ENOUGH

EUROPEAN FOOD CHAIN SUPPLY TO REDUCE GHG EMISSIONS BY 2050

PROJECT ACRONYM	ENOUGH
PROJECT TITLE	European Food Chain Supply to reduce GHG emission by 2050
PROJECT COORDINATOR	SINTEF OCEAN
PROJECT DURATION	2021-2025

D5.7

Full release of the SDS

PROJECT NO: 101036588

TYPE OF ACTION: IA (INNOVATION ACTION)

CALL: H2020-LC-GD-2020-4

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101036588

DOCUMENT INFORMATION

DELIVERABLE NO.	5.7
DISSEMINATION LEVEL	Public
WORK PACKAGE	WP5
TASK	T5.6
LEAD BENEFICIARY	UNIVPM
CONTRIBUTING BENEFICIARY(IES)	SO, LSBU, UoB, NTNU, INRAE, KUL, TUGraz, ELET, EMIL
DUE DATE OF DELIVERABLE	31/03/2025
ACTUAL SUBMISSION DATE	23/09/2025

DOCUMENT HISTORY

Version	DATE	CHANGE HISTORY	AUTHOR	BENEFICIARY
1.0	19/03/2025		Massimiliano Pirani, Luca Spalazzi, Alessandro Carbonari, Alessandro Cucchiarelli, Tariq Naeem, Felipe Maté, Luca Cesarini, Manuel Verdnik, Fernando Negozi, Enrico Gambini, Matteo Di Stefano, Andrea Fadda, Dìlan Durmus, Loris Bottegoni, Leonardo Cambiotti, Valerio Crocetti	UNIVPM, TUGraz, ELET

QUALITY ASSURANCE, STATUS OF DELIVERABLE		
ACTION	PERFORMED BY	DATE
Reviewed	Reviewed internally by the team, reviewed by Advisory Board Member Dr. Sebastian Scholze	16/09/2025
Approved	Massimiliano Pirani	23/09/2025
Uploaded to SyGMa	Kristina N. Widell	23/09/2025

Table of Contents

EXECUTIVE SUMMARY

This deliverable presents the developments and implementations achieved in the context of work package 5 (WP 5 - Smart data systems). Major outcome of the activities is a framework called SDS (Smart Data System).

The SDS is a business process that comprehends a platform, specifications, a set of digital tools and applications, and paradigms aimed to the long-term exploitation of the results of the ENOUGH project.

In practical terms, during the developments of the ENOUGH project, four major results have been achieved:

1. specification of the SDS multi-party and multi-sided business for the creation of sustainable food supply chain markets and businesses;
2. design and development of the SDS infrastructure and relative platform, able to manage and conduct the life cycle of Supply Chain Entities, which constitute a set of digital twins of relative food supply chains;
3. a layer of the SDS application that introduces advanced blockchain technologies in the SDS business and in the food supply chain realities;
4. a layer for distributed decision making, in particular a demonstrative component for holistic decision making in complexity, featuring Mixed Reality and holonic management methods.

Taken altogether, the work of WP5 of the ENOUGH project, culminating in the presentation of this report, has demonstrated the possibility of a paradigm shift in the use and conception of information technologies. A space and a pathway have been created that facilitate the achievement of the objectives proposed by the context of Industry 5.0, both at the European and global level, so that the reality of practice and research for improved sustainability of food supply chains can be advanced through inclusiveness, openness, participation, and innovation.

ABBREVIATIONS

AR	Augmented reality
BC	Blockchain
BPMN	Business Model Language Notation
CI/CD	Continuous Integration/Continuous Delivery and Deployment
CIP	Cleaning-in-place
DBMS	Database management system
DCS	Distributed Control System
FSC	Food supply chain
GHG	Green House Gases
HMT	Holonic Management Tree
HTHP	High temperature heat pump
JSON	JavaScript Object Notation
KPI	Key performance indicator
MR	Mixed Reality
MVP	Minimum Viable Product
OEE	Overall equipment effectiveness
OTE	Overall throughput effectiveness
RMAS	Relational-model Multi-Agent System
SC	Supply chain
SCE	Supply Chain Entity
SDG	Sustainable Development Goals of the United Nations
SDS	Smart Data System
SDS DB	SDS Database management system
SDS DDM	SDS Distributed Decision Making
UML	Unified Modelling Language
WP	Project work package

Full release of the SDS

1 INTRODUCTION

This report concludes and summarizes the activities performed in Task 5.6 in the context of Work Package 5 (WP5) of the ENOUGH project. The overarching ambition of WP5 has been to design, implement, and demonstrate the Smart Data System (SDS), a key innovation aimed at enabling sustainable digital transformation across food supply chains.

The primary objective was to establish a strong and distinctive and unique concept of the SDS that could go beyond the numerous digital platforms and reference architectures already funded by the EU in past decades. Many of these initiatives delivered valuable technological results, yet often remained siloed, overly specialized, or insufficiently adaptable to the complexity of real-world supply chains. WP5 therefore set out to conceive the SDS not as "just another digital platform," but as a systemic enabler: a holistic and evolving digital backbone of services capable of interlinking stakeholders, harmonizing heterogeneous data, and fostering decision-making that is both collaborative and nudge-driven—using behavioural insights and subtle incentives to steer participants toward more sustainable or optimal choices, without coercion.

At its core, the SDS was envisioned as a digital and virtual backbone of run-time services that can interoperate with stakeholder components in a unified and co-creative manner. This architecture integrates sustainability "by design," reflecting the momentum of the emerging Enterprise 5.0 paradigm, where human-centric, resilient, and inclusive practices complement industrial digitalization. The SDS concept therefore deliberately balances technological sophistication with usability and accessibility, ensuring that stakeholders across domains, sizes, and levels of maturity can adopt and benefit from it.

Developing such a system within an evolving and highly complex domain—namely the digitalization of food supply chains—required methodological rigor. Concepts such as Digital Twins, which remain fluid and continuously redefined, demand an open and adaptive framework. For this reason, the SDS was designed as both stable and flexible: its mission and backbone remain solid, while its interfaces and modules remain open to iteration and innovation.

To achieve this balance, WP5 adopted an Agile development approach, aligned with scientific and industrial best practices. Agile principles—collaboration, adaptability, frequent deliveries, and simplicity—ensured that stakeholders were actively engaged in the process. In practice, the approach emphasized progressive validation with end-users, minimized reliance on static documentation, and treated the codebase and prototypes as the primary references for design. Agile was complemented by Kanban, which provided a lightweight yet effective framework for managing workflows, prioritizing tasks, and accommodating the unpredictable rhythms of collaborative research and industrial engagement.

Within this context, requirement elicitation relied heavily on user stories: concise, stakeholder-focused descriptions of functional needs that express increments of business value. Represented as Kanban cards, user stories became the backbone of communication between stakeholders, designers, and developers. This method supported early validation, reduced unnecessary complexity, and enabled continuous prioritization, while also aligning with the pull-production logic that underpins many modern supply chain practices.

Through this process, the SDS gradually evolved from a conceptual framework into a business-ready system. Its design emphasizes not only technical robustness but also long-term exploitability, ensuring medium- and long-term impacts beyond the duration of the ENOUGH project. By combining digital twins, blockchain-enabled trust mechanisms, marketplaces for innovation exchange, and decision-support paradigms rooted in systemic thinking, the SDS positions itself as both an infrastructure and a catalyst for business and societal transformation.

The body of this deliverable documents this evolution in detail. Section 2 introduces the vision and the methodology adopted for the development of the SDS concept along with major details on the resulting framework. Section 3 introduces the driving business cases, outlining the central role of Supply Chain Entities (SCEs), the importance of open and inclusive marketplaces, and the embedding of systemic paradigms for sustainability. Sections 4 to 7 expand on four major business use cases: the life cycle management of SCEs, the purchasing of services from the SDS Marketplace, the provision of third-party solutions into the marketplace, and the channeling of research-driven innovations into practice. These use cases illustrate how the SDS framework can enable configuration, monitoring, negotiation, interoperability, and systemic innovation.

Section 8 then presents the lessons learned during the development and application of the SDS. It provides recommendations tailored for industry (focusing on adoption and competitiveness), policymakers (emphasizing systemic interoperability and regulation), and society (stressing trust, transparency, and sustainability). Together, these recommendations highlight the SDS as a socio-technical enabler rather than merely a technical platform.

Overall, this report offers a landscape that spans conceptual design, implementation, business validation, and policy implications. It demonstrates how the SDS contributes not only to the specific objectives of the ENOUGH project but also to the broader aspirations of Industry 5.0: enabling sustainable, human-centric, and resilient food supply chains through inclusive and innovative digital transformation.

Terminology and Definitions

To avoid ambiguity, key terms are consistently used throughout this deliverable.

– **SDS (Smart Data System, concept):** the overarching paradigm that combines methodologies, governance principles, and digital infrastructures to enable sustainable and systemic management of food supply chains. The concept is being materialized and enforced with a framework that consists in a workflow and a set of specifications to be satisfied to participate in the SDS context.

– **SDS Platform:** the technical instantiation of the SDS concept, including software components, connectors, and interfaces. In particular, the SDS Platform is the environment in which the entities that digitalize a supply chain are configured, created, and operated.

– **SCE (Supply Chain Entity):** the informational and digital representation (digital twin) of a real supply chain or part thereof.

– **Marketplace:** the digital environment within the SDS where stakeholders can exchange services, algorithms, or datasets.

– **CI/CD:** Continuous Integration/Continuous Delivery, applied in the SDS Platform to the life cycle management of SCEs.

2 THE PATHWAY TO SDS BUSINESS, PARADIGMS, AND TECHNOLOGY

2.1 Design thinking approach for the development of the SDS concept and specification

The concept of the Smart Data System (SDS) began with the idea of a platform and infrastructure capable of uniting and integrating various levels of knowledge, information, and data from all stakeholders involved in a supply chain, but above all, a living repository of the results and research from all the work packages of the ENOUGH project.

Primarily objective was to gain a strong concept of the SDS that would go beyond the already vast set of digital platforms and reference architectures that the EU has funded in the last decades, through many work programmes.

The idea of work package 5 (WP5) was to design and develop the SDS as a digital and virtual backbone of services that constitute a holistic set of run-time services, in order to interoperate stakeholders' components in a unified and co-creative way, with sustainability inside (looking towards the growing momentum around the Enterprise 5.0 initiative worldwide). This could be achieved only if the SDS takes the form of a collaborative framework in which stakeholders can find convenient use, suitable for their language, domain, and scale.

The main objective of WP5 was to create the SDS as a complex system within an equally complex and constantly evolving domain. As the digitalization of the food supply chain progresses and concepts such as the Digital Twin remain fluid, the SDS must stay open to change while maintaining a solid core and mission.

An Agile development approach was chosen, in line with best scientific and industrial practices, prioritizing collaboration, adaptability, frequent deliveries, and simplicity, following the twelve principles of the Agile Manifesto. In practice, this entails strong interaction between stakeholders and developers for progressive validations, reduced reliance on formal diagrams in favor of the code as the primary reference, and the use of diagrams mainly during the requirements translation phase. Considering the distribution of roles in the ENOUGH project, Kanban complements Agile in managing the workflow.

In WP5, the primary tool for requirement elicitation in Agile development was the user story, as a piece of functionality understood by the potential customer or product owner and representing an increment of business value. These stories captured the needs and expectations of stakeholders for the SDS and were represented as Kanban cards. This approach fostered collaboration among stakeholders, designers, and developers, provided an initial but effective view of requirements without excessive detail, and facilitated both planning and prioritization. The use of Kanban aligned with the pull production concept, enabling delivery in a continuous flow, which proved particularly effective in situations where work arrived unpredictably and where it was beneficial to deploy as soon as items were ready.

The core idea is that stakeholders interact with what was called the SDS Marketplace, generating stories that serve as the main drivers for requirements and specification refinements. These inputs inform and guide developers in producing components and artefacts that can meet the evolving needs. Stakeholders can also act as third party software producers, integrating their own solutions into the SDS. Kanban additionally acts as a support for a design thinking approach, enabling co-creation, iterative refinement, and adaptation of the SDS, with a vision extending beyond the project's official end.

The process of construction of the SDS business process framework began with an Elicitation phase, during which stories were collected, reordered, and refined until they reached a sufficient level of maturity. Once this threshold was met, the Definition phase started, producing sequence diagrams to capture behavioural and communication flows, and an architecture description defining the structural constraints and capabilities. These two artefacts influenced one another and sometimes required a return to elicitation for further refinements.

A third key element of the SDS specification was the data model, whose features and priorities were determined by the feasibility and cost assessments emerging from the Definition phase. These priorities were influenced by business strategies, market demands, and the broader SDS vision. The data model was inherently open-ended, with its content shaped by evolving strategic and contextual needs, and it occasionally introduced new requirements that led to further cycles of definition and refinement. In Figure 2.1.1, are shown the major steps and workflow that led to the specifications of the SDS by design thinking and elicitation of requirements.

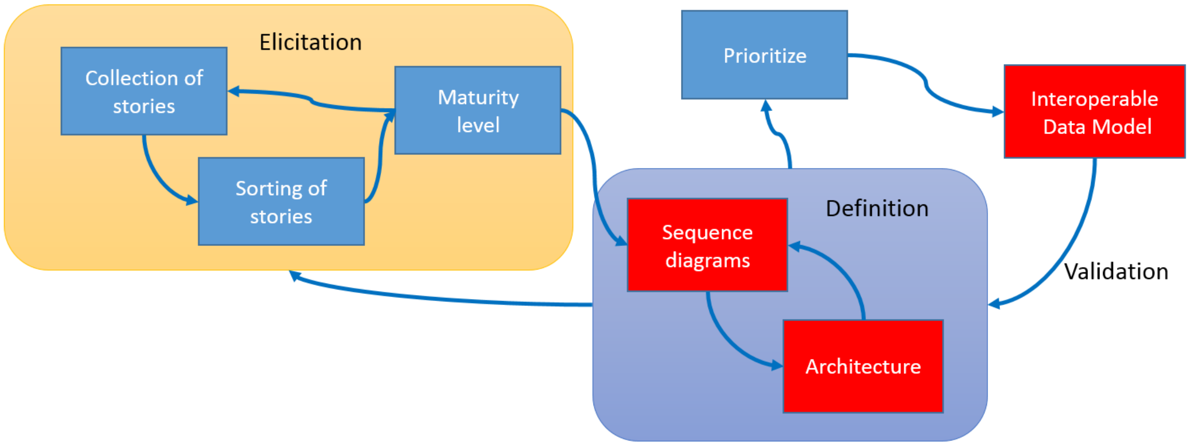

Figure 2.1.1: Workflow from requirement analysis to specifications and their iterative refinement in the SDS design process.

Throughout this process it came out that the essence of the SDS system is indeed its capability to include and interoperate software modules and components from third parties and coordinate them into a purposeful holistic goal of continuous improvement.

By design, the SDS can grow further in many directions and include many other kinds of modules. Nonetheless, the modules here already present constitute a minimum viable set of components that can showcase the potential of the concept.

2.2 Stable concept of the SDS

A first concept and requirement shared during the design phase was that the SDS sees the domain of the food supply chain as a complex one, with a fractal structure of the problems. This view comes from the UNIVPM research and previous experience in the industrial and construction automation playground. In addition, it is an outcome of previous EU-funded research conveyed as background in the ENOUGH project. In essence, the supply chain in the Farm2Fork context and discourse is seen as a system of systems in which a distributed control and decision making is overlaid as a cyber-physical system of systems[1,2]. The picture that was discussed in meetings to this point is shown in Figure 2.2.1. It expresses rather abstractly that the whole circular process of the Farm to Fork can be decomposed into a self-recurrent problem of complexity "decomposition", although not in a reductionistic way but holistic. The dashed boxes in the figure refer to the process of decomposing the whole system into subsystems of systems. The SDS needs the inherent capability in the design to support such a holistic and systemic view of the supply chain.

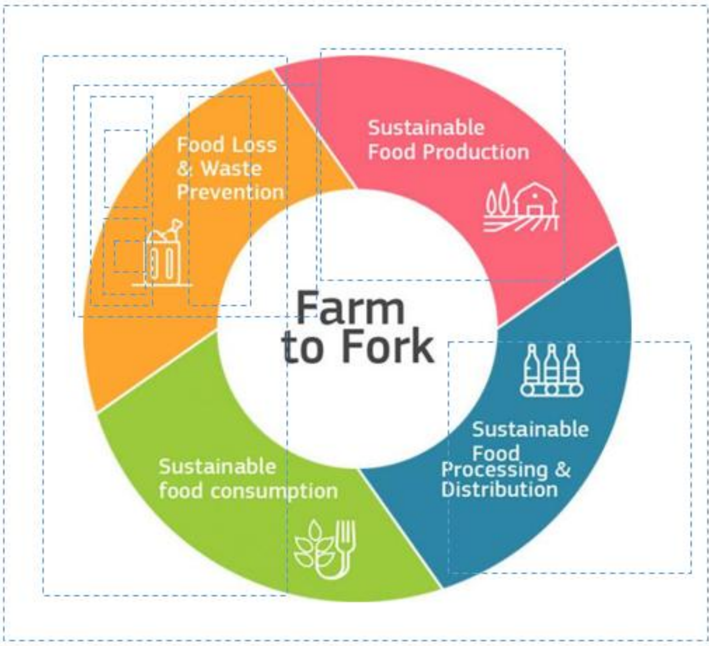

Figure 2.2.1: Farm2Fork represented as a complex system of systems, with decomposition into subsystems illustrated by dashed boxes.

