



(REVIEW ARTICLE)



Multi-agent orchestration architectures for real-time enterprise decision intelligence systems

Swaroop Suresh Borukar *

Independent Researcher, Workday Inc., San Jose, CA.

World Journal of Advanced Engineering Technology and Sciences, 2026, 19(03), 176-184

Publication history: Received on 05 May 2026; revised on 13 June 2026; accepted on 16 June 2026

Article DOI: <https://doi.org/10.30574/wjaets.2026.19.3.0319>

Abstract

The architectures that enable the coordination of distributed analytics, enterprise events, machine agents, human judgment, and governance constraints under low latency conditions are becoming more common for real-time enterprise decision intelligence systems. Multi-agent orchestration is an architectural approach that distributes intelligence across software agents, which perform the functions of sensing, reasoning, negotiating, executing and explaining tasks or actions, instead of having it centralized in a single analytical core. In this review, we focus on peer-reviewed journal articles published since 2015 that are directly relevant to multi-agent orchestration architecture in real-time enterprise decision intelligence systems. The literature reviewed shares a number of common themes: cyber-physical decentralization, event-driven analytics, hybrid human-AI decision structures, explainability, and digital-twin mediated coordination. Although reported studies demonstrate the benefits of agent-based orchestration in terms of adaptability, local responsiveness, and resilience, agent-based orchestration still faces several challenges, including interoperability, verification, cross-agent accountability, organizational integration, and empirical evaluation at enterprise scale. The field is important, because future decision intelligence systems demand predictive accuracy, in addition to coordinated, explainable, auditable, and context-sensitive decision execution in complex enterprise environments.

Keywords: Agent orchestration; Decision intelligence; Digital twins; Enterprise analytics; Multi-agent systems; Real-time architecture

1. Introduction

Multi-agent orchestration architectures for real-time enterprise decision intelligence systems are at a technical intersection of distributed AI, enterprise analytics, cyber-physical systems and organizational decision design, and constitute a technically focused yet increasingly significant research area. Multi-agent systems offer coordination schemes for distributed software entities that observe local environments, share knowledge, negotiate actions, and adapt to environmental changes [1]. In enterprise decision intelligence, such capabilities are relevant because decisions need to be made in real-time, from multiple sources, on various platforms, with a volatile supply chain, cyber-physical devices, and human supervision. The architectural challenge comes not only in the form of selecting the right model for analysis, but also in the way the different decision agents need to act in a way that accounts for latency, regulatory, and organizationally contextual constraints.

The technical basis for this topic has been provided by manufacturing and industrial systems research, as it is in these environments that distributed sensing, low latency control, and adaptive coordination between the computational and physical assets is necessary [2]. However, enterprise decision intelligence extends beyond the factory floor. It comprises risk management in the supply chain, resource prioritization for maintenance, resource allocation, demand response,

* Corresponding author: Swaroop Suresh Borukar

compliance triage, and executive decision support. However, value generated by the capability of analytics in such environments depends on its incorporation into processes, roles, platforms and performance logic [3]. Therefore, multi-agent orchestration should be interpreted as an architectural method to connect analytical intelligence with coordinated enterprise action, not as an AI technique.

There is a significant gap between the current research landscape of technical agent architectures and organizational decision studies. Technical studies focus on, for example, agent communication, decentralization, reconfigurability, cyber-physical integration, and that of information systems and management studies on, for example, business value, governance, human-AI collaboration, and decision rights. The difference is significant because real-time enterprise decision intelligence systems need both layers. A system of orchestration that is technically efficient can be unsuccessful if there is a lack of accountability, interpretability, escalation logic, or managerial integration. On the other hand, an analytics system focused on governance would be too slow and too centralized in dynamic environments. This is a structural tension that is confirmed by organizational decision-making research in the era of artificial intelligence, which reveals that automation is not just about the speed of the decision, but also about the delegation, supervision, and control of decision-making [4].

