A Laboratory Testbed for Benchmarking Orchestration Logic in Collaborative Environments

Leopoldo O. Silva[1,2(✉)], Diogo L. Rodrigues[3], Francisco J. Duarte[1,2], and Luís S. Dias[1,2]

[1] CCG/ZGDV Institute, University of Minho, Azurém Campus, 4800-058 Guimarães, Portugal
leopoldo.silva@ccg.pt
[2] ALGORITMI Research Centre, University of Minho, 4710–057 Braga, Portugal
[3] Colep Packaging, Rua Comendador Arlindo Soares de Pinho 1977, Vale de Cambra, Portugal
https://www.ccg.pt

Abstract. The need for agile and decentralized supply chains, demanded by Industry 4.0, requires seamless integration among several Enterprise Information Systems. On this paper we present an architecture and laboratory exploration tool designed to test and benchmark orchestration algorithms within a collaborative environment. By conceptualizing the Enterprise Information Systems as digital representations of the ecosystem enterprises and using a modern context broker for synchronization, our laboratory tool acts as a decoupled orchestration architecture. We validate this tool through a First-Come, First-Served baseline algorithm and quantify the architecture performance via a set of laboratory experiments measuring the orchestration latency, throughput, fidelity and architecture resilience. The results demonstrate that our architecture provide a low-latency, scalable and resilient foundation for coordinating supply chains experiments, confirming the viability of the proposed transactional shared data model and offering a pathway for further tests with advanced planning algorithms in diverse industrial contexts.

Keywords: Supply Chain Orchestration · Planning Algorithms · Industry 4.0 · Digital Twin

1 Introduction

Industry 4.0 requires agile, decentralized supply chains synchronized across heterogeneous Enterprise Information Systems (EIS) [1]. However, to achieve the needed operational agility, several challenges remain unsolved, being data silos and limited visibility between organizations two of those [2]. This research addresses these challenges by proposing a collaborative environment where each

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026, LNCS 16769, pp. 177–192, 2027.
https://doi.org/10.1007/978-3-032-30494-0_12

participant is an enterprise, represented by its Digital Twin of the Organization (DTO). By leveraging the FIWARE Orion Context Broker and the NGSI-LD standard, a shared digital context is established by the framework, allowing orchestration among the participant enterprises. The main contribution is a laboratory tool designed to explore and benchmark orchestration within distinct industrial scenarios. This tool was validated through a First-Come, First-Served (FCFS) baseline within a fragmented enterprise ecosystem, providing a pathway for future tests on supply chain synchronization in distinct industrial scenarios.

2 Related Work

One the main goals of Industry 4.0, is to integrate Cyber-Physical Systems and the Internet of Things into manufacturing [3], with the Digital Twin serving as a real-time virtual mirror of physical assets [4]. The synchronization of decentralized value chains, with their own data silos, using a network of twins to achieve seamless collaboration, requires event-driven mechanisms capable of keeping low-latency responses during high-frequency orchestration tasks, a requirement that is not well quantified in decentralized industrial contexts [5].

Using EISs as the source of truth, DTOs abstract complete enterprises and their processes by keeping operational states up to date. [6]. However, even if these twins are already able to provide capacity and constraints models, there is still a research gap in how to federate this models for cross-company orchestration, particularly regarding mechanisms to persist shared digital states and maintain synchronization after failures. [7].

The FIWARE ecosystem [8] and the NGSI-LD standard [9] support the technical interoperability required by these federated environments. Context brokers manage high-density data flows by decoupling ecosystem participants through standardized Interfaces. The asynchronous notification architecture implemented by these technologies is essential for maintaining the fidelity of the DTOs and their digital shadow [10], ensuring that state changes, such as production progress, are propagated across the ecosystem without the delays associated with traditional methods.

To plan a decentralized supply chain, trade-offs between local optimization and global efficiency are required [11]. As such, a baseline logic, like the FCFS, provides a neutral foundation for benchmarking more advanced algorithms [12]. There is also a clear need for evidence-based testbeds that can empirically measure the throughput and reliability of these algorithms in diverse industrial scenarios.

