



Edge-Driven Predictive Maintenance Architecture for High-Stakes Industrial Networks

Sophia Martinez

Asst Professor

Abstract

A distributed cognitive maintenance framework for mission-critical networks is proposed that integrates existing automation and supervisory systems in order to improve their maintenance. Maintenance is neither an end-user activity nor an operation, but a supported and automated activity that enhances the operation of the mission-critical systems. Maintenance activities include predictive and proactive maintenance, real-time fault detection and diagnosis, resource-aware scheduling and allocation, and human-actor collaboration with supervisory control. To achieve this, data-sensing systems issue data in a sentient fashion to cognitive agents, and reasoning agents send information for collective decision making. Control systems, especially for closed-loop operations, are taken into consideration.

Mission-critical networks serve production and product operations in real time. They are deployed and operated in an automation mode that may also ensure adequate redundancy for fault tolerance. Nevertheless, systems and facilities remain without any of fundamental principles on how to perform obsolescence management, preventive maintenance or even fault detection and diagnosis. Communication between such systems is commonly limited to supervisory control and the exchange of operational parameters. A different view considers maintenance as a mission-supported activity closely integrated with the operation.

Keywords : Distributed cognition, Cognitive maintenance, Mission-critical networks, Industrial network resilience, Predictive maintenance systems, Autonomous fault diagnosis, Edge intelligence, Industrial IoT reliability, Adaptive network orchestration, Cyber-physical system maintenance.

1. Introduction

Simulation-based testing cannot replace real-life experience, but realistic simulations are necessary prerequisites to performing a field test. Digital twin technology is the natural solution to address such problem. Connecting the physical assets to a digital model will enable virtual testing of equipment and procedures, identifying weak points and potential problems before they manifest in the real world. While tests in a digital, virtual environment will be very beneficial, it is still necessary to perform a limited number of field tests. No matter how realistic the simulation, unintended issues or interaction with other equipment can cause unexpected failures.

Other staff-related requirements such as the amount of staff required to perform a test and supervisor's approval also need to be prepared and arranged in advance. Legit position of those test have to be clearly communicated to ensure that no other event is prioritised. Risk of failure must be communicated to have mitigations already arranged in the case of an unforeseen failure. Validated procedures need to be created prior to the test, pointing out possible mistakes that might happen and how those should be avoided. All steps are useful to guarantee the best possible outcome of a field event that is still a unique chance for achievement in a real environment. Truly Predictive substituting the tasks focused on knowing when to act on degraded systems with those focused on detecting before failure and act accordingly.




Fig 1: Distributed Cognitive Maintenance Framework

1.1. Simulation and Digital Twin Integration

The internal structure of complex systems such as transportation networks or power generation and distribution systems is generally not visible in real-time. Thus, in addition to the deployment of physical sensors, realistic simulation can support performance assessment. Simulation as a service can be offered to application stakeholders. Suitable simulator abstractions reside in the cloud and can be instantiated by typical end-user devices. The interface between the physical and simulation systems are the real-time sensor data (for feeding the simulator) and the resulting predictions (for feeding the physical system). Basic application-aware simplifications can also be realised. For example, for an electrical distribution network, specific equipment states may be replaced by their actual performance metrics. Alarms and critical status notifications generated by the event-based digital twin mapping over the physical system can be monitored in simulation mode to test possible decision-making actions.

Equation 1: Predictive Maintenance Remaining Useful Life (RUL)

This equation estimates the Remaining Useful Life of industrial assets using degradation monitoring:

$$RUL(t) = T_f - t$$

Where:

- $RUL(t)$ = Remaining Useful Life at time t
- T_f = Predicted failure time
- t = Current operational time

This relates directly to predictive maintenance and degradation analysis discussed in the framework.

Inline visualization:

$$RUL(t) = T_f - t$$

1.2. Benchmarking Scenarios and Metrics

Several benchmarking scenarios for a distributed cognitive framework in mission-critical industrial systems are described along with associated metrics. Service- and response-time-related measures support framework implementation performance assessment. Metrics associated with cognitive maintenance cover improvements achieved through a distributed approach. Efficiency-related measures, reflecting resource consumption and quality, are defined for maintenance operations.

Mission-critical industry networks consist of resources that are responsible for the generation and distribution of energy for other industrial processes. Therefore, the operational execution of these resources is extremely time-sensitive and any planned or unplanned maintenance activity has associated completion time and effort specifications. Deviating from these specifications results in wasted (or unaimed) efforts and surplus consumption of resources (and/or specialists). The following benchmarking scenarios represent a framework implementation in an industrial environment. Although these scenarios relate directly to an integrated proof-of-concept implementation, they can also serve as general evaluation measures for the underlying cognitive maintenance approach.




Table 1: Core Components of the Distributed Cognitive Maintenance Framework (DCMF)

Component	Description	Key Functions
Distributed Cognition Engine	Coordinates intelligent decision-making across industrial networks	Collective reasoning, coordination, adaptive response
Digital Twin System	Virtual replica of physical infrastructure	Simulation, predictive testing, fault forecasting
Predictive Maintenance Module	Forecasts failures before occurrence	Remaining Useful Life (RUL) estimation, condition monitoring
Real-Time Fault Detection System	Detects and diagnoses failures during operation	Fault isolation, anomaly detection
Resource-Aware Scheduler	Optimizes resource allocation	Workforce scheduling, spare-part allocation
Human-Supervisory Control	Human-in-the-loop oversight	Validation, emergency intervention

1.3. Field Trials and Deployment Considerations

Although, as a primary criterion, the approach must be applicable to scenarios where the cognitive maintenance system is widespread, integrated, tested, and validated, the dwelling on all aspects of a real-world validation is neither feasible nor useful. A balanced selection of elements is provided here, ensuring that the exposition is both concise and clear, without compromising coverage and attention to detail.

