

Towards a Supply Chain Digital Twin

George Georgiades

Marinos Kariolis

Rafaellos Leontiou

Savvas G. Loizou

Abstract—This work proposes a novel concept of digital twin, the Supply Chain Digital Twin (SCDT), aimed at supply chain operations in the manufacturing sector. The paper describes the development, architecture and implementation of SCDT, with the capability of both monitoring near real time data from various factories and simulating various scenarios of operations, recycling schemes and logistics between manufacturing factories in a supply chain. The proposed approach to factory representation, allows a variable degree of detail in the operational models and their visualizations, providing flexibility and scalability. An appropriate visual representation of the supply chain(s) is provided, that includes concurrent visualizations of both the inter- and intra- factory logistics. The material flows are produced by logistic tasks, transferring materials from one factory to another, and manufacturing processes, asynchronously converting quantities of a combination of materials into scrap, waste and products. The operations can be arbitrary (continuous/discrete), and sequential and/or concurrent. The complexity of the system can vary from a single factory to a group of defined supply chains exchanging and sharing resources. The SCDT tool resides on a server, and can accept multiple requests from multiple users. The user can request to visualize either real time data from a platform with sensor data, or to simulate the time evolution of a supply chain based on the current factory's (or supply chain's) state as a starting point or to study the evolution of hypothetical scenarios. In addition to the operational process models, the SCDT develops an interactive 3D visualization of the requested setup of factories and their supply chain. The paper showcases the various plugins/tools that were developed, as well as the capabilities of the SCDT tool. Example case studies are presented and analyzed.

Keywords: Supply Chain Digital Twin, Manufacturing Supply Chains, Near Real-Time monitoring, Scenario-based simulation, 3D Visualization

I. INTRODUCTION

Digital Twin technologies have recently become the focus of massive attention as main utilities for thoroughly exploring, refining, and optimizing complex industrial systems[1]. In the frame of manufacturing and supply chains, such technologies support and improve decision making, as they reveal system behavior, utilization of resources, and the performance of both real and hypothetical operating conditions[2].

The synergy between virtual representation and simulation holds particular relevance as the combination of these two in a single tool creates a seamless link between real-time monitoring and predictive analysis, thus providing for complete decision support [3][4].

Overall, digital representations reflect the current state of operations for instant awareness, whereas simulation takes

it a step further by considering "what-if" scenarios, thus forecasting results such as supply shortages or demand spikes [5]. Such a mixture is highly necessary to manage the intricacies of contemporary supply chains effectively, as the ability to quickly adapt and remain resilient heavily relies on both reactive monitoring and proactive forecasting [6][7][8].

In manufacturing and supply chains, embedding such technologies within decision-making processes can greatly aid and enhance them. This is because such technologies unearth the overall system behavior, resource utilization, and performance of both actual and speculative scenarios [2].

Several case studies show how these types of models are used to locate bottlenecks in operations, increase the sustainability of processes by tracking emissions, and reduce the impact of disruptions, like during the COVID-19 pandemic, where digital twin-based simulations helped to reconfigure networks for resilience [3].

In the work presented here, the terms "factory" and "supply chain" are used in a loose sense for simplicity. "Factory" represents a unit accepting one or more inputs, with one or more processes, converting those input(s) to one or more outputs. "Supply chain" is a group of those units sharing some resources. The SCDT tool adopts this generic approach to the modeling of a factory. This broadens the bandwidth of modeling capabilities, and can easily be scaled up in complexity to model one or multiple supply chains. The SCDT considers the developed visualization as a series of inputs, processes and outputs. The inputs and outputs are represented as material buffers, while the processes timely convert quantified inputs to outputs. The level of visualized detail can vary, according to provided information. Material buffers and processes can be daisy chained to accomplish higher degrees of complexity and scale. In this way, in the modeling of a supply chain, one of the factories can have a detailed representation of all its processes and materials, while another factory, for which limited details are available, can be represented as a generic input-output process following a reduced order model abstraction. For the flow of materials between factories, the concept of "Logistics Factory" is used. This is defined as a factory where each process represents a transfer of material from one factory (or a supplier model) to another factory. Each process in the "Logistics Factory" could represent an actual Logistics organization or the logistics department of one of the visualized factories.