3 Research Methodology

On this work we develop a laboratory testbed for benchmarking decentralized planning strategies, using the Design Science Research Methodology [13]. This methodology allows us to interactively design and evaluate artifacts in controlled environments, which is essential for bridging theoretical supply chain models

with the practical interoperability requirements of Industry 4.0. On the design phase, four critical requirements were identified: semantic interoperability to unify diverse EIS, state representation for digital twin synchronization, architectural scalability for concurrent requests, and system resilience to ensure state consistency during service initialization and restoration. The FIWARE ecosystem was selected as the technological stack for context-aware data management, since it promises to facilitate the resolution of those requirements. It also results from an European project and will allow the technological sovereignty that is sought off in current European directives.

A single-hub multi-spoke model was used on the experimental design to simulate coordination across decentralized participants while ensuring the isolation of variables. Õn this framework, several planning algorithms serve as independent variables, being our future goal to test distinct operational philosophies, ranging from basic sequencing to complex collaborative optimization [14]. At this step, the study uses a FCFS baseline to establish empirical benchmarks for system connectivity and asynchronous state machine transitions.

To validate our architecture we realized several simulations that cross-reference the orchestration logic with specific business pressures. A multi-metric evaluation matrix was used to quantify performance, focusing on orchestration latency, system throughput, twin convergence, and initialization recovery time. This systematic approach provided empirical evidence to assess the suitability of the orchestration logic for fragmented participant ecosystems. With the quantification of these parameters, the research demonstrates the utility of our foundational architecture for further experimentation on synchronized and resilient supply chain management.

4 The Experimental Simulation Framework

We have designed the experimental simulation framework as a modular and interoperable architecture, illustrated in Fig. 1, with the goal of facilitating the digital representation of complex supply chain dynamics. On it, each participant, represented by their EISs, can interact through a unified digital interface. By conceptualizing these systems as DTOs, the framework enables the mapping of internal operational states to a shared global context. This approach ensures that participant's information, such as production capacity, request status, and resource availability, is consistently maintained and synchronized across the ecosystem, providing a solid foundation for automated decision-making and strategic planning.

The integration of distributed data sources through a centralized context broker is the core of this framework and it serves as the synchronization backbone for all digital entities within the Digital Twin Layer. Using standardized communication protocols and semantic data models, the framework decouples the internal complexity of individual EISs, in the Enterprise layer, from the orchestration logic in the Service layer. This decoupling is essential for creating a flexible testbed where different planning paradigms can be evaluated against various business constraints.

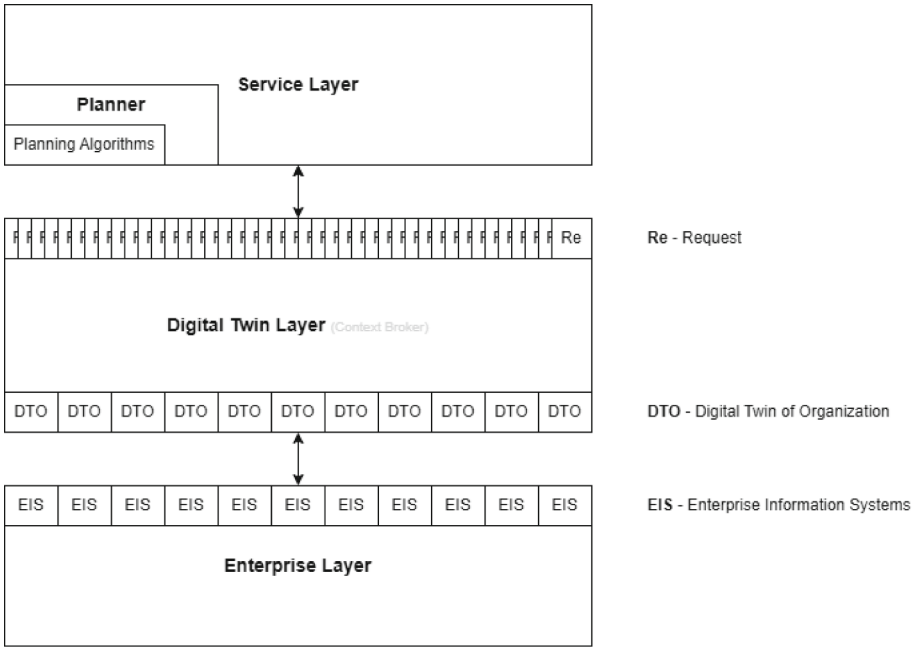

Fig. 1. High-level system architecture of the collaborative simulation framework.

The following sections detail the structural components of the framework, including the technological infrastructure, it's layers and the shared data model.