Depending on the application, domain, and chosen architecture, prototypes or solutions that fall under the parameters of the proposed approach may have already been designed, implemented, tested, or validated. Therefore, detailed experimental results are rarely reported yet with a complete focus on cognitive maintenance—as envisaged here—rather on adaptive, computerized, and autonomous maintenance. The operation set tends to include, but is not limited to, concepts of predictive, proactive, and real-time maintenance, or combinations thereof, such as predictive-control-mission-assisted maintenance. Therefore, pioneering and supporting proposals that—and the effort aims at combining and achieving all of these principles—should ultimately focus on fulfilling the three highlighted conditions supporting the full definition of cognitive maintenance. Where these conditions are met, activity implementation should fit the definition as well.”



2. Background and Motivation

Power Generation and Distribution Networks

Mature industrial societies are critically dependent on their infrastructure and the systems that manage it. Within this context, power generation, distribution, and control networks are vital for any industrial plant. Exploiting cogeneration with other adjacent manufacturing processes can also facilitate a reduction of the power footprint and optimize production cost.



Generating electrical energy using hydro turbine technology is often proposed. Canal-based reservoirs drive the hydraulic turbines and the generated electrical energy is mainly used for supplying industrial plants located nearby, with an average production of around 3 MW. An electronic controller adjusts the rotational speed of the generators (parallel with the grid) that are used for power distribution. For industrial plants, a local automatic voltage regulator, combined with blackberry technology, maintains the power quality by injecting corrective power in response to local demand. The local dedicated digital twin presents sudden requests for power following spontaneous behaviors of some manufacturing processes, while the long-term electrical demand is well known.

Manufacturing Execution and Plant Automation

In the manufacturing domain, large groups of workers are employed inside factory plants executing a variety of operations. Flexible manufacturing systems that allow for reconfiguration of the assembly line integrate equipment of different kinds. Maintenance and logistics processes are known to be the most important problems during operation. A group of collaborative robots or other intelligent agents operating in the same area can manage material streams, overcoming missing components in the automation cells and correcting potential errors. The dedicated digital twin synthetically represents these behaviors. Recent industrial trials conducted inside a big industrial plant confirmed the complexity of the process and the importance of the human supervisor. Real human-technology collaboration can take place through the interaction of supervisors with the equipment that is recreating a representation during the main output phases, recreating actual teamwork.

Critical Infrastructure Monitoring

Public services, such as water distribution networks, are also provided from city management departments. Very big companies usually own the structures for Internet connections. High security levels are guaranteed inside these critical infrastructures, normally allowing just a limited group of authorized personnel to access the operation area. A

dedicated two-action embedded device is proposed for online security control. An expected human-technology interaction is integrated into the process. Collaboration occurs when the device warns a nearby authorized person to take corrective actions. Environmental accidents are considered in a preventive way. Whenever a big catastrophic event is recorded, the automatic property detection activity should be operational.

Equation 2: Reliability Function for Mission-Critical Systems

A common reliability equation for industrial networks:

$$R(t) = e^{-\lambda t}$$

Where:

- $R(t)$ = Reliability at time t
- λ = Failure rate
- t = Operating time

This supports resilience and fault-tolerant architecture evaluation in the paper.

$$R(t) = e^{-\lambda t}$$

A

k

$$y = Ae^{-kt} = 6e^{-0.6t}$$

2.1. Power Generation and Distribution Networks

Current models for these networks frequently employ centralized control strategies but cognitive approaches support the deployment of distributed decision-making processes. Interfaces with other mission-critical networks, such as those present in the monitoring of dams or electrical distribution systems at a lower level, offer further deployment avenues. Such interfaces enable the conservation of human-mactor expertise while relying on the supplementary



capacities offered by introducing data-driven analysis methods.

Current distributed control solutions focused on fault detection, isolation, and recovery rely on a combination of rehearsed plans executed with limited human oversight. However, deployment in environments with limited historical data, such as embarkation on new vessels, highlights the limitations of such systems. A cognitive architecture that draws inspiration from socially enhanced learning enables the retention of human supervisory capacity while adding resource-aware scheduling and allocation of recovering resources. The architecture incorporates the aspects of distributed modelling, maintenance needs identification, prediction of remaining life, strategy selection, and real-time diagnosis through the use of Data-Driven Digital Twin-assisted Decision Support Simulators.




Fig 2: Electrical power system performing generation

2.2. Manufacturing Execution and Plant Automation

While Manufacturing Execution Systems (MES) and Plant Automation Platforms (PAP) are vital components of manufacturing, related research has largely focused on either the information systems layer or the automation layer. In practice, both are inseparable; since the automation layer must execute schedules from the information systems layer, the real-time conditions measured by the automation layer should feed back into the information systems layer to fulfill its responsibilities seamlessly, yet this information is often

underutilized. These two information sources are independently handled by MES and PAP products, leading to redundant installations generating separate data silos that hinder an integrated view of the overall position and operational state of the plant.

An integrated cognitive maintenance framework bridges this gap, addressing issues such as trigger-based inspection scheduling using real-time plant status information, human-actor coordination using live video feeds from unmanned machines being investigated, what-if scenario creation and monitoring system safety during inspection and maintenance activities. Deterioration models should also be evaluated for transitioning from predictive to proactive maintenance. Given the safety-critical nature of MES, any disturbances created during diagnostic and maintenance activities should be reconciled by the MES, hence the integrated nature. When an inspection is triggered, the trigger condition itself, along with the scheduled task, should be fed into the MES to ensure safe execution of other related manufacturing tasks.

Table 2: Maintenance Types and Their Characteristics

Maintenance Type	Objective	Techniques Used	Benefits
Predictive Maintenance	Predict future failures	AI/ML models, sensor analytics	Reduced downtime
Proactive Maintenance	Prevent degradation before failure	Trend analysis, environmental monitoring	Extended equipment life
Condition-Based Maintenance	Trigger maintenance using live data	Real-time sensing	Reduced maintenance costs
Real-Time Fault Detection	Immediate identification of failures	Diagnostics, anomaly detection	Faster recovery




2.3. Critical Infrastructure Monitoring

The emergence of the Internet of Things (IoT) is triggering a paradigm shift in the way critical infrastructures are monitored. Data are harvested, stored, and analyzed in real-time and at all system levels to feed analysis algorithms and human operators with information useful for detecting and mitigating short-term or long-term threats. World-wide, the adoption of smart meters and intelligent sensors enables utilities to analyze the distribution grid performance and power quality while tracking resource consumption. In energy generation, condition-based monitoring of safety- and mission-critical assets enables operators to intelligently allocate resources and prioritize actions in a way such that these resources have the highest impact on safety of personnel, protection of the environment, and business continuity. In factories, the accurate modeling of production assets, the definition of Key Performance Indicators (KPIs), and the implementation of scalable real-time monitoring systems contribute to making manufacturing systems adaptive and down-time risk analysis more effective.

