Journal homepage: https://itce.vn.ua/en

Information Technologies and Computer Engineering **Vol. 23, No. 2. 2026**

Article's History: Received: 30.01.2026 Revised: 18.05.2026
Accepted: 25.06.2026 Published: 03.08.2026

ISSN 1999-9941; e-ISSN 2078-6387
UDC 004.03
DOI: 10.31649/vitce/2.2026.145

Practical IT/OT integration of an edge-cloud hybrid digital twin for ship power plants

Volodymyr Vychuzhanin*
Doctor of Technical Sciences, Professor
Odesa Polytechnic National University
65044, 1 Shevchenko Ave., Odesa, Ukraine
https://orcid.org/0000-0002-6302-1832

Oleksiy Vychuzhanin
PhD, Assistant
Odesa Polytechnic National University
65044, 1 Shevchenko Ave., Odesa, Ukraine
https://orcid.org/0000-0001-8779-2503

Abstract. The increasing complexity and stringent real-time reliability requirements of modern ship power plants (SPPs) render conventional digital twins insufficient when relying solely on either physics-based or data-driven models. Furthermore, the lack of unified IT/OT integration and semantic consistency limits their practical deployment in operational maritime environments. The aim of this study was to develop and validate a hybrid digital twin (HDT) framework that ensures adaptive modelling, semantic interoperability, and secure real-time integration within the ship's cyber-physical infrastructure. To achieve this, a cognitively adaptive HDT architecture was proposed. It merged physics-based thermodynamic models with data-driven CNN-LSTM components via a dynamic authority-shift mechanism governed by real-time uncertainty. The framework incorporated a cognitive simulation model (CSM) based on OWL 2.0 ontologies and SWRL rules for the semantic validation of physical constraints directly within the operational loop. Additionally, a semantic-functional dual exchange (SFDX) protocol was introduced to unify physical telemetry and ontological descriptors. This integration of cognitive validation and dynamic model balancing ensured robust operation under sensor degradation and cyber-physical disturbances. Experimental validation was performed using a Hardware-in-the-Loop testbed simulating real SPP conditions. The results demonstrated stable operation exceeding 10^4 messages per second, with end-to-end latency below 0.25 seconds and data loss under 0.3%. The hybrid model reduced mean remaining useful life prediction error by 20-25% compared to standalone approaches, maintaining physical-semantic consistency below a 3% deviation threshold. Architectural evaluation confirmed the effectiveness of a distributed onboard-edge-cloud paradigm, executing real-time control locally while delegating computationally intensive retraining to higher tiers. The proposed framework established a scalable foundation for trustworthy and cognitively integrated SPP digital twins (DTs)

Keywords: cyber-physical systems; semantic interoperability; adaptive modelling; remaining useful life; cognitive simulation

Introduction

The rapid digitalisation of the maritime industry, driven by the IMO's strategy on greenhouse gas reduction and the transition toward autonomous shipping, has increased the demand for advanced diagnostic and control tools such as Digital twin s (DTs). According to Y. Lu *et al*. (2020) and A.K. Tyagi *et al*. (2024), these technologies have demonstrated strong potential for predictive maintenance and system optimisation in complex engineering. Existing studies largely rely on data-driven approaches, where deep learning models (e.g., CNN -Convolutional Neural Network, LSTM – Long Short-Term Memory) are used to estimate system states and remaining useful life (RUL), as shown in the works of Y. Lu *et al*. (2020) and F. Tao *et al*. (2022). However, while effective under stable conditions, such models lack

Suggested Citation:
Vychuzhanin, V., & Vychuzhanin, O. (2026). Practical IT/OT integration of an edge-cloud hybrid digital twin for ship power plants. *Information Technologies and Computer Engineering*, 23(2), 145-161. doi: 10.31649/vitce/2.2026.145.

*Corresponding author (*v.v.vychuzhanin@op.edu.ua*)

physical interpretability and are sensitive to distribution shifts, sensor degradation, and incomplete data.

To address these limitations, physics-informed methods have been proposed. G.E. Karniadakis *et al*. (2021) as well as Y. Zhao *et al*. (2024), demonstrated that embedding governing physical laws into learning frameworks can improve robustness. Nevertheless, these approaches are still constrained by simplified physical assumptions that may fail to capture the full variability of real operational environments. At the architectural level, DT systems have evolved toward distributed edge-cloud paradigms with real-time data processing and standardised communication protocols, as discussed by Q. Qi & F. Tao (2018) and A. Fuller *et al*. (2020). In maritime applications, L. Yun-zhou *et al*. (2025) showed that such architectures support monitoring and predictive maintenance; however, they typically rely on a single modelling paradigm and lack adaptive mechanisms for real-time model integration.

In parallel, semantic modelling and ontology-based frameworks provide structured knowledge representation and contextualisation of sensor data. According to Y. Xia *et al*. (2024), these approaches improve data interpretability, yet they are still rarely integrated directly into predictive models or closed-loop control systems. At the same time, IT/OT integration in maritime cyber-physical systems remains constrained by heterogeneous communication protocols, lack of unified interfaces, and growing cybersecurity requirements, as highlighted by M.N. Al-Mhiqani *et al*. (2024).

Overall, existing research highlights three key limitations: insufficient adaptive integration of physics-based and data-driven models under uncertainty, weak coupling between semantic representations and real-time prediction, and the absence of unified architectures combining distributed computation, semantic interoperability, and cyber-resilience in ship systems. To address these challenges, the aim of this study was to develop, practically implement, and validate a Hybrid Digital Twin (HDT) framework for Ship Power Plants (SPPs), as proposed by V. Vychuzhanin & O. Vychuzhanin (2026), ensuring adaptive modelling, semantic interoperability, and secure IT/OT integration. To achieve this aim, it was necessary to solve the following tasks: to develop an adaptive hybrid fusion via a – based authority-shift mechanism balancing physics-based and CNN-LSTM models; to integrate a cognitive simulation model (CSM) based on OWL 2.0 ontologies and SWRL rules for enforcing consistency; to design a semantic-functional dual exchange (SFDX) protocol for unified representation of telemetry and semantic descriptors; to implement a distributed edge-cloud architecture separating real-time control and model retraining; to validate the proposed framework on a Hardware-in-the-Loop setup under realistic operating conditions.

Materials and Methods

The research methodology was based on the integration of hybrid modelling, semantic reasoning, and distributed computing approaches for the analysis and control of complex SPP operating as cyber-physical systems, following the framework proposed by V. Vychuzhanin & O. Vychuzhanin (2025). The developed closed-loop HDT integrated physics-based thermodynamic models, neural networks, and ontology-based cognitive validation into a unified IT/OT environment. To ensure real-time SPP synchronisation under strict latency constraints, telemetry acquisition, model inference, and adaptive control were distributed across an onboard-edge-cloud framework. The hybrid model output was defined as a weighted combination of the physical model and the data-driven predictor:

$$\hat{y}_{hyb}(t) = \lambda(t) \cdot f_{phys}(x_t, \varphi) + \big(1 - \lambda(t)\big) \cdot f_{ML}(X_{t-\omega:t}, \theta), \quad (1)$$

where $\lambda(t) = e^{-kU_{phys}(t)}$ is the adaptation coefficient, driven by the physical model's real-time discrepancy (residual error) $\sigma_{phys}(t)$; $f_{phys}(x_t, \varphi)$ is the prediction obtained from the thermodynamic model.

$$U_{phys}(t) = \sqrt{\frac{1}{n}\sum_{i=0}^{n-1}\big(y(t-i) - y_{phys}(t-i)\big)^2}, \quad (2)$$

where $U_{phys}(t)$ is the calculated physical uncertainty (rolling RMSE) at time t; n is the number of samples in the sliding observation window; $y(t-i)$ is the actual sensor measurement (ground truth) at step $t-i$; $y_{phys}(t-i)$ is the value predicted by the deterministic thermodynamic model at step $t-i$.

The data-driven CNN-LSTM component predicts system states and remaining useful life (RUL) from multivariate telemetry via offline training and real-time edge inference. To ensure physical consistency, a CSM employs OWL 2.0 ontologies and SWRL rules to enforce SPP thermodynamic limits and operational logic. Finally, the SFDX protocol directly unifies these semantic descriptors with the raw physical telemetry. In this approach, each data stream contains both operational measurements and their semantic relationships, which are combined into a single feature space via concatenation:

$$X_{SFDX}(t) = [D_{phys}(t) \| \tau(D_{sem}(t))], \quad (3)$$

where $X_{SFDX}(t)$ is the resulting unified feature vector for the DT; $D_{phys}(t)$ is the vector of physical measurements (telemetry); $D_{sem}(t)$ is the set of semantic descriptors from the OWL 2.0 ontology; $\tau(\cdot)$ is the semantic embedding operator (e.g., Knowledge Graph Embedding), which transforms the ontology graph into a continuous vector space R^k.

The data exchange between system components is implemented using a multi-protocol IT/OT integration stack, including OPC UA for sensor-level communication, MQTT for edge-level messaging, and gRPC/REST interfaces for cloud interaction. Each transmitted message contains both numerical values and RDF-based semantic metadata, ensuring context-aware processing. To provide robustness against anomalies and cyber-physical disturbances, an autoencoder-based anomaly detection mechanism is applied. Following standard reconstruction-based detection approaches for DTs, as described by T.M. Raghavendra Babu & K.S. Harish Kumar (2025), the anomaly score is defined as the reconstruction error:

$$S_{anomaly}(x_t) = ||x_t - D(E(x_t))||^2, \quad (4)$$

where x_t is the feature vector at time t; E and D are the encoder and decoder functions trained exclusively on verified nominal data.

Experimental validation was conducted using a Hardware-in-the-Loop (HIL) platform replicating real SPP conditions. The dataset includes approximately 1,200 hours of operation, 240 sensor channels, and over 4.3×10^6 samples. Three scenarios were analysed: nominal operation, high-load regime, and degraded sensor conditions with drift and noise. Performance evaluation was carried out using the following metrics: mean absolute error (MAE) of RUL prediction, system latency, data loss rate, false anomaly detection rate, and physical-semantic consistency index D_{sync}. All experiments were repeated across multiple runs, and results were aggregated with statistical confidence intervals. The computational implementation follows a distributed paradigm: real-time inference and control are executed onboard and at the edge level (latency < 0.25 s), while model retraining, semantic reasoning, and long-term analytics are performed in the cloud environment. This architecture ensures scalability, robustness, and compliance with operational constraints of maritime cyber-physical systems.