4.1 The Testbed Environment

To provide a context-aware and robust infrastructure for our industrial ecosystem, the FIWARE Orion context-broker was used as the foundation of the framework, function as a high-performance mediator of the testbed and maintaining a consistent global context across the digital twins. This approach promises to eliminate information gaps by ensuring that all participating entities share a synchronized view of the industrial state.

The broker maintains an up-to-date digital representations through the NGSI-LD standard, where participants and their requests, are modeled as semantic entities. This linked-data approach creates a comprehensive digital data set that serves as the single source of truth for the network, facilitating the complex inter-organizational coordination required for our experiments.

Furthermore, the architecture decouples data producers from consumers through an asynchronous publish and subscribe mechanism. The participants interact with each-other, sololely by the mediation of their digital twins and the broker, which automatically trigger notifications to the service layer when the right conditions are met. This event-driven interaction should allow high system

responsiveness and scalability, supporting the rapid state transitions that we expect to observe during experimental validation.

4.2 Data Models

To ensure seamless interoperability, a semantic approach based on the NGSI-LD standard is used. A common digital language is defined with the goal of describing the participants and their capacities. The data architecture centers on two primary linked entities: the Participant, representing the physical ecosystem with their EISs, and the Request, representing the lifecycle of the transactions among participants.

Table 1. Data Model Attributes for the *Participant* Entity

Attribute	Data Type	Description	Orchestration Utility
`id`	URI	Unique system identifier	Entity referencing
`name`	Property	Human-readable identifier	Identity and logging
`roles`	List	Consumer/Producer status	Capability filtering
`location`	GeoProperty	Geographic coordinates	Transport latency
`maxThroughput`	Property	Absolute capacity limit	Scalability assessment
`currentWorkload`	Property	Active request count	Fairness distribution
`availableCapacity`	Property	Remaining processing space	Concurrency control
`operationalStatus`	Property	Current availability state	Resilience/Re-planning
`avgProcessingTime`	Property	Manufacturing time per unit	Total Lead Time

The Participant entity is modeled as a multi-role object that supports industrial prosumer behaviors, by allowing participants to act as producers, consumers, or both. Attributes summarized in Table 1 provide the identity, spatial context, and operational states necessary for orchestration. The current experiments only use a subset of these fields, others such as `location` and `avgProcessingTime` are maintained to allow the extensibility of the platform for future logistical scenarios.

The Request was designed as a consolidated transactional artifact that integrates planning and execution data directly into the entity (Fig. 2). This model integrates the planning and execution phases into the request itself, defining sequential responsibilities for participants. The `status` attribute functions as a dynamic state machine, transitioning through stages such as `PENDING`, `ALLOCATED`, `IN_PRODUCTION`, and `COMPLETED` and allowing all involved participants to update their own internal states and with their own logic, by monitoring the evolution of the request (Table 2).

Our hope is that this merging of planning, execution, and lifecycle status into a single entity, allows the framework to ensure high data consistency and provide

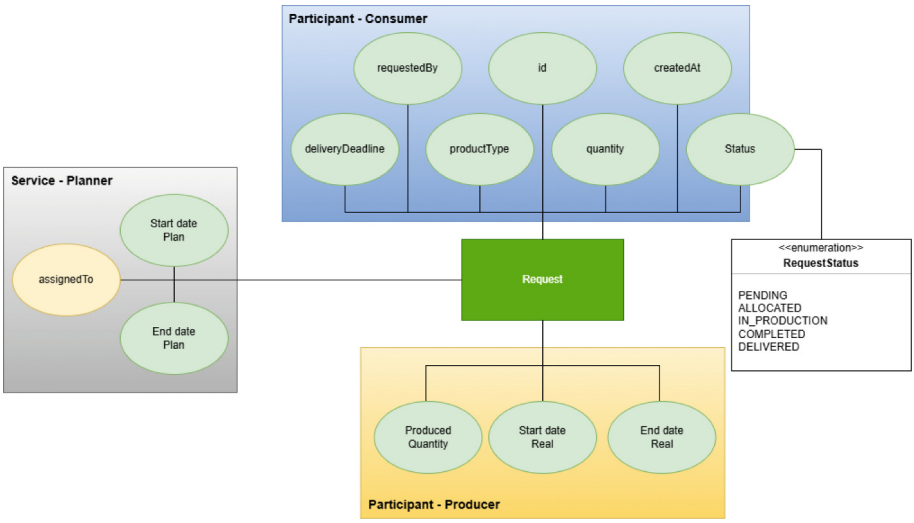

Fig. 2. Request entity data model: attributes, stakeholder responsibilities, and lifecycle status transitions.

