A Semantic Framework for IT-OT Integration in Industrial Environments

Anand Todkar[1], Mrinmoy Sarkar[2], Jitendra Solanki[3]

[1,2,3] *Siemens Foundation Technology, Princeton, New Jersey, 08540, USA*
anand.todkar@siemens.com
mrinmoy.sarkar@siemens.com
jitendra.solanki@siemens.com

Abstract

This paper presents a semantic framework to bridge the gap between IT-OT integration in industrial environments. The proposed solution addresses fundamental challenges of PHM (prognostics and health management) by providing contextualized semantic information from the shop floor to enterprise IT systems. Built upon an OPCUA (Open Platform Communications Unified Architecture) aggregation server architecture, the framework leverages OPCUA Information Models and companion specifications as its foundation for semantic representation. By transforming these models into knowledge graphs stored in RDF format, the system enables sophisticated semantic information retrieval through SPARQL-based semantic queries that can traverse complex relationships between equipment, processes, and operational parameters. The framework further implements GraphQL to automatically generate a Type schema derived from OPCUA types, creating a unified query interface that facilitates IT-like interaction with industrial data. This semantic approach significantly improves fault diagnostics, predictive maintenance, and anomaly detection by preserving contextual relationships that are often lost in traditional data integration methods. Furthermore, the GraphQL schema provides a structured foundation for generative AI applications to formulate contextually appropriate queries, extract relevant maintenance insights, and generate human-interpretable explanations of equipment health patterns, all while maintaining semantic fidelity across the IT-OT boundary. The vertical integration capability ensures that domain-specific models remain coherent across organizational levels such as line, area, floor, etc., enabling PHM practitioners to implement more effective condition-based maintenance strategies with improved visibility into causal factors affecting equipment reliability and performance.

1. Introduction

Prognostics and Health Management (PHM) has become a critical discipline for industrial systems, focused on predicting equipment failures, optimizing maintenance strategies, and extending asset life cycles (Jardine, Lin, & Banjevic, 2006). However, the effectiveness of PHM systems is fundamentally dependent on the seamless integration between Operational Technology (OT) that generates data and Information Technology (IT) that processes these data for decision making (Wollschlaeger, Sauter, & Jasperneite, 2017). This integration challenge has persistently limited the potential of PHM implementations, particularly in complex manufacturing environments where contextual understanding of data is crucial.

Industrial systems generate a large amount of heterogeneous data between different equipment, controllers, and processes. Traditional data integration approaches have focused primarily on data transport rather than semantic and contextual information transfer, resulting in ”data lakes” that lack the contextual relationships needed for effective root cause analysis and predictive maintenance (Wollschlaeger et al., 2017). When equipment failures occur, maintenance engineers often struggle to correlate symptoms between systems due to the absence of semantic continuity between operational data and enterprise knowledge (Tao, Qi, Wang, & Nee, 2019).

The increasing complexity of modern manufacturing systems further inflates these challenges. As Industry 4.0 initiatives drive greater interconnectivity, the volume and variety of data continues to grow exponentially (Xu, Xu, & Li, 2018). Without a semantic foundation that preserves relationships between equipment, processes, and operating parameters, industry use cases around prognostics face difficulties in extracting actionable insights for health management (Xia et al., 2018).

This paper introduces a solution that addresses these fundamental challenges by providing a unified interface to represent, access, and analyze industrial data across the IT-OT boundary. Using OPCUA information models transformed into knowledge graphs and providing SPARQL and GraphQL interfaces, the framework allows more effective root cause analysis, predictive maintenance, and integration with emerging generative AI capabilities.

2. RELATED WORK

Early efforts to bridge the IT-OT gap focused primarily on physical connectivity and basic protocol translation. In the 1990s, proprietary middle-ware solutions emerged to facilitate data exchange between control systems and enterprise applications. These approaches typically involved custom connectors that mapped data points between specific systems but failed to preserve semantic context across organizational boundaries (Williams, 1994). Furthermore, the proprietary and closed nature of legacy solutions resulted in interoperability challenges, severely hindering their adoption in prognostics and health management applications, where a comprehensive end-to-end view of process cycles is essential for effective system monitoring and maintenance.