A highly targeted journal corpus which covers the entire title is still limited. Decision intelligence is less well defined in peer-reviewed journals compared to other terms like business analytics, cyber-physical systems, smart manufacturing, artificial intelligence decision support, digital twins, and hybrid intelligence. Therefore, only peer-reviewed journal articles published from 2015 onward were considered for this work that directly impact the architecture/orchestration, real-time analytics and enterprise decision dimensions of the topic. The intent is not to offer a general coverage of the subject of AI in organizations. The aim is to critically analyze the literature that is most relevant to multi-agent orchestration architectures for real-time enterprise decision intelligence systems, compare reported results, highlight methodological limitations and clarify the research priorities into scalable, trustworthy and adaptive decision infrastructures.

2. Literature Review

The existing research related to multi-agent orchestration architectures of real-time enterprise decision intelligence systems can be categorized into four overlapping categories: agent-based cyber-physical coordination, smart manufacturing and real-time data pipelines, enterprise analytics value creation and human-centred decision-governance. The first cluster sets the architectural foundation. According to smart-agent studies, agent-based cyber-physical systems can enhance the modularity, local decision capabilities, self-configuration and responsiveness of the systems in industrial environments [5]. While related work in cyber-physical automation shows technical feasibility of prototype implementations, it also shows that there are ongoing barriers to standardization and interoperability, migration from legacy control systems, and certification [6]. Distributed sensing, cloud connection, and intelligent control, are also identified as enablers for adaptive production in smart factory research, but are usually considered as one-off studies or at the research scale without attention to enterprise-scale deployment [7]. In contrast, Kusiak's representation of smart manufacturing is more integrative, presenting analytics, sensors, optimization, and enterprise connectivity as a system level agenda and not a single technology stack [8].

The second cluster deals with real-time enterprise data flow. Industrial analytics pipelines illustrate how streaming infrastructure can support predictive maintenance and operational analysis when data ingestion, cleaning, storage, and modeling are coordinated as a pipeline rather than treated as sequential, stand-alone modeling processes. This has direct implications on multi-agent orchestration as agents need to have timely and context-rich state representations. But, when it comes to pipeline studies, data engineering and model execution are often the focus rather than agent negotiation, distributed autonomy, or cross-agent conflict resolution. The enterprise analytics literature also demonstrates that analytics capability has a positive impact on firm performance only in combination with strategy, dynamic capability, and process redesign [10, 11]. Strategic value frameworks also warn that data volume and analytical sophistication do not guarantee the value of the decisions; decisions must be integrated and governed in the organization.

The third cluster relates to contexts where decisions are riskier to make, particularly in the context of supply chain risk. AI research focused on supply-chain risk management reveals potential prediction, classification, planning and disruption response capabilities, but also weak validation across complex operational networks and fragmented evidence [13]. These results are relevant to the topic, as multi-agent orchestration is most useful if there are multiple actors, constraints and objectives that can interact in real time. However, there are numerous articles in the supply-chain AI space which focus on algorithms, with questions remaining regarding the role of agents, protocols for coordination, and who is responsible for what. The fourth cluster is about the collaboration between humans and AI.

Research on human–AI symbiosis suggests that organizational decision-making processes can be enhanced by bringing together computational pattern recognition with human contextual judgment [14]. Opaque AI systems are not suitable for accountable enterprise decision intelligence when decisions involve ethics or customers or have strategic impact [15]. Hybrid intelligence research considers the problem in the light of joint human-machine cognitive performance, and it is appropriate to the architectures of orchestration that involve escalation agents, explanation agents, and human approval paths [16].

Selected studies from the corpus are summarized in Table 1. These findings are presented in the table and reveal that the highest architectural evidence is found in industrial cyber-physical journals, and the highest organizational evidence is found in the analytics, information systems and decision-governance journals. The literature as a whole corroborates the topic while also demonstrating that there is a fragmentation of technical and managerial traditions in the literature.

