Diana Ibraheem

Feasibility Study of the Remote Operation Centre (ROC) for Autonomous Power Plants

School of Technology and Innovations
Master's thesis
Sustainability and Autonomous system

Vaasa 2025

UNIVERSITY OF VAASA
School of Technology and Innovations
Author: Diana Ibraheem
Title of the thesis: Feasibility Study of the Remote Operation Centre (ROC) for Autonomous Power Plants
Degree: Master of Computing Science
Degree Programme: Sustainability and Autonomous System
Supervisor: Petri Välisuo, Amir-Mohammad Shamekhi, Mahmoud Elsanhoury

Year: 2025 **Pages:** 112

ABSTRACT:

Autonomous operation of distributed power assets requires a remote operation centre (ROC) that integrates sensing, control and assurance with auditable execution. This thesis defines and implements a ROC reference architecture for autonomous power plants, emphasising deterministic supervision, explainable autonomy and operator-in-the-loop authority. The design draws on established standards and technologies: measurement and calibration workflows, time synchronisation (PTP/NTP), interoperable interfaces, and the data plane that supports multi-rate fusion, digital-twin synchronisation and embedded-grade cybersecurity. The control stack combines model-predictive control and reinforcement learning, bounded by safety envelopes, action-admissibility checks and traceable logging to preserve transparency and rollback capability. The architecture is evaluated in the EPS/VEBIC laboratory against timing and data-integrity constraints, demonstrating end-to-end timing fidelity, fault-tolerant acquisition and the feasibility of deterministic supervisory control as a base for higher autonomy. The main contribution is an implementable ROC blueprint that links measurement integrity, secure communications and hierarchical control to an auditable route from remote supervision to safe autonomy. The work is limited to architectural design and laboratory-based validation; a formal safety case and comprehensive verification and validation of AI components remain future work, alongside scaling to multi-asset supervision.

KEYWORDS: power plants; autonomous systems; artificial intelligence; control systems; reinforcement learning; control rooms; synchronising; cyber security.

Foreword

This Master`s thesis was carried out at the University of Vaasa as part of the research on autonomous systems. The work contributes to the study and design of a Remote Operation Centre (ROC) for autonomous power plants, conducted within the framework of the Efficient Powertrain Solutions (EPS) laboratory and the Vaasa Energy Business Innovation Centre (VEBIC).

Firstly, I would like to thank my instructor, Dr. Amir-Mohammad Shamekhi, for his guidance and committed support throughout this work. I am also grateful to my supervisor, Prof. Petri Välisuo, for his academic supervision and constructive comments. I extend my thanks and appreciation to my second supervisor, Dr. Mahmoud Elsanhoury. I would also like to thank the EPS and laboratory staff for their assistance during the research phase.

Finally, I wish to thank my family for their continuous support throughout my studies.

Contents

Figures

Tables

Abbreviations

AAS	Asset Administration Shell
ABB	ASEA Brown Boveri
ADC	Analog-to-digital converter
BESS	Battery energy storage system
CAD	Crank-angle degree
CA50	Crank angle at 50% heat release
CAN	Controller Area Network
CLD	Chemiluminescence detector
CO	Carbon monoxide
DAQ	Data acquisition
DAS	Distributed acoustic sensing
DED	Directed energy deposition
DER	Distributed energy resources
DMPC	Distributed model predictive control
DRL	Deep reinforcement learning
DT	Digital twin
DTS	Distributed temperature sensing
EKF	Extended Kalman filter
EMS	Energy management system
EPS	Efficient Powertrain Solutions
FFT	Fast Fourier transform
FTIR	Fourier transform infrared
FTC	Fault-tolerant control
GOOSE	Generic Object Oriented Substation Event
HIL	Hardware-in-the-loop
HMI	Human–machine interface
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IMEP	Indicated mean effective pressure
IRKA	Iterative Rational Krylov Algorithm
LSTM	Long short-term memory
MHE	Moving horizon estimation
MIL	Model-in-the-loop
MOR	Model order reduction
MPC	Model predictive control
MQTT	Message Queuing Telemetry Transport

NIST	National Institute of Standards and Technology
NOX	Nitrogen oxides
NTP	Network Time Protocol
OOTL	Out-of-the-loop
OPC UA	Open Platform Communications Unified Architecture
PHM	Prognostics and health management
PMU	Phasor measurement unit
POD	Proper orthogonal decomposition
PTP	Precision Time Protocol
RIAPS	Resilient Information Architecture Platform for Smart Grid
RL	Reinforcement learning
ROC	Remote Operation Centre
ROM	Reduced-order model
ROHR	Rate of heat release
RUL	Remaining useful life
SCADA	Supervisory Control and Data Acquisition
SV	Sampled Values
THC	Total hydrocarbons
TIDE	Time-series Dense Encoder
UAV	Unmanned aerial vehicle
UKF	Unscented Kalman filter
VEBIC	Vaasa Energy Business Innovation Center
VPP	Virtual power plant
XDT	Executable digital twin

1 Introduction

1.1 Background

The global energy sector is currently undergoing a major transformation, driven by three interrelated developments: decarbonisation, digitalisation, and increasing autonomy. Decarbonisation introduces requirements for efficiency, operational flexibility, and integration of variable renewable energy sources. Digitalisation provides new capabilities in data acquisition, machine learning, and predictive analytics. Together, these development trends are directing the energy sector towards more autonomous operation, where supervisory intelligence gradually replaces traditional operator-driven coordination. This transition the Remote Operation Centre (ROC) has become a sort of key enabler, providing the needed infrastructure for centralized monitoring, coordination and optimisation of power plants that are often scattered across wide areas. Industrial examples already show how the ROC concept is actually used in practice. Wärtsilä's WISE programme applies these same principles to power plant management, focusing mainly on predictive optimisation, lower emissions and better operator support especially in remote or hard to access environments. In a similar way ABB's Ability™ platform goes in this direction too, by bringing automation systems together into central control hubs designed for both microgrids and larger industrial setups.

In the maritime sector remote operation centres are also becoming quite common, used for supervising and controlling autonomous vessels. As reported by DNV (2022), shore-based control centres are now used for navigation, mission planning, and emergency response of vessels operating with reduced or no onboard crew. These examples indicate that the ROC should not be viewed solely as a monitoring layer. Instead, it acts as an essential component of the overall control architecture, having responsibilities related to coordination, safety supervision, and high-level decision support.

Despite the rapid industrial deployment, academic research has not yet produced a unified theoretical framework for ROCs in autonomous power plant applications. Existing supervisory systems are still limited by deterministic yet rigid optimisation methods, low adaptability to uncertainty, and fragmented data exchange. Advances in fields such as

Artificial Intelligence, Modelling, and Control provide new opportunities. These include reinforcement learning for adaptability, digital twins for predictive modelling, and advanced model predictive control for constraint-aware optimisation. However, these methods are often examined separately and are rarely integrated into a comprehensive supervisory structure.
Meanwhile the growing autonomy of industrial systems brings along several challenges in how humans and automation actually work together, like maintaining operator trust, keeping situational awareness, and avoiding the typical out-of-the-loop performance problem (Mooneyham & Schooler, 2013). The motivation for this thesis sort of lies at the crossing point of these industrial and academic developments. While the industry already shows that ROCs can work as real enablers of autonomy, there is still a missing systematic, research-based design framework that limits their broader use in autonomous power plants. This work tries to fill that gap by developing an AI-enhanced ROC architecture, which combines deterministic optimisation with adaptive learning, predictive digital modelling and a structured data management concept into one coherent supervisory setup.

1.2 Remote Operation Centres (ROC)

Remote Operation Centres (ROCs) are fast becoming central parts in the digitalisation of critical infrastructures. In the energy field, a ROC basically works as a supervisory hub that gathers real-time data from distributed assets, runs optimisation and predictive algorithms, and gives decision-support interfaces for operators. Unlike local controllers that operate in millisecond timescales to keep deterministic feasibility, the ROC takes care of the long-term goals such as efficiency, emission reduction and asset health management across several plants at once.

Industrial examples already show why this idea is important. Wärtsilä’s WISE platform brings ROCs into autonomous power plant operation by offering predictive maintenance, fleet-level optimisation and better operator support (Wärtsilä, 2024). ABB’s

Ability™ system follows a similar path in microgrids and industrial automation, providing remote monitoring and event detection and centralised optimisation (ABB, 2024). While in the maritime sector, DNV (2022) describes how shore-based ROC are used to supervise remote and autonomous vessels, coordinating navigation, mission planning and also emergency handling when is needed, similar trends can be seen in the oil and gas industry, remote operation is already seen as a step toward fully autonomous sites. Devold and Fjellheim (2019) point out that these setups are enabled by Internet of Things technologies used for sensing and control, connectivity for information exchange, and AI-based analytics. Jointly these cases make it clear that ROCs have moved past the trial phase and are now quite essential in safety-critical work.

Nonetheless, despite the steady industrial progress, today`s ROC systems remain somewhat fragmented. Synchronising distributed telemetry is tricky because of different standards and protocols (like IEEE 1588 PTP, IEC 61850, OPC UA, MQTT) which often makes integration more difficult than it should be, in many current setups, supervisory decision-making still relies on fixed optimisation logic. These methods are dependable but miss the flexibility of reinforcement learning or the predictive value that comes with digital twins. More automation inside supervisory systems also brings up its own set of problems with human-automation interaction. Issues like reduced situational awareness, uneven operator workload and the well-known out-of-the-loop (OOTL) effect (Mooneyham & Schooler, 2013) tend to appear. The OOTL effect refers to how operator engagement and awareness drop when most control authority is handed over to automation, when human operator mainly monitors instead of directly controlling, the ability to notice anomalies or react quickly to abnormal events decreases, which can weaken the reliability and safety of remote supervision in critical systems (Merat et al., 2019).

In this thesis the ROC is observed as a centralised supervisory control hub for autonomous power plants. The concept includes four main functions: (i) real-time data collection and fusion, (ii) supervisory intelligence through hierarchical MPC, reinforcement learning and digital twin modelling, (iii) runtime assurance and validation of commands for safety, and (iv) operator decision-support for trust and transparency. With this, ROC is not just a monitoring tool but an AI-supported supervisory architecture that follows industrial practice while also tackling the open academic challenges still ahead.

1.3 Research Problem and Questions

The global energy sector is advancing towards autonomous and data-driven operational concepts, where supervisory intelligence and predictive analytics are increasingly integrated into plant-level decision-making (ABB, 2024; Wärtsilä, 2024). The industrial initiatives have been discussed in earlier to demonstrate this development trend. However, despite the significant industrial progress the academic research has not yet formulated a unified framework for ROC design within the context of autonomous power plants.

The supervisory role of a ROC demands the integration of several advanced methodologies within a single coherent architecture, Model Predictive Control (MPC) enables constraint-aware optimisation but remains restricted by its computational demands and modelling accuracy requirements (Rawlings et al., 2020), Reinforcement Learning (RL) provides adaptability under uncertain conditions though it lacks deterministic safety assurance (Yu et al., 2024), Digital Twins (DTs) support predictive and condition-aware modelling; however, their performance depends strongly on accurate synchronisation and model fidelity, especially during transient operation (Hautala et al., 2022). Simultaneously, the interoperability of communication standards for example IEEE 1588 PTP, OPC UA, and MQTT continues to be a challenge in heterogeneous industrial environments (Gupta et al., 2023). Furthermore the growing autonomy of supervisory systems introduces human-automation interaction issues with including operator trust, situational awareness, and the out-of-the-loop (OOTL) problem, which have been observed in both the energy and maritime sectors (Mooneyham & Schooler, 2013; DNV, 2022).

The increasing autonomy of supervisory systems also introduces human-automation concerns, including operator trust, situational awareness, and the out-of-the-loop (OOTL) problem, which have been observed in both energy and maritime domains (Mooneyham & Schooler, 2013; DNV, 2022).

The objective of this thesis is to design and evaluate an AI-enhanced ROC framework for autonomous power plants, the framework is integrating the deterministic optimisation, adaptive learning, predictive modelling and robust data processing pipelines into a coherent supervisory architecture, By consideration is given to computational feasibility and operator interaction to ensure both operational safety and trust. The proposed design is validated the concept through literature and practically against what available infrastructure in the EPS laboratory at the University of Vaasa.

The research aims to answer the following questions:

1. What are the functional and architectural requirements for a Remote Operation Centre to manage an autonomous power plant?
2. How can hierarchical MPC, reinforcement learning and digital twins be integrated into a supervisory control layer that remains both adaptive and safe?
3. What data collection, fusion and communication pipelines are required to enable ROC operation across heterogeneous assets?
4. How can runtime assurance and human-AI collaboration be embedded into the ROC to ensure reliability, transparency, and operator trust?

The scope of this thesis is focused on the supervisory design of a Remote Operation Centre (ROC) intended for autonomous power plant applications. The main emphasis is placed on AI-based decision-making, predictive modelling, and data synchronisation at the supervisory level. Mechanical design aspects of individual plants and the implementation of local controllers are excluded from consideration. The main contribution of this work is the development of an integrated ROC architecture that combines hierarchical model predictive control, reinforcement learning, and digital twins into a coherent supervisory framework, supported by a structured and reliable data pipeline.

1.4 Thesis Structure

The thesis is organised in 6 chapters. After Introduction and problem definition in chapter 1, Chapter 2 is dedicated to *Literature Review* to establish the theoretical foundation for the ROC design. It examines supervisory intelligence methods, including model predictive control, reinforcement learning, and digital twins, as well as data communication and human-automation interaction. Rather than serving only as a survey, this chapter identifies the functional requirements that a ROC must fulfil.

Chapter 3, *EPS Laboratory*, introduces the experimental context of the Efficient Powertrain Solutions facility at the University of Vaasa. The laboratory infrastructure, consisting of real-time controllers, data acquisition equipment, and communication systems, defines the practical boundary conditions against which the ROC design must be developed. Chapters 2 and 3 together establish both the conceptual requirements and the practical limitations that define the basis for the ROC development.

Chapter 4, reflects *ROC Design and Architecture Proposal*, which forms the central part of this thesis. It synthesises to find from the literature review and laboratory investigations within a conceptual framework for an AI-enhanced ROC. The proposed design integrates the hierarchical Model Predictive Control, Reinforcement Learning and Digital Twin technologies within a supervisory control layer that supported by a real-time data pipeline architecture, this chapter demonstrates how theoretical progress and experimental conditions can be led to unified ROC framework works for autonomous power plant operation.

Chapter 5, *Feasibility and Challenges*, assesses the proposed ROC in terms of computational scalability, communication reliability and human-operator interaction, It highlights the main discrepancies between the conceptual design and the practical industrial implementation thereby defining the direction for continued development.

Chapter 6, *Conclusions and Future Work*, presents the summary of the key contributions of the thesis and reflects on its limitations, furthermore, it suggests the possible extensions such as multi-plant coordination and enhanced human-AI interfaces in order to improve supervisory autonomy and decision support.

The methodological structure of the thesis is organised so that the literature review (Chapter 2) and the laboratory framework (Chapter 3) establish the foundation for the ROC concept presented in Chapter 4, which is subsequently evaluated in Chapter 5 and concluded in Chapter 6.

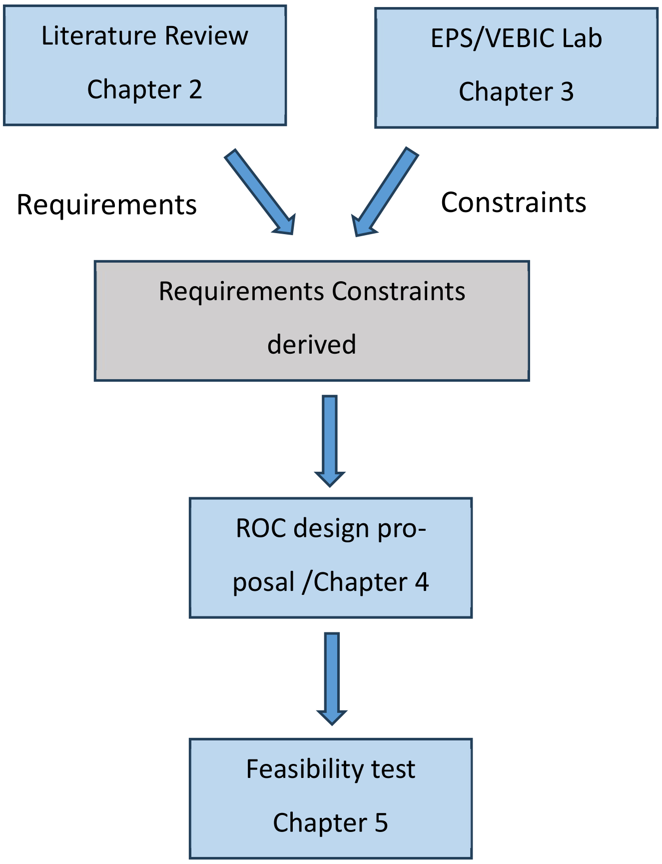

Figure 1. Thesis structure: LR and EPS baseline → derived requirements & constraints → ROC design (Ch. 4) → feasibility & gaps (Ch. 5).

2 Literature Review and Technology Landscape

2.1 Evolution and State of the Art in Autonomy for Maritime and Power Sectors

The shifting towards autonomous operation has evolved in different trajectories within the maritime and power sectors. In the maritime domain the progress has been systematic and regulation-driven that guided by international frameworks and phased implementation roadmaps, in contrast, the power sector has faced more gradually through ongoing digitalisation and the progressive centralisation of supervisory functions, in this section reviews the development trends within both sectors and examines their significance for the design and implementation of Remote Operation Centres (ROCs) intended for autonomous power plant applications.

2.1.1 Maritime Autonomy

The maritime sector has improved towards autonomy, and that through staged roadmaps that combine technological innovation and with regulatory guidance, an early developments has included the adoption of autopilot systems, radar and GPS-based navigation, which provided the foundation for more advanced capabilities, since 2010, the integration of sensor fusion, real-time control and remote operation platforms has accelerated the development of Maritime Autonomous Surface Ships (MASS), demonstrator projects such as the *Yara Birkeland*, the vessel that represents the world`s first fully electric and zero-emission autonomous container, to illustrate the gradual transition from manual operation to remote control and eventually to full autonomy (Munim, 2022).

Finland has contributed significantly in this trajectory, the Azipod propulsion system that developed collaboratively by Wärtsilä Marine, Stromberg and the Finnish National Board of Navigation was the first installation on the icebreaker *Seili* in 1990. System introduced a high degree of manoeuvrability later supported the development of autonomous navigation platforms (ABB Marine, 2020). More recently the integration of Vessel Traffic Services (VTS) with coastal surveillance systems has enhanced situational awareness through the fusion of radar and AIS data, also improving the target detection and tracking the performance by leveraging varied information fusion methods, such as Dempster, Shafer evidence theory (Wu, Wu, Ma, & Wang, 2023)., A central enabler of maritime autonomy is the International Maritime Organization (IMO), which has introduced a regulatory framework for the Maritime Autonomous Surface Ships (MASS). The framework defines four levels of autonomy, ranging from Level 0 (manual operation) to Level 4 (fully autonomous operation without human intervention) (Goerlandt, 2020; Klein et al., 2020). The existence of a regulatory framework ensures interoperability and phased adoption and distinguishing maritime development from the more fragmented progress seen in other sectors.

2.1.2 Power Sector and Autonomous Power Plants

The transition moving to the autonomy in the power sector has progressed more gradually than in the industrial domains, while most modern generation units, including wind, solar and hydropower plants, already employ advanced automation for local process control and optimisation, this automation represents only a partial form of autonomy, as decision making remains rule based and limited within individual plant boundaries. Anatomy system by adaptive coordination and predictive control across multiple assets, is still at an early stage of development.

Traditional Supervisory Control and Data Acquisition (SCADA) systems and Energy Management Systems (EMS) have improved monitoring and stability at the transmission level, but they have provided only limited autonomy within distribution networks. As stated by Di Silvestre et al. (2020, p. 3), “Smart Grid infrastructures are still fragmented across different voltage levels and ownerships, limiting the real-time coordination of distributed energy resources and the achievement of full autonomy in system operation." This fragmentation continues to restrict the scalability of decentralised control and highlights the need for supervisory integration through Remote Operation Centres.

Industrial practice during the last decade demonstrates clear progress toward remote and predictive operation. Wärtsilä's Expert Insight platform enables predictive maintenance and lifecycle optimisation, with more than ninety-six percent of reported issues resolved remotely without on-site intervention (Wärtsilä, 2023a)., through support centres in Houston and Trieste, Wärtsilä supervises over two hundred fifty power plants, by providing continuous diagnostics and remote lifecycle management (Wärtsilä, 2023b). Valmet also has achieved similar benefits by consolidating six biomass power plants into a single control room, which has improved operational efficiency and coordination (Valmet,2021). The Automation in renewable energy systems has also evolved rapidly. Chen et al. (2021) presented a deep learning aided model predictive control framework for wind farms that enhances automatic generation control through dynamic wake interaction modellng. This approach illustrates how intelligent supervisory algorithms can extend automation towards more adaptive and cooperative control.

At the grid level, the Siemens Finland Vibeco project integrates smart buildings, energy storage, and flexible loads into a digital platform that supports grid balancing and enables participation in energy markets (Siemens, 2022).

Despite these advancements, full autonomy at the multi plant and system coordination level remains uncommon. Most industrial plants still operate within predefined supervisory limits that correspond to intermediate stages of autonomy. Fragmented goverance and the lack of unified supervisory frameworks continue to slow down the progress towards large scale autonomous operation (Liu et al., 2020; Zürn & Faude, 2013).

2.1.3 Synthesis and Implications for ROC

The difference between the two sectors appears quite evident. In the maritime domain, autonomy has evolved under structured roadmaps that have been strongly supported by regulatory clarity and coordinated international efforts, which altogether have allowed a phased and rather predictable integration of new technologies. The power sector, on the other hand, has advanced in a slower and more organic manner, driven mostly by ongoing digitalisation and the steady centralisation of supervisory functions, yet its progress has often been limited by heterogeneous infrastructures and fragmented forms of governance that still vary across operators and regions.

Even with these differences both sectors show a convergence towards Remote Operation Centres (ROCs) as supervisory hubs. In maritime applications, ROCs function as transitional nodes between manned, remote, and autonomous operation. In the power sector, they provide multi-plant supervision, predictive optimisation, and reduced dependence on on-site intervention, for autonomous power plants this convergence means that ROCs must integrate constraint aware supervisory control using model predictive control, adaptive intelligence through reinforcement learning, predictive modelling through digital twins, and reliable communication infrastructure. These requirements identified from the state of the art form the basis for the ROC design proposal presented in Chapter 4.

2.2 Data Foundations: Sensors, Acquisition, and Computational Topologies

Autonomous power plants rely on sensing as the main source of state information for both local deterministic control and supervisory decisions in the Remote Operation Centre (ROC). When compared with more traditional SCADA-based systems, ROC-oriented installations usually require a higher sensor density, tighter temporal alignment, stronger observability, and also more reliable data quality. These requirements are needed to support real-time optimisation, predictive maintenance, and different kinds of fault management. In this way, sensors are not only reporting devices but they actively influence state awareness, which at the same time restricts and enables safe control actions.

This section discusses sensing and observability in the supervisory framework of autonomous power plants. First, basic concepts of observability are introduced together with performance aspects of different sensing modalities. After this, redundancy, calibration, sensor fusion, and time synchronisation are considered, since they are important for dependable state estimation. The discussion then moves to data acquisition and handling of time-series, where different computational topologies such as edge, fog, and cloud are taken into account. Cybersecurity aspects that are directly related to sensors are also noted, followed by short remarks on self-powered and passive sensing technologies. At the end, some industrial cases are presented and the key requirements for ROC design are summarised.