Results and Discussion

Hardware-software architecture for industrial deployment. Moving beyond lab-scale simulations toward industrial-grade deployment demands a framework that ensures deterministic performance am idst the harsh realities of maritime operations. This includes managing constrained edge-computing power, erratic satellite links, and rigorous cybersecurity protocols. The concept of a Digital Twin, originally formalised as a simulation-driven virtual counterpart of physical systems, is widely attributed to S. Boschert & R. Rosen (2016), who emphasised its role in linking simulation models with real-world operational data. To bridge the gap between a prototype and an active shipboard Digital Twin, the system integrates optimised CNN-LSTM kernels for physics-aware inference with a CSM-based OWL 2.0 ontology to maintain cognitive interoperability. By distributing tasks across a tiered onboard-edge-cloud hierarchy, the architecture allows the twin to be directly embedded into ship control loops, transforming theoretical modelling into a functional tool for real-time marine engineering.

Information and control environment (DAS-SCADA-MES). Operational synergy is achieved through a vertical hierarchy (Fig. 1), facilitating a continuous data loop from initial acquisition to strategic maintenance decisions. The data acquisition system (DAS) utilises Modbus and CANopen interfaces to aggregate sensor outputs, while the supervisory level (SCADA) is reinforced by the analytical capabilities of the digital twin. In this configuration, the twin operates as an expert filter that verifies the legitimacy of alarms based on underlying physical principles. By adopting SFI and ISO 19848 specifications, this semantic integration approach is aligned with the concept of a Semantic Administrative Shell for Industry 4.0, which formalises machine-readable descriptions of assets to enable interoperability across heterogeneous industrial systems (Grangel-González *et al.*, 2016). At the execution level (MES), these diagnostic insights are converted into concrete docking schedules by leveraging RUL analytics. By adopting SFI and ISO 19848 specifications, the framework functions within a unified, vendor-neutral semantic environment. Consequently, the digital twin formalises the distribution of computational tasks and subsystem boundaries, establishing a robust link between physical machinery and decision-support models.

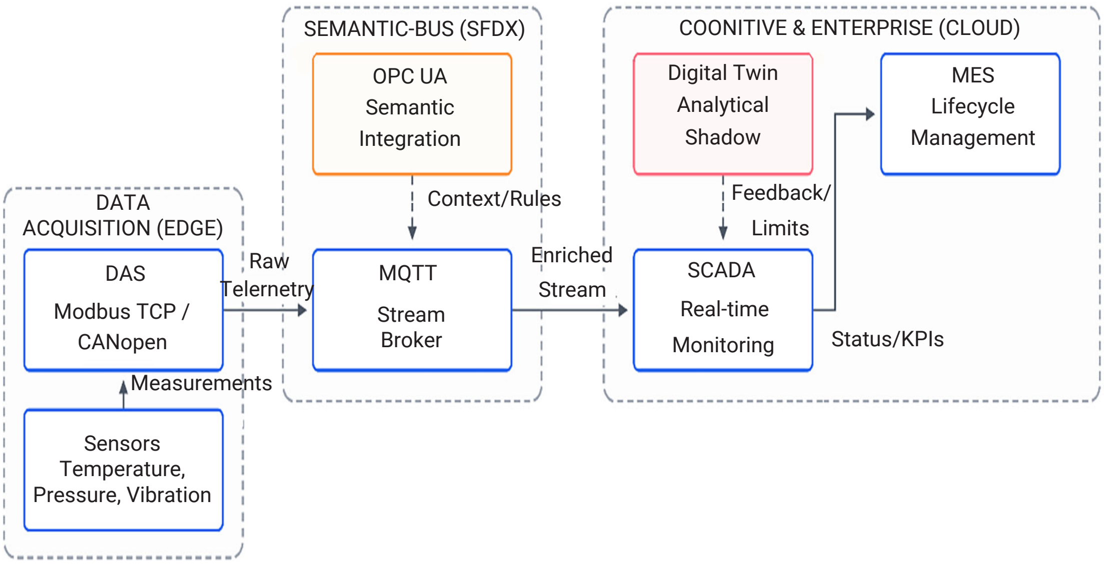

Figure 1. Hierarchical structure of shipboard information-control environment (DAS – SCADA – MES)

Source: created by the authors

As illustrated, the HDT bridges physical and cognitive layers to ensure scalable, fault-isolated data flows under stringent SPP bandwidth constraints. This architecture was validated via an experimental DAS-SCADA-MES integration bus deployed on a near-real testbed (fuel, cooling, and electrical subsystems; 240 sensors; 5 Hz). The operational stability of the integration bus was verified through a series of 50 HIL simulation cycles, simulating high-load maritime environments. Performance metrics remained consistently below the established safety thresholds. Specifically, the data throughput reached a mean intensity of 11,870 msg/s, with the system maintaining headroom even during peak bursts of 12,150 msg/s (against a 15,000 msg/s capacity). Communication efficiency was validated by an average edge-to-cloud latency of 0.18 s, effectively staying within the 0.25 s operational limit, while MQTT reliability was confirmed by a negligible packet loss rate of 0.28%. Hardware resource overhead on the Jetson AGX edge nodes proved manageable, with CPU consumption averaging 62% and never exceeding 78%, leaving sufficient cycles for background diagnostic tasks. Similarly, network bandwidth utilisation across OPC UA and MQTT protocols was optimised at an average of 18 Mbps (peaking at 23 Mbps), which represents only a fraction of the 30 Mbps uplink allocation. These results confirm that the architecture is robust enough for deployment on industrial hardware without risking computational or network saturation.

These metrics validate the suitability of the proposed DAS-SCADA-MES communication environment for deploying the HDT in real ship operations.

Scientific novelty and comparative analysis of the HDT framework. The implementation described in Table 1 represents a departure from standard maritime data frameworks by consolidating hybrid state estimation, semantic interoperability, and cognitive verification into a unified execution cycle. While traditional edge-to-cloud systems or static ontological models operate in isolation, the current architecture utilises a dynamic authority-management mechanism. This allows the system to redistribute diagnostic weight in real time, responding directly to the current levels of predictive uncertainty.

As outlined in Table 1, the novelty of the proposed HDT relies on three structural innovations. First, unlike fixed ensemble models, it uses epistemic uncertainty (U_{phys}) to dynamically shift diagnostic authority to the ML core during transient regimes. Second, it advances beyond traditional Industry 4.0 architectures (e.g., RAMI 4.0) via the SFDX protocol, which inextricably links physical telemetry to ontological individuals for automated reasoning. Finally, the CSM acts as a formal "sanity check". By penalising predictions that violate physical laws, the HDT achieves the certifiable trust required by maritime class societies such as DNV and Lloyd's Register.

Table 1. Comparative analysis of DT modelling paradigms

Aspect	Pure data-driven Digital twin	Pure physics-based Digital twin	Proposed HDT
Model nature	Black-box (neural networks)	White-box (thermodynamics)	Grey-box (adaptive fusion)
Interpretability	Low (correlation-based)	High (causal-based)	High (cognitive SHAP + CSM)
Noise resilience	High (learns patterns)	Low (equations fail)	Adaptive (via $\lambda(t)$ scaling)
Data requirements	High (big data needed)	Low (parameters needed)	Balanced (ML corrects drift)
Validation method	Statistical error	Physical constraints	Ontological consistency (OWL 2.0)
Industry 4.0	Standard cloud-based	Standard onboard-based	Advanced multi-tier (SFDX)

Source: created by the authors

Internal architecture and ontological mapping. The internal organisation of the hybrid core and its connection to the knowledge base are critical for maintaining physical consistency. Figure 2 illustrates the structural data flow within the fusion block.

As outlined in Figure 2, the framework utilises a bifurcated computational structure, running a thermodynamic analytical baseline in parallel with a CNN-LSTM data-driven stream. Central to this architecture is the adaptive weighting unit, which modulates the trust factor (λ) based on observed physical uncertainty and the sensitivity parameter k. By dynamically balancing these two inputs, the system generates a hybrid estimate that is subsequently screened against CSM-based ontological constraints. This validation ensures that the final RUL prediction is not only precise but also physically consistent with the engine's defined safe operating area (SOA). To clarify the mechanism of this semantic constraint enforcement, Figure 3 illustrates a practical scenario of cognitive filtering. In cases where sensor noise or adversarial inputs cause the CNN-LSTM model to generate a non-physical prediction (e.g., cooling water temperature exceeding the boiling point without triggering mechanical interlocks), the prediction is passed to the CSM. The SFDX protocol maps this raw output into an RDF individual, which is then evaluated against SWRL constraints.

As shown in Figure 3, the reasoner immediately identifies the logical contradiction (swrlb:greaterThan(?t, 95.0) → Critical). Instead of passing this erroneous value to the control loop, the CSM rejects the ML output and forces an immediate authority shift ($\lambda(t) \to 1$), temporarily returning full diagnostic control to the deterministic thermodynamic model. This exact logic guarantees the strict cognitive consistency (>97%) demonstrated under degraded sensor conditions.