Table 1 Summary of key findings

Ref	Focus	Key Findings
[5]	Smart agents in industrial cyber-physical systems	Agent-based control improves modularity, local responsiveness, and distributed coordination, but industrial validation remains constrained by interoperability and dependability demands.
[6]	Cyber-physical automation prototypes	Prototype architectures show feasibility of distributed automation, yet migration, standardization, and certification limit enterprise deployment.
[7]	Smart factory architecture	Connected sensors, cloud platforms, and intelligent control support adaptive production, but orchestration logic is often under-specified.
[8]	Smart manufacturing	Enterprise-scale manufacturing intelligence requires integrated analytics, sensing, optimization, and process connectivity.
[9]	Industrial big data pipeline	Real-time analytics depends on coordinated data ingestion, transformation, storage, and model execution.
[10]	Analytics capability and strategic alignment	Analytics value increases when analytical capability aligns with strategy and operational objectives.
[11]	Big data analytics and dynamic capability	Analytics contributes to performance through dynamic capability, not through data accumulation alone.
[12]	Strategic value from big data analytics	Business value requires governance, process embedding, and decision-oriented resource configuration.
[13]	Artificial intelligence in supply-chain risk	Artificial intelligence supports risk prediction and mitigation, but transparency and network-level validation remain limited.
[14]	Human–AI organizational decision-making	Effective decision systems combine computational analysis with human judgment and role redesign.

This fragmentation is extended in recent review-oriented literature. Artificial intelligence (AI) applications for supply-chain management are wide but not well-integrated, a vast majority of which look at forecasting, inventory management, transportation and supplier selection but do not focus on orchestration of coordinated agents [17]. The research of digital supply-chain twin has a better architecture bridge since digital twins can share the state representation of distributed agents, simulate scenarios, and support decision-making with resilience [18]. When the literature is combined, a multi-agent orchestration architecture for real-time enterprise decision intelligence systems is not yet a solid, consolidated research area. Rather, there are currently compatible, but somewhat disjointed fields that are becoming integrated via architectural-level investigation.

3. Methodology

Methodologically, the literature reviewed is based on conceptual architectures, prototype demonstrations, design frameworks, survey reviews, empirical analytics models and qualitative organizational theorization. In the studies on cyber-physical and smart manufacturing, a typical layered architecture is used, where all layers of entities share information via standard interfaces [5]–[8]. The power of this approach is that it provides architectural clarity in terms of where sensing, reasoning, execution, and coordination occur. The disadvantage is that enterprise decision intelligence

may involve multi-functional objectives that involve more than production control alone. There may be a production agent optimizing throughput, a risk agent optimizing resilience and a finance agent optimizing cost containment. The majority of the architectures reviewed do not describe completely how such objective conflicts are mediated in real time.

The methodological approach of pipeline studies is different. Industrial big data pipeline research focuses on the movement, transformation and deployment of the model as the operational requirements to make decisions with analytics [9]. The empirical models or conceptual frameworks used to link analytics resources with performance outcomes are commonly used in analytics capability studies [10]–[12]. This literature helps to understand the need for an organizationally embedded orchestration, but fails to analyze the technical architecture of interacting agents. Human–AI and explainability studies bring in techniques based on organizational design, decision theory and accountability [14]–[16]. This is vital to enterprise decision intelligence as real-time recommendations can be a source of operational risk if they are made without escalation, explanation, or override logic.

Table 2 presents a comparison of methodological approaches in the literature reviewed. The comparison highlights their respective strengths as they cover a different aspect of the topic and none of them provides a comprehensive description of multi-agent orchestration architectures for real-time enterprise decision intelligence systems. This provides a methodological opportunity for the integration of simulation, enterprise process modelling, field deployment and human-centered assessment in mixed architecture evaluations.

Table 2 Method comparison

Ref	Method	Strengths	Limitations
[5]	Agent-oriented cyber-physical architecture review	Specifies agent roles, coordination needs, and industrial control relevance.	Limited enterprise-level evaluation beyond industrial automation contexts.
[6]	Prototype-oriented cyber-physical implementation analysis	Connects architectural concepts to implementation barriers.	Prototype evidence remains difficult to generalize across sectors.
[7]	Smart factory architecture analysis	Links sensing, connectivity, and intelligent control.	Agent orchestration protocols receive limited analytical treatment.
[9]	Industrial data pipeline design	Clarifies real-time data dependencies for decision agents.	Coordination among multiple decision agents is not central.
[12]	Strategic value framework	Explains process embedding and value creation mechanisms.	Lacks low-latency orchestration and agent-level architectural detail.
[14]	Human–AI organizational decision analysis	Identifies judgment allocation and collaboration requirements.	Provides limited computational specification for orchestration.
[16]	Hybrid intelligence conceptualization	Frames joint human–machine cognition as a design target.	Requires operational translation into enterprise architecture.
[18]	Digital supply-chain twin framework	Offers shared state representation for resilient decision coordination.	Validation at real-time multi-agent enterprise scale remains limited.