Parallel to developments in industrial communication standards, the semantic web community began exploring applications of ontologies and knowledge graphs to industrial data. The authors in (Beden, Cao, & Beckmann, 2021) proposed one of the first frameworks for semantic description of manufacturing assets, while the authors in this paper (TITAH, BOUCHAALA, AMIRATE, & BOULGHOBRA, 2023) demonstrated how semantic technologies could enhance equipment diagnostics through improved context representation.

Therefore, several researchers attempted to bridge OPCUA with semantic web technologies. The authors in (Graube, Pfeffer, Ziegler, & Urbas, 2011) and (Schiekofer, Grimm, Brandt, & Weyrich, 2019) proposed a method to map OPCUA information models to OWL ontologies, while the authors in (Cavalieri, Salafia, & Scroppo, 2019) developed a semantic middleware layer to improve OPCUA with reasoning capabilities. These approaches demonstrated the potential value of semantic integration, but typically required custom mappings for each information model and lacked scalable query mechanisms across heterogeneous systems. Existing approaches lack a unified IT like interface that supports semantics as well as access to live data from sensors.

More recent research has focused on knowledge graphs as a unifying paradigm for industrial data representation. The authors here (Zhou, Yu, & Zhang, 2015) proposed a knowledge graph-based approach to equipment failure analysis that demonstrated improved diagnostic accuracy, but required expert knowledge and engineering. Similarly, the authors in the paper (Liu et al., 2022) developed a domain-specific knowledge graph for predictive maintenance that showed promising result for specific types of equipment but lacked a general approach in diverse industrial environments.

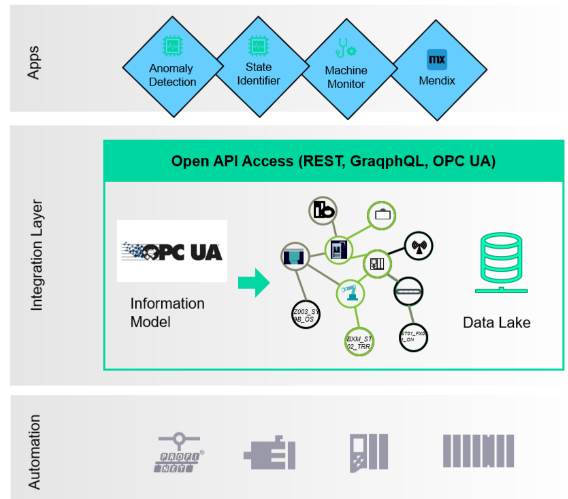

Figure 1. Overview of Industrial Information Hub Semantics

Besides the rapid evolution of large language models (LLMs) and generative artificial intelligence represents a paradigm shift for Prognostics and Health Management (PHM) systems in industrial environments. These sophisticated AI architectures, capable of understanding context, generating natural language explanations and identifying complex patterns across diverse datasets, fundamentally change the way maintenance professionals interact with industrial systems and interpret equipment health data. Recent research by (Zhang, Sun, Chen, & Jiang, 2024) has demonstrated promising applications of transformer-based architectures for equipment failure prediction, achieving significant improvements in both accuracy and lead time compared to traditional machine learning approaches. These models have shown particular strength in identifying subtle precursors to equipment degradation that conventional algorithms often miss, especially in scenarios involving multiple interacting subsystems with complex dependencies.

However, the full potential of generative AI in PHM remains constrained by critical limitations in data representation. Unlike traditional analytics that can operate on structured tabular data, advanced generative models require semantically rich, contextually complete representations that capture the complex interrelationships between equipment states, operational parameters, and historical performance patterns. Current industrial data architectures typically fail to provide these rich representations that can be effectively queried by AI systems.

3. PROPOSED SOLUTION: INDUSTRIAL INFORMATION HUB SEMANTICS

As shown in Figure 1, the semantic framework presented in this paper differs from previous approaches in several key aspects. The framework leverages the inherent structure of OPCUA information models, including companion specifications, as the basis for semantic representation, thus eliminating the need for custom ontology development [1], and automates the transformation of these OPCUA models into knowledge graphs while fully preserving their semantic richness, such as type hierarchies, references, and metadata, as described in (Schiekofer et al., 2019). It provides automatically generated SPARQL and GraphQL interfaces, facilitating semantic queries for advanced analytics and structured application integration without requiring specialized query expertise, and supports vertical integration of domain-specific semantic representations across shopfloor and plant boundaries (Schmied, Großmann, Mathias, & Banerjee, 2020). Additionally, the automatically generated GraphQL schema derived from OPCUA types offers a structured foundation that enables generative AI systems to formulate contextually appropriate queries even without domain-specific training, creating a distinctive capability suitable for dynamically changing production processes. Unlike earlier approaches, which focused on either data connectivity mechanisms or isolated semantic models, this integrated framework maintains contextual semantic integrity throughout the entire data journey: from operational sensor networks through enterprise analytics to AI-driven decision support systems, thus achieving unprecedented continuity between physical operations and high-level business intelligence applications.