Since the first Elicitation interactions and exchanges between project partners and the beneficiaries directly involved in the WP5 activities, the following set of keywords have arisen:

- low carbon footprint / low GHG emission;
- real-time;
- trust;
- human-centric;
- socio-technical problems;
- circularity and sustainability;

Those concepts created a further step for UNIVPM partner's view of the problem and the pathway for effective solutions. In Figure 2.2.2, a structure of the problem and a vision of the supply chain as a chain of trust has been put forth in the Elicitation phase. In this figure it is still clear how the supply

[1] Bonci, A., Pirani, M., Carbonari, A., Naticchia, B., Cucchiarelli, A., & Longhi, S. (2018, September). Holonic overlays in cyber-physical system of systems. In 2018 IEEE 23rd International Conference on Emerging Technologies and Factory Automation (ETFA) (Vol. 1, pp. 1240-1243). IEEE.

[2] Pirani, M., Carbonari, A., Cucchiarelli, A., Giretti, A., & Spalazzi, L. (2024). The Meta Holonic Management Tree: review, steps, and roadmap to industrial Cybernetics 5.0. *Journal of Intelligent Manufacturing*, 1-42.

chain is seen as a composition of steps that bear the same problem structure, in a self-similar fashion. For each of the steps we have to consider a short-term layer where real time issues are driving the processes. Over that layer, a holistic management and control system has to be developed across the many nodes constituting the actors and the process components. As an overall, all these elements have to convey into a "chain of trust" that permeates the organized and healthy behaviours of an innovative supply chain.

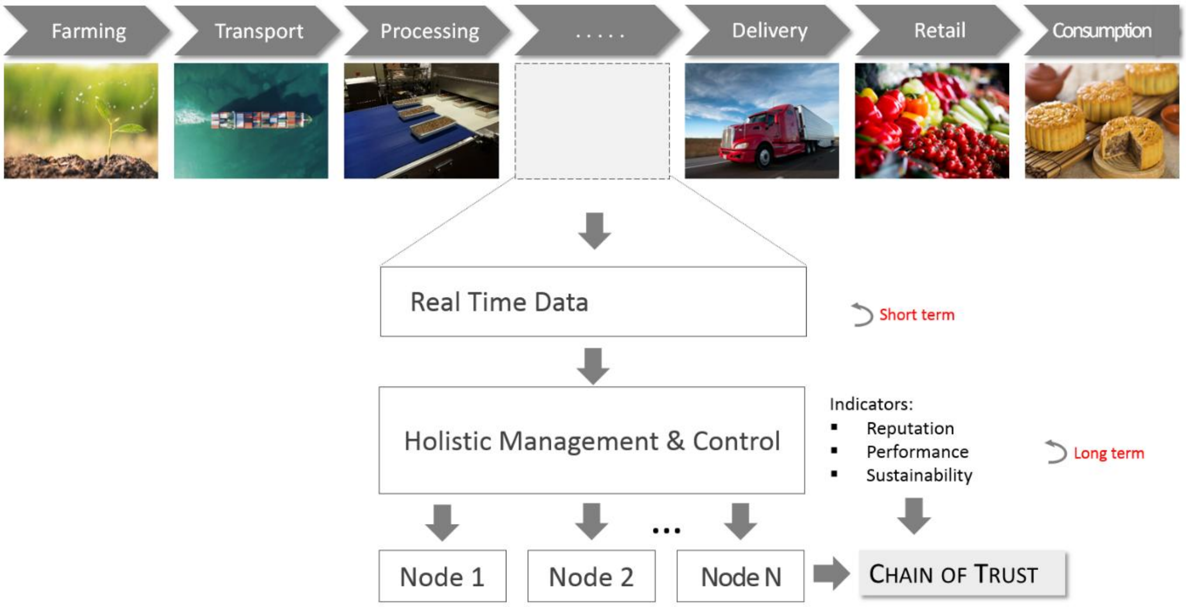

Figure 2.2.2: Vision of the supply chain as a "chain of trust," showing systemic layering of real-time and holistic management perspectives.

The management and control layer envisioned has to establish suitable indicators and has to conceive and use a clear methodology that links total quality to the chain of trust. This concept has been discussed among the partners with the slide shown in Figure 2.2.3. This figure iterates on the concept that for each of the systematic processes (self-similar in nature) that composes the supply chain a continuous improvement has to be pursued in order to reach the several technological, social, and organizational dimensions that sum up to a global quality indicator that can be conveyed into the chain of trust concept. This concept has been elicited in turn from the exchanges with project partners as a fundamental concept in their food field experience and pervading their technology solutions.

A novelty that this project has to introduce is also a mutual role of humans and machines, as peer active contributors to the processes in the new hybrid reality realm of symbiotic autonomous systems that shape the new kind of sustainable business in the digital era and beyond[3,4,5]. We referred to that vision as "the Kaizen and the Bazaar," a play on words meant to highlight its parallelism with the revolutionary economic and production models—such as the gift economy—spawned by the open-

[3] Raikov, A. N., & Pirani, M. (2022). Human-Machine Duality: What's Next in Cognitive Aspects of Artificial Intelligence?. IEEE Access, 10, 56296-56315.

[4] Raikov, A. N., & Pirani, M. (2022). Contradiction of modern and social-humanitarian artificial intelligence. Kybernetes, 51(13), 186-198.

[5] Pirani, M., Cucchiarelli, A., Naeem, T., & Spalazzi, L. (2025). A Blockchain-Driven Cyber-Systemic Approach to Hybrid Reality. *Systems*, *13*(4), 294.

source movement in recent decades, an evolution first analysed in depth in E. S. Raymond's best-selling book[6].

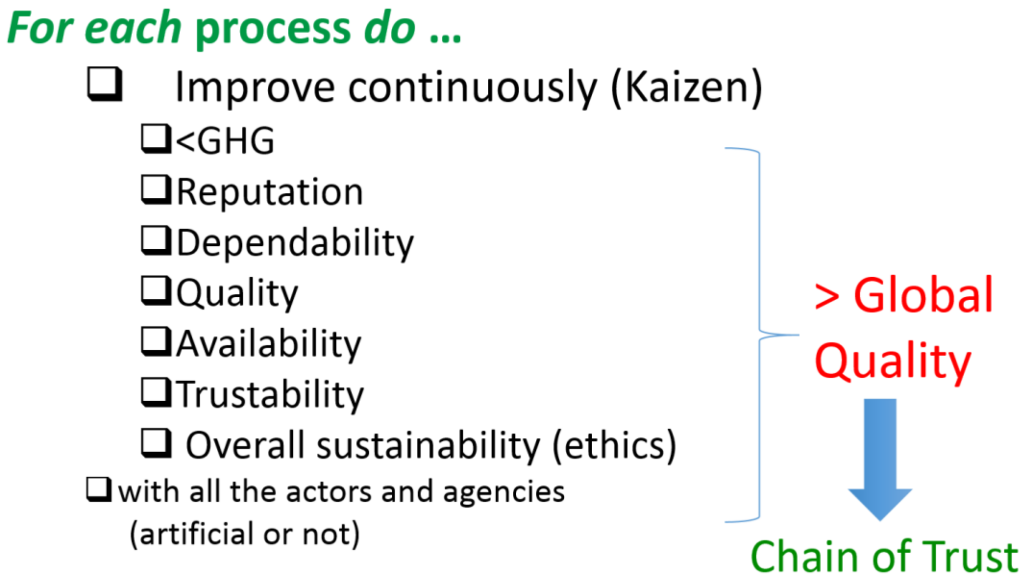

Figure 2.2.3: Concept of "The Kaizen and the Bazaar," illustrating mutual roles of humans and machines as peer contributors to processes.

From these underlying ideas, a pragmatic architecture for the SDS had to be generated at the end of the day. The first schema discussed with the partners was the SDS as a "wrapper". The SDS should be able to connect all the dimensions from data to knowledge and organize it whilst considering new technologies like Blockchain among the main fundamental blocks.

The SDS must act as a cognitive interoperability hub that allows overall control and organization of the processes in the food supply chain context. The idea was that the SDS should be an intelligence layer across all the elements, being those services, data sources, database management systems (DBMS), humans, processes, and management. It should sport also an inherent idea of flexibility and adaptation in order to represent an inclusion means for all the technological and socio-technical aspects of the food supply chain. Nonetheless, it remained a tall order to fulfil for the SDS system. As such, the active contribution of all the stakeholders was essential for its realization. It required that most of the knowledge and of the technologies already present in the food industrial playground to be harmonized and included from the bottom up. At the same time, it requires a top-down strength in order to let this process be purposeful and organized.

At the end of the requirements iterations, and as a ground of the technological specifications, the SDS concept took a clearer but maybe increasingly ambitious scheme and meaning.

What is a Smart Data System then? We can connote the three words Smart, Data, and System, with concepts that would express better all the intensions and the final purpose rather than going for a mere technical and analytical definition. The connotations achieved are shown pictorially in the Figure 2.2.4. We briefly discuss these points:

- **Smart**. The acceptation of the term can be many. This has become quite an abused adjective in recent years. Surely, the first connotation is intelligence, in many contexts, has the capability to It is the ability to meet the needs of all stakeholders. But at the same time, it is the ability to make it easy for a wide audience to participate and use the latest technologies. This entails

[6] Raymond, E. S. (1999). The cathedral and the bazaar. O'Reilly.

an essential capacity for simplification, particularly in a vast context with the participation of the most diverse stakeholders and in the face of ever-increasing complexity.

- **Data**. Today the capability to collect data is huge. The Big Data is both a reality and a frontier. In industry often, there are more data than the actual possibility of handling them purposely and effectively. Nevertheless, there are still big holes where they should be provided and completely lacking. Thus, the current problem at hand is a rational and sustainable way to analyse the data whether they are available, but also there is the need of analysis of their absence or low quality. Pervading the world with sensors is unfeasible. Brute-force digitalization is unsustainable economically, energetically, environmentally, ethically, and socially. This requires a more intelligent ways of sensing, dynamically defined depending on the context and on the actual effectiveness of it. Moreover, most of current efforts on Big Data are about collection and analysis. But the control loop to the controlled system to be improved is usually closed by feeble actions, not in real time, but mostly completely absent. Hence, more strength must be put on the control side that is a complementary aspect, which provides a self-regulating framework on quality and quantity of actually necessary data acquisitions.
- **System**. A system is to be seen as in the General Systems Theory acceptation. A system is every reality that can objectively be modelled and then controlled. In addition, Systems Thinking is brought about as soon as complexity sets in, and a holistic view of the problems is required. This means including aspects, even apparently distant between them, and related spread dimensions of information. In addition, the natural, social, and technological dimensions must be included at once in the system and controlled at the same time. An extended definition of systems determines more inclusiveness, openness and, at the end of the day, trust for all the stakeholders called to participate in and include in the system. While inclusion determines the scope of a system, trust is the organizing force that determines the fundamental structures for its viability, resilience, and survival.

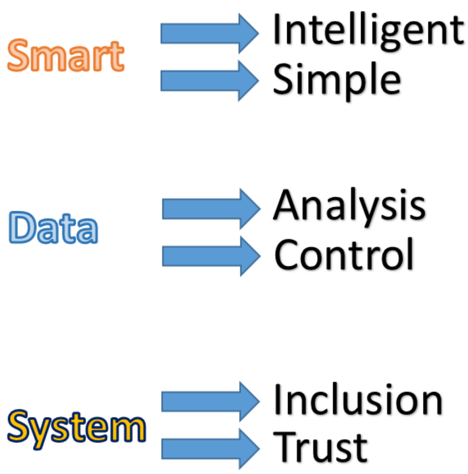

Figure 2.2.4: SDS definition and connotations derived from the requirements elicitation phase (Smart, Data, System).

The SDS concept has taken a shape as a digital and virtual backbone of services, a holistic ecosystem of run-time services, in order to interoperate ENOUGH stakeholders in a unified and co-creative way, with sustainability inside (with a look to the new vision of Enterprise 5.0).

In particular, it will be required that:

- Starting from TRL 4-6 background components provided by the project partnership will be the first step to arrive to express sustainable innovation in food supply chain.
- SDS will do that with something clearly distinguished with respect to other digital platforms with the following guiding keywords:
    - Impact
    - Uniqueness

As from the project baseline the lowering of GHG has to be the primary goal. This objective can be pursued with many routes. The choice of the "optimal" route has to consider sustainability as a side effect in the wider sense defined by the United Nations' SDGs, as expressed with the Figure 2.2.5.

Figure 2.2.5: Overall aims of the Smart Data System (SDS) in the ENOUGH project, aligned with UN Sustainable Development Goals.

The advances expected with respect to the state of the art are then multiple and multidimensional.

Beyond the advancements foreseen, which concern more the research part of the project, the question that is steadily standing beyond the definition of the SDS is: what are we (ENOUGH) going to sell with the SDS? The clear answer now is:

- Not just another digital platform gadget. – A huge set of digital platforms, more or less successful, have been already funded by EU. None of them is general purpose, and difficult to be adapted in complex contexts like the ENOUGH one.
- We do not reinvent the wheel. – But use as much as possible existing technologies and components. Innovation is using something that already exists in a new way. This means bringing about results and experience from past EU and company projects and drive them up to innovation.
- Express our uniqueness. – Uniqueness is the key to any business model worthy of the name. We focus on the new worldviews emerging from ENOUGH.

Nevertheless, this stance for the SDS has to be pursued with complying and interpreting some lessons in learnt in affine contexts like the "Smart Data Models" initiative (led by FIWARE[7]), an initiative started with EU funding. The actions in the FIWARE follow the Agile Standardization Manifesto[8], which reads as follows:

[7] FIWARE: Open Source Platform for the Smart Digital Future. https://www.fiware.org/

[8] https://github.com/smart-data-models/data-models/blob/master/MANIFESTO.md

1. Don't just standardize, be agile and standardize (agile standardization, velocity should be measured in terms of days or weeks)
2. Do not reinvent the wheel
3. Normalize real cases (agile standardization should only be based on real cases)
4. Be open
5. Don't be overly specific
6. Flat not Deep (be local and self-contained)
7. Sustainability is key

Under these principles, the SDS is going to exploit:

- A "DevOps" rather than a specific digital platform.
- A set of practices that combines software development, management, and operational technology.
- The shortening of the systems development life cycle and provide continuous delivery with high software quality.
- A multipart and multisided business. In which buyers, sellers and developers are interchanging and peer roles.

Definitely the idea is that ENOUGH is not selling only software. We are rather selling a platform, within a framework, along with its *DevOps*. *DevOps* here is prototypical because we are continuously adapting due to the novelty of the domain and the collaborative-research-oriented realm.

Our *DevOps* is based on Kanban that means a well-defined process:

1. collection of narrative requirements;
2. production of sequence diagrams (or in the future maybe starting even from use cases);
3. validation of sequence diagrams with users (and crosscheck with SDS platform capabilities);
4. pull request to developers (ENOUGH app shop?) for:
    - UIs (user interfaces)
    - configurators
    - planners
    - connection to platform (through APIs – application program interfaces);
    - other services (in Cloud);
    - etc.

Under the technological aspects, the smartness expected in the SDS has to consider the new frontiers of AI (artificial intelligence) inside. Today, AI is the buzzword of digitalization. Nonetheless, the definition of AI and its boundaries and definitions are continuously changing and evolving. The SDS cannot avoid or exempt itself from the inclusion of artificial intelligence. However, the complexity it seeks to address requires an approach that is both pragmatic and yet advanced. One concept to which SDS can adhere with all its dimensions is that of HABA-MABA (Humans Are Better At – Machines Are Better At). This concept comes from afar[9], yet it is still a frontier in the field of the interface between processes governed in symbiosis between humans and machines. SDS candidate itself as a matchmaker between the categories that are better tackled (and for a long time if not forever) by humans and the corresponding activities that are worth (and should) be done by means of automations.

[9] Fitts, P. M. (Ed.). (1951). Human engineering for an effective air-navigation and traffic-control system. National Research Council, Div. of.

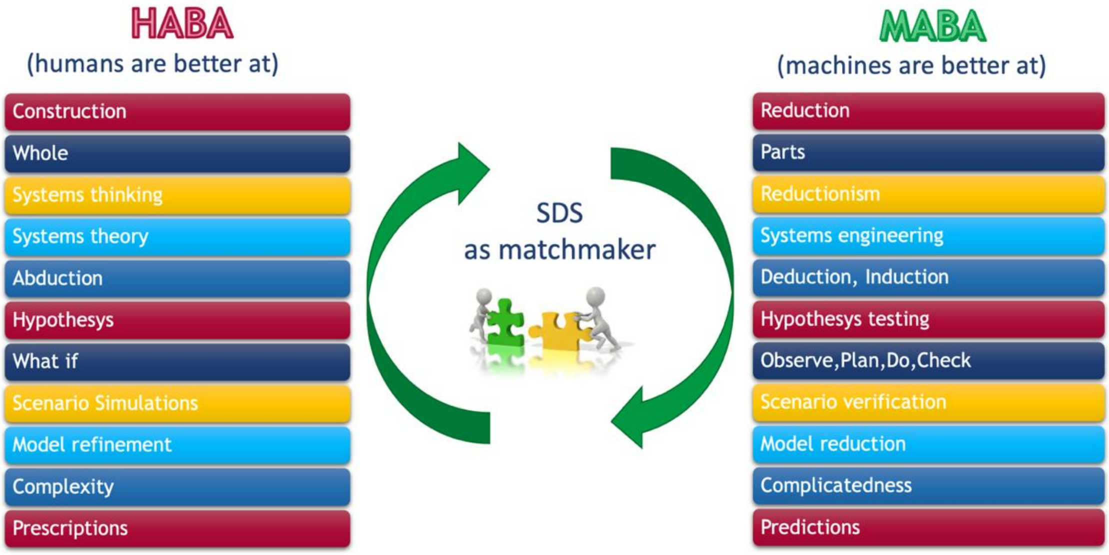

Figure 2.2.6: Ambitious role of the SDS as a cognitive interoperability hub between humans and machines.

