



EXCERPT FROM THE
PROCEEDINGS
OF THE
TWENTY-THIRD ANNUAL
ACQUISITION RESEARCH SYMPOSIUM AND
INNOVATION SUMMIT

VOLUME III
“ACCELERATING WARFIGHTING CAPABILITIES”

**Spectrally Efficient Peer-Peer Networking for Enhanced
Distributed Predictive Maintenance**

Published: April 30, 2026

Approved for public release; distribution is unlimited.

Prepared for the Naval Postgraduate School, Monterey, CA 93943.

Disclaimer: The views expressed are those of the author(s) and do not reflect the official policy or position of the Naval Postgraduate School, US Navy, Department of Defense, or the US government.



The research presented in this report was supported by the Acquisition Research Program, Graduate School of Defense Management at the Naval Postgraduate School.

To request defense acquisition research, please contact:

Acquisition Research Program
Department of Defense Management
Naval Postgraduate School
E: arp@nps.edu
www.acquisitionresearch.net

Copies of Symposium Proceedings and Presentations; and Acquisition Sponsored Faculty and Student Research Reports and Posters may be printed from the **NPS Defense Acquisition & Innovation Repository** at <https://dair.nps.edu/>.



ACQUISITION RESEARCH PROGRAM
DEPARTMENT OF ACQUISITION, FINANCE, AND MANPOWER
NAVAL POSTGRADUATE SCHOOL

Spectrally Efficient Peer-Peer Networking for Enhanced Distributed Predictive Maintenance

F. Dell Kronewitter, PhD—Director of Technology at Trabus Technologies (TRABUS) specializing in wireless communications, sensor fusion, and complex systems engineering. Dell has more than 20 years of experience developing embedded communications, wireless technologies, and advanced system solutions. He helps guide TRABUS's technical strategy, supports new program development, and provides engineering expertise across software, network, and hardware domains to deliver innovative, high-quality capabilities for our customers. [dellk@trabus.com]

Art Salindong, Deng—Founder and CEO of Trabus Technologies (TRABUS). As an entrepreneur and technologist, Art leads TRABUS in driving innovation across multiple industries including wireless technologies, artificial intelligence, and cybersecurity. Prior to founding TRABUS, Art served as an executive at Qualcomm, overseeing a portfolio of highly technical programs for the DoD and other federal agencies. He led development efforts in wireless device technology integration, signal processing, and software defined applications. Prior to Qualcomm, Art completed a naval career having served as a nuclear submarine officer, engineering duty officer, and government program manager. [art@trabus.com]

Abstract

Spectrally Efficient Peer-to-Peer Distributed Predictive Maintenance (SEPP-DPM) is a decentralized architecture designed for resilient, scalable system health monitoring in communication-constrained or contested environments. Unlike traditional predictive maintenance frameworks that rely on centralized data aggregation and processing, SEPP-DPM distributes both learning and inference across a network of edge nodes. Each node trains local autoencoder-based health models on sensor data and exchanges compact model updates through spectrally efficient, peer-to-peer communications rather than raw data. Coordination is achieved via the Kademlia Distributed Hash Table (DHT), enabling asynchronous and fault-tolerant model synchronization across heterogeneous nodes.

The system integrates a dual-model approach consisting of the Distributed System Health Model (DSHM) and the Distributed Communication Channel Model (DCCM), jointly capturing equipment behavior and RF channel dynamics. A digital twin simulation environment is employed to evaluate performance under non-stationary, regime-dependent degradation conditions, demonstrating that reconstruction loss from decentralized autoencoders reliably tracks system degradation and supports real-time Remaining Useful Life (RUL) estimation.

Experimental results show that SEPP-DPM maintains high diagnostic fidelity and operational awareness even under degraded connectivity, providing fleet-level health indicators without centralized coordination. This approach aligns with modern defense and industrial objectives—such as Joint All-Domain Command and Control (JADC2)—by enabling autonomous, edge-based predictive maintenance across distributed assets. SEPP-DPM represents a significant advancement in resilient machine learning architectures, coupling spectrum-efficient communications with decentralized intelligence for predictive maintenance in next-generation tactical and industrial networks.

Keywords: Predictive Maintenance, Distributed Predictive Maintenance, Resilient Condition-Based Maintenance, Autonomous Sustainment, Degraded-Communications Operations, Cross-Layer System/Channel Adaptation, Digital Twin Synchronization, Expeditionary Logistics, Tactical Edge Monitoring

Introduction

Predictive Maintenance is a maintenance strategy that uses data, sensors, and analytics to predict when equipment is likely to fail so maintenance can be performed just in time before failure occurs. Predictive maintenance is critical to ensuring the reliability and operational readiness of complex systems, particularly in naval and defense applications, reducing



downtime, lowering maintenance costs, extending asset life and improving reliability. Traditional predictive maintenance architectures rely on centralized data collection and processing, which introduces vulnerabilities in contested or communication-constrained environments. In such settings, systems must operate with limited connectivity while still maintaining awareness of system health. This work presents a decentralized approach to predictive maintenance that eliminates reliance on centralized infrastructure. By distributing both data processing and model training across edge nodes, the system enables continuous learning and anomaly detection even under degraded communication conditions. The proposed architecture integrates peer-to-peer communication with decentralized machine learning to provide a scalable and resilient solution. To learn, validate, and experiment with our system we have created a rudimentary generic digital twin which models a time data input from a set of homogeneous spatially distributed sensors each with several parameters (e.g., electrical current, vibration, etc.). Each will experience location dependent behavior qualitatively unique for each parameter. Additionally, our digital twin models the time varying communication between the sensors. While regime-dependent behavior is modeled within the digital twin to capture realistic system dynamics, the deployed SEPP-DPM system does not rely on explicit regime identification. Instead, it infers degradation directly from reconstruction error trends, enabling robust operation in uncertain and non-stationary environments. In this work, the term *digital twin* refers specifically to a system-level simulation environment used to model asset behavior and evaluate predictive maintenance algorithms under controlled conditions. This environment enables injection of regime changes, variation in sensor characteristics, and evaluation of distributed learning performance across heterogeneous nodes. The primary contributions of this work are as follows: (a) a distributed predictive maintenance architecture based on peer-to-peer learning, (b) a regime-based degradation model for evaluating system performance under non-stationary conditions, and (c) experimental validation demonstrating that model reconstruction loss correlates with system degradation and can serve as a real-time health indicator.