4. VALIDATION AND PROTOTYPE OF INDUSTRIAL INFORMATION HUB SEMANTICS

The development of IIH Semantics began with an in-depth evaluation of existing technologies and approaches for industrial IT-OT integration. Based on understanding of the state-of-the-art solutions available, the problem of having a unified plant-wide interface that serves semantic information as well as an access to live values from machines and sensors was identified to be the core use case to focus on. Based on this, the following starting points were identified.

- Identified OPCUA as the most promising foundation due to its information modeling capabilities and its close connection with industrial communication.
- Team recognized the potential of knowledge graphs to preserve semantic relationships.
- GraphQL was selected as the query interface based on its compatibility with modern application development paradigms and its capacity to represent industrial information using structures and patterns readily consumable by enterprise IT systems such as MES and ERP platforms.

[1]https://reference.opcfoundation.org/Core/Part5/v104/docs/

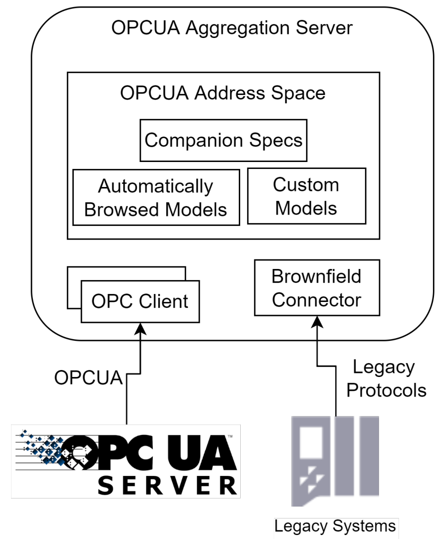

Figure 2. Internal Components Of OPCUA Aggregation Server

4.1. OPCUA Aggregation Server

The evaluation phase began with an assessment of OPCUA Aggregation capabilities as the foundation for the semantic integration layer. As shown in Figure 2 the server would serve as the connection point to source OPCUA Servers as well as to devices which don't have OPCUA capabilities. While the southbound connectivity with devices is out of the scope for this paper, the important design decision was to have connectivity with all types of devices including greenfield which have OPCUA servers and brownfield systems which do not have this capability available.

4.2. Semantic Transformation: OPCUA Information Model To Knowledge Graph

The second major component focused on transforming structured OPCUA information models into a semantic knowledge graph representation that preserves relationships while enabling sophisticated queries. After extensive evaluation of potential mapping methodologies, the team adopted the formal mapping approach proposed in (Schiekofer et al., 2019), utilizing Apache Fuseki as the RDF store to realize the knowledge graph implementation. Initially developed in Java, the transformation process enabled programmatic conversion of OPCUA models into RDF, successfully addressing several technical challenges encountered during implementation, such as handling complex reference types, preserving inheritance relationships between object types, managing hierarchical relationships without sacrificing query performance, and accurately interpreting modeling rules defined in companion specifications. Overall, the mapping approach successfully converted OPCUA's hierarchical structures into