Table 2. Data Model Attributes for the *Request* Lifecycle

Responsible	Attribute	Function
Consumer	`id, productType, quantity`	Demand Definition
	`deliveryDeadline, requestedBy`	Constraints and Origin
Planner	`assignedTo`	Allocation Decision
	`startDatePlan, endDatePlan`	Optimized Schedule
Producer	`producedQuantity`	Execution Monitoring
	`startDateReal, endDateReal`	Real-time Feedback
Shared	`status` (Enumeration)	Workflow Synchronization

an up-to-date view of the underlying processes. This integrated structure simplifies interoperability requirements for the participants while providing a scalable foundation for synchronized and resilient supply chain orchestration. Similarly to the participant entity model, the request model has specific attributes, such as planned and real execution dates, that are only partially used on the current baseline tests. These attributes were kept to enable future evidence-based decision-making through precise deviation analysis.

4.3 The Enterprise and Digital Twin Layer

The Digital Twin layer interfaces internal organizational logic with the collaborative network by representing each participant as a DTO. By utilizing the EIS as the foundational source of truth, this layer abstracts each participant business

complexities into standardized digital states. Functionally, this layer ensures that transactional changes in the EIS are reflected in the digital counterpart through an event-driven mechanism, allowing digital twins to act as mirrors of organizational constraints and facilitating the state convergence we wish to attain during experimental validation.

Finally, this layer supports information flows bi-directionally, acting both as a sensor and actuator for the participants EISs, either updating the status on the digital representation of the participant or sending orchestration commands to the EISs, This command injection into organizational workflows should allow the integration of the centralized assignments into local constraints.

4.4 The Service Layer

The service layers provides the specialized coordination processes of our testbed, representing the strategic command center of the environment. While the Orion Broker provides a dynamic mirror of the network's overall state, including request fulfillment and participant availability, this layer adds the centralized governance for the interactions between the participants of our framework.

At the moment and for the purpose of our experiments, the provided services are exclusively request fulfillment planners and allocators, allowing this layer to govern the industrial transaction lifecycle. Upon detecting a new request, the planner initiates the allocation process and triggers the necessary updates to the shared entities. Such governance ensures that interactions remain synchronized and aligned with the strategic objectives of the specific industrial scenario being simulated.

The service layer ensures system resilience without requiring direct integration with the EISs and this modular design allows for the seamless replacement of planning algorithms while maintaining a trail of all activities. By acting as the primary agent of change, the service layer transforms independent participants into a reactive and synchronized ecosystem.

5 Experimental Design and Setup

On this section we detail the laboratory-based evaluation we did to quantify the performance and reliability of the implemented framework. The goal was to validate the feasibility for synchronization and functional benchmarking by establishing a controlled simulation environment. Subsections 5.1 and 5.2 outline the FCFS orchestration baseline and the participant ecosystem parameterization, while Sect. 5.3 presents the experimental procedures used to measure orchestration latency, throughput, digital fidelity, and architectural resilience. With this approach we want to ensure that the results are based on empirical evidence and provide an assessment of the framework's suitability for decentralized supply chain experiments.

5.1 Implementation Baseline

Our functional baseline, the FCFS algorithm, uses event-driven sequential logic to process requests based on arrival time. This algorithm serves as a transparent gatekeeper for decentralized allocation, establishing essential connectivity between participants through a neutral and predictable task-release mechanism. It also provides a clean and neutral logic baseline, able to measure the performance limits of the NGSI-LD synchronization layer before introducing another algorithms in later experiments.