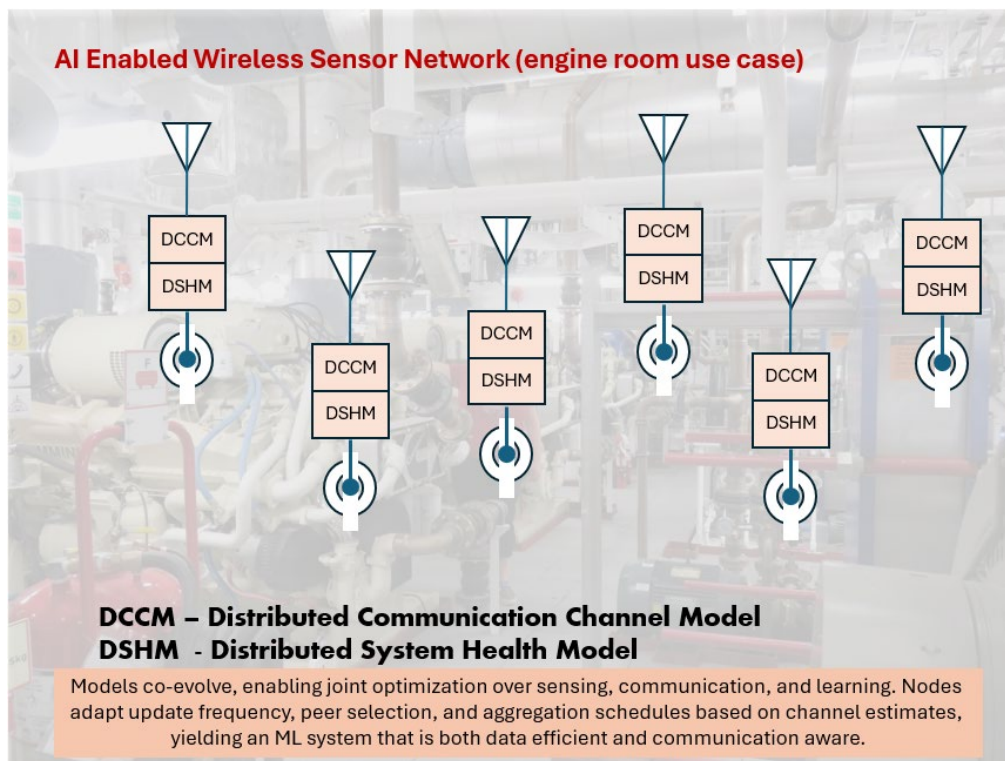


Figure 1. Marine Engine Room Case DPM



Importance to Support U.S. Navy Operations

SEPP DPM enables resilient, distributed predictive maintenance across naval platforms by pushing both analytics and coordination out to the edge, so ships, submarines, unmanned systems, and shore facilities can monitor their own health and share only compact model updates—even when bandwidth is constrained or connectivity is intermittent. This lets the Navy estimate Remaining Useful Life (RUL) for critical systems and components across a fleet, detect emerging anomalies early, and maintain fleet-level health awareness without relying on a single cloud or shore-based data center.

In U.S. Navy applications, SEPP DPM is well-suited for monitoring rotating machinery (pumps, turbines, motors), power systems, HVAC, and other mission-critical mechanical and electrical subsystems where distributed sensors already exist or can be added. Each node locally learns nominal behavior using an autoencoder and reports reconstruction loss trends as a health indicator, while the peer-to-peer network aggregates these into a global view; this supports use cases like shipboard condition-based maintenance, expeditionary logistics planning, and “fight-through” capability in contested EM environments. The same architecture extends naturally to unmanned platforms and remote sensor networks—such as distributed maritime or port surveillance systems—where heterogeneous, non-IID sensor data and degraded RF links are the norm, aligning with JADC2 priorities for edge analytics and decentralized decision-making.

Because SEPP DPM couples a spectrally efficient mesh waveform with distributed learning, it also enables cross-layer use cases where RF channel behavior and system health are monitored together—supporting detection of adversary activity, RF interference correlated with equipment stress, or environmental regimes that drive accelerated wear. In practice, this supports scenarios such as damage progression tracking after battle damage, sensor-network health monitoring across a strike group, and autonomous sustainment of forward-deployed or infrastructure-denied assets, all while avoiding single points of failure and allowing operations to continue under partitioned or intermittent connectivity.

Background

Predictive maintenance is a high-value application of machine learning, enabling early degradation detection and reducing unscheduled failures. Naval and industrial assets, however, often operate in distributed, communication-limited environments where centralized training or cloud inference is impractical. Traditional siloed sensors introduce single points of failure. Across distributed, multi-asset environments—including rotating machinery, HVAC systems, batteries, and unmanned platforms—ML-driven programs report 18%–25% reductions in sustainment cost and up to 40% reductions in downtime in mature systems (Pruckovskaja et al., 2023).