One expression of this capability is potentially in the provisions available from some of the partners. The SDS framework is then able to enable a new kind of interaction between humans and the machines used as extension of their intelligence. New tools of extended and hybrid reality are the new frontier between the human and the artificial. This seems not new, but the novelty is in the way they are used in an innovative way. In the spirit of Figure 2.2.6. This means getting new insights in the frontier between the digital and the convergent and collaborated decision-making.

Eventually, the requirements elicitation process brought about also the importance of the contextualization of the SDS into the biggest trends in industry. In particular, SDS cannot avoid considering the major standardisation processes currently underway. Now, speaking of industry in Europe without mentioning the financial and research efforts around the Industry 4.0 and 5.0 would be unrealistic. Sure, SDS intervenes in food that is one of the most important and relevant sectors for this new vision of the smart industry, in particular for the EU. In the Figure 2.2.7, the major framework that the SDS has to consider as a necessity are listed. Among them, not mentioned so far, we find the RAMI 4.0 standardization process and the Asset Administration Shell, which is a technology that is gaining a big momentum for both researchers and practitioners in the field.

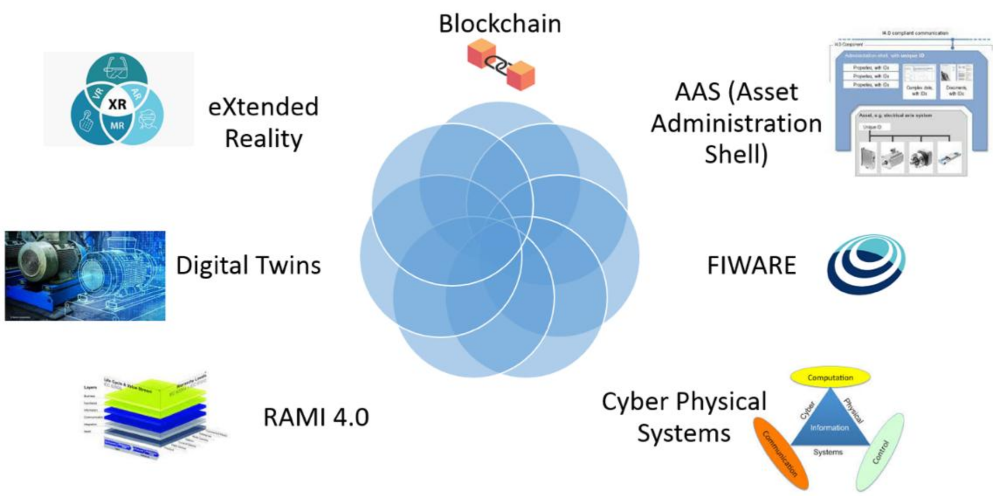

Figure 2.2.7: SDS specifications aligned with Industry 4.0/5.0 frameworks, including RAMI 4.0 and Asset Administration Shell

2.3 The SDS as an evolving digital twin of the SCE

One of the major concepts achieved in the last version of the SDS Infrastructure prototype is the fact that by means of the SDS system, the host of modules and informational nodes present in the system, render it a de-facto digital twin of a supply chain process. There are many definitions and nuances of digital twin definition in literature. We do not bring academic controversy in here. The main concept here put forth is that the supply chain is a set of information and physical processes, along with human decisions and actions. Part of the process is prone to digitalization and virtualization, while most of it has to remain (even for the mere problem of computing sustainability) in the physical realm and not digitizable.

Therefore, the SCE, which is for us the informational instance of a supply chain, has to include in itself a physical inseparable counterpart that is continuously intertwined and in connection with the digital information part. Moreover, this connection acts bidirectionally. The whole SCE works well if the flux of information from the physical realm to the digital is at real-time. Real-time does not necessarily mean "fast" or "immediate", it means timely before a control act or decision is made based on this information. The same happens in the other direction. The decisions or actions commanded by the digital environment due to simulations or predictions, or simply communications and automations between the digital counterpart of the Actors[10], must happen before a significative change of conditions in the environment happens, which respect to the models.

The SCE system has the infrastructure designed having these requirements in mind. Mostly it is event-based both form the physical to digital, but also from the digital to physical. Event-based computation paradigms are more suitable to fit real-time constraints than other type of computation. At least, they do it in a more sustainable and natural way[11].

In the SDS approach, the model of the SCE does not simply reside somewhere as a virtual central oracle (although something like that can be a Node in the distributed computation, provided by someone in the SDS marketplace as seen above). The model is distributed between the Participants. Parts of this model is private to the Participants, and part of it is public. The relevant thing is that the action of any of the Participant follow a double global and local view: they see the whole and they see their own process, in view of the whole. Thus, the model, used for predictions and control actions is shared and hybrid, in the sense that some of the Participant (usually most) do not have digital twin themselves of their processes, but a model anyway.

Being the model linked to the belief, knowledge and behaviour of the Participants; it is prone to evolve continuously as the search for a continuous improvement is a built-in feature of the SDS application.

---

[10] We use capital letter for Actor as we mean a specific definition of it in the context of SDS. The SDS Actor is an entity (human or automaton, a person or a hardware/software application) that participates in the SCE process and information processing.

[11] Pirani, M., Dragoni, A. F., & Longhi, S. (2021, October). Towards sustainable models of computation for artificial intelligence in cyber-physical systems. In *IECON 2021–47th annual conference of the IEEE industrial electronics society* (pp. 1-8). IEEE.

The digital part of this evolution is recorded in the data that trace all the exchanges between the actors. This constitutes the history and the dynamics of the SCE, which remains stored and embedded in the SDS DB.

2.4 CI/CD for the SCE

CI/CD, which stands for Continuous Integration and Continuous Delivery (or Continuous Deployment), is a modern software development practice that automates the process of building, testing, and releasing code. With Continuous Integration, developers frequently merge code changes into a shared repository, where each change triggers an automated build and tests to ensure compatibility and catch issues early. Continuous Delivery takes this further by automatically packaging and preparing the software for release, keeping it always in a ready-to-deploy state. Continuous Deployment goes one step beyond by automatically pushing every validated change directly into production. Together, CI/CD speeds up development, improves quality through automated testing, reduces risks with small frequent updates, and supports agile and DevOps collaboration.

The production of an SCE from the SDS Platform must follow this approach to be effective.

The production of an SCE is the core application and business for the SDS Platform. At least it is the core application that fulfils the target and the objectives of the ENOUGH project with the management of the entire lifecycle of an informational object that represents the digital counterpart of a food supply chain in all its processes and complexities. As remarked in section 2.3, the SDS system renders an SCE as an evolving digital twin of the supply chain entity. This is a new concept in the field, which we consider a breakthrough. However, it is a complex digital object itself, and many informational and software dimensions are combined and comprehended in the SCE.

The creation, deployment and run of an SCE is a collection of a host of multiple activities, that require both human actions and automations to be conducted. The workflow of this process must be systematic to work and to be improved and maintained in the future.

The inner detail of every step of the SCE LIFE CYCLE is very technical and beyond the scope of this very report. Nevertheless, we provide here the activity diagram (in UML 2.5) that was used by the developers as a specification for the realization and verification & validation of all the steps of the life cycle of the SCE under CI/CD concept.

In Figure 2.4.1, the activity diagram expresses the life cycle of the SCE in its essential steps. This diagram shows mostly the major systematized procedures that now have been perfectioned and validated for this version of the prototype. They are highlighted in coloured shapes with bold text, namely:

1. **SCECreate.** The SCE is created by the SDS Developers considering the information that come from a group of stakeholders that would be interested to commit themselves in such an initiative. This is done by means of the SCE Configuration Tool. The process of creation is incremental. The information and specifications from the stakeholders can come in different separate chunks and may require several iterations between the SDS Developers and the Stakeholder.
2. **ActorIntegration.** This is a process that produces instructions and specifications of requirements to a Stakeholder when it is their turn to provide information about the variables that they will produce and the variables that they will need to receive to implement their Actor behaviour. A Stakeholder (by our own definition in this context, see section 3.1) may own one

or more Actors. In case one Actor commits to appear in the ***SCE Choreography view***. This view is achieved by a BPMN Choreography diagram[12],[13]. The SCE Choreography is the "public" agreement or contract among the SCE participants. Other kind of contracts maybe be "private" and not visible to other supply chain participants — for example the usual interconnection between two organizations at the OT level (remote control and automation of processes). When an Actor has a public role, it is to be considered a Participant in the SCE Choreography. The SCE configurator, in the current version, can only handle Actors that appear in the Choreography. This was a choice and a priority as the Choreography is the actual added value in which the inherent collaboration and continuous improvement is elicited. For the other case, of "private" connections, there are many commercial tools available for interoperability from many vendors — creating a new one does not add much value, but integrating some of them on demand is part of the business. In the Actor type connection, only a typical DCS (distributed control system) is achieved if the Choreography is not active. However, in the practice both the modalities should be active at the same time. In most cases, the more general actor case would require a finer, manual procedure.

3. **SDSConnect.** After all the information for the integration of the Actors from a Stakeholder is received, this phase is to produce four major artefacts about the Actor:
    3.1. **Secure Actor ID**, is the set of authorization and authentication data, which are used when the Actor communicates in the SDS Infrastructure. This data is maintained and managed by the Security layer of the SDS infrastructure .
    3.2. **SDS Connector**. It consists in the suite of software that the Stakeholder has to install in order to be connected to the SDS infrastructure. Part of this process might be in charge to the SDS Developers or directly by the Stakeholder's resources. It depends on case by case from the nature of the skills available by the Stakeholder staff or simply by their make-or-buy policy. Usually, it is done by means of two major technologies: RMAS (Relational-model Multiagent System)[14,15,16]or GrapQL — provided in four forms by ELET (Eletica S.r.l) ranging from bare-metal embedded software up to desktop or cloud. The open technology of the SDS Connector enables any participant, regardless of technological level, to connect to the SDS infrastructure. Neither specific operating systems nor particular computing power requirements are needed. This approach effectively lowers or even removes barriers to participation in the SDS, especially for SMEs.
    3.3. **Actor Web Interface**. This is optional, as not strictly required for an Actor to join the SCE. However, this provision could be useful to an Actor (human in this case) to have a native dashboard on the SDS to control their own behaviour and data.

---

[12] Allweyer, T. (2016). *BPMN 2.0: introduction to the standard for business process modeling*. BoD–Books on Demand.

[13] OMG (2014). BPMN™ — Business Process Model And Notation. Version 2.0.2. https://www.omg.org/spec/BPMN/2.0/

[14] Pirani, M., Cucchiarelli, A., & Spalazzi, L. (2024, November). A Role of RMAS, Blockchain, and Zero-Knowledge Proof in Sustainable Supply Chains. In *IECON 2024-50th Annual Conference of the IEEE Industrial Electronics Society* (pp. 1-4). IEEE.

[15] Pirani, M., Cucchiarelli, A., Naeem, T., & Spalazzi, L. (2025, May). Verifiable Actor Model Systems Through Relational-Model Multi-Agent System and Zero-Knowledge Proofs. In *2025 IEEE 8th International Conference on Industrial Cyber-Physical Systems (ICPS)* (pp. 01-06). IEEE.

[16] Pirani, M., Bonci, A., & Longhi, S. (2022). Towards a formal model of computation for RMAS. *Procedia Computer Science*, *200*, 865-877.

3.4. **Actor Wallet**. Optional. . This application allows an Actor to express a reputation judgment on the conduct and behavior of other Actors within the SCE in which they participate. This is an incentive and nudging mechanism overlaid on the SCE within the Blockchain infrastructure, activated in parallel with the SDS Infrastructure to fully express the potential of the chain of trust.

4. **SCEInfrasturctueTest.** After all the provisions for the Actors are set, there is an internal test in the infrastructure to verify that all the communications across the Publish/Subscribe scheme of the SDS Infrastructure is correct. It is done by an automation that produces and routes fake and test data for the running of the SCE. For this test the Actors are not involved.
5. **SCESecureConnectionTest.** For this test the Actors are involved. It is request to any of them to test the secure connection and transmission of some data and to access to their provisions on the SDS system. This is the verification of the correct configuration of the Secure Actor ID layer.
6. **SCEDryRun.** A further final integration test will involve all the Actors simultaneously. At the end of this test, the SCE is made ready to start and the Actors are requested to be ready.
7. **SCERun.** Is the, usually long, phase in which the SCE is used.
8. **SCEStop.** On arrival of a configuration request, a setback decision, or the integration of another Stakeholder, the SCE has to be stopped. At the moment, it is considered to complex and brittle to handle changes of configuration of the SCE on the fly. Nonetheless, it could be a decision for a new feature in the future. This of course requires due communication to all the Stakeholders and a new commitment for anyone.
9. **SCEArchive.** This very important step concludes the life cycle of the SCE. Being all the data circulated between the Actors recorded, the trace of the SCE run is archived for subsequent process mining of lesson learnt. Still a business case on this phase has to be developed. It is not clear how much and where are the costs of the creation of the repository of the SCEs.

Taken as a whole, Figure 2.4.1 illustrates, through a UML 2.5 activity diagram, the systematic articulation of the SCE life cycle within the SDS framework, conceived and implemented according to the principles of Continuous Integration and Continuous Delivery (CI/CD).

The diagram highlights how the creation, deployment, operation, and termination of an SCE are organized as a coherent workflow, combining automated procedures with stakeholder involvement in a controlled and verifiable sequence of steps. The cycle begins with the definition and incremental construction of the SCE, informed by the requirements and commitments of a consortium of stakeholders. This preliminary stage is followed by the integration of actors, where specifications of data exchange and behavioural variables are formalized to ensure interoperability within the supply chain choreography.

Once actors are integrated, the SDS Connector provides the technical means for secure and transparent interaction, establishing a foundation that does not impose restrictive technological prerequisites and thus reduces barriers to entry, especially for SMEs.

Subsequently, the infrastructure undergoes an internal validation phase, ensuring that the communication backbone of the SDS, based on a publish/subscribe scheme, operates correctly. This is followed by secure connection testing, which actively involves the actors in verifying the robustness of their authentication credentials and data exchange channels. A dry run phase then simulates the coordinated participation of all actors, providing a final systemic check before the SCE is released into its operational phase.