Under the proposed framework, a number of components are grouped as a factory, a number of factories are grouped as a supply chain, and several supply chains can be represented with material flows between components, factories

<pre> "Inputs": { "value": ["ByProduct_B4", "Consumable_B1", "Electricity_B1", "Water_B1"]}, "InputsQ": { "value": [1.0, 0.5, 1.85, 1.5]}, "InputsU": { "value": ["u", "kg", "kwh", "Lt"]}, "Outputs": { "value": ["Product_B1", "Waste_B3", "Waste_B2"]}, "OutputsQ": { "value": [1.0, 0.4, 0.4]}, "OutputsU": { "value": ["u", "kg", "kg"]}, "Timing": { "value": {"duration": 1.0 }}, </pre>	<pre> "quantity": { "value": 0.0, "unitCode": "kg/h"}, "isActive": {"value": 1}, "isEnabled": {"value": 1}, "isFaulty": {"value": 0}, >Status": {"value": 0}, "Inputs": {"value": [] }, "Outputs": { "value": ["Process_B5", "Process_B4", "Process_B3", "Process_B2", "Process_B1"]}, </pre>	<pre> "Inputs": { "value": [["SC_One->Factory_A->Product_A1"]}, "InputsQ": { "value": [5.0]}, "InputsU": { "value": ["kg"]}, "Outputs": { "value": [["SC_One->Factory_B->Material_B1"]}, "OutputsQ": { "value": [5.0]}, "OutputsU": { "value": ["kg"]}, "Timing": { "value": {"duration": 1.0}}, </pre>
(a) Manufacturing Process	(b) Material Buffer	(c) Logistics Process

Fig. 1: Partial definitions for representation

and supply chains.

The rest of the paper is organized as follows: Section II describes the factory and logistics representation framework developed for the SCDT. Section III presents the overall SCDT system architecture and simulation workflow. Section IV focuses on factory and supply chain visualization and interaction. Section V introduces representative supply chain case studies and discusses observed bottlenecks and sustainability-related insights. Finally, Section VI concludes the paper and outlines directions for future work.

II. FACTORY REPRESENTATION

For the development of a supply chain digital twin, multiple manufacturing units/factories need to be represented. Factory representation in the SCDT is achieved via the use of JSON-LD templates. Factory data is transferred to a JSON-LD template, that in turn can be loaded directly to an Orion-LD server. Each factory is defined in a JSON-LD file that is used for both, real time data and simulation of predefined scenarios. The JSON-LD file contains various entities that in their totality describe the factory and its components. A “Factory” entity describes general factory parameters including, type of factory, supply chain name to which the factory belongs to, and list of factory components that constitute the factory. The factory components are divided into two main categories. Processes and material buffers. Each “FactoryComponent” is described in the JSON-LD file as a separate entity. Figure 1 shows partial definitions of factory components as described below. Some of the fields’ attributes necessary for ngsl-d representation were removed for clarity.

Each process has a timing setting that is associated with the time required for the process to convert inputs to outputs. The inputs and outputs are listed in the process representation (JSON-LD file) along with the units and quantities required to produce the quantity output. Figure 1(a) shows the partial definition of a factory process. The field “Inputs”, lists the material buffers that are inputs to the process, the field

“InputsQ” lists their corresponding quantities and the field “InputsU” defines their units of measurement. The process converts those inputs to the corresponding “Outputs”, “OutputsQ” and “OutputsU” according to the “Timing” definition.

Each material buffer is either an input to a process, an output of a process or both (processes in series). Depending on this, they are classified as inputs to a factory, products of a factory or intermediate components. Intermediate components can be available to market (can connect to other factories) or not. Figure 1(b) shows the partial definition of a material buffer. The inputs and outputs of a material buffer are the processes that produce the material and use the material respectively. The shown material buffer is an input to the factory. It is not produced by a factory process (empty “Inputs” list). The “Outputs” field lists the processes that use this material buffer. In the example definition there are 5 processes that use this material buffer for manufacturing.

In order to transfer quantities between material buffers of different factories, a “Logistics Factory” type can be defined to model supply-chain level logistics processes. A “Logistics Factory” does not have its own materials buffers. It only has processes. These processes have inputs and outputs from material buffers of a different factory that could be in the same or different supply chain. Figure 1(c) shows the partial definition of a process of a “Logistics Factory”. The process removes material from a material buffer in supply chain “SC_One” factory “Factory A”, and adds the material to a material buffer in supply chain “SC_One” factory “Factory B”.