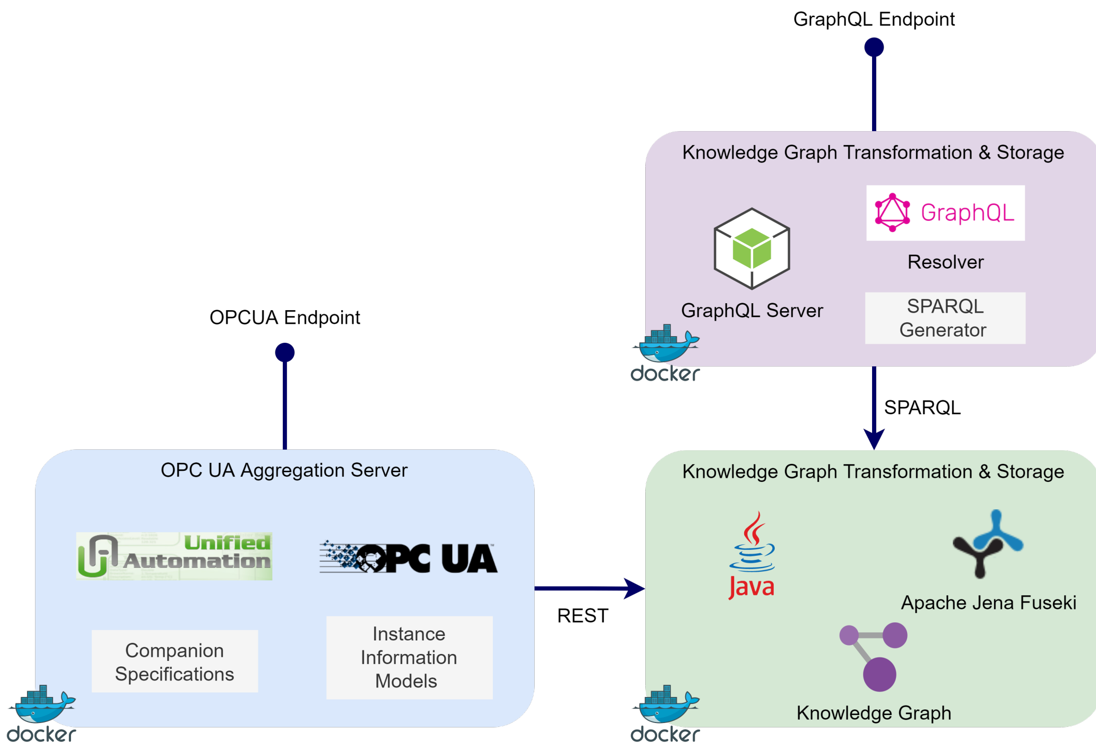

Figure 3. Architectural overview of the IIH Semantics evaluation system showing containerized components and their interactions

a graph-based representation while fully maintaining their semantic integrity.

4.3. GraphQL Schema And Query Resolution

This component focused on creating a dynamic and type-safe query interface based on the underlying OPCUA types to make semantic information accessible to applications. The initial implementation employed a hard-coded GraphQL schema, with SPARQL queries automatically generated based on schema queries corresponding to OPCUA types, and utilized Node.js along with GraphQL libraries to ensure ease of use. As the approach evolved, the schema generation transitioned to a dynamic process, enabling type schemas to be generated directly from the knowledge graph and empowering complex SPARQL queries to automatically extract type information immediately following the OPCUA-to-knowledge-graph translation. This advancement eliminated the need for manual schema maintenance, becoming a fundamental feature of the application, and enabled the resulting system to dynamically and seamlessly adapt to changes in the underlying OPCUA information models.

4.4. Integration Architecture

The proof-of-concept deployment as shown in Figure 3 used a containerized architecture with three separate containers for modularity and scalability. The system architecture adopted a modular containerization strategy, implementing each core component as an independent container: an OPCUA Aggregation Server for efficient data collection, a Knowledge Graph transformation and storage container for semantic data processing, and a dedicated GraphQL schema generation and query server container providing flexible API access. Communication among containers was facilitated through HTTP-based interfaces, enabling independent scaling, streamlined maintenance, and tailored deployments that responded to varying performance requirements in different environments.

4.5. Insights Derived from the Evaluation Stage

The comprehensive evaluation of the proof-of-concept implementation delivered key insights, confirming that the proposed semantic integration framework effectively bridged semantic gaps typically existing across industrial organizational boundaries, with OPC UA Aggregation Server reliably providing operational data and preserving essential contextual metadata, while knowledge graph transformation processes successfully retained semantic relationships and supported flexible querying, complemented by dynamic GraphQL schema generation which eliminated manual schema maintenance overhead. However, evaluation results revealed significant performance issues arising from

HTTP-based communication boundaries between components, as inter-container communication introduced latency that impeded real-time data delivery, resulting in inconsistent sub-millisecond latency performance and network overhead bottlenecks during high-throughput operations. Consequently, an architectural reassessment identified the modular monolith pattern as more suitable to substantially reduce inter-container communication latency by eliminating unnecessary network interactions, preserving logical module separation, increasing overall responsiveness, boosting throughput capacity, streamlining deployment and configuration management, and effectively satisfying stringent performance demands typical of industrial monitoring and control scenarios. Though container-based deployment provided notable operational flexibility and portability advantages, evaluation indicated considerable performance tuning and resource optimization are required to meet the reliability and throughput criteria for industrial production usecases, thus critical architectural and performance enhancements are considered in the subsequent design phase.