Monitoring of Critical Infrastructures (CI M) can go even further. In addition to detecting risks when they occur and predicting when the deteriorating condition of a monitored asset could lead to a failure, it can also enable equally distributed and well-planned actions to care for the integrity and availability of all monitored assets based on their critical role. In other words, all relevant assets could constantly be under careful attention and predefined, planned, and controlled actions may be taken before preventive maintenance becomes necessary. This is similar to the preservation of a City Security and Surveillance Control Room (CSSC) that smartly combines data coming from an advanced collection of environmental sensors to preserve the integrity of its main structure and support a service level for all its citizens.

Pie Chart: Maintenance Function Distribution



3. Foundational Concepts

A comprehensive understanding of mission-critical industrial networks demands not only an exploration of machining resources, plant conditions, and power supplies, but also a conceptual background that considers networks for maintenance tasks, particularly the four recognized maintenance nodes. These nodes include a cognitive maintenance management network with data about incoming sensors and data for forward actions and alerts, a real-time fault detection and diagnosis network with data for presently breaking components along with possible future affected components and resource-allocation information, a predictive and proactive maintenance network with condition-monitoring data and predicted component-state information, and a human-mactor collaboration network, also with a supervisory control function. The cognition elements on these networks need a clear, intuitive grouping to address cognizance more efficiently.

Operating principles of distributed cognition in industrial environments and fields of application where cognitive maintenance concepts can bring benefits offer supportive insights. Such maintenance, even for business-impacting

critical, safety-related systems, can be enabled in operator networks or cross-network relations, with recognition that safety-critical systems often interact with non-safety-critical systems. The risk hierarchy of the maintenance cycle enables evolution in prevailing habits, separating normal corrective actions from critical risk-oriented actions. This enhances efficiency and reduces the knowledge gap for maintenance teams beyond just the safety-critical environment, once communication aspects are fulfilled. Communication aspects may require additional attention and detail with maintenance decision processes, though.

Equation 3: Sensor Data Fusion Model

For distributed cognition and multi-source sensing:

$$D_f = \sum_{i=1}^n w_i d_i$$

Where:

- D_f = Fused data estimate
- d_i = Sensor data from source i
- w_i = Weight associated with source reliability
- n = Number of sensors

This equation reflects the data collection and fusion mechanisms described in the DCMF architecture.

$$D_f = \sum_{i=1}^n w_i d_i$$

3.1. Distributed Cognition in Industrial Environments

Distributed Cognition in Industrial Environments

Industrial networks and systems are becoming increasingly distributed, flexible, complex, and interactive, as well as mission or business critical. An extension of distributed cognition theory can be applied where the actors and the system, collectively known as a social or ecological niche, are

dispersed but coordinated and supported by information and communication technologies. Attention allocation control strategies can alter task execution in a data-driven manner. The approach enhances the creation and use of knowledge in action, and is particularly suited for networks subject to stress, overload, or environmental shocks. Control strategies enable the bottom-up integration and interaction of local and global knowledge, while improving performance.

In mission-critical networks, the dynamic capability to adjust attention and task execution is a key factor for achievement and operational resilience. The collaborative feature of distributed cognition allows attention to deviate from strict routing control to focus on anomalous system properties. Communication bottlenecks can also be eliminated. By exploiting the data generated by the network components, distributed cognition allows local and short-term operational priorities to interact with global decision-making processes. Therefore, the interaction of operation and maintenance functions, in line with the principles of cognitive maintenance, can be suitably supported by distributed cognition theory.




Fig 3: Distributed Cognition in Industrial Environments

3.2. Maintenance Frameworks for Mission-Critical Networks



Mission-critical networks such as power generation and distribution systems, critical infrastructure monitoring systems, and large physical and chemical plants require maintenance to avoid unplanned downtimes, hazardous situations, and damage to the environment. Maintenance in these networks is often considered from the perspective of either predictive or condition-based approaches. Predictive maintenance seeks to predict incipient faults within a network with enough advance notice that actions can be taken to avoid unscheduled downtime, using diagnostic or prognostic approaches to estimate Remaining Useful Life on the critical components of a system. Condition-based maintenance focuses on scheduled actions using information gathered in real-time by the system-sensing components regarding operational parameters that affect fault occurrence, while performing the various corrective actions only when the conditions in the real environment signal the need, thus reducing costs without introducing risks.

The maintenance of mission-critical networks is indeed a combination of protective, predictive, and condition-based actions governed using risk analysis of the different components and, above all, ensuring the required redundancy in the architecture so that the system can continue to operate even with one or more component failures. Redundancy can be implemented in various ways: the dedicated processing resources for Fault Detection and Diagnosis modules can be made redundant and provide majority voting to enhance the final determination; additional Fault Detection and Diagnosis resources can monitor the Response Time in their own processing and, based on a failure prediction, not use their final conclusion during the control, prepare a replacement or control the network in a brake mode. In terms of human-actor collaboration, humans act as supervisory controllers, and the Critical Infrastructure Monitoring domain plays an important role in defining the critical components.

Table 3: Industrial Domains and DCMF Applications

Industrial Domain	Application of DCMF	Expected Outcome
Power Generation Networks	Distributed fault recovery and predictive diagnostics	Increased grid reliability
Manufacturing Plants	MES/PAP integration and maintenance orchestration	Improved operational efficiency
Critical Infrastructure Monitoring	IoT-based monitoring and automated alerts	Enhanced infrastructure safety
Transportation Systems	Digital twin simulation and predictive monitoring	Reduced operational disruptions
Smart Factories	AI-driven adaptive maintenance	Autonomous industrial optimization


3.3. Cognitive Maintenance Principles

The implementation of Distributed Cognition principles in industrial environments allows collective operations of intelligence networks and serves as the foundation for Cognitive Maintenance, covering the entire decision-making lifecycle of predictive, preventive, real-time, and proactive activities in maintaining system health and resilience.