III. ARCHITECTURE

The SCDT module has a continuously running monitoring node. This node is responsible for monitoring and capturing user requests from a remote server (data platform). The request includes data that describes the visualization requested. The visualization can be of a single factory in a supply chain, multiple factories in the same supply chain




Fig. 2: SCDD Architecture

or multiple factories in multiple supply chains. The visualization can utilize real-time data supplied from the remote server, or generate model driven data based on internal simulations of models created from information obtained from factories. The monitoring node sends an update to the data platform, setting the request status as “Received”, and proceeds to developing a dedicated visualization. Every dedicated visualization is developed using its own docker containers (including a local OrionLD container). The factory and supply chain information of the request is loaded to the dedicated Orion LD server. A ROS2 node then scans the information on the local Orion LD server, develops the Gazebo URDF files required for the visualization and starts a gazebo server. A second ROS2 node converts information from the Orion LD server to Gazebo visualization. In the case of real time monitoring, the information on the local container is updated by a node that links information on the Data Platform to values inside the container. In the case of simulated visualization, a ROS2 node inside the container reads the properties of the processes and material buffer quantities and updates asynchronously the local Orion LD with simulated data based on the information given. The updates are asynchronous, since every manufacturing process runs independently. This enables the seamless scale up to larger factories and supply chains, while ensuring the consistency of the simulation. When the simulation is ready, the monitoring node updates the request entity on the data platform. It provides a URL link and password for the user and sets status to ready. The user can view and interact with the Gazebo visualization through a web interface using their

browser. When the user logs out of the simulation in their browser, the dedicated visualization container is destroyed. The request is set to status closed and is disabled. Figure 2(a) shows an Overview of the architecture adopted for the SCDD as described above. Each user request produces local container(s). Figure 2(b) shows the information/command flow of the local containers of the SCDD. Each container(s) represents a dedicated visualization.




Fig. 3: Factory Visualization

IV. FACTORIES VISUALIZATION

The factory representations are implemented by ROS2 Nodes as detailed in Section III. The material buffers are not all either inputs or outputs of a factory. Some are waste by-products that are not available to the market. Some could be intermediate components, produced by a factory process to be used by other processes within the same factory for product manufacturing. Nevertheless, these intermediate components could be available to market or not. In order to

reflect this, each manufacturing factory in the visualization has four rows of components as shown in Figure 3.

Figure 3 shows the visualization of a factory. From right to left, the four rows shown are:

- Row 1: the input material buffers coming from suppliers to the factory
- Row 2: the processes of the factory,
- Row 3: the material buffers that are produced by processes and are not products of the factory,
- Row 4: the final products of the factory

The material buffers in Row 3 are divided into two categories. The ones that are available to market shown in blue texture to indicate market availability and blinking green to indicate activity, and the ones that are not available to market shown in brown texture to indicate no availability to market and blinking green to also indicate activity. The ones that are available to market could be scrap/byproducts that could be used by another factory as material, or material that is used by an internal process but can also be used by another factory. In simulation mode, the processes convert quantities from material buffers to output material buffers according to the process definition of quantities and timing. While not restricted by the proposed architecture, in the presented implementation, the simulation time to actual execution time mapping is as follows: 1 second in simulation time maps to 1 hour in actual execution time. In the event that one of the input material buffers does not have enough quantity needed by the manufacturing process, the process halts. The material buffer and the process of the factory turn red, indicating the process pause. Depending on what is selected by the user, information panels can display factory, component and scenario information.

The scenario information is shown in two information panels on the top left and middle of every figure. The information panels are not shown in Figure 3 for clarity. The information details are shown in Figures 5a and 5b. The actual position of the information panels is shown in Figure 7

The top left information panel (Figure 5a) shows information about the Factory selected and the component selected. In the indicative example, a Factory with Name "Factory_C" was selected that has 10 components. Out of the 10 components a component named "Process_C1" was selected. In the "Factory Properties" section of Figure 5a (second column- top) the properties of the components are listed. All the available quantities and their units are displayed for the material buffers, while the duration and any emissions are displayed for the processes. For the selected component ("Process_C1") more details are shown on the second column at the bottom. The list shows all input and output material buffers along with the required quantities for conversion of inputs to outputs.