5. Production Implementation Based on Evaluation Assessment

Based on the critical insights gained during the evaluation phase, a significant architectural redesign was implemented for the production version of IIH Semantics. This new approach addressed the performance constraints and complexity challenges identified during proof-of-concept.The revised architecture maintained the core semantic processing capabilities while implementing structural optimizations that significantly enhanced system responsiveness, deployment simplicity, and operational reliability in industrial environments.

5.1. Consistent Technology Stack

The production architecture, illustrated in Figure 5, transitioned to a cohesive technology stack entirely based on the C# programming language, replacing the earlier heterogeneous environment. Specifically, the Unified Automation SDK used previously was replaced by the OPCFoundation SDK for the aggregation server, providing direct access to native OPCUA functionality, closer alignment with the official OPCFoundation specifications, and reduced licensing constraints through direct consortium collaboration. The OPC-to-RDF transformation was redesigned as a dedicated C# library, improving algorithm implementation and enabling seamless in-process integration with other C# modules, thus facilitating easier OPC server context sharing and achieving a 50% performance improvement over the initial proof-of-concept. Finally, the previously employed Node.js GraphQL implementation was replaced by dotnetRDF, which streamlined schema generation directly from C# types, improved query-resolution

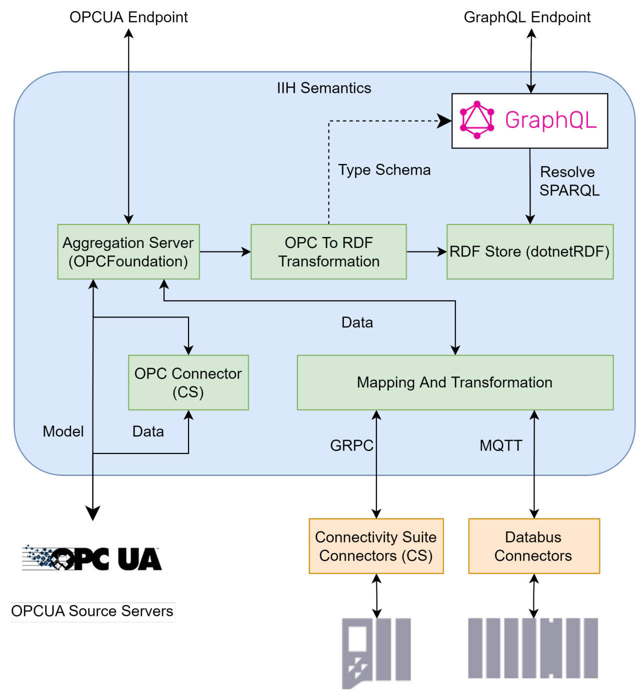

Figure 4. Industrial Information Hub Semantics Application Architecture

performance with in-process operation, and eliminated the overhead associated with cross-language data serialization.

5.2. Modular Monolith Architecture

The updated implementation transitioned to a modular monolith architecture pattern to effectively address performance limitations identified during the evaluation phase. All components were integrated within a single Docker container, maintaining logical separation through clearly defined internal interfaces rather than isolated deployments. HTTP-based intercomponent communication was replaced by direct method calls, and shared memory access eliminated serialization and deserialization overhead. As a result, runtime performance improved substantially, particularly for operations involving multiple components.

5.3. Enhanced Mapping and Transformation Engine

A significant enhancement in the production implementation was the introduction of a comprehensive mapping and transformation engine, enabling dynamic transformation of live data values tailored specifically to consumer needs, including capabilities like unit conversion and normalization, coordinate system transformations, mathematical transformations, and context-based data enrichment. Additionally, the engine supported bidirectional mapping between different semantic representations, facilitating seamless translation between industry-specific models and adaptation to diverse IT system data formats while carefully preserving semantic integrity across transformations. Furthermore, the implementation allowed runtime configuration of transformation rules without necessitating system restarts, greatly improving flexibility and responsiveness to evolving operational requirements.