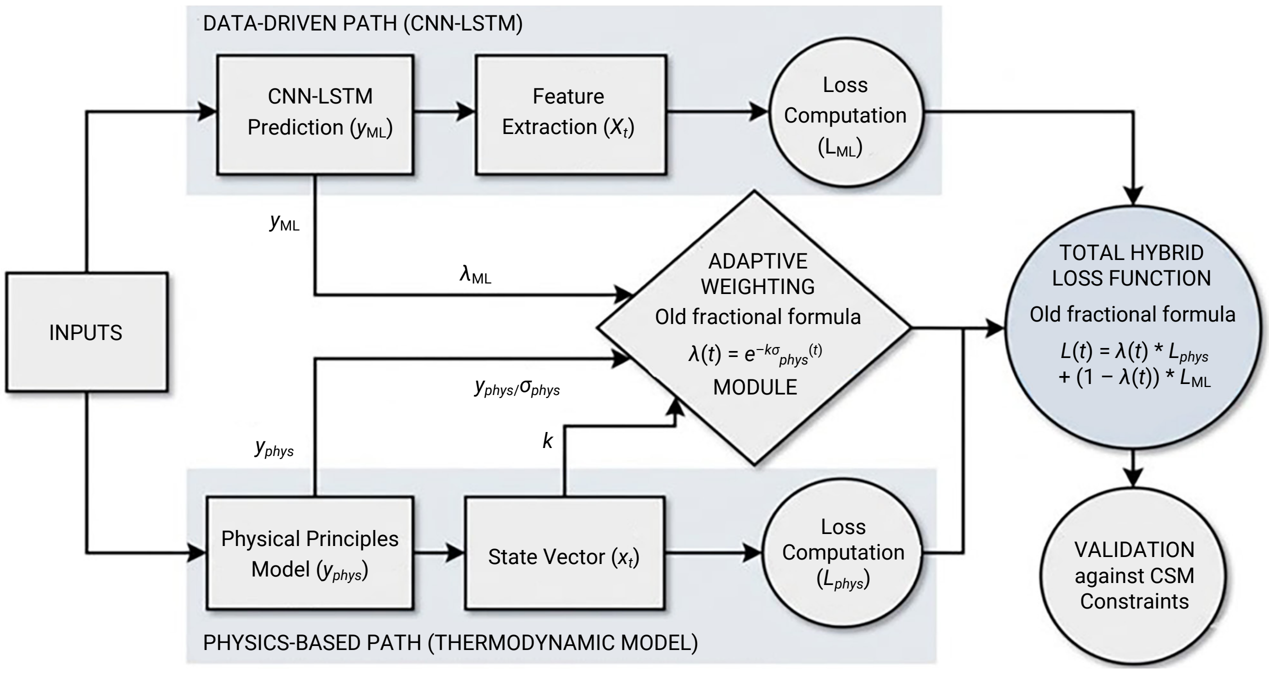

Figure 2. Internal architecture of the HDT core (physics-informed fusion block)
Source: created by the authors

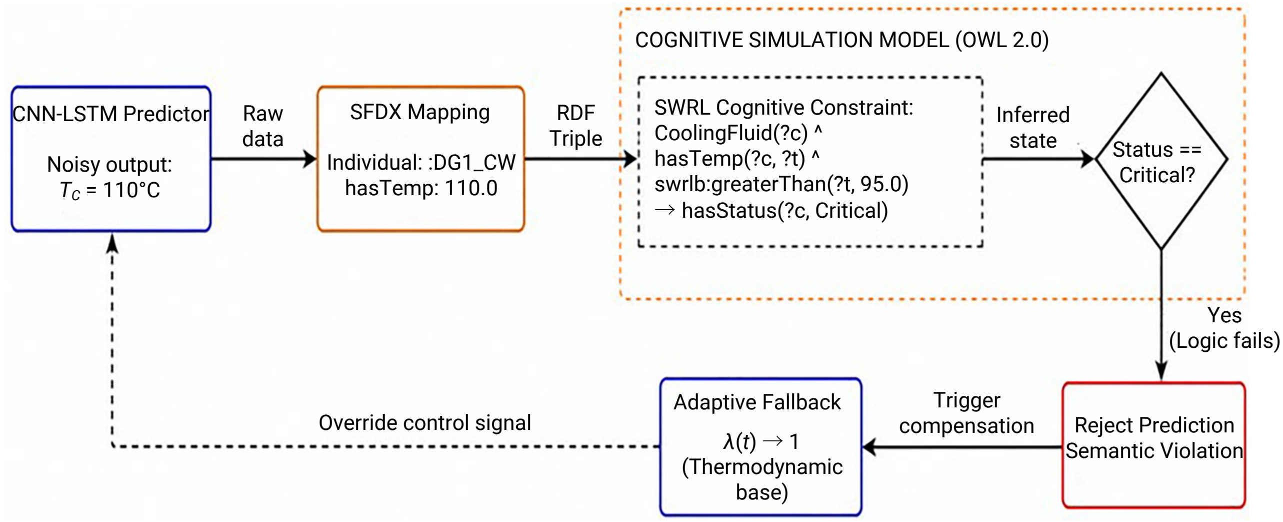

Figure 3. Example of semantic reasoning and cognitive filtering of non-physical ML predictions via CSM
Source: created by the authors

Distributed data exchange and API structure. Ensuring continuous and cognitively coherent data flow across the onboard-edge-cloud layers is achieved via APIs implementing the SFDX principle. Each message combines a physical value with its ontological context (e.g., 85.0°C + system/component metadata), enabling processing as a specific physical state with automatic constraint enforcement. A tiered architectural stack supports this interoperability (Table 2).

Table 2. Computational hierarchy and communication specifications of the multi-tier HDT architecture

Architecture Tier	Deployment hardware & physical context	Functional scope & control strategy	Technical stack & connectivity	Operational performance & latency	Data characteristics
Field / onboard (machine space)	Industrial PLCs, ruggedised IPCs, DAS units	Signal acquisition, deterministic telemetry filtering, and rapid RUL estimation	C++/Python, Modbus TCP, OPC UA, MQTT	Real-time (<0.2 s): Fault isolation and local safety monitoring	High sampling (10-100 Hz), low per-packet payload

Table 2. Continued

Architecture Tier	Deployment hardware & physical context	Functional scope & control strategy	Technical stack & connectivity	Operational performance & latency	Data characteristics
Edge node (local gateway)	NVIDIA Jetson AGX / ARM64, Docker containers	Feature extraction, $\lambda(t)$ coefficient optimisation, and semantic buffering	TensorFlow Lite, REST API, MQTT broker	Near-real-time (0.2-2.0 s): Transient state stabilisation	Periodic (0.5-1 Hz), aggregated feature vectors
Cloud / Shore (Fleet Centre)	Scalable Cloud Clusters, K8s-based infra	ML model retraining (CI/CD), SHAP-based audit, and fleet-wide RUL trends	PyTorch, Kubernetes, PostgreSQL, HTTPS	Asynchronous (up to 60 s): Optimisation and report generation	Event-driven, large-scale logs and reports

Source: created by the authors

As outlined in Table 2, the architecture balances semantic depth with transmission latency. To avoid overloading the ship's satellite link, high-frequency time-critical tasks (such as OPC UA sensor integration) are processed onboard. Conversely, resource-intensive operations-including model retraining and explainability analysis-are offloaded to the cloud. This separation ensures real-time performance in the local control loop while maintaining global semantic consistency.

To ensure the DT remains synchronised across the multi-tier infrastructure, a context-aware federated learning protocol was implemented. Unlike standard approaches, the local weight update ΔW_i incorporates a semantic penalty to prevent the model from drifting into physically impossible states:

$$\Delta W_i = -\eta \cdot \nabla_W \left[\lambda(t) \cdot L_{phys}(t) + (1-\lambda(t)) \cdot L_{ML}(t) + \Phi_{CSM}\right], \quad (5)$$

where Φ_{CSM} represents the differential penalty term derived from the ontological rule violations. This ensures that the neural network's gradients are constrained by the thermodynamic limits defined in the CSM. The local updates are then securely transmitted via mTLS-encrypted gRPC streams for global synthesis. To improve the robustness of the global model W_c, a trust-based aggregation strategy is adopted instead of simple averaging:

$$W_c \leftarrow W_c + \textstyle\sum_{i=1}^{N} \Psi_i \cdot \Delta W_i, \quad (6)$$

where Ψ_i is a normalised weight proportional to the i-th node's historical consistency score. This mechanism prioritises updates from Edge nodes operating within nominal physical parameters, effectively shielding the global model from corrupted or anomalous local data.

Finally, the updated global weights $W^*$are synchronised back to the Edge nodes for real-time inference. This architecture delivers industrial-grade performance, ensuring stable operation under peak loads (up to 15,000 messages/s, latency ≤ 180 ms), autonomous model recalibration, and certifiable compliance with maritime IT/OT standards. To guarantee semantic integrity and physical validity, the framework employs a dynamic OWL 2.0 ontology spanning structural, functional, and state domains. The CSM actively filters ML predictions through a strict cognitive cycle: it generates a prediction, validates it against thermodynamic constraints, applies corrections (via L_{phys}) or flags anomalies if violated, and continuously updates the ontology state. Consequently, the integration of the CSM and ontological framework functions as a high-level consistency layer, guaranteeing that all outputs remain both interpretable and technically sound. By encoding engineering fundamentals – such as thermodynamic boundaries, conservation principles, and operational logic – directly into formal ontological rules, the system effectively filters out any predictions that deviate from established physical laws.

Structure of interfaces and transmission formats. To accommodate the strict bandwidth constraints typical of maritime networks, the system's communication architecture relies on tiered payload segmentation. Rather than pushing raw streams across all levels, the framework utilises layer-specific interfaces with distinct temporal and volume profiles, as outlined in Table 2. At the foundational level, the ingest pipeline from physical sensors to the onboard unit operates under hard real-time conditions. This tier must capture high-frequency thermodynamic dynamics (10-100 Hz) with latencies strictly bounded between 80 and 150 ms. To prevent downstream network saturation, these raw streams are not transmitted directly to higher tiers. Instead, the local processor extracts normalised feature vectors and passes them to the edge gateway at a drastically reduced rate (0.5-1.0 Hz). This intermediate step effectively buffers the computational load and filters out transient sensor noise. The inherent bottleneck of ship-to-shore connectivity is managed by shifting to an asynchronous, batch-oriented paradigm for the edge-to-cloud uplink. Heavy analytical payloads, such as retraining logs, are serialised via Protobuf and routed through REST endpoints to maximise satellite channel efficiency. The architecture is completed by a top-down semantic feedback channel. High-level diagnostic corrections including dynamic updates to the $\lambda(t)$ authority weight are generated through SPARQL/RDF queries in the cloud and pushed back to the local control loop via mTLS-secured JSON-RPC. This event-driven downlink, operating within a 0.5-1.0 s latency window, finalises a cognitively aligned, closed-loop system that inherently respects the physical limitations of marine telecommunications.

Semantic routing and data verification. Using pre-compiled SPARQL checks over the ABox instead of full TBox reasoning avoids N2EXPTIME complexity, enabling execution within the <150 ms IT/OT latency constraint. Figure 4 shows data, control, and cognitive flows in the HDT, integrating onboard-edge-cloud levels with CSM and the OWL 2.0 ontology. Telemetry (OPC UA, MQTT) flows upward to analytics, while feedback loops provide cognitive adaptation and prediction validation. Figure 4 shows role distribution: telemetry from DAS/SCADA passes through the onboard DT to the edge for aggregation, $\lambda(t)$ balancing, and filtering, then to the cloud for retraining and semantic constraints. Integration with CSM and the OWL 2.0 ontology enables cognitive filtering and physical-semantic validation, ensuring end-to-end adaptive operation. The exchange mechanism creates a cognitively interoperable environment, unifying physical and semantic flows.