2.2.1 Multimodal and Redundant Sensor Architectures

Following LaValle (2011), a sensor can be defined as a mapping
$\mathrm{h}: \mathrm{X} \rightarrow \mathrm{Y}$ where X denotes the physical state space, Y the observation space, and $\mathrm{y} = \mathrm{h}(\mathrm{x})$
the observation. Since h is typically many-to-one, different states in X may correspond to the same observation y. The preimage $h^{-1}(y)$, therefore represents the set of states that are consistent with a given observation. The usefulness of a sensing configuration depends on how effectively it reduces the information space (I-space), defined as the set of states consistent with all past and current observations. A smaller I-space improves state discriminability and enhances the quality of subsequent decision-making.
The core relations are (LaValle, 2011):

$$y = h(x) \tag{1}$$

the sensor maps state $x \in X$to an observation $y \in Y$

$$h^{\{-1\}}(y) = \{ x \in X \mid h(x) = y \} \quad (2)$$

The preimage $h^{-1}(y)$ collects all states consistent with the observation y.

$$\mathcal{I}_{k} = \bigcap_{i=0}^{k} h_{\varepsilon}^{-1}(y_i) \quad (3)$$

$\mathcal{I}_{k}$ means better state discriminability.

Definitions in Eqs. (1)-(3) follow LaValle (2011). In practice, the informativeness of a sensor is constrained by device physics and signal processing. As described by Fraden (2016), the most relevant parameters include resolution (the minimum detectable change), sensitivity (output change per unit input), accuracy (deviation from the true value), drift (long-term deviation), hysteresis (direction-dependent output), repeatability (variation under identical conditions), and the transfer function $\mathrm{y} = \mathrm{f(x)}$. These parameters collectively determine how reliably measurements reduce uncertainty in the I-space, especially under conditions of noise, saturation, and environmental variability.

Table 1. Sensor Performance Characteristics (Fraden, 2016)

Property	Description
Resolution	Minimum detectable change in input signal
Sensitivity	Output change per unit input
Accuracy	Deviation between measured and true value
Drift	Long-term deviation due to aging or environment
Hysteresis	Output variation depending on input direction
Repeatability	Variability in repeated measurements under identical conditions
Transfer function $\mathrm{y} = \mathrm{f(x)}$	Functional sensor response

These parameters reflect how much the sensor helps the system to understand its internal state, especially when facing noise, signal overload, or changing conditions.

Application-oriented sensing modalities in power plants cover thermal, mechanical, and flow processes. Temperature monitoring is typically performed using resistance temperature detectors (RTDs) and thermocouples, which provide input for thermal management and combustion optimisation. Pressure sensing in fuel, steam, and hydraulic circuits relies on piezoelectric or capacitive devices, while flow measurement is achieved using ultrasonic and differential-pressure flowmeters to support mass balance and cooling system control. Mechanical condition monitoring is enabled by triaxial MEMS accelerometers, which detect imbalance and misalignment in rotating machinery (Dhanraj et al., 2020). These fixed modalities are complemented by non-contact and spatially distributed approaches. Electro-optical and thermal imaging systems support remote inspection and hotspot detection, LiDAR provides three-dimensional geometry for UAV-based inspections, millimetre-wave radar maintains robustness under low-visibility conditions, and fibre-optic sensing (DTS/DAS) enables distributed monitoring along cables, pipelines, and offshore assets with immunity to electromagnetic interference (Oliveira et al., 2024; Zhu et al., 2022).

Redundancy forms an essential mechanism for improving observability, and reliability in supervisory systems. Modal redundancy refers to measuring the same variable by applying different physical principles, for example ultrasonic and differential pressure flowmeters, spatial redundancy employs identical sensors installed at different locations to capture gradients and identify local faults. Of time redundancy is based on repeated sampling to reduce noise and follow drift under fast varying thermomechanical conditions (LaValle, 2011; Fraden, 2016). These mechanisms enable cross validation, probability checking and gradual degradation behaviour, all of which are important when ROC decisions need to be made under delayed or incomplete data conditions.
Sensor fusion transforms redundancy into actionable state information. In linear Gaussian settings, Kalman filtering provides minimum variance estimates by recursively

combining predictions with new observations (Kalman, 1994). For nonlinear dynamics or non Gaussian noise, Bayesian approaches such as extended or unscented Kalman filters and particle filters propagate posterior distributions (LaValle, 2011). When noise statistics are uncertain but physical constraints are known, set-based filtering eliminates inconsistent states without requiring probabilistic assumptions (LaValle, 2011). In ROC-oriented deployments, fusion typically operates at the edge or fog level to enable rapid closure of local control loops, while higher-level fusion at ROC or cloud scale supports fleet-wide estimation, fault detection, and digital twin synchronisation.

2.2.2 Self-Powered and Passive Sensors

In autonomous and remotely supervised energy systems, such as those coordinated through Remote Operation Centres (ROCs), sensors are often deployed in locations where access is limited or where stable power supply cannot be guaranteed. In these cases, self-powered and passive technologies provide a practical solution. Examples include piezoelectric, inductive, and RFID-based devices that operate with minimal or no local energy storage. They either harvest energy from the surrounding environment or draw interrogator power through backscatter. Such technologies are particularly relevant in hazardous, enclosed, or remote areas where battery replacement or cabling is not feasible. By converting vibration, strain, thermal gradients, airflow, or electromagnetic fields into electrical power, these sensors reduce maintenance demands, improve safety, and extend the monitoring reach of autonomous systems (Bai, Jantunen, & Juuti, 2018).

Several modalities have been explored in power and industrial settings. Piezoelectric harvesters (PZT/ZnO) capture mechanical energy from rotating machinery and vibrating structures (Raja, Umapathy, Uma, & Usharani, 2023). Triboelectric nanogenerators (TENGs) use contact electrification to power gas sensors and low-power telemetry under variable operating conditions (Yu et al., 2024). Thermoelectric generators (TEGs) convert heat gradients near engines, turbines, or exhausts into usable electricity. Wind and fluid harvesters make use of airflow or pressure variations in ducts and pipelines. Field trials have shown that wind-powered carbon monoxide monitoring can reach transmission distances of about 1.5 km without the need for batteries (Liu et al., 2020). In maritime applications, TENG-powered ammonia sensors placed in engine rooms have demonstrated detection limits close to 0.2 ppm, with vibration-powered wireless telemetry allowing continuous leak detection without wired connections (Yu et al., 2024). Cross-sector demonstrations also indicate transferability. For example, wood-based TENG harvesters combined with CNT gas sensors for ammonia monitoring in cold-chain logistics suggest a pathway for distributed leak detection at plant scale, where ultra-low-power nodes and extended service autonomy are required (Zhang et al., 2023).

At the ROC level, self-powered sensors contribute to early fault detection, emissions and leak surveillance, and predictive maintenance while reducing the need for frequent human intervention. Their small size, local energy autonomy, and event-driven operation lower energy consumption at the sensing edge and decrease transmission frequency, which is important in bandwidth-constrained environments (Di Silvestre et al., 2018). Integration into industrial cyber-physical and IIoT networks is possible, although protocol stacks such as OPC UA and MQTT can be too demanding for ultra-low-power endpoints. This has motivated the use of lighter protocol adaptations and hybrid gateways at the fog layer. Energy-aware sampling policies, supported by machine learning, allow the adjustment of sampling and reporting rates according to harvested energy reserves and the relevance of detected events. This approach extends operational lifetime while maintaining diagnostic fidelity (Bai, Jantunen, & Juuti, 2018; Liu et al., 2020). Security remains a challenge for these systems, since conventional cryptography requires high computational effort. Lightweight security methods (see 2.2.6) are therefore needed to maintain confidentiality and authenticity within the limited energy budgets of self-powered nodes.

2.2.3 Calibration, Synchronisation (PTP/NTP), and Reliability

Traceable calibration, precise time synchronisation, and device reliability are prerequisites for trustworthy autonomous operation in ROC-managed systems. Calibration must account for temperature coefficients, ageing, and installation-related deviations, and should follow recognised standards such as IEC 62828-1 (2017) and the NIST Handbook 44 (2021). Emerging practice includes peer-to-peer sensor comparisons, scheduled recalibration during low-demand periods, and the use of online health indicators as part of predictive maintenance strategies.

Time synchronisation requirements vary according to function. For supervisory logging, the Network Time Protocol (NTP) provides millisecond-level accuracy, which is typically sufficient, by contrast, the control-critical and analytics streams often demand submicrosecond precision, achieved through the Precision Time Protocol (PTP) specified in IEEE 1588, IEEE C37.238, and the IEC 61850-9-3 time profile (Eidson, 2006; IEEE, 2008). Architectures commonly employ GPS-disciplined grandmasters, boundary clocks integrated in time-aware switches, and continuous monitoring of quality-of-time indicators, these synchronisation becomes a critical enabler for maintaining the digital-twin updates and for achieving PMU-like analytics, event correlation and coordinated operation of hierarchical MPC structures, also including the validation of twin-in-the-loop configurations (Zhang & Liu, 2019).

Ensuring reliability under demanding environmental conditions, for instance high temperature, continuous vibration and electromagnetic interference, depends largely on compliance with established safety standards like IEC 61508. Alongside these additional mechanisms like the redundant sensing is built-in self-testing, plausibility verification and metadata quality flags are used to guarantee trustworthy data for estimation and control. The use of modular sensor packaging further helps to reduce the mean time to repair and allows fast replacement of units even in remote or otherwise resource-limited sites

Industrial deployments highlight the role of synchronisation in practice, while ABB's Longmeadow microgrid in Johannesburg coordinates photovoltaic generation, to battery energy storage, and diesel generators using sub-microsecond PTP alignment within the Microgrid Plus system. This synchronisation stabilises inverter, generator interactions during rapid transitions between grid-connected and islanded operation, while maintaining load balancing and generator coordination through consistent, time-aligned communication with distributed equipment (ABB, 2015). Similarly, the Secure Scalable Microgrid Test Bed at Sandia National Laboratories applies IEEE 1588-compliant PTP to synchronise sensors, controllers, and hardware-in-the-loop simulators across photovoltaic, storage, and diesel subsystems (Eidson, 2006; IEEE, 2008; Glover, Neely, Lentine, & Finn, 2012; Sandia National Laboratories, 2022). The resulting common time base enables real-time testing of autonomous control strategies, fault-tolerance mechanisms, and digital-twin integration; it can be concluded that high-fidelity timing is therefore a functional prerequisite for ROC-orchestrated autonomy (Zhang & Liu, 2019).

In summary, traceable calibration, high-precision time synchronisation with quality-of-time monitoring, and reliability engineering at both device and data-quality levels are essential prerequisites for autonomous power plant operation. These conditions ensure that measurements remain coherent, trustworthy, and time-aligned across supervisory functions. In Chapter 4 the ROC in design developing, by providing the requirements foundation for digital-twin-assisted MPC and reinforcement learning, where the quality of decision-making depends directly on the fidelity of sensed data

2.2.4 Data Acquisition and Time-Series Handling

Autonomous power plants depend on continuous and high-frequency measurement streams, which is typically in the millisecond or microsecond range, to detect disturbances, stabilise control loops and enable predictive diagnostics, within the ROC, the performance of autonomous decision-making is determined by how efficiently of measurement data are acquired, stored an processed, and how it visualised, wide sensing base covering electrical, thermal, mechanical, and environmental variables forms the foundation of this process. When high-resolution signals are correctly managed, they enable rapid detection of changes, support advanced control strategies such as Model Predictive Control (MPC), and provide input for digital-twin simulations. Previous studies have shown that well-structured time-series management improves maintenance prediction, decreases unplanned downtime, and enhances situational awareness across the system (Kejser et al., 2017; Grzesik & Mrozek, 2020).

Three frequent challenges can be identified in practice, first, heterogeneous measuring devices that introduce irregular sampling and temporal misalignment, which it must be corrected before signals and can be compared or fused (Zhang & Liu, 2019). Second, the volume of persistent measurements can reach several million points per second that requiring the data platforms to scale without compromising reliability (Jensen et al., 2017). Third, the anomaly detection must be performed within sub-second latency in order to react to events such as generator malfunctions or fuel system leaks, since the delayed response may endanger both safety and availability (Grzesik & Mrozek, 2020; Dhanraj et al., 2020).

Table 2. Main challenges in time-series data management

Challenge	Description
Irregular sampling and time alignment	Heterogeneous sensors may acquire at slightly different times, complicating comparison and fusion; alignment techniques are required to correct temporal mismatches (Zhang & Liu, 2019).
High volume and scalability	Modern deployments can generate millions of data points per second; platforms must scale while maintaining reliability (Jensen et al., 2017).
Fast anomaly detection	Detection and response should operate at sub-second latencies for events such as generator anomalies or fuel leaks (Grzesik & Mrozek, 2020; Dhanraj et al., 2020).

Different approaches have been developed to address these requirements at the edge and data storage layers. At the edge level, embedded devices such as Raspberry Pi or Speedgoat controllers are commonly used to reduce network load by correcting timestamps, validating data formats, and applying compression. These devices also enable local fault detection at an early stage (Grzesik & Mrozek, 2020). In the storage layer the use of time-series databases like InfluxDB, TimescaleDB and Alibaba Lindorm has become common, as these systems are particularly designed for handling ordered high-frequency data streams. They offer efficient data persistence together with fast retrieval and configurable retention management, features that are essential for tracking both historical trends and the real-time operating conditions within the ROC (Jensen et al., 2017). Benchmarking studies like TSM-Bench show that InfluxDB and TimescaleDB generally reach the high performance in both absorption and querying, while Lindorm, though does not reflect the quickest in terms of raw throughput, and brings built-in machine learning tools that can directly assist ROC diagnostics (Jensen et al., 2017).

The Real-time analysis is most often carried out through Complex Event Processing, either built inside the database itself or by using separate stream-processing platforms such as Apache Flink. This approach allows continuous pattern detection, sliding-window trend evaluation and near-instant alerting, there is growing evidence leads to that stream-based analysis can reveal operational problems faster and with better accuracy than the conventional SCADA-based alarm systems, especially in smart-grid and off-grid cases (Grzesik & Mrozek, 2020; Dhanraj et al., 2020).

Visualisation works as the final layer that ties these functions together, refreshing operator screens with only minimal a delay, ideally within about a second, proper visualisation focuses attention on key parameters, short-term behaviour and selected historical comparisons so that both human operators and automated agents can read the situation and react to changes without delay (Grzesik & Mrozek, 2020; Dhanraj et al., 2020).

The processes of measurement acquisition and time-series handling are closely connected to control and modelling tasks. The live data feeds become direct input for control algorithms and digital-twin models, which then allow online adjustment of parameters, real-time simulation of possible actions and constant monitoring of system performance against expected trajectories. In ROC-supervised autonomy this tight integration helps to recognise and mitigate emerging faults early, often before any human intervention is required.

2.2.5 Computational Topologies: Edge, Fog, and Cloud

Intelligent control for autonomous power systems imposes stringent requirements on real-time execution, low latency, and high processing capacity (Vidal et al., 2022). three-tier computing structure consisting of edge, fog and cloud layers has proved to be a practical basis for meeting the performance constraints of ROC-supervised distributed energy systems. In such setups the sensor data is gathered through SCADA connections and local edge devices, processed as near to the source as possible and then sent forward in an optimised form to the upper layers for broader coordination (Vidal et al., 2022). Within this arrangement the edge level works almost like a local extension of the cloud, just one hop away from the sensors or actuators, which helps to cut latency and enables fast actuation whenever time-critical reactions are required (Shi et al., 2016). An experience from deployments in microgrids, autonomous vehicles and industrial automation has shown that these edge units are capable of validating control commands, detecting anomalies and carrying out the first level of sensor fusion already on-site, without the need to rely on remote servers (Satyanarayanan, 2017)., The fog layer was introduced by Vidal et al. (2022) as an intermediate step between edge and cloud. Fog nodes are typically installed in substations or in site-level gateways and provide more computational capacity than edge devices while maintaining lower latency than cloud services. According to Shi & Dustdar. (2016), fog computing enables coordination across local networks, buffering of critical signals during connectivity disruptions, and pre-processing of data before cloud transmission.

The cloud layer is located in remote data centres and provides the largest computational and storage resources. Satyanarayanan (2017) state that cloud computing is mainly applied to resource-intensive tasks such as large-scale data analytics and training of

artificial intelligence models. Due to its higher latency, cloud computing is unsuitable for real-time control functions, but it supports predictive maintenance, historical analysis, and system-wide optimisation.
A comparative overview of the three layers is presented in Table 3.

Table 3. Comparative characteristics of edge, fog, and cloud computing layers in autonomous control environments adapted from (Shi et al., 2016; Satyanarayanan, 2017).

Layer	Latency	Compute type	Example applications
Edge	< 10 ms	Embedded, real-time	Combustion control; actuator response
Fog	10–100 ms	Site-level gateways	Load balancing; anomaly detection
Cloud	> 100 ms	Virtualised clusters	Fleet analytics; AI training; long-term planning

It was noted by Eidson (2006) and IEEE (2008) that accurate time synchronisation is essential for coordinated performance across edge and fog layers. Network Time Protocol (NTP) provides millisecond-level accuracy but is insufficient for substation-level or real-time control. Precision Time Protocol (PTP), defined in IEEE 1588 and profiled in IEC 61850-9-3, achieves sub-microsecond precision and is therefore adopted in systems where reliable communication and accurate event correlation are required. Without this synchronisation, phasor tracking, fault detection, and sensor-data fusion tasks may be compromised.

Each layer introduces advantages and limitations. Control executed at the edge minimises latency but restricts the use of complex models. Cloud-based execution allows large-scale optimisation but adds delay and increases dependency on wide-area connectivity. Edge and fog computing improve local autonomy and resilience, but their distributed structure raises exposure to physical and cyber security risks. Shi and Dustdar (2016) highlight the role of adaptive middleware that dynamically reallocates tasks across layers according to operating context and network conditions.

In ROC-managed deployments, the mapping of layers follows a clear division of roles. The edge layer typically consists of programmable logic controllers, embedded controllers, and real-time prototyping platforms such as Speedgoat for critical control functions. The fog layer is commonly realised through substation or hub-level equipment, which coordinates local decision-making, executes backup logic, and performs diagnostics. The cloud layer connects multiple plants, enabling long-term planning, artificial intelligence training, and performance benchmarking (Yıldırım et al., 2025).

Within the EPS Laboratory, elements of this layered model are already present. Real-time control at the edge is provided by Speedgoat and embedded devices, while fog-level tasks are implemented within the SCADA system for test sequencing and procedure management. Cloud-level functions are supported through remote monitoring and collaborative analysis enabled by VPN connections and shared storage. Strengthening the interfaces between layers through semantic middleware, mirrored decision logic, and secure communication tools would align the laboratory with the computational substrate described in Chapter 4 and move it closer to full ROC-based autonomy.

2.2.6 Sensor-Centric Cybersecurity in the Data Plane

In recent years, the extensive deployment of sensor networks within ROCs has markedly expanded the exposure of such systems to a variety of cyber threats, most what notably at the sensing layer, in this regard, vulnerabilities including weak device authentication, false-data injection, and the exploitation of low-power wireless communication links may significantly compromise operational stability, data integrity, and eventually, the quality of supervisory decision-making, these issues become particularly critical in constrained devices that possess only limited computational and energy resources, where the implementation of conventional security stacks proves impractical (Al-Quayed et al., 2024; Gupta et al., 2023; Qiu et al., 2020). Therefore, research has increased in emphasizing in the need for a layered defence strategy, one that effectively combines hardware-rooted identity mechanisms with physical-layer anti-spoofing and lightweight cryptographic methods suitable for embedded environments (Beaulieu et al., 2015; Maes, 2013).

One promising approach for strengthening authentication and data integrity is the application of Physically Unclonable Functions (PUFs). In this regard, PUFs provide a hardware-rooted basis of trust by exploiting unavoidable manufacturing variations to create unique challenge–response pairs, allowing each device to be identified without relying on stored long-term keys. This becomes particularly beneficial in wireless sensor networks and industrial IoT nodes where the available secure memory is limited (Maes, 2013). Recent studies have combined PUFs with lightweight authentication protocols built on XOR operations, hashing functions, and session keys derived from real-time responses. Such designs enable mutual authentication at very low energy cost. Moreover, public-key cryptographic schemes have been developed in which elliptic-curve keys are produced directly from PUF outputs, thus removing the need for pre-provisioned secrets and simplifying the onboarding of unattended ROC devices (Akhundov et al., 2020; Gupta & Varshney, 2023).

At the same time the sensor networks must to be defended against spoofing and replay attacks that can disturb time-critical tasks such as event correlation, fault detection, and control. Physical-layer authentication methods, including RF fingerprinting, carrier-frequency-offset checking, and channel-state analysis, that have been shown to distinguish legitimate transmissions from forged ones even under constrained wireless conditions (Liu et al., 2022). These techniques will perform best in when supported by secure time synchronisation. Implementations of the IEEE 1588 Precision Time Protocol (PTP) that use timestamped packets and cryptographic watermarks provide guarantees of both freshness and origin, and when paired with machine-learning classifiers lead to these solutions to enhance detection of replay attempts and reduce false-alarm rates (Qiu et al., 2020).

Lightweight cryptography forms the third key element of this defence strategy. Conventional RSA- or TLS-based mechanisms impose heavy memory, computation, and energy demands on embedded sensors, limiting their practicality. Alternatives such as the SIMON and SPECK block ciphers achieve confidentiality with much lower resource use, making them suitable for harsh or remote deployments (Beaulieu et al., 2015). Likewise, compact elliptic-curve settings allow efficient key exchange and message authentication. Hybrid schemes that combine ECC with PUF-derived keys further strengthen security by enabling devices to create secure sessions without storing long-term secrets (Akhundov et al., 2020; Gupta & Varshney, 2023).

Deploying these mechanisms in industrial settings still presents challenges. Standard stacks such as OPC UA and IEC 61850 are typically bound to PKI and TLS infrastructures, which exceed the capacities of ultra-low-power devices (Gupta et al., 2023). A practical solution is to delegate heavy cryptographic and identity-management tasks to fog nodes or middleware, allowing constrained devices to participate securely with minimal overhead. Behaviour-based adaptive trust models that depend on traffic patterns and signal features can add complementary protection by reducing dependence on static credentials (Qiu et al., 2020). Before implementation, candidate protocols should be validated through formal analysis methods such as AVISPA or BAN logic to confirm resilience against replay, desynchronisation, and man-in-the-middle attacks (Gupta & Varshney, 2023).

To summarise the hardware-rooted identity, physical-layer anti-spoofing supported by trusted time, and lightweight cryptographic schemes together form a coherent foundation for protecting sensor data in ROC-managed autonomous power systems, these same components define the security assumptions that underpin the models discussed in Chapter 4

2.2.7 Industrial deployments demonstrating ROC readiness

This subsection presents six industrial deployments that illustrate how integrated sensing and remote operations support ROC readiness in different contexts.

Wärtsilä has developed Remote Support and Data Management services that utilise real-time information from sensors embedded in engines, fuel systems, turbochargers, emissions units, and thermal circuits. Typical variables include vibration, cylinder temperature, turbocharger speed, oil pressure, and emission levels. These are streamed to global

Expertise Centres for live monitoring and analysis. A practical example was reported where repeated night-time shutdowns were resolved when remote specialists identified a faulty speed-sensor cable using live feeds and historical logs. Guided replacement on site restored full operation in less than two hours, avoiding prolonged downtime and travel expenses (Wärtsilä, 2023a).
In the turbomachinery domain, MAN Energy Solutions reports the use of more than 40,000 sensors that provide input to an AI-based Model Predictive Controller. The system continuously tracks flow behaviour, pressure, temperature, and vibration. Company figures indicate that this approach has reduced unplanned downtime by approximately 30 percent and achieved annual cost savings of up to €2 million through lower maintenance requirements, extended equipment life, and improved efficiency (MAN Energy Solutions, 2023).

