Cyber-Physical System Integration for Real-Time Monitoring and Control in Industrial Engineering

Pedro R. Dunn
School of Computing, Clemson University, Clemson, SC, USA.
dunnpedro@clemson.edu

Pradeep C. Mukherjee
Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
pradeep.mukherjee931@missouri.edu

Francesco Willis
Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
willis736@colostate.edu

Bastian Riley
Department of Computer Science, University of New Hampshire, Durham, NH, USA.
bastianr@unh.edu

Abstract

The convergence of operational technology and information technology through cyber-physical system integration is reshaping industrial engineering, enabling real-time monitoring and closed-loop control across complex manufacturing and process environments. This paper presents a system-level examination of the architectural paradigms, communication infrastructures, governance frameworks, and deployment models that underpin such integrated systems. We discuss the evolution from rigid hierarchical automation pyramids to distributed, service-oriented architectures where edge intelligence, time-sensitive networking, and cloud-based analytics coexist. The analysis emphasises structural trade-offs between centralised coordination and decentralised autonomy, between determinism and flexibility, and between security and real-time performance. Special attention is given to the lifecycle implications of large-scale integration, including resilience engineering, sustainability metrics, and socio-technical governance. Without resorting to formal equations or algorithmic representations, the paper develops a conceptual apparatus that connects industrial real-time control to broader concerns of fairness, explainability, and regulatory alignment. The discussion is illustrated through cross-domain comparisons drawn from discrete manufacturing, continuous process industries, and energy systems, revealing common patterns of complexity that transcend sectoral boundaries. By synthesising recent advances in middleware design, network slicing, digital twin orchestration, and standardisation efforts such as the Asset Administration Shell, the paper provides a forward-looking perspective on how industrial cyber-physical systems can be designed not only for efficiency but also for robustness, sustainability, and equitable value distribution. The conclusion identifies critical research gaps, particularly around verifiable safety in multi-stakeholder environments and the integration of circular economy principles into real-time control loops.

Keywords

cyber-physical systems; industrial engineering; real-time monitoring; control systems; system integration; edge-cloud architecture; time-sensitive networking; digital twin; governance; resilience.

1. Introduction

Industrial engineering has long relied on layered automation architectures that segregate physical processes, control loops, supervisory systems, and enterprise planning. The emergence of cyber-physical systems (CPS) blurs these boundaries by embedding computation, networking, and storage directly into physical assets, enabling continuous real-time observation and actuation across entire value chains [1], [2]. While early conceptions of CPS focused on tightly bounded applications such as engine control units or robotic workcells, the current vision encompasses plant-wide and even supply-chain-wide integration, driven by the availability of low-cost sensors, high-bandwidth wireless communication, and scalable cloud analytics [3]. This shift promises substantial gains in productivity, quality, and energy efficiency, yet it simultaneously introduces unprecedented complexity that challenges traditional industrial engineering methodologies.

The integration of CPS for real-time monitoring and control demands a systemic perspective that goes beyond component-level optimisation. In a typical discrete manufacturing setting, for instance, a programmable logic controller may execute a hard real-time task to position a cutting head, while a higher-level supervisory control and data acquisition (SCADA) node monitors tool wear and an edge analytics engine predicts remaining useful life. Integrating these layers into a coherent system requires careful consideration of timing semantics, data consistency models, failure modes, and the division of responsiblity among human operators, automated controllers, and business optimisation algorithms [4]. Similar challenges appear in continuous process industries, where thermal dynamics and fluid flows impose strict latency bounds on control adjustments, and in energy systems, where distributed generation units must coordinate in near-real time to maintain grid stability. Across all these domains, the central question is not merely whether integration is technically feasible, but how to design integration architectures that are scalable, evolvable, trustworthy, and aligned with broader societal objectives.

This paper addresses that question through a multi-level analysis. We proceed from architectural foundations to communication infrastructures, then examine system-level trade-offs in monitoring and control. Subsequent sections explore governance, resilience, sustainability, deployment models, and the often-overlooked issues of fairness and policy. Throughout, the discussion remains at a conceptual and systems level, intentionally avoiding mathematical formalisms to foreground structural reasoning and design rationale. The intent is to equip industrial engineers, systems architects, and policy makers with a coherent framework for reasoning about CPS integration as a socio-technical undertaking that intertwines physics, computation, organisation, and regulation.