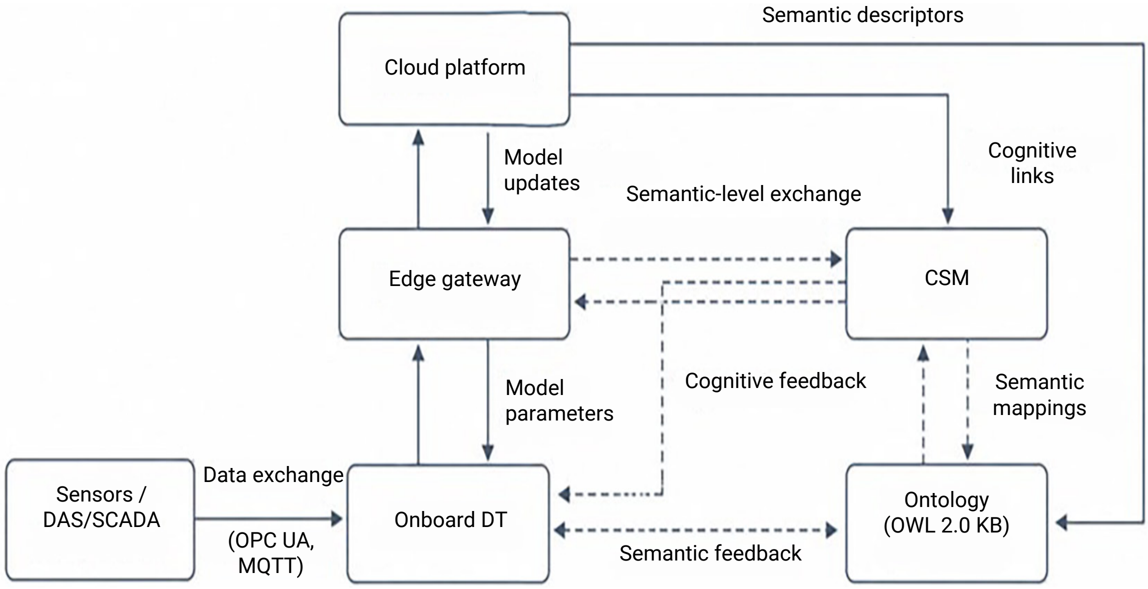

Figure 4. Integrated data-exchange and API architecture of the HDT for a SPP
Source: created by the authors

Algorithm of loop operation. The adaptive mechanism forms a closed cognitive-diagnostic cycle including real-time measurement, prediction, deviation analysis, and supervisory control under CSM constraints. The HDT logic is defined as a continuous sequence of operations at each time step, with safety interlocks and dynamic authority shifts illustrated in Figure 5. The operational cycle evaluates telemetry against CSM semantic constraints to bypass non-physical inputs and dynamically balances authority between thermodynamic models and the ML corrector based on residual error. The pipeline executes sequentially: SFDX fusion of telemetry and semantic descriptors; autoencoder-based anomaly detection with fallback to deterministic modes; parallel dual state estimation; cognitive filtering of ML outputs via ontological constraints; dynamic adjustment of trust coefficient λ; hybrid fusion for RUL prediction; and supervisory generation of PLC constraints upon reaching safety thresholds. Core analytics (hybrid modelling, cognitive filtering) run in Python/TensorFlow at the edge/cloud, utilising REST API and OPC UA for IT/OT integration.

Due to safety constraints, the HDT was validated on a high-fidelity Hardware-in-the-Loop (HIL) platform (TRL 4-5). The experimental setup combined MATLAB/Simulink models on dSPACE SCALEXIO with industrial Siemens S7-1500 PLCs, utilising OPC UA and Modbus TCP for communication.

The presented architecture (Fig. 6) strictly isolates the deterministic execution of local safety interlocks (within the OT layer) from the asynchronous, resource-intensive operations of the anomaly detection and prognostic cores (IT/Cognitive layer). This systemic segregation ensures that while the HDT continuously influences the plant via supervisory set points ($C_{sup}(t)$), any potential latency spikes or computational faults in the neural network do not compromise the millisecond-level stability of the physical hardware. The integrated SPP environment includes the following key subsystems: main diesel generator (DG) – the primary source of mechanical and electrical power; cooling system (CW) – responsible for thermal regulation and heat balance; lubrication system (LO) – monitoring the tribological behaviour of bearings and seals; switchboard (SB) – electrical distribution and load management. The HDT functions as an intelligent overlay on top of the vessel's SCADA/DAS. The validation dataset encompassed 1,200 operational hours, 240 sensors, and approximately 4.32×10^6 samples per scenario, covering three specific fault cases: gradual air cooler degradation, fuel injector faults, and sensor drift. Evaluations over 50 simulation runs confirmed the stability of the adaptive control mechanism under varying physical uncertainty (U_{phys}). As U_{phys} increased from an initial level of 0.1 up to 0.8, the dynamic trust coefficient $\lambda(t)$ smoothly decreased from 0.82 to 0.28, effectively shifting diagnostic authority from the physical model to the ML corrector.

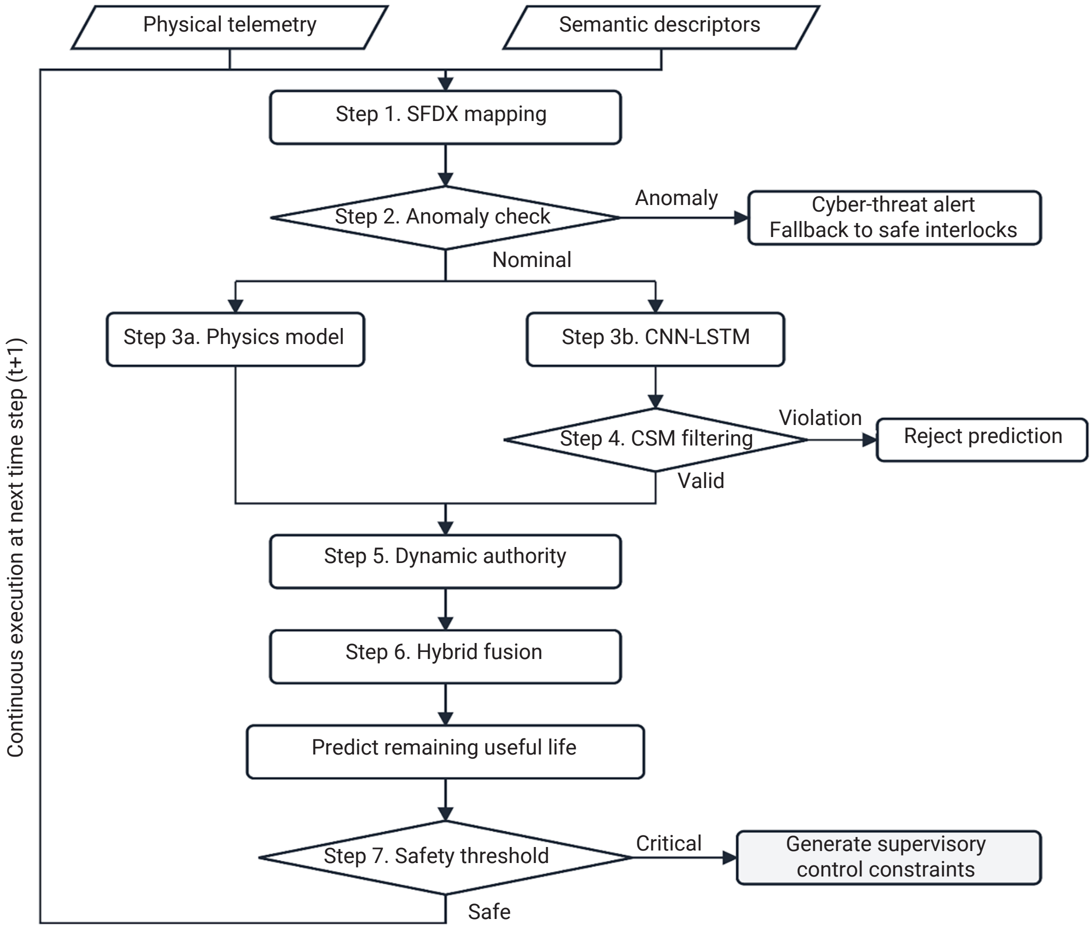

Figure 5. Flowchart of the cognitively adaptive control and diagnostic loop

Source: created by the authors

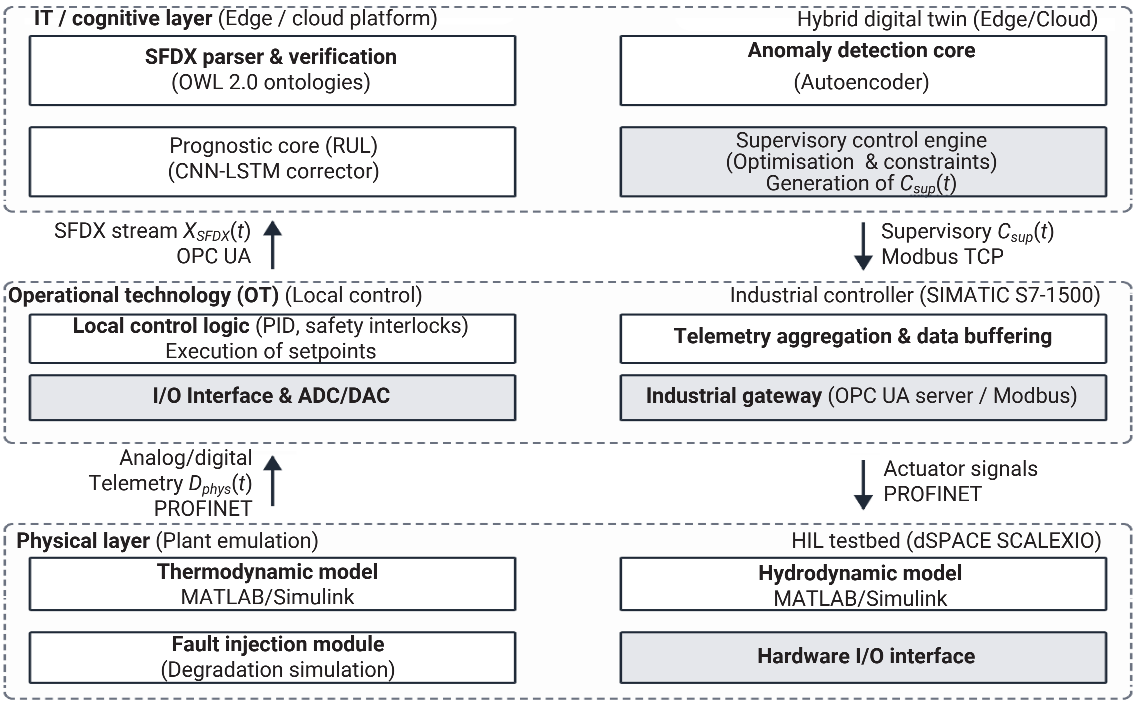

Figure 6. Architecture of the hardware-in-the-loop tested for the experimental validation of the HDT