5.4. Performance Improvements

The architectural changes delivered substantial performance improvements across multiple key metrics, significantly enhancing the overall system effectiveness. Latency reduction was achieved by eliminating HTTP overhead, resulting in more than a 90% reduction in internal communication latency, consistently maintaining sub-millisecond response times for live data values, and accelerating end-to-end query response times by 60-80% for common operational scenarios. Concurrently, throughput enhancement was notable, as the shared memory architecture facilitated handling of 10 times more concurrent requests, knowledge graph operations executed 3-4 times faster due to efficient in-memory processing, and GraphQL query resolution speed improved by approximately 70%. Additionally, substantial resource utilization benefits included a 40% reduction in overall memory footprint due to optimized resource sharing, a 20% decrease in CPU usage for equivalent workloads, and container startup time improving from minutes down to mere seconds.

5.5. Southbound Connectivity

In addition to the greenfield connectivity with OPCUA sources, IIH Semantics supports two additional types of southbound connectivity layers named as GRPC based Connectivity Suite and MQTT based Databus respectively.

5.5.1. Connectivity Suite

The Connectivity Suite represents a specification based on Protocol Buffers (Protobuf) and gRPC technologies [2] that defines a consistent communication framework. This specification establishes how connectors share metadata and live data with consuming applications, creating a unified interface layer that abstracts the underlying complexity of diverse industrial protocols. IIH Semantics utilizes the Connectivity Suite as a foundational component, consuming structured data streams through this standardized interface and transforming them into semantically rich information. Details on internals of Connectivity Suite Connectors is beyond the scope of this paper.

5.5.2. Databus

The Databus specification provides a standardized messaging infrastructure for Industrial Edge applications based on publish-subscribe communication patterns, offering high-throughput, reliable message transport for operational technology environments. Databus connectors serve as integration points between this messaging backbone and various data sources and consumers, supporting MQTT protocol. IIH Semantics integrates with the Databus connectors, consuming data streams over MQTT communication. Detailed architecture of Databus is out of scope for this paper.

[2]https://support.industry.siemens.com/cs/document/109817512/connectivity-suite

6. IIH SEMANTICS DEPLOYMENT ON SIEMENS INDUSTRIAL EDGE

Siemens Industrial Edge [3] is a managed platform for deploying containerized applications in industrial environments, bringing IT capabilities directly to the shop floor. It consists of three main components:

- **Industrial Edge Management (IEM)**: A centralized platform for managing applications, devices, and configurations across distributed manufacturing sites.
- **Industrial Edge Devices (IED)**: Hardened hardware or virtual machines that run containerized applications near production equipment.
- **Industrial Edge Applications**: Containerized software that provides specific functionality for industrial use cases.

Industrial Edge facilitates the use of modern IT technologies within operational technology (OT) environments, ensuring robust security boundaries and compliance with industrial-grade performance standards. The platform provides a unified interface for device lifecycle management, streamlined application deployment, and comprehensive monitoring capabilities. Although Industrial Edge encompasses a wide range of functionalities, this paper primarily discusses the platform from the app developer's standpoint, emphasizing how these capabilities are leveraged specifically for deploying the IIH Semantics application.

Deployment on Industrial Edge utilizes a structured and streamlined process tailored specifically for industrial settings, starting with applications being packaged as Docker containers along with a Docker Compose file that clearly defines application structure, configuration options, and resource requirements. These packaged containers are published to the Industrial Edge Hub, serving as a centralized repository that makes applications available for direct deployment. Users can browse this centralized catalog via Industrial Edge Management to acquire applications intended for specific edge devices. Prior to deployment, standardized interfaces support comprehensive configuration management, enabling users to adjust parameters as specified by the application developer. Industrial Edge Management orchestrates

[3]https://docs.eu1.edge.siemens.cloud/

the secure and reliable transfer of both container images and their configurations to the targeted edge devices, managing versioning and providing rollback capabilities as required. Following deployment, applications are continuously monitored within the Industrial Edge platform, supported by integrated logging, metrics collection, and health tracking capabilities. This well-defined process ensures consistent, secure, and governed application delivery across distributed manufacturing sites, substantially reducing operational complexity and overhead.