2. Architectural Paradigms for Industrial CPS Integration

The traditional functional hierarchy of industrial automation, often depicted as a pyramid with field devices at the base and enterprise resource planning at the apex, presupposes a strictly vertical flow of information. Real-time monitoring and closed-loop control, however, increasingly demand horizontal interactions across devices and subsystems, giving rise to architectures that resemble distributed service meshes rather than monolithic stacks [5]. Three complementary paradigms characterise current integration efforts: the continuation of layered

control with enhanced interconnectivity, the shift toward edge-centric computing fabrics, and the adoption of asset-oriented information models that decouple data semantics from execution platforms.

In the layered approach, standardised fieldbus protocols such as EtherCAT and PROFINET still guarantee deterministic periodic communication within a control island, while publish-subscribe middleware like the Data Distribution Service enables cross-island data sharing with configurable quality-of-service policies [6]. This preserves the real-time guarantees of local loops while opening pathways for supervisory analytics and fleet-level optimisation. The architectural challenge lies in managing the semantic gap between time-triggered control loops and event-driven enterprise services, a gap that often necessitates gateway nodes performing protocol translation and temporal decoupling. Such gateways become critical points of reliability concern and potential bottlenecks, demanding rigorous failure mode analysis and redundancy design.

A second paradigm pushes significant computation to the edge, close to physical processes, using heterogeneous hardware ranging from field-programmable gate arrays to embedded general-purpose processors running containerised microservices. Edge-centric architectures reduce backhaul latency, limit the volume of raw data sent to cloud data centres, and allow local loops to maintain partial functionality during connectivity outages [7]. The trade-off, however, is that edge nodes have constrained computational and energy budgets, which limit the sophistication of local analytics and constrain the deployment of advanced machine learning models. Partitioning workloads across edge, fog, and cloud layers therefore becomes a core architectural decision that must account for the dynamics of data rates, inference accuracy requirements, and the cost of model updates. This partitioning problem has stimulated considerable research into adaptive orchestration frameworks that dynamically migrate tasks based on load, available energy, and network conditions.

The third paradigm centres on information modelling standards such as the Asset Administration Shell promoted by the Platform Industrie 4.0 and comparable initiatives worldwide. These models abstract physical assets as digital representations with well-defined submodels that capture identification, condition monitoring, capability, and compliance information [8]. By standardising the semantics of asset data, such frameworks enable integration across vendor-specific controllers and legacy equipment without requiring wholesale replacement. They also lay the foundation for digital twins that co-evolve with physical counterparts throughout the asset lifecycle, supporting applications from commissioning to decommissioning. The effectiveness of such models depends critically on governance processes that keep the digital representations synchronised with physical reality, a challenge that grows with the scale and heterogeneity of the deployed asset base.

3. Communication Infrastructures and Real-Time Data Flow

Real-time monitoring and control are fundamentally contingent on the ability of communication networks to deliver data with bounded latency, controlled jitter, and high reliability. Industrial environments, however, are characterised by electromagnetic interference, multipath propagation, and mechanical obstructions that degrade wireless links, while wired solutions, though more predictable, constrain physical topology and reconfiguration [9]. The ongoing convergence of operational technology networks with standard Ethernet and fifth-generation cellular technologies heralds a new era in which deterministic behaviour can be provisioned on shared infrastructure, yet it also introduces new

dimensions of complexity related to resource isolation, security policy enforcement, and lifecycle management.

Time-sensitive networking (TSN), standardised by the IEEE 802.1 working group, extends conventional Ethernet with mechanisms such as time-aware shaping, frame preemption, and stream reservation to guarantee worst-case latency bounds for critical traffic classes while coexisting with best-effort enterprise traffic on the same physical links [10]. TSN enables a converged network that can simultaneously carry hard real-time control messages, video surveillance streams, and bulk data transfers for cloud-based analytics. The configuration of TSN networks, however, is a non-trivial combinatorial scheduling problem, especially in large-scale brownfield installations where existing wiring and traffic patterns impose constraints. Automating TSN configuration through centralised network management controllers that compute gate control lists offline and distribute them via NETCONF or RESTCONF interfaces is an active area of standardisation and research.