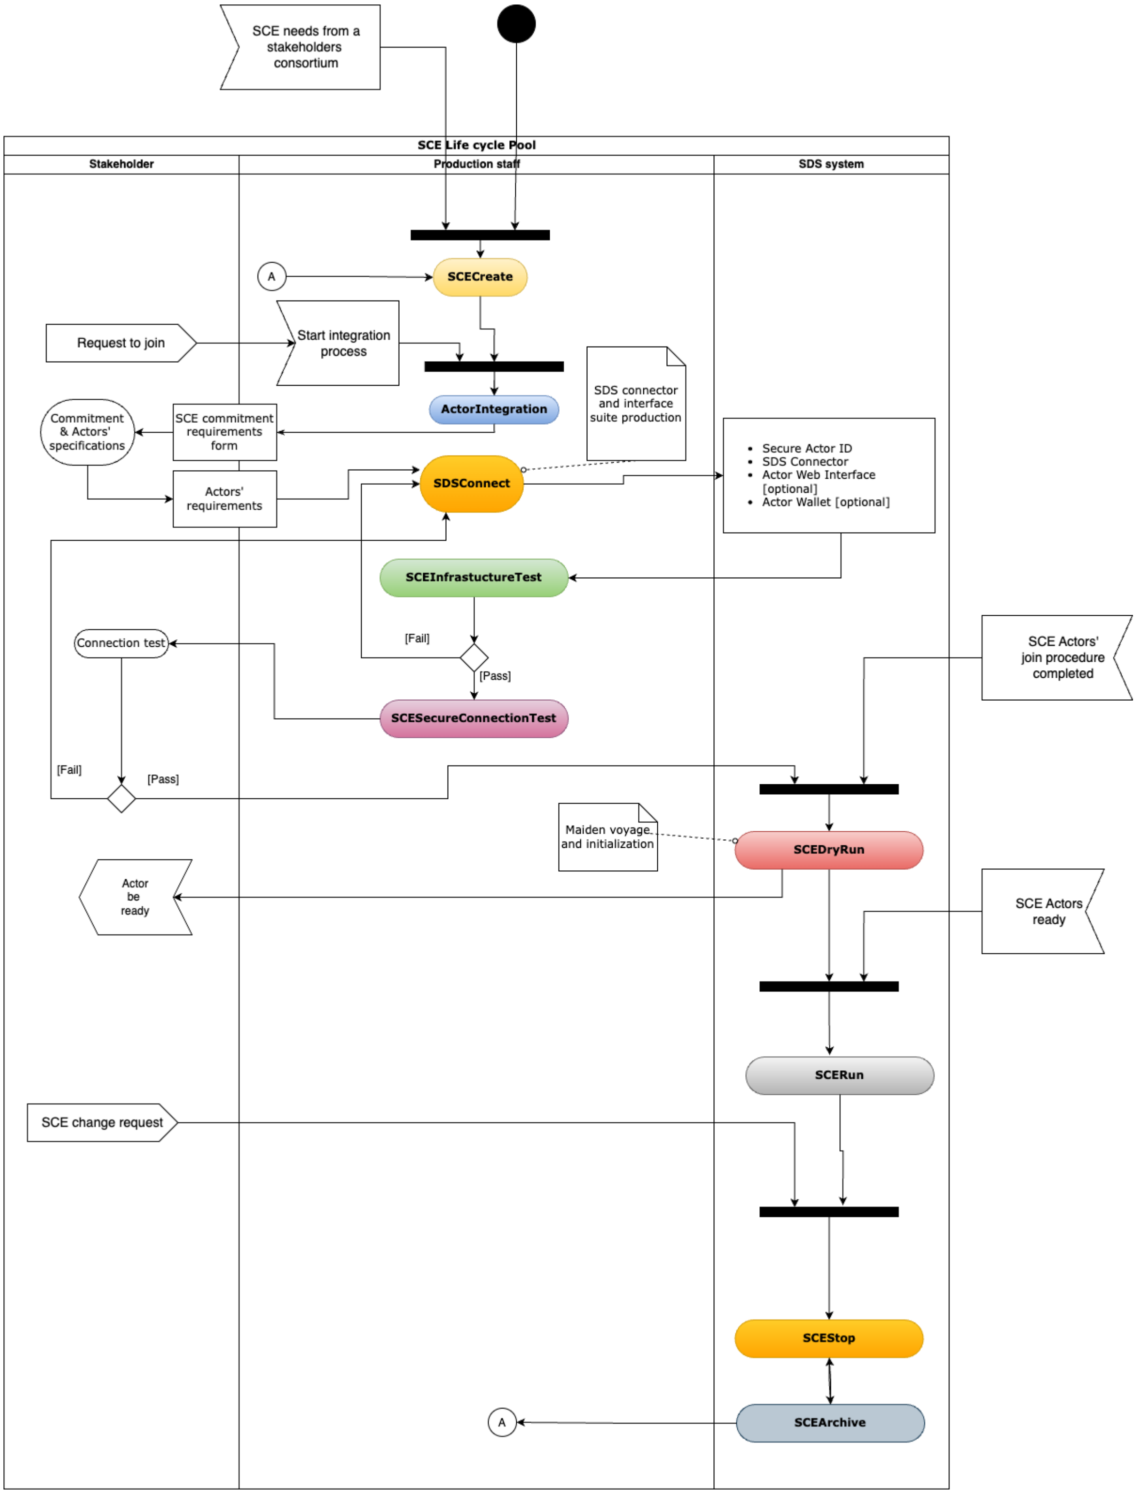

Figure 2.4.1: UML 2.5 activity diagram illustrating the systematic life cycle of an SCE under CI/CD principles.

The SCERun phase is typically of long duration and constitutes the core functionality of the SCE, enabling it to act as a dynamic and evolving digital twin of the supply chain entity. During this period, continuous interactions among actors and the SDS infrastructure not only facilitate operational goals but also create a data-rich environment for monitoring, learning, and process improvement. At the conclusion of its operational cycle, the SCE may be stopped due to reconfiguration requests, stakeholder changes, or strategic decisions. This interruption is followed by archiving, a critical step that ensures the preservation of all generated data and the possibility of conducting ex post analyses, process mining, and the extraction of lessons learned.

The structured progression of these phases reflects the essence of CI/CD in the context of digital supply chains: each step integrates verification and validation, ensuring that errors are detected early, interoperability is guaranteed, and stakeholders maintain trust in the system.

Moreover, the archival of each SCE run provides a longitudinal perspective, allowing cumulative knowledge to inform future implementations.

In this way, the SDS life cycle is not only a technical procedure for managing digital entities but also a methodological scaffold that embodies systemic learning, continuous improvement, and socio-technical integration in line with the objectives of sustainable and resilient supply chain management.

3 THE DRIVING BUSINESS CASES OF THE SDS

The realization of the Smart Data System (SDS) concept can be seen as a fundamental step in the evolution of digital infrastructures and digital transformation supporting sustainability, supply chain management, and business innovation. The SDS vision does not stop at being a simple technological layer; instead, it aims to integrate business models, governance strategies, and systemic paradigms into a single coherent framework. The concept embraces the need to combine technological advancement with holistic approaches that recognize the interdependencies between economic, environmental, and social dimensions. This realization has taken the structure allowed by at least **4 business use cases.**

At its core, the SDS can be boiled down to a business framework, articulated through four major use cases that represent concrete entry points for industrial and societal adoption. These use cases are not abstract exercises but are grounded in real business contexts where the availability, management, and valorization of data play a crucial role. They embody the promise of shifting from fragmented and siloed practices to an integrated model where digital ecosystems become enablers of transparency, traceability, and continuous improvement.

The first pillar of the SDS framework is a platform dedicated to the life cycle management of Supply Chain Entities (SCE). These entities, when represented as **digital twins of supply chains**, provide an unprecedented ability to monitor, simulate, and optimize the flow of resources, products, and information across multiple stages of production and distribution. By embedding digital twins into the SDS, supply chains can evolve into living systems, capable of self-assessment and adaptation to changing market conditions, regulatory requirements, and sustainability targets. The life cycle view ensures that the management of resources is not restricted to short-term efficiency but extends to long-term resilience and sustainability.

The second component of the SDS framework is a **digital networking infrastructure** designed to lower barriers to entry. Many small and medium-sized enterprises (SMEs) currently face challenges in adopting advanced digital solutions due to costs, complexity, and lack of expertise. The SDS proposes an inclusive infrastructure that minimizes these barriers, enabling even smaller actors to participate fully in digitalized ecosystems. Such inclusivity not only democratizes access to technology but also strengthens the overall supply chain by integrating diverse participants who bring agility and innovation capacity.

The third pillar is the creation of an **open marketplace** for the digital transformation of sustainable supply chains. This marketplace is envisioned as a dynamic environment where stakeholders can exchange services, data, algorithms, and solutions. By functioning as a trusted intermediary, the SDS marketplace enables collaboration across organizational boundaries, reduces duplication of effort, and accelerates innovation. Moreover, it supports the development of new business models where data becomes an asset that can be shared, traded, and co-created responsibly. In this way, the marketplace embodies the principles of openness, transparency, and value-sharing.

The fourth and final dimension of the SDS framework concerns **holistic and systemic paradigms** for the continuous improvement of food supply chains. Unlike conventional approaches that focus on isolated optimizations, this paradigm views the supply chain as a complex adaptive system, where interactions among components generate emergent behaviors. Holistic paradigms ensure that improvements in efficiency do not come at the expense of sustainability or social well-being. Instead, the emphasis is on aligning performance metrics with environmental goals (such as reducing greenhouse gas emissions), social equity (ensuring fair labor and consumer safety), and economic viability. This continuous improvement cycle reflects the cyber-systemic tradition, where feedback loops and adaptive strategies are essential for long-term success.

Taken together, these four pillars define the essence of the SDS as both a technological infrastructure and a business enabler. The SDS is not merely a tool for digitalization; it is a framework for systemic transformation. By linking supply chain entities through digital twins, by lowering the threshold for digital participation, by fostering open marketplaces, and by embedding holistic paradigms of improvement, the SDS sets the stage for a new generation of supply chains that are intelligent, adaptive, and sustainable.

The Smart Data System can be seen as a convergence point between technology, business, and systemic thinking. It offers a roadmap for organizations to transition toward Industry 5.0 principles, where human-centric, sustainable, and resilient practices are integrated with cutting-edge digital tools. The realization of this concept will require coordinated efforts across industries, academia, and governance bodies, but the potential rewards in terms of efficiency, sustainability, and inclusivity make the SDS a cornerstone for the future of digital and sustainable supply chain management.

3.1 Underlying business concepts

The capabilities and features of the SDS design were conceived with a (though preliminary) business model in mind. This approach originated from the design thinking process, as recalled in the previous section.

On hindsight we can tell that the methodological approach was effective. With such an approach we could definitely harmonize the requirements from the field experts and practitioners with the actual knowledge and skills of systems designers.

The envisioned business for the SDS framework (not only the platform) can be summarized with a few items as follows.

At the foundations of the (expected) SDS business model there will be a DevOps and multi-sided multi-part engineering, and a set of practices that combine software development, management, and operational technology.

Multi-sided, multi-part business engineering refers to the structured process of designing, integrating, and optimizing complex business models that involve multiple stakeholders, interconnected components, and dynamic interactions across various sides of a market or ecosystem. This perspective is particularly relevant in modern, interconnected economies, where many organizations no longer operate as isolated entities, but rather as platforms, ecosystems, or networks that enable and orchestrate value exchange among diverse participants.

A **multi-sided** business involves multiple distinct groups of stakeholders—such as customers, suppliers, partners, and service providers—whose interactions generate value for all parties involved. This is typical of platform models, where the core business activity is to facilitate exchanges between groups. For example, Uber connects drivers and riders, while Amazon links sellers with buyers, ensuring that each group benefits from the presence and activity of the other.

The **multi-part** dimension refers to the internal structure of the business model itself, which is made up of several interdependent components or subsystems. These may include distinct value propositions, operational workflows, revenue streams, or customer segments. Each part contributes to the overall functionality and competitive advantage of the business, and changes in one part often have ripple effects throughout the rest of the system.

Business engineering in this context is the discipline of systematically designing, analyzing, and refining such models. It blends principles from systems engineering, business strategy, and operations research to ensure that the business model is both robust and adaptable.

Real-world applications include **platform economies** (e.g., Airbnb mediating between hosts and guests), **ecosystem businesses** (e.g., Apple's App Store enabling developers and users to co-create value), **supply chain integration** (coordinating suppliers, manufacturers, and distributors into a cohesive network), and **digital transformation** (unifying multiple digital tools, data flows, and customer interfaces into one integrated system).

The core components of this approach can be summarized as:

- **Value creation**, defining how each stakeholder group benefits.
- **Value exchange**, mapping the flows of products, services, payments, and data across the system.
- **Interdependencies**, understanding how the performance of one group or component affects the others.
- **Revenue and cost structures**, ensuring sustainable profitability through mechanisms such as subscriptions, transaction fees, or advertising.

When effectively implemented, multi-sided, multi-part business engineering offers benefits such as greater innovation through co-creation, scalability through network effects, and resilience against market fluctuations.

However, it also presents challenges, including the need to manage complexity across diverse stakeholder groups, balance competing interests fairly, and navigate regulatory or ethical constraints. The success of this approach lies in maintaining both a systemic vision and the flexibility to adapt to evolving market conditions while ensuring that the value generated is distributed in a way that sustains engagement and trust across the ecosystem.

This approach is increasingly relevant as businesses move towards more interconnected, collaborative, and platform-oriented models.
The SDS framework has been designed in order to elicit and implement this kind of business.
In particular, the elements of this SDS business framework can be found summarized in Figure 3.1.1.
There are 4 essential blocks:

1. **Stakeholders**. They are in general individuals, groups, or entities that have an interest or stake in the outcomes, decisions, or operations of a project, organization, or system. They can affect or be affected by the entity's objectives and activities. In the specific case of the SDS, they are all the external parties (with respect to the SDS group) that are interested in selling, buying, and developing products and services, provide and comply to regulations, create a community with common financial, operational, or social interests.
2. **SDS Marketplace**. Is a market and a collection of tools, products and services available to the stakeholders.
3. **SDS Business Unit**. This is a unit or group of core owners of the SDS technology and business. This unit should be one of the outcomes of the ENOUGH project towards permanent exploitation of project results.
4. **SDS Shared Space**. Is a central application that constitutes the engine and the hub of all the integrated components. It sports an infrastructure for the connection of all the nodes in the SDS network of actors, a Choreography level that constitutes the view of a supply chain entity (SCE), and essential dashboards for the management and configuration of the system.

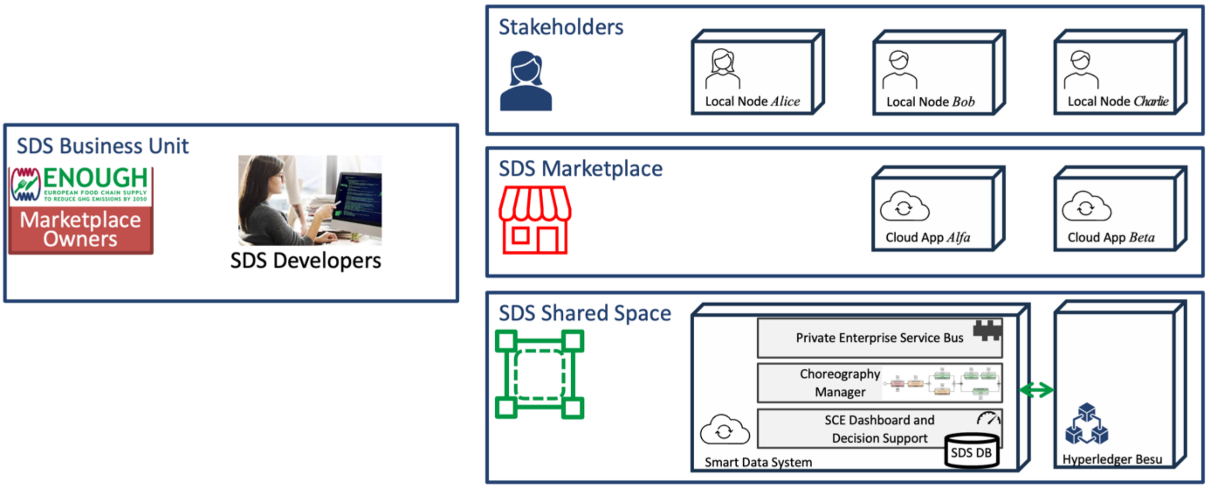

Figure 3.1.1: Building blocks of the SDS business framework: Stakeholders, SDS Marketplace, SDS Business Unit, SDS Shared Space.

Given these business components, four typical business processes are allowed in this framework.

3.2 A summary of the four main SDS business cases

Given the former SDS business playground, four major business use cases have been identified. These are the driving ones for the moment, and they represent a fil rouge across all the technological developments achieved around the SDS concept in this project. However, this does not impede some other cases to spring up in the future, due to the openness in the reference architecture of the SDS.

The first business case is shown in Figure 3.2.1. Clients use some or all of the provisions of the SDS Shared Space. These include the Private Enterprise Service Bus, the Choreography Manager and the SCE Dashboard, and Decision Support[17]. These applications can be at the IT (Information Technology) level or at the OT (Operational Technology) level. This means that there is a continuum between the

[17] Pirani, M., Cucchiarelli, A., & Spalazzi, L. (2024, November). A Role of RMAS, Blockchain, and Zero-Knowledge Proof in Sustainable Supply Chains. In *IECON 2024-50th Annual Conference of the IEEE Industrial Electronics Society* (pp. 1-4). IEEE.

operations like the shop floor and the planning and information processing at the highest level, as the last reference architectures like RAMI 4.0 are establishing in Industry 4.0/5.0 since decades[18,19].

This first business use case is the basic case of holistic commitment of Participants in an SCE, for their self-improvement together with the improvement of the whole SCE.

Typically, this will be the configuration that is showcased in the Demo 1 of this ENOUGH project[20]. An SCE is created through a commitment act by a group of clients (and stakeholders). Typically, this group will be constituted by some (or all) the actors in certain supply chain process. The SCE is run and controlled by the SDS systems, with all it provisions and features.

The Participant (or Actor) clients are provided with suitable SDS Connectors for the automation of data and messages transmission across the Actors network. Moreover, they are provided with access to a SCE Dashboard for the monitoring of what is going on, also in terms of performance, in the SCE as a whole.

Beyond monitoring, the Actors are enabled some control in the SCE as the choice of exchanged data and messages should move the SCE process towards improvement.

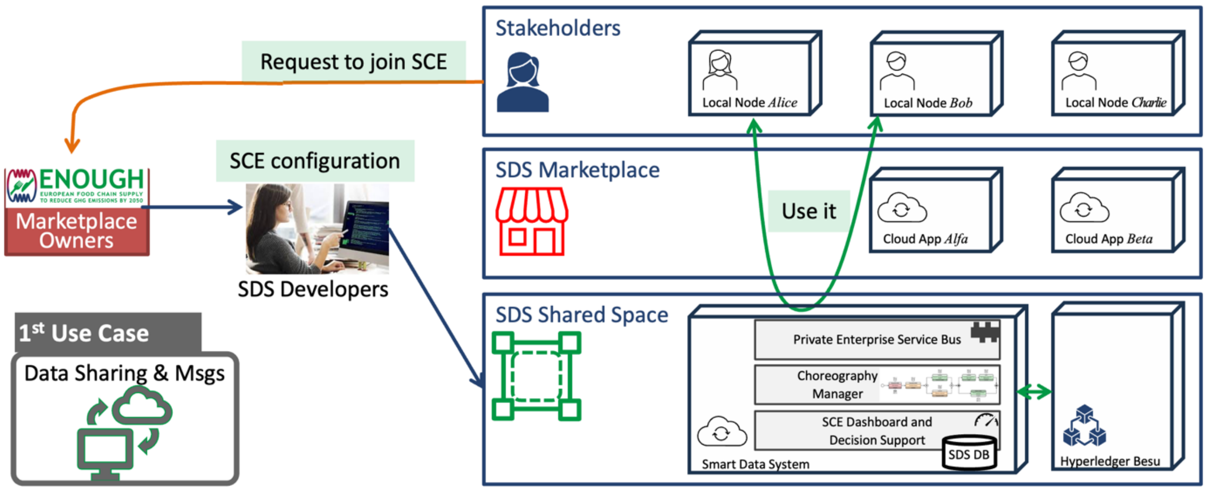

Figure 3.2.1: Business case 1 – Clients use SDS Shared Space provisions at IT (information technology) and OT (operational technology) levels for holistic SCE management.