Shell has introduced autonomous operations in a chemical-plant setting. In this deployment, optical sensors and robotic systems perform monitoring tasks that would normally require on-site personnel. The plant operates continuously with only two remote operators. Robotic units equipped with thermal and optical cameras conduct routine inspections and diagnostics, reducing the need for human presence in hazardous environments while maintaining operational continuity (Chemical Processing, 2023).
Autonomous aerial inspection has been applied by Sensorem in remote Australian mining regions. The company operates UAVs equipped with electro-optical and LiDAR sensors, which are used for inspections in areas with difficult terrain. More than 11,000 beyond-visual-line-of-sight missions have been reported. Real-time sensor feedback supports obstacle avoidance, infrastructure assessment, and transmission of inspection data to central control hubs. This demonstrates how mobile sensing platforms can extend the reach of ROC supervision (FlytBase, 2023).

The Calistoga Resiliency Centre represents an example of hybrid microgrid implementation. The facility integrates hydrogen fuel cells with battery storage and applies fibre-optic distributed temperature sensing together with AI-assisted controllers. These elements provide load balancing, thermal monitoring, and equipment protection. During external outages, the microgrid is able to supply power independently for up to 48 hours, with sensor-driven automation playing a central role in ensuring resilience (Calistoga City Council, 2022).

Predictive analytics is exemplified by ROTEC`s PRANA platform which is aggregating extensive vibration and the temperature measurements collected from thousands of sensors distributed across multiple thermal power plants, then at this regard will notice the system continuously evaluates the operational condition of critical components such as turbines, pumps and generators that enabling early detection of performance degradation and potential mechanical faults. Since its deployment, then the platform has reportedly prevented more than 300 major equipment failures, clearly demonstrating the tangible benefits of predictive analytics supported by high-fidelity sensor data (ROTEC, 2022). These industrial implementations illustrate the dense and the well-integrated sensing architectures when it is complemented by remote diagnostics, the autonomous inspection and predictive analytics can yield measurable enhancements in plant availability and operational safety and overall efficiency. Furthermore, they are providing concrete design patterns and functional insights that directly inform ROC development considerations that discussed in Chapter 4.

2.3 Control and Decision-Making Frameworks

The necessitate of Autonomous energy systems is advanced the control and decision-making methodologies capable of addressing complex system dynamics, stringent real-time constraints, inherent safety requirements, and long-term operational objectives. In this regard, while the recent research has evolved beyond traditional reactive feedback mechanisms by incorporating elements of prediction and optimisation by learning, then and resilience. With all these components are providing a comprehensive basis for maintaining reliable performance under varying conditions of uncertainty and environmental change. Accordingly, this section outlines the principal strategies employed in remotely operated and autonomous hybrid energy systems and establishes the conceptual foundation for the Remote Operation Centre (ROC) design that is further developed in Chapter 4.

Proportional-integral-derivative (PID) control has long been the most widely used method in industrial automation. It is valued for its simple structure and its ability to work across a broad range of processes. However, tuning its parameters becomes difficult in systems with nonlinear behaviour or in cases where operating conditions vary over time, or strong coupling between control variables. PID continues to perform well for single-input, single-output loops close to steady-state operation. However, its effectiveness decreases in multivariable systems with interactions or in cases where strict constraints must be respected (Qin & Badgwell, 2003). To address these challenges, Model Predictive Control (MPC) has been widely adopted. MPC computes constrained optimal actions across a receding horizon using a system model. It is capable of explicitly handling multivariable couplings and state/input constraints and has been applied in areas such as microgrids, engine systems, and building energy control (Rawlings et al., 2020).

With the growth of data availability and the increasing difficulty of first-principles modelling, AI-based controllers, particularly reinforcement learning (RL), have become more prominent. These methods derive policies from data, either offline or during interaction with the system. They are particularly suitable when the plant is only partially observable or when non-stationary disturbances affect performance (Yu et al., 2024). In energy applications, RL and learning-enabled controllers have been studied for adaptive energy management, fault handling, and incremental performance improvement (Zahra & Singh, 2022).

Fault-tolerant control (FTC) contributes to safe operation when components fail or when abrupt changes occur. Passive FTC provides robustness by design without changing the control structure, while active FTC detects faults and reconfigures the controller in response (Ghosh et al., 2020). In addition, decentralised and distributed control schemes have become important for modular or geographically dispersed assets such as islanded microgrids and hybrid maritime power systems. These methods allow local autonomy while still maintaining coordination at higher levels (Luo et al., 2023).

In practice, these control approaches are combined in hierarchical structures. Local MPC enforces the fast constraints and ensures the feasibility, and an upper-level learner, for example RL can adapts supervisory policies to evolving conditions. Meanwhile an FTC layer monitors the system and activates reconfiguration when necessary. Such layered control reconciles the demands of real-time responsiveness with long-term optimisation in ROC-supervised deployments. The subsections that follow examine each method's principles, application contexts, and main challenges. These insights inform the architectural design, module interfaces, and validation strategy applied in the ROC proposal of Chapter 4.

2.3.1 Model Predictive Control (MPC)

Model Predictive Control (MPC) has became one of the most extensively applied methods for the supervisory control in systems that what demand predictive, multi-objective, and constraint-aware decision-making. Regarding, its significance to autonomous energy platforms arises from the capability to system anticipation the behaviour and optimise control actions while adhering to the underlying physical limitations. This makes MPC particularly well-suited for hybrid energy systems that exhibit nonlinear load profiles, to discrete the switching behaviours, and tightly coupled thermal, electrical, and mechanical subsystems.

The fundamental principle of MPC is established upon a discrete-time representation of the controlled system, which is most commonly formulated in a linear state-space structure (Rawlings et al., 2020):

$$x_{\{k+1\}} = A\, x_k + B\, u_k \tag{4}$$

Here, X_k represents the system state at time step $k1$, and U_k is the control input. This model serves as the basis for forecasting the future behaviour of the system and for solving an optimisation problem that balances the defined performance objectives against operational and physical constraints. In the context of energy systems, such objectives typically encompass minimising fuel consumption, maintaining battery state-of-charge within desired limits, reducing NOx emissions, and constraining mechanical stress or component wear to acceptable levels.

Several industrial implementations have demonstrated the practicality of MPC in real-time operational environments, time-varying of MPC has been applied to wind turbines for dynamically adjusting blade pitch and generator torque in response to continuously changing wind conditions (Dickler et al., 2021). In islanded microgrids, distributed MPC (DMPC) has been employed to coordinate local generation and storage assets while ensuring that voltage and frequency remain within prescribed tolerances (Liu et al., 2022). At the embedded level, explicit MPC (eMPC) achieves sub-millisecond control by relying on precomputed control laws, thereby enabling real-time operation on resource-constrained hardware platforms, such as those used in remote or mobile environments (Rawlings et al., 2020; Dickler et al., 2021).

Despite these demonstrated benefits, conventional MPC frameworks can exhibit sensitivity to model uncertainties and unmeasured system dynamics. Consequently, hybrid control architectures have emerged, integrating learning-based techniques such as reinforcement learning (RL) alongside MPC. In such configurations, RL agents are often employed to adapt control parameters or formulate long-term strategies, whereas MPC continues to enforce safety and feasibility within the short-term prediction horizon. This combined approach has been successfully implemented in energy systems characterised by high renewable penetration and variable fault conditions (Zhang et al., 2021).

Within ROC-managed infrastructures, MPC operates across several hierarchical layers. At the edge level, it provides rapid and deterministic control for localised subsystems.

At the fog layer, it performs supervisory coordination by leveraging updated models, constraint information, and prioritisation schemes. Its intrinsic capability to manage constraints, resolve multi-objective trade-offs, and integrate seamlessly with estimation and diagnostic tools establishes MPC as a central component of supervisory control design. However, the overall effectiveness of MPC remains dependent on the fidelity of the underlying system model, the accuracy of state estimation, and the computational efficiency of the solver implementation. These factors become especially critical in distributed autonomous systems, where plant dynamics are complex, nonlinear, and continuously evolving.

2.3.2 AI-Based and Reinforcement Learning Controllers

The growing complexity and geographical spread of modern power systems, driven by the integration of distributed energy resources (DERs), makes it necessary for ROCs to use controllers that can adapt on their own while still keeping safety, reliability, and efficiency under uncertain conditions, in this context, Artificial Intelligence (AI) and especially reinforcement learning (RL) are used to support data-driven decision-making and to build control strategies that work with only limited direct supervision.

Deep reinforcement learning (DRL) has been applied in a variety of power-system contexts. By learning policies through interaction with the environment, DRL is well suited to time-varying tasks such as voltage regulation, power balancing, fault handling, and demand-supply management. In one grid-reconfiguration study, a duelling double deep Q-network (D3QN) successfully adjusted the power-flow topology in the IEEE 14-bus system, demonstrating the feasibility of fully automated high-level decision-making (Damjanović et al., 2022). For emergency control purposes, deep reinforcement learning (DRL) has been applied to dynamic braking and under-voltage load shedding in the IEEE 39-bus test system, where it has demonstrated the capability to respond to faults without relying on predefined threshold rules (Huang et al., 2019). These studies suggest that

DRL can strengthen the autonomy of ROCs and improve overall system stability under nonlinear and nonstationary operating conditions.
Because the trial and the error of exploration can introduce risks in safety-critical environments, recent research has placed increasing attention on the concept of safe reinforcement learning. This approach embeds safety directly into both the training and deployment stages by incorporating mechanisms such as reward shaping, safety filters, constrained policy updates, and the inclusion of expert demonstrations. Contemporary surveys highlight the use of safe RL in frequency control, voltage regulation and microgrid management, with particular focus on integrating operational limits and safety constraints into the learning process itself (Bui et al., 2024; Yu et al., 2024). These techniques are especially relevant for ROC-managed systems, where ensuring safe and reliable operation remains an absolute requirement.

Another active area of development is multi-agent reinforcement learning. The physical distribution of DERs and communication limitations motivate decentralised and coordinated architectures. PowerNet is an example of a multi-agent DRL framework that coordinates distributed generators using location-aware rewards and structured communication. This enables scalable team-based control across multiple sites, which is consistent with ROC supervision of microgrid clusters and DER fleets (Chen et al., 2020). To support standardised evaluation, Recent surveys outline evaluation practices and stress-testing protocols that expose agents to diverse operating conditions prior to deployment (Bui et al., 2024; Yu et al., 2024).

In summary, DRL contributes to autonomy, safe RL ensures constraint-aware learning in safety-critical contexts, and multi-agent architectures provide scalability for geographically distributed assets. Together, these strands form a coherent basis for AI-based supervisory control in ROC environments. They also inform the design of supervisory AI modules in Chapter 4, including policy adaptation under uncertainty and integration with MPC for runtime-assured decision-making.

2.3.3 Fault-Tolerant and Decentralized Control

Autonomous power systems are increasingly distributed, which makes decentralised control and fault tolerance essential for stability and reliability, particularly under ROC supervision. Heavy reliance on central command structures brings several risks. These include communication delays, hardware failures, and cyber disturbances. To reduce such vulnerabilities, recent approaches shift more authority to local controllers. They also build resilience so that operation can continue even when part of the system fails.
Decentralised control lets each subsystem act on its own measurements with only limited coordination. This improves both responsiveness and scalability, especially in systems with low inertia and variable energy supply. In these situations, decentralised designs support real-time regulation and demand coordination without creating a single point of failure (Zahra & Singh, 2022; Ahsan et al., 2023).

Industrial middleware provides practical examples of this concept, the Resilient Information Architecture Platform for Smart Grid (RIAPS) offers resource monitoring, event management, and distributed fault detection across nodes. When a task or device fails, RIAPS performs local recovery actions. As a result, microgrid controllers and data-processing tasks can keep operating even if individual components are lost (Ghosh et al., 2020). Fault-tolerant control (FTC) contributes to performance preservation in the presence of sensor, communication, or actuator faults by detecting anomalies and reconfiguring control actions, for example, in the permanent-magnet machines, a detect-and-reconfigure scheme was able to maintain stability under fault conditions. This technique is transferable to applications such as wind farms, offshore platforms, and remote substations, where high availability is critical (Hao, Li, & He, 2011). Complementary approaches based on consensus further enhance recovery. Combining wavelet analysis with local agreement protocols allows groups of generators to remain synchronised despite sensor dropouts or communication failures, thereby reducing dependence on a central coordinator and exploiting device-to-device communication (Han, 2023).

From an architectural perspective, fault handling should be designed into communication middleware and control software from the outset. Beyond node-level resilience, system-level services also need to handle degradation detection, mode switching, and safe reconfiguration. The RIAPS platform again provides an example of this integration. It supports machine-learning modules and fault-diagnosis functions as core system components (Ghosh et al., 2019). In ROC-managed deployments, these capabilities help local units maintain safe operation and continue exchanging critical state information even when higher layers are degraded or temporarily unreachable.

Decentralised control, combined with explicit fault-tolerant mechanisms, allows scalable autonomy that can sustain safe operation during disturbances and partial failures. Together with predictive and learning-based strategies, these methods form the mathematical foundation for supervisory intelligence in ROC-managed systems. Their practical performance, however, it still depends on accurate state and parameter estimation, as well as careful management of model-plant mismatch and integration of predictive diagnostics. These factors create the motivation for adopting digital twin technologies, which are discussed further in Section 2.5.

2.4 Communication and Integration Protocols

Remote Operation Centres (ROCs) depend on communication infrastructures that can support real-time and secure data exchange across heterogeneous field devices and wide geographical areas. Legacy, centrally oriented links are often insufficient, as they do not provide the peer-to-peer interaction, rapid fault recovery, or interoperability required at scale. Modern deployments therefore integrate a range of standards, including OPC UA, IEC 61850, MQTT, Modbus, and CAN. Each is applied according to its role and layer within the system, with growing emphasis on decentralised and resilient communication (Tebekaemi & Wijesekera, 2018).

In addition, recent studies emphasise that communication and control must be closely coupled to maintain stability under conditions of renewable variability and high levels of automation. Effective communication architectures therefore not only provide data exchange but also enable dynamic coordination and robust supervisory control (Sajadi et al., 2019; Wang et al., 2022).

2.4.1 Data plane and time base

Several communication standards are applied in ROC-managed systems, each contributing specific capabilities at different layers of the architecture. OPC UA provides interoperable client-server and publish-subscribe communication patterns together with rich semantic data models. Companion specifications extend its use to smart-grid contexts and allow integration with IEC 61850-based environments (Tebekaemi & Wijesekera, 2018; Wang et al., 2022). IEC 61850 itself defines substation object models in the form of logical nodes and supports time-critical messaging through GOOSE and Sampled Values. These mechanisms enable protection and control actions with response times below four milliseconds and support progressive digital upgrades in utility environments (Rehtanz, 2003; Wang et al., 2022).

MQTT provides a lightweight publish-subscribe transport that is particularly suited to constrained communication links in remote microgrids or isolated facilities. Although it lacks native semantic modelling, it can be integrated through middleware layers that supply the required data context (Sajadi et al., 2019). At the equipment edge, Modbus and CAN continue to be widely used for simple and robust signalling. In ROC deployments, these protocols are typically gatewayed into higher-level namespaces for interoperability across the wider system (Bumiller et al., 2010). According to latency measurements by Mathiesen (2024), CAN and its extensions exhibit transmission delays from tens of microseconds up to several milliseconds, corresponding to a medium-to-high real-time performance classification in Table 4.

A common time base for these heterogeneous communication channels is established elsewhere in the stack through synchronisation protocols such as Precision Time Protocol (PTP) or Network Time Protocol (NTP). This ensures coherent data fusion and reliable event correlation across system layers, as discussed in Section 2.2.3.

Table 4. Protocol Feature Comparison for ROC Applications

Feature	OPC UA	IEC 61850	MQTT	Modbus	CAN
Communica-tion Model	Client/Ser-ver, Pub/Sub	Event-driven, Client/Server	Pub-lish/Sub-scribe	Re-quest/Res-ponse	Broad-cast/Mes-sage
Semantic Data Modelling	Yes	Yes	No	No	No
Real-Time Per-formance	Medium	High (<4 ms)	Medium	Low	Medium-High
Interoperabi-lity	High	High	Medium	Low–Me-dium	Low
Resource Effi-ciency (Edge)	Medium	Low	High	High	High
Native Security Features	TLS/PKI	IEC 62351	Basic	None	None

Note. Real-time performance categories are derived from typical latency ranges observed in industrial deployments. Latency classification for CAN protocols is based on measurements reported by Mathiesen (2024), where transmission delays ranged between approximately 70 µs and 7267 µs depending on frame type and network load.

The information shown in Table 4 comes from a careful review of trusted sources in communication and automation research. Features such as latency, support for data modelling, and performance at the edge level were compared based on findings from Tebekaemi & Wijesekera (2018), Rehtanz (2003), Sajadi et al. (2019), and Bumiller et al. (2010). The comparison shows that each protocol brings different strengths and challenges when used in real-time systems. It also points out that no single protocol works best for every task. Because of this, Remote Operation Centres (ROCs) benefit from using a layered mix of protocols that can work together across different parts of the system.

2.4.2 Data lifecycle and latency

For ROC applications, communication infrastructures must be capable of scaling with plant size while remaining regulation-aware and closely integrated with control functions (Sajadi et al., 2019; Wang et al., 2022). High- and low-frequency measurement streams from data acquisition units are normalised within an archive that applies consistent schema and quality flags. In parallel, streaming channels provide timely data ingestion for analytics and decision-making.

The end-to-end data path typically begins at the plant, passes through local control and acquisition systems, and proceeds to ROC ingestion services. From there, data are processed in AI modules, validated, and subsequently dispatched as control actions. Percentile latency and jitter targets are defined for each stage of this chain to ensure reliable timing performance. To maintain continuity during partial outages, modern designs increasingly employ peer-to-peer communication paths that allow local subsystems to exchange critical information without relying solely on central coordination (Tebekaemi & Wijesekera, 2018).

2.4.3 Cybersecurity Layers and Secure Interfaces

In ROC-based autonomous power systems, communication infrastructures must meet strict cybersecurity requirements in addition to performance objectives. As edge devices, substations, and cloud platforms become interconnected, every interface constitutes a potential attack surface. Without layered protection, such systems remain exposed to unauthorised access, data manipulation, replay attempts, and denial-of-service attacks (Sajadi et al., 2019; Wang et al., 2022).

OPC UA provides a structured security framework that includes Transport Layer Security (TLS) for encryption, certificate-based authentication using X.509 certificates, secure

session handling, and role-based access control. When configured correctly, these features support compliance with IEC 62443 (Wang et al., 2022). Field experience, however, shows that weak configurations such as open endpoints or outdated cipher suites can reduce their effectiveness. This underlines the need for secure defaults and strict configuration management in operational environments (Sajadi et al., 2019).
In IEC 61850-based systems, protection is specified in the IEC 62351 standard. The framework defines digital signatures, message authentication codes, and public-key infrastructure for key management. Network-level measures, including VLANs and zoning, are also used to separate critical substation traffic. These mechanisms help protect time-critical communication services such as GOOSE and Sampled Values, while still maintaining the sub-cycle performance needed for protection and control (Rehtanz, 2003). Even with these standards in place, insecure practices are still widespread. MQTT brokers are often deployed without authentication, and Modbus messages are still sent in clear text. OPC UA can also become a weak point if its security settings are left at default or incorrectly applied. In practice, the actual level of security depends more on configuration, maintenance, and update discipline than on the protocol's built-in capabilities (Bumiller et al., 2010; Sajadi et al., 2019; Wang et al., 2022).

To include intrusion detection systems tailored to Modbus, MQTT, and GOOSE traffic; one-way gateways or data diodes to separate critical assets from internet-facing segments; and micro-segmentation aligned with ISA/IEC 62443. Together, these measures reduce both the likelihood and the impact of cyberattacks and strengthen the resilience of communication infrastructures in distributed autonomous power systems.

2.5 Digital Twins and Predictive Diagnostics

Digital twins (DTs) are real-time, virtual counterparts of plant subsystems that are able to anticipate future behaviour and reflect the evolving physical state. They could also ingest time-aligned system measurements to adapt and modify themselves . In ROC-supervised autonomous power plants, DTs underpin three core functions: Predictory support, where the DT estimates the variables for which there are not sensors; supervisory support for control, where the twin supplies fast predictions to assist set-point selection and constraint handling; and predictive diagnostics, where the twin tracks degradation, detects incipient faults, and informs maintenance and risk-aware operation. Industrial studies consistently argue that future power facilities will require a robust DT architecture to achieve high reliability, availability, and maintainability at lower lifecycle cost (Sleiti, 2022). Within such an architecture, hybrid DT approaches fuse physics-based structure with data-driven learning to predict performance degradation more accurately than either method alone (Hartmann et al., 2021).

From a diagnostics point of view DTs make it possible to keep continuous, watch over system behaviour, and give practical insight for detecting faults and tracing their root causes. This is achieved by maintaining a real-time reflection of the system's dynamic state (Pasupuleti, 2025). Examples from different industrial sectors show how this works in practice. In thermal power plants, for instance, turbine-rotor twins demonstrate abilities in virtual monitoring, physical-state tracking, and predicting abnormal operating conditions for critical assets (Li et al., 2022).

At a larger scale, plant-level decision-support systems that integrate DTs with machine learning have been proposed to forecast operational trends and deviations in engine-based thermal facilities. These systems improve operator awareness and allow timely actions before failures develop (Pasupuleti, 2025). Meanwhile, studies from the wind-energy sector point to a growing maturity in predictive DTs that merge physics-based, data-driven, and hybrid modelling methods to support operation and maintenance. Altogether, these developments show a convergence of modelling approaches that can serve different types of generation technologies (Sleiti et al., 2022).

Meeting ROC timing requirements demands execution-efficient twins. Model-order reduction (MOR) is widely recognised as a key enabling technology. By projecting high-fidelity models onto reduced spaces, MOR accelerates simulation while preserving predictive accuracy to a degree sufficient for online supervision and control. This makes applications feasible that would otherwise be constrained by computational cost (Hartmann et al., 2021). In practice, reduced-order or surrogate twins provide short-horizon predictions for supervisory control, while higher-fidelity counterparts operate asynchronously for validation and policy improvement, an arrangement consistent with the multi-tier execution used elsewhere in this thesis.

2.5.1 Digital Twin-as-Observer and Predictor for Model Predictive Control (MPC)

In supervisory control architectures, the digital twin can extend beyond its role as a simulation tool to act as a real-time observer for unmeasured variables. These both could serve as system-states estimator for Model Predictive Control (MPC). In this configuration, the twin continuously estimates both the system state vector $\mathbf{X_k}$ and key parameters of the physical plant, providing the optimiser with plant-consistent information at every control interval. This directly addresses one of the core challenges of MPC: maintaining alignment between the predictive model and the evolving, uncertain dynamics of the plant. y reducing model and plant mismatch, observer-based digital twins improve constraint handling and raise the accuracy of short-horizon predictions. Built on established MPC frameworks (Rawlings et al., 2020), digital twins act as real-time observers that rebuild hidden states and parameters when direct measurements are limited or unreliable. This improves predictive control under uncertain or changing plant conditions. Traditional estimation methods such as Kalman filtering and moving-horizon estimation offer structured ways to estimate internal states, but they are often limited by model mismatch, unmeasured dynamics, and sensitivity to noise.

Digital twins extend these methods by combining physics-based models with data-driven updates. This hybrid gray-box idea allows the model to stay aligned with live operational data. In this role, the twin works as a virtual sensor and at the same time adjusts its estimation accuracy when operating conditions change. Arulampalam et al. (2002) highlight the use of Bayesian filtering techniques, including EKF, UKF, and particle filters, for reconstructing nonlinear and noisy state variables. Zhang and Liu (2019) show that combining EKF with MHE improves robustness under multi-rate sampling and measurement noise. These studies confirm the role of digital twins as observers that provide MPC with steady and timely state estimates in uncertain environments.

The architectural role of digital twins as observers can also be seen in the wider system where they operate. Figure 2 shows a digital twin framework adapted from ABB (ABB, 2020). In it, simulation, information, and machine-learning models are linked with real-time data streams, enterprise systems, and the Asset Administration Shell (AAS). Within this structure, the twin works as a hub for state and parameter estimation. It uses inputs from devices, installed base records, and master data. The AAS provides standardised links between the twin and supervisory or external systems, keeping consistency across the asset lifecycle.