Predictive maintenance has traditionally relied on centralized analytics platforms, where sensor data is aggregated and processed to identify anomalies and predict failures. Even existing federated approaches still rely on centralized aggregation (Ahn, 2023). While effective in stable environments, this approach is not well-suited for distributed systems operating in contested or bandwidth-limited conditions. This work proposes a fully decentralized approach. Recent advances in decentralized machine learning have enabled collaborative model training without centralized coordination. Peer-to-peer learning frameworks allow nodes to exchange model updates directly, reducing dependence on centralized servers and improving resilience. These approaches are particularly relevant for distributed systems such as naval fleets, unmanned platforms, and remote sensor networks. In addition, modeling non-stationary system behavior is essential for realistic predictive maintenance. Many systems exhibit regime changes due to wear, environmental factors, or faults, requiring models that can adapt to evolving data



distributions. This work incorporates a regime-based degradation model to simulate such conditions and evaluate system performance.

Approach

The proposed SEPP-DPM architecture consists of a distributed network of edge nodes, each equipped with sensing, computation, and communication capabilities. Each node processes local sensor data and trains a lightweight machine learning model, specifically an autoencoder, to learn nominal system behavior.

Nodes participate in decentralized training through peer-to-peer communication, exchanging model updates rather than raw sensor data. This enables collaborative learning across the network without reliance on centralized infrastructure, allowing the system to maintain situational awareness even under degraded, partitioned, or intermittent connectivity.

A key challenge in such systems is enabling scalable communication and coordination across large numbers of nodes without overwhelming limited RF resources. To address this, SEPP-DPM integrates spectrally efficient peer-to-peer communication with decentralized learning algorithms, forming a unified architecture for distributed predictive maintenance in contested environments.

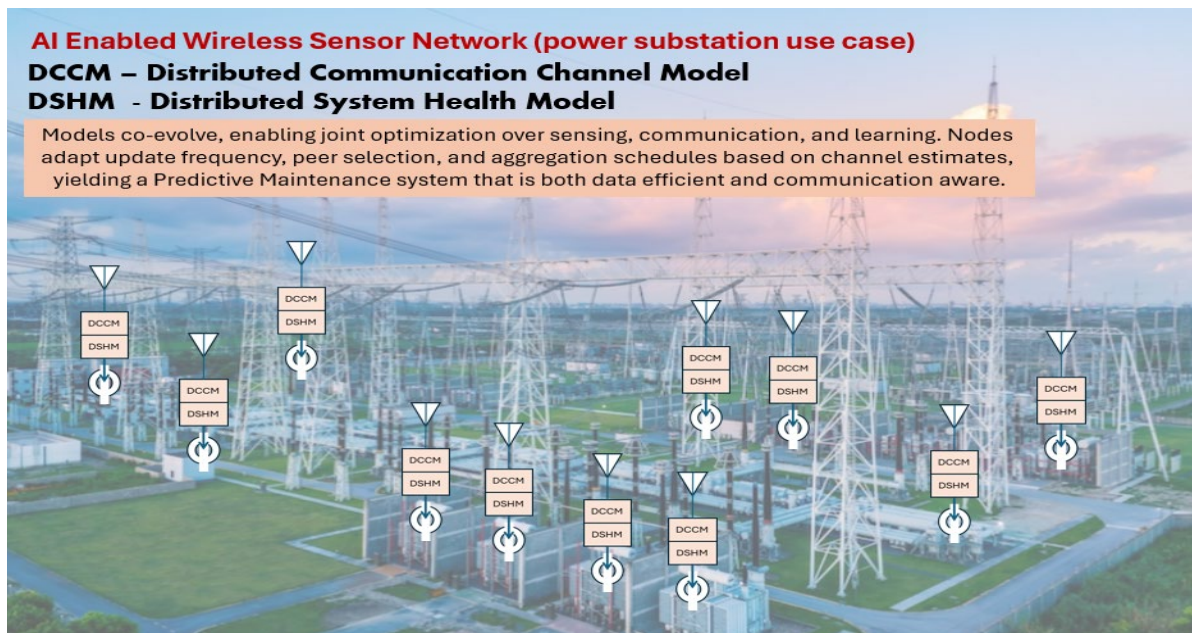


Figure 2. Power Substation Use Case DPM

A key innovation is SEPP-DPM's dual-model architecture:

- Distributed System Health Model (DSHM), an online trained representation of equipment behavior; and
- Distributed Communication Channel Model (DCCM), data channel which captures time varying RF conditions that influence model-exchange strategies.

Current SEPP Wireless network SEPP Communications Architecture

The SEPP network waveform was designed to enable decentralized wireless communication between peers in a spectrally efficient manner while optimizing the standard key



performance parameters in mesh networking, like time synchronization, data rate, and reliability in challenging RF environments.

SEPP is in a class of distributed wireless communication systems in which multiple nodes share a common spectral resource without strict orthogonalization. In these systems, nodes coordinate transmission timing, adapt transmission parameters, and manage mutual interference through local interactions rather than centralized control. This approach shifts the design paradigm from interference avoidance to interference management, enabling improved spectral utilization in dense or resource-constrained environments. Such systems are particularly relevant in infrastructure-denied or contested settings, where centralized scheduling is impractical or infeasible.

This class of communication architectures naturally induces a tightly coupled networked system in which each node's transmission decisions affect the performance of its neighbors. As a result, the network can be interpreted as a collection of interacting agents operating under shared constraints, where local observations provide only partial information about global system state. This perspective aligns closely with frameworks from distributed optimization and learning over networks, where agents iteratively adapt their behavior based on local measurements and limited peer-to-peer exchange. Rather than explicitly solving a centralized resource allocation problem, the system evolves toward efficient operating points through decentralized adaptation.