18 https://www.zvei.org/fileadmin/user_upload/Presse_und_Medien/Publikationen/2016/januar/GMA_Status_Report__Reference_Archtitecture_Model_Industrie_4.0__RAMI_4.0_/GMA-Status-Report-RAMI-40-July-2015.pdf

19 Leitao, P., & Barbosa, J. (2025, May). Alignment of ADACOR Holonic Architecture with RAMI4. 0 and Industry 5.0 Principles. In *2025 IEEE 8th International Conference on Industrial Cyber-Physical Systems (ICPS)* (pp. 1-6). IEEE.

20 https://enough-emissions.eu/demonstrator/demo-1-holistic-supply-chain-management-and-control/

The second business case (or business use case), as shown in Figure 3.2.2, is similar to most of the components.

The added feature here is the possibility to have access for clients to additional applications or third-party services available in the SDS Marketplace. The SDS Marketplace can be defined as the set of applications and services provided along with the SDS system. In this specific use case the application is provided by Marketplace Owners but, in general they can come from OEM (Original Equipment Manufacturer) or third parties under due diligence agreements. Nonetheless, these applications have to undergo a step for their integration with the SDS system. The integration is analogous to the connection that clients experience in the former case. No major differences. The application or service supplier has to install the SDS connector and provide the list of exchanged variables with the SDS system. They might want to provide a special Web interface to their application in the SDS system as well, which is possible.

Examples of these applications can be: chatbot systems for decisions; AI systems for data analysis; decision support systems; databases and knowledge bases; special dashboards for data visualization and analysis; and others.

An example can be the tool developed in work package 4 (WP4) for supply chain configurations that can be easily (though with some additional effort) participate in this marketplace.

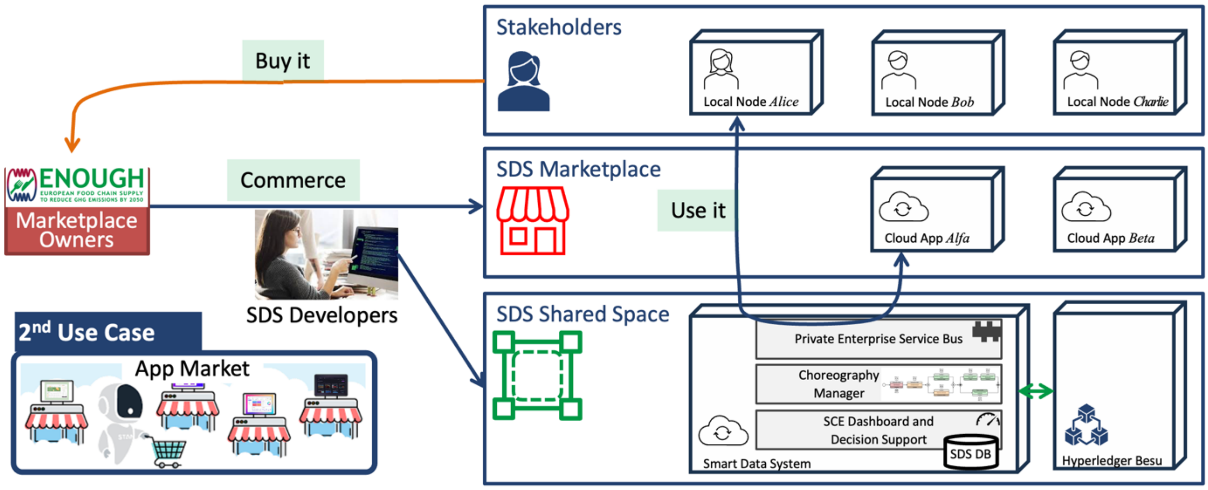

Figure 3.2.2: Business case 2 – Clients access additional third-party applications and services via the SDS Marketplace.

In Figure 3.2.3, it is shown how some third-party stakeholders have the possibility to sell their own solutions and services in the SDS Marketplace. They would request to integrate their software or service to the offer that the ENOUGH Marketplace Owners could already provide by means of the previous business cases.

After due agreements, the SDS Developers are involved in the collection of the requirements for the integration. This is done by providing the client with a suite of software that allows to connect their data with the SDS Infrastructure (through the SDS Connector) and to access their application (a specific SCE) directly from a Web interface that authenticates and authorizes clients in the SDS Shared Space.

Finally, as in Figure 3.2.4, the Owners of the SDS (from the consortium) might want to push in the marketplace new solutions and modules to let them available and sold to clients. In this case the requirements and the integrations come from actors that are knowledgeable of the integration process. In this case the process of integration is similar, but the requirements collection is negligible as well-known by the internal producers. In this business case it is also comprised the continuous maintenance and upgrade of the SDS system, towards unforeseen new functional and nonfunctional requirements elicited by the ENOUGH business unit.

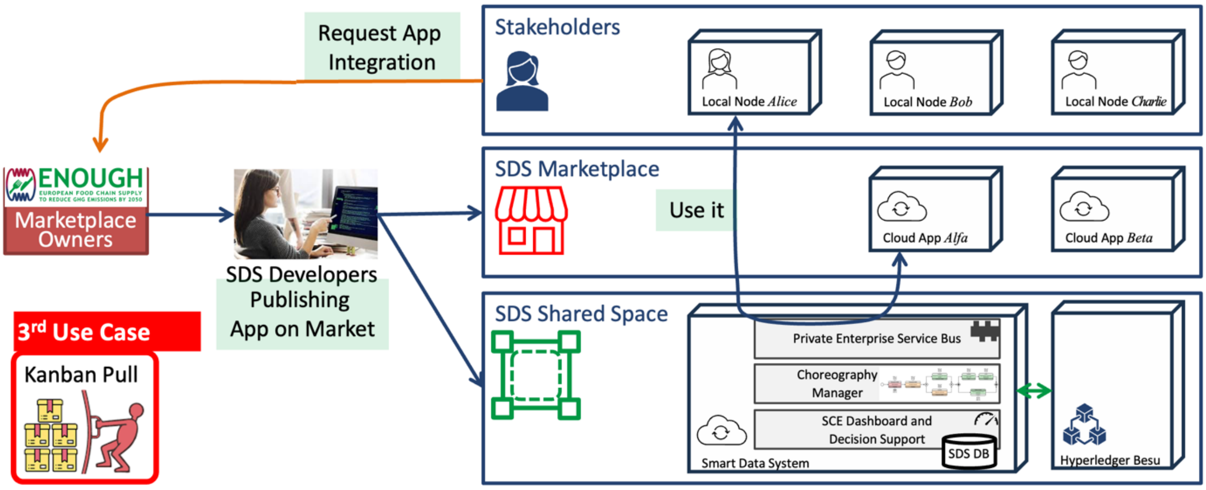

Figure 3.2.3: Business case 3 – Stakeholders integrate and sell their solutions in the SDS Marketplace through SDS connectors.

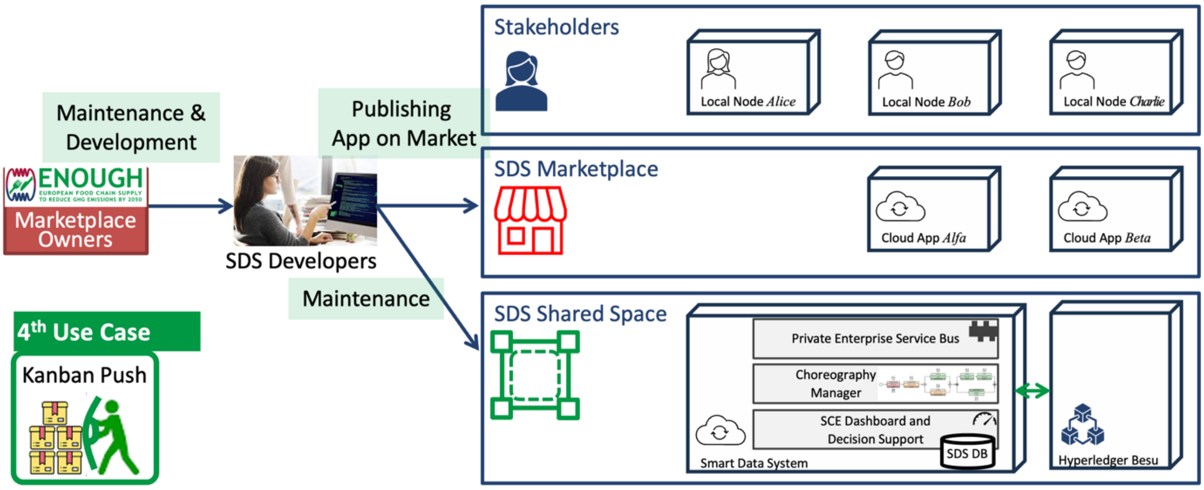

Figure 3.2.4: Business case 4 – SDS Owners introduce new modules into the marketplace, ensuring continuous upgrades.

Overall, the four business cases provide a coherent view of how the SDS framework can operate as both an enabling infrastructure and a catalyst for innovation. The first case demonstrates the foundational scenario, where clients engage directly with the SDS Shared Space to manage and improve their supply chain entities through holistic participation. The second extends this capacity by opening access to the SDS Marketplace, thereby allowing clients to complement core functionalities with external applications and decision-support tools. The third case highlights the reciprocal opportunity for third-party stakeholders to integrate and commercialize their own solutions within the SDS ecosystem, expanding the diversity and value of available services. Finally, the fourth case illustrates how SDS Owners themselves can introduce new modules, ensuring continuity, scalability, and responsiveness to evolving business and technical requirements.

Overall, these cases reveal a layered and evolutionary strategy: starting from basic integration, moving through marketplace participation, and culminating in continuous system growth and co-creation. This progression emphasizes not only the technological openness of the SDS architecture but also its potential to foster collaborative dynamics among different actors, from end-users to developers and system owners. In this sense, the SDS framework is not a static platform but a living business ecosystem, designed to sustain innovation, support interoperability, and address the complexity of modern supply chains in a structured yet flexible manner.

In the following sections, we will illustrate, through one example for each business case, how these concepts have been concretely materialized during the course of the project activities. Inevitably, certain limitations have emerged, largely due to constraints in time, resources, and scope. Nevertheless, the realizations presented here provide not only a tangible demonstration of the SDS business concept but also a valuable means of showcasing the skills, methodologies, and research outcomes generated within WP5 by its core participants. In this way, the section serves as both an operational proof-of-concept and a narrative of the collaborative expertise that underpins the development of the SDS framework.

4 BUSINESS USE CASE 1: SCE LIFE CYCLE MANAGEMENT

In this business use case, the creation, configuration, conduction and management of a Supply Chian Entity (SCE) are realized.

This is the primary business for the SDS, as it was particularly conceived to handle the SCE as a special digital twin of the supply chain, aiming at its holistic control and improvement.

In this first business case, the clients of the SDS framework make use of some of the core provisions of the SDS platform. These applications are accessible at both the IT (Information Technology) level—encompassing enterprise systems for planning, analytics, and decision support—and the OT (Operational Technology) level, which includes shop-floor operations, automation, and process control. The dual availability ensures that information flows seamlessly across the entire hierarchy of the supply chain, from operational processes to strategic planning, aligning with paradigms established in frameworks such as RAMI 4.0 and Industry 5.0 reference models.

This configuration represents the basic and foundational case of SDS adoption. Here, the participants in a Supply Chain Entity (SCE) engage in a holistic commitment, where each actor benefits from access to shared services while simultaneously contributing to the collective performance of the SCE. The holistic dimension is crucial: clients are not only improving their own sustainability posture and efficiency, but they are also reinforcing the overall resilience, transparency, and effectiveness of the supply chain ecosystem to which they belong.

In practice, this means that actors use the SDS Shared Space—including tools like the Enterprise Service Bus, the Choreography Manager, and the SCE Dashboard—to coordinate activities, exchange data, and align performance goals. By committing to this collaborative approach, participants ensure that their individual improvement trajectories are tightly coupled with the continuous improvement of the SCE as a whole, thus turning sustainability and digital transformation into shared and systemic achievements rather than isolated initiatives.

The workflow of the business use case follows the activity diagram presented by Figure 2.4.1 of section 2.4. The SDS Business Unit is involved in collecting and providing requirements for the Stakeholders in order to create an SCE contract, with specific features depending on the particular supply chain. Note that the SCE represents the digital twin of what a specific consortium of stakeholders can provide in terms of information, capabilities, ambitions, and their individual goals.

In particular, for the primary scope of the ENOUGH project, the goal is to lower and control GHG emissions, but other and multiple objectives can be pursued in general with the same business case. It depends only by the SCE participants aims and scope. We recall in Figure 4.1, for convenience of the reader, the business use case picture. After the workflow is started, and the configurations achieved, the users participate to the SCE life cycle and get the benefits of it.

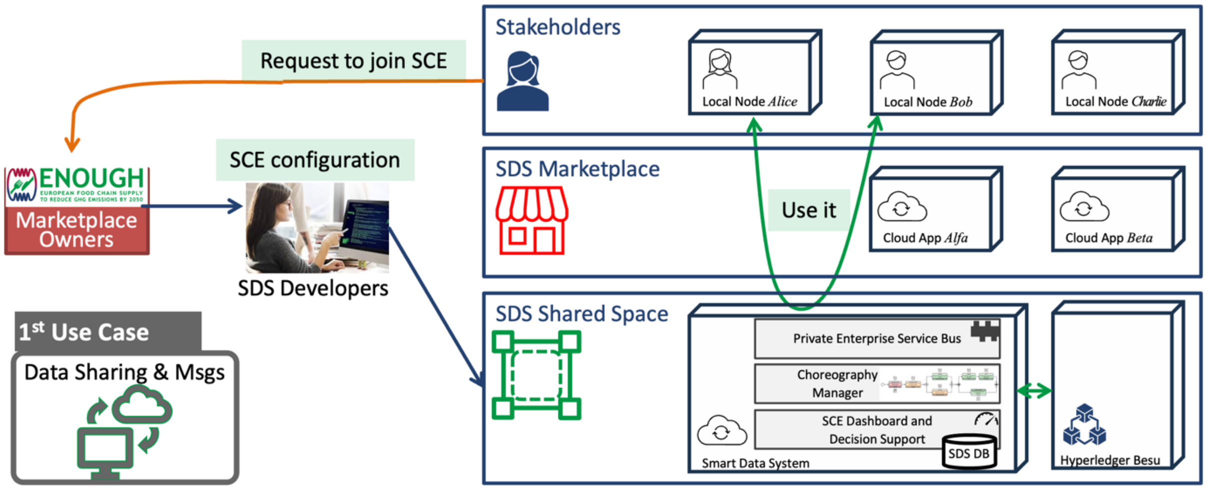

Figure 4.1: Business case 1 representation – Creation and conduction of the SCE (see section 3 for more details).

At the current state of the realization of the SDS business framework, the shared view of the SCE is achieved by a BPMN Choreography. During the activities of the ENOUGH project, two versions and relative experiments have been tested: the **SCE_D5.3** and the **SCE_D5.7**. They are connected to two slightly different narrations of the SCE, and served as a bench for the test of the SDS concept and potential.

In the following 4.1 and 4.2 subsections we describe the major characteristics of both the SCEs.

In Figure 4.1.1, and 4.2.1 we show respectively the two BPMN Choreographies used for the ENOUGH activities and demonstrations.

4.1 A summary of the SCE_D5.3 supply chain entity

The BPMN choreography depicted in the Figure 4.1.1 describes the full cycle of an order process as a coordinated interaction among six main participants: the *Retailer*, the *Coolstore*, the *Dairy Plant*, the two carriers *Transporter 1* and *Transporter 2*, and the *Display Cabinet Manager*. The process begins with the *Retailer* initiating an order request towards the two suppliers *Coolstore* and *Dairy Plant*. From this starting point, the choreography unfolds along two parallel directions.