This setup underlines the observer function. By combining different model views with operational data, the twin gives MPC accurate and timely estimates of unmeasured states and slowly changing parameters. In ROC-managed systems, this ensures that observer-based control is built into a scalable and interoperable environment rather than developed in isolation. Even with these advances, challenges remain for ROC-focused systems. Twin observers need to merge data streams that arrive at different rates, which calls for reliable multi-rate sensor fusion. Their performance must also be tested under real edge-fog latency, since communication delays can affect how quickly estimation results are delivered. Handling these issues is important for using observer-based twins in autonomous power systems, where supervisory MPC depends on fast, accurate, and consistent state reconstruction.

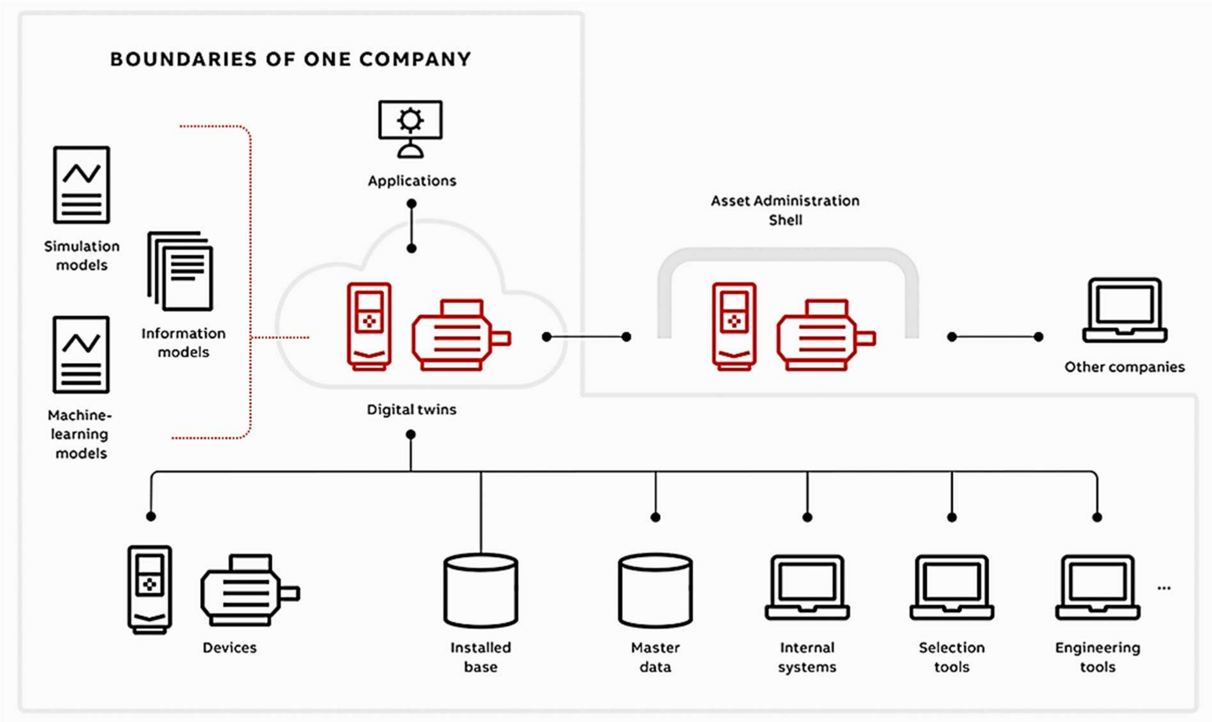

Figure 2. Digital twin framework integrating multi-model data through the Asset Administration; The digital twin: From hype to reality (ABB, 2020)

2.5.2 Twin-in-the-Loop Model Predictive Control

Digital twins can also be integrated directly into the optimisation loop of Model Predictive Control (MPC), where they act as surrogate models for rapid trajectory prediction. This twin-in-the-loop approach allows the MPC to evaluate candidate control actions across the full prediction horizon without relying on computationally intensive simulation at each step.

A practical implementation of this concept is the Time Series Dense Encoder (TiDE) neural network, trained offline as a multivariate surrogate for plant dynamics. Once embedded in the control loop, the TiDE model produces complete state trajectories in a single forward pass, supporting fast multi-step forecasting. In a Directed Energy Deposition (DED) application, this setup improved temperature regulation, reduced signal oscillations, and lowered defect rates when compared to conventional PID controllers, while still maintaining the required geometric and quality specifications (Chen et al., 2025; Das et al., 2023).

Comparable outcomes have been reported in related process industries. Karkaria et al. (2024) used LSTM-based models to predict temperature in real time during additive manufacturing, enabling dynamic adjustment of laser power. Similarly, Karkaria et al. (2024) developed surrogate twins based on multi-physical simulation data and Neural ODEs, achieving faster-than-real-time forecasting suitable for MPC embedding (Karkaria et al., 2024). These examples show that surrogate models can reduce computational load while preserving the level of prediction accuracy required for supervisory control.
However, the use of surrogate twins presents several challenges. Prediction accuracy may deteriorate under conditions not represented in the training data. To maintain model relevance, ongoing data collection and periodic retraining or online adaptation are required, especially in systems with non-stationary dynamics.

When applied in this configuration, twin-in-the-loop architectures support predictive optimisation within the time constraints typical of ROC-managed environments. By embedding efficient data-driven models such as TiDE, LSTM networks, or Neural ODEs into the MPC framework, it becomes possible to execute real-time control decisions without compromising on prediction depth or constraint handling.

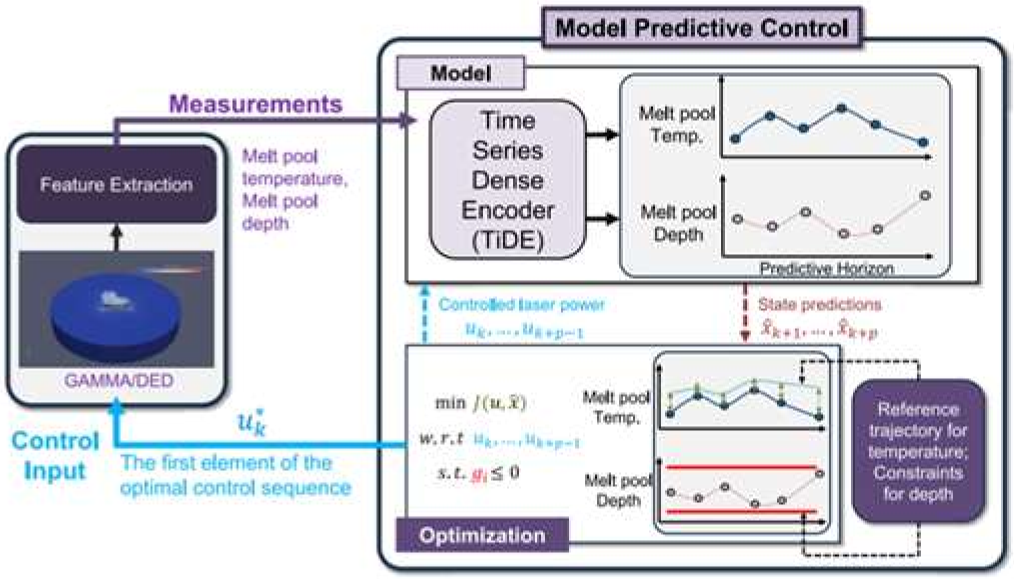

Figure 3. System overview of Twin-in-the-Loop MPC framework using a Time Series Dense Encoder (TiDE) as a digital twin surrogate for Directed Energy Deposition (Chen et al., 2025; Das et al., 2023).

2.5.3 Hybrid Health Modeling Twin for MPC Constraint Management

The raising of distribution of autonomous power systems calls for control strategies that combine decentralisation and fault tolerance to maintain stable and reliable operation. Relying too much on centralised command structures brings several points of vulnerability, with communication delays, hardware problems or cyber incidents can easily interrupt system operation and cause wider disturbances, to limit these risks and newer control architectures move toward decentralised layouts that include local fault-handling and recovery functions. This setup makes it possible for the system to keep running even when some parts fail or drop offline.

In a decentralised arrangement each subsystem acts on its own local measurements while exchanging limited coordination data through distributed communication layers. With this structure enhances the responsiveness and scalability, and is particularly valuable in systems with a high share of renewable generation. Dörfler et al. (2019) demonstrated that decentralised coordination could support overall stability under low-inertia and variable-demand conditions. The Resilient Information Architecture Platform for Smart Grid (RIAPS) applies these principles in practice, providing runtime monitoring, event management, and distributed fault detection across nodes (Ghosh et al., 2020). Such implementations illustrate how distributed autonomy strengthens resilience in modular energy environments.

The Fault-tolerant control (FTC) complements decentralisation by monitoring and maintaining the acceptable performance in the presence of component or communication faults. Hao et al. (2011) demonstrated that adaptive reconfiguration in permanent-magnet motor drives can preserve system stability following fault detection. Consensus-based designs extend this idea further. Han (2023) showed that generator groups can stay synchronised using wavelet-based signal analysis and local consensus, which reduces dependency on a central coordinator. The results underline the importance of embedding redundancy and to adaptive recovery within the overall control architecture.

At the system level, decentralisation and fault tolerance must also be supported through middleware and supervisory functions. Integrating diagnostic and recovery routines directly within the communication stack allows operation to continue even when upper layers are degraded. RIAPS provides such functionality by supporting modular inclusion of anomaly-detection and machine-learning components within local nodes (Ghosh et al., 2019). This is particularly important in ROC-managed environments, where situational awareness and operational continuity depend on resilient performance across all layers of the system.

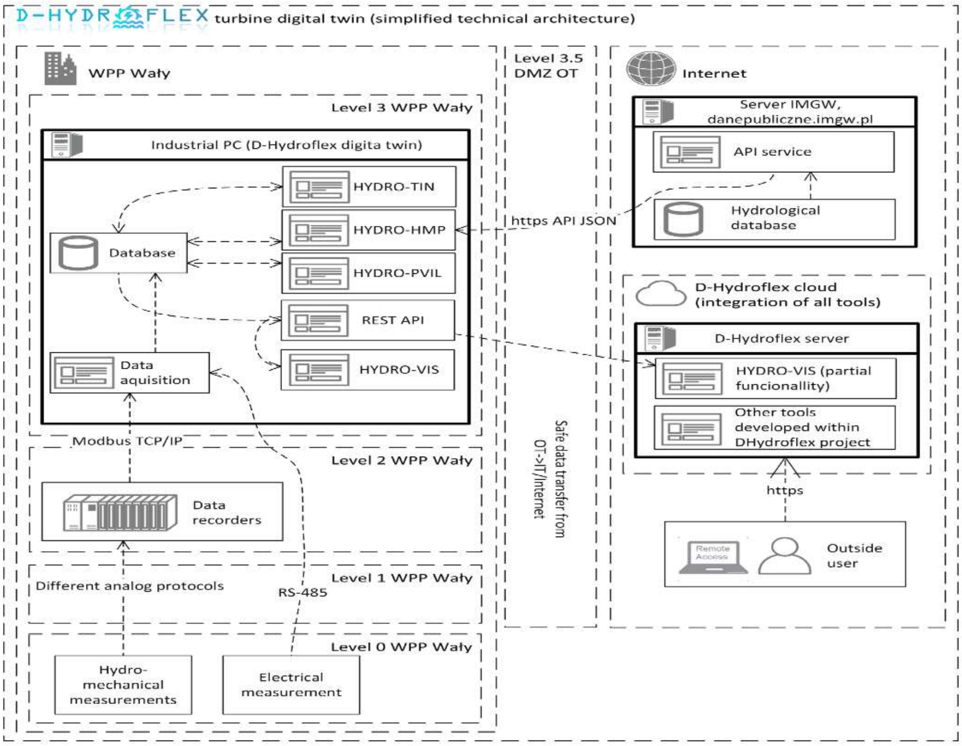

Figure 4. Simplified technical architecture of the D-Hydroflex turbine digital twin at Wały Śląskie Hydroelectric Power Plant (Machalski et al., 2025).

2.5.4 Reduced-Order Real-Time Twin for MPC Feasibility

One of the key challenges in using digital twins for real-time supervisory control is managing their computational requirements. While high-fidelity models can simulate system behaviour with high accuracy, they are typically too slow to meet the time constraints of predictive control. To address this, reduced-order modelling (ROM) techniques are used to simplify complex models into more efficient forms that still capture the dominant dynamics of the system.

Several established methods are used to generate ROMs. Proper Orthogonal Decomposition (POD) identifies the most significant patterns in the system response to create a simplified model. The Iterative Rational Krylov Algorithm (IRKA) maintains frequency-domain accuracy while reducing the system order. Balanced state reduction techniques further condense the model by retaining only those states that contribute most to control performance. These approaches enable reduced models that can operate faster than real time (RTF < 1), allowing predictive controllers to simulate multiple future trajectories within the available execution window (Lorenzetti et al., 2020; Hartmann et al., 2021).

When integrated into Model Predictive Control (MPC), ROMs improve constraint handling and prediction capabilities while keeping computation times within acceptable limits. However, reduced models may lose accuracy when the system enters operating regions that were not well represented in the training data. To mitigate this issue, hybrid approaches combine reduced-order models (ROMs) with anomaly detection and fallback procedures, when the ROM output deviates beyond a defined threshold, the system can revert to a full-order model or initiate a re-identification process until the reduced model is once again synchronised with real system behaviour. This concept has already been demonstrated in industrial applications. Siemens` Simcenter platform, for instance, employs executable digital twins (xDTs) that integrate ROMs directly into real-time control environments (Siemens, 2023). The presented models operate with limited sensor inputs and provide simulation along with diagnostic and optimisation capabilities in real time. In one reported case, a reduced-order neural network model was applied to predict air-conditioning performance within a vehicle cabin. The model achieved rapid computation while maintaining the necessary predictive accuracy. These examples indicate that ROM-based digital twins are not limited to research studies but are practical tools for operational control settings as well. Within the context of the Remote Operation Centre (ROC), this layered modelling strategy holds clear practical importance. Full-order twins can be located in the cloud, where they support offline diagnostics, AI model training, and fleet-level performance analysis. Reduced-order models, on the other hand, are implemented at the edge and fog layers to provide fast and deterministic control responses ROM-based digital twins as a key element within the ROC’s supervisory framework, which is described further in Chapter 4.

3 Laboratory Infrastructure and Suitability for ROC Development

3.1 Engine and Auxiliary Systems: EPS Laboratory Platform

The Efficient Powertrain Solutions (EPS) Laboratory at the University of Vaasa provides a dedicated environment for research on combustion and hybrid powertrain systems under realistic operating conditions, at the centre of the facility is a medium-speed Wärtsilä 4L20 multifuel engine, capable of producing approximately 800 kW of mechanical power at maximum load. The engine is connected to a 1200 kW ABB generator through a frequency converter and a 2 MVA transformer rated at 690 V/20 kV, which allows to both grid-connected and islanded modes of operation. This setup offers the flexibility to investigate system behaviour under various load profiles and grid-interaction scenarios. The main technical specifications of the Wärtsilä 4L20 engine are listed in Table 5.

Table 5. Wärtsilä 4L20 engine specifications (Valkjärvi, 2022).

Engine	Wärtsilä 4L20
Cylinder configuration	4-cylinder, inline
Bore diameter [mm]	200
Stroke [mm]	280
Connecting rod [mm]	510
Displacement per cylinder [dm^3]	8.8
Compression ratio	15.8:1
Number of valves per cylinder	4

The laboratory has auxiliary subsystems for fuel delivery, cooling, lubrication, and emission control. These make it possible to carry out experiments with different fuels and combustion strategies. Hybridisation studies are supported by a battery energy storage system (BESS) that can run together with the engine or on its own, allowing both transient and steady-state operation. This setup gives a way to study load sharing, thermal efficiency improvement, and emission reduction. Wide sensor network forms the base of the laboratory`s measurement system, it includes thermocouples, pressure transducers, flow meters, emission analysers and vibration sensors, that they provide detailed data on combustion, thermal, and mechanical behaviour. Data are collected using modular Dewesoft data acquisition (DAQ) units, which support high sampling rates and crank-angle-resolved acquisition. Deterministic real-time control is achieved with a Speedgoat rapid control prototyping system, while supervisory monitoring is carried out through a SCADA environment. Together, these tools allow simultaneous low-latency feedback control and higher-level monitoring, ensuring coherence between experimental data and supervisory analytics.

The laboratory currently supports both manual and semi-automated operation. Its modular instrumentation and flexible control structure, however, make it well suited for extending into advanced communication technologies, cloud-based analytics, and AI-driven supervisory modules. In this way, the EPS laboratory provides a robust foundation for developing and validating Remote Operation Centre (ROC) functions in the study of autonomous power plants.

3.2 Current Control and Automation Topology

EPS Laboratory at the University of Vaasa, there a layered control and automation setup is used to manage both manual and automated operation of a hybrid power platform built around the Wärtsilä 4L20 medium-speed engine. The system is built for real-time testing, development of control strategies and system-level diagnostics in hybrid energy applications, this section describes the setup through its main parts, including control structure, data acquisition, communication protocols, supervisory automation and digital twin integration..

3.2.1 Control Hierarchy and System Components

The automation system in EPS Laboratory is organised around two main platforms, consisting of a Speedgoat real-time target machine and a network of programmable logic controllers (PLCs). This structure follows the established design principles of the Vaasa Energy Business Innovation Centre (VEBIC), where experimental flexibility is combined with a strong focus on safety management (Mikulski et al., 2025).

Speedgoat platform functions is the computational core of the setup supporting both model-in-the-loop (MiL) and hardware-in-the-loop (HiL) simulations, through full integration with MATLAB/Simulink, that what enables the implementation of control-oriented models for engine and hybrid powertrain studies. As demonstrated by McBreen (2024), hard real-time operation was achieved on a Speedgoat Mobile target machine at sample rate equivalents of 14 crank-angle degrees (CAD) for the standalone engine model and 18 CAD for the combined electrical and engine configuration, then can be the performance levels confirm that the platform is capable of running computationally intensive simulations under real-time constraints, while maintaining sufficient fidelity for control and system-level testing 1000 RPM. These sampling intervals are providing a sufficient precision for crank-angle-resolved analysis are routinely applied in EPS experiments. The PLC network complements this by managing auxiliary subsystems and carrying out safety-critical functions. These tasks include control of lubrication, cooling, and fuel delivery, as well as the execution of emergency shutdown routines. Valkjärvi (2022) notes that the PLCs communicate with external equipment through Modbus TCP/IP alongside standard analogue and digital I/O, ensuring deterministic operation of protection functions independently of experimental control processes.

A centralised supervisory framework integrates the two platforms, linking Speedgoat-based simulation and real-time control with the process and safety routines implemented by the PLCs. This architecture provides a stable basis for experimental work while also positioning EPS and VEBIC as national research platforms supporting the development of hybrid propulsion technologies and clean energy systems (Mikulski et al., 2025).

3.2.2 Data Acquisition and Synchronization

EPS Laboratory uses a dual-path for data acquisition to setup that collects both high-frequency combustion data and lower-frequency process signals. High-frequency measurements are taken with a Dewesoft SIRIUS system which records crank-angle resolved signals at 0.1 CAD resolution by using a Kistler encoder. The sampling rate can reach 100 kHz. This setup makes it possible to analyse combustion in detail by providing in-cylinder pressure traces, ignition timing, and vibration data (Hautala et al., 2022; Valkjärvi, 2022). Lower-frequency data includes a temperature values with emission levels, and supervisory feedback, these are sampled at 6 CAD with the Speedgoat real-time target system and CANape software, then the signals give information about thermal behaviour and control dynamics which are used for model validation and controller development (McBreen, 2024). Synchronisation between the two data paths is done through crank-angle referencing or real-time clock alignment so that results from combustion, thermal, and emission domains can be compared.

Data are stored on dedicated PCs during measurements, and backup copies are kept on separate hard drives in the laboratory. Version control for datasets and model setups is handled through a BitBucket repository, which helps maintain traceability and consistent data management between experiments.

3.2.3 Automation, and Communication Architecture and SCADA Integration

The communication and supervisory automation system at the EPS Laboratory has been developed to support both modern and legacy equipment in one hybrid testing environment. The main design idea has been to keep data transfer deterministic for control-critical functions while still allowing flexibility for experimental instruments and remote operation, at the device level, the laboratory combines conventional analogue interfaces (4-20 mA and 0-10 V) analogue signal to serial bus converters RS-232 and RS-485 serial connections. This mix provides compatibility with the wide range of analysers and precision instruments that form part of the test cell instrumentation (Valkjärvi, 2022; Hautala et al., 2022). At the controller level, data exchange relies primarily on Modbus TCP/IP and CAN bus. Both have been validated in previous VEBIC campaigns, where they were shown to provide deterministic and reliable communication between the Speedgoat real-time target, the PLC network, and auxiliary subsystems under real-time constraints (Valkjärvi, 2022). Supervisory functionality is provided by a SCADA system, which aggregates process data from the engine, generator, and battery subsystems into a single human–machine interface (HMI). This environment allows continuous monitoring of process variables and integrates alarm handling, fault logging, and protective interlocks. Automated test procedures can be executed through parameter-based scripting, while embedded PID controllers regulate essential process parameters such as fuel pressure and coolant temperature (Hautala et al., 2022).

In addition, the laboratory supports remote access through ABB software and CANape tools. These enable parameter tuning, control adjustments, and diagnostic activities to be carried out from outside the test cell. Such capability reflects the broader VEBIC vision of developing remote operation and hybrid system validation, positioning the EPS laboratory as a national platform for advancing clean propulsion research (Mikulski et al., 2025).

3.2.4 Integration with Digital Twin and Simulation Frameworks

The EPS laboratory enables the development and validation of digital twin models that are directly connected to the experimental infrastructure. These twins are implemented as one-dimensional flow and system-level models of the Wärtsilä 4L20 engine and its

hybrid subsystems in MATLAB/Simulink, and are deployed to the Speedgoat real-time target for hardware-in-the-loop (HiL) operation (McBreen, 2024).
In HiL experiments, the digital twin communicates in real time with sensors and actuators in the testbed, linking simulated states with measured signals. McBreen (2024) was able to meet hard real-time requirements on a Speedgoat Mobile target machine at sample rate equivalents for 14 CAD for an engine model and 18 CAD for a combined electrical and engine model, both at 1000 RPM. Comparison with measured cylinder pressure traces confirmed that these resolutions capture combustion pressure rise and transient load changes with sufficient accuracy. These results show that crank-angle resolved control strategies can be evaluated under realistic conditions in the EPS laboratory.

The work further demonstrates that physics-based one-dimensional models can be applied in real-time simulation. As described by McBreen (2024), this bridges the gap between detailed offline simulations and models that are computationally feasible for MiL and HiL applications. Earlier studies at VEBIC reached similar conclusions, showing that reduced-order models calibrated against experimental data can provide predictive capability for Reactivity-Controlled Compression Ignition (RCCI) control research (Hautala et al., 2022; Valkjärvi, 2022).

Through the integration of experimental systems and real-time simulation, the EPS Laboratory offers a complete platform for controller validation and digital twin-assisted studies. According to Mikulski et al. (2025) the facility plays an important role in research on hybridisation, supervisory control intelligence, and the development of Remote Operation Centre (ROC) functions

3.3 Data Lifecycle and Workflow in the EPS Laboratory

The EPS Laboratory at the University of Vaasa has a data management system that ensures reliable handling of experimental and control signals, the system is used for real-time control validation, measurement analysis, and for developing automated diagnostics for hybrid power systems. Data management follows a clear sequence that starts with signal an acquisition and preprocessing and continues to storage and analysis. During acquisition, then sensor and control data are checked for signal quality and calibrated according to predefined parameters. The processed data are stored on a central server, where each measurement is organised by test campaign and date. This arrangement makes it possible to compare results between experiments and provides traceable documentation for all laboratory measurements.