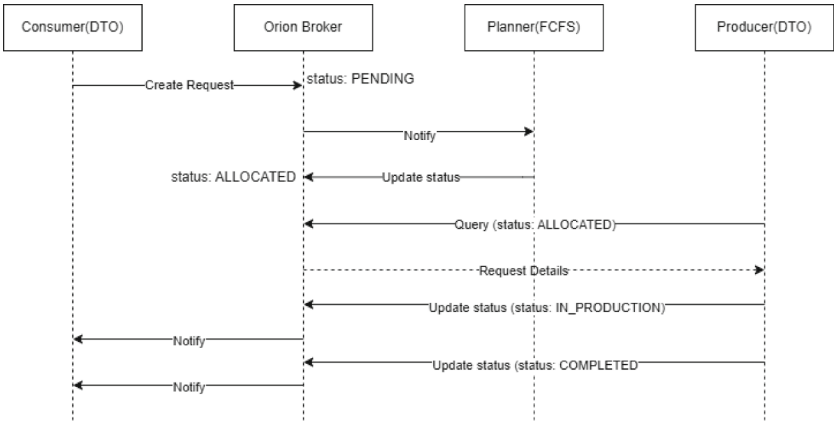

Fig. 3. Sequence diagram of the FCFS orchestration workflow.

Technical implementation relies on an asynchronous listener using the Orion Broker subscription system(Figure 3). The planner monitors requests in `PENDING` state and, upon notification, performs a state transition to `ALLOCATED`. This transition is the primary coordination signal, allowing producers to identify and claim tasks through their own decoupled notifications.

This shared context doesn't require point-to-point integration and enables a continuous feedback loop via automated status update notifications. As producers transition requests to `IN_PRODUCTION` and `COMPLETED`, notifications synchronize the consumers that created those requests. This asynchronous interaction among participants maintains a consistent state representation across all of them.

Furthermore, during the planner initialization, a synchronization routine executes to identify backlog requests created during orchestrator downtime. This query ensures that no transaction is overlooked, maintaining data consistency and a predictable state machine regardless of component startup timing or blackouts.

5.2 Simulation Environment

The simulation environment replicates a fragmented SME ecosystem, mirroring sectors where production is distributed across specialized low capacity par-

ticipants. Each SME is modeled as a DTO within the Orion Broker, featuring restricted processing windows and specific throughput constraints(resource modeling is parameterized by setting `maxThroughput` and `availableCapacity` to low threshold values). With this experiment we want to challenge the service layer to synchronize multiple small-scale units to satisfy collective demand.

5.3 Experimental Procedures and Metrics

Four laboratory experiments where done to evaluate the performance of the FCFS baseline within the SME ecosystem. The following sub-subsections detail protocols for assessing orchestration latency, system throughput, digital fidelity, and architectural resilience, providing a systematic framework to quantify the efficiency of our architecture for supply chain coordination.

Experiment 1: Orchestration Latency and Responsiveness. On this experiment we establish a baseline for orchestration latency, defined as the difference between initial request in a `PENDING` state and the following `ALLOCATED` update by the planner and received by the consumer. This metric assesses the responsiveness of the asynchronous notification mechanism, quantifying the architectural overhead. On this experiment there are only two participants, a consumer making requests and a producer receiving them.

The procedure involves creating 100 requests from a consumer, recording timestamps for broker processing, network latency, and logic execution. The mean latency and it's standard deviation are then calculated, allowing to evaluate the viability of the NGSI-LD standard for coordination. A low and consistent response time demonstrates that the decoupled architecture keeps agility and that the ecosystem can remain synchronized under strict temporal requirements.

Experiment 2: System Throughput Under Network Load. On this experiment we evaluate the scalability of the orchestration layer under high-density transactional flows. The goal is to assess the capacity of the Orion Broker and the FCFS planner to process bursts of information, ensuring that requests transition from `PENDING` to `ALLOCATED` without compromising context integrity.

The procedure tests the management of parallel threads and their database updates, simulating peak loads with a burst of 100 concurrent requests from multiple consumers. Using system throughput (measured in requests per minute) and the success rate of state transitions we identify potential bottlenecks or race conditions.

The framework's suitability for large-scale scenarios is determined by these limits. If the framework is able to achieve a high throughput rate, it may support the intense transactional demands of decentralized supply chains while maintaining operational efficiency.

Experiment 3: Digital Convergence and Fidelity. On this experiment we validate digital synchronization during the execution phase by focusing on the propagation of `producedQuantity` updates. The procedure simulates a production cycle where a producer advances the `producedQuantity` attribute, while the consumer updates is own state when notifications are received. The main metrics are synchronization jitter, measuring propagation time variance, and state fidelity, ensuring state consistency between participants (the consumer and the producer). By quantifying the convergence between local and perceived states, this experiment want to demonstrate the efficacy of the shared model for evidence-based decision-making.