The top middle information panel (Figure 5b) shows information about the running scenario. The indicative example shows a simulation of a scenario named "Case 1" The panel also shows the total emissions (all processes involved in the scenario) and active emissions (from active processes). The

panel also displays a "Product Emissions Indicator" for every product of the factory. This can be derived from separate tools (e.g. LCA analysis). All indicators and units are for display purposes and they are not used for any calculations in this version of the SCDT. The processes asynchronously convert arbitrary quantities of inputs to quantities of outputs[9].




Fig. 4: Factory Material Flow

Figure 4 shows the visualization of a different factory in the same supply chain. When the user selects a component (mouse click), arrows appear showing the material flow for that particular component. In the specific example, "Process_C1" was selected and arrows in Figure 4 show the input and output material buffers for the process. The selected process is shown in red indicating inactivity. This is because after operating for a certain amount of time, Factory_B run out of Consumable_C2 which is also shown in red. When the material buffer has its quantity increased (i.e through a logistics schedule), the process will resume, and start flashing green indicating activity. The available and required by the process quantities are shown on the left information panel (Figure 5a). In the case of networked suppliers, e.g water and electricity, material buffers can be defined as such, allowing negative values for the buffer quantity, and thus capturing the total consumption during the simulated duration. The user can reset the simulation time and quantities to the initial values using a "Reset" button (Figure 5c), thus retesting the defined scenario.

Figure 6 shows a "Logistics Factory" with 3 processes. Each process represents one material transfer of a specific quantity from one factory to another. The source factory could also be any supplier not shown in the visualization. As an example, the selected logistic process transfers a certain amount of units (e.g. 10 items) every a certain amount of time (e.g 5 hours), of Factory B product "Product_B1" to Factory C to be used as input for manufacturing (arrow to Factory_C, Material_C1). The process turned red to indicate inability to process, since there are not sufficient units available to complete the scheduled transfer, thus highlighting a bottleneck (inability to process transfer) in the supply chain. However, the factory is still producing. When the quantity of the product of the factory reaches the required amount, as defined by the logistics process, it will return to active status, and the transfer will be performed.




Fig. 5: Information Panels and Simulation Reset




Fig. 6: Logistics Factory

V. SUPPLY CHAIN CASE STUDY

In order to demonstrate the operation of SCDT a case study with two scenarios was developed. The case study uses artificial as well as anonymized factory data obtained through the CIRCULOOS project, such as process timings, material quantities, and conversion ratios. The full list of implemented data is too extensive to be included in the current, and for the needs of the presentation a set of key data is listed in Table I.

The test case considers three hypothetical factories ("Factory_A", "Factory_B" & "Factory_C"). All factories have inputs, processes, and outputs. Outputs include one product for each factory, scrap material that can be made available to other factories, and waste materials (i.e. not available on the supply chain). A fourth factory named "Logistics_Factories"

has the processes responsible for transferring materials between factories and suppliers. Figure 7 shows the four factories that constitute the visualization of a supply chain.

Table I lists the number of each type of components in each Factory.

TABLE I: Factory Components

	Input Buffers	Processes	Output Buffers
Factory_A	4	1	2
Factory_B	4	5	8
Factory_C	6	1	3
Logistics_Factories	-	4	-

The "Logistics_Factories" on the top right of Figure 7 is shown to have four processes. In more detail, from left to right, the logistics processes are:

- **ProcL1.ProdA1.to.MatIB1:** Five units of "Product_A1" are transferred from "Factory_A" every hour to "Factory_B" to be used as input "Material_B1"
- **ProcL2.ProdB1.to.MatIC1:** Five units of "Product_B1" are transferred from "Factory_B" every hour to "Factory_C" to be used as input "Material_C1"
- **ProcL3.Suppl.to.MatIA1:** Fourty units of "Material_A1" are transferred from an external supplier (external to the visualization) every 40h to "Factory_A" to be used as input "Material_A1",
- **ProcL3.ScrapC1.to.MatIA1:** Five units of "ByProduct_C1" are transferred from "Factory_C" every hour to "Factory_A" to be used as input "Material_A1".

All other factory processes have a timing of 1hr converting defined quantities of inputs to outputs.