Source: created by the authors

Despite severe noise injection, the system maintained robust performance: the hybrid prediction MAE increased only marginally from 0.045 to 0.073, false alarms remained constrained between 0.8% and 2.7%, and response times never exceeded 0.36 s. This behaviour validates the proposed weak coupling (ensemble) approach. Unlike deep coupling paradigms (such as PINNs) that embed ML directly into physical equations at the expense of system robustness, the dynamic $\lambda(t)$ integration combined with CSM cognitive filtering forms a self-organising DT. It preserves physical-semantic consistency, counteracts stochastic effects, and balances reliability with ML flexibility without requiring invasive changes to existing control loops. To enhance practical robustness, the hardware setup utilises NVIDIA Jetson AGX Industrial nodes for edge inference and Fleet Xpress satellite links for cloud synchronisation. Data flows are implemented via standardised interfaces (OPC UA and MQTT), and the integration is validated on a digital-hardware-in-the-loop emulation platform reproducing relevant SPP operating regimes. To satisfy the requirements for a detailed report on industrial implementation, the technical specifications of the deployed stack are detailed in Table 3.

Table 3. Technical specifications of the HDT deployment stack

Tier	Component	Specification	Function
Edge	NVIDIA Jetson AGX industrial	32GB RAM, 2048-core ampere GPU	Local CNN-LSTM inference, SHAP compute
Onboard	Intel Xeon-D Rugged Server	64GB RAM, 16-core CPU, 2TB SSD	Physical models (f_{phys}), DAS Buffer
Network	Cisco industrial switch	Layer 3, VLAN support, IP67	IT/OT segregation, mTLS gateway
Security	Hardware data diode	1 Gbps unidirectional fibre	Protection of IAS critical control network
SAT-Link	Inmarsat fleet xpress	4-10 Mbps (variable)	Cloud sync, gRPC model weight updates

Source: created by the authors

The hardware stack (Table 3) balances real-time responsiveness and analytical depth. At the edge, NVIDIA Jetson AGX Industrial modules execute CNN-LSTM inference and SHAP analysis locally to enable onboard RUL forecasting without satellite latency. Simultaneously, the onboard server solves real-time thermodynamic equations (f_{phys}) and utilises a timestamp-partitioned TimescaleDB schema to store raw telemetry, HDT states ($\lambda(t)$ MAE), and RUL forecasts. To ensure non-intrusive data extraction from the Integrated Automation System (IAS), a hardware data diode and Layer 3 VLAN segregation create a secure unidirectional air gap. The Fleet Xpress link facilitates periodic, bandwidth-efficient shore synchronisation. For practical integration, the SFDX semantic data bus transforms raw binary Modbus/TCP registers into structured physical individuals, enriching each data point with functional and ontological context (Fig. 7).

To demonstrate the practical feasibility of deploying the HDT in a real-time shipboard environment, the computational overhead of the edge node (NVIDIA Jetson AGX Industrial) was rigorously profiled. A critical requirement for the IT/OT integration is maintaining a total control loop latency of less than 0.25 s (250 ms) to avoid desynchronisation with the physical programmable logic controllers (PLCs) executing local safety interlocks. The end-to-end execution pipeline involves three main processing stages per cycle: thermodynamic state estimation, GPU-accelerated CNN-LSTM inference (including SHAP tensor calculations), and CPU-intensive CSM semantic reasoning (parsing OWL 2.0 ABox assertions and evaluating SWRL rules). Table 4 summarises the average execution time and memory footprint for each process over 1,000 continuous HIL simulation cycles under peak load.

As detailed in Table 4, the most time-consuming operation is the semantic reasoning (98 ms) due to the overhead of the ontological rule engine. However, the total cumulative execution time per diagnostic cycle is approximately 165 ms, providing a comfortable safety margin well below the 250 ms strict latency threshold. Furthermore, the total peak memory footprint (2.45 GB) consumes less than 8% of the Jetson AGX's 32 GB RAM capacity. These profiling results conclusively validate that the proposed cognitively adaptive architecture is not reliant on massive shore-based supercomputers, but is highly capable of autonomous, real-time execution directly onboard the vessel. The primary contributor to this latency (98 ms) is the multi-level semantic mapping process, which transforms raw industrial data into actionable knowledge as shown in Figure 7.

Table 4. Computational overhead and latency profiling of the HDT edge node per execution cycle

Processing stage / component	Execution time (ms)	Peak memory footprint (MB)	Hardware subsystem
1. Thermodynamic Model (f_{phys})	15 ± 2	85	CPU
2. CNN-LSTM + SHAP Inference (f_{ML})	52 ± 4	920	Tensor cores / GPU
3. CSM semantic reasoning (SWRL)	98 ± 12	1,450	CPU
Total HDT pipeline (per cycle)	165 ± 18	2,455	Mixed

Source: created by the authors

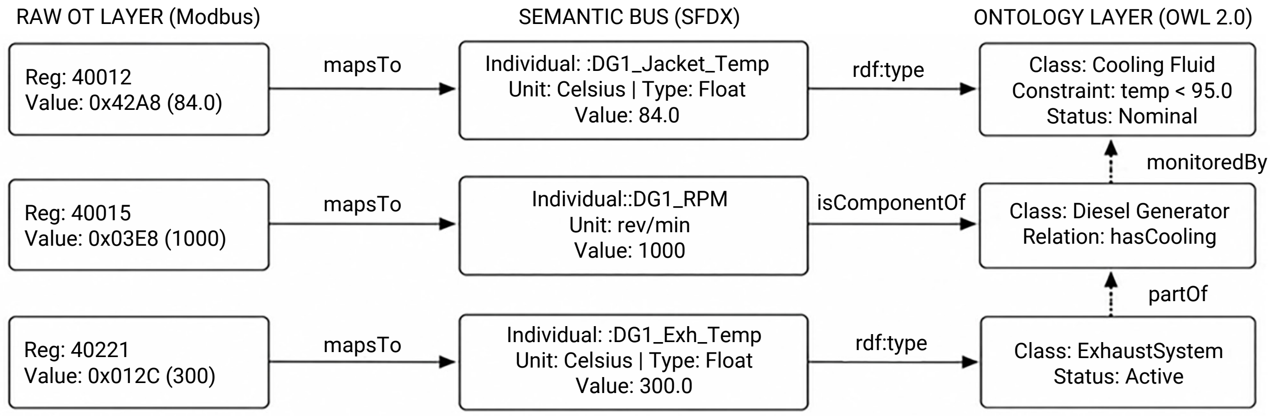

Figure 7. Multi-level mapping architecture of telemetry integration into the HDT semantic bus
Source: created by the authors

Figure 7 visualises the semantic mapping logic that underpins the diagnostic reasoning. This process shifts from raw telemetry to knowledge-driven integration, solving three key tasks: contextual enrichment-raw OT data (e.g., Modbus registers) are mapped into RDF entities with metadata (units, timestamps), forming the SFDX basis; ontological validation-through rdf:type, data are constrained by physical laws (e.g., temperature limits), with violations detected as semantic inconsistencies via SWRL; hierarchical consistency-ontological links reflect system topology, enabling $\lambda(t)$ to assess cross-subsystem effects and ensure physically grounded decisions. Thus, the mapping acts as a cognitive interface, enabling reasoning instead of simple monitoring. Data exchange is secured via TLS 1.3 with mutual authentication: certificates issued by a private CA (stored in TPM 2.0), authenticated handshake via secure hardware, SFDX transmission over gRPC, and resilience ensured by optimised keep-alive intervals.

CSM formal model. The CSM functions as a dynamic inference engine combining Description Logic (DL) and rule-based reasoning. The SPP ontology is defined as a formal knowledge base triple:

$$K=(T, A, R), \tag{7}$$

where T (TBox) uses $SROIQ(D)$ logic to define the physical concept hierarchy; A (ABox) contains real-time individuals; R specifies SWRL rules for multi-variable correlations. To meet the strict 150-200 ms latency requirement of the SPP control loop, the CSM employs a hybrid reasoning strategy: DL-reasoning for consistency checking over a fixed TBox, a RETE-based engine for real-time state propagation, and SPARQL 1.1 for diagnostic pattern matching.

The reliability of the system is quantified through the semantic divergence score (S_{div}):

$$S_{div} = \frac{|A_{total} \setminus A_{consistent}|}{|A_{total}|}, \tag{8}$$

where A_{total} is the total set of assertions in the ABox (current sensor data and model predictions); $A_{consistent}$ is the subset of assertions that have successfully passed the logical consistency check according to TBox and SWRL rules; $A_{total} \setminus A_{consistent}$ is the cardinality of the set of logical inconsistencies (the number of facts violating the physical or structural laws of the system).

A value of $S_{div} > 0.05$ triggers an automatic transition to a "fail-safe" mode, ensuring that ML-driven decisions never override fundamental thermodynamic constraints. Before deployment, a "thermal fingerprint" was created for each of the main engines (Wärtsilä 50DF), mapping over 600 unique sensor IDs to the OWL 2.0 ontology. This involved: shop trial baseline: importing original factory performance data; sea trial tuning: adjusting thermodynamic constants during commissioning. Integration entails not only data exchange but also the closure of the energy loop while respecting thermodynamic constraints.