The pivotal role of Maintenance and Resilience in mission-critical systems across diverse industries is widely acknowledged. Mission-critical systems must deliver a service fulfilling system-level SLAs. In this context, the cognitive maintenance paradigm collectively oversees predictive and proactive maintenance activities as well as real-time fault detection and diagnosis, resource-aware scheduling and allocation, and collaborative human-



supervisory control – thus addressing the full decision-making life-cycle for maintenance, recovery, and resilience. The concept of cognitive maintenance focuses on mission-critical functions of maintenance – those in which ultimate SLAs drive the need for high speed, reliability, and resilience, thus inspiring the outsourcing of efficiency to resource-aware decision-support and automation (see Section 3.1).



4. Architecture of the Distributed Cognitive Maintenance Framework

Residential and commercial AFDs as protective devices for power distribution systems have gained considerable momentum in recent years aiming to reduce the number of electric fires due to failure in home appliances. Simulations of the complete electrical distribution circuit including appliances as well as the functioning of AFDs involves complex models that consume a huge amount of computational time. AFDs have not yet found their way into a practical built environment in real installations. Adaptive virtual variations of AFDs which naturally assist internal fire simulation in unoccupied space have emerged as a viable alternative for residential, commercial, or anywhere AFDs are required but have not yet found Memorial installations. Complete natural fire load modelling along with normal building simulation will lessen the simulation time and effort.

Digital twin technology has evolved with developing infrastructure & configuration in parallel mode till final

deployment of physical assets. Digital twins of the deployed assets help mission assurance by simulating critical scenarios. Smart Digital Twin setup combines simulation capabilities and functions of Physical Twin and assists in providing data of trained algorithms required for vehicle system Health Management. Cyber-managed digital duality created for Mission Assurance have equivalent roles both for the Human Factors and for the operations of their commanding Vehicle System. The concept of smart dynamic virtual twins forms the basis of advanced intelligent supervision and virtual duals of virtual command missions. Such dynamic emulators created using the concept of dynamic virtual motion fail to find acceptance in the present built environment.

Smart dynamic virtual twins created in dynamic virtual environment along with their physical assets form the smart set-up, both twins integrated into one network. The advancement in science and technology including vehicle system engineering gives major emphasis to applications of spatial and virtual systems along with their consolidation with cyber world. Moving into the realm of Intelligent cyber control offers playway style of operation and is enhancing all critical applications. Natural Fire Detection in enclosed habitation is fast emerging as a new area of research.




Fig 4: Distributed cognitive architecture



4.1. System Layering and Roles

The proposed architecture is represented as a distributed set of systems vertically layered into five operating layers, similar to the Open Systems Interconnect model. Each layer encapsulates and abstracts neighboring layers, establishing zones of control, coordination, and influence for the involved artificial agents. Each layer also performs a specific type of cognitive activity. Layers I to IV represent a standard decomposition of supervisory control systems for monitoring mission-critical processes and infrastructure. Layer V implements a formalized cognitive process that coordinates high-level cognition throughout the entire architecture. Layers I to IV classify as classical supervisory control systems for mission-critical industrial applications, focusing on at least one specific source. Redundant copies of these layers may exist, whether decentralized or fully replicated, to achieve fault tolerance.

The purpose of layer V is to embed oversight cognition throughout the architecture, supporting all vertical information paths with appropriate data indications, in addition to overseeing the lower layers' demands on the monitoring and retrieval components of layers I, II, and III. A dedicated human supervisor is assumed available for maintaining higher-level oversight and controlling physical means to intervene in cases of automation failure, communications breakdown, or unforeseen situations not covered by the automated processes. Layers III and IV are intimately coupled as a cognition-and-action system whose activity controls the deployment of resources for layer I, integrating within the layer's supervision to meet the resources management and optimization goals.

Equation 4: Fault Detection Probability

Used in real-time diagnosis systems:

$$P(F | S) = \frac{P(S | F) \cdot P(F)}{P(S)}$$

Where:

- $P(F | S)$ = Probability of fault given sensor observation
- $P(S | F)$ = Probability of sensor behavior under fault
- $P(F)$ = Prior probability of fault
- $P(S)$ = Total sensor observation probability

This Bayesian model aligns with probabilistic fault diagnosis mechanisms mentioned in the framework.

$$P(F | S) = \frac{P(S | F) \cdot P(F)}{P(S)}$$

4.2. Communication and Coordination Mechanisms

Various typical communication and coordination mechanisms span the different layers of the Distributed Cognitive Maintenance Framework. Interface nodes provide an interface to external systems, such as supervisory control. As these systems have inherent limitations concerning the level of detail they can capture, the information provided for external dimensions of agency must inevitably be abstracted, framed, and communicated. A two-way information flow is necessary: external systems need to be informed of important events during execution, while changes in context must also be communicated to the system.

Lower layers disseminate information when action is required or when an internal understanding of context changes. Deliberative modules maintain a knowledge base that can be openly queried by any module requiring context information. Synchronisation mechanisms maintain mutual knowledge of system state and, when necessary, require joint actions to maintain system integrity. High-level, less frequent data from predictive and preventative maintenance can be filtered, summarised, or abstracted into a form usable by lower layer processes. When constituent modules need to consider high-frequency data originating from disparate sources, suitable communication mechanisms should ensure that these channels remain manageable.

Table 4: Five-Layer Distributed Cognitive Architecture

Layer	Function	Role in DCMF
Layer I	Physical Monitoring Layer	Collects sensor and operational data
Layer II	Supervisory Control Layer	Coordinates device-level operations
Layer III	Resource Management Layer	Allocates resources and schedules tasks
Layer IV	Coordination & Communication Layer	Synchronizes distributed agents
Layer V	Cognitive Decision Layer	Performs high-level reasoning and oversight

4.3. Sensing, Reasoning, and Actuation

Sensor and actuator configurability and resource awareness play key roles throughout DACOMI. Resources are defined in a broad sense, encompassing the storage, communication links, sensing and actuation capabilities, signal-processing and management competencies, and energy supplies needed to execute the functions of DACOMI, as well as the temporal, spatial, or social availability of human-actor resources.