Figure 1 shows a conceptual architecture centered on real-time enterprise decision intelligence, with streaming enterprise events being processed by specialized agents that communicate with each other via an orchestration layer to generate governed enterprise actions. This is important because the literature reviewed suggests that decision intelligence involves having to interact with agents, share state, govern, and give feedback on execution – not just prediction models. The figure is based on conceptual coding of architectural themes in [5]–[9], [14], [16], [18].




Figure 1 Conceptual architecture for multi-agent orchestration in real-time enterprise decision intelligence systems

The concept is that orchestration must be placed in the first-class architectural layer. With this missing layer, enterprise decision intelligence becomes a collection of individual analytics services that cannot resolve conflict, maintain auditability, or continue to take action when things change.

4. Discussion

Table 3 Results comparison

Ref	System / Context	Metric / Basis of Comparison	Outcome
[5]	Industrial cyber-physical agents	Architectural capability for modular control	Distributed agents supported modular coordination, with dependability and standardization unresolved.
[6]	Cyber-physical automation prototypes	Implementation feasibility and migration barriers	Prototypes validated distributed automation concepts, but legacy integration remained a major constraint.
[7]	Smart factory systems	Connectivity and adaptive production capability	Sensor-cloud-control integration supported adaptivity, while orchestration mechanisms stayed underspecified.
[9]	Smart manufacturing maintenance analytics	Data pipeline completeness for real-time analytics	Coordinated ingestion and processing enabled maintenance analytics more reliably than isolated model deployment.
[10]	Analytics capability in firms	Strategy alignment and performance linkage	Analytics improved performance when aligned with strategic priorities.
[11]	Big data analytics capability	Dynamic capability pathway	Performance effects depended on capability transformation rather than data scale alone.
[13]	Supply-chain risk management	Artificial intelligence contribution to risk detection and response	Predictive tools aided risk management, but validation and transparency gaps persisted.
[16]	Hybrid intelligence	Human-machine complementarity	Joint human-machine cognition improved decision framing but required explicit task allocation.
[18]	Digital supply-chain twin	Shared simulation and resilience basis	Digital twins offered a state layer for disruption analysis, with limited live multi-agent validation.

There is a major paradox regarding both decentralization and control. Distributed autonomy is an important factor in the literature on Multi-agent systems, as it allows for more flexibility and responsiveness at the local level [1, 5]. Accountability, traceability and control of the decision process are instead the focus of enterprise decision governance literature [4], [14], [15]. This is a tension in the architecture. More agent autonomy can shorten response times and increase adaptability, but can also make it more difficult to hold the agents accountable if pathways to action are not logged, explained, and bounded. Explainable AI is not fully a solution, since the explanations can be at different levels: model level or orchestration level. For enterprise decision intelligence, there needs to be explanatory information – why did a particular agent take a particular action? Why did other agent recommendations get rejected? Why did escalation happen? Why did the final action meet policy constraints?

The reviewed studies were qualitatively coded according to architectural function: agent coordination, real-time data, enterprise value, governance and digital-twin state representation, as presented in Figure 2. The distribution shows that the evidence base is fragmented, often studying technical orchestration, analytics value and governance individually and as separate decision intelligence architectures. Dominant contributions in [5]–[18] were coded qualitatively using a binary classification.




Figure 2 Evidence distribution across architecture functions in reviewed studies

The coding focuses on the extent to which architectural building blocks are represented in the literature, but limited in integrated assessment. While agent coordination is a concept in industrial systems research, enterprise value is in analytics, and governance in human–AI studies, real-time enterprise decision intelligence mandates a combined evaluation across all functions.