The modular monolith architecture of IIH Semantics was packaged within a single Docker container, incorporating OPCFoundation SDK for aggregation functionality, dotnetRDF for knowledge graph processing, and GraphQL.NET to facilitate query interfaces. Resource requirements were clearly outlined through a Docker Compose file tailored to the industrial context, defining specifications such as memory allocation optimized for knowledge graph operations, network port mappings for OPCUA (port 62520), and configurations for a reverse proxy to securely expose HTTP interfaces externally beyond the Industrial Edge device. Furthermore, the Industrial Edge App Publisher tool was employed to package the IIH Semantics application, validating compatibility with the Industrial Edge runtime, automatically generating appropriate metadata and documentation, and creating a deployable application package. The resulting application was published to the production hub, enabling customers to conveniently download it into their respective Industrial Edge Management environments for subsequent deployment to edge devices.

7. VERTICAL INTEGRATION USING IIH SEMANTICS

The implementation of IIH Semantics addresses key semantic integration challenges for individual industrial edge installations. However, modern manufacturing environments typically span multiple facilities with distributed assets that require coordinated data exchange. This necessitates exploring how IIH Semantics fits within a broader federated information architecture.

IIH Sync functionality in IIH Semantics app enables federated information architecture. IIH Sync extends the capabilities of IIH Semantics by enabling synchronized connections between edge devices in a hierarchical architecture. Current implementations use a straightforward duplication approach where child devices (with their own IIH Semantics instances) synchronize complete OPCUA information models to a parent device, creating a centralized aggregation point as explained in paper (El Kalach, Solanki, & Todkar, 2024). This simple yet effective model enables organizations to collect semantic information from distributed sources while maintaining a unified view at the enterprise level.

The federated information system proposed by authors in (El Kalach et al., 2024) offers a more sophisticated approach

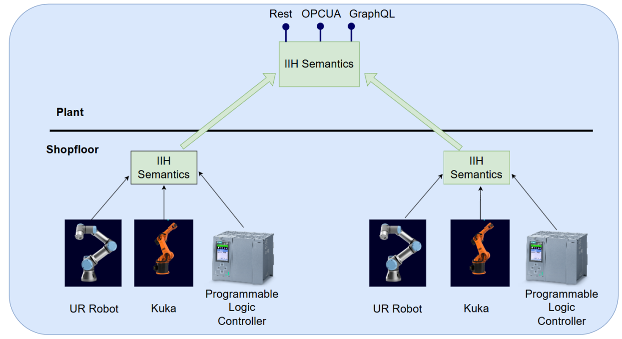

Figure 5. Industrial Information Hub Vertical Aggregation

that goes beyond the current implementation. Their framework envisions autonomous components that can independently perform semantic integration while participating in a coordinated federation. Unlike the current model where complete information models are duplicated, the proposed abstraction-level synchronization would allow selective sharing of semantic information based on contextual relevance and query demand.

The current IIH Semantics implementation relies on complete model duplication between child and parent devices, which, while functional, faces several critical limitations: limited support for autonomous semantic reasoning at each node, difficulties in selective information sharing based on context, and problematic access to historical data. Future development will focus on implementing the abstraction model synchronization approach proposed in (El Kalach et al., 2024), which would enable more efficient federation by synchronizing only the required semantic abstractions rather than entire models. Such an approach would better support scenarios where different organizational units need to maintain autonomy while still contributing effectively to a collective semantic understanding, ultimately creating a more flexible and scalable federated semantic architecture.

8. CONCLUSION

The convergence of Information Technology (IT) and Operational Technology (OT) offers transformative opportunities for industrial environments, yet integration efforts continue to be hampered by fundamental semantic gaps where contextual meaning is lost despite data flow between systems. This semantic discontinuity creates information silos that obscure crucial relationships between equipment, processes, and parameters, significantly limiting industrial digitalization initiatives, particularly in complex manufacturing settings. Semantic information provides the essential contextual framework that transforms raw data into actionable knowledge by preserving relationships, hierarchies, and domain-specific meanings—without which maintenance engineers struggle to correlate symptoms across systems, resulting in extended down-time and reduced equipment effectiveness. As manufacturing systems grow increasingly complex with Industry 4.0 initiatives, establishing this semantic foundation becomes critical for extracting meaningful insights from industrial data and achieving the full potential of IT-OT integration.

