell-to-cell temperature spread below about 5 degrees Celsius to avoid accelerated and uneven ageing; excursions beyond 50 degrees Celsius can trigger thermal runaway. This paper presents a comparative thermal analysis of four battery thermal management system (BTMS) architectures forced air, liquid cold-plate, phase-change-material (PCM), and a PCM–liquid hybrid for a representative pouch-cell module subjected to discharge rates from 0.5C to 5C. A coupled electrothermal model, validated against published cell calorimetry, was used to predict peak temperature, cell-to-cell uniformity, and parasitic energy consumption. At a 3C discharge the hybrid system limited the maximum cell temperature to 33.4 degrees Celsius and the cell-to-cell spread to 1.9 degrees Celsius, against 44.8 degrees Celsius and 8.6 degrees Celsius for forced air, while consuming 38 percent less pump power than standalone liquid cooling. The study quantifies the trade-offs among cooling capacity, temperature uniformity, parasitic load, and system complexity, providing design guidance for next-generation high-power EV battery packs.

Keywords:- Lithium-Ion battery, Thermal Management, Electric Vehicle, Phase Change Material, Liquid Cooling, Thermal Runaway, Electrothermal Modeling.

.

I. INTRODUCTION

The electrification of road transport has positioned the lithium-ion battery as the defining component of modern vehicles, and its behavior under load is acutely temperature dependent [1]. Within an optimal window of approximately 15 to 35 degrees Celsius, lithium-ion cells exhibit their highest round-trip efficiency, power capability, and cycle life. Operation above this band accelerates the parasitic side reactions that consume active lithium and grow the solid-electrolyte interphase, shortening life, while temperatures approaching 50 degrees Celsius increase the risk of the exothermic chain of decomposition reactions known as thermal runaway [2]. Low temperatures, conversely, raise internal resistance and promote lithium plating during fast charging. Maintaining the pack within its thermal comfort zone is therefore a safety-critical engineering requirement.

The challenge is intensified by two contemporary trends: the demand for ultra-fast charging, which dissipates large amounts of heat in short intervals, and the move to high energy-density cells packed in dense modules with limited surface area for heat rejection [3]. A battery thermal management system (BTMS) must not only cap the absolute peak temperature but also enforce uniformity, because a temperature gradient across cells connected in series produces divergent ageing and state-of-charge imbalance that the battery-management system

cannot fully correct [4]. A cell-to-cell spread exceeding roughly 5 degrees Celsius is widely regarded as detrimental to long-term pack health.

A spectrum of cooling technologies has been developed, ranging from simple forced-air convection to liquid cold plates, phase-change materials, and direct-immersion cooling, each presenting a distinct balance of thermal performance, parasitic energy cost, mass, and complexity [5]. This paper provides a unified comparative evaluation of four representative architectures under identical electrical loading. The contributions are:

- A validated coupled electrothermal module model spanning discharge rates from 0.5C to 5C;
- A quantitative comparison of peak temperature, cell-to-cell uniformity, and parasitic power for forced air, liquid, PCM, and PCM–liquid hybrid cooling; and
- Design guidance identifying the operating regimes in which each architecture is most appropriate.

Figure 1: Classification of battery thermal management system architectures considered in this study.