Experiment 4: Resilience and Initialization Bootstrapping. On the final experiment we want to assess the resilience of the orchestration layer during component downtime, where resuming operations without data loss is crucial for reliability. The procedure tests the planner initialization routines by identifying requests missed during inactivity. For this, a participant injects a backlog of requests that remain in a `PENDING` state, during a controlled interruption of the planner. Upon restart, the initialization routine should obtain those requests and transition them to `ALLOCATED`. The main metrics are recovery time and chronological state consistency, which are important parameters for architectural robustness and the value of the broker's persistence layer during automated recovery in unpredictable environments.

6 Results

With the experimental procedures defined, as described in the previous section, we conducted the experiments and evaluation of our framework. The data collected with these experiments provided the empirical evidence regarding the performnace of the framework and the FCFS orchestration baseline within the fragmented SME ecosystem, On the following sections we present the quantitative results and technical observations obtained from each experiment, focusing on the efficiency of the asynchronous state transitions and the fidelity of the digital representations present on both the broker and the EIS simulations.

6.1 Experiment 1: Orchestration Latency and Responsiveness

The results for Experiment 1 quantify the end-to-end latency of the orchestration process, that is the elapsed time from the initial request creation by the consumer to the receipt of the request `ALLOCATED` status notification. This metric allows for the validating the reactivity of our framework. As shown in the representative sample (Table 3), the system demonstrated high temporal stability throughout the execution phase.

The system achieved an average end-to-end latency of 10.93 ms and the low standard deviation of 0.99 ms suggests that the asynchronous notification mechanism provided by the Orion Context-Broker, is highly predictable for coordination, ensuring that the overhead introduced by the digital twin and the service

Table 3. End-to-End Orchestration Latency Samples (ms)

Request ID	Latency (ms)
Request 01	9.98
Request 02	11.95
Request 03	10.13
Request 04	10.93
Request 05	10.28
Request 06	12.31
Mean Latency	**10.93**
Standard Deviation	**0.99**

layer remains well within the requirements for supply chain synchronization scenarios.

The logs also reveal that the primary component of this latency is on the internal dispatching logic of the context broker and the network round-trip between the distributed participants and services. The internal processing time of the planner service was consistently recorded as being below 2 ms, confirming that the FCFS logic is not a bottleneck. The consistency of these results validates the feasibility of the architecture for experimental scenarios and demonstrates that the decoupling of participants through a shared digital context maintains a high level of responsiveness.

6.2 Experiment 2: System Throughput Under Supply Chain Load

The objective of Experiment 2 was to evaluate the scalability and performance of the framework when subjected to a high volume of concurrent requests. The laboratory environment was configured with five consumers and ten producers, simulating a network of participants. The results (Table 4) show the system's capacity to stay under operational stability during peak load conditions.

Table 4. System Throughput and Load Performance Summary

Metric	Observed Value
Total Processed Requests	100
Simulation Duration (s)	80
Average Throughput (Requests/min)	75.0
Peak Throughput (Requests/s)	1.25
State Transition Success Rate	100%

These results show that the framework was able to hndle simultaneous injection of requests from multiple sources. As the consumers create requests, the

service layer processed the resulting notifications in a consistent way, ensuring that the shared context remained synchronized. A critical observation is the high degree of parallelism achieved at the execution level even if the FCFS planner acts sequentially, the decentralized participant layer enabled up to ten simultaneous producing cycles, effectively demonstrating some of the advantages of the decoupled architecture.

On the top of that, there is no evidence of race conditions or database deadlocks on the results. The stability of the throughput, which remained consistent at approximately 75 requests per minute, along with the previous findings, suggest that our architecture is capable of supporting the transactional density required for other decentralized scenarios and provides a robust foundation for more complex optimization algorithms in future research iterations.

6.3 Experiment 3: Digital Convergence and Fidelity

The objective of Experiment 3 was to validate the accuracy and synchronization of the digital representation during the execution phase. In our framework, fidelity is measured by the system's ability to propagate, with minimal delay, state changes from the producer to the consumer. By cross-referencing the timestamps of the updates done, by the producers, to the requests with the notifications at the consumers, we quantify the temporal gap in the digital representation of the shared context.

The results (Table 5), derived from the simulation involving 15 participants, indicate a high degree of temporal convergence, with a mean state synchronization latency of 8.42 ms. This suggests that the use of the Orion Context Broker for shared state management eliminates the information lag traditionally associated with inter-organizational communication. Even under the persistent load of many simultaneous production cycles, the jitter remained low and, as a result, the consumer's view of the producer's progress (`producedQuantity`) remained consistent and reliable.