3.3.1 Data Acquisition and Sampling Strategy

The data collection setup in the EPS Laboratory is divided into two streams. One handles high-speed signals, and the other focuses on slower, system-wide measurements. This separation makes it possible to capture both the fast combustion events and the slower processes happening in the hybrid system. The high-speed side are collecting signals in-cylinder pressure, crankshaft angle, injector current and valve timing. Records at 0.1 crank angle degree (CAD) resolution by using encoders and piezoelectric sensors, to having this level of detail is important for studying advanced combustion modes like Reactivity-Controlled Compression Ignition (RCCI) and dual-fuel operation (McBreen, 2024). The lower-speed stream records values like temperature, exhaust gas concentration, intake air conditions, and battery system feedback. Sampling rates of up to 6 CAD are used, which is sufficient for system monitoring, diagnostics, and energy flow analysis (Valkjärvi, 2022).

Both streams can keep in sync in the system relies on hardware-based timing, signals that depending on crank angle are matched with encoder pulses, while in time-based signals follow the shared system clock. This alignment allows data from both streams to be analysed together and helps in identifying short-term changes and verifying control performance.

3.3.2 Calibration and Preprocessing Pipeline

Calibration is an essential part of making sure that all measurements are accurate, and can be repeated, before each test campaign and the laboratory performs systematic calibration routines for all sensors and actuators involved. For instance, the Gas Valve Unit (GVU) runs automated sequences that check for leaks and carry out purge cycles. These actions confirm that the valves are working properly and help maintain operational safety. Instruments such as the Fourier-transform infrared (FTIR) spectrometer and the chemiluminescence detector (CLD) are calibrated using zero and span checks with certified reference gases (Valkjärvi, 2022). Analog input channels on the Dewesoft modules and the Speedgoat systems are verified with precise voltage standards that are available in the lab.

After calibration the raw data goes through preprocessing, this includes normalising signals, aligning units, and resampling when necessary. MATLAB scripts and CANape preprocessors are used to convert the collected measurements into structured formats. These steps make sure that the next stages, whether in control development or machine-learning analysis, are based on reliable and clean data (McBreen, 2024).

3.3.3 Data Storage and Retention Architecture

The EPS Laboratory uses a layered data storage approach to keep experimental results accurate and easy to access, during each test run, measurement data are first stored on local computers that are directly connected to the acquisition systems. DewesoftX is used for crank-angle, resolved combustion signals, CANape for controller interfacing and calibration, and Calcmet for exhaust gas analysis, to maintain consistency between measurements, high- and low-frequency data streams are aligned in time using either crank-angle referencing or shared system clocks (Valkjärvi, 2022; Hautala et al., 2022).

After each session, the recorded files are copied to external drives that act as offline backups. This redundancy protects against possible hardware faults or operator mistakes. Before the data are archived, basic validation checks are done to make sure the files are complete and correctly transferred.

For long-term storage and shared work, some datasets and post-processing scripts are uploaded to a BitBucket repository that is managed together with Wärtsilä's R&D team. The repository keeps version control for files and scripts and makes it easier to develop models and study hybrid systems in a shared research setting.
With the mix of local storage, backup drives, and version-controlled repositories, the EPS Laboratory keeps a reliable way to manage data. This setup protects experimental results and supports ongoing work in digital twin validation and ROC development.

3.3.4 Post-Processing and Feature Extraction

Post-processing routines in the EPS Laboratory turn raw measurement signals into useful features that can be used in diagnostics, control tuning and model development. In the combustion analysis, cycle-based metrics such as Indicated Mean Effective Pressure (IMEP), the crank angle at 50% heat release (CA50), and the Rate of Heat Release (ROHR) are calculated from in-cylinder pressure data. These values are taken by using the Dewesoft combustion analysis module, where baseline corrections and phasing are referenced to crank-angle measurements to keep results consistent under different operating conditions (Valkjärvi, 2022).

Vibration signals are analysed in the frequency domain an most often using Fast Fourier Transform (FFT) methods, to detect spectral shifts that show how the rotating parts behave. This approach has been used in EPS and VEBIC test campaigns as part of condition monitoring and mechanical fault detection work (Hautala et al., 2022; Valkjärvi, 2022), the mission data are corrected for time delay and filtered using rolling averages to make the signals stable for supervision and compliance studies. These corrections also allow the data to be aligned with combustion results when doing combined analysis (Valkjärvi, 2022).

For data-driven studies, curated datasets are prepared to support machine learning applications. Cleaned and labelled signals have been used in fault classification, thermal limit estimation, and digital twin validation. Data preparation routines are implemented in Python (Pandas, SciPy) and MATLAB, with workflows developed in collaboration with Wärtsilä's R&D teams to support joint experimental and modelling tasks (McBreen, 2024).

3.3.5 Remote Access and Workflow Automation

To support activities similar to those in Remote Operation Centres, the EPS Laboratory includes tools for remote system access and test automation. These features are designed to enable around-the-clock operation and fast diagnostic feedback. Secure VPN connections provide external access to Speedgoat real-time units and ABB PLCs, allowing parameter tuning, system diagnostics, and monitoring tasks to be performed without the need for on-site presence (Valkjärvi, 2022). This capability extends the flexibility of the testbed and supports collaborative use across research teams. Test operations in the EPS Laboratory are managed through automation scripts written in MATLAB and Python. These scripts handle setpoint adjustments, check safety conditions, record measurements, and log any abnormal events, all routines are stored in a version-controlled repository, which allows collaborative development and ensures full traceability of every experimental campaign (Hautala et al., 2022; McBreen, 2024). The collected data are also available through structured APIs that follow the interoperability guidelines from the National Institute of Standards and Technology (NIST, 2017). These interfaces make it possible to connect directly with co-simulation tools, digital twin models, and larger data platforms.

The remote and automated features increase the EPS Laboratory’s ability to run unattended test campaigns, carry out remote troubleshooting, and perform model-based optimisation. In this way, the setup supports the wider goals of VEBIC as a national platform for digital twin development and remote operation research (Mikulski et al., 2025). Overall, the EPS Laboratory operates in a semi-automated mode with supervisory monitoring and remote access. This configuration provides a dependable foundation for experimental work, although the system still relies on operator supervision during active testing. In its current form, the facility can be characterised as pre-autonomous, offering a foundation for future development toward ROC-ready operation.

4 ROC Vision

4.1 Functional Requirements of the Proposed ROC

ROC is expected to provide the supervisory functions that make autonomous plant operation safe and dependable. These functions can be divided into three main groups of requirements. The first group relates to supervisory intelligence. In this area, Model Predictive Control (MPC) should be used together with Reinforcement Learning (RL) and supported by a digital twin, together, with these tools allow optimisation tasks to be carried out while maintaining adaptability and awareness of the system's condition. The second group relates to communication. In order to avoid interruptions or inconsistencies, data exchange has to be secured with time-synchronised and interoperable standards. Examples of such protocols include PTP, NTP, OPC UA, MQTT, and IEC 61850. The third group focuses on the human operator. Clear and explainable outputs are needed to maintain trust in the system, predictive alarm handling is required to reduce unnecessary reactions, and authority transfer should follow a staged procedure rather than an abrupt change. These requirements, taken as a whole, provide the framework for the control hierarchy, the artificial intelligence modules, and the supporting data architecture presented in the next sections.

4.1.1 Supervisory Intelligence (MPC, RL, DT)

The supervisory function of the ROC needs to go beyond traditional SCADA-type monitoring and deterministic control. It must include advanced control intelligence that can handle constraints, adapt control policies, and predict system behaviour. Model Predictive Control (MPC) provides the basis for supervisory optimisation because it allows multi-variable, horizon-based optimisation problems to be formulated while taking process limits into account and reducing operational costs or other objectives (Qin & Badgwell, 2003). Industrial applications have already shown that MPC can be scaled to complex systems and used in a wide range of process environments, for complex energy systems where efficiency, emissions, and equipment degradation must be balanced (Rawlings et al., 2020). EPS laboratory context, the MPC provides a structured means of coordinating engine, converter, and battery energy storage interactions at supervisory timescales (0.5-5 s), complementing the deterministic 1–10 ms loops executed by Speedgoat RT controllers and PLCs described in Chapter 3.

While MPC ensures constraint satisfaction, it remains limited in handling unmodeled dynamics and rapidly changing operational regimes. Reinforcement Learning (RL) provides an adaptive element by allowing the ROC to learn control policies in conditions of uncertainty and partial observability, as noted in control-oriented research on safe RL (Yu et al., 2024). The ability of RL to build, an adaptive supervisory strategies has been improved through its recent integration, by using nonlinear MPC and digital twin environments, these combinations make it possible to handle the constraints under uncertainty in more flexible way (Zhang et al., 2021; Yu et al., 2024). In this hybrid setup, RL produces candidate control policies that are then checked and refined through MPC to ensure safety while keeping adaptability.

Digital twin (DT) has included to complete this supervisory structure. The twin provides a continuously updated model that helps with state estimation, surrogate prediction, and system health tracking. Digital twins have been shown to improve both observability and prediction range, especially when trained together with RL agents (Hartmann et al., 2021; Lorenzetti et al., 2020). While in ROC, the DT is expected to act as an observer for unmeasured states such as combustion dynamics, a surrogate model for fast MPC optimisation, and a prognostic tool for including Remaining Useful Life (RUL) estimates in supervisory decisions. Together, MPC, RL, and DT form the base of supervisory intelligence for autonomous ROC operation, creating a balance between predictability, adaptability, and health awareness that conventional supervisory systems lack.

4.1.2 Reliable and Secure Communication (PTP/NTP, OPC UA, MQTT)

The effectiveness of supervisory intelligence in the ROC depends on reliable, time-synchronized, and secure communication across all plant layers, PTP standardized in IEEE 1588 and IEC C37.238 can enable sub-microsecond synchronization of distributed measurements and has been validated in both microgrid and remote operation testbeds (Eidson, 2006; ABB, 2015; Zhang & Liu, 2019). With regard to slower supervisory data streams, Network Time Protocol (NTP) gives enough accuracy to keep operator logs, alarms, and decision records consistent. In the EPS Laboratory, where crank-angle-resolved engine data are combined with low-frequency SCADA signals, the use of both PTP and NTP is needed to maintain time coherence between local control functions and supervisory analytics, as described in Chapter 3. While the Communication protocols also have to support interoperability and scalability. OPC Unified Architecture (OPC UA) provides a structured, model-based way to represent process data, which allows the ROC to connect different subsystems under one supervisory layer (Wang et al., 2022). At the same time, MQTT offers a lightweight publish–subscribe setup that works well for event-driven messages such as alarms or state changes, and it has been shown to fit systems with limited device resources (Gupta et al., 2023). In industry, the two are often used together: OPC UA keeps the meaning and structure of the data clear, while MQTT makes event transmission fast and scalable.

At present, the EPS Laboratory mainly uses Dewesoft DAQ and Modbus/CAN connections. Moving toward ROC readiness, the system will need to include OPC UA and MQTT with PTP and NTP time synchronisation. This integration will make data flows between local controllers and supervisory AI modules both consistent and secure.

4.1.3 Operator Support (Explainability, Alarms, Trust in Automation)

Higher levels of autonomy that are added to ROC operations did not cover the needs to Human interference, while still the human operators play an essential role in supervisory decision-making, studies in human factors show that in systems with high automation, operators can become out of the loop, which often leads to lower alertness and slower or less effective reactions when abnormal events occur (Mooneyham & Schooler, 2013). This issue is especially relevant to remote or autonomous systems, where situational awareness is already limited and recovery actions are harder to perform.

For this reason, ROC functions is needed to include explainability features in their supervisory AI outputs, the recommendations produced by reinforcement learning and model predictive control should be shown with clear and understandable explanations, for example by identifying the main input features or the reasoning that leads to a certain action, while this helps operators follow how the system reaches its conclusions and verify that its proposals make sense. Alarm handling also needs to move beyond the simple threshold-based systems still used in many EPS operations. Predictive and prioritised alarms can reduce unnecessary alerts, make information clearer, and help operators keep focus during demanding conditions.

An Experience from industrial remote operation shows that the goal should be to support operators, not replace them. A staged transfer of authority, moving from Observe to Advise, then to Constrained Execute, and finally to Autonomous Execute, gives a controlled way to shift responsibility between human operators and the automation system (DNV, 2022; ABB, 2024). This gradual approach builds trust and helps keep human oversight active within the supervisory process.

.

4.2 Control Hierarchy and Execution Path

The proposed control architecture for the Remote Operation Centre (ROC) is based on the well-established concepts of hierarchical and layered control. These principles have been used for many years in critical infrastructure systems, where multi-level control frameworks have shown to be effective in managing complexity and ensuring reliability.

Examples of such applications can be found in microgrids (Luo et al., 2023; Di Silvestre et al., 2018), in large-scale power networks, and in supervisory systems for advanced nuclear plants. (Rehtanz, 2003; Dörfler et al., 2019). In each of these cases separating control functions by timescale, scope, and authority has led to better scalability, high resilience and safer operation (Liu et al., 2022; Rawlings et al., 2020; Vidal, V. F., et al. 2022).

The ROC applies the same logic to autonomous power generation, mapping functions across Edge, Fog, and Cloud computing layers. At the field level, the architecture covers the main physical assets of the plant: engines, generators, converters, battery energy storage, and auxiliary systems. These assets are equipped with sensors and actuators that provide the basis for real-time monitoring and actuation. Data acquisition here focuses on high-frequency signals such as crank-angle-resolved combustion parameters, the temperature processing, pressures, and electrical variables. These measurements are typically collected through dedicated acquisition hardware or through protection devices with embedded sensing functions. Devices in this layer carry out only basic preprocessing, for example filtering or timestamping. Deterministic control and higher-level decision-making are performed further up the hierarchy.

4.2.1 Local Control Layer

The local control layer, implemented on Speedgoat real-time targets and PLCs, executes deterministic feedback control with 1–10 ms cycles. These loops include conventional PID controllers, actuator PWM generation, and feedforward compensation. In hierarchical control taxonomies, this corresponds to the primary control layer (Luo et al., 2023; Qin, S. J., & Badgwell, T. A., 2003), ensuring local stability under disturbances. In the proposed ROC, local controllers enforce actuator-level commands derived from validated supervisory trajectories but remain independent of AI-driven decision-making.

4.2.2 Supervisory Layer: Hierarchical MPC and RL Integration

At the supervisory level, the ROC applies a hierarchical combination of Model Predictive Control (MPC) and Reinforcement Learning (RL). The design assumes that MPC is not a single, centralised function, but distributed across the Edge and Fog layers, with overall coordination maintained at the Cloud-based ROC. At the Edge, fast MPC units is running on embedded hardware that close to the plant equipment. These units enforce feasibility constraints and perform short-horizon optimisation at sub-second intervals, providing rapid disturbance rejection and ensuring compliance with operational limits (Rawlings et al., 2020; Dickler et al., 2021), while at the fog level, distributed MPC units coordinate several edge controllers that manage engines, converters, and battery energy storage systems. This layer handles the coupled dynamics and plant-wide constraints while keeping latency lower than cloud-based optimisation approaches (Liu et al., 2022), and at the highest level, the cloud-hosted ROC runs supervisory MPC with horizons between 0.5 and 5 seconds. The forecasts are happening at this layer, from the digital twin are combined with guidance from reinforcement learning to achieve plant-wide optimisation that balances efficiency, emissions, and component degradation, while also supporting long-term planning.

Reinforcement learning supports the control framework by adding adaptability when dealing with uncertainty, nonlinear dynamics, and this behaviour that the model does not describe, in recent studies have shown that using RL together with MPC in digital twin environments can improve supervisory control by combining learning capability with constraint handling (Zhang et al., 2021; Yu et al., 2024). In the ROC setup, RL agents are mostly located at the cloud level, where they produce adaptive strategies, and these strategies are then checked and limited by MPC at the fog and edge levels before being used in the system.

This layered MPC and RL structure follows the same principle of hierarchical supervisory control that has been used in nuclear systems (Rehtanz, 2003; Dörfler et al., 2019). In these systems, the supervisory intelligence increases adaptability and optimisation, while the lower layers keep deterministic functions to ensure safety and reliability during autonomous operation.

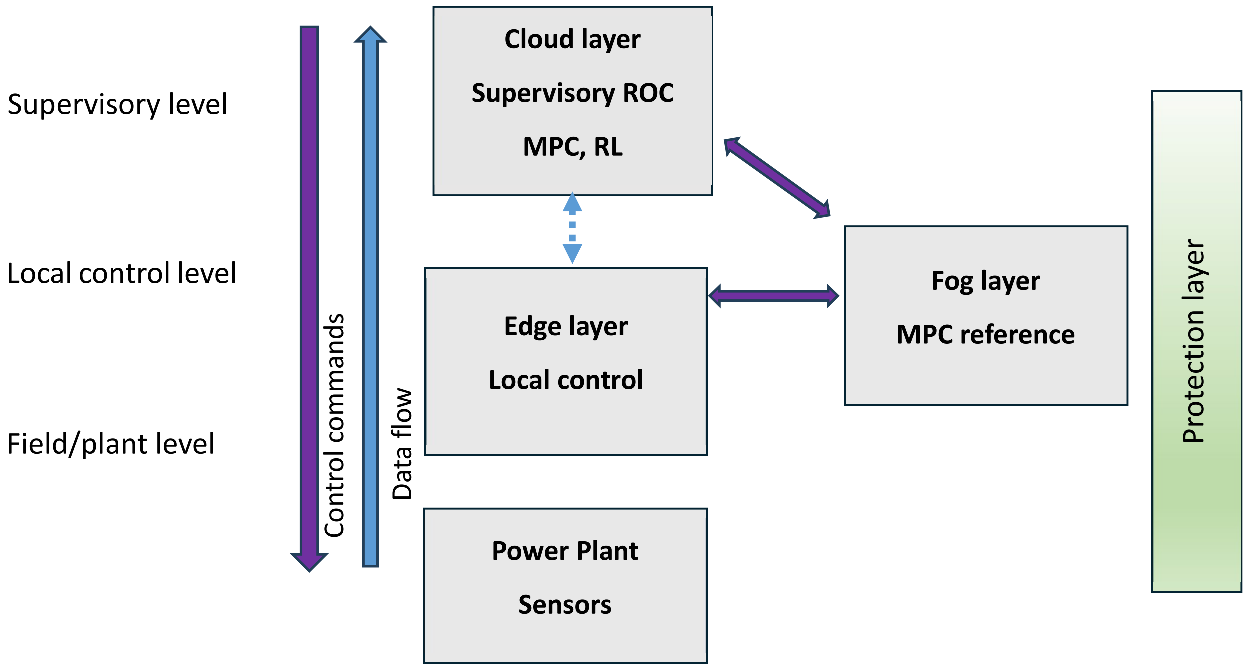

Figure 5. Control Hierarchy, Execution Path and Data Flow.

4.2.3 Protection and Safety Layer

In parallel with all control layers, the protection layer enforces hard safety constraints through Safety Integrity Level (SIL) rated relays, hardwired trip circuits, and IEC 61850-compliant Intelligent Electronic Devices (IEDs). This layer operates independently of ROC intelligence, guaranteeing sub-50 ms fault isolation. Its autonomy ensures that cascading failures or AI-driven anomalies cannot compromise plant integrity, consistent with cyber-physical safety frameworks for industrial systems (International Electrotechnical Commission, 2019; Hao et al., 2011).

4.2.4 Execution Path and Timing

The execution path between the layers follows a bottom-up flow for data and a top-down flow for control. Telemetry from field devices is time-stamped at the edge and preprocessed at the fog nodes before being sent to the ROC, an the supervisory MPC and RL modules at the Cloud level generate candidate setpoints, which are filtered through runtime assurance and command validation before dispatch to Fog- or Edge-level MPCs. Edge MPC then ensures real-time feasibility, while local controllers translate supervisory commands into actuator signals.

Protection systems operate independently by monitoring raw field variables directly and reacting with immediate action when predefined limits are reached. This layered design follows the defence-in-depth principle described in hierarchical grid control (Vidal et al., 2022) and in supervisory systems used for nuclear applications (Dörfler et al., 2019).

4.2.5 Computational Topology Alignment: Edge-Fog-Cloud Mapping

The control hierarchy follows a computational structure that spans the Edge, Fog, and Cloud layers. Starting with Edge computing that runs local and fast MPC on deterministic controllers placed close to the equipment, allowing sub-second enforcement of constraints. Fog computing manages distributed MPC coordination, buffering, and multi-sensor data fusion, which lowers latency and supports plant-level optimisation without depending on external networks (Vidal et al., 2022; Kumari et al., 2021). At the highest level, Cloud computing, represented by the ROC, hosts supervisory MPC, RL agents, and digital twin synchronisation for long-horizon optimisation and coordination across multiple sites (Dörfler et al., 2019; Shi & Dustdar, 2016; Vidal et al., 2022).

layered setup ensures that control tasks are matched with computing resources suited to their timescales and critical importance. It also provides scalability, since new plants or subsystems can be added through extra Fog or Edge nodes. In addition, it increases resilience, because Fog-level MPC can keep essential functions running even if the

Cloud connection is lost, by embedding AI modules within this topology, the ROC realizes a hierarchical, distributed supervisory architecture that balances adaptability, safety, and reliability.

Table 6. Control and Computational Layer Alignment

Control Layer	Timescale	Computing Substrate	Core Function	Literature Basis
Field	µs–ms	Sensors, IEDs	Data capture, actuation	Luo et al. (2023)
Local Control	1–10 ms	Edge (RT, PLCs)	Fast deterministic control	Qin, S. J., & Badgwell, T. A. (2003)
Fast MPC	<1 s	Edge (embedded MPC)	Constraint enforcement, short-horizon optimization	Dickler et al. (2021)
Distributed MPC	0.1–1 s	Fog nodes	Multi-asset coordination	Liu et al. (2022)
Supervisory MPC + RL	0.5–5 s	Cloud (ROC servers)	Plant-wide optimization, adaptive policies	Rehtanz (2003); Nature SR (2024)
Protection	<50 ms	Hardware relays, IEDs	Fault isolation, shutdown	International Electrotechnical Commission, 2019; Rehtanz, 2003

4.3 AI Modules in the ROC Supervisory Layer

The ROC supervisory layer realises adaptive, constraint-aware decision-making by combining Model Predictive Control (MPC) with Reinforcement Learning (RL) and a sensor-driven digital twin. The modules are deployed hierarchically across computational substrates to meet timing and reliability requirements: fast MPC at the Edge, distributed MPC at the Fog, and supervisory MPC and RL at the Cloud. This arrangement preserves real-time feasibility locally while enabling plant-wide optimisation and policy adaptation at the ROC. Runtime assurance governs all outward control actions to maintain safety and auditability across the hierarchy.

4.3.1 Model Predictive Control (MPC)

MPC constitutes the optimisation backbone of the supervisory layer. In accordance with receding-horizon principles (Rawlings et al., 2020), at each ROC tick the controller solves a finite-horizon programme that minimises an economic objective-fuel use, emissions and degradation proxies, subject to process and actuation constraints derived from the plant and the digital twin.

4.3.2 Reinforcement Learning (RL)

Reinforcement Learning (RL) complements Model Predictive Control (MPC) by adapting supervisory policies under uncertainty, regime changes, and model mismatch. In energy management applications, RL has shown performance gains by learning allocation and scheduling strategies directly from interaction data (Zhang et al., 2021). Recent studies demonstrate that combining deep RL with nonlinear MPC and digital-twin rollouts can produce adaptive controllers that respect operational constraints while improving disturbance response (Nature, 2024).