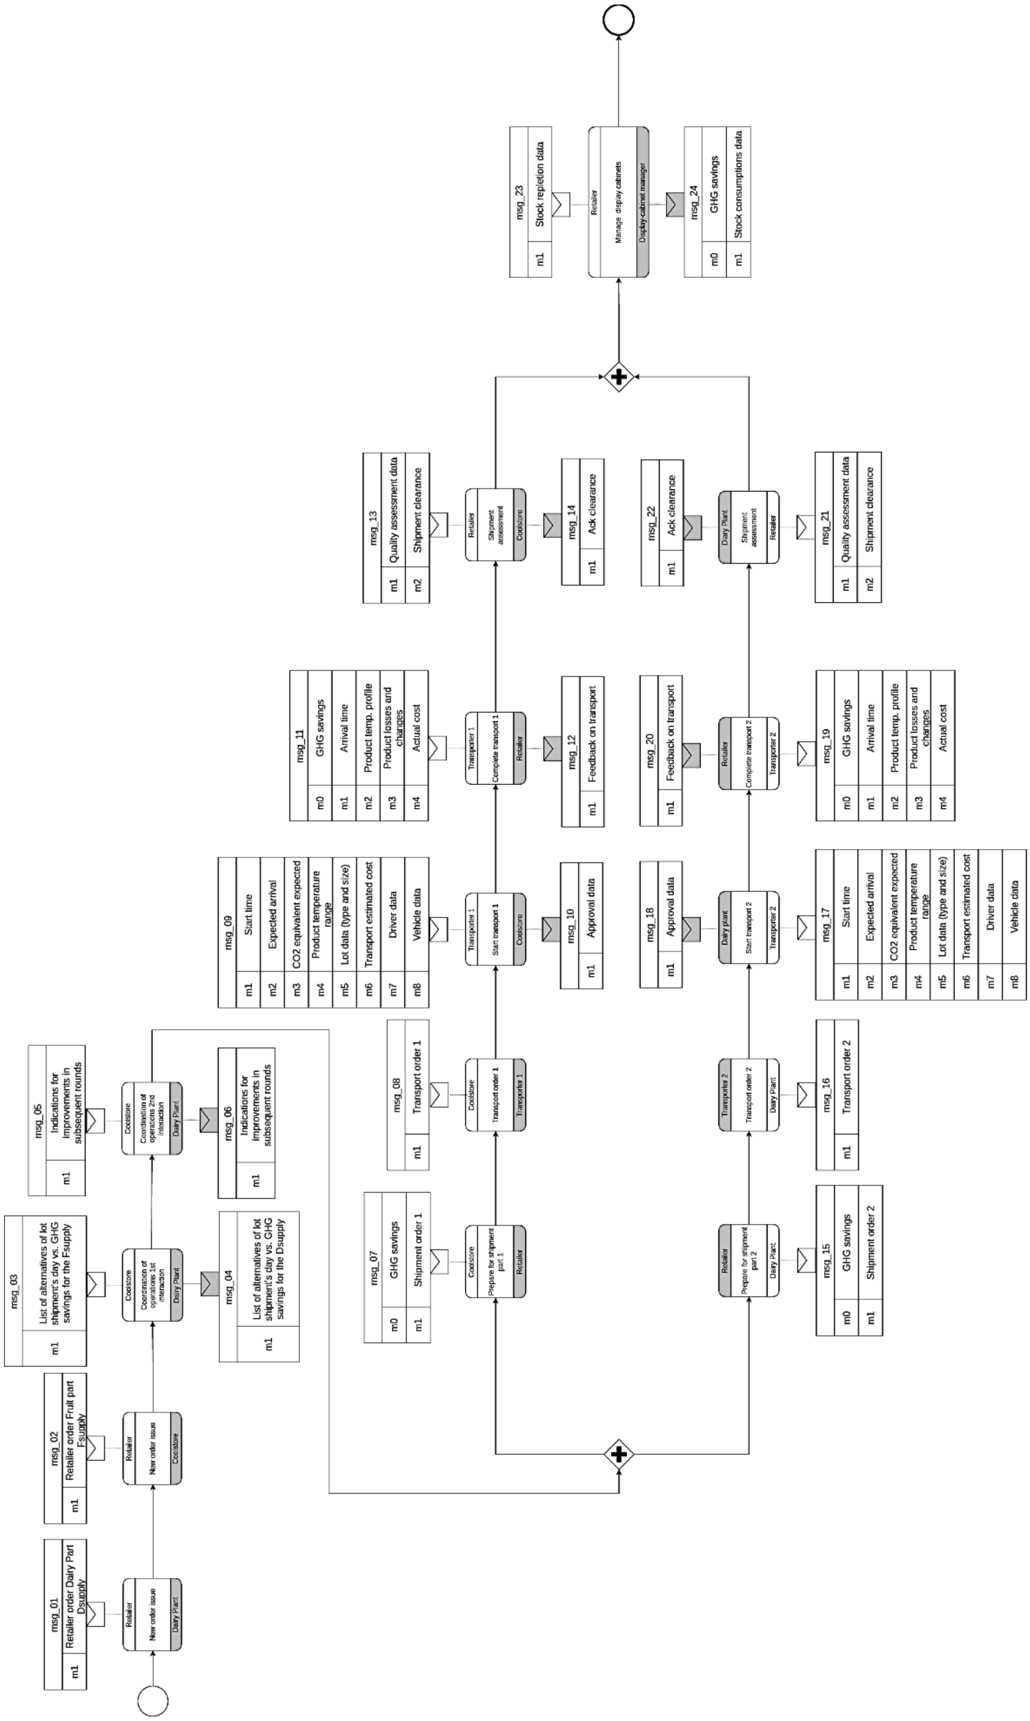

Figure 4.1.1: BPMN choreography of the SCE_D5.3 supply chain entity.

The short narration that explains the structure of the SCE_D5.3 is as follows:

- A **Retailer**, (virtual Participant in this experiment) issues every week an order composed in part of fruits (*Fsupply*) and a part of dairy products (*Dsupply*).
- The **Retailer** provides a section of the shop with display cabinets selling the "ENOUGH" certified low-GHG products.
- These products will be coming from the SCE_D5.3 in which at least two Participants **Coolstore** (Demo 5 of the ENOUGH project[21]), and **Dairy Plant** (Demo 3 of the ENOUGH project[22]) will negotiate to achieve jointly the best GHG savings for the shipments requested by the **Retailer**.
- While working to improve their local performance and gain control through the use of renewables, the two Participants exchange information through the SDS infrastructure in order to align their purposes toward a common goal.

In this figure, the contents of the messages exchanged in the choreography are explicitly reported. However, there is also a field, not shown everywhere and left implicit when not much relevant, that pervades every message: **m0**. This field is included in each message to indicate the amount of GHG savings achieved in the specific task. It provides the foundation for the performance indicators that drive improvements within this particular SCE.

Here the Retailer Actor is the driver of the SC.

The Retailer weekly starts the choreography of the SCE_D5.3 by issuing an order of:

- A certain lot of fruit, called hereafter the *Fsupply*.
- A certain lot of dairy products; called here the *Dsupply.*
- A comprehensive order for the Retailer is the combination of the *Fsupply* plus the *Dsupply*; they form the *FDsupply*, which is the final "thing" that Retailer wants to achieve from this SC weekly. The rate of *Fsupply* and *Dsupply* within the *FDsupply* is something that the Retailer updates weekly based on consumers' demand, food waste accounting, and other possible constraints that are generated by the Display Cabinet Manager, an Actor that is kept virtual for the present scope.

This Retailer owns a bank of display cabinets and shelves dedicated to the SCE_D5.3, in which the marketing puts on a big advertisement: "Products from the ENOUGH supply chain: low GHG guarantee!". In these dedicated cabinets only the *Fsupply* and the *Dsupply* each week are displayed in plain sight, ready on late Friday afternoon for the maximum influx of customers on Saturday.

Every Sunday morning the SCE_D5.3 restarts a new session. Every Sunday morning the chain starts again the Choreography, because the Retailer is able to assess the *FDsupply* for the subsequent week, having all the data and information available to do that.

The *msg_00* communicates to the Dairy Plant the *Dsupply* and the *msg_01* communicates the *Fsupply* to the *Coolstore*.

Now, the *Coolstore* and the Diary Plant start a collaborative negotiation. At that very moment, they are the core of the ENOUGH SC spirit as advertised, and they give example even to the others with this first collaboration experiment. To perform the experiment all the Actors of the SCE_D5.3 should have collaborated; but as happens often in reality someone at the beginning goes in tow of the others.

[21] https://enough-emissions.eu/demonstrator/rq-based-dynamic-controlled-atmosphere-storage-of-fresh-fruit/

[22] https://enough-emissions.eu/demonstrator/hthp-dairy-austria/

Nonetheless, if the incentive inner mechanism is working, in the medium term more participants will be more active for the good of the SC and society.

Therefore, the first contract associated to the SCE_D5.3, gives plenty of freedom to all the Actors involved, but Coolstore and Diary Plant have one constraint to satisfy. Recalling section 2.1, these two Actors are committed to:

1. With message *msg_02*, Coolstore has to (within Sunday 1:00 pm) communicate to Dairy Plant its list of GHG emissions (estimate) depending on the delivery day, from Monday to Thursday (at 11 pm).
2. With message *msg_03*, Dairy Plant has to (within Sunday 1:00 pm) communicate to Coolstore its list of GHG emissions (estimate) depending on the delivery day, from Monday to Thursday (at 11 pm).
3. On arrival of these messages, Coolstore and Diary Plant must choose the same shipment day in order to contribute to the same *FDsupply* timely and to minimise the sum of the respective GHG emissions for the whole SC in the *k*-th week. They obey to a policy that is shared between them since their signing and commitment into the SCE_D5.3.

In Figure 4.1.2, the processes here recalled are shown in BPMN notation.

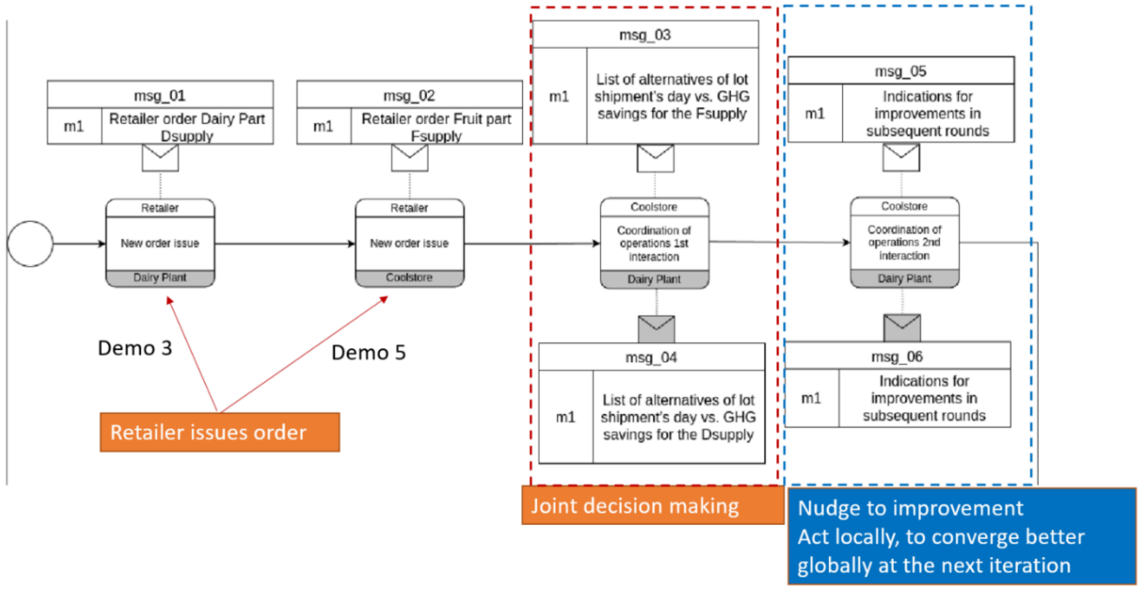

Figure 4.1.2: Crucial part of the SCE_D5.3, in which the negotiation happens through *msg_03* and *msg_04*. By means of *msg_05* and *msg_06* actions can be undertaken towards a convergent behaviour in subsequent session iterations.

Each of these two Actors, on each *k*-th round of the BPMN, store their information about performance locally — a private account and trace of the data about the SCE process. Moreover, the two actors thereto may use a digital twin to predict the emission, cost and quality for the coming week, and communicate the daily outputs. The performance data of the two Actors may be based on the following internal actions' set:

- **Coolstore**:
  - Prediction of daily energy use depending on outside temperature forecast
  - Prediction of use of renewable energy for the coming week based on weather forecast
  - Prediction of cooling cost
  - Prediction of quality change and remaining shelf life

- **Dairy Plant**:
  - Prediction of daily energy use depending on exogenous factors like production and CIP sessions
  - Prediction of use of renewable energy for the coming week based on exogenous factors like production and CIP sessions
  - Prediction of natural gas cost

These internal processes will produce local data that are used for internal management of the Participant. Recording the value function data opens up new perspectives. Participants can not only certify their performance but also, in the long term, demonstrate their virtuous behavior with respect to the SCE_D5.3 commitments and receive credit for it. Of course, self-certification requires mechanisms based on shared trust (or decentralized) as achivable by means of the Blockchain framework. These aspects have been studies and part of the research in this project as well[23,24].

Since medium to long-term credit is an incentive mechanism, control mechanisms can also be established in the shorter term. In particular, the messages *msg_04* and *msg_05*, can be used to provide desired "day of the week" for the next round of shipments. For example, with *msg_04* the *Coolstore* Actor expresses their preferred day of the week for the next one, basing on their own predictions and simulations. This could have an effect on the *Dairy Plant* management, as they could tweak something in their process in order to match this requirement for the next iteration, and avoid extra prices or losses necessary to abide by the SCE_D5.3 policy. In turn, with *msg_05*, the *Dairy Plant* can provide similar information to *Coolstore*. Convergence is not guaranteed due to many factors, but an infrastructure for continuous improvement is set.

In Figure 4.1.2 above, a particular of the SCE_D5.3 choreography is shown. With msg_01 and msg_02 the order from *Retailer* is issued (transferred by communication). From that event on both *Coolstore* and *Dairy Plant* have the data to process their joint shipment solution, optimal with respect to the agreed policy of pursuing the maximum sum of the GHG savings for the sake of the whole SC (and maybe not optimal in local sense).

They may hope that, in the next iteration, outcomes will improve and that the global optimum will align with the local optimum, thereby minimizing losses and enhancing both local and overall processes through better organization. To achieve this, they have another opportunity by communicating a nudge to the other Participant, aiming to foster convergence in the medium to long term.

This nudging provokes an internal re-organisation of the processes. For example, the Dairy Plant might act by adjusting the production schedule.

---

[23] Pirani, M., Cucchiarelli, A., & Spalazzi, L. (2024, November). A Role of RMAS, Blockchain, and Zero-Knowledge Proof in Sustainable Supply Chains. In *IECON 2024-50th Annual Conference of the IEEE Industrial Electronics Society* (pp. 1-4). IEEE.

[24] Pirani, M., Cucchiarelli, A., Naeem, T., & Spalazzi, L. (2025, May). Verifiable Actor Model Systems Through Relational-Model Multi-Agent System and Zero-Knowledge Proofs. In *2025 IEEE 8th International Conference on Industrial Cyber-Physical Systems (ICPS)* (pp. 01-06). IEEE.

Discussion about incentives can be continued considering another long-term effect. We notice that in the case of the Dairy Plant Actor controlled by TUGraz partner, the production cannot currently be planned or controlled. The schedules of the CIP sessions and the actual shipment schedules are currently completely independent variables with respect to the use of the HTHP, which is the value-added part of the process in considering the use of renewables in the process. The use of the HTHP guarantees lower GHG emissions, but its activation is currently beyond control of the ENOUGH experimental settings. This fact can be interpreted favourably in any case. The commitment in the SCE_D5.3 implies that with the aforementioned negotiation, if Dairy Plant achieves to produce the *Dsupply* in the favourable HTHP activity period for the joint shipment with *FDsupply*, it means that they have behaved well in a specific part of their production site when the SCE_D5.3 is concerned. The *Dsupply* lots will reach the ENOUGH-rank quality. Other productions of theirs will still participate to other supply chains that cannot sport the low GHG emissions as in this case. If the ENOUGH experiment is successful, the Dairy Plant management might rearrange their productions in order to achieve similar performance even on other supply chains and then "virally" express the ENOUGH values and its impacts.

For the sake of this first proof of concept about the role of the SDS in improving the lowering of GHG emissions for a SC, we will establish a simple but handy specification on the SCE_D5.3. In general, we search for KPIs to be expressed within the messages. Possibly computed by a specific operational Actor (not necessarily participating to the choreography). In this initial experimental setup, the choice was to dedicate a slot of every message (**m0**) exchanged in the choreography for a declaration from the Actors about the CO_2 equivalents saved by their own processes. This will be, as a lumped value that considers all kind of possible emission generators (even the digital computing itself). This position will permit a first simple dashboard to compare the rounds of the SC and then appreciate if the collaboration among at least two Actors produces increasing benefits to the whole.

For example: in the SCE_D5.3 the *m1* field of *msg_04* will contain the GHG level after negotiation for *Coolstore*; the same *m1* field of *msg_05* will contain the GHG level after negotiation for the *Dairy Plant*. When not relevant the value can be left empty (or zero) in other messages. Negative values can be admitted when a carbon capture is allowed by a process.

In some detail:

- **Coolstore** and **Dairy Plant** exchange GHG savings estimation (prediction) with messages *msg_03* and *msg_04*.
- Both of them obtain the graph of Figure 4.1.3, and commit themselves to use the **Tuesday** as the best day for the shipment, although Coolstore would have preferred Monday and the Dairy Plant Thursday, if they were acting individually. This is because both aim at maximizing the collective performance, even at the cost of sacrificing immediate individual gains. Nevertheless, this sacrifice is documented and recorded, remaining as both an intangible credit (such as reputation) and a tangible one in cases where explicit compensation or reimbursement is provided.