Wireless technologies are equally transformative, particularly for mobile assets such as automated guided vehicles, rotating equipment, and wearable devices used by maintenance personnel. Industrial-grade Wi-Fi 6, private 5G networks, and emerging Wi-Fi 7 solutions promise microsecond-level synchronisation and ultra-reliable low-latency communication through features like uplink scheduling, multi-access edge computing, and network slicing [11]. A private 5G network, for instance, can dedicate a slice with guaranteed bandwidth and latency characteristics to a robotic assembly cell, while another slice serves augmented reality terminals for remote expert assistance. The operational challenge is to ensure that the configuration of radio resources, core network functions, and edge computing nodes remains consistent with the safety integrity levels required by the application. This requires a tight coupling between the orchestration of communication slices and the orchestration of industrial workloads, a domain where cross-disciplinary expertise is still scarce.

Beyond the physical and link layers, higher-level protocols determine how data is structured, timestamped, and distributed. The OPC Unified Architecture (OPC UA) has evolved from a simple access protocol into a comprehensive framework supporting publish-subscribe patterns, deterministic transport over TSN, and companion specifications for vertical domains [12]. OPC UA’s information modelling capabilities allow it to serve as a semantic bridge between the asset administration shell and real-time data streams, thereby reducing integration overhead. Nevertheless, the coexistence of multiple middleware protocols in a single plant, driven by equipment vendor preferences and legacy constraints, inevitably creates islands of interoperability. Designing integration fabrics that can seamlessly stitch together OPC UA, DDS, MQTT, and legacy protocols while preserving end-to-end timing guarantees remains a significant systems engineering challenge.

4. System-Level Trade-Offs in Monitoring and Control Design

The design of an integrated CPS for industrial monitoring and control is inherently a multi-objective exercise in which performance, reliability, security, and cost must be balanced against one another. One fundamental trade-off involves the tension between centralised and distributed control architectures. Centralised approaches, exemplified by a cloud-hosted model predictive controller that optimises an entire production line, can exploit global knowledge to achieve near-optimal efficiency. However, they introduce a single point of failure and are vulnerable to communication delays and denial-of-service attacks [13]. Distributed architectures, in which local controllers collaborate through peer-to-peer coordination, offer greater resilience and lower latency but may converge to suboptimal

operating points or exhibit oscillatory behaviour if coordination protocols are not carefully designed.

A related trade-off concerns the granularity and freshness of monitoring data. Increasing sensor density and sampling frequency improves state estimation accuracy and enables finer defect detection, yet it also elevates communication load, storage costs, and energy consumption. In large-scale facilities, many sensing points cannot be continuously powered and must rely on energy-harvesting techniques, creating a direct coupling between data freshness and the available energy budget. System designers must therefore make deliberate choices about which variables to monitor at high resolution, which can be sampled intermittently, and how data compression and event-triggered communication schemes can reduce resource demands without compromising control quality.

The security dimension introduces additional tensions. Hard real-time systems have historically been isolated on physically separate networks, a practice that provided a de facto security boundary. Integration with enterprise IT and cloud services erodes this boundary, exposing control devices to a wider attack surface. The adoption of strong cryptographic authentication and encryption can protect data integrity and confidentiality, but cryptographic operations add latency and jitter that may violate control loop deadlines [14]. Emerging lightweight security protocols and hardware-accelerated encryption attempt to mitigate this overhead, but a residual trade-off persists between the level of security assurance and the achievable real-time performance. Balancing these factors requires a risk-informed approach that matches security mechanisms to the criticality of each data flow and control action.

Fairness and resource allocation constitute a further systemic tension. In a shared computing and networking infrastructure, multiple control applications compete for bounded resources, and scheduling policies must adjudicate among them. A purely strict-priority scheme risks starving lower-criticality functions, while a purely fair-share approach may fail to meet the hard deadlines of safety-critical loops. Hybrid schemes that combine time-triggered and priority-based scheduling with admission control and runtime monitoring are necessary to provide differentiated guarantees while maintaining high resource utilisation.

5. Governance, Resilience, and Sustainability

The integration of cyber-physical systems at industrial scale necessitates a governance framework that extends beyond technical configuration to encompass organisational responsibilities, lifecycle processes, and alignment with environmental and social objectives. Governance in this context refers to the set of policies, standards, auditing mechanisms, and incentives that guide how systems are designed, operated, maintained, and eventually retired. Effective governance ensures that the behaviour of an integrated CPS is not only efficient in steady state but also resilient in the face of perturbations, sustainable over extended time horizons, and accountable to human stakeholders.