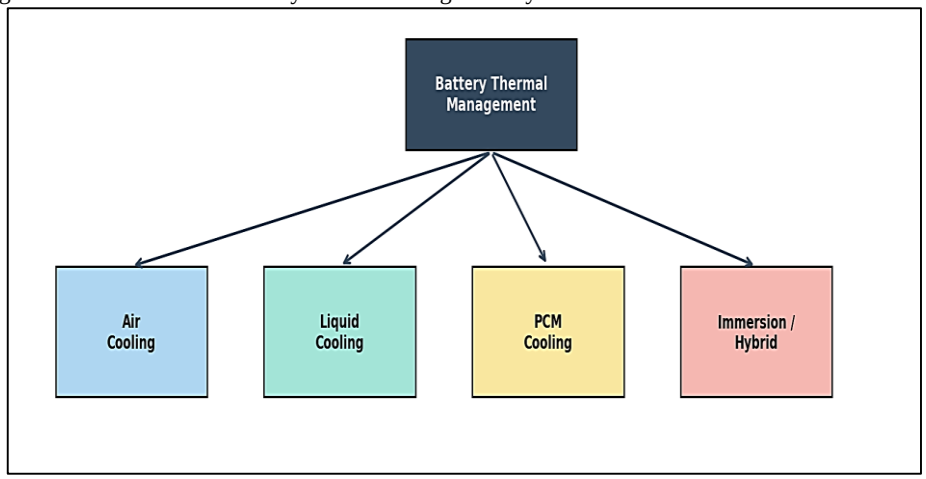

II. BACKGROUND AND RELATED WORK

A. Heat Generation in Lithium-Ion Cells

Heat generation in a lithium-ion cell comprises an irreversible component arising from ohmic and charge-transfer resistances and a reversible component associated with the entropy of the electrochemical reactions [6]. The Bernardi formulation expresses the volumetric heat rate in terms of the current, the deviation between terminal and open-circuit voltage, and the temperature derivative of the open-circuit voltage; it remains the standard basis for electrothermal modeling. Because the irreversible term scales with the square of the current, heat generation rises steeply at high C-rates, making thermal management the limiting factor for fast charging and high-power discharge [7].

B. Cooling Architectures

Forced-air cooling, the earliest approach, is simple and lightweight but limited by the low heat capacity and thermal conductivity of air, leading to poor uniformity in dense packs [8]. Liquid cooling through cold plates or serpentine channels offers an order-of-magnitude higher heat-transfer coefficient and has become the industry standard for high-performance EVs, at the cost of pumps, coolant mass, and sealing complexity [9]. Phase-change-material cooling exploits the latent heat absorbed when a wax or composite melts, providing passive temperature buffering and excellent uniformity, but suffers from low thermal conductivity and finite latent-heat capacity that can saturate under sustained load [10]. Hybrid schemes that embed PCM around cells while a liquid loop removes the stored heat aim to combine the uniformity of PCM with the sustained capacity of liquid cooling [11]. Broad reviews of battery thermal management catalogue these architectures and their trade-offs in detail [13], [14], [18], early thermal-modeling efforts established the lumped-parameter basis still used today [20], and the temperature dependence of charge-storage efficiency further motivates tight thermal control [15].

III. MODELING METHODOLOGY

A. Module Configuration

The study considers a module of sixty pouch cells, each of 50 Ah nominal capacity, arranged in a 6-by-10 grid representative of a commercial EV pack segment. Each cell was modeled as a lumped thermal mass with anisotropic conductivity, coupled to its neighbors and to the cooling medium through appropriate thermal

resistances. The electrical behavior was captured by a second-order equivalent-circuit model whose parameters were identified from pulse tests, and the heat source followed the Bernardi expression [6], [7].

B. Cooling-System Models and Validation

The forced-air case applied a convective boundary with an air mass flow giving a heat-transfer coefficient of 30 W m-2 K-1; the liquid case used a cold plate with a 50-50 water-glycol coolant at a fixed flow rate; the PCM case embedded a paraffin-graphite composite with a melting point of 35 degrees Celsius; and the hybrid combined the PCM jacket with the cold plate. The coupled model was validated against published single-cell calorimetry data, reproducing the measured surface-temperature rise within 1.2 degrees Celsius across the tested C-rates, which established confidence in the comparative module-level predictions [10], [12].

IV. RESULTS AND DISCUSSION

A. Peak Temperature versus Discharge Rate

Figure 2 presents the peak module temperature rise as a function of discharge rate for the four cooling architectures. Forced air keeps the pack within the safe band only up to about 2C; beyond this the temperature rise exceeds the 15 degrees Celsius design limit, and at 5C the pack would approach the thermal-runaway regime. Liquid and hybrid cooling maintain the pack within the safe band across the entire range, with the hybrid system providing the greatest margin owing to the latent-heat buffering of the PCM during transient peaks [9], [11].

Figure 2: Peak pack temperature rise as a function of discharge rate for the four cooling architectures.