Interpretation is affected by technical differences between studies. Many existing works on cyber-physical systems operate under assumptions of deterministic/semi-structured operational domains where agents interact with equipment, sensors, and control services [2], [5], [6]. The literature of enterprise analytics typically focuses on more general performance goals, but not necessarily on what decisions are made, or how, second by second [10]–[12]. As for judgment and responsibility, human–AI work tackles these issues yet rarely has architectures that can be executed [14]–[16]. Digital-twin research contributes simulation capability and shared state representation, while real-time synchronization among multiple agents across enterprise platforms remains insufficiently validated [18]. Thus, there is not yet enough agreement to support a general architecture. A more defensible conclusion is that architecture must be bound by decision latency, domain volatility, regulatory exposure, interoperability burden, and acceptable autonomy level.

In Figure 3, four architectural archetypes are analyzed for their suitability in terms of latency fit, their coordination complexity, explainability, integration burden, and resilience, in the context of real-time enterprise decision intelligence. It is important as the literature review shows that there is not a specific architecture that dominates, but hybrid digital-twin and multi-agent orchestration provides a strong overall fit but increases coordination and integration complexity. Findings from [5]–[18] are interpreted and assigned a score of 1- 5 for comparative review visualization.




Figure 3 Comparative fit of architectural archetypes for enterprise decision intelligence

The key design implication from the review is illustrated in Figure 3. While centralized analytics may be easier to manage, it is less well suited for low-latency distributed action, edge Multi-agent systems can provide greater speed and resilience but with greater explanation and coordination loads, the potential for digital-twin-plus-agent architectures is enhanced for centralized governance, and human-governed hybrid architectures can further enhance accountability but may be less suited to high-frequency action domains.

The following are some research gaps. In the first place, there are few comparative benchmarks of orchestration architectures. Most of the studies focus on prototypes or frameworks, or on comparing a capability relationship to other models, but not on comparing centralized, hierarchical, federated and agent-based orchestration in a control framework. Secondly, decision intelligence metrics are still underdeveloped. These attributes accuracy, latency, throughput, resilience, explainability, escalation frequency, policy compliance, and human trust should be measured together, but studies rarely do so. Third, there is a dearth of enterprise-scale empirical evidence. While industrial prototypes offer much technical insight, enterprise decision intelligence extends beyond departments, platforms, incentives and risk regimes. Fourthly, orchestration-level accountability is weak. Future systems need to store logs and explanations for the outputs of the system, but also for agent negotiation, arbitration, and for rejected alternatives.

Future Directions

For real-time enterprise decision intelligence systems, future studies of multi-agent orchestration architectures should focus on the architecture level instead of further individual algorithm-based studies. One of the first priorities should be the development of benchmark environments that enable a comparison of centralized analytics, hierarchical control, federated orchestration, and multi-agent digital-twin architectures, operating on the same enterprise decision scenarios. Measures of latency, decision quality, frequency of conflict, recovery time, completeness of explanation, compliance of policy and human intervention rate should be included in such benchmarks. Though there are some cyber-physical and analytics studies, they only offer partial measures, while integrated metrics still do not exist [5], [9], [12].

Orchestration accountability is a second priority. The research on explainability and hybrid intelligence should be further extended to agent-level explainability [15] and [16]. Traceable documentation of agent observations, proposed actions, negotiation processes, arbitration rules, human overrides and final decisions are required for real-time decision intelligence systems. In regulated enterprise domains this is particularly crucial, where speed cannot be used as a

substitute for responsibility. The third priority is digital-twin based coordination. Shared state and simulation capability for distributed agents can be achieved with digital twins, but further research is required in order to validate live synchronization, drift detection, multi-agent scenario evaluation, and resilience under disruption [18].

The fourth priority is related to interoperability and migration. Legacy ERP, MES, CRM, SCM, cloud, edge and compliance systems are found within enterprise environments. Therefore, architecture patterns that allow for the gradual adoption instead of full platform replacement are worthy of study for multi-agent orchestration research [6, 7]. An interdisciplinary evaluation is a fifth priority. Technical research should include organizational decision rights, incentives, workflow redesign, and human oversight; organizational research should include latency constraints, data pipelines, agent protocols, and system verification. The most useful future work will consider real-time enterprise decision intelligence as a coupled socio-technical architecture where agents, humans, data infrastructure, governance, and digital twins are all components that interact.