In other words, "Factory_A" supplies its product "Product_A1" to "Factory_B", that uses it to manufacture its own product "Product_B", that in turn is supplied to "Factory_C" to produce its own product "Product_C". This establishes an inter-dependency of the three factories. Factory_B depends on the production of Factory_A and Factory_C depends on the production of Factory_B and subsequently to the




Fig. 7: Logistics Factory with Disabled Recycling Process

production of Factory_A. "Factory_C" also has the option to supply or not, one of its scrap outputs ("ByProduct.C1") back down the supply chain to "Factory_A" to be used as manufacturing material "Material_A1". The test case considers two scenarios. For both scenario visualizations, the factory definition files are the same with the same process timings and initial quantities. The difference is that in the first scenario with no recycling of "ByProduct.C1", the Logistics Process ("ProcL3_ScrapC1_to_MatrlA1") responsible for transferring the produced by-product is disabled, while for the second scenario where "Factory_C" recycles "ByProduct.C1" to "Factory_A", the process is enabled.

Figure 7 shows the scenario where Factory_A does not recycle its scrap "ByProduct.C1". The process responsible for the recycling in the "Logistics.Factories" is shown in black (disabled process). This indicates that even though the specific process is included in the factory definition, it is not used in the scenario visualization.

This first scenario is used to test the reset and repeatability characteristics of the SCDT. A test run of over 22000 production hours is shown in Figure 8. The visualization was reset nine times, as can be seen from the quantity of "Product.C1". This is a monotonic value since no other process in the supply chain uses it, and thus is a good candidate for comparing repeated runs and scenarios. The reset in the repeatability test is achieved by stopping all processes, reinitializing buffers to initial quantities, and restarting the scenario processes.

Figure 9 shows the "Product.C1" quantities of Figure 8 plotted one on top of each other for comparizon. Any small deviations (shown in zoomed insert) are associated with the differences between recording intervals and process timings

intervals (x-axis timing shift).

Any value or property of any of the components simulated can be recorded, allowing processing and analysis of manufacturing data. For the purpose of comparing the two tested scenarios, and showcasing the capabilities of the SCDT, one output one process and one input for every one of the manufacturing factories was recorded for the two cases of "No Recycling" and "Recycling" of "ByProduct.C1".

Figure 10 shows the plotted results. The curves show the quantities of "Product.C1" for both cases of recycling (Product.C1.Scen2.quantity) and no-recycling (Product.C1.Scen1.quantity) of "ByProduct.C1". Initially, due to initial quantities in the material buffers of all factories, the production in "Factory_C" is similar for the two scenarios. However, after approximately 300 hours of operation, the production of "Factory_C" starts to deviate, with the "Recycling" scenario giving higher production rate. This is because the various factories have run out of initial material quantities in their buffers and the supply chain is only sustained by the logistic processes. The recycling covered bottlenecks of factories "Factory_A" and in turn "Factory_B" down the supply chain, that in turn helped "Factory_C" to overcome its own bottleneck issues of the supply of "Material.C1" from "Factory_B" (logistics process "ProcL2_ProdB1_to_MatrlC1"). The SCDT is shown to be stable in the process timings and quantity calculations as can be seen from the lines of constant gradient (after initial quantity usage) along the simulated timeline.

The quantities plotted in Figure 10, show the material buffers levels as recorded in simulated 5h intervals. After approximately 3250 hours of simulation, "Factory_C" produced 800 units of "Product.C1" without Recycling, while in




Fig. 8: Resetting a Visualization




Fig. 9: Repeatability Comparizon




Fig. 10: Quantity Recording for both Scenarios

the recycling scenario produced 1055 units during the same simulation hours. For both material buffers and processes, a status flag was also recorded. This flag indicates whether a component is 0:idle, 1:active, 2:inactive and 3:disabled. During a simulation, whenever one of the input (to a process) material buffer does not have enough quantity for the process to continue, the processes status flag goes from 1:active to 2:inactive. The material buffer also goes inactive (both visualized with red coloring of the component) to indicate which of the inputs is the one that halts the process (bottleneck). If the material buffer quantity increases to the necessary levels, the processes status flag goes from 2:inactive to 1:active, and the process continues.

Figure 11 shows an example of selected process flags toggling between 1 (active:normal operation) and 2 (inactive:

not enough material quantity for the process). Figure 11 shows the complexity of the manufacturing operations in the simple supply chain example test case. The flags shown are only for three of the processes for eac of the two scenarios. There is one process flag for each of the three manufacturing factories. The top three plots are for the "NO Recycling" Scenario1, and the bottom three plots the "Recycling" Scenario2, so that each flag is shown for both scenarios.