Resilience engineering provides a useful lens for examining how industrial CPS can gracefully handle unexpected disruptions, whether they stem from equipment failures, cyberattacks, or extreme environmental events. Rather than attempting to enumerate all possible failure modes and design specific barriers, a resilience-oriented approach focuses on building capacities for monitoring, responding, learning, and anticipating [15]. In practice, this means designing systems that can degrade functionality in a controlled manner when resources are constrained, that provide operators with clear situation awareness during transients, and that support rapid recovery through pre-planned reconfiguration strategies.

Microservice architectures, for instance, allow individual control functions to be independently restarted or relocated without bringing down the entire system, thus enhancing overall graceful degradation. However, such architectures also multiply the number of inter-component dependencies that must be monitored, raising the bar for observability tooling and anomaly detection.

Sustainability considerations are becoming increasingly central to industrial CPS design. Real-time monitoring enables fine-grained energy metering and load scheduling that can reduce the carbon footprint of production processes, yet the digital infrastructure itself consumes energy and material resources throughout its lifecycle [16]. A holistic sustainability analysis must account for the embodied energy of edge devices and network equipment, the operational energy consumed by data centres processing industrial telemetry, and the end-of-life disposal or recycling of electronic waste. The concept of circularity can be embedded into CPS monitoring loops by tracking material flows and equipment health in ways that facilitate remanufacturing and reuse, aligning operational control objectives with broader circular economy goals. Such integration requires aligning real-time production schedules with the availability of refurbished components and the energy availability profiles of on-site renewable generation.

The governance of integrated CPS also raises questions of algorithmic transparency and human agency. As machine learning models are increasingly deployed within control loops for tasks such as predictive maintenance or quality inspection, their opacity can undermine trust and complicate safety certification. Even when model accuracy is high, operators may hesitate to rely on recommendations they do not understand, particularly in safety-critical contexts where erroneous predictions can cause costly shutdowns or hazards. Explainability techniques that generate post-hoc justifications can partially address this issue, but they must be carefully integrated into operational dashboards without overwhelming operators with extraneous information. Moreover, governance structures must define clear protocols for human override, especially when automated decisions conflict with experiential knowledge or ethical considerations.

6. Deployment Models and Scalability Challenges

Bridging the gap between laboratory prototypes and production-grade deployments of integrated CPS remains a formidable challenge. Industrial environments are rarely greenfield sites; instead, they consist of heterogeneous fleets of legacy equipment, each with its own communication interfaces, data formats, and maintenance histories [17]. Brownfield integration strategies must therefore provide incremental pathways for adding connectivity and intelligence without disrupting existing production flows or invalidating safety certifications. Retrofitting legacy machines with external sensor suites and edge gateways is a widely adopted tactic, but it introduces additional latency layers and can create inconsistencies between the data reported by the machine's internal controller and the externally sensed data. Managing these semantic and temporal inconsistencies requires careful data fusion and timestamp reconciliation, often mediated by a local fog node that serves as a digital proxy for the physical asset.

Scalability is not solely a matter of hardware capacity; it is equally a challenge of configuration management and model maintenance. A plant with thousands of sensors and hundreds of control loops generates an immense combinatorial space of possible interactions, and the human effort required to manually engineer each integration point quickly becomes unsustainable. Model-based engineering approaches, supported by domain-specific languages

and automated code generation, can alleviate this burden by enabling system designers to specify structural and behavioural models that are automatically translated into deployable configurations [18]. However, these models must be kept consistent with the evolving physical plant, necessitating continuous model validation pipelines and version-controlled digital twins that can be subjected to virtual commissioning and what-if analysis before changes are pushed to production.

Data governance adds another dimension of scalability concern. Real-time monitoring generates torrents of time-series data that must be ingested, stored, and analysed under strict latency constraints. Adopting a data lakehouse architecture that combines the schema flexibility of data lakes with the transactional guarantees of data warehouses can provide a unified platform for both operational analytics and historical batch analysis [19]. Yet such architectures must be carefully tuned to the specific access patterns of industrial workloads, which often involve windowed aggregations, anomaly detection over streaming data, and federated queries across geographically distributed plants. The cost of data egress from multiple sites to a central cloud can be prohibitive, incentivising a hybrid model in which initial pre-processing and noise reduction are performed on edge nodes, and only aggregated or exceptional data is transmitted to central repositories.