Dedicated modules handle the sensing, reasoning, and actuation of each layer. Sensing modules discover the status of the monitored systems and identify possible emergencies. Reasoning modules aggregate the data of the connected systems and assess resource consumption efficiency while ensuring that real-time requirements are satisfied. Acting modules allocate the available resources and supervise the activation of nondedicated modules within a DACOMI instance: for example, deciding to manually recharge an electric vehicle in the charging-docking area of a set of one or more charging-docking areas.



5. Data Management and Provenance

Data management plays a pivotal role in enabling the decision-making activities of the various maintenance entities participating in the DCMF. Data from heterogeneous sources—sensing, simulation, external, or acting directly upon a system—must be collected in a quality-assured manner and made available to the involved entities. Fusing measurements from several sources, some of which may be unreliable or affected by noise, increases the accuracy of the resulting data products. Traceability is also a major issue, especially in contexts with strict legal compliance requirements, where both the integrity and the origin of data must be guaranteed. In that framework, a powerful provenance-subsystem, ensuring compliance with traceability constraints, is needed. Security and privacy, also of paramount importance in mission-critical systems, must be safeguarded at all levels.

Data collected and created during the normal operation of the mission-critical industrial system, throughout its lifecycle, constitute one of its major assets. Properly managed and treated, this data invisible for most of the time contributes to the design, deployment, maintenance, and upgrading activities, including certification if needed. Consequently, besides covering conventional data-management aspects, proper data-provenance subsystems allow regulators and supervisory actors to ascertain the integrity of the data and its support for the above-mentioned missions.

Equation 5: Resource-Aware Scheduling Cost Function

A generalized optimization equation for maintenance resource allocation:

$$C = \sum_{i=1}^n (T_i + E_i + R_i)$$

Where:

- C = Total maintenance scheduling cost
- T_i = Time cost for task i
- E_i = Energy/resource consumption
- R_i = Risk factor associated with delayed maintenance

5.1. Data Collection, Fusion, and Quality Assurance

The DCMF depends on accurately reasoning about the state of the system and its environment to achieve a suitable maintenance solution. This requires quality-assured data from many heterogeneous sources. Data quality assurance follows a layered approach. First, the internal data collectors enforce basic and format checks and minimize noise. Second, the DCMF's data quality service—a part of the reasoning engine—purges low-confidence or contradictory information. Third, when external data is received from the original data producers or data aggregators, it is evaluated against the declared quality criteria. Failing these criteria, it is filtered out or flagged for manual inspection.

Data fusion combines information from multiple sources, including different types of data repositories. Examples are digital twins used in realistic and critical scenarios, operational data repositories such as MES and logs from IAM systems. Fusion considers multiple aspects, including time, space, and data quality. Correct temporal alignment is crucial to associating the right sources of different types of data.




Fig 5: Data Collection, Fusion, and Quality Assurance

5.2. Provenance, Traceability, and Compliance

Providing origin and handling information for collected data ensures their quality, particularly in regulated sectors, such as energy production and distribution. Provenance allows for tracking quality degradation throughout the data chain with the aim of preventing violations of the integrity, timeliness, and other requirements associated with the data's eventual uses.

Monitoring data provenance improves timeliness, as timely resource provisioning decreases production delays and improves the chance of on-time delivery. Similarly, timeliness is improved when requested data are available and correctly addressed to the requester. Complying with environmental and safety regulations is essential regardless of the sector, as violations incur huge penalties. Data-traceability mechanisms bind data to individual products, providing crucial information on an industrial entity's environmental footprint and the nature of any potential safety hazards.


Compliance is largely a matter of clear data traceability throughout the considered manufacturing value chain; the ability to establish whether some data are fulfilled during production, either explicitly or through depletion, is essential for the final products.

Table 5: Communication and Coordination Mechanisms



Mechanism	Purpose	Example
Interface Nodes	Connect external supervisory systems	SCADA integration
Knowledge Base Sharing	Shared contextual intelligence	Distributed reasoning
Synchronization Protocols	Maintain system integrity	State consistency checks
Event Dissemination	Notify system changes	Fault alerts
Data Abstraction	Simplify high-frequency data	KPI summarization

group and lead to a tailor-made solution. Security- and privacy-related LBS acceptance testing together with threat modelling can unveil vulnerable aspects. Risk assessments can distinguish the crucial items before detailed considerations about the others commence.



5.3. Security and Privacy Considerations

Securities regarding the communication and storage of sensitive data should be stated. Confidentiality should be ensured by integrity checking with hashing, validation against tampering and encryption. Authentication helps guarantee provenance and trust in the DSS's model as well as that of the individual actors, while key management preserves the secrecy of the keys used for encryption and signature generation. Anonymization can protect user identities; even so, the ability to identify data donors becomes a prerequisite for compliance with legal responsibilities. The treatment of personally identifiable information can be assisted by the ready-made service provided by the PAN European eID programme.

Provisions to incorporate all security and privacy points from the Logic-based Specification (LBS) through the software development life cycle (SDLC) is strongly suggested. Embedded requirements can further support practical decision-making about a system's lifecycle. Detailed visualizations can raise Snyder and Wulf's communicational gestalt and succinctly emphasize the importance of considering security and privacy from the start. User-centric design principles (UCDP) communicate with the target user

6. Decision-Making Processes for Maintenance

Maintenance actions are typically driven by decision-making processes that include predictive maintenance, proactive maintenance, real-time fault detection and diagnosis, resource-aware scheduling and allocation, and human-actor collaboration and supervisory control.

Predictive maintenance employs mathematical models to derive estimates of system behaviour. These estimates can be further analysed using the models' interpretable parameters to assess the likelihood of malfunctions or failures during a specified forecasting horizon. Proactive maintenance enhances predictive maintenance by augmenting the pure model-driven approach with sensor data that increases predictive performance and reduces uncertainty. Data-driven models quantify the current status of the examined system through sensor data and issue early warnings according to the assigned Service Level Agreements (SLAs). These models do not account for unknown malfunctions and rely on substantial historical data for proper functioning.