Within this context, distributed learning techniques can be leveraged to enable nodes to jointly model system behavior and anticipate future states. Each node maintains a local model, such as an autoencoder or forecasting model, trained on its own sensor observations and operational context. Instead of transmitting raw data, nodes exchange compact model updates, gradients, or embeddings with neighboring peers. These updates are aggregated across the network using decentralized protocols, allowing the system to converge toward a shared representation of normal and degraded behavior without requiring centralized data collection. This approach reduces communication overhead, improves scalability, and enhances resilience to node or link failures.

This integration of distributed communication and learning is particularly well-suited for predictive maintenance applications in geographically dispersed systems. In such settings, assets are often heterogeneous, non-identically distributed, and subject to localized environmental and operational conditions. By combining interference-aware, peer-to-peer communication with decentralized model training, the network can detect emerging degradation patterns, estimate remaining useful life, and propagate insights across nodes in near real time. The resulting system enables collaborative, data-efficient predictive maintenance in environments where connectivity, bandwidth, and infrastructure are constrained.

Current SEPP Instantiation. SEPP network development has been conducted over the last few years and been tested in the lab and in the field for a variety of use cases including first responder, tactical applications, and general consumer communication. The SEPP Module has been developed over the last several years to host our SEPP networking software. The current SEPP Module supports a variety of end user devices via Wi-Fi and Ethernet including the type of multi-modal sensors that are the subject of this research. A low-rate initial production run has been conducted and the Air Force Global Strike Command is currently evaluating the capabilities of the network.

By unifying decentralized learning with efficient communications, SEPP-DPM provides a scalable framework for distributed fault detection, anomaly prediction, and resilience in naval environments, with higher availability and ML driven decision superiority at the tactical edge.





Figure 3. SEPP Module



Figure 4. Sample SEPP Network

Decentralized Coordination via Kademia DHT

Coordination in our distributed system is implemented using a distributed hash table based on the Kademia protocol. Kademia is designed for peer-to-peer decentralized computer networks and was introduced in 2002 (Maymounkav & Mazières, 2002). Kademia is currently used in popular highly distributed applications like BitTorrent, Ethereum, and Inter-Planetary File System. The look up algorithm is particularly interesting, and we describe it here.

The Kademia lookup algorithm is a method used in decentralized peer-to-peer networks to locate nodes or data efficiently without relying on any central coordinator. Each node in the network is assigned a unique identifier, and distance between nodes is defined using the XOR metric. This distance function has useful mathematical properties—most importantly, it creates a well-structured space where nodes can progressively route queries closer to a target ID.

When a node initiates a lookup, it begins by selecting a small set (typically $\alpha = 3$) of the closest nodes it already knows about relative to the target key. It then sends parallel queries to these nodes, asking either for the value associated with the key or for nodes that are even closer to the target. Each responding node returns a list of the closest peers it knows. The querying node merges these responses into its candidate list and continues querying nodes that are progressively closer to the target in XOR space. This iterative, parallel process continues until no closer nodes can be found or the closest known nodes have all been queried.




Figure 5. Look-up Buckets

The algorithm is efficient because it reduces the search space exponentially at each step, resulting in an average lookup time of $O(\log(N))$ hops in a network of N nodes. Its use of limited parallelism (controlled by α) improves both latency and resilience to node failures. Additionally, Kademia nodes cache results along the lookup path, which helps speed up future queries and reduces overall network load.



Overall, Kademlia’s lookup mechanism can be thought of as a greedy search over a binary key space, where each step moves closer to the target based on XOR distance. This design allows the network to remain scalable, fault-tolerant, and efficient even under high churn, making it a foundational component in many distributed systems, including modern decentralized storage and learning frameworks.

Table 1. Kademlia Distributed Hash Table Features

Feature	Details
XOR Metric	The distance between two nodes or keys is calculated using the exclusive or (XOR) operation on their 160-bit identifiers, which allows for consistent routing and finding nodes
Efficient Lookup	Nodes only need to contact a small number of nodes ($O(\log(N))$) to locate data within a network of participants, significantly improving search efficiency
K-buckets	Nodes maintain routing tables (k-buckets) containing contact information for other nodes, optimizing for low-latency paths and network resilience.
Self-Healing	Kademlia automatically recovers from nodes joining or leaving, making it highly resistant to denial-of-service attacks.
Application	It is used in numerous decentralized protocols and networks, including BitTorrent, Ethereum, IPFS, and Storj.

Decentralized Learning Using Hivemind

Our software implementation leverages the Hivemind library, an open-source framework developed by the Learning@home project, to enable decentralized, peer-to-peer training of neural networks without reliance on a centralized parameter server. Each participant (peer) runs a local training loop on its own data and maintains a full copy of the model. Instead of synchronizing through a central coordinator, peers periodically form temporary groups to average their model parameters or gradients. This makes the system resilient to node failures, network instability, and heterogeneous compute environments—well suited for edge or distributed settings.

Coordination between peers is handled through a distributed hash table based on Kademlia. The DHT acts as a decentralized control plane where peers publish their availability, discover others, and organize into averaging groups. When a node is ready to synchronize, it performs a lookup for a group identifier (a key in the DHT), finds other peers with the same objective, and joins a collaborative averaging round. This process is dynamic and asynchronous, groups form and dissolve continuously as peers come and go, with no global schedule required.

Data exchange in Hivemind is optimized for efficiency and robustness. Large model updates are split into chunks for transmission, and averaging is performed using bandwidth-aware, fault-tolerant protocols that can proceed even if some peers drop out mid-round. Because each node retains a full model replica, there is no single point of failure, and training can continue even under high churn. The result is a system that trades some communication overhead for strong decentralization, scalability, and adaptability, particularly valuable in environments like distributed edge systems or unreliable networks.