The human dimension of deployment is frequently underestimated. Successful integration projects require close collaboration between control engineers, IT network architects, data scientists, and plant operators, each with distinct mental models and vocabularies. Organisational silos can impede the flow of knowledge and lead to integration failures that are not purely technical in nature but stem from misaligned incentives or insufficient training. Investing in cross-functional teams, common visualisation platforms, and collaborative digital engineering environments is essential to building the socio-technical fabric that sustains large-scale CPS integration over time.

7. Fairness, Trust, and Policy Implications

While much of the technical literature on industrial CPS focuses on performance optimisation, the societal dimensions of real-time monitoring and control merit equally rigorous examination. As integrated systems become more pervasive, decisions that were once made by skilled operators are increasingly delegated to algorithms, raising concerns about the fairness of those decisions, the distribution of their costs and benefits, and the erosion of human expertise [20]. For example, an optimisation algorithm that prioritises throughput might systematically delay maintenance tasks on less profitable product lines, thereby concentrating equipment degradation in areas that disproportionately affect certain worker groups or communities. Such feedback loops are not captured by conventional engineering metrics, yet they can have tangible consequences for occupational safety, job quality, and regional economic stability.

Trust in automated systems is a multi-faceted construct involving reliability, predictability, transparency, and value alignment. In industrial settings, trust is built over time through repeated interactions and shared experiences, but the introduction of black-box models can disrupt this process. Policy interventions that mandate minimum levels of explainability, audit trails of control actions, and independent third-party certification of safety-critical components can help re-establish a basis for trust. However, these requirements must be calibrated to avoid imposing disproportionate compliance burdens that stifle innovation or disproportionately disadvantage small and medium-sized enterprises. Finding the appropriate

regulatory balance is a delicate task that calls for close dialogue between industry, standardisation bodies, labour representatives, and government agencies.

The international dimension of policy further complicates governance. Supply chains span multiple jurisdictions with divergent regulations on data privacy, cybersecurity, and environmental protection. A CPS deployment that processes worker location data for safety monitoring in one country may run afoul of privacy laws in another, creating complex legal exposures. International harmonisation efforts, such as those under way through the International Electrotechnical Commission and the Institute of Electrical and Electronics Engineers, aim to produce globally acceptable reference architectures and conformance profiles, but the pace of standardisation often lags behind the rapid evolution of technology. Forward-looking firms are therefore developing internal governance frameworks that are robust to regulatory fragmentation, embedding privacy-by-design principles and ethical risk assessments into their systems engineering processes.

Looking ahead, the integration of industrial CPS with energy markets and carbon trading systems will further entangle engineering decisions with economic and policy considerations. Real-time control that adjusts production rates in response to dynamic energy prices or carbon intensity signals can improve sustainability, but it also exposes manufacturers to new financial risks and regulatory dependencies. Designing control policies that are both economically optimal and robust to policy shifts requires a combination of stochastic optimisation, scenario planning, and flexible contractual arrangements, a domain where interdisciplinary research is urgently needed.

8. Conclusion

The integration of cyber-physical systems for real-time monitoring and control in industrial engineering represents a profound shift in how physical production processes are conceived, managed, and governed. This paper has argued that successful integration demands a system-level perspective that transcends individual technological components and addresses architectural design, communication infrastructure provisioning, governance mechanisms, and broader socio-technical dynamics. We have examined the trade-offs inherent in centralised versus distributed control, in deterministic versus best-effort networking, and in security versus real-time determinism, highlighting that no single architectural pattern is universally optimal. Rather, effective integration results from a deliberate, risk-informed process of matching architectural choices to application-specific requirements, legacy constraints, and long-term strategic objectives.

Resilience, sustainability, and fairness have emerged as critical yet underexplored dimensions of industrial CPS design. As systems become more interconnected and algorithmically driven, the importance of building adaptive capacity, minimising lifecycle environmental impact, and ensuring equitable outcomes will only grow. These objectives cannot be achieved through technical means alone; they require the evolution of organisational practices, regulatory frameworks, and professional education to cultivate a workforce capable of navigating the interplay between automation and human judgment.