Within the proposed ROC architecture, RL is placed at the Cloud layer and is constrained by hierarchical MPC. The RL agent generates trajectories or setpoint adjustments, which are projected onto the admissible set defined by Edge and Fog layer MPC limits. These candidate actions are then checked by a runtime assurance mechanism before being dispatched. In this way, RL contributes adaptability while the overall system ensures that safety and feasibility are maintained at each stage of the hierarchy.
The input provided to the agent consists of estimates filtered through the digital twin, multi-rate features extracted from the archiv, and recent command histories. The agent's action space encodes supervisory setpoint increments or allocation ratios. Training follows an offline-online procedure: extensive offline pre-training using past-data and twin-

based simulations, followed by cautious online fine-tuning. During online adaptation, strict confidence bounds and drift monitors are applied to prevent unsafe learning behaviour. As a result, the Cloud-level RL agent remains aligned with the evolving plant dynamics while the deterministic guarantees of the lower control layers are preserved.

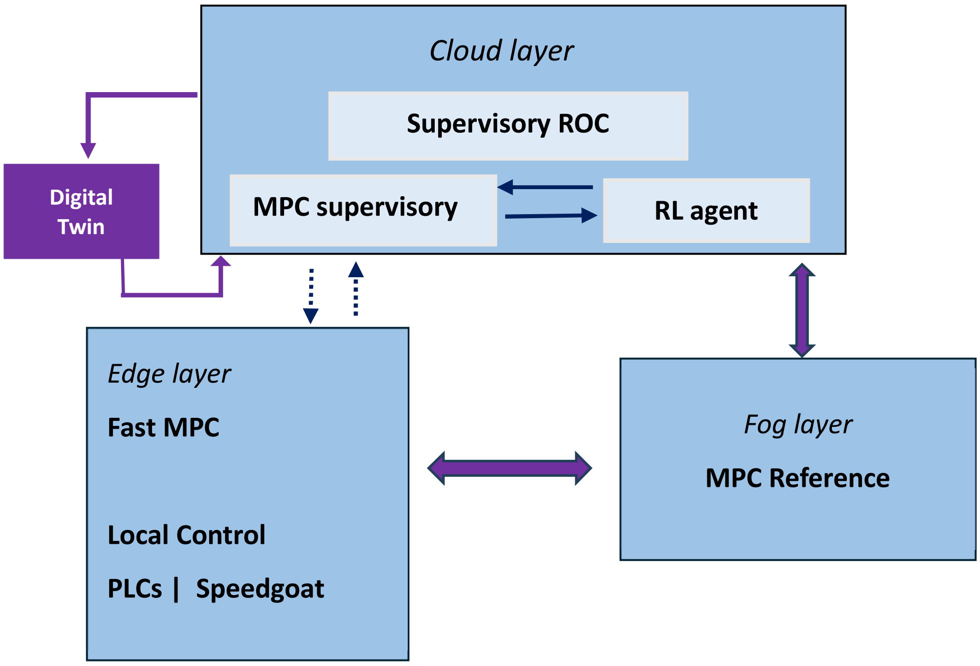

Figure 6. Hybrid integration MPC and RL modules across Edge, Fog, and Cloud layers in the ROC supervisory architecture.

4.3.3 Digital Twin Synchronization

The effectiveness of supervisory Model Predictive Control (MPC) and Reinforcement Learning (RL) in the ROC depends on the capability of the digital twin (DT) to provide accurate state estimates, surrogate predictions, and health-aware constraint adjustments in real time. To satisfy the computational demands of predictive control, the twin must operate at or faster than real-time execution while retaining sufficient fidelity to represent the dominant system dynamics. The one approach for reconciling these requirements is by using of Reduced-Order Models (ROMs), hence these models approximate the high-fidelity plant representations at significantly lower computational cost. Balanced State Reduction and Proper Orthogonal Decomposition (POD), and the Iterative Rational Krylov Algorithm (IRKA) that have been applied successfully to compress large-scale models while preserving the essential input-output behaviour and dominant frequency response (Lorenzetti et al., 2020; Hartmann et al., 2021). In practice, the computational fluid dynamics or thermodynamic models can be simplified into lower-dimensional spaces by using methods such as (POD) or Galerkin projection. The reduced-order models (ROMs) that result from this process provide predictive dynamics in smaller state spaces that are easier to handle in real time. This makes it possible to use longer optimisation horizons in supervisory MPC without exceeding the limits of real-time computation. In operation, the ROM-based digital twin serves two main purposes. The first is state and parameter estimation which is achieved by combining high-rate sensor data with the reduced-order dynamics through methods such as the Extended Kalman Filter (EKF) or by Moving Horizon Estimation (MHE) (Arulampalam et al., 2002). Unmeasured states that including combustion phasing or shaft torque, in real time. Second, the ROM serves as a surrogate prediction engine within the MPC solver, where fast rollouts of system trajectories are required to evaluate constraint satisfaction. Studies on twin-in-the-loop MPC have shown that such surrogates can maintain constraint handling performance while improving feasibility under real-time execution limits (Hartmann et al., 2021; Lorenzetti et al., 2020). Figure 7 summarises the synchronisation path used here: a POD-based FRM for the engine, IRKA-reduced electrical submodels, Simulink integration, and real-time/HIL coupling with the generator-converter chain, as demonstrated on the EPS/VEBIC platform (Söderäng, Hautala, Mikulski, Storm, & Niemi, 2022).

A challenge with reduced-order representations is the potential loss of accuracy when the physical plant operates outside the regimes represented in the reduced basis. Unusual transients or component-level faults can cause the ROM to diverge from the real system. To mitigate this, ROC implementations integrate anomaly detection and adaptive twin updating. If deviations exceed defined thresholds, the twin either re-projects onto an updated reduced basis or temporarily reverts to a high-fidelity model until recalibration is complete. This dynamic synchronisation method keeps the digital twin reliable as a predictive tool during both steady-state and transient operating conditions. This has been shown in recent studies on sensor-based, real-time synchronisation of digital twins (Chen et al., 2025).

Besides observability and prediction, the digital twin also includes prognostics and health management (PHM) features. Degradation models and Remaining Useful Life (RUL) estimators are built into the twin to support constraint handling. The results from these models are turned into dynamic operating limits, such as tighter thermal boundaries for converters or modified ramp-rate settings as components age. With this logic the digital twin supports real-time estimation and prediction while enabling condition-based supervisory optimisation. In this way it connects monitoring prediction and health management in one continuous process then the design targets specify that the observer shall operate at a sampling rate of ≥100 Hz with end-to-end latency <10 ms and a normalized root mean square error (NRMSE) ≤5% for in-cylinder pressure p_{cyl} and shaft torque (Söderäng et al., 2022; McBreen, 2024). The supervisory MPC shall use horizon $N = 20$and meet a solve time ≤50 ms at a 10 Hz control rate, with feasibility provided by ROMs (Rawlings et al., 2020; Hartmann et al., 2021; Lorenzetti et al., 2020).

4.3.4 Runtime Assurance and Command Validation

A central requirement for supervisory AI in the ROC is the preservation of operator situational awareness while avoiding out-of-the-loop (OOTL) degradation, which is a well-documented risk in highly automated systems (Mooneyham & Schooler, 2013). To mitigate this, alarm management strategies are designed around rationalisation and prioritisation rather than simple threshold-based triggering. In practice, alarms are derived from Model Predictive Control (MPC) residuals, digital twin constraint violations, and Reinforcement Learning (RL) confidence metrics. These are filtered to emphasise events of clear operational relevance. Evidence from industrial deployments shows that predictive and prioritised alarms can reduce operator overload and improve response quality (ABB, 2024).

Another key requirement is explainability of supervisory decisions. The operators must be able to follow the reasoning behind AI-generated recommendations, whether the cause is related to tighter constraints, forecast changes or health-based operating limits. That needs to fit with staged autonomy frameworks, where the level of authority moves step by step from Observe to Advise, then to Constrained Execute, and finally to Autonomous Execute (DNV, 2022). Giving clear and understandable reasons for each recommendation helps operators judge the system`s decisions and when needed, override supervisory actions with confidence, by combining structured alarm management with explainable supervisory outputs, the ROC keeps operators actively involved in decision-making. At the same time, this approach avoids excessive cognitive load, helping ensure that human supervision stays balanced and consistent with the intended level of autonomy.

4.3.5 AI-based Anomaly Detection and Prognostics

The ROC is responsible not only for optimisation and runtime assurance but also for the health of critical assets. To keep autonomous plants safe and reliable, the ROC needs AI-based anomaly detection and prognostics. These functions track equipment condition, anticipate failures, and support predictive maintenance. In this way, supervisory intelligence becomes asset-aware, cuts unplanned downtime, and improves availability. Reviews show that combining digital twins with AI improves both fault detection and prognostic accuracy compared with conventional monitoring (Sleiti et al., 2022; Machalski et al., 2025; Li et al., 2022). Plant-validated twins that run in real time and co-simulate with the electrical system provide the short-horizon predictions and residuals a ROC needs for surveillance (Söderäng, Hautala, Mikulski, Storm, & Niemi, 2022).

Anomaly detection should operate across the edge-fog-cloud hierarchy. At the edge, lightweight detectors compare live sensor signals to twin-predicted states, producing early alarms that fixed thresholds often miss. Engine-generator twins demonstrated on the EPS/VEBIC platform, with real-time execution on a Speedgoat target and I/O to plant equipment, show how such residual checks can run under tight timing constraints (Söderäng et al., 2022). Residual monitoring has also detected subtle anomalies in other real-time systems (Ghosh et al., 2020; Han, 2023). At the fog layer, detection benefits from multi-sensor fusion and cross-checks between several twins; comparing outputs across twins improves the identification of irregular behaviour in complex assets (Karkaria et al., 2024; Hartmann et al., 2021). At the cloud layer, more capable models are feasible; recent work indicates that sequence architectures such as sequence models can learn degradation patterns across fleets and support scalable fault identification under weak supervision (Das et al., 2023).

Prognostics and remaining useful life estimation forecast degradation and schedule maintenance before faults occur. Fog nodes can host local prognostic models for engines, batteries, and converters, while the ROC aggregates results from multiple sites to plan fleet maintenance. Case studies report DT-enabled prognostics with accuracy sufficient for deployment in energy equipment (Sleiti et al., 2022; Machalski et al., 2025). Data-driven methods, including LSTM autoencoders and sequence models, also improve availability and cost efficiency (Sleiti et al., 2022; Qiu et al., 2020).

The digital twin serves two roles. First, it is a reference model: continuous comparison between measured signals and twin predictions yields residuals that form robust anomaly indicators. Second, it generates data: by simulating fault scenarios, the twin provides the diverse labelled examples that real plants rarely produce. In the EPS/VEBIC context, a detailed 1D engine model was reduced to a fast-running model, imported to Simulink, and coupled with the generator, converter, DC link, and optional BESS for real-time co-simulation. The twin achieved validation accuracy suitable for supervisory residuals and short-horizon forecasts, supporting ROC needs (Söderäng et al., 2022). Executable twin workflows used in industry follow the same idea of embedding surrogate models for predictive maintenance (Siemens Digital Industries Software, 2023; Karkaria et al., 2024).

The general workflow is shown in Figure 7. It links four stages: build the twin, acquire and preprocess data, train predictive models, and deploy them for anomaly detection and prognostics. In the EPS/VEBIC setting, the flow starts from a detailed crank-angle-resolved engine model in GT-Suite, reduces it to a fast-running model, imports it into Simulink, couples it to the electrical plant (generator, converter, DC link, optional BESS), and executes the twin in real time on a Speedgoat target for hardware-in-the-loop co-simulation. The figure also notes practical constraints such as solver settings and CPU overload that matter for real-time execution (Söderäng et al., 2022).

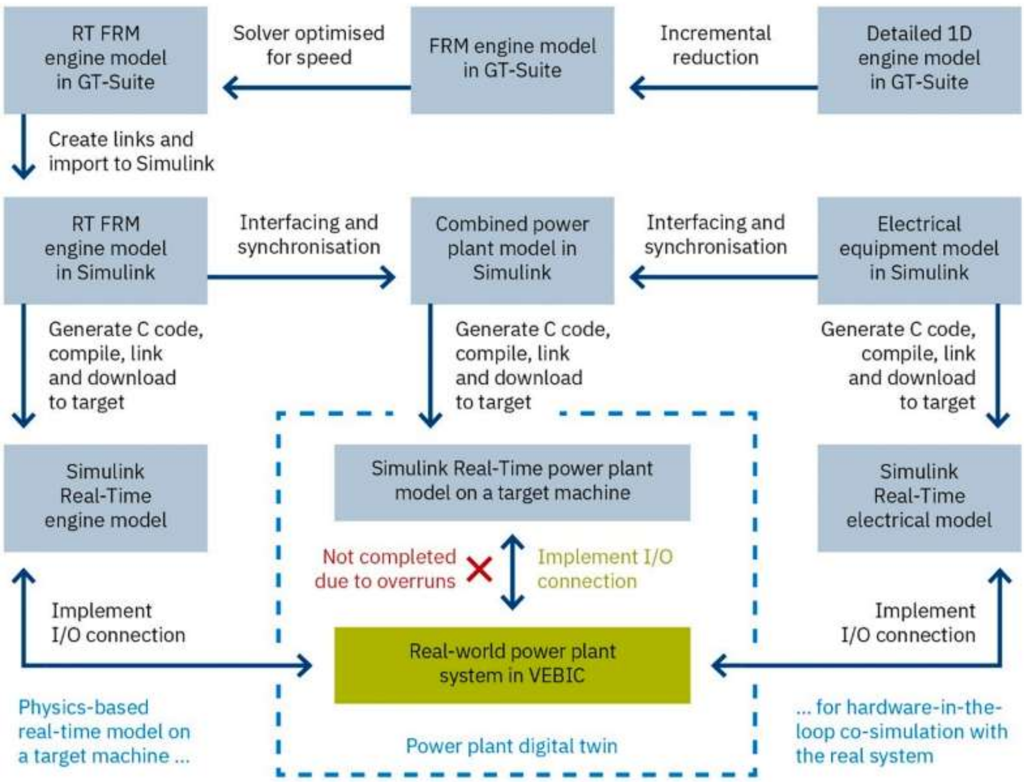

Figure 7. Digital-twin workflow: reduced engine model, real-time/HIL coupling, residuals and RUL features (Söderäng et al., 2022).

4.4 Data and Communication Architecture

For autonomous plants to operate reliably, data must be transferred between the plant and the Remote Operation Centre (ROC) with both accuracy and structure. Different functions place different demands on synchronization. High-frequency measurements, such as crank-angle pressure traces, demand timing precision within the sub-millisecond range to maintain coherence. In contrast, supervisory data, including operator inputs and event logs, can tolerate lower temporal resolution, with mixing these requirements would risk inconsistent data streams and delayed control responses, which would weaken both responsiveness and safety, and to prevent this risk, the system applies a dual synchronisation strategy. PTP is used for control-critical measurements that require microsecond accuracy, while NTP provides the millisecond precision needed for supervisory functions.

Process data and event information are transmitted through separate communication channels. The Continuous variables like the power, temperature and pressure are transferred using a structured client, server protocol that guarantees orderly and reliable communication. Event-driven information, including alarms and fault reports, is handled through a lightweight publish, and subscribe protocol designed for fast message delivery. This division ensures that continuous monitoring remains efficient, while time-critical events are reported without delay. As outlined in Section 4.1.2, the communication framework separates continuous and event-based data exchange to preserve efficiency and maintain deterministic behaviour across the network.

Data acquisition and data storage are combined in a single and integrated system while the acquisition system records both high and low frequency signals and aligns them by using the deterministic timing, then selected measurements are transferred to an archival database for long-term use to including model training, benchmarking, and operational analysis. To reduce network load, feature extraction is carried out at the edge before transmission, improving scalability while retaining diagnostic value.

Through this combination of synchronised acquisition, structured data handling, and historical storage, the architecture provides a dependable foundation for supervisory intelligence. The flow of information begins at the sensors, continues through local control and acquisition systems, and is processed at the ROC before being returned as verified commands to the plant. Each stage operates within defined timing constraints, supporting deterministic behaviour and reinforcing both safety and reliability in autonomous plant operation.

4.4.1 Time Synchronization (PTP/NTP)

Reliable time synchronization is fundamental to autonomous plant operation. Supervisory intelligence, reinforcement learning, and digital twins rely on coherent time-stamped data to link high-frequency signals with slower process measurements. Without sub-microsecond synchronization at the plant level and millisecond alignment at the supervisory level, latency and jitter would reduce the accuracy of model predictive control and weaken the reliability of AI-driven decisions.

The proposed architecture applies a dual-domain synchronization strategy. NTP provides millisecond-level accuracy for supervisory tasks such as logging, operator interactions, and event timestamps. PTP, is used for control-critical measurements, including crank-angle encoders, in-cylinder pressure sensors, and high-speed emission analyzers. Sub-microsecond precision is required at this level to preserve cycle-resolved accuracy. The Dewesoft DAQ system, described in Chapter 3, functions as the timing backbone, receiving PTP signals from a GPS-disciplined grandmaster clock and distributing synchronized time to high-frequency instruments.

A bridging mechanism ensures consistency between NTP and PTP domains. The grandmaster clock outputs both protocols in parallel, while boundary clocks and time-aware switches distribute synchronized time and provide translation between domains.

Latency auditing agents monitor drift and generate quality-of-time indicators, ensuring trust in the synchronized data streams. Hybrid synchronization of this kind has been validated in practice. ABB's Longmeadow microgrid applies PTP for inverter-generator synchronization and NTP for SCADA logging (ABB, 2015). The Sandia Secure Microgrid Test Bed integrates PTP across sensors, controllers, and hardware-in-the-loop simulators, while NTP is used for test management functions (Zhang & Liu, 2019). Substation automation systems following IEC 61850 apply PTP on the process bus and NTP on the station bus, fully consistent with IEEE 1588 standards (Rehtanz, 2003; Eidson, 2006; IEEE, 2008).

combining PTP and NTP within a single time-distribution architecture, the system ensures that plant-level measurements and supervisory ROC data remain aligned and coherent. This synchronization forms the timing foundation needed for model predictive control, AI modules, and digital twins to operate with reliability and safety.

4.4.2 Telemetry and Messaging Protocols

Reliable communication between the autonomous plant and the Remote Operation Centre (ROC) requires protocols that can handle both structured process data and event-driven information. If these data streams are not separated, supervisory systems and AI modules can receive delayed or inconsistent inputs. This would weaken responsiveness and lower the reliability of control decisions. A dedicated telemetry layer is therefore needed to link different field devices while keeping latency low and communication secure and compatible.

In the proposed system, the process and event data are handled through different communication paths. The process variables, such as electrical power, state of charge and thermal measurements are transferred using OPC Unified Architecture (OPC UA). The protocol supports hierarchical data models, semantic tags and secure transfer through TLS and PKI, it also connects directly with archive databases for storing results from long-term tests. Event-based data, including alarms, health information, and fault reports, are sent using MQTT. This publish subscribe protocol is lightweight and provides quality-of-service options that make it suitable for fast and decoupled event transmission. Studies have shown that OPC UA works well for structured supervisory control, while MQTT fits better for quick event communication (Tebekaemi and Wijesekera, 2018; Wang et al., 2022).

For protection functions the deterministic protocols are also required while the architecture uses IEC 61850 GOOSE and MMS messaging to send interlock and trip commands with low delay to ensuring that protection actions are executed immediately. OPC UA, MQTT, and IEC 61850 together follows practices already recommended for mixed power system environments (Rehtanz, 2003; Sajadi et al., 2019).

To stay compatible with older equipment, Modbus TCP/IP is still used inside the test cell, though it is not connected to the ROC levels. Security is applied across all communication layers. OPC UA channels are secured with TLS and PKI-based authentication, MQTT brokers use TLS with topic-level access rules, and IEC 61850 traffic is protected according to IEC 62351 to reduce spoofing or replay risks. Together, these settings reduce protocol vulnerabilities and keep communication between field and supervisory levels reliable.

OPC UA, MQTT, and IEC 61850, the architecture establishes a reliable telemetry backbone for the ROC. This ensures that AI modules, reinforcement learning agents, and digital twins receive consistent, time-aware data streams. In doing so, the supervisory layer is able to perform accurate state estimation and deliver responsive control in autonomous plant operation.

4.4.3 Data Acquisition and Integration

Autonomous power plants depend on the continuous collection and long-term storage of both high-frequency and low-frequency process data. If these streams are not stored and aligned within a common structure, reinforcement learning agents, digital twins, and anomaly detection systems cannot be trained, validated, or applied effectively during operation.

The acquisition and subsystem addresses this by combining different data streams into a single, time-aligned framework. The Dewesoft DAQ system presented in Chapter 3 records cycle-resolved pressure data, injection signals, and slower auxiliary measurements, all synchronized with PTP. From this collection, selected variables are transferred to past database for long-term storage. This database provides the foundation for reinforcement learning training, digital twin calibration, and operational analysis. Recent benchmarking studies show that time-series databases such as InfluxDB and Lindorm are well suited to handling large volumes of mixed-frequency, time-stamped data efficiently (Jensen et al., 2017; Grzesik & Mrozek, 2020).

Synchronisation across different DAQ sources is achieved using deterministic clocks, which solve the challenge of combining high-speed measurements with supervisory-level data (IEEE, 2008; Eidson, 2006; ABB, 2015; Zhang & Liu, 2019). To reduce bandwidth use, feature extraction is performed at the edge before transmission. Extracted indicators, such as combustion phasing or maximum pressure rise, preserve diagnostic value while lowering communication demands. This method supports early detection of faults (Grzesik & Mrozek, 2020) and enables sub-second anomaly detection through streaming pipelines (Vidal et al., 2022). Latency is managed so that storage does not delay access to variables required for supervisory AI.

Integrating DAQ and archive systems, the architecture provides a coherent data foundation for AI-based supervisory control. Reinforcement learning agents and digital twins use this repository to retrain models, validate decisions, and identify anomalies. In this way, past dataset ensures that supervisory intelligence operates on accurate, synchronised, and time-consistent process data.

4.4.4 Edge-Fog-Cloud Data Flows

Autonomous power plants depend on a layered flow of data between the edge, fog, and cloud. This arrangement allows fast local control, intermediate processing, and large-scale learning to work together. Real-time control tasks are kept at the edge, close to the equipment, where timing is most critical. The fog handles aggregation and feature processing, while the cloud is used for training and updating models. Without this distribution, delays and bandwidth limits would restrict how well supervisory intelligence could respond.

In the ROC design, measurements from sensors and controllers are collected and time-aligned at the edge using the Dewesoft DAQ, which provides a consistent data basis for model predictive control. At the fog layer, these data streams are aggregated, key features are extracted, and cleaned information sets are forwarded to the ROC. This hierarchical arrangement allows reinforcement learning, anomaly detection, and streaming analytics to be executed locally at the fog layer, while the ROC utilises the processed outputs for supervision, optimisation, and coordination without overloading the communication network. The cloud provides the computing resources for long-term training and for updating digital twins. Updated model components and optimisation weights, generated through cloud-based training, are periodically synchronised with the ROC to enhance supervisory decision-making. The approach follows established edge-fog-cloud frameworks (Shi et al., 2016; Satyanarayanan, 2017). The prior studies reflect that distributing tasks across layers is needed to keep latency predictable, in power applications, AI execution also requires clear limits on delay and jitter (Vidal et al., 2022; Yıldırım et al., 2025). In the proposed system the high-frequency signals remain time-aligned at the DAQ through hardware-based synchronization. ROC-level assumption is completed within hundreds of milliseconds, while runtime assurance filtering operates with a response time below 10 milliseconds. These timing constraints work to ensure reliable supervisory coordination and preserve safe plant operation.

By dividing tasks across the three layers, the system combines immediate responsiveness with long-term adaptability. Edge-level MPC reacts quickly to disturbances, fog-level processing ensures efficient data handling, and cloud-level retraining supports continuous improvement. This makes it possible to run reinforcement learning agents and digital twins within the ROC while still respecting real-time requirements.

4.5 Security and Safety Envelope

Autonomous power plants need protection both from external cyber threats and from faults that could compromise safe operation, for this reason, the framework is built with two layers. The outer layer focuses on cybersecurity that working to detect intrusions and block harmful activity. The inner layer focuses on safety, making sure that even if a breach occurs, the plant can still operate within acceptable limits through runtime checks and fallback control actions.