Autoencoder-Based Health Monitoring

With the machinery described above we can present our SEPP-DPM system in detail. The fundamental tool we use is well known in the Predictive Maintenance space and has been shown to be an effective anomaly detection technique. The autoencoder is a neural network that learns to compress input data into a lower-dimensional latent representation and then reconstruct it as accurately as possible. It consists of an encoder that maps inputs to the latent space and a decoder that reconstructs the original data from that representation. By minimizing reconstruction error, autoencoders learn the underlying structure of the data and are often used for dimensionality reduction, anomaly detection, and feature learning. The standard autoencoder diagram is shown below.

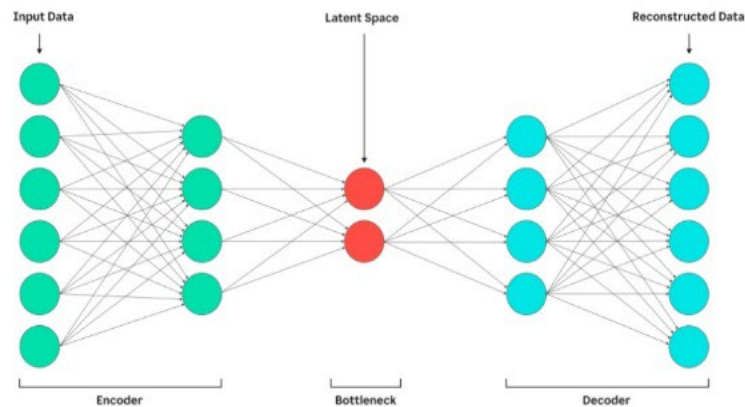


Figure 6. Standard Autoencoder

SEPP-DPM utilizes a hybrid of a local learning model and distributed (global) learning model where the DHT is used to maintain the global neural weights. The resulting decentralized architecture then contains a mix of global and local hyper-parameters as shown in Table 2. Further research will determine optimal values for these parameters. The experiments described in this initial effort used somewhat arbitrary values.

Table 2. Hyper Parameters for Our SEPP-DPM Inference System

Hyper-parameter Scope	Parameter	Description
Global	Aggregated Model Weights	The shared neural autoencoder parameters resulting from decentralized averaging across nodes. Each peer periodically contributes its local model updates, and the network converges toward a consensus model representing global system behavior.
Global	Reconstruction Loss	The network-wide aggregate of reconstruction error, computed by combining local reconstruction losses from participating nodes (e.g., mean or weighted mean). This provides a system-level health indicator that is more robust than any single node's measurement.
Global	RUL mapping Curve	Nodes A function or model that maps the global health indicator (reconstruction loss) to an estimated RUL. This may be linear, exponential, or learned offline (e.g., via your digital twin), and is shared across nodes to ensure consistent interpretation of degradation.
Global	Latent dimension	Kademia The size of the autoencoder's latent space, which defines how input data is compressed. This is fixed across all nodes to ensure compatibility of model updates and consistent feature representation across the distributed system.
Global	Aggregation frequency	The rate at which nodes participate in collaborative averaging (e.g., every N seconds or after N local updates). This controls the tradeoff between: Fresh global consensus (high frequency) Lower bandwidth / compute overhead (low frequency)
Global	Consensus smoothing factor	A parameter (e.g., EMA coefficient) used to stabilize global metrics or model updates over time. It reduces sensitivity to noise, node dropout, or transient anomalies by smoothing the evolution of global reconstruction loss or weights.
Local	Reconstruction loss	The autoencoder error computed on local sensor data at a node. This reflects how well the global model explains local observations and serves as the primary input to anomaly detection and contribution to global metrics.
Local	σ_i	A node-specific estimate of sensor noise or variability, capturing non-IID behavior across spatially distributed sensors. This can be used to: Weight contributions to global loss Normalize reconstruction error Improve robustness to noisy nodes
Local	Local optimizer state	The internal state of the training algorithm (e.g., Adam moments, gradients, learning rate state) maintained locally at each node. This is typically not shared directly, but influences the updates each node contributes during aggregation.
Local	Anomaly threshold	A node-specific threshold applied to reconstruction loss (or normalized loss) to flag anomalies. This may be: Fixed (based on baseline statistics) Adaptive (based on local variance or history)



Local	Data window size	The number of recent samples used for training and inference at a node. This defines the temporal context for: Reconstruction loss computation Smoothing / trend estimation Local model updates
--------------	------------------	--

System-Level Health Estimation and RUL

Fundamentally our SEPP-DPM system tracks the state of the system and estimates the health of the system. In particular we would like to estimate a RUL where we can estimate the time left. (This is similar to Weibull Analysis.) System RUL becomes an emergent property of a decentralized learning network, where each node contributes to and benefits from global degradation intelligence. Operationally, RUL is inferred from the autoencoder’s reconstruction error, which quantifies deviation from learned healthy system dynamics. The global reconstruction error is used as a fleet-level health indicator. Remaining useful life is estimated by projecting the smoothed error trajectory forward and computing the time required to reach a predefined failure threshold.

While the above algorithms provide a reconstruction error, what would prove more operationally relevant is the system RUL or the estimated time to failure. The global reconstruction error $E(t)$ is the system-wide degradation signal. As the fleet drifts away from nominal behavior, $E(t)$ should trend upward. By determining a threshold “fail” value E_{fail} we can provide a RUL

$$RUL(t) = \frac{E_{fail} - \widehat{E}(t)}{\dot{\widehat{E}}(t)}$$

where

$\widehat{E}(t)$ - Estimated growth rate

\hat{E} - Smoothed global reconstruction

E_{fail} – Failure threshold

to provide an estimate of time-to-failure. In the next section we will provide some graphs which show how the SRUL trends with a synthetically generated world as it degrades over time, a distributed system being monitored.