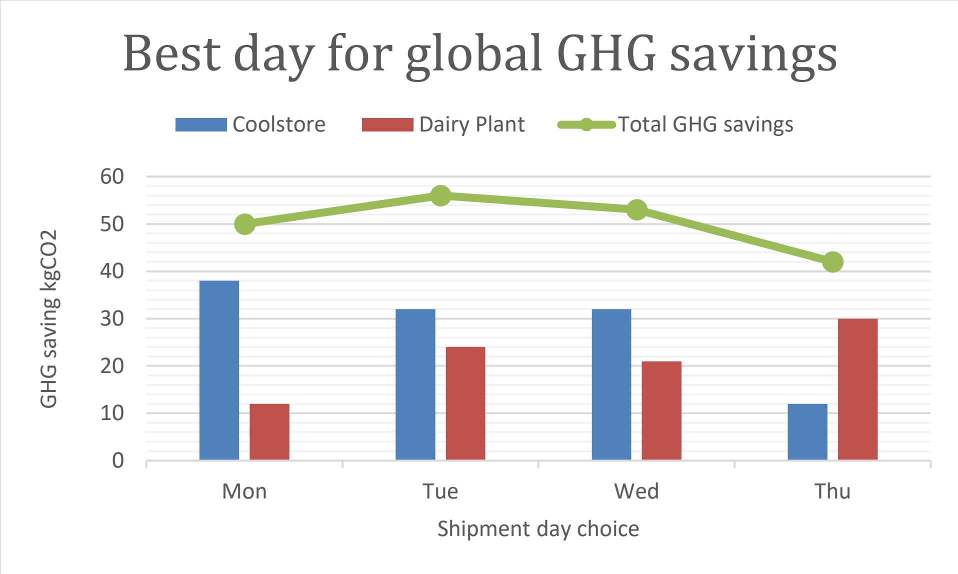

Figure 4.1.3: Data available to Coolstore and Dairy Plant Participants after msg_03 and msg_04. The selected shipment day corresponds to the maximum of the total GHG savings.

An important remark must be made about the data that finally will be used to achieve a KPI on GHG savings for the SCE through its many sessions, and so to check the improvements eventually obtained. Note that the contents of the *msg_03* and *msg_04* contain only forecasts or predictions. These predictions rely on the quality of the simulators and digital twins available to the Participants. Thus, a comparison between the two values (before and after measurement) can give also a hint on the performance of the predictors themselves. This can constitute a secondary performance indicator that will trigger improvements in this respect.

Nonetheless, in this example, the primary objective is to focus on a KPI that is directly linked to the GHG saving actually achieved, for the sake of simplicity, and minimum interaction. For a first assessment of the experiments these values will be the same as produced by simulations, but in a real case a measurement system for final values must be in place.

With this first experiment, a rehearsal of the outcome of the demonstration of Demo 1 in WP6 was made. The expected Final Demo 1 results will demonstrate how the participation to the SDS framework of the stakeholders in an ENOUGH SCE can produce continuous improvement of the supply chain involved.

4.2 A summary of the SCE_D5.7 supply chain entity

With this second layout, a more realistic and slighlty complex BPMN was achieved. This second step was made to adhere more to the actual developments of the demonstration activities and the associated research activities. In addition, this was occasion to refine and upgrade most of the bugs and uncertain requirements experienced with the first SCE instance.

In this version three different BPMN gateways have been used to comply with the new narration as follows.

As the Coolstore and Dairy Plant have usually different timings, this new BPMN considers that one of the branches might be repeated several times and the other might stall unless the timings of the two orders from the Retailer happen in the same week. As before, the Retailer typically issues weekly orders to the Dairy Plant, because the nature of the dairy orders and production complies to this frequence. The Coolstore process acts on different time scales. Typically, the opening of cool store rooms are opened once or twice a year. The controlled atmosphere storage is made to plan with accuracy the most appropriate period for a lot of fruits to be used and shipped. Once a store is opened, all the goods have to be shipped. Some of them will be used for the SCE and the Retailer's orders, which in this case might have a period of more than one week.

This behaviour is possible with the Inclusive *OR gateway* (the one after *msg_14*), which allows one or more of the choreography branches to be followed independently of each other.

The choreography begins with the Retailer sending a *new order base* message simultaneously to the Coolstore and the Dairy Plant. Both branches then initiate their own negotiation process with the Retailer. The Coolstore negotiates order requirements, while the Dairy Plant does the same for its own production planning. These negotiations are separate but run in parallel, reflecting the different operational rhythms of the two supply sources.

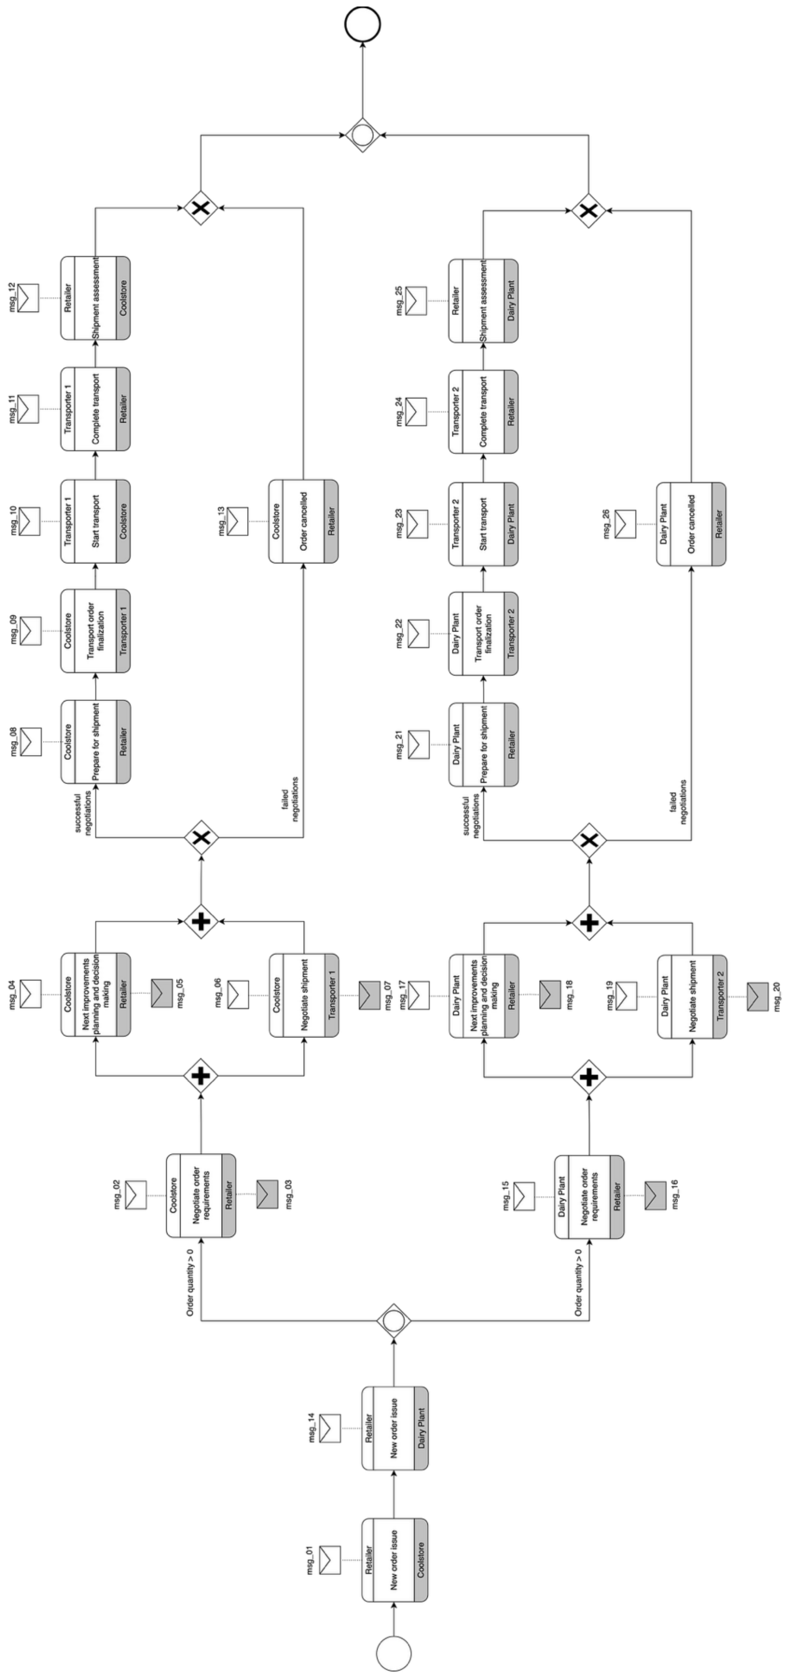

Figure 4.2.1: The SCE_D5.7 BPMN Choreography. Inclusive OR gateway after message msg_14 in the choreography, allowing asynchronous execution of the Coolstore and Dairy Plant branches.

After this negotiation phase, both the Coolstore and the Dairy Plant confirm the order to the Retailer. At this point, the workflow diverges through an *inclusive OR gateway*: either or both branches may proceed, depending on the timing and availability of goods.

On the Coolstore side, once the order is confirmed, the store prepares the shipment and transmits shipping information to Transporter 1, which starts and then completes the transport to the Retailer. The Retailer subsequently acknowledges the shipment assessment and closure of the process. This reflects the bulk and seasonal nature of fruit storage: once a store is opened, a full shipment cycle is triggered.

On the Dairy Plant side, the order confirmation leads to preparation for shipment, followed by notification to Transporter 2. As with the Coolstore branch, Transporter 2 carries out the shipment and delivers it to the Retailer, who again performs shipment assessment and closes the process. This reflects the weekly, cyclical rhythm typical of dairy logistics.

The choreography accommodates the fact that the Dairy Plant's weekly cycles may repeat multiple times before the Coolstore branch is activated again. This mechanism allows the SCE to manage asynchronous processes while still converging on the Retailer as the final recipient.

Another difference with the previous SCE is that the main negotiations do not occur here between the Coolstore and the Dairy Plant, but more realistically between each of them and the Retailer, who manages the two suppliers as independent entities, together with their respective transporters. It is the Retailer who, particularly in managing the display cabinets for dairy products, coordinates their presentation and promotion to customers alongside other products from the same low-GHG emission supply chain, such as the fruit and vegetable products supplied by the Coolstore.

4.3 The SCE configuration

For both of the SCEs described above, the SDS Platform must receive a configuration as outlined in the workflow of Figure 2.4.1 in Section 2.4. Most of this process cannot be fully automated, as it depends on requirements, agreements, and specifications involving a heterogeneous and numerous set of human actors. Nonetheless, the SDS Business Unit has been designed for this purpose, and a dedicated toolchain has been developed within the SDS Platform to support these back-office operations, and so the production and test of the SCE.

The SCE Dashboard and SCE Configurator is a suite of Web tools available both to administrators and to users, in particular to the Participants, to monitor the proceeding of the SCE during its work. It is also a major means for decision making about the collaborative goal of lowering overall GHG emissions in the SCE.

In particular, in the SDS Configurator, Supply Chain Models serve as templates for creating SCE instances, specifying core components like activities, Participants, and Messages. The model editor allows tasks and workflows to be defined, validated, and saved.

Supply Chain Instances are derived from models. These instances capture specific supply chain configurations. Editors enable fine-tuning tasks and Participant assignments while ensuring validation of roles like initiators and receivers. Published instances undergo automated tests and checks.

This application framework exemplifies a well-structured approach to managing supply chains, tasks, and stakeholders while integrating robust monitoring, validation, and user management capabilities.

Currently, this configuration tool is dedicated to Participants rather than to generic Actors. This was a deliberate choice. Although the SDS Infrastructure technology, based on a publish/subscribe scheme of RMAS[25], can in principle host and integrate any type of format, the number and diversity of potential Actors is extremely broad. Moreover, there are already numerous platforms and products that specifically address interoperability and configuration in the context of industrial automation.

The requirements for Actors cannot be fully specified and must be handled on a case-by-case basis. The case of considering also this configuration will depend on the actual nature and use of the SCE part, which for the moment can be left manual. The focus at the moment and for the current status of business plan is to focus the valuable properties of the Choreography, while the standard OT (operations technology) level of interaction is considered a special and specific case.

A **Supply Chain Model** represents the basic structure of a supply chain, where components such as start points and end points, individual activities, and messages must be specified. It has a 0-N relationship (from none up to a number N) with Supply Chain Instances (Instances of SCEs), where the Participants in the Choreography must be specified.

A Stakeholder can have multiple participants, each with at least one type (and possibly more).

Tasks must specify a source participant type and a target participant type. Each task must define a source message, with an optional target message.

The SCE Configurator provides a structured environment for defining, instantiating, and monitoring supply chain entities within the SDS Platform. At its core, it operates through a sequence of pages that progressively guide stakeholders from high-level models to fully operational instances. A sense of these pages can be seen in Figure 4.3.1. The process begins with the Home Page, which serves as the entry point and overview dashboard. Here, users can immediately access global statistics, manage stakeholders, and drill down into the details of participants. Stakeholders can create and configure multiple participants, each classified by type, which will later be used in defining the roles of supply chain tasks. This ensures that the representation of the network mirrors the complexity of real actors in the chain (Figure 4.3.1).

[25] Pirani, M., Bonci, A., & Longhi, S. (2022). Towards a formal model of computation for RMAS. *Procedia Computer Science*, *200*, 865-877.

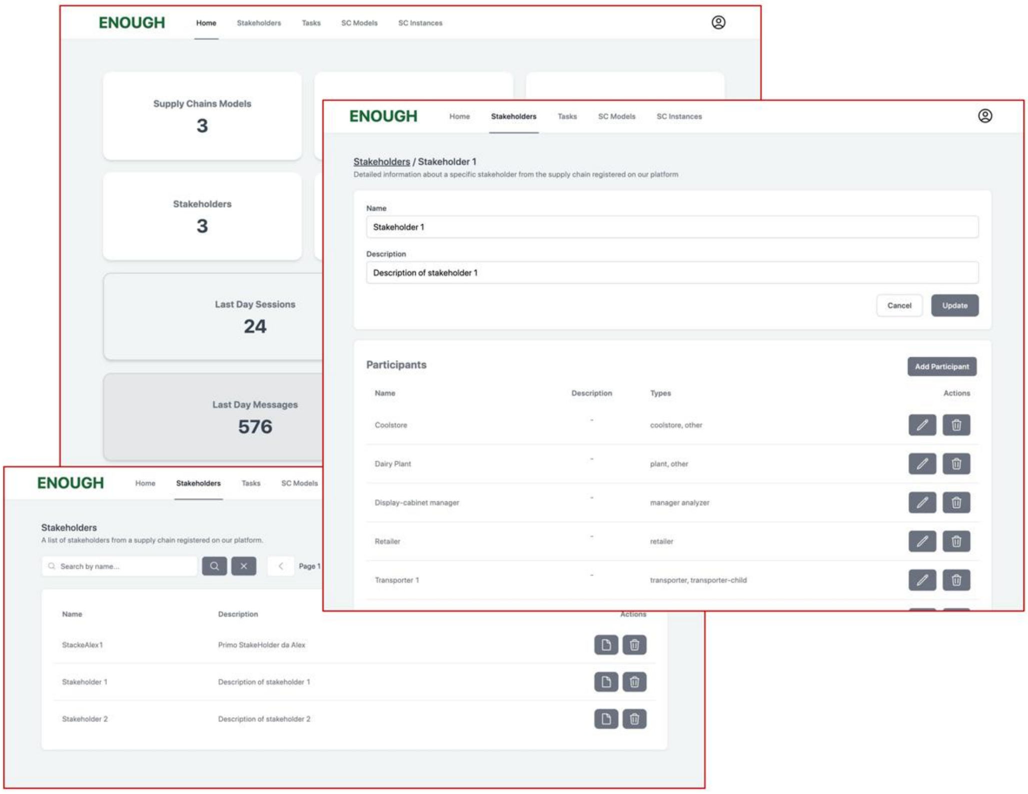

Figure 4.3.1: SCE Configuration: snapshots of Home and Stakeholder widgets and GUIs.

Next, the configurator turns to the definition of **Tasks and Messages** (Figure 4.3.2). Tasks describe interactions between participants, specifying a source and a target, while messages capture the content exchanged. Particular attention is given to the property *m0*, which records GHG savings and later aggregates into KPIs. Through this mechanism, every operational message becomes not just a coordination signal but also a data point for sustainability monitoring.

Once tasks are established, the system advances to the construction of **Supply Chain Models** (Figure 4.3.3). These models act as reusable templates that define the choreography of the supply chain. In the editor, tasks and participants are arranged into workflows, and validations ensure coherence—such as the mandatory presence of an initial message. Publishing a model locks its structure, transforming it into a reliable reference for generating instances.

From models, users then move to **Supply Chain Instances** (Figure 4.3.4). Instances are concrete realizations of models, tailored with specific participants and task assignments. Editors allow initiator and receiver roles to be explicitly bound to each activity, with visual cues highlighting incomplete configurations. Once published, the instance is linked to a dedicated database created automatically by the SDS engine, ensuring operability.