Proofs and studies from several fields show and present the benefit of this combined approach. In shipboard microgrids, for example the studies that report an attack like data manipulation or rootkits can interfere with normal operation, and If there is no detection at the security level and no fallback at the safety level then the system may start to behave in unsafe ways (Gupta et al., 2023; Al Quayed et al., 2024; Qiu et al., 2020). While similar findings have been reported in nuclear power plants, where missing runtime checks have allowed security events to grow into safety incidents (Chen et al., 2020). Observations from renewable energy systems point in the same direction are showing that strong authentication and secure communication are needed to stop cyber problems from turning into operational risks (Maes, 2013; Liu et al., 2022). In general these studies present that cybersecurity and safety work best when they are treated as linked parts of the same defence structure, the cybersecurity layer targets risks at the level of devices, data, and communication. Hardware-based measures such as physically unclonable functions (PUFs) can provide unique device identities, while lightweight encryption methods protect confidentiality without exceeding the capacity of embedded controllers (Maes, 2013; Beaulieu et al., 2015). Demanding more in cryptographic operations can be handled by fog-level nodes that have higher processing capacity (Beaulieu et al., 2015). The communication protocols used in the system are also secured. OPC UA is protected through TLS and PKI, while IEC 62351 adds security to IEC 61850 messages, reducing the risk of replay or spoofing attacks (Rehtanz, 2003; Wang et al., 2022). MQTT and Modbus, which do not include strong authentication by default, can be strengthened with role-based access control, topic filtering, and intrusion detection (Gupta et al., 2023; Sajadi et al., 2019). Even time synchronisation streams such as PTP are monitored for abnormal behaviour to detect tampering or injection (Eidson, 2006; IEEE, 2008). Together, these measures create several layers of defence across devices, networks, and supervisory control.

The safety layer acts when cybersecurity barriers are bypassed. All commands from the ROC are checked against predefined constraints before they are applied. If a command is unsafe, the system reverts to baseline logic or model predictive control setpoints that are known to be stable (Rawlings et al., 2020). Reinforcement learning agents are filtered in real time so that unsafe policies are never passed to execution (Bui et al., 2024; Yu et al., 2024). Fault tolerance is further improved through redundancy, with local controllers sharing responsibility for safety and handling faults independently when required (Han, 2023; Ghosh et al., 2020). In this way, safe operation is maintained even under compromised communication or unexpected disturbances.

In General, combining cybersecurity with safety creates a balanced protection strategy. The cybersecurity layer protects data integrity and communication paths, while the safety layer enforces operational limits during normal operation and acts when those limits or trust boundaries are crossed. Research in microgrids, nuclear systems, and renewable plants shows that this mix of preventive and reactive measures lowers the chance of cyber incidents developing into unsafe conditions

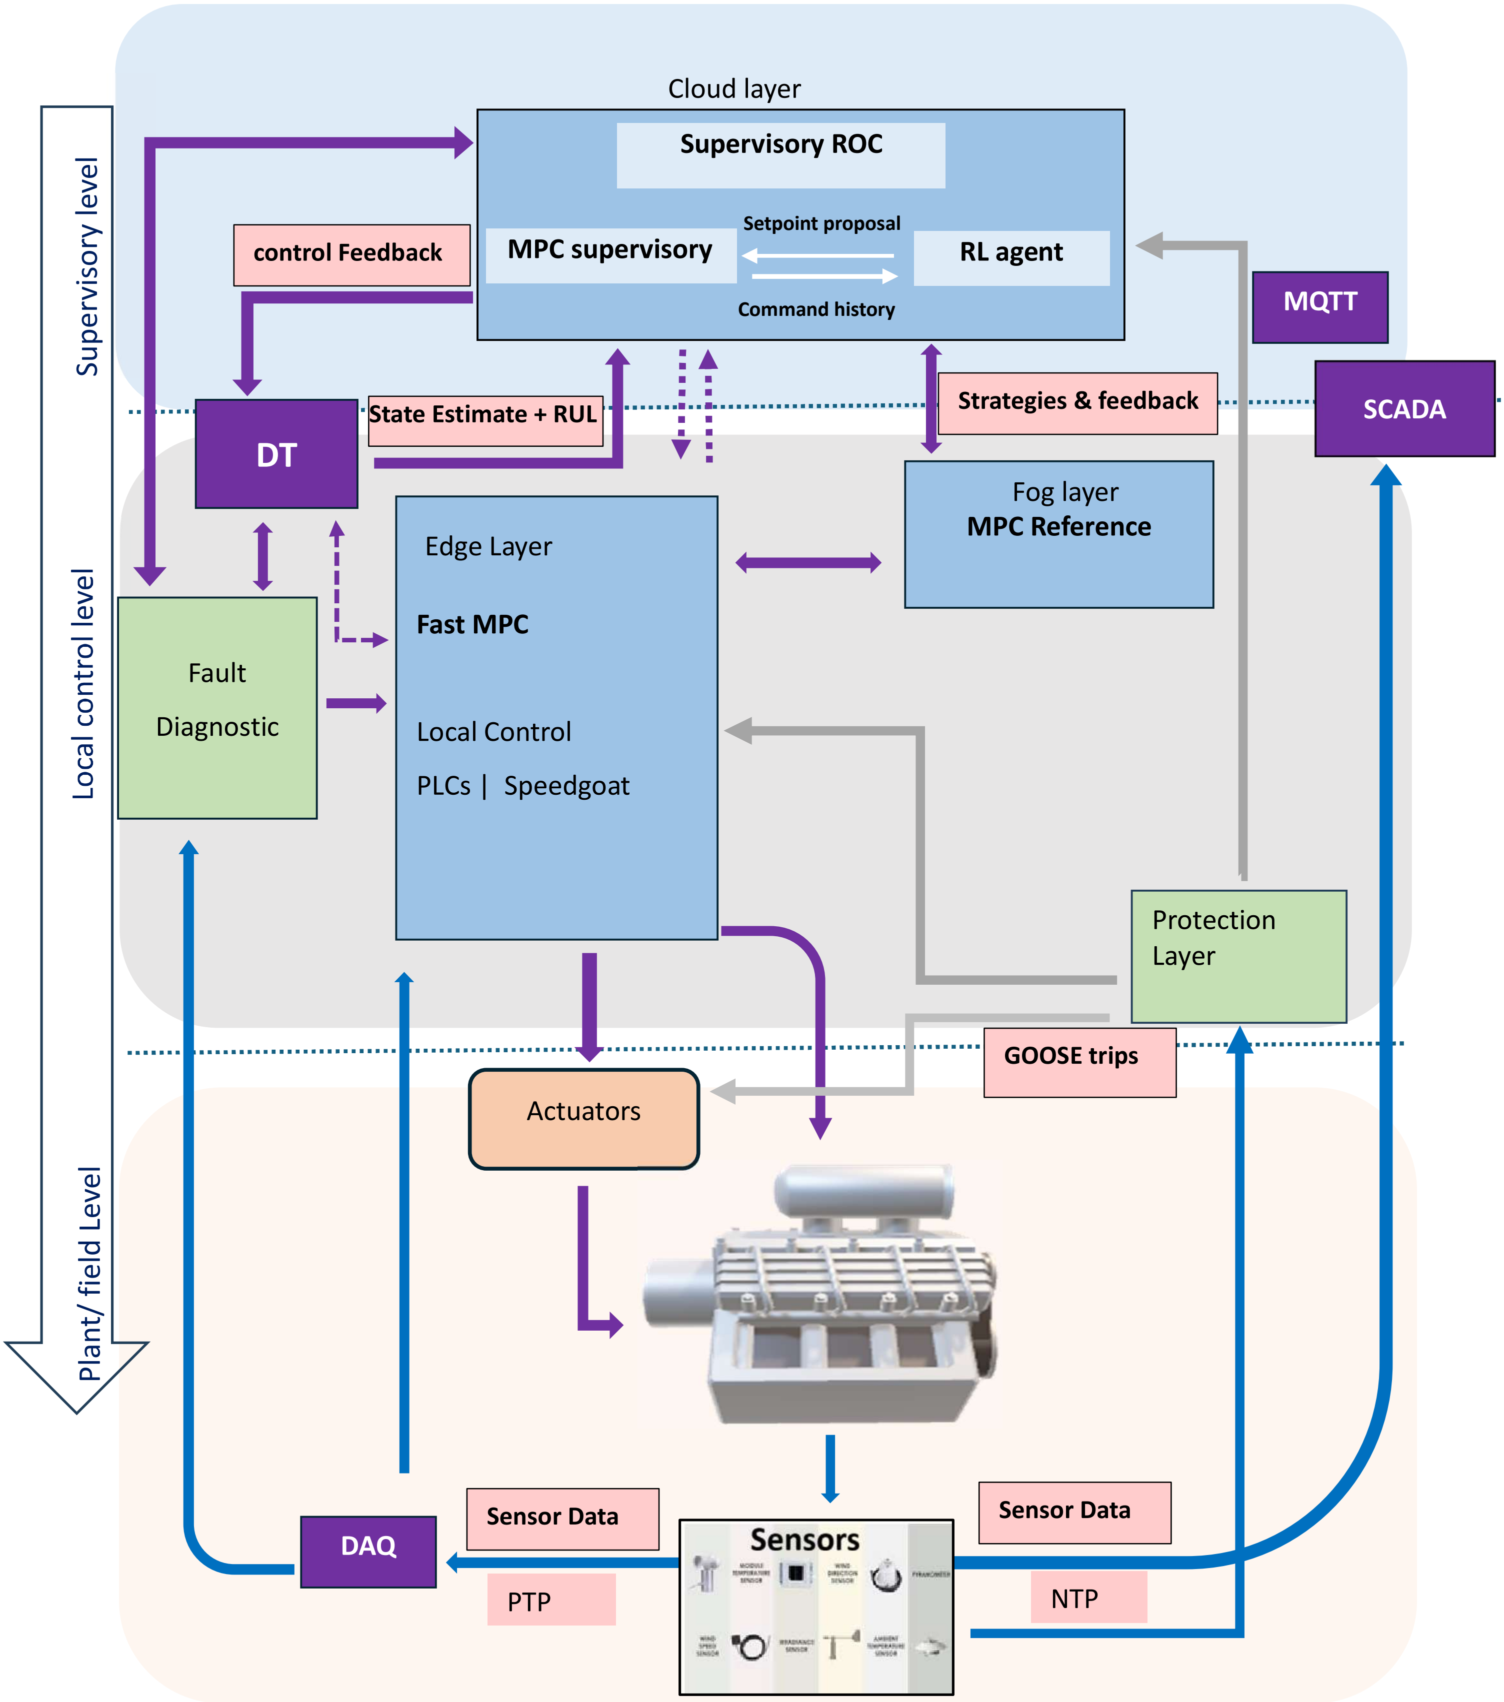

Figure 8. Supervisory ROC proposal Design.

Table 7. Legend table of Figure 8.

Label	Meaning
ROC	Remote Operation Centre (Cloud-level supervisory control)
RL	Reinforcement Learning
MPC	Model Predictive Control
DT	Digital Twin
DAQ	Data Acquisition system (Dewesoft-based, time-synchronised)
GOOSE	Generic Object Oriented Substation Events (IEC 61850 protocol)
IED	Intelligent Electronic Device (protection interface)
PTP / NTP	Precision Time Protocol / Network Time Protocol for time synchronisation
Edge	Field-level hardware
Fog	Intermediate computing layer for distributed MPC and data preprocessing
Cloud	ROC-level servers for supervisory MPC, RL, and twin synchronisation
Trip	Deterministic protection signals
Setpoints	Control commands or reference trajectories generated by ROC
→	Control Commands
→	Data Flow

5 Feasibility Analysis and Integration Challenges

5.1 Technical Feasibility and Gap Mapping

Evaluating the feasibility of the ROC concept involves comparing the EPS Laboratory setup with recognised automation frameworks and the requirements outlined in Chapter 4, the ISA-95 hierarchy (International Electrotechnical Commission, 2003) and the NIST Cyber-Physical Systems Framework (2017) are used as reference models for this assessment. Both frameworks distinguish between local deterministic control, defined as Level 2, and supervisory coordination, defined as Level 3. At present, the EPS Laboratory shows strong Level 2 performance through its PLC-based routines and Speedgoat real-time targets. These systems run deterministic feedback loops that manage combustion, safety interlocks, and actuator control (Mikulski, 2025). Operator-driven actions such as start-up, purging, and fuel switching are executed automatically once initiated, providing reliability while keeping human oversight in the process(Valkjärvi, 2022). Supervisory monitoring is provided by a SCADA interface, which allows operators to observe system states, follow transitions, and respond to alarms (Hautala, 2022). In some cases, such as mode transitions from diesel to RCCI, the switch is triggered automatically under defined conditions, showing a degree of conditional automation (Valkjärvi, 2022).

These features indicate that the laboratory already extends into selected Level 3 functionality, particularly in SCADA dashboards, past-style data collection, and semi-automated procedures. However, it does not represent a complete Level 3 environment. Workflow scheduling, coordinated resource management, and enterprise integration are absent, and supervisory authority remains operator-driven rather than system-driven (ISA-95; NIST, 2017; McBreen, 2024). For this reason, the EPS Laboratory is best described as Level 2 with partial Level 3 features, a solid basis for deterministic control but still below the requirements for ROC readiness.

The current condition and the ROC vision are separated by a number of technical issues. Without adaptive optimization layers like MPC or RL that enable constraint-aware decision-making under uncertainty, control stays fixed and rule-based (Rawlings et al., 2020; Yu et al., 2024). Fault detection relies on threshold alarms, lacking predictive capabilities demonstrated in AI-enabled prognostics (Sleiti et al., 2022; Machalski et al., 2025; Li et al., 2022). Data acquisition is strong in sampling rate but fragmented across combustion, control, and auxiliary channels without semantic integration or deterministic time alignment, the communication is limited to Modbus, CAN and VPN access whereas ROC requirements include PTP/NTP synchronisation and interoperable protocols such as OPC UA, MQTT, and IEC 61850 (Wang et al., 2022; Gupta et al., 2023; Rehtanz, 2003). The at present, cybersecurity and safety provisions are not yet aligned with international standards such as IEC 61508, ISO 26262, and IEC 62351. This leaves assurance gaps that must be addressed before fully autonomous operation can be realised.

Even with these limitations, the EPS Laboratory provides a practical foundation for further development. The combination of Speedgoat hardware and Simulink workflows forms a controlled testbed for model-in-the-loop and hardware-in-the-loop integration of MPC, RL, and digital twin modules (Rawlings et al., 2020; Dickler et al., 2021; Liu et al., 2022). The Dewesoft data acquisition platform already supports high-rate measurement and can be expanded with PTP-based synchronisation to enable supervisory AI applications (ABB, 2015; Zhang and Liu, 2019). SCADA monitoring can also be extended toward cloud dashboards, allowing ROC-level diagnostics and remote operation (Dörfler et al., 2019; Vidal et al., 2022). Increasing redundancy in sensors, adding middleware for data fusion, and incorporating structured historical datasets would further improve system observability and reliability.

Table 7 presents a structured comparison between the current capabilities of the EPS Laboratory and the requirements defined for ROC implementation. It shows where alignment already exists, such as modular energy system design and deterministic local control, and identifies remaining gaps in adaptive intelligence, predictive fault management, communication interoperability, and compliance with standards. This mapping establishes the technical feasibility baseline and guides the operational integration analysis discussed in Section 5.2.

Table 8. Comparison of EPS Laboratory capabilities and ROC requirements.

Dimension	EPS Laboratory	ROC Requirements	Gap / Recommendation
Primary Energy System	Wärtsilä 4L20 multifuel engine + BESS hybrid	Modular hybrid generation system	Aligned
Sensors Installed	Thermocouples, flow, pressure, vibration, emission	Multi-modal, redundant, calibrated sensors	Add redundancy and smart sensors
Data Acquisition	Dewesoft modular DAQ, high sampling rates	Time-synchronised, scalable DAQ	Add PTP/NTP timestamping
Real-time Control	Speedgoat RT controller	Edge computing with adaptive control	Integrate AI modules for ROC scenarios
Monitoring Interface	SCADA + DAQ software	Cloud-based ROC with analytics dashboard	Extend SCADA to predictive diagnostics
Data Storage	Local storage, limited archiving	Cloud/distributed storage for ROC analytics	Integrate time-series DB + cloud sync
Communication	Modbus, CAN, analog I/O	OPC UA, MQTT, IEC 61850	Upgrade to standard interoperable protocols
Remote Access	Limited VPN	Secure bidirectional ROC interface	Enhance secure remote access capability
Cybersecurity	Not defined	Hardened architecture with intrusion detection	Add encryption, access control, logging
Autonomy Readiness	Semi-automated; operator supervision	Autonomous decision-making with fallback layers	Implement autonomy framework and fail-safes

5.2 Operational Integration and Interoperability

The ROC requires local control, supervisory functions, and combined data flows into a consistent operational framework. In the EPS (VEBIC) Laboratory this integration is still limited. Control logic is divided between PLCs and Speedgoat targets, while acquisition channels for combustion, control, and auxiliary data are retained separate. Communication relies on Modbus, CAN, and VPN links. These arrangements prevent supervisory modules from accessing time-aligned information and limit transparency at the operator level.

ROC operation assumes that integration is performed on a standards-based layer. As mentioned, time synchronisation must follow PTP at the field level, and NTP at the supervisory level. Telemetry should be exchanged through OPC UA, while event information is handled through MQTT. These protocols support semantic consistency and scalable event delivery, and they are already applied in industrial practice (Wang et al., 2022; Gupta et al., 2023; Rehtanz, 2003). Integration also extends to operator interfaces. Current SCADA dashboards support monitoring but not explainable decision support. For ROC deployment, supervisory recommendations from MPC or RL modules must be presented with interpretable reasoning, while alarms need to be prioritised and predictive. This staged approach to automation strengthens trust and keeps the operator in the supervisory loop (DNV, 2022; ABB, 2024).

Finally, interoperability must be designed across sites. ROC frameworks assume that local nodes continue basic operation under cloud disconnection, while cloud functions manage long-horizon optimisation and learning. This edge-fog-cloud distribution is consistent with recent industrial architectures (Dörfler et al., 2019; Vidal et al., 2022).

In summary, the EPS Laboratory must move from isolated subsystems to a synchronised and standards-based structure. This integration forms the basis for the economic and scalability analysis discussed in Section 5.3. The operational readiness of the EPS Laboratory can also be summarised through an automation matrix. Table 8 reflects the current control and supervisory practices against ROC requirements across five dimensions. The mapping shows that while deterministic control logic is established, gaps remain in predictive fault handling, semantic data fusion, ROC-ready interfaces, and compliance with safety standards.

Table 9. EPS Laboratory automation readiness matrix

Dimension	**Current state**	**Target ROC readiness**	**Gap / limitation**
Control logic	Rule-based (Simulink, PLC)	Supervisory MPC with RL guidance; runtime-assured	No adaptive supervisory layer yet
Fault response	Threshold/reactive alarms	Predictive, prioritised alarms; autonomous mitigation paths	No residual/DT/RL-based indicators
Data fusion	Separate DAQ channels (combustion/process/aux.)	Unified, PTP-aligned telemetry; OPC UA + archive scheme	No semantic layer or common clock
Interface	Local SCADA + VPN scripts	ROC HMI with explainability, command validation, rollback	Limited bi-directional supervision
Safety & standards	Formal alignment not documented	IEC 61508 / ISO 26262; IEC 62351; V&V for AI	Missing safety case and test artefacts

6 Conclusions

This thesis investigates Remote Operation Centre (ROC) for autonomous power plants and examined its feasibility in the EPS Laboratory. The ROC was treated as a supervisory layer that combines time-synchronised data handling, model predictive control (MPC), reinforcement learning (RL) and a digital twin, with command validation and operator support. The architecture was developed in Chapter 4 and assessed against laboratory constraints in Chapter 5.

Concerning research question 1 - the functional and architectural requirements for the ROC, it can be concluded that a layered structure is required. Deterministic control remains local. Supervisory optimisation and learning are hosted at the ROC. Protection operates independently of supervisory intelligence. The mapping follows the edge-fog-cloud distribution and the timing and authority limits stated in the architecture chapter.

Concerning research question 2 - the integration of MPC, RL and the digital twin, it can be concluded that the roles are complementary. RL provides adaptive guidance. MPC enforces feasibility and constraints. The digital twin supplies state estimates, surrogate dynamics and health-aware limits. A runtime-assurance gate is required so that only admissible actions are dispatched.

Concerning research question 3 - the data collection, fusion and communication pipelines, it can be concluded that time coherence and semantic structure are prerequisites. PTP is needed at field level and NTP at supervisory level. OPC UA should carry structured process telemetry and past links, and MQTT should carry event streams. These choices support multi-rate fusion, auditability and remote supervision.

Concerning research question 4 - reliability, transparency and operator trust, it can be concluded that command validation, prioritised predictive alarms, staged authority and explainable recommendations are necessary to keep the operator in the loop while enabling higher autonomy.

The feasibility analysis shows that the EPS Laboratory provides deterministic local control (Speedgoat/PLC), SCADA-based supervision, high-rate data acquisition and remote access. These features are suitable for developing the ROC data plane and twin-in-the-loop tests. Operations management functions are not present, and supervisory authority remains operator-supervised. Gaps remain in adaptive supervisory intelligence, predictive alarms and diagnostics, semantic data fusion with PTP alignment, interoperable telemetry (OPC UA/MQTT/IEC 61850) and documented safety and cybersecurity artefacts. It can be concluded that the main contribution is an implementable ROC framework that links supervisory MPC/RL/twin, runtime assurance and a time-synchronised data plane to the EPS Laboratory context. This provides a practical route from deterministic supervision to auditably safe autonomy.

The work is limited to architectural design and laboratory-based feasibility. No closed-loop field validation of cloud-hosted supervisory MPC/RL was executed, and a complete safety case and V&V package for AI components remain future work.

The outlook is incremental. First deploy PTP/NTP and an OPC UA information model with archive schema and add MQTT for events. Then implement residual-twin-based alarm logic with a runtime-assurance gate and decision logging. Next, a reduced-order digital twin with embedded state estimation should be introduced to support supervisory MPC validation and health-aware constraint formulation. Reinforcement learning can complement this by enabling long-term optimisation and adaptive decision making beyond predefined operating modes. Health-related functions, such as predictive maintenance and condition-based scheduling, can further enhance system resilience and lifecycle performance. Finally, staged autonomy should be piloted through an ROC human-machine interface that supports explainability, rollback, and multi-asset supervision.

Acknowledgement

AI-assisted tools were used for non-substantive support only. The tools of ChatGPT assisted in generating and checking abbreviations, formatting equation code, polishing table layout, and scanning for grammar and punctuation. No AI system was used to produce research questions, methods, results, analyses, or citations. All interpretations, argumentation, and source verification were performed by the author.