Digital Twin Simulation Environment

To investigate and validate our approach we have created a generic environment where we can create a realistic degraded environment which includes models of both system behavior and wireless channel effects. We expect these two domains to be correlated in the most challenging scenarios where adversaries or resultant RF noise accompanies system health failure. Our virtual environment must include this level of correlation fidelity. Our initial virtual test environment simply provides stochastic sensor samples to all the various sensor parameters which are generated at various geolocations. An example of a SEPP-DPM system multi-modal sensor would provide the following sensor values:



Table 3. Sample Multi-Modal Sensor

Channel Number	Sensor	Description
0	Current	Electrical current at sensor node
1	Voltage	Measured voltage at sensor node
2	Ax	Acceleration along x-axis direction
3	Ay	Acceleration along y-axis direction
4	Az	Acceleration along z-axis direction
5	RMS Acceleration	Root mean square acceleration magnitude
6	Peak acceleration	Maximum observed acceleration amplitude
7	Kurtosis/crest	Signal impulsiveness via crest factor
8	Temp internal	Internal sensor node temperature measurement
9	Temp external	External ambient temperature at sensor

Our initial characterization of the anomaly under investigation is that it is at one point and sensors can sense the anomaly to various degrees possibly depending on their distance from the anomaly. For instance, one can imagine a broken bolt holding down an assembly and the associated acceleration is highest closest to the bolt and lowest farthest away from the bolt. Furthermore, physical industrial systems typically fail in stages (or behavior regimes) so one regime could be a single bolt fails and the next regime could be when two bolts have failed. While our SEPP-DPM technique should work well under a gradual degradation scenario, these spiky failures are simpler to model and generate distinct failure modes which we can easily map into RUL estimates. With this in mind, we have created sensor data generating system which generates normal data where the standard deviation of the sensed signal increases randomly over uniform intervals but the standard deviation of the signal is monotonically decreasing with the sensor id.

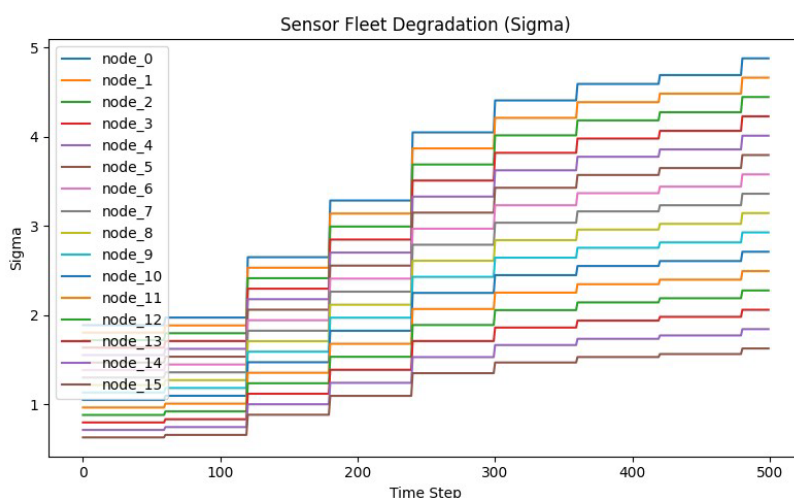


Figure 7. Standard Deviation of Sensor Data Over Time



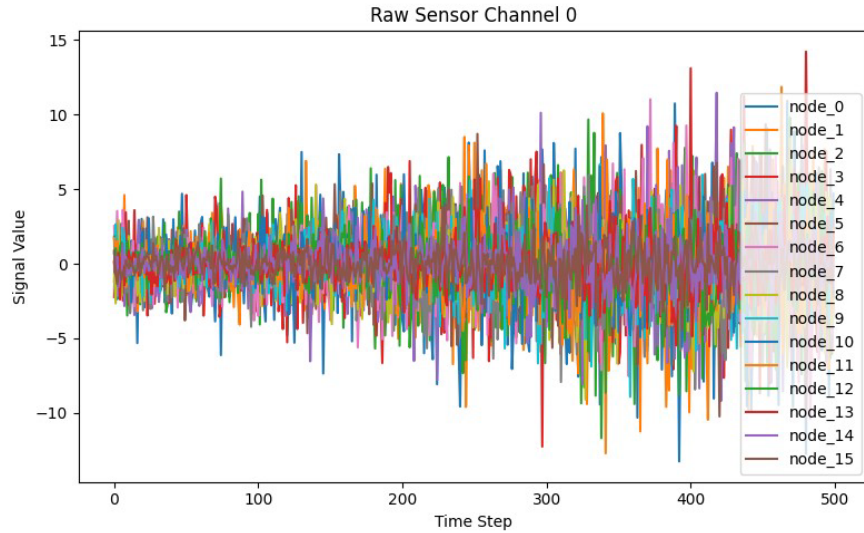


Figure 8. Raw Sensor Data

To simulate realistic system conditions, a regime-based degradation model is introduced. This model generates time-varying noise levels that represent system degradation. Regime changes occur periodically, causing shifts in the underlying data distribution. The autoencoder is trained to reconstruct input signals, and reconstruction loss is used as a proxy for system health. The combination of decentralized learning and regime-based modeling enables evaluation of the system’s ability to track gradual degradation and respond to abrupt changes. Our digital twin is Non-IID, non-stationary and partially correlated. All these design choices increase the fidelity of our simulation and allow us a more meaningful investigation of effective hyper-parameters and indeed the design as a whole.