5. Conclusion

Multi-agent orchestration architectures for real-time enterprise decision intelligence systems are a promising yet largely uncoordinated research field. The architectural value of distributed agents in cyber-physical and smart manufacturing contexts is supported with the strongest evidence, as is the operational importance of real-time data pipelines, the organizational importance of analytics alignment and the governance importance of human-centered and explainable AI. Though the literature suggests that agent orchestration can benefit enterprise responsiveness, adaptability, modularity and resilience, enterprise decision intelligence cannot just be based on decentralized computation. It also needs to include shared state representation, conflict resolution, auditing, human oversight, and integration with decision rights within the organisation.

The academic value of the topic is related to its ability to connect enterprise action with AI architecture under time-sensitive conditions. Challenges remain in comparative evaluation, explainability at the orchestration level, enterprise scale validation, integrating with legacy platforms and metrics that together account for latency, quality, accountability, and resilience. The next step is to shift from stand-alone demonstrations and conceptual architecture to integrated, empirically validated architectures. There is a clear need to explore the role and interaction of agents, digital twins, analytics pipelines and human governance in making reliable real-time decisions across complex enterprise environments as part of a mature research agenda.

References

- [1] Dorri, A., Kanhere, S. S., & Jurdak, R. (2018). Multi-agent systems: A survey. *IEEE Access*, 6, 28573–28593.
- [2] Monostori, L., Kádár, B., Bauernhansl, T., Kondoh, S., Kumara, S. R. T., Reinhart, G., Sauer, O., Schuh, G., Sihn, W., & Ueda, K. (2016). Cyber-physical systems in manufacturing. *CIRP Annals*, 65(2), 621–641.
- [3] Seddon, P. B., Constantinidis, D., Tamm, T., & Dod, H. (2017). How does business analytics contribute to business value? *Information Systems Journal*, 27(3), 237–269.
- [4] Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4), 66–83.
- [5] Leitão, P., Karnouskos, S., Ribeiro, L., Lee, J., Strasser, T., & Colombo, A. W. (2016). Smart agents in industrial cyber-physical systems. *Proceedings of the IEEE*, 104(5), 1086–1101.
- [6] Leitão, P., Colombo, A. W., & Karnouskos, S. (2016). Industrial automation based on cyber-physical systems technologies: Prototype implementations and challenges. *Computers in Industry*, 81, 11–25.
- [7] Wang, S., Wan, J., Li, D., & Zhang, C. (2016). Implementing smart factory of Industrie 4.0: An outlook. *International Journal of Distributed Sensor Networks*, 12(1), 3159805.
- [8] Kusiak, A. (2018). Smart manufacturing. *International Journal of Production Research*, 56(1–2), 508–517.
- [9] O'Donovan, P., Leahy, K., Bruton, K., & O'Sullivan, D. T. J. (2015). An industrial big data pipeline for data-driven analytics maintenance applications in large-scale smart manufacturing facilities. *Journal of Big Data*, 2(1), 25.
- [10] Akter, S., Wamba, S. F., Gunasekaran, A., Dubey, R., & Childe, S. J. (2016). How to improve firm performance using big data analytics capability and business strategy alignment? *International Journal of Production Economics*, 182, 113–131.

- [11] Wamba, S. F., Gunasekaran, A., Akter, S., Ren, S. J. F., Dubey, R., & Childe, S. J. (2017). Big data analytics and firm performance: Effects of dynamic capabilities. *Journal of Business Research*, 70, 356–365.
- [12] Grover, V., Chiang, R. H. L., Liang, T. P., & Zhang, D. (2018). Creating strategic business value from big data analytics: A research framework. *Journal of Management Information Systems*, 35(2), 388–423.
- [13] Baryannis, G., Validi, S., Dani, S., & Antoniou, G. (2019). Supply chain risk management and artificial intelligence: State of the art and future research directions. *International Journal of Production Research*, 57(7), 2179–2202.
- [14] Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577–586.
- [15] Rai, A. (2020). Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science*, 48(1), 137–141.
- [16] Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61(5), 637–643.
- [17] Toorajipour, R., Sohrabpour, V., Nazarpour, A., Oghazi, P., & Fischl, M. (2021). Artificial intelligence in supply chain management: A systematic literature review. *Journal of Business Research*, 122, 502–517.
- [18] Ivanov, D., & Dolgui, A. (2021). A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. *Production Planning & Control*, 32(9), 775–788.