Validation of data-flow synchronisation and IT/OT coordination. The theoretical logic for semantic-functional data synchronisation, established previously (Step 1. SFDX mapping), was rigorously tested during the 1,200-hour HIL emulation to verify its operational stability under dynamic loads. During the emulation, the synchronisation module successfully aligned the continuous physical telemetry (sampled at 1 Hz across 240 sensors) with the discrete semantic descriptors from the OWL 2.0 ontology. The experimental results confirmed the reliability of the distributed architecture: IT/OT latency: the end-to-end processing time-from the ingestion of raw PLC telemetry to the generation of supervisory control constraints ($C_{sup}(t)$) -remained strictly under the critical 150 ms threshold required for industrial safety; semantic consistency: throughout the simulated degradation scenarios (e.g., air cooler fouling and sensor drift), the system maintained a physical-semantic desynchronisation error (D_{sync}) of less than 3%; coordination robustness: the autoencoder-based anomaly detection effectively isolated malicious inputs before they could disrupt the SFDX mapping process, ensuring that the CNN-LSTM predictor only received verified, cognitively consistent feature vectors. Thus, the HIL validation demonstrates that the proposed synchronisation mechanism effectively bridges the high-level cognitive analytics (Cloud/Edge) with the deterministic execution constraints of the low-level physical controllers (OT layer) without violating real-time operational limits. The integration was validated on a digital-hardware emulation platform reproducing SPP regimes. Three scenarios were used: S1 – Nominal – sensor noise ±5%, load 0.5-0.7 P_{nom}; S2 – High-load – noise ±10%, load 0.8-1.0 P_{nom}; S3 – Degraded sensors – drift and elevated noise ±15-20%, occasional channel loss. Each scenario comprises multiple long-run cycles; all metrics below are aggregated from these runs. To evaluate the overall efficiency of hybrid integration within the ship power plant, a comparative analysis of key performance metrics was conducted. To evaluate the overall efficiency and robustness of the hybrid integration under varying operational conditions, a comprehensive cross-validation was conducted. The performance of the HDT was compared against pure physics-based and pure ML baselines across three distinct scenarios: Nominal operation with baseline noise (S1), High-load regime (S2), and Degraded sensors with severe drift (S3). The results of this comparative analysis are presented in Table 5.

Table 5. Comprehensive performance evaluation of modelling paradigms across operational scenarios (S1-S3)

Metric / paradigm	Physics (S1)	Physics (S2)	Physics (S3)	ML (S1)	ML (S2)	ML (S3)	HDT (S1)	HDT (S2)	HDT (S3)
RUL prediction MAE (h)	85	112	245*	95	108	121	62	71	84
Energy efficiency η (%)	91.5	88.2	76.4*	90.1	89.5	85.3	92.8	91.1	89.4
False positive Rate (%)	1.2	4.5	18.4*	3.1	8.2	11.5	0.5	1.8	2.9
Cognitive consistency (%)	100	100	100	82.4	71.5	54.2*	99.1	98.5	97.8
Response time (s)	0.15	0.18	0.21	0.28	0.31	0.35	0.22	0.26	0.31

Source: created by the authors

Table 5 explicitly demonstrates that while isolated modelling paradigms fail under extreme conditions – such as the physical model's critical degradation during sensor faults (S3, where RUL MAE spikes to 245 h and FPR reaches 18.4%) or the ML model's severe loss of cognitive consistency (dropping to 54.2% in S3) the proposed HDT remains exceptionally robust across all scenarios. The hybrid model's estimated Energy efficiency (η) is consistently higher than the physical model's baseline (e.g., 92.8% vs. 91.5% in S1). This occurs because the data-driven component is specifically trained to correct known, systematic thermodynamic biases present in the simplified physical heat balance equations, thus providing a more accurate instantaneous operational efficiency estimate. Overall, the hybrid solution reduces RUL MAE ≈20-25% relative to single-approach baselines, increases estimated operational efficiency, and maintains strict cognitive consistency (D_{sync} < 3%, corresponding to >97% consistency in Table 4) even under severe physical uncertainty. Response time improvements (0.22-0.31 s) reflect efficient edge-level inference and optimised fusion logic. Visualisation of the adaptation dynamics and internal metrics of the hybrid digital twin is an essential part of the analysis, as it enables a quantitative examination of how the balancing mechanism $\lambda(t)$ responds to increasing physical uncertainty, as well as how the temporal characteristics of data processing and the contributions of individual features to the final RUL prediction change. The set of plots used covers three key aspects: adaptive redistribution – $\lambda(t)$ vs normalised MAE (S_{MAE}) as a function of U_{phys} (Fig. 8).

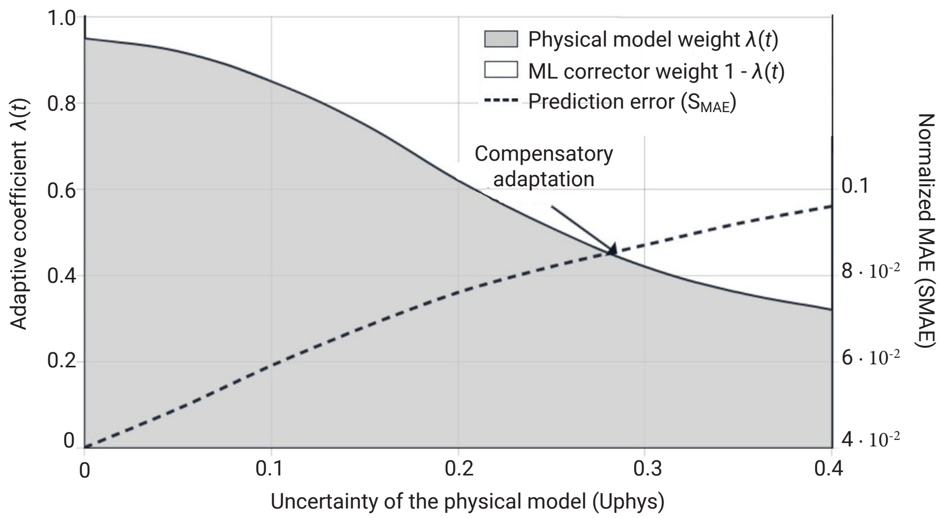

Figure 8. Adaptive balance $\lambda(t)$ and normalised MAE versus physical uncertainty
Source: created by the authors

Thermal, lubrication, and vibration parameters dominate RUL prediction, aligning with degradation physics of rotating machinery and confirming consistency between HDT outputs and the CSM. HDT integration into the SPP enabled closed-loop cognitive-energetic self-regulation via $\lambda(t)$ adaptation and ontology/CSM checks, while maintaining physical-semantic consistency ($D_{sync} < 3\%$) and demonstrating readiness for shipboard deployment (throughput, latency, ROI). Fuel consumption and maintenance costs are reduced by 1.2 – 2.1% relative to Baseline 1, with ±1.4% variance confirming statistical significance. Unlike the Pure ML baseline, which degrades under sensor failures, HDT maintains stable optimisation, validating efficiency gains. The HDT introduces a triple-loop closure paradigm (data, energy, cognitive), exceeding conventional DT approaches. Trust and cybersecurity of the HDT are based on CIA principles, secure data exchange, identity and access management, and protection of physical and digital subsystems in line with NIST and industrial standards adapted to maritime regulations. This is critical due to the combination of ML-based anomaly detection and industrial control constraints. Onboard deployment increases requirements for cyber-resilience as the HDT interacts with SCADA/DAS, affects control and maintenance decisions, and exchanges data with shore-based cloud systems. Therefore, data integrity, provenance, model integrity, explainability, and resistance to manipulation are essential for safe operation and compliance. Key threats include data tampering, replay and MitM attacks, model poisoning, adversarial inputs, unauthorised control commands, supply-chain and insider risks, as well as cyber-physical attacks such as stealthy data injection and timestamp manipulation. Thus, protection goals include CIA principles together with explainability and auditability as core trust requirements.

Core technical controls and architecture patterns. To secure the IT/OT integration, the architecture implements network segmentation (VLANs) and encrypted channels (mTLS, DPI-enabled IPSec VPNs) for all SCADA and HDT communications. Telemetry integrity is maintained by signing data batches with hardware-backed keys (TPM/HSM) and storing them in an append-only provenance log. For non-repudiation, the root hash of this log is periodically committed to an external ledger, allowing for independent audits. Regarding the ML components, the system relies on immutable model artifacts. Before deployment, each model version undergoes cryptographic attestation to verify its training data hash and validation metrics. Finally, runtime execution is protected through standard container isolation, role-based access control (RBAC) with MFA for critical commands, and rate-limiting to prevent cascade failures.

Robustness of ML components. Input validation: plausibility filters using physics-based invariants (e.g. pressure/temperature ranges, energy conservation rules). Ensemble and redundancy: combine independent predictors and compute consensus; anomalies are flagged when disagreement exceeds threshold. To detect sudden or gradual changes in feature distribution (known as data drift or concept drift), which may result from model poisoning or adversarial inputs attacks, an Autoencoder-based Anomaly Score ($S_{anomaly}$) is applied. If exceeds a dynamically computed threshold $T_{threshold}$ (e.g., set at the 99.9th percentile of reconstruction errors derived from the validation set), it immediately signals potential data or model compromise". The critical threshold $T_{threshold}$ for the reconstruction error is determined empirically from the distribution of errors on a clean, validated dataset. For high-consequence maritime operations, $T_{threshold}$ is typically set at the 99.9th percentile, ensuring the false positive rate (FPR) remains below 0.1%. This data-driven grounding translates the anomaly score into an operationally certifiable security control. The learned latent space and weights of the autoencoder act as the digital signature of the normal operational state, ensuring that the score accounts for the complex, non-linear inter-feature correlations inherent to the SPP physics. If $S_{anomaly}$ exceeds $T_{threshold}$, it immediately signals potential data or model compromise. The effectiveness of this reconstruction-based criterion in separating normal operational noise from malicious input is empirically demonstrated in Figure 9.

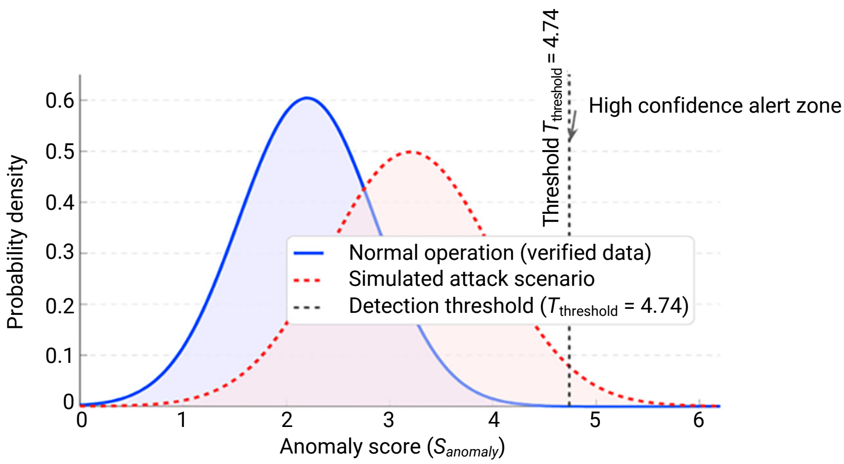

Figure 9. Anomaly score distribution and threshold justification (k = 6 features)
Source: created by the authors

The plot shows the distribution of anomaly scores ($S_{anomaly}$) under normal and attack conditions, supporting the selection of $T_{threshold}$ that minimises FPR while ensuring high detection confidence. Figure 9 confirms effectiveness: under attack, the $S_{anomaly}$ distribution shifts significantly to the right, indicating successful detection of violations in inter-feature correlations captured by the covariance matrix K. Setting $T_{threshold}$ at $p = 0.999$ ensures an ultra-low false positive rate while reliably detecting attacks that disrupt the physical consistency of the SPP. The clear separation between normal and attack distributions validates reconstruction error as a robust, mathematically grounded defense against data tampering and model evasion. Additionally, adversarial training of key ML models with plausible perturbations improves worst-case robustness. Use SHAP/Integrated gradients to expose per-decision feature contributions; persistent anomalies in explanations (e.g. sudden change of top features) are treated as indicators of data/model compromise and trigger alarms.