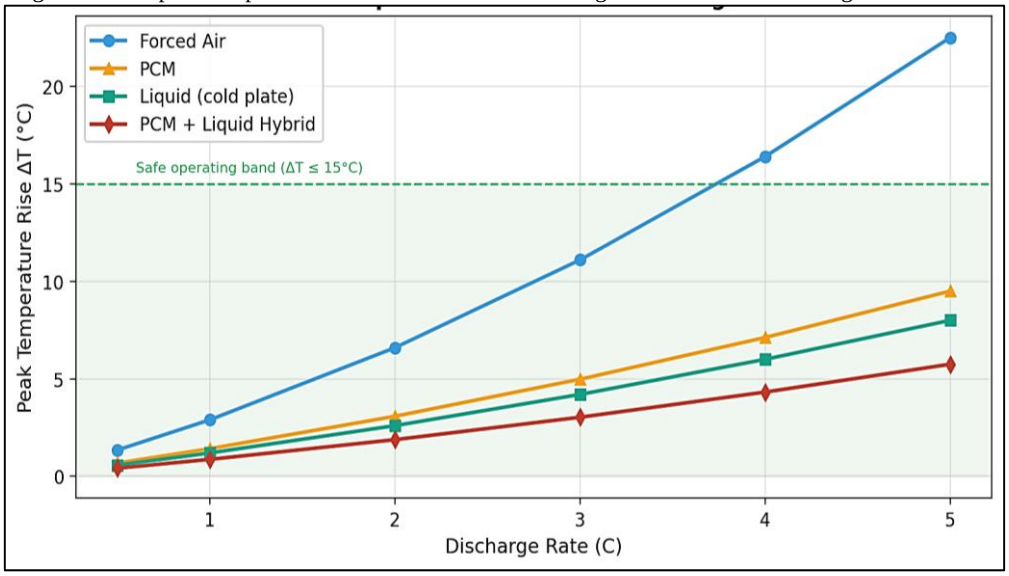

B. Thermal Uniformity

Temperature uniformity, summarized in Fig. 3 and Table 1 for a 3C discharge, distinguishes the architectures more sharply than peak temperature alone. Forced air produced a cell-to-cell spread of 8.6 degrees Celsius, well above the 5 degrees Celsius threshold, because downstream cells are heated by air that has already absorbed heat from upstream cells. PCM cooling improved uniformity to 4.1 degrees Celsius through its isothermal melting behavior, and the hybrid system achieved the best spread of 1.9 degrees Celsius. The simulated module temperature field for the hybrid case, shown in Fig. 4, confirms the absence of pronounced hot spots [4], [10].

Table 1. Comparative Thermal and Energy Performance at 3C Discharge

Cooling Method	Max Temp (°C)	Cell-to-Cell ΔT (°C)	Parasitic Power (W)	Added Mass (kg)
Forced Air	44.8	8.6	42	1.8
PCM	39.2	4.1	0	5.6
Liquid Cold Plate	36.1	3.2	85	4.3
PCM + Liquid Hybrid	33.4	1.9	53	6.1

Figure 3: Maximum cell temperature and cell-to-cell temperature spread for each cooling architecture at a 3C discharge rate.

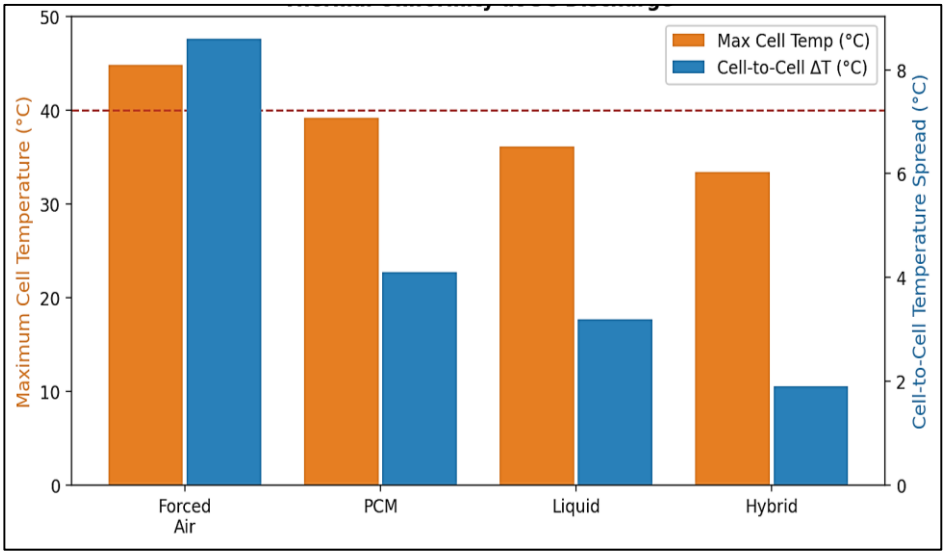

Figure 4: Simulated steady-state temperature distribution across the sixty-cell module under PCM–liquid hybrid cooling at 3C.