Real-time fault detection and diagnosis detect and diagnose faults when they occur to trigger immediate corrective



actions, typically by the maintenance crew. These models mainly rely on sensor data; knowledge for fault detection can be automatically learnt from historical sensor data but can also be represented by a set of logical rules. Resource-aware scheduling and allocation combine allocation of limited resources with maintenance activities. Human-actor collaboration and supervisory control provide human operators or other higher-level controllers with information about the system's status, maintenance actions, detected faults, and decisions taken by automated processes.

Table 6: Sensing, Reasoning, and Actuation Functions


Function	Description	Operational Role
Sensing	Collects operational data	Detects abnormal conditions
Reasoning	Processes and analyzes data	Decision support
Actuation	Executes maintenance actions	Resource deployment
Resource Awareness	Evaluates available resources	Scheduling optimization
Human Interaction	Supports supervisory control	Manual override and approval

6.1. Predictive and Proactive Maintenance

Users analytically predict the remaining useful life of different assets and recommend preventive actions for time-based maintenance. Predictive analytics examines data collected by machine sensors to forecast failure events, enabling near-time maintenance scheduling. Sufficient advance notice allows the operator to schedule time-based maintenance before asset failure. An example of predictive maintenance is the use of condition monitoring of lift motors and predictive analytics based on temperature and other sensor data. Proactive maintenance goes beyond predictive

maintenance to address earlier-stage problems caused by deteriorating operations or environments, replacing road surface materials well before they fail, for instance.

In the process of intelligent piping-segment-condition management, a model learns from historic data to predict the most likely segment conditions. The neural network models segment states within a time range, with each state recognised by a status name such as good, very good, and so on. Piping in the worst condition can then be restored via camera inspections and rectifying actions, reducing traffic disruption and costs and increasing transport safety. A monitoring system for fluorescent-lamp burnouts incorporates machine-learning technology that factors in seasonality, non-operating days, and clustering. It can be further advanced for intelligent preventive maintenance by predicting other failure types, quantifying the effects of preventive maintenance, modelling users' operational pattern, and optimally allocating maintenance resources.



6.2. Real-Time Fault Detection and Diagnosis

Faults in mission-critical industrial networks should be detected, diagnosed, and rectified at runtime, minimizing any disruption to the network's normal operation. Unlike traditional centralized approaches, the cognitive maintenance framework diffuses these tasks across the network's maintenance agents, coordinating their operation through the network's multi-agent capabilities.



The detection of faults is achieved using a collection of detection agents that monitor a specific network component or its environment in an akin fashion. Each such monitoring agent is responsible for identifying a subset of faults/abnormal conditions based on information received from the component being monitored. The fault-detection technique and the monitored entity are application dependent, and the monitoring agents exploit network-local information. Local measurement-based techniques with well-defined limits require little formal knowledge representation and reasoning. More advanced techniques using approaches such as machine learning, case-based reasoning, or remote-sensing techniques are also supported by the framework. The output of every fault-detection agent is a probabilistic distribution over the possible faults/abnormal conditions affecting their monitored entity, based on the available information.

The diagnosis of faults is performed by a collection of diagnosis agents that maintain an activity model of the network and track its runtime evolution. These diagnosis agents reside at the boundaries shared with other networks, where deviations in the expected activity may indicate faults that must be rectified. Each such agent uses the discrete-event representation of the activity model to update the state of all component activity indicators when the relevant input is received. Where necessary, the updated indicators are then interrogated to determine the set of detected faults, displayed as a similar probabilistic distribution.





Fig 6: Real-Time Fault Detection and Diagnosis

Resource-aware scheduling and allocation are key decision-making functions for DCM-enabled mission-critical industrial networks. These functions leverage human and physical resources with required capabilities and availability to fulfil scheduled maintenance tasks on time. Failure to perform tasks on time may lead to system-level failures and therefore additional measures that affect the mission-critical operation of the network.

Resources include people, machinery, tools and equipment, transporters, spare parts, and consumables. A task requires the capabilities and availability of its resources to be at least at the min-allocation level. The allocation decision takes the current and future capability levels of resources into account. Resource-aware scheduling provides a candidate schedule and a requirement profile for each task, based on a combination of decision functions. It must determine the availability of the required capabilities in a way that the fulfilment of the resulting assignments is guaranteed, based on normal monitoring and alert-level supervision.

The allocation model can be flat or hierarchical, but sufficient local availability of capabilities must be ensured in case of highly distributed operations. The assignment of candidates with long travel times to successive, close-distance tasks is avoided and opportunities for additional fellowships are pointed out. Allocated resources can be fully engaged, partially engaged for performing policy-related requests, or engaged but not able to fulfil assignments. The resource allocation algorithm manages the availability of human and physical resources in response to changes in situation awareness and checks whether the specified capability levels are met.

6.3. Resource-Aware Scheduling and Allocation




Resilience and fault-tolerance considerations flow into most layers of the cognitive maintenance framework. A redundancy strategy protects digitally generated information. Interaction with the simulated environment confers risk-limiting properties to human-machine collaboration and supervisory control. Furthermore, carefully delineated human-machine agreements prepare the ground for a security-informed, resource-conscience, and decision-resilient Hutchins loop.

Digital-twin integration confers safety assurance and certification advantages to the DSSDD system. Resilience and FT analysis leverage the identified fault modes and redundancy strategies of physical deployment, supporting identification of exploitable failure modes and generation of robust resource-aware SSE routes. These analyses furthermore inform selection of an appropriate recovery-time objective from the perspective of the associated critical service. Finally, cognitive maintenance interaction with a complete replica of the resource entails that safety experts can focus on automated fault-removal procedures and their assurance, reducing their workload and facilitating their mastery of the relevant materials.

DCMF Architecture

7. Resilience, Fault Tolerance, and Safety Engineering

Resilience is the ability of a system to absorb disturbances, adapt to changing conditions, and retain essential functionality. Fault tolerance denotes the internal properties enabling resilience. Formal redundancy affords protection against specific failure events. Safety combines design practices that minimize hazardous conditions with supporting analyses that identify and evaluate remaining risks.