Looking forward, several research frontiers warrant attention. The verification of safety properties in dynamically composed, multi-tenant edge environments remains an open problem. The integration of circular economy metrics into real-time control objectives calls for new optimisation frameworks that reconcile economic, environmental, and social performance indicators. The governance of federated learning across competing industrial

actors raises nuanced questions about intellectual property protection and data sovereignty. Addressing these challenges will require sustained collaboration across the disciplines of systems engineering, computer science, industrial ecology, law, and the social sciences, reinforcing the inherently interdisciplinary nature of cyber-physical system integration.

References

1. Monostori, L., Kádár, B., Bauernhansl, T., Kondoh, S., Kumara, S., Reinhart, G., ... & Ueda, K. (2016). Cyber-physical systems in manufacturing. CIRP Annals, 65(2), 621–641.

2. Rajkumar, R., Lee, I., Sha, L., & Stankovic, J. (2010). Cyber-physical systems: The next computing revolution. In Proceedings of the 47th Design Automation Conference (pp. 731–736). ACM.

3. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405–2415.

4. Derler, P., Lee, E. A., & Sangiovanni-Vincentelli, A. (2012). Modeling cyber–physical systems. Proceedings of the IEEE, 100(1), 13–28.

5. Colombo, A. W., Karnouskos, S., Kaynak, O., Shi, Y., & Yin, S. (2017). Industrial cyberphysical systems: A backbone of the fourth industrial revolution. IEEE Industrial Electronics Magazine, 11(1), 6–16.

6. Pfrommer, J., Grüner, S., & Palm, F. (2016). Hybrid OPC UA and DDS: Combining the best of two worlds to provide a service-oriented middleware for the Industrial Internet of Things. In Proceedings of the 21st IEEE International Conference on Emerging Technologies and Factory Automation (pp. 1–8). IEEE.

7. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637–646.

8. Plattform Industrie 4.0. (2019). The Asset Administration Shell: Implementing digital twins for use in Industrie 4.0. Federal Ministry for Economic Affairs and Energy (BMWi).

9. Ferrari, P., Flammini, A., Rinaldi, S., & Sisinni, E. (2021). On the use of LoRaWAN for the monitoring and control of distributed energy resources in industrial applications. IEEE Transactions on Industrial Informatics, 17(12), 8421–8430.

10. Nasrallah, A., Thyagaturu, A. S., Alharbi, Z., Wang, C., Shao, X., Reisslein, M., & ElBakoury, H. (2019). Ultra-low latency (ULL) networks: The IEEE TSN and IETF DetNet standards and related 5G ULL research. IEEE Communications Surveys & Tutorials, 21(1), 88–145.

11. Aijaz, A. (2020). Private 5G: The future of industrial wireless. IEEE Industrial Electronics Magazine, 14(4), 136–145.

12. Leitner, S. H., & Mahnke, W. (2006). OPC UA – Service-oriented architecture for industrial applications. Softwaretechnik-Trends, 26(3), 19–30.

13. Mittal, S., & Tolk, A. (2019). Complexity challenges in cyber physical systems: Using modeling and simulation (M&S) to support intelligence, adaptation and autonomy. In Proceedings of the Winter Simulation Conference (pp. 1450–1463). IEEE.

14. Sinha, R. S., Wei, Y., & Hwang, S. H. (2017). A survey on LPWA technology: LoRa and NB-IoT. ICT Express, 3(1), 14–21.

15. Hollnagel, E., Woods, D. D., & Leveson, N. (Eds.). (2006). Resilience engineering: Concepts and precepts. Ashgate Publishing.

16. Kusiak, A. (2018). Smart manufacturing must embrace big data. Nature, 544(7648), 23–25.

17. Fleischmann, H., Kohl, J., & Franke, J. (2016). A modular architecture for the design of condition monitoring processes. Procedia CIRP, 57, 410–415.

18. Vogel-Heuser, B., Fay, A., Schaefer, I., & Tichy, M. (2015). Evolution of software in automated production systems: Challenges and research directions. Journal of Systems and Software, 110, 54–84.

19. Armbrust, M., Ghodsi, A., Xin, R., & Zaharia, M. (2021). Lakehouse: A new generation of open platforms that unify data warehousing and advanced analytics. In Proceedings of the 11th Conference on Innovative Data Systems Research (pp. 28–40). CIDR.

20. Müller, J. M., Kiel, D., & Voigt, K. I. (2018). What drives the implementation of Industry 4.0? The role of opportunities and challenges in the context of sustainability. Sustainability, 10(1), 247.