Table 5. Digital State Convergence Latency (ms)

Event Type	Min Latency	Max Latency	Mean	Std. Dev.
Status: `STARTED`	4.12	14.56	9.15	1.34
Status: `PROGRESS`	3.85	12.10	7.82	0.95
Status: `COMPLETED`	4.30	13.92	8.29	1.12
Overall Average	–	–	**8.42**	**1.14**

These findings demonstrate the efficacy of the shared transactional data model, as the EISs of each organization may be synchronized across the network without the need for direct point-to-point integration. This may enable stakeholders to make evidence-based decisions, such as adjusting logistics or

downstream production schedules, based on the high-fidelity digital representation of their suppliers.

6.4 Experiment 4: Resilience and Initialization Bootstrapping

The final experiment evaluated the resilience of the architecture by assessing its ability to recover and synchronize pending requests after a period of downtime. This scenario is critical for environments where component failures cannot result in data loss or permanent supply chain desynchronization.

Table 6. System Resilience and Initialization Recovery Metrics

Metric	Observed Value
Backlog Size (Requests)	50
Total Recovery Time (ms)	421.0
Bootstrap Throughput (Req/s)	118.7
Transition to Live State	Seamless
Data Consistency Post-Recovery	100%

The results (Table 6) demonstrate that the initialization routine was able identify all requests created during the FCFS service inactivity. The framework shown a rapid recovery ability, having cleared the backlog with a high-performance burst and achieving a throughput of over 118 requests per second. This allowed the DTOs and the corresponding EISs simulations to be brought back into alignment with the global state almost instantaneously.

The experiment also shown a seamless transition between the initialization sweep and the asynchronous notification mechanism. Requests created just after the planner's restart were processed with the baseline latency from Experiment 1. These findings validate the robustness of our architecture, demonstrating that the use of a persistent context broker as a single source of truth provides the necessary resilience for experimental industrial deployments.

7 Discussion and Analysis

The empirical evaluation validates that our framework is ready for future experiments in decentralized supply chain orchestration. The following analysis interprets this results, specifically assessing the architectural trade-offs and strategic value of the collaborative environment within Industry 4.0 scenarios.

First, the conceptualization of EIS as DTOs provides a semantic context that unifies distinct data sources usually residing in their own data silos. By decoupling the internal complexity of the EISs from the orchestration logic, the framework avoids the cost of point-to-point integration while enabling participants synchronization. This transition changes the industrial ecosystem into a

unified system of systems, where participant states are available for coordinated decision-making.

FCFS results indicate that the context broker overhead is irrelevant when compared to industrial lead times. The observed 10 ms end-to-end latency proves the responsiveness of the architecture for real-time orchestration. Paired with the throughput of 75 requests per minute, it also validates system scalability for maintaining agility in fragmented networks.

Another key finding is the temporal convergence between physical execution and digital representation. The experiments on fidelity confirm a consistent digital representation across the participants, with synchronization jitter remaining in low ranges. This may enable evidence-based decision-making by eliminating operational blind spots and shows that shared transactional artifacts provide a reliable source of truth for producing life-cycles.

The strategic value of shared states for system resilience is shown by the bootstrapping results. With the achievement of a recovery throughput of 118 requests per second, the framework demonstrates the ability to handle backlog surges while preserving chronological integrity. This robust transition between recovery and notification-driven operations ensures that the framework survive the intermittent failures typical of distributed industrial environments.

Finally, even if FCFS establishes essential connectivity, its sequential logic lacks the logistical awareness required for complex supply chains. Our architecture addresses this by acting as a testbed for future advanced algorithms experiments, including Concurrent and Collaborative Planning. In this study, FCFS acts as a controlled baseline that allows us to isolate the architectural overhead from algorithmic complexity, ultimately validating that the framework is ready to host more advanced algorithms. We do believe that the underlying data models enable these kinds of algorithms and their application to other sectors and applications, where synchronized digital states are paramount for long-term operational sustainability.