Maintenance Decision Flow

7.1. Redundancy Strategies and Failure Modes

Redundant systems and components can prevent failure when some elements fail, and their application is a widely accepted principle in safety engineering. Essential industrial operations often require redundancy, ranging from simple, cheap, and specific equipment to elaborate architectures relying on costly components and elaborate protocols. Computer systems controlling critical processes and operations are optimized for their availability and uptime and include different redundancy strategies. Nonetheless, ensuring redundancy remains a matter of resource-awareness.

A first-level check of a distributed cognitive maintenance framework and its implementation needs to reveal critical

components and functions. Then, a failure mode effect analysis can determine for each item whether re-execution can recover the primary task—namely, data collection, fusion, reasoning, actuation, and the supervisory control that allows delegating perceptual and executive responsibilities to intelligent agents in the infrastructure and evolving process areas. Cyber security is of critical concern because any compromise can easily affect the availability and reliability of the entire system and, hence, that of the industrial operation being monitored or controlled.

7.2. Safety Assurance and Certification Aspects

Safety assurance for mission-critical environments is crucial at both individual and network levels. Failure of any element within the network can jeopardize the entire infrastructure. Consequently, appropriate safety measures need to be adopted for each element of the distributed cognitive maintenance system. Traditional certification methodologies based on fault tree analysis are limited due to the set-theoretic approach employed. Therefore, PGS and related techniques provide a more suitable approach because they offer a model-based strategy for analysis, validation, and certification of safety properties—especially where redundancy is employed. Furthermore, for certified safety, development processes must obey strict rules, as required by IEC 61508, for programmable electronic systems. Consequently, safety certification becomes excessively costly when the system must be resiled after each change due to safety-related updates.

Within the highway ecosystem, Adaptive Cruise Control (ACC) systems—mainly based on control applications—can improve safety and user comfort by safely following the vehicle ahead—thereby reducing accidents related to speed discrepancies. However, the validation of ACC requires a significant computational effort that can be reduced using formal verification methods. Nevertheless, depends on reduction strategies relying on falsification as a complement. In the railway domain, model-based safety evaluation of a Traffic Management System on a simplified communication data-flow model is presented to support safety change management. Safety assurance is performed using structured



argument with autoregressive-exponential moving average models for Black-Zhukovsky formula prediction.

7.3. Recovery and Recovery-Time Objectives

As a continuation of the preceding section, consideration of recovery addresses the ability to return to a correctly functioning state after an abnormal incident while still supplying a service or fulfilling a mission. Formal definitions for recovering from a fault are in the context of a system model that treats the faults and the means of recovery as given. Once again, timing may appear in the definitions of the fault-tolerance strategy but not in the definitions of the recovery-time objectives, which are stated and considered separately. Another area in which the contract between a system vendor and a consumer is sometimes made explicit is in the definition of recovery-time objectives (RTOs), which set an upper time limit on the restoration of normal systemic operations after an abnormal incident.

The RTO is not intended to apply to the control of the telecommunications networks hosted, or in part hosted, by the mission-critical network, but simply to their normal operation. The distinction is blurred, however, as some of the components of the telecommunications networks may perform a sensing function, generating and relaying alarms for the early detection of developing conditions that may threaten the normal operation of the mission-critical network.

8. Conclusion

The exploration of a collaboratively Cognition Engine, a Digital Twin of the Field, and the consequent Maintenance Architecture has revealed an architecture that consolidates the characteristics of such industrial and mission-critical environments. The minimum characteristics of the proposed Distributed Cognition Engine rely on a cognitive ability that serves all cognition requests from the interconnected sub-networks, allowing data synthesis, communication, coordination, and remotely executing other entities' decision plans.

The Distributed Cognitive Maintenance and Supervisory Control for Mission-Critical Industrial Networks encompasses the principal Cognition Engine, termed the Cognitive Maintenance and Supervisory Control Engine, and the Digital Twins Architecture of each monitored industrial process, allowing for a distributed team of business players synthesising a broader cognitive ability. By acting almost autonomously for daily operations, such digital twins enable specialised operators and supervision, leaving normal functioning without human presence. Such functionality allows continuous data observation, providing the required input for the Distributed Cognitive Maintenance by detecting and warning anomalous behaviours.

References

1. Alassery, F. (2022). Predictive maintenance for cyber physical systems using neural network based on deep soft sensor and industrial internet of things. *Computers and Electrical Engineering*, 101, 108062.
2. Alotaibi, B. (2023). A survey on industrial internet of things security: Requirements, attacks, AI-based solutions, and edge computing opportunities. *Sensors*, 23(17), 7470.
3. Kolla, S. H. (2022). Strategic Information Integration Models for Cross-Functional Service Optimization in Large-Scale Enterprises. *International Journal of Emerging Trends in Engineering and Management Research*, 7(3), 11811.



4. Ayvaz, S., & Alpay, K. (2021). Predictive maintenance system for production lines in manufacturing: A machine learning approach using IoT data in real-time. *Expert Systems with Applications*, 173, 114598.
5. Bampoula, X., Siaterlis, G., Nikolakis, N., & Alexopoulos, K. (2021). A deep learning model for predictive maintenance in cyber-physical production systems using LSTM autoencoders. *Sensors*, 21(3), 972.
6. Mangalampalli, B. M. Generative AI Applications In Healthcare Data Mart Design And Optimization.
7. Behera, S., & Misra, R. (2021). Generative adversarial networks based remaining useful life estimation for IIoT. *Computers & Electrical Engineering*, 92, 107195.
8. Chen, C., Fu, H., Zheng, Y., Tao, F., & Liu, Y. (2023). The advance of digital twin for predictive maintenance: The role and function of machine learning. *Journal of Manufacturing Systems*, 71, 581–594.
9. Inala, R. (2023). AI-powered investment decision support systems: Building smart data products with embedded governance controls. *Journal for ReAttach Therapy and Developmental Diversities*, 6(10), 2251-2266.
10. Elattar, H., El-Brawany, M., Elminir, H. K., Ibrahim, D. A., & Ramadan, E. A. (2023). Artificial intelligence-based data-driven prognostics in industry: A survey. *Computers & Industrial Engineering*, 184, 109605.
11. Elgendy, I. A., Muthanna, A., Hammoudeh, M., Shaiba, H., Unal, D., & Khayyat, M. (2021). Advanced deep learning for resource allocation and security aware data offloading in industrial mobile edge computing. *Big Data*, 9(4), 265–278.
12. Ferreira, C., & Gonçalves, G. (2022). Remaining useful life prediction and challenges: A literature review on the use of machine learning methods. *Journal of Manufacturing Systems*, 63, 550–562.
13. Amistapuram, K. (2023). Privacy-Preserving Machine Learning Models for Sensitive Customer Data in Insurance Systems. *Educational Administration: Theory and Practice*, 29(4), 5950-5958.
14. Feng, X., Wu, J., Wu, Y., Li, J., & Yang, W. (2023). Blockchain and digital twin empowered trustworthy self-healing for edge-AI enabled industrial Internet of things. *Information Sciences*, 642, 119169.
15. Fordal, J. M., Schjølberg, P., Helgetun, H., Skjermo, T. Ø., Wang, Y., & Wang, C. (2023). Application of sensor data based predictive maintenance and artificial neural networks to enable Industry 4.0.