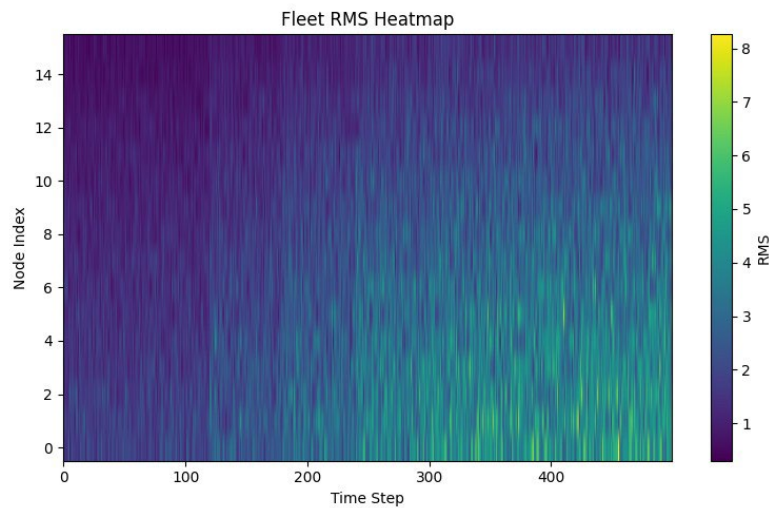


Figure 9. Heatmap of Error Magnitude Over Time

Results

Experimental results demonstrate that the proposed SEPP-DPM approach effectively captures system degradation through reconstruction loss and enables real-time estimation of RUL. Figure 10 illustrates the evolution of the RUL estimate over time for a representative simulated degradation scenario.



As the system operates in a nominal regime, reconstruction error remains low and relatively stable, resulting in large or unbounded RUL estimates. This behavior reflects the model's ability to accurately represent normal system dynamics. As degradation begins to emerge, the global reconstruction error exhibits a sustained upward trend, leading to a rapid decrease in the estimated RUL.

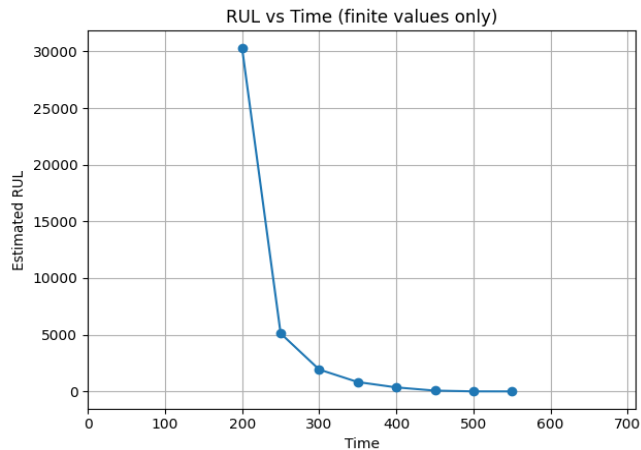


Figure 10. Estimated RUL Over Time

Notably, the transition from nominal operation to degraded behavior is characterized by a sharp drop in RUL, indicating the system's sensitivity to early deviations from learned baseline conditions. Following this transition, the RUL estimate decreases in a smooth and monotonic fashion, tracking the progression of system degradation. As the simulated system approaches failure, the RUL estimate converges toward zero, consistent with the expected behavior of a time-to-failure metric.

In scenarios involving gradual degradation, the autoencoder adapts to slowly evolving system dynamics, maintaining relatively low reconstruction error while still capturing long-term trends. In contrast, abrupt regime transitions produce transient spikes in reconstruction loss, reflecting a mismatch between the learned model and the new system state. These spikes contribute to accelerated reductions in RUL, enabling rapid detection of significant system changes.

The results further demonstrate that reconstruction loss serves as a reliable and interpretable proxy for system health. When aggregated across nodes, the global reconstruction error provides a consistent fleet-level indicator that supports robust RUL estimation. Importantly, the decentralized training framework maintains stable performance across distributed nodes, even in the presence of non-IID data and asynchronous updates, highlighting the scalability and resilience of the approach.

Discussion

The results demonstrate that decentralized predictive maintenance is a viable and effective approach for monitoring distributed systems operating in contested and communication-constrained environments. By distributing both computation and learning across edge nodes, the SEPP-DPM architecture eliminates reliance on centralized infrastructure, thereby removing critical single points of failure. This enables continued operation under degraded, partitioned, or intermittent network conditions, which are common in naval and defense applications.



A key finding of this work is that reconstruction loss derived from a decentralized autoencoder model provides a robust and interpretable system-level health indicator. As the system deviates from learned nominal behavior, reconstruction error increases in a manner that correlates with underlying degradation. When aggregated across nodes, this global reconstruction error captures both local anomalies and system-wide trends, enabling estimation of RUL without requiring explicit fault models. This is particularly advantageous in environments where failure modes are complex, poorly characterized, or evolving over time.

The decentralized learning framework plays a critical role in enabling this capability. By leveraging peer-to-peer coordination and asynchronous model averaging, the system is able to converge toward a shared representation of system behavior while accommodating non-IID data distributions across nodes. This reflects real-world conditions in which sensors experience heterogeneous operating environments, noise characteristics, and degradation patterns. The ability to learn effectively under such conditions represents a significant advantage over centralized or federated approaches that assume more uniform data distributions or require centralized aggregation.

The integration of the SEPP communication layer further strengthens the architecture by enabling scalable model exchange in bandwidth-constrained environments. Unlike traditional communication systems that rely on orthogonal channel allocation, SEPP supports simultaneous transmissions within a shared spectral resource, allowing large numbers of nodes to participate in distributed learning without excessive communication overhead. This coupling of communication efficiency with decentralized learning is essential for scaling predictive maintenance to dense sensor deployments in operational settings.