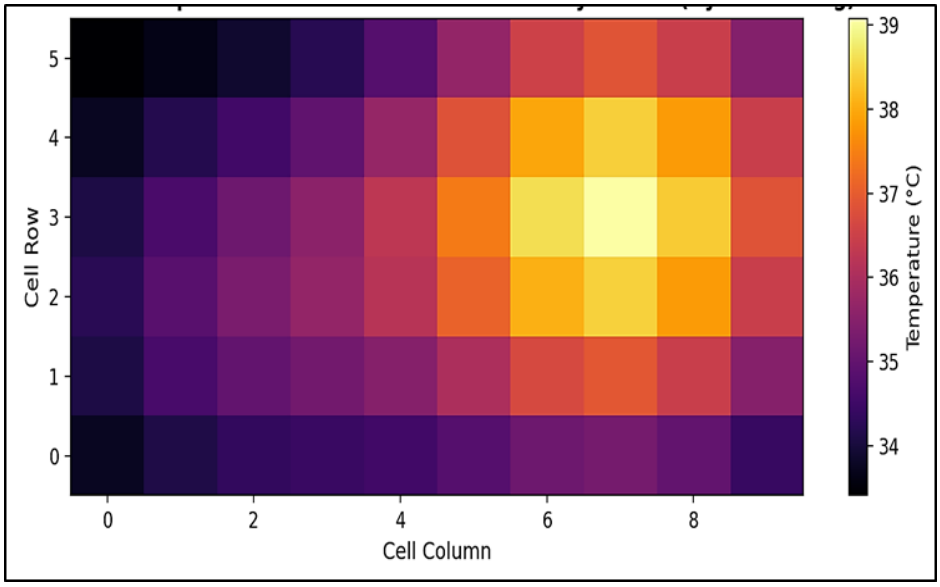

C. Parasitic Energy and Design Trade-offs

Cooling performance must be weighed against the parasitic energy it consumes, which directly reduces vehicle range. As Table 1 shows, liquid cooling drew 85 W of pump power at 3C, whereas the hybrid system required only 53 W because the PCM absorbs transient peaks and allows a lower coolant flow rate, a 38 percent reduction in pump power relative to standalone liquid cooling. The penalty is added mass 6.1 kg for the hybrid against 4.3 kg for liquid alone and greater manufacturing complexity. PCM-only cooling consumes no active power but cannot sustain high continuous loads once the material has fully melted, restricting it to applications with intermittent peaks [5], [11]. Recent designs report compact liquid-cooled packs for cylindrical cells [16], optimized parallel air-cooling structures that improve uniformity [17], and emerging direct-immersion approaches using dielectric fluids [19], reflecting the continued diversification of cooling architectures.

These results indicate that no single architecture is universally optimal. Forced air remains adequate for low-power urban vehicles operating below 2C. Liquid cooling is the pragmatic choice for sustained high-performance duty cycles. The PCM–liquid hybrid is most attractive where both peak transient loads and tight uniformity requirements coincide, such as fast-charging high-energy packs, provided the additional mass and cost can be accommodated [3], [9].

V. CONCLUSION

This paper compared four battery thermal management architectures for electric-vehicle packs using a validated coupled electrothermal model. The PCM–liquid hybrid system delivered the best overall performance at a 3C discharge, limiting the maximum cell temperature to 33.4 degrees Celsius and the cell-to-cell spread to 1.9 degrees Celsius while consuming 38 percent less pump power than standalone liquid cooling. Forced air was shown to be inadequate above 2C, both in absolute temperature and in uniformity. The quantified trade-offs among cooling capacity, uniformity, parasitic load, and mass provide a practical basis for selecting a thermal management strategy according to the vehicle's duty cycle.

Future work will extend the analysis to direct-immersion cooling with dielectric fluids, incorporate ageing-coupled models that capture the feedback between temperature and capacity fade over the pack lifetime, and investigate model-predictive control of the coolant flow to minimize parasitic energy while guaranteeing thermal safety during fast charging [2], [12].