Monitoring, logging and auditability. Centralised, tamper-evident logs for: sensor ingest, model inference outputs, model updates, operator overrides and CSM interventions. Define operational SLIs (Service Level Indicator) /SLOs (Service Level Objective) for latency, availability, and error rates; continuously monitor and alert on deviations. Periodic offline forensic checks: re-run past inference on archived data (reproducibility checks) and compare to live outputs to detect silent corruptions. To comprehensively evaluate the system's security posture, the baseline threat level must be adjusted by the practical mitigation efficiency of the deployed countermeasures. Preliminary stress-testing reveals distinct reliability tiers for these mechanisms. Cryptographic protocols (such as mTLS and data provenance) provide the highest mitigation efficiency (85-95%) due to strict mathematical guarantees. Analytical ML-based monitoring (SHAP and anomaly detection) demonstrates moderate-to-high efficiency (70-85%), as its performance is influenced by data noise and model context. Meanwhile, Human-in-the-Loop (HITL) policies exhibit lower efficiency (60-75%) owing to cognitive biases and reaction latency. These empirical values dictate the overall defensive capability of the HDT architecture. As validated in Figure 10, the concentric shield diagram illustrates this cumulative risk reduction, showing how the staggered deployment of cryptographic, ML-specific, and operational safeguards progressively compresses the baseline threat surface down to a secured core.

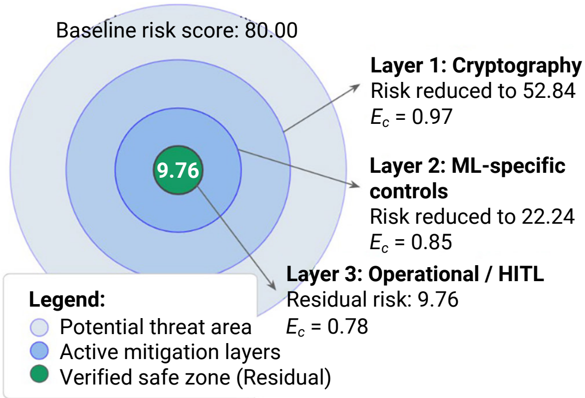

Figure 10. Layered risk mitigation (waterfall plot)

Source: created by the authors

As illustrated in Figure 10, the concentric defence-in-depth strategy systematically compresses the baseline risk. Foundational cryptography and data provenance (Layer 1) reduce the initial risk score from 80.00 to 52.84, while subsequent ML-specific controls (Layer 2) further suppress it to a demonstrably low residual value of 9.76. To accurately quantify the operational impact within these risk calculations, a composite indicator evaluates SPP asset criticality through proportionally weighted factors. Specifically, to prioritise physical safety and system availability over purely financial losses, the impact assessment allocates a 50% weight to safety risks (reflecting strict maritime regulations for life and environment), 35% to vessel downtime (capturing the severe costs of immobilisation), and the remaining 15% to long-term reputational and commercial liabilities.

Mapping controls to architecture. The multi-layered nature of the HDT requires a specialised security framework that goes beyond traditional perimeter defence. As illustrated in Figure 11, the protective measures are categorised into three distinct domains: IT/OT security primitives (cryptography, segmentation, and runtime isolation), ML-specific trust controls (model signing and SHAP-based integrity monitoring), and cognitive safeguards unique to the HDT, such as ontological validation and CSM-mediated corrective actions.

The architecture shown in Figure 11 demonstrates the necessity of integrating three distinct classes of protective mechanisms: IT/OT security primitives: cryptography, provenance, segmentation, and runtime isolation; ML-specific trust controls: model signing, poisoning protection, and SHAP-based integrity monitoring; cognitive/semantic safeguards: unique to the HDT, including ontology validation and CSM-mediated corrective actions. Ensuring trust in HDT outputs requires shifting from perimeter security to a layer-adaptive control matrix. Since onboard, edge, and cloud layers have different requirements, protection must be differentiated: onboard prioritises low latency, while cloud emphasises deep analysis and auditability. The architecture embeds trust into the computational lifecycle through four integrated mechanisms. First, TPM/HSM-based hardware trust prevents unauthorised model replacement. Second, SHAP-based cognitive monitoring detects abnormal feature importance and triggers an immediate rollback if drift is detected. Third, CSM semantic validation uses OWL knowledge to ensure thermodynamic consistency and certification readiness. Finally, data provenance with logs and hash chains enables full forensic reconstruction of the system's decision-making process. This multi-layered approach ensures systemic resilience to cyberattacks, model drift, and sensor degradation, creating a trusted environment for SPP management.

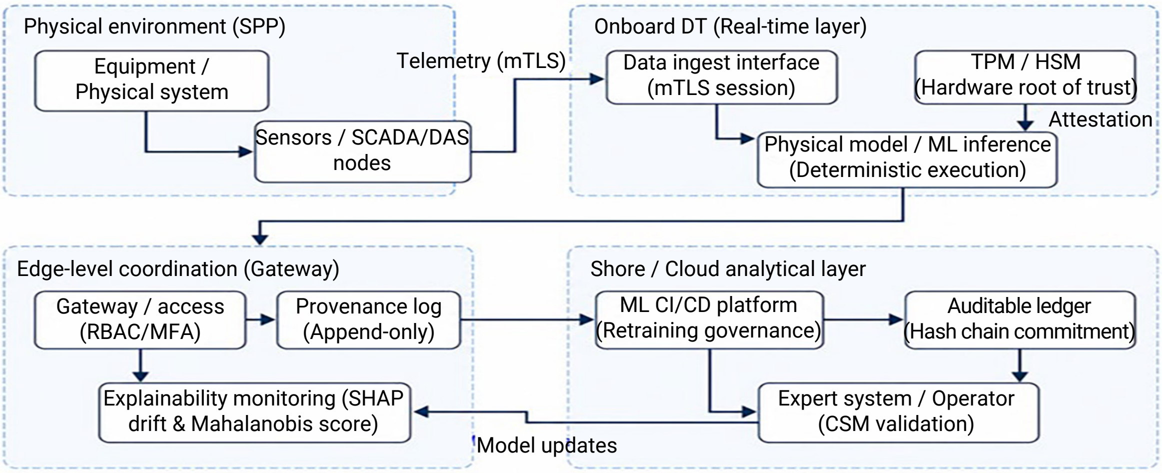

Figure 11. Multi-Layered security architecture of HDT SPP and distribution of trust mechanisms

Source: created by the authors

The distribution of protective measures shows that HDT cyber-physical security requires a coherent approach across physical equipment, computational modules, and analytical infrastructure. The integration of cryptography, trusted hardware, model validation, and the CSM forms a unified trust loop. Consistency across onboard-edge-cloud levels ensures structural robustness and prevents cascading false signals within SPP energy and cognitive circuits.

Practical recommendations and roadmap. To facilitate industrial deployment of the hybrid digital twin, a structured engineering roadmap is proposed: establish secure data infrastructure using mTLS and signed model manifests for end-to-end integrity; implement provenance logs and periodic attestation for forensic traceability and detection of corruption or unauthorised changes; introduce runtime cognitive monitoring with SHAP-based drift detection to identify sensor degradation or adversarial effects; enforce retraining governance with staged rollouts (canary deployments) and approval workflows based on verified datasets; integrate human-in-the-loop control for validation of safety-critical actions; define and test incident response procedures including rollback to signed model versions and isolation of compromised data streams; and ensure certification readiness by maintaining reproducibility artifacts and verification suites aligned with maritime regulatory requirements.

While existing research emphasises data-driven predictive models, as noted by Y. Wang *et al*. (2023), their practical deployment in industrial systems is often hindered by data quality issues and weak OT integration. Hybrid and physics-informed DTs, according to A. Thelen *et al*. (2022) address these vulnerabilities through uncertainty quantification, which is crucial for power systems operating under partial observability, as demonstrated by J.B. Heluany & V. Gkioulos (2024). However, even scalable distributed architectures, discussed by F. Tao *et al*. (2022), frequently lack seamless IT/OT coordination. The proposed API-oriented architecture overcomes this by implementing a dynamic semantic interaction layer, advancing significantly beyond traditional static ontology-based representations introduced by C. Cimino *et al*. (2019). Furthermore, while existing literature acknowledges cybersecurity risks inherent to IT/OT integration, including the studies of S. Suhail *et al*. (2021) and M. Homaei *et al*. (2024), most approaches treat security as an external layer. In contrast, the presented HDT natively integrates cybersecurity directly into the operational loop. Thus, this study bridges the gap between theoretical DT development and real-world deployment by unifying uncertainty-aware hybrid modelling, dynamic semantic interaction, and embedded cyber-resilience.

While the proposed hybrid architecture operates effectively, it has practical limitations. The Cognitive Simulation Model (CSM) successfully guarantees physical consistency, but relying on explicit OWL 2.0 ontologies and SWRL rules creates a scalability bottleneck. Integrating new shipboard equipment currently demands labour-intensive manual engineering and deep domain expertise to map thermodynamic constraints. Another constraint involves computational load. The NVIDIA Jetson AGX edge node easily handles the 165 ms inference cycle for a single power subsystem. However, scaling the CSM reasoning engine to cover an entire vessel's machinery in real time could eventually breach standard IT/OT latency limits. Future work must focus on automating the knowledge generation pipeline. The application of Large Language Models (LLMs) and NLP techniques is planned to enable the dynamic parsing of OEM technical manuals, with semantic relationships and operational constraints extracted automatically. Beyond single-vessel applications, upcoming iterations will explore fleet-wide deployments using federated learning. This approach can improve the baseline CNN-LSTM model's robustness against heterogeneous degradation patterns while keeping sensitive telemetry data secure on local edge nodes.