8 Conclusion and Future Work

This research validates our experimental laboratory testbed designed to experiment with fragmented supply chains Industry 4.0 scenarios. By assuming that many of the current interconnection issues with EISs may be solved by conceptualizing them as DTOs and using FIWARE Orion context broker as a backbone to keep those DTOs data, our framework seems to enable real-time orchestration through a shared digital context. Our experimental validation using the FCFS baseline confirmed that the architecture as the needed responsiveness, scalability and resilience, proving that our common data model and respective instantiated entities maintain the needed fidelity for synchronizing independent organizations EISs. This laboratory testbed serves as a functional artifact for evidence-based decision-making, by providing a single source of truth via shared transactional artifacts and offering an agile, interoperable alternative to brittle point-to-point integrations.

Preliminary integration of logistically aware heuristics is already showing promising results, further demonstrating the modularity of the Service Layer for multi-criteria benchmarking. In future research we plan to experiment with advanced planning algorithms, such as Service Fairness for equitable load balancing, Concurrent Planning for high-density environments, and Collaborative Planning for multi-criteria matchmaking. These modules will leverage the transport and scheduling attributes, already in place in the data model, to address coordination challenges beyond sequential task release. Furthermore, the framework will expand to encompass new ecosystem scenarios like verticalized industries or global supply chains. Such extensions intend to demonstrate the versatility of our framework, positioning it as a foundational technology for resilient, synchronized, and sustainable global value chain experimentation.

Acknowledgments. This work has been supported by the European Union under the Next Generation EU, through a grant of the Portuguese Republic's Recovery and Resilience Plan (RRP) Partnership Agreement, within the scope of the project PRODUTECH R3 – "Agenda Mobilizadora da Fileira das Tecnologias de Produção para a Reindustrialização", Total project investment: 166.988.013,71 Euros; Total Grant: 97.111.730,27 Euros.

References

1. Hofmann, E., Rüsch, M.: Industry 4.0 and the current status as well as future prospects on logistics. Comput. Ind. **89**, 23–34 (2017)
2. Muñoz, M., Torres, M., Gil, J., Guzmán, J.: An internet of things platform for heterogeneous data integration: methodology and application examples. J. Netw. Comput. Appl. **240**, 104197 (2025)
3. Lasi, H., Fettke, P., Kemper, H.-G., Feld, T., Hoffmann, M.: Industry 4.0. Bus. Inf. Syst. Eng. **6**, 239–242 (2014)
4. Tao, F., Zhang, H., Liu, A., Nee, A.Y.C.: Digital twin in industry: state-of-the-art. IEEE Trans. Industr. Inf. **15**(4), 2405–2415 (2019)
5. Fuller, A., Fan, Z., Day, C., Barlow, C.: Digital twin: enabling technologies, challenges and open research. IEEE Access **8**, 108952–108971 (2020)
6. Riss, U.V., Maus, H., Javaid, S., Jilek, C.: Digital twins of an organization for enterprise modeling. In: Grabis, J., Bork, D. (eds.) PoEM 2020. LNBIP, vol. 400, pp. 25–40. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-63479-7_3
7. Kierans, D., Pleiter, D.: Realising distributed digital twins within federated digital infrastructures. In: DiDit 2024 - Proceedings of the 1st International Workshop on Distributed Digital Twins, co-located with DisCoTec 2024 - 19th International Federated Conference on Distributed Computing Techniques, DisCoTec 2024: CEUR-WS, 2024. QC 20250715 (2024)
8. FIWARE Foundation. Fiware ecosystem archives (2025). Accessed 30 Mar 2026
9. ETSI. ETSI GS CIM 009 V1.9.1: Context Information Management (CIM); NGSI-LD API (2025). Accessed 30 Mar 2026
10. Kritzinger, W., Karner, M., Traar, G., Henjes, J., Sihn, W.: Digital twin in manufacturing: a categorical literature review and classification. IFAC-PapersOnLine **51**(11), 1016–1022 (2018). 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018

11. Saharidis, G.K., Dallery, Y., Karaesmen, F.: Centralized versus decentralized production planning. RAIRO - Oper. Res. **40**, 113–128 (2006)
12. Hesselink, W.H., Buhr, P.A., Parsons, C.A.: First-come-first-served as a separate principle. ACM Trans. Parallel Comput. **11** (2024)
13. Peffers, K., Tuunanen, T., Rothenberger, M., Chatterjee, S.: A design science research methodology for information systems research. J. Manag. Inform. Syst. **24**, 45–77 (2007)
14. An, Y., Zhang, Y., Zeng, B.: The reliable hub-and-spoke design problem: models and algorithms. Transp. Res. Part B Methodol. **77**, 103–122 (2015)