- Advances in Manufacturing, 11(2), 248–263.
16. a, S. K., & Reddy, V. A. R. (2023). Deep Learning Architectures For Multimodal Medical Data
 17. Fu, H., & Liu, Y. (2022). A deep learning-based approach for electrical equipment remaining useful life prediction. *Autonomous Intelligent Systems*, 2, Article 16.
 18. Gupta, V., Mitra, R., Koenig, F., Kumar, M., & Tiwari, M. K. (2023). Predictive maintenance of baggage handling conveyors using IoT. *Computers & Industrial Engineering*, 177, 109033.
 19. Hung, Y.-H. (2021). Improved ensemble-learning algorithm for predictive maintenance in the manufacturing process. *Applied Sciences*, 11(15), 6832.
 20. Inala, R. (2023). Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective. *International Journal of Finance (IJFIN)-ABDC Journal Quality List*, 36(6), 579-606.
 21. Jiang, Y., Dai, P., Peng, F., Zhong, R. Y., et al. (2022). A²-LSTM for predictive maintenance of industrial equipment based on machine learning. *Computers & Industrial Engineering*, 172, 108560.
 22. Kang, Z., Catal, C., & Tekinerdogan, B. (2021). Remaining useful life (RUL) prediction of equipment in production lines using artificial neural networks. *Sensors*, 21(3), 932.
 23. Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. Zenodo.
 24. Liu, Z., Liu, H., Jia, W., Zhang, D., & Tan, J. (2021). A multi-head neural network with unsymmetrical constraints for remaining useful life prediction. *Advanced Engineering Informatics*, 50, 101396.
 25. Mo, H., Custode, L. L., & Iacca, G. (2021). Evolutionary neural architecture search for remaining useful life prediction. *Applied Soft Computing*, 108, 107474.
 26. Inala, R. Designing Scalable Technology Architectures for Customer Data in Group Insurance and Investment Platforms.
 27. Mohan, T. R., Roselyn, J. P., Uthra, R. A., Devaraj, D., & Umachandran, K. (2021). Intelligent machine learning based total productive maintenance approach for achieving zero downtime in industrial machinery. *Computers & Industrial Engineering*, 157, 107267.
 28. Mourtzis, D., Angelopoulos, J., & Panopoulos, N. (2022). Design and development of an edge-computing platform towards 5G technology adoption for improving equipment predictive



- maintenance. *Procedia Computer Science*, 200, 611–619.
29. Nikfar, M., Bitencourt, J., & Mykoniatis, K. (2022). A two-phase machine learning approach for predictive maintenance of low voltage industrial motors. *Procedia Computer Science*, 200, 111–120.
 30. Resende, C., Folgado, D., Oliveira, J., Franco, B., Moreira, W., Oliveira-Jr, A., Cavaleiro, A., & Carvalho, R. (2021). TIP4.0: Industrial Internet of Things platform for predictive maintenance. *Sensors*, 21(14), 4676.
 31. Sharma, D. K., Brahmachari, S., Singhal, K., & Gupta, D. (2022). Data driven predictive maintenance applications for industrial systems with temporal convolutional networks. *Computers & Industrial Engineering*, 169, 108213.
 32. KollIntegration. *South Eastern European Journal of Public Health*, 248–260.
 33. Shaheen, B., Kocsis, Á., & Németh, I. (2023). Data-driven failure prediction and RUL estimation of mechanical components using accumulative artificial neural networks. *Engineering Applications of Artificial Intelligence*, 119, 105749.
 34. Steurtewagen, B., & Van den Poel, D. (2021). Adding interpretability to predictive maintenance by machine learning on sensor data. *Computers & Chemical Engineering*, 152, 107381.
 35. Taşcı, B., Omar, A., & Ayvaz, S. (2023). Remaining useful lifetime prediction for predictive maintenance in manufacturing. *Computers & Industrial Engineering*, 184, 109566.
 36. Uhlmann, E., Polte, J., & Koutrakis, N.-S. (2021). Holistic concept towards a reference architecture model for predictive maintenance. *Procedia CIRP*, 104, 1430–1433.
 37. Vermesan, O., Coppola, M., Bahr, R., Bellmann, R. O., Martinsen, J. E., Kristoffersen, A., Hjertaker, T., Breiland, J., Andersen, K., Sand, H. E., & Lindberg, D. (2022). An intelligent real-time edge processing maintenance system for industrial manufacturing, control, and diagnostic. *Frontiers in Chemical Engineering*, 4, 900096.
 38. Yang, L., Zou, W., Wang, J., & Tang, Z. (2022). EdgeShare: A blockchain-based edge data-sharing framework for Industrial Internet of Things. *Neurocomputing*, 485, 219–232.
 39. Yu, W., Liu, Y., Dillon, T. S., & Rahayu, W. (2023). Edge computing-assisted IoT framework with an autoencoder for fault detection in manufacturing predictive maintenance. *IEEE Transactions on Industrial Informatics*, 19(4), 5701–5710.