From an operational perspective, the proposed approach aligns closely with the objectives of Joint All-Domain Command and Control (JADC2), which emphasize resilient, distributed data processing and decision-making across contested environments. The SEPP-DPM system supports these objectives by enabling edge-based analytics, decentralized coordination, and robust operation under limited connectivity. Key aspects of this alignment are summarized below.

Concept	JADC2 Relevance
Decentralized Learning	Eliminates reliance on centralized infrastructure
Non-IID Distributed Sensors	Reflects heterogeneous, real-world operating conditions
Regime Detection	Enables detection of evolving system states
Edge processing	Supports operation in disconnected and contested environments

Despite these advantages, several limitations should be noted. First, the current evaluation relies on a simulated digital twin environment, which, while designed to capture non-IID and non-stationary behavior, cannot fully replicate the complexity of real-world systems. Second, the RUL estimation approach is based on trend extrapolation of reconstruction error and assumes a consistent relationship between error growth and degradation progression. In practice, this relationship may vary across systems and operating conditions, requiring calibration or adaptation. Third, communication-performance coupling, particularly the use of learned models to optimize SEPP network behavior, remains an area for future integration and validation.



Future work will focus on extending the architecture in several directions. The incorporation of more advanced model architectures, such as transformer-based or temporal sequence models, may improve the system's ability to capture complex temporal dependencies and regime transitions. Evaluation of real-world datasets and hardware deployments will be critical for validating system performance and robustness. Additionally, incorporating spatial and temporal correlations across nodes may enable detection of fault propagation and system-level failure modes that are not observable from individual sensors. Finally, tighter integration between communication and learning layers, such as adaptive model exchange strategies informed by channel conditions, may further enhance system efficiency and resilience.

Conclusion

This paper presents a decentralized predictive maintenance architecture, SEPP-DPM, that integrates peer-to-peer communication, distributed machine learning, and regime-based system modeling to enable robust operation in contested and communication-constrained environments. By distributing both computation and learning across edge nodes, the proposed approach eliminates reliance on centralized infrastructure and enables continuous monitoring and anomaly detection under degraded network conditions.

A central result of this work is the demonstration that reconstruction loss from a decentralized autoencoder can serve as an effective and interpretable proxy for system health. When aggregated across nodes, this metric provides a global view of system degradation that supports real-time estimation of RUL. This approach allows the system to detect both gradual degradation trends and abrupt regime changes without requiring explicit fault models, making it well-suited for complex and evolving operational environments.

The integration of decentralized learning with the SEPP communication layer further enables scalable operation across dense, distributed sensor networks. By exchanging compact model updates rather than raw data, the system reduces bandwidth requirements while maintaining collective awareness across nodes. This combination of communication efficiency and distributed intelligence represents a key advancement over traditional centralized and federated predictive maintenance architectures.

From an operational and acquisition perspective, SEPP-DPM provides a pathway toward resilient, scalable predictive maintenance capabilities aligned with modern defense priorities. The ability to operate autonomously at the edge, tolerate intermittent connectivity, and adapt to non-IID and non-stationary conditions makes the architecture particularly relevant for naval platforms, unmanned systems, and other distributed assets operating in contested environments.

In summary, this work demonstrates that decentralized, communication-aware learning architectures can provide an effective foundation for predictive maintenance in next-generation distributed systems. Future work will focus on validation in real-world deployments, integration of advanced learning models, and tighter coupling between communication and learning to further enhance system performance and operational impact.

References

- Ben Ali, J., Chebel-Morello, B., Saidi, L., Malinowski, S., & Fnaiech, F. (2015). Accurate bearing remaining useful life prediction based on Weibull distribution and artificial neural network. *Mechanical Systems and Signal Processing*, 56–57, 150–172. <https://doi.org/10.1016/j.ymssp.2014.10.014>



- Jain, A. K., Chouksey, P., Parlikad, A. K., & Lad, B. K. (2020). *Distributed diagnostics, prognostics and maintenance planning: Realizing industry 4.0*. Elsevier. <https://doi.org/10.17863/CAM.51888>
- Lou, G., Lin, W., Huang, G., & Xiang, W. (2023). A two-stage online remaining useful life prediction framework for supercapacitors based on the fusion of deep learning network and state estimation algorithm. *Engineering Applications of Artificial Intelligence*, 123, 106399. <https://doi.org/10.1016/j.engappai.2023>
- Maymounkov, P. & Mazières, D. (2002). Kademia: A peer-to-peer information system based on the XOR metric. In *Revised Papers from the First International Workshop on Peer-to-Peer Systems (IPTPS '01)*. Springer-Verlag, Berlin, Heidelberg, 53–65.
- Pruckovskaja, V., Weissenfeld, A., Heistracher, C., Graser, A., Kafka, J., Lepusch, P., Schall, D., & Kemnitz, J. (2023). Federated learning for predictive maintenance and quality inspection in industrial applications. *2023 Prognostics and System Health Management Conference*. <https://ieeexplore.ieee.org/document/10159323>
- Ryabinin, M., Borzunov, A., Diskin, M., Gusey, A., Mazur, D., Plokhotnyuk, V., Bukhtiyarov, A., Samygin, P., Sinitsin, A., & Chumachenko, A. (2020). *Hivemind: Decentralized deep learning in PyTorch*. <http://github.com/learning-at-home/hivemind>
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014a). Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1), 1929–1958.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014b). A two-stage online remaining useful life prediction framework for supercapacitors based on the fusion of deep learning network and state estimation algorithm. *The Journal of Machine Learning Research*, 15(1), 1929–1958.





ACQUISITION RESEARCH PROGRAM
DEPARTMENT OF ACQUISITION, FINANCE, AND MANPOWER
NAVAL POSTGRADUATE SCHOOL
555 DYER ROAD, INGERSOLL HALL
MONTEREY, CA 93943

WWW.ACQUISITIONRESEARCH.NET