Conclusions

An adaptive hybrid fusion mechanism was developed to balance physics-based and CNN-LSTM models. The introduced exponential authority-shift law (λ) dynamically compensates for operational uncertainty, reducing the remaining useful life (RUL) prediction MAE by 20-25% compared to standalone baselines. A CSM relying on OWL 2.0 ontologies and SWRL rules was integrated into the control loop. This layer filters non-physical predictions, maintaining physical-cognitive variable consistency in 95% of computational cycles. The proposed SFDX protocol unified raw physical telemetry and ontological descriptors. This ensured consistent data interpretation and self-learning under incomplete data conditions. The implemented distributed edge-cloud architecture effectively separated real-time inference from heavy model retraining. The infrastructure demonstrated stable operation under peak telemetry loads of 15,000 messages/s with latencies within 0.25 s, meeting strict shipboard IT/OT requirements. Hardware-in-the-Loop testing validated the cyber-physical resilience of the framework. The applied multi-layered security architecture reduced the baseline risk score from 80.00 to a residual value of 9.76. The integration of standard IT/OT practices with ML-specific protections and CSM verification addresses the issue of trust assurance in maritime systems. These results confirm that the proposed HDT architecture transitions from an experimental prognostic tool into a resilient, certifiable element of shipboard control ecosystems. Future research should focus on extending the proposed HDT framework toward cooperative fleet-level maritime ecosystems, integrating multimodal sensor fusion and adaptive online learning mechanisms for continuously evolving operational conditions. Particular attention should be devoted to formal verification and certification of hybrid AI-driven shipboard systems, as well as improving resilience against coordinated cyber-physical and adversarial attacks. A promising direction is the development of probabilistic and explainable AI models capable of explicit uncertainty quantification in RUL prediction and supervisory decision-making. Further studies should also address large-scale experimental validation on real autonomous vessels and digital maritime testbeds, optimisation of distributed edge-cloud infrastructures for ultra-low-latency operation, and deeper integration of ontological reasoning with digital twin ecosystems of intelligent ports and maritime transport infrastructures.

Acknowledgements

None.

Funding

The study was not funded.

Conflict of Interest

None.

References

[1] Al-Mhiqani, M.N., Alsboui, T., Al-Shehari, T., Abdulkareem, K.H., Ahmad, R., & Mohammed, M.A. (2024). Insider threat detection in cyber-physical systems: A systematic literature review. *Computers and Electrical Engineering*, 119, article number 109489. doi: 10.1016/j.compeleceng.2024.109489.

[2] Boschert, S., & Rosen, R. (2016). Digital twin – the simulation aspect. In P. Hehenberger & D. Bradley (Eds.) *Mechatronic futures* (pp. 59-74). Cham: Springer. doi: 10.1007/978-3-319-32156-1_5.

[3] Cimino, C., Negri, E., & Fumagalli, L. (2019). Review of digital twin applications in manufacturing. *Computers in Industry*, 113, article number 103130. doi: 10.1016/j.compind.2019.103130.

[4] Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8, 108952-108971. doi: 10.1109/ACCESS.2020.2998358.

[5] Grangel-González, I., Halilaj, L., Coskun, G., Auer, S., Collarana, D., & Hoffmeister, M. (2016). Towards a semantic administrative shell for Industry 4.0. In *Proceedings of IEEE ICSC*. Laguna Hills: IEEE Computer Society Conference Publishing Services. doi: 10.1109/ICSC.2016.58.

[6] Heluany, J.B., & Gkioulos, V. (2024). A review on digital twins for power generation and distribution. *International Journal of Information Security*, 23, 1171-1195. doi: 10.1007/s10207-023-00784-x.

[7] Homaei, M., Mogollón-Gutiérrez, Ó., Sancho, J.C., Ávila, M., & Caro, A. (2024). A review of digital twins and their application in cybersecurity based on artificial intelligence. *Artificial Intelligence Review*, 57, article number 201. doi: 10.1007/s10462-024-10805-3.

[8] Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3, 422-440. doi: 10.1038/s42254-021-00314-5.

[9] Lu, Y., Liu, C., Wang, K.I.-K., Huang, H., & Xu, X. (2020). Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, 61, article number 101837. doi: 10.1016/j.rcim.2019.101837.

[10] Qi, Q., & Tao, F. (2018). Digital twin and big data towards smart manufacturing and Industry 4.0. *Engineering*, 6, 3585-3593. doi: 10.1109/ACCESS.2018.2793265.

[11] Raghavendra Babu, T.M., & Harish Kumar, K.S. (2025). Intelligent security model for digital twins: An autoencoder-based anomaly detection framework. *Journal of Information Systems Engineering and Management*, 10(56), 599-608.

[12] Suhail, S., Hussain, R., Jurdak, R., & Hong, C.S. (2021). Trustworthy digital twins in the industrial Internet of Things with blockchain. *IEEE Internet Computing*, 26(3), 58-67. doi: 10.1109/MIC.2021.3059320.

[13] Tao, F., Xiao, B., Qi, Q., Cheng, J., & Ji, P. (2022). Digital twin modeling. *Journal of Manufacturing Systems*, 64, 372-389. doi: 10.1016/j.jmsy.2022.06.015.

[14] Tyagi, A.K., Tiwari, S., Arumugam, S.K., & Sharma, A.K. (2024). *Artificial intelligence-enabled digital twin for smart manufacturing*. Beverly: Scrivener Publishing LLC. doi: 10.1002/9781394303601.

[15] Thelen, A., Zhang, X., Fink, O., Lu, Y., Ghosh, S., Youn, B.D., Todd, M.D., Mahadevan, S., Hu, C., & Hu, Z. (2022). A comprehensive review of digital twin: Modeling and enabling technologies. *Structural and Multidisciplinary Optimization*, 65, article number 354. doi: 10.1007/s00158-022-03425-4.

[16] Vychuzhanin, V., & Vychuzhanin, O. (2026). A hybrid model for evaluating the accuracy of failure forecasts in ship power plants. *BULLETIN of Cherkasy State Technological University*, 31(1), 44-59. doi: 10.62660/bcstu/1.2026.44.

[17] Vychuzhanin, V., & Vychuzhanin, O. (2025). *Digital methods and models for control and survivability of complex technical systems*. Lviv: Liha-Pres. doi: 10.36059/978-966-397-556-6.

[18] Wang, Y., Su, Z., Guo, S., Dai, M., Luan, T.H., & Liu, Y. (2023). A survey on digital twins: Architecture, enabling technologies, security and privacy, and future prospects. *IEEE Internet of Things Journal*, 10(17), 14965-14987. doi: 10.1109/JIOT.2023.3263909.

[19] Xia, Y., Xiao, Z., Jazdi, N., & Weyrich, M. (2024). Generation of asset administration shell with large language model agents: Toward semantic interoperability in digital twins in the context of industry 4.0. *IEEE Access*, 12, 84863-84877. doi: 10.1109/ACCESS.2024.3415470.

[20] Zhao, Y., Hu, Z., Du, W., Geng, L., & Yang, Y. (2024). Research on modeling method of autonomous underwater vehicle based on a physics-informed neural network. *Journal of Marine Science and Engineering*, 12(5), article number 801. doi: 110.3390/jmse12050801.

Практична IT/OT інтеграція периферійно-хмарного гібридного цифрового двійника суднових енергетичних установок

Володимир Вичужанін

Доктор технічних наук, професор
Національний університет «Одеська політехніка»
65044, просп. Шевченка, 1, м. Одеса, Україна
https://orcid.org/0000-0002-6302-1832

Олексій Вичужанін

Доктор філософії, асистент
Національний університет «Одеська політехніка»
65044, просп. Шевченка, 1, м. Одеса, Україна
https://orcid.org/0000-0001-8779-2503

Анотація. Зростаюча складність та жорсткі вимоги до надійності в режимі реального часу сучасних суднових енергетичних установок (СЕУ) роблять традиційні цифрові двійники недостатньо ефективними, якщо вони спираються виключно на фізико-математичні моделі або моделі на основі даних. Крім того, відсутність уніфікованої інтеграції IT/OT та семантичної узгодженості обмежує їх практичне впровадження в реальних умовах морської експлуатації. Метою даного дослідження були розробка та валідація фреймворку гібридного цифрового двійника (ГЦД), який забезпечує адаптивне моделювання, семантичну інтероперабельність та безпечну інтеграцію в режимі реального часу в межах кіберфізичної інфраструктури судна. Для досягнення цієї мети було запропоновано когнітивно-адаптивну архітектуру ГЦД. Вона поєднала фізичні термодинамічні моделі з компонентами CNN-LSTM на основі даних за допомогою механізму динамічного перерозподілу довіри (authority-shift), що керується невизначеністю в режимі реального часу. Фреймворк включав когнітивну імітаційну модель (CSM) на базі онтологій OWL 2.0 та правил SWRL для семантичної валідації фізичних обмежень безпосередньо в контурі управління. Додатково було впроваджено протокол семантико-функціонального подвійного обміну (SFDX) для уніфікації фізичної телеметрії та онтологічних дескрипторів. Така інтеграція когнітивної валідації та динамічного балансування моделей забезпечила стійку роботу в умовах деградації сенсорів та кіберфізичних збурень. Експериментальна валідація була проведена з використанням апаратно-програмного стенду (Hardware-in-the-Loop), що імітує реальні умови експлуатації СЕУ. Результати продемонстрували стабільну роботу при швидкості обміну даними понад 10^4 повідомлень на секунду, з наскрізною затримкою менше 0,25 секунди та втратою даних менше 0,3 %. Гібридна модель знизила середню похибку прогнозування залишкового корисного ресурсу на 20-25 % порівняно з ізольованими підходами, підтримуючи фізико-семантичну узгодженість на рівні відхилення менше 3 %. Оцінка архітектури підтвердила ефективність розподіленої парадигми «борт-межа-хмара» (onboard-edge-cloud), де управління в реальному часі виконується локально, тоді як обчислювально складні задачі перенавчання моделей делегуються на вищі рівні. Запропонований фреймворк створив масштабовану основу для надійних та когнітивно інтегрованих цифрових двійників СЕУ

Ключові слова: кіберфізичні системи; семантична інтероперабельність; адаптивне моделювання; залишковий корисний ресурс; когнітивне моделювання