References

ABB Marine & Ports. (2020, August 30). *Sustainable shipping: Autonomous and zero-emission vessels powered by ABB* [White paper]. ABB. https://new.abb.com/docs/librariesprovider15/gov/usv-abb-white-paper-20200830.pdf

ABB. (2020, November 23). *The digital twin: From hype to reality. ABB News.* https://new.abb.com/news/detail/80770/the-digital-twin-from-hype-to-reality

ABB. (2015). *ABB microgrid reference case – Longmeadow facility, Johannesburg, South Africa*. ABB Group. https://resources.news.e.abb.com/attachments/published/38167/sv-SE/CBEF037E121A/Data-sheet-Longmeadow-FINAL.pdf

ABB. (2024). *Integrated report 2024*. ABB Group. https://library.e.abb.com/public/ab32fda35c6f4a3cb62f8c5e2e4476c6/Integrated%20Report%202024_ENGLISH.pdf

ABB. (2024). *Products have a digital twin—and you can find it too!* https://new.abb.com/news/detail/116709/products-have-a-digital-twin-and-you-can-find-it-too

Ahsan, M., Rodriguez, J., & Abdelrahem, M. (2023). Distributed control algorithm for DC microgrid using higher-order multi-agent system. *Sustainability, 15*(10), 8336. https://doi.org/10.3390/su15108336

Akhundov, H., van der Sluis, E., Hamdioui, S., & Taouil, M. (2020). *Public-key based authentication architecture for IoT devices using PUF*. In Proceedings of the 19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN) (pp. 249–250). https://doi.org/10.5121/csit.2019.91328 SSRN

Al-Quayed, F., Ahmad, Z., & Humayun, M. (2024). A situation-based predictive approach for cybersecurity intrusion detection and prevention using machine learning and deep learning algorithms in wireless sensor networks of Industry 4.0. *IEEE Access, 12*, 34800–34819. https://doi.org/10.1109/ACCESS.2024.3372187

Arulampalam, M. S., Maskell, S., Gordon, N., & Clapp, T. (2002). A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing, 50*(2), 174–188.

Bai, Y., Jantunen, H., & Juuti, J. (2018). *Hybrid, multi-source, and integrated energy harvesters*. Frontiers in Materials, 5. https://doi.org/10.3389/fmats.2018.00065

Beaulieu, R., Shors, D., Smith, J., Treatman-Clark, S., Weeks, B., & Wingers, L. (2015). *The SIMON and SPECK lightweight block ciphers*. Cryptology ePrint Archive. http://eprint.iacr.org/2013/404.pdf

Bui, V.-H., Das, S., Hussain, A., Hollweg, G. V., & Su, W. (2024). A critical review of safe reinforcement learning techniques in smart grid applications. *arXiv*. https://doi.org/10.48550/arXiv.2409.16256

Bumiller, G., Lampe, L., & Hrasnica, H. (2010). Power line communication networks for large-scale control and automation systems. *IEEE Communications Magazine, 48*(4), 106–113. https://doi.org/10.1109/MCOM.2010.5439083

Calistoga City Council. (2022). *Calistoga Resiliency Center project update*. Retrieved from https://www.ci.calistoga.ca.us

Chemical Processing. (2023). *Shell showcases success with autonomous plant*. ChemicalProcessing.com. https://www.chemicalprocessing.com/automation/automation-it/article/11299394/shell-showcases-success-with-autonomous-plant-chemical-processing

Chen, D., Chen, K., Li, Z., Chu, T., Yao, R., Qiu, F., & Lin, K. (2020). PowerNet: Multi-agent deep reinforcement learning for scalable powergrid control. *arXiv*. https://doi.org/10.48550/arXiv.2011.12354

Chen, K., Lin, J., Qiu, Y., Liu, F., & Song, Y. (2021). Deep learning-aided model predictive control of wind farms for AGC considering the dynamic wake effect. Control Engineering Practice, 116, 104925. https://doi.org/10.1016/j.conengprac.2021.104925

Chen, Y.-P., Liu, Y., Zhang, Y., Wu, B., & Cao, J. (2025). Real-time decision-making for digital twin in additive manufacturing: MPC with TiDE surrogate. *Computers &*

Chemical Engineering. (Advance online publication). https://arxiv.org/abs/2501.07601

Damjanović, I., Pavić, I., Puljiz, M., & Brčić, M. (2022). Deep reinforcement learning-based approach for autonomous power flow control using only topology changes. *Energies, 15*(19), 6920. https://doi.org/10.3390/en15196920

Das, A., Kong, W., Leach, A., & Xu, H. (2023). Long-term forecasting with TiDE: Time-series dense encoder. *arXiv*. https://arxiv.org/abs/2304.08424

Det Norske Veritas (DNV). (2022). Remote operation centres for autonomous and remotely operated ships. DNV Maritime Report. https://www.dnv.com/maritime/autonomous-remotely-operated-ships/

Devold, H., & Fjellheim, R. (2019, September). Artificial intelligence in autonomous operation of oil and gas facilities. In *SPE Offshore Europe Conference and Exhibition*. Society of Petroleum Engineers. https://doi.org/10.2118/197399-MS

Dewesoft. (2023, May 22). Data acquisition synchronisation and timing methods. https://dewesoft.com/blog/data-acquisition-synchronisation dewesoft.com

Dhanraj, R., Parmar, M., Shah, H., Patel, S., & Pandya, V. (2020). MEMS accelerometer based vibration analysis for predictive maintenance of rotating machines using IoT. *2020 IEEE International Conference on Advent Trends in Multidisciplinary Research and Innovation (ICATMRI)*, 1–5. https://ieeexplore.ieee.org/document/9273212

Di Silvestre, M. L., Favuzza, S., Riva Sanseverino, E., & Zizzo, G. (2018). How decarbonization, digitalization and decentralization are changing key power infrastructures. *Renewable and Sustainable Energy Reviews, 93*, 483–498. https://doi.org/10.1016/j.rser.2018.05.068

Di Silvestre, M. L., Graditi, G., & Sanseverino, E. R. (2020). A survey of the role of Smart Grids in enabling the integration of renewable energy sources. *Renewable and Sustainable Energy Reviews, 134*, 110400. https://doi.org/10.1016/j.rser.2020.110400

Dickler, S., Wintermeyer-Kallen, T., Zierath, J., Bockhahn, R., Machost, D., Konrad, T., & Abel, D. (2021). Full-scale field test of a model predictive control system for a 3

MW wind turbine. *Forschung im Ingenieurwesen*, 85, 313–323. https://doi.org/10.1007/s10010-021-00467-w

Dörfler, F., Jovanović, M. R., Chertkov, M., & Bullo, F. (2019). Distributed control and optimization for autonomous power grids. *Nature Communications, 10*, 400. https://doi.org/10.1038/s41467-019-11400-1

Eidson, J. C. (2006). *Measurement, control, and communication using IEEE 1588*. Springer. https://doi.org/10.1007/1-84628-251-9

FlytBase. (2023). *Autonomous docked drones for industrial sites*. Case Studies. https://www.flytbase.com/case-studies/sensorem-autonomous-docked-drones-deployment-australia

Fraden, J. (2016). *Handbook of modern sensors: Physics, designs, and applications* (5th ed.). Springer. https://doi.org/10.1007/978-3-319-19303-8

Ghosh, P., Eisele, S., Dubey, A., Metelko, M., Madari, I., Völgyesi, P., & Karsai, G. (2020). Designing a decentralized fault-tolerant software framework for smart grids and its applications. *Journal of Systems Architecture, 109*, 101759. https://doi.org/10.1016/j.sysarc.2020.101759

Ghosh, S., Abdelwahed, S., & Zonouz, S. (2019). On the design of fault-tolerance in a decentralized software framework for smart grid control applications. In *Proceedings of the 15th ACM International Conference on Embedded Software (EMSOFT)* (pp. 1–10). https://doi.org/10.1145/3349266.3351372

Ghosh, S., Pierre, J. W., Zonouz, S., & Abdelwahed, S. (2020). Designing a decentralized fault-tolerant software framework for smart grid control applications. *ACM Transactions on Cyber-Physical Systems, 4*(3), 1–24. https://doi.org/10.1145/3364222

Goerlandt, F. (2020). Maritime autonomous surface ships from a risk governance perspective: A review. *Safety Science, 128*, 104758. https://doi.org/10.1016/j.ssci.2020.104758

Gupta, C., & Varshney, G. (2023). *A lightweight and secure PUF-based authentication and key-exchange protocol for IoT devices*. SSRN. https://doi.org/10.2139/ssrn.4646453 SSRN+1

Han, Y. (2023). Wavelet analysis and consensus algorithm-based fault-tolerant control for smart grids. *Frontiers in Energy Research, 11*, 1160256. https://doi.org/10.3389/fenrg.2023.1160256

Hao, S., Li, W., & He, Y. (2011). Fault-tolerant control of linear permanent-magnet synchronous motor systems based on decentralized sliding-mode observers. *IEEE Transactions on Industrial Electronics, 58*(5), 1803–1811. https://doi.org/10.1109/TIE.2010.2046002

Hartmann, D., Herz, M., Paffrath, M., Rommes, J., Tamarozzi, T., Van der Auweraer, H., & Wever, U. (2021). Model order reduction and digital twins. In P. Benner, S. Grivet-Talocia, A. Quarteroni, G. Rozza, W. H. A. Schilders, & L. M. Silveira (Eds.), *Model order reduction. Volume 3: Applications* (pp. 379–430). De Gruyter. https://doi.org/10.1515/9783110499001-012

Hautala, S., Mikulski, M., Söderäng, E., Storm, X., & Niemi, S. (2022). Toward a digital twin of a mid-speed marine engine: From detailed 1D engine model to real-time implementation on a target platform. *International Journal of Engine Research*. https://doi.org/10.1177/14680874221106168

Huang, Q., Huang, R., Hao, W., Tan, J., Fan, R., & Huang, Z. (2019). Adaptive power system emergency control using deep reinforcement learning. *arXiv*. https://doi.org/10.48550/arXiv.1903.03712

IEEE. (2008). IEEE standard for a precision clock synchronization protocol for networked measurement and control systems (IEEE Std 1588-2008). Institute of Electrical and Eletronics Engineers. https://doi.org/10.1109/IEEESTD.2008.4579760

IEEE. (2011). *IEEE standard for a precision clock synchronization protocol for networked measurement and control systems* (IEEE Std 1588-2008, Rev. of IEEE Std 1588-2002). IEEE. https://doi.org/10.1109/IEEESTD.2008.4579760

International Electrotechnical Commission (IEC). (2019). IEC 61850: Communication networks and systems for power utility automation. IEC Standards. https://webstore.iec.ch

International Electrotechnical Commission. (2003). *IEC 62264-1: Enterprise-control system integration – Part 1: Models and terminology* [Standard]. IEC.

https://cdn.standards.iteh.ai/samples/35480/106d6c212cc9472bbc89d6a1450ffd7d/IEC-62264-1-2003.pdf

International Electrotechnical Commission. (2017). *IEC 62828-1:2017 – Reference conditions and procedures for testing industrial and process measurement transmitters – Part 1: General procedures*. https://cdn.standards.iteh.ai/samples/19833/8f1d959ca5474f759cf7bff1228e231c/IEC-62828-1-2017.pdf

International Electrotechnical Commission. (2020a). *IEC 62541-3: OPC Unified Architecture – Part 3: Address Space Model* [Standard]. IEC.

International Electrotechnical Commission. (2020b). *IEC 62541-5: OPC Unified Architecture – Part 5: Information Model* [Standard]. IEC.

iu, D., Chen, B., An, J., Wang, Z., Xu, C., Li, X., Liu, Y., Xu, J., & Wang, Z. L. (2020). Wind-driven self-powered wireless environmental sensors for Internet of Things at long distance. *Nano Energy, 73*, 104819. https://doi.org/10.1016/j.nanoen.2020.104819

Jensen, S. K., Pedersen, T. B., & Thomsen, C. (2017). *Time Series Management Systems: A Survey.* IEEE Transactions on Knowledge and Data Engineering, 29(11), 2581–2600. https://doi.org/10.1109/TKDE.2017.2740932

Grzesik, P., & Mrozek, D. (2020). Comparative analysis of time series databases in the context of edge computing for low power sensor networks. *Lecture Notes in Computer Science*, 12412, 340–354. https://doi.org/10.1007/978-3-030-50420-5_26

Kalman, R. E. (1994). Randomness reexamined. *Modeling, Identification and Control, 15*(3), 141–151. https://doi.org/10.4173/mic.1994.3.3

Karkaria, V., Goeckner, A., Zha, R., Chen, J., Zhang, J., Zhu, Q., Cao, J., Gao, R. X., & Chen, W. (2024). Towards a digital twin framework in additive manufacturing: Machine learning and Bayesian optimization for time-series process optimization. *Journal of Manufacturing Systems*. https://doi.org/10.1016/j.jmsy.2024.03.010

Klein, N., Guilfoyle, D., Karim, M. S., & McLaughlin, R. (2020). *Maritime autonomous vehicles: New frontiers in the law of the sea*. International & Comparative Law Quarterly, 69(3), 719–734. https://doi.org/10.1017/S0020589320000226

LaValle, S. M. (2011). Sensing and filtering: A fresh perspective based on information spaces. In *Motion Strategy in Information Spaces* (pp. 281–314). Cambridge University Press. https://lavalle.pl/papers/Lav11.pdf

Li, L., Xu, Y., Wang, Z., & Yan, J. (2022). Thermal power plant turbine rotor digital twin: Automation construction and monitoring system. *Journal of Mathematics*, 2022, 8527281. https://doi.org/10.1155/2022/8527281

Liu, R., Liu, Y., & Jing, Z. (2020). Impact of industrial virtual power plant on renewable energy integration. *Global Energy Interconnection, 3*(6), 545–552. https://doi.org/10.14171/j.2096-5117.gei.2020.06.004

Liu, X., Feng, X., Wang, B., Zeng, Z., Luo, G., & Li, H. (2022). Voltage optimization control strategy for islanded microgrid based on collaborative distributed MPC. *Frontiers in Energy Research, 10*, 880825. https://doi.org/10.3389/fenrg.2022.880825

Liu, Y., Zhang, P., Liu, J., et al. (2022). *Physical layer authentication in MIMO systems: a carrier frequency offset approach*. Wireless Networks, 28, 1909–1921. https://doi.org/10.1007/s11276-022-02916-y

Lorenzetti, J., McClellan, A., Farhat, C., & Pavone, M. (2020). Linear reduced-order model predictive control. *arXiv*. https://arxiv.org/abs/2012.03384

Luo, X., Liu, J., & Wang, J. (2023). A comprehensive review on control strategies for microgrids with distributed energy resources. *Wiley Interdisciplinary Reviews: Energy and Environment*, e12607. https://doi.org/10.1002/2050-7038.12607

Machalski, A., Cichoń, M., Twardochleb, M., & Gutowski, Ł. (2025). The concept of a digital twin for the Wały Śląskie Hydroelectric Power Plant: A case study in Poland. *Energies, 18*(8), 2021. https://doi.org/10.3390/en18082021

Maes, R. (2013). *Physically Unclonable Functions: Constructions, Properties and Applications*. Springer. https://doi.org/10.1007/978-3-642-41395-7

MAN Energy Solutions. (2023). *Learning from machines: Autonomous turbomachinery optimization*. https://www.man-es.com/discover/learning-from-machines

McBreen, J. (2024). *Matching 1-dimensional engine flow solution with digital twin requirements for system-level simulations: A model-in-the-loop study of an engine-*

based stationary power plant. University of Vaasa. https://urn.fi/URN:NBN:fi-fe2024060444777

Merat, N., Louw, T., Wilbrink, M., Schieben, A., & Cabrall, C. (2019). The “Out-of-the-Loop” concept in automated driving: Proposed definition, measures and implications. *Cognition, Technology & Work, 21*(1), 87–98. https://doi.org/10.1007/s10111-018-0525-8

Mikulski, M., et al. (2025). Clean propulsion technologies: Securing technological dominance for the Finnish marine and off-road powertrain sectors. *Energies, 18*(5), 1240. https://doi.org/10.3390/en18051240

Mooneyham, B. W., & Schooler, J. W. (2013). The costs and benefits of mind-wandering: A review. Canadian Journal of Experimental Psychology, 67(1), 11–18. https://doi.org/10.1037/a0031569

Munim, Z. H. (2022). Autonomous ships for container shipping in the Arctic routes. Journal of Marine Science and Technology, 27(4), 888–898. https://doi.org/10.1007/s00773-021-00836-8

National Institute of Standards and Technology. (2017). *Framework for cyber-physical systems: Volume 1, Overview (NIST SP 1500-100).* https://doi.org/10.6028/NIST.SP.1500-100

National Institute of Standards and Technology. (2021). *NIST Handbook 44: Specifications, tolerances, and other technical requirements for weighing and measuring devices* (2021 ed.). https://www.nist.gov/system/files/documents/2022/11/30/2023%20NIST%20Handbook%2044.pdf

Oliveira, A., Dias, A., Santos, T., Rodrigues, P., Martins, A., & Almeida, J. (2024). LiDAR-based unmanned aerial vehicle offshore wind blade inspection and modeling. *Drones, 8*(11), 617. https://doi.org/10.3390/drones8110617

Pasupuleti, S. (2025). *Digital twin technology for integration and optimization of manufacturing systems* [Preprint]. ResearchGate. https://www.researchgate.net/publication/391667377_Digital_Twin_Technology_for_Integration_and_Optimization_of_Manufacturing_Systems

Qin, S. J., & Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice, 11*(6), 733–764. https://doi.org/10.1016/S0967-0661(02)00186-7

Qiu, W., Tang, Q., Wang, Y., Zhan, L., Liu, Y., & Yao, W. (2020). *Multi-View Convolutional Neural Network for Data Spoofing Cyber-Attack Detection in Distribution Synchrophasors*. IEEE Transactions on Smart Grid, 11(4), 3457–3468. https://doi.org/10.1109/TSG.2020.2971529

Raja, V., Umapathy, M., Uma, G., & Usharani, R. (2023). Performance enhanced piezoelectric rotational energy harvester using reversed exponentially tapered multi-mode structure for autonomous sensor systems. *Journal of Sound and Vibration, 544*, 117429. https://doi.org/10.1016/j.jsv.2022.117429

Rawlings, J. B., Mayne, D. Q., & Diehl, M. (2020). *Model predictive control: Theory, computation, and design* (2nd ed.). Nob Hill Publishing.

Rehtanz, C. (2003). *Autonomous systems and intelligent agents in power system control and operation*. Springer. https://doi.org/10.1007/978-3-642-59355-5

ROTEC. (2022). *PRANA predictive analytics platform*. ROTEC Official Website. https://www.rotc.ru/en/prana/

Sajadi, A., Strezoski, L., Strezoski, V., Prica, M., & Loparo, K. A. (2019). Integration of renewable energy systems and challenges for dynamics, control, and automation of electrical power systems. *WIREs Energy and Environment, 8*(1), e321. https://doi.org/10.1002/wene.321

Sandia National Laboratories. (2022, May 11). Powering the moon: Sandia researchers design microgrid for future lunar base. https://newsreleases.sandia.gov/lunar_mircogrid

Satyanarayanan, M. (2017). The emergence of edge computing. *Computer*, 50(1), 30–39. https://doi.org/10.1109/MC.2017.9

Shi, W., & Dustdar, S. (2016). The promise of edge computing. *Computer*, 49(5), 78–81. https://doi.org/10.1109/MC.2016.145

Siemens Digital Industries Software. (2023). *Executable digital twin (xDT) – Simcenter*. https://plm.sw.siemens.com/en-US/simcenter/integration-solutions/executable-digital-twin/

Siemens. (2022). Siemens expands green energy potential with virtual power plants [Press release]. Siemens AG. https://press.siemens.com/global/en/pressrelease/siemens-expands-green-energy-potential-virtual-power-plants

Sleiti, A. K., Kapat, J. S., & Vesely, L. (2022). Digital twin in energy industry: Proposed robust digital twin for power plant and other complex capital-intensive large engineering systems. *Energy Reports, 8*, 3704–3726. https://doi.org/10.1016/j.egyr.2022.02.305

Tebekaemi, E., & Wijesekera, D. (2018). *A communications model for decentralized autonomous control of the power grid.* In 2018 IEEE International Conference on Communications (ICC) (pp. 1–6). IEEE. https://doi.org/10.1109/ICC.2018.8422911

Valkjärvi, P. (2022). *Integrated data acquisition for state-of-the-art large-bore engine test cell*. University of Vaasa. https://urn.fi/URN:NBN:fi-fe2022110264145

Valmet. (2021). *Valmet Annual Review 2021*. Valmet Corporation. Retrieved from https://www.valmet.com/globalassets/investors/reports--presentations/annual-reports/2021/valmet-annual-review-2021.pdf

Vidal, V. F., de Mello Honório, L., Pinto, M. F., Dantas, M. A. R., Aguiar, M. J. R., & Capretz, M. A. M. (2022). *An edge-fog architecture for distributed 3D reconstruction and remote monitoring of a power plant site in the context of 5G. Sensors, 22*(12), 4494. https://doi.org/10.3390/s22124494

Wang, X., Mutlu, I., Rani, F., Drowatzky, L., & Urbas, L. (2022). *A comparative study to evaluate the performance of communication protocols for process industry*. In 2022 32nd International Telecommunication Networks and Applications Conference (ITNAC) (pp. 170–177). IEEE. https://doi.org/10.1109/ITNAC55475.2022.9998327

Wärtsilä. (2023a). *Expert Insight service*. Retrieved from https://www.wartsila.com/energy/services/lifecycle-solutions/remote-support-and-data-management

Wärtsilä. (2023b). *Remote support and data management / lifecycle solutions*. Retrieved from https://www.wartsila.com/energy/services/lifecycle-solutions/remote-support-and-data-management

Wärtsilä. (2024). *WISE / Name / Title 23.4.2024*. WISE Ecosystem. https://www.wiseecosystem.com/wp-content/uploads/2024/04/WISE_Webpage.pdf

Wärtsilä. (2024, February 16). *Wärtsilä-led EUR 200 million collaboration ecosystem to develop autonomous zero-emission balancing solutions*. Retrieved from https://www.wartsila.com/media/news/16-02-2024-wartsila-led-eur-200-million-collaboration-ecosystem-to-develop-autonomous-zero-emission-balancing-solutions-for-the-energy-transition-3404703

Wu, Y., Wu, G., Ma, C., & Wang, Q. (2023). A heterogeneous information fusion method for maritime radar and AIS based on D-S evidence theory. *Journal of Signal and Information Processing, 14*(2), 24–40. Scientific Research Publishing (SCIRP). https://doi.org/10.4236/jsip.2023.142002

Yıldırım, B., Erdal, H., Cakir, A., & Akbulut, A. (2025). A comprehensive review of edge computing for power systems: Architectures, applications, and challenges. *Applied Sciences*, 15(8), 4592. https://doi.org/10.3390/app15084592

Yu, P., Wang, Z., Zhang, H., & Song, Y. (2024). Safe reinforcement learning for power system control: A review. *IEEE Access, 12*, 34800–34819. https://doi.org/10.1109/ACCESS.2024.10457539

Zahra, A., & Singh, B. (2022). A review of decentralized and distributed control approaches for microgrids and distributed energy resources. *Electric Power Systems Research, 203*, 107861. https://doi.org/10.1016/j.epsr.2022.107861

Zhang, H., Xu, C., Zhang, R., Wang, S., & Li, H. (2021). Data-driven model predictive and reinforcement learning-based control for building energy management: A survey. *arXiv:2106.14450*. https://arxiv.org/abs/2106.14450

Zhang, Q., Chen, M., & Li, Y. (2021). Reinforcement learning-based hybrid energy control for autonomous microgrids. *IEEE Transactions on Industrial Electronics, 68*(4), 3441–3452. https://doi.org/10.1109/TIE.2020.2992961

Zhang, Q., Zhang, B., Jin, L., Jiang, W., Han, X., Zhang, Z., & Liu, Y. (2023). Self-powered ammonia sensors based on TENGs for gas leakage monitoring in marine engine rooms. *Sensors, 23*(3), 1329. https://doi.org/10.3390/s23031329

Zhang, Y., & Liu, Y. (2019). Real-time data capture and alignment for digital twin based dynamic modeling of manufacturing systems. *IEEE Transactions on Industrial Informatics, 15*(4), 2253–2261. https://doi.org/10.1109/TII.2018.2873186

Zhu, H.-H., Liu, W., Wang, T., Su, J.-W., & Shi, B. (2022). Distributed acoustic sensing for monitoring linear infrastructures: Current status and trends. *Sensors, 22*(19), 7550. https://doi.org/10.3390/s22197550

Zürn, M., & Faude, B. (2013). On fragmentation, differentiation, and coordination. *Global Environmental Politics, 13*(3), 119–130. https://doi.org/10.1162/GLEP_a_00183