The Supply Chain and Factories were setup in such a way to have exaggerated shortage of material buffers to showcase the capabilities of the SCDT Tool. As detailed earlier, the only difference in the two Scenarios is the extra transfer of material from "Factory.C" recycling its "ByProduct.C1", back down the supply chain to be used as




Fig. 11: Status Recording for both Scenarios

input "Material_A1" of "Factory_A". This extra transfer of material is the one enabling "Process_A1_Scn2", i.e. the first process of Factory_A in Scenario 2, to stay always active (6th plot from the top of Figure 11) The corresponding Scenario 1 plot (3rd plot from the top) shows continuous interruptions of the manufacturing process. This down-time of the manufacturing process translates to increased down-time to the process of "Factory_B" that uses the final product of "Factory_A" as an input material. In turn, the increased down time translates to increased down-time to the process of "Factory_C" that uses the final product of "Factory_B" as an input material. Comparing the two Scenarios for process "Process_C1_ScnX_Status", for the duration of Factory Operation, Scenario 2 showed a calculated increase in Up-time from 640 hours to 840 hours, reflecting the increased production shown in Figure 10.

VI. CONCLUSIONS

The SCDT tool is a versatile scalable platform for supply chain digital twinning that can be used for scenario analysis, bottleneck identification, and sustainability oriented assessment. It can utilize factory representations of various abstraction levels and detail, to represent one or more interacting supply chains.

It is capable of near real-time monitoring of factory operations in conjunction with enabling simulation of various scenarios, such as logistics and recycling schemes.

The consistency of the SCDT was demonstrated through a high degree of repeatability. The inbuild asynchronous nature of the system, allows SCDT to be versatile and scalable, since adding more factories and components is straight forward. Adding factories, disabling/enabling components and modifying logistic transactions can be achieved while running a simulation. In the supply chain test case, the hypothetical arrangement of three different factories (A, B, and C) demonstrated how the software could be used to find

the operational bottlenecks and evaluate the sustainability impacts.

Future work concentrates on developing and augmenting automation and optimization features for the SCDT.

ACKNOWLEDGMENT

This work was developed under the CIRCULOOS project that was funded from the European Union's HORIZON –CL4-2022-TWIN-TRANSITION-01-06 Innovation Actions programme under grant agreement No 101092295

REFERENCES

- [1] M. Grieves and J. Vickers, *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems*. Springer, 2017, pp. 85–113.
- [2] F. Tao and M. Zhang, "Digital twin shop-floor: A new shop-floor paradigm towards smart manufacturing," *IEEE Access*, vol. 5, pp. 20 418–20 427, 2017.
- [3] L. Wang, T. Deng, M. Shen, H. Hu, and Y. Qi, "Digital twin-driven smart supply chain," *Frontiers of Engineering Management*, vol. 9, 01 2022.
- [4] D. Ivanov and O. Gusikhin, "Supply chain digital twin design and implementation at scale: A case study at the ford motor company and generalizations," *Omega*, vol. 139, p. 103447, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0305048325001732>
- [5] Ö. Tozanli and M. J. Saénz, "Unlocking the potential of digital twins in supply chains," *MIT Sloan management review*, vol. 63, no. 4, pp. 1–4, 2022.
- [6] S. A. H. Zaidi, S. A. Khan, and A. Chaabane, "Unlocking the potential of digital twins in supply chains: A systematic review," *Supply Chain Analytics*, vol. 7, p. 100075, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2949863524000189>
- [7] A. T. Wasi, M. A. Anik, A. Rahman, M. I. Hoque, M. S. Islam, and M. M. Ahsan, "A theoretical framework for graph-based digital twins for supply chain management and optimization," 2025. [Online]. Available: <https://arxiv.org/abs/2504.03692>
- [8] F. Freese and A. Ludwig, "A conceptual framework for supply chain digital twins – development and evaluation," *International Journal of Logistics Research and Applications*, vol. 28, no. 6, pp. 676–698, 2025. [Online]. Available: <https://doi.org/10.1080/13675567.2024.2324895>
- [9] V. Xiros, J. M. Gonzalez Castro, F. Fernandez-Pelaez, B. Magoutas, and K. Christidis, "Supply chain data analytics for digital twins: A comprehensive framework," *Applied Sciences*, vol. 15, no. 12, 2025. [Online]. Available: <https://www.mdpi.com/2076-3417/15/12/6939>