IIH Semantics addresses this challenge by providing a comprehensive semantic integration layer that bridges the IT-OT gap. By transforming OPCUA information models into knowledge graphs and providing both SPARQL and GraphQL interfaces, the solution creates a unified semantic representation that preserves contextual relationships while making this information accessible to both specialized analytics tools and mainstream IT applications. The modular monolith architecture delivers the performance required for real-time industrial applications while maintaining logical separation between components.

The integration with Industrial Edge provides a deployment mechanism that respects the operational requirements of manufacturing environments while enabling centralized management. By creating a foundation where context-rich industrial knowledge flows seamlessly between operational systems and enterprise applications, IIH Semantics enables more effective root cause analysis, predictive maintenance, and operational optimization—representing the missing link that allows manufacturers to fully capitalize on their industrial data for truly intelligent manufacturing systems.

REFERENCES

Beden, S., Cao, Q., & Beckmann, A. (2021). Semantic asset administration shells in industry 4.0: A survey. In *2021 4th ieee international conference on industrial cyber-physical systems (icps)* (pp. 31–38).

Cavalieri, S., Salafia, M. G., & Scroppo, M. S. (2019). Integrating opc ua with web technologies to enhance interoperability. *Computer Standards & Interfaces*, *61*, 45–64.

El Kalach, F., Solanki, J., & Todkar, A. (2024). A federated information system framework for vertical integration. *Manufacturing Letters*, *41*, 1192–1199.

Graube, M., Pfeffer, J., Ziegler, J., & Urbas, L. (2011). Linked data as integrating technology for industrial data. In *2011 14th international conference on network-based information systems* (pp. 162–167).

Jardine, A. K., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical systems and signal processing*, *20*(7), 1483–1510.

Liu, M., Li, X., Li, J., Liu, Y., Zhou, B., & Bao, J. (2022). A knowledge graph-based data representation approach for iiot-enabled cognitive manufacturing. *Advanced Engineering Informatics*, *51*, 101515.

Schiekofer, R., Grimm, S., Brandt, M. M., & Weyrich, M. (2019). A formal mapping between opc ua and the semantic web. In *2019 ieee 17th international conference on industrial informatics (indin)* (Vol. 1, pp. 33–40).

Schmied, S., Großmann, D., Mathias, S. G., & Banerjee, S. (2020). Vertical integration via dynamic aggregation of information in opc ua. In *Intelligent information and database systems: 12th asian conference, aciids 2020, phuket, thailand, march 23–26, 2020, proceedings 12* (pp. 204–215).

Tao, F., Qi, Q., Wang, L., & Nee, A. (2019). Digital twins and cyber–physical systems toward smart manufacturing and industry 4.0: Correlation and comparison. *Engineering*, *5*(4), 653–661.

TITAH, M., BOUCHAALA, M. A., AMIRATE, A., & BOULGHOBRA, C. (2023). Semantic technology for intelligent industrial maintenance: A knowledge-driven approach. *The 1st National Conference on Industrial Production and Maintenance Engineering*.

Williams, T. J. (1994). The purdue enterprise reference architecture. *Computers in industry*, *24*(2-3), 141–158.

Wollschlaeger, M., Sauter, T., & Jasperneite, J. (2017). The future of industrial communication: Automation networks in the era of the internet of things and industry 4.0. *IEEE Industrial Electronics Magazine*, *11*(1), 17-27. doi: 10.1109/MIE.2017.2649104

Xia, T., Dong, Y., Xiao, L., Du, S., Pan, E., & Xi, L. (2018). Recent advances in prognostics and health management for advanced manufacturing paradigms. *Reliability Engineering & System Safety*, *178*, 255–268.

Xu, L. D., Xu, E. L., & Li, L. (2018). Industry 4.0: state of the art and future trends. *International journal of production research*, *56*(8), 2941–2962.

Zhang, X., Sun, W., Chen, K., & Jiang, R. (2024). A multimodal expert system for the intelligent monitoring and maintenance of transformers enhanced by multimodal language large model fine-tuning and digital twins. *IET Collaborative Intelligent Manufacturing*, *6*(4), e70007.

Zhou, A., Yu, D., & Zhang, W. (2015). A research on intelligent fault diagnosis of wind turbines based on ontology and fmeca. *Advanced Engineering Informatics*, *29*(1), 115–125.