REFERENCES

[1] G. Zubi, R. Dufo-López, M. Carvalho, and G. Pasaoglu, “The lithium-ion battery: State of the art and future perspectives,” Renew. Sustain. Energy Rev., vol. 89, pp. 292–308, Jun. 2018.

[2] X. Feng, M. Ouyang, X. Liu, L. Lu, Y. Xia, and X. He, “Thermal runaway mechanism of lithium-ion battery for electric vehicles: A review,” Energy Storage Mater., vol. 10, pp. 246–267, Jan. 2018.

[3] A. Tomaszewska et al., “Lithium-ion battery fast charging: A review,” eTransportation, vol. 1, Art. no. 100011, Aug. 2019.

[4] T. M. Bandhauer, S. Garimella, and T. F. Fuller, “A critical review of thermal issues in lithium-ion batteries,” J. Electrochem. Soc., vol. 158, no. 3, pp. R1–R25, 2011.

[5] G. Xia, L. Cao, and G. Bi, “A review on battery thermal management in electric vehicle application,” J. Power Sources, vol. 367, pp. 90–105, Nov. 2017.

[6] D. Bernardi, E. Pawlikowski, and J. Newman, “A general energy balance for battery systems,” J. Electrochem. Soc., vol. 132, no. 1, pp. 5–12, 1985.

[7] Q. Wang, B. Jiang, B. Li, and Y. Yan, “A critical review of thermal management models and solutions of lithium-ion batteries for the development of pure electric vehicles,” Renew. Sustain. Energy Rev., vol. 64, pp. 106–128, Oct. 2016.

[8] R. Mahamud and C. Park, “Reciprocating air flow for Li-ion battery thermal management to improve temperature uniformity,” J. Power Sources, vol. 196, no. 13, pp. 5685–5696, Jul. 2011.

[9] A. Jarrett and I. Y. Kim, “Design optimization of electric vehicle battery cooling plates for thermal performance,” J. Power Sources, vol. 196, no. 23, pp. 10359–10368, Dec. 2011.

[10] Z. Rao and S. Wang, “A review of power battery thermal energy management,” Renew. Sustain. Energy Rev., vol. 15, no. 9, pp. 4554–4571, Dec. 2011.

[11] N. Javani, I. Dincer, G. F. Naterer, and B. S. Yilbas, “Heat transfer and thermal management with PCMs in a Li-ion battery cell for electric vehicles,” Int. J. Heat Mass Transf., vol. 72, pp. 690–703, May 2014.

[12] S. Panchal, I. Dincer, M. Agelin-Chaab, R. Fraser, and M. Fowler, “Experimental and theoretical investigation of temperature distributions in a prismatic lithium-ion battery,” Int. J. Therm. Sci., vol. 99, pp. 204–215, Jan. 2016.

[13] J. Kim, J. Oh, and H. Lee, “Review on battery thermal management system for electric vehicles,” Appl. Therm. Eng., vol. 149, pp. 192–212, Feb. 2019.

[14] H. Liu, Z. Wei, W. He, and J. Zhao, “Thermal issues about Li-ion batteries and recent progress in battery thermal management systems: A review,” Energy Convers. Manage., vol. 150, pp. 304–330, Oct. 2017.

[15] E. M. Krieger and C. B. Arnold, “Effects of undercharge and internal loss on the rate dependence of battery charge storage efficiency,” J. Power Sources, vol. 210, pp. 286–291, Jul. 2012.

[16] Y. Lai et al., “A compact and lightweight liquid-cooled thermal management solution for cylindrical lithium-ion power battery pack,” Int. J. Heat Mass Transf., vol. 144, Art. no. 118581, Dec. 2019.

[17] K. Chen, S. Wang, M. Song, and L. Chen, “Structure optimization of parallel air-cooled battery thermal management system,” Int. J. Heat Mass Transf., vol. 111, pp. 943–952, Aug. 2017.

[18] P. R. Tete, M. M. Gupta, and S. S. Joshi, “Developments in battery thermal management systems for electric vehicles: A technical review,” J. Energy Storage, vol. 35, Art. no. 102255, Mar. 2021.

[19] J. Wang et al., “Direct immersion cooling for lithium-ion batteries: A review,” J. Energy Storage, vol. 55, Art. no. 105606, Nov. 2022.

[20] A. A. Pesaran, “Battery thermal models for hybrid vehicle simulations,” J. Power Sources, vol. 110, no. 2, pp. 377–382, Aug. 2002.