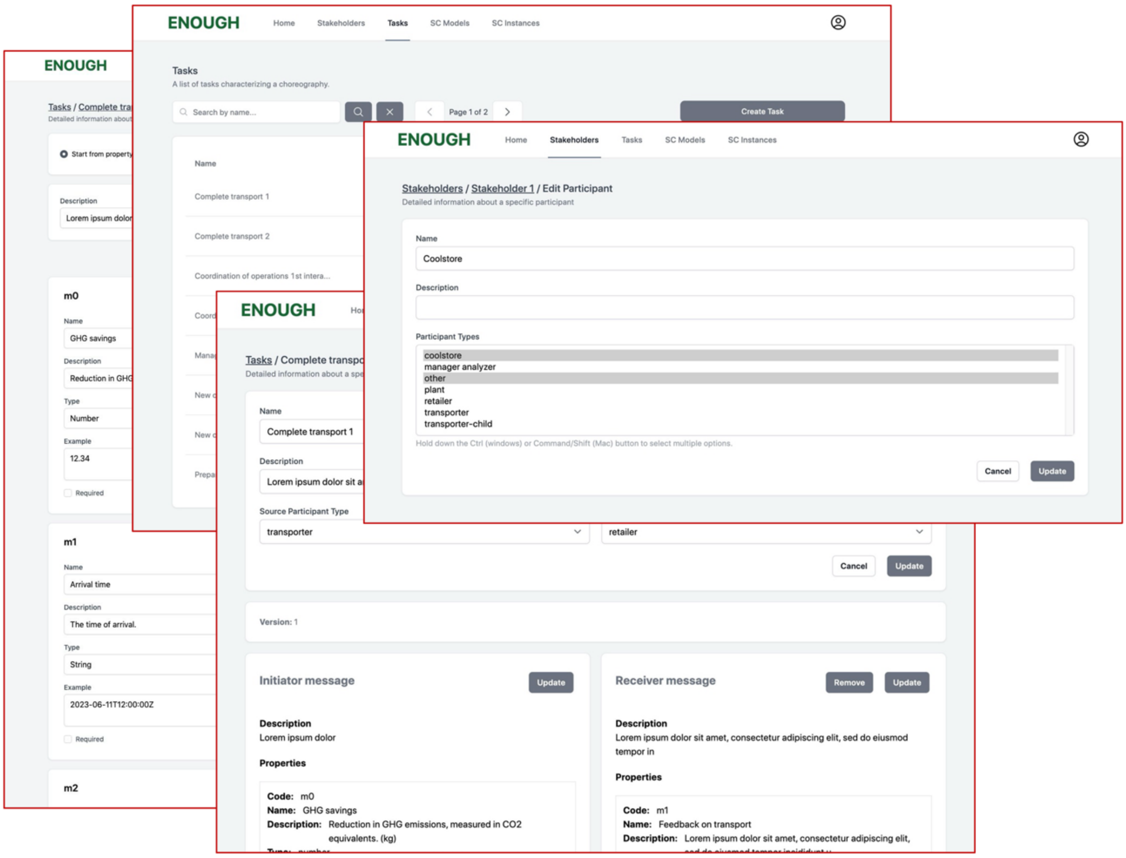

Figure 4.3.2: SCE Configuration: snapshots of Participants, Tasks, and Messages widgets and GUIs.

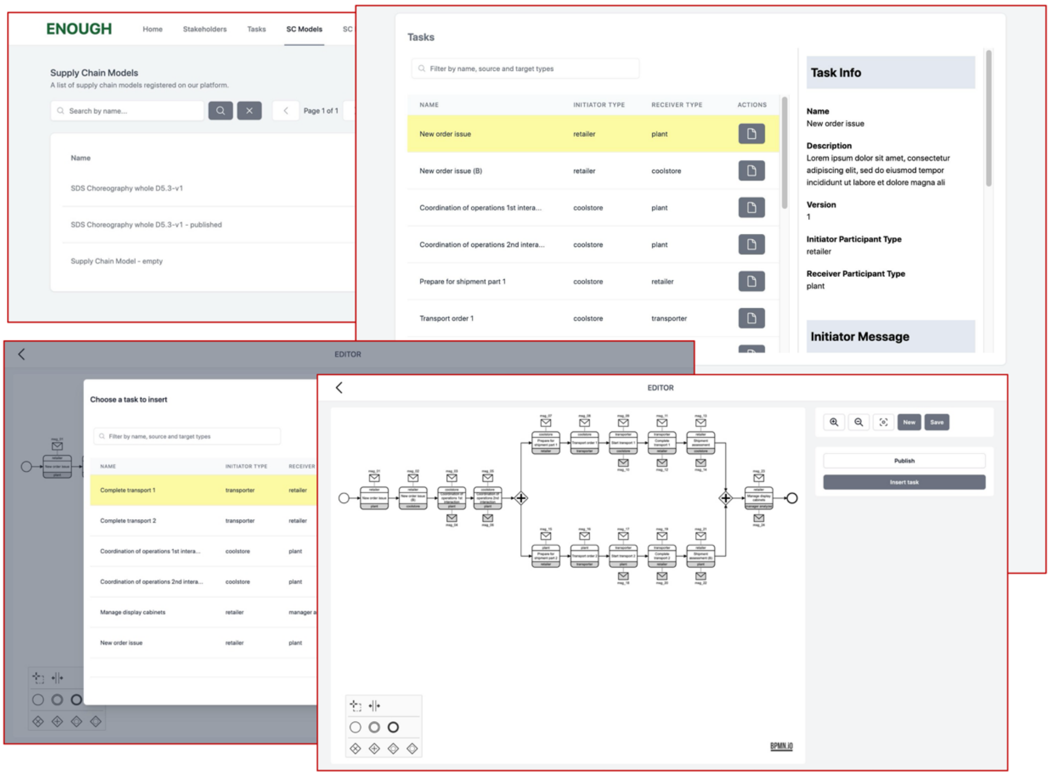

Figure 4.3.3: SCE Configuration: snapshots of Supply Chain Models widgets and GUIs.

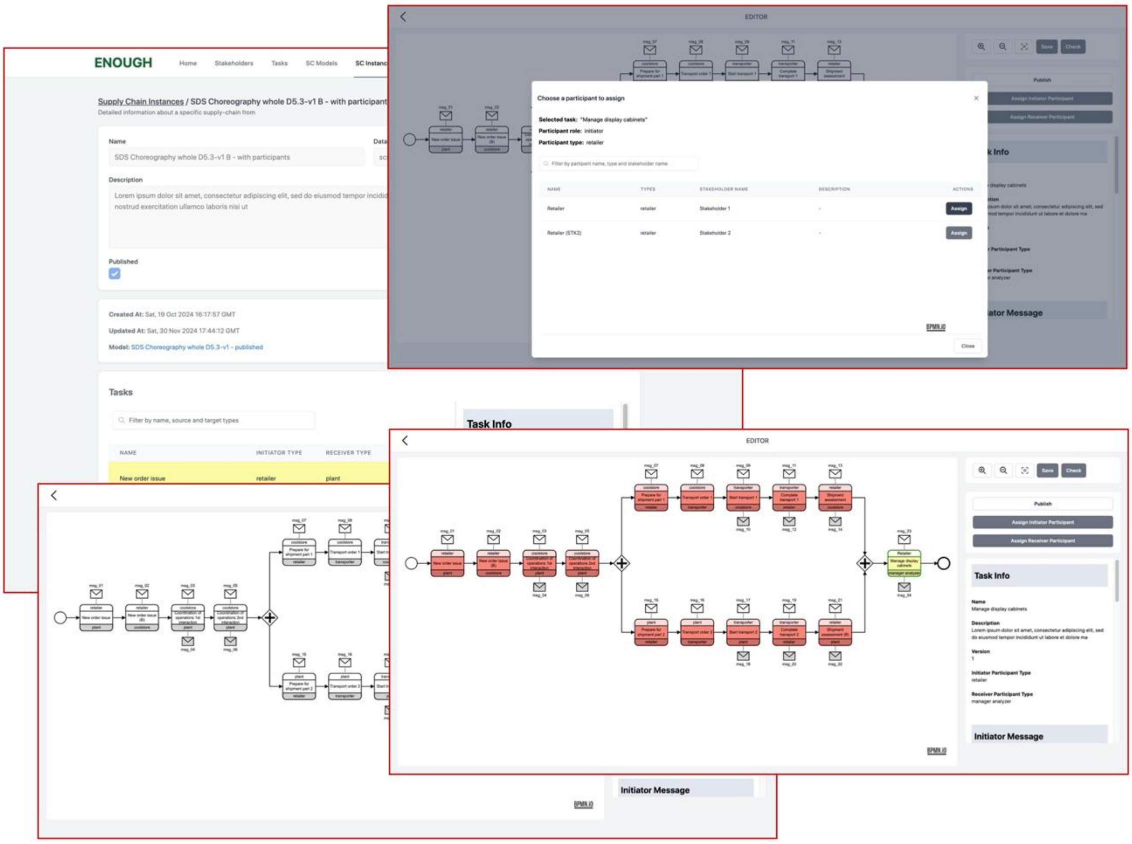

Figure 4.3.4: SCE Configuration: snapshots of Supply Chain Instances widgets and GUIs.

4.4 Running and monitoring of the SCE.

The main access for the Participants to the SCE Choreography management and monitoring is by means of a web page like in Figure 4.4.1. This should be understood as a purely experimental and provisional GUI. In an actual business phase, the graphical design and user experience can be customized for each specific SCE and for each Participant.

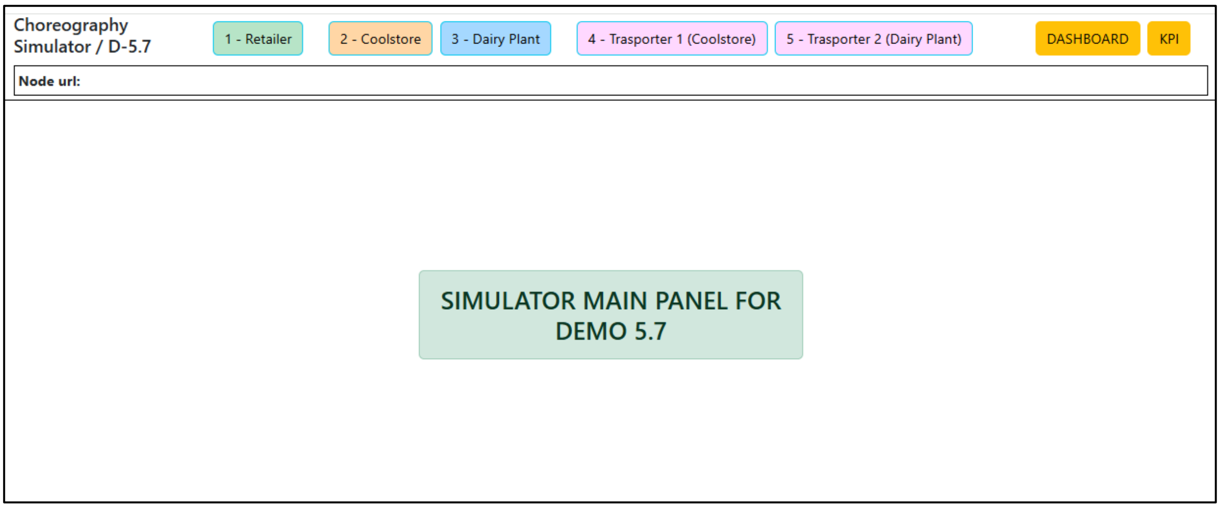

Figure 4.4.1: SCE Dashboard Web container page.

Indeed, the interaction page may be different for each of the Participants and not accessible by any other. In this way any Participant could have a customized interface reflecting their understanding and semantics of the data received through the choreography Messages; indeed, this is a key of interoperability and autonomy at the same time. Participants might want to use distinct web applications, developed with different development tools and on different operational platforms, all integrated into the SDS system. Not least, the different costs of the technologies can be tailored, while still allowing participants to join the SCE with a very low-profile and essential interface.

For the sake of the easiness of this test, the links to the five Participant interfaces have been collected into this collector web page at the top. Five buttons with the name of the Participant allow to access each Participant interface. In the real field each of them has a different access point and address. Through this panel, we can easily call up the user interfaces of the various applications to follow the evolution of the operational flow during the simulation of the SCE here achieved.

At the top of the panel are found five buttons, namely “Retailer”, “Coolstore”, “Dairy Plant”, “Transporter 1”, and “Transporter 2”. Pressing on these buttons the interface view for each of the Participants is displayed. There are also two other links to the monitoring dashboard and to the KPI dashboard.

The monitoring Dashboard of the SDS system allows any Participant to view the public data exchanged among them through Messages, analyze their details, and provides a BPMN choreography diagram to monitor the process. This will be the shared interface to the SCE Choreography. Dashboard will be the view that allows holistic information of the SCE to be read and used for the overall organization and decision making at a global level for the supply chain.

After the start phase of the SCE, the Participants can monitor the status of messages within the choreography and, through the other dedicated GUIs, perform the negotiation and collaboration actions necessary for the functioning of the choreography itself and thus for the advancement of the SCE process. In Figure 4.4.2, we show examples of the monitoring dashboard available to all the Participants and for the two SCEs here experienced.

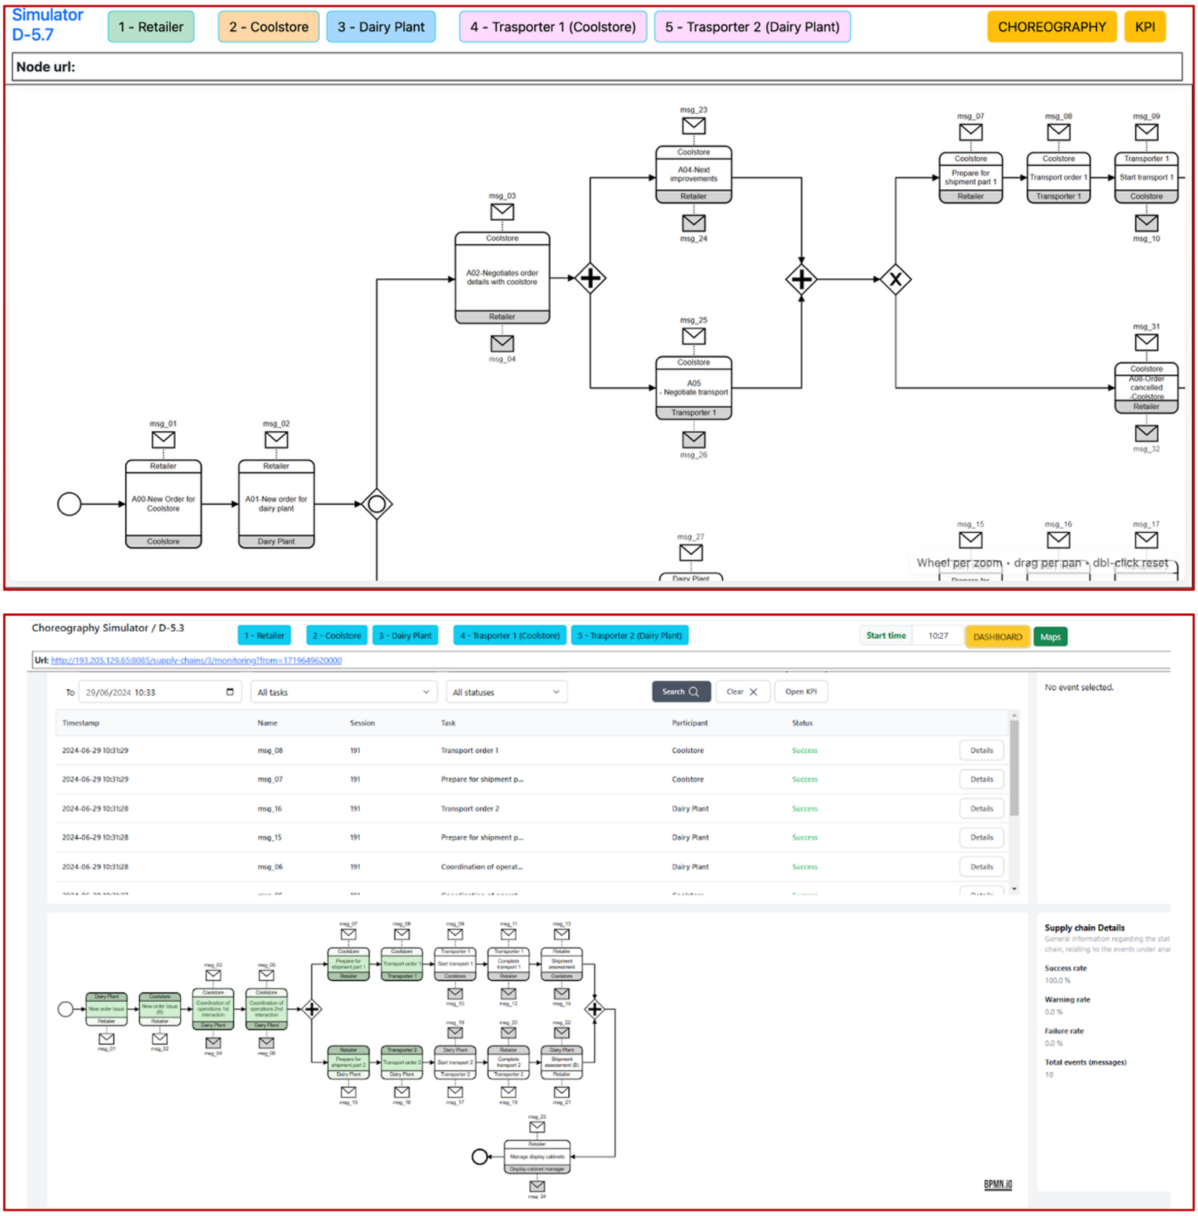

Figure 4.4.2: SCE Dashboard (top) for the monitoring of the SCE Choreography (bottom).

The Retailer – Create Requests interface (Figure 4.4.3) is the entry point for initiating the choreography of supply chain collaboration. From this screen, the Retailer specifies the product, quantity, measurement unit, and requested delivery date for each order. The interface is divided into two sections, reflecting the two independent suppliers in the choreography: the Coolstore, which provides fruit such as apples, and the Dairy Plant, which provides milk. For each supplier, the Retailer can toggle the *Enable request* option, which corresponds to the inclusive gateway in the BPMN model, allowing one or both branches to be activated depending on the timing and demand.

Once the details are entered, pressing the *Send orders* button generates and transmits the appropriate messages through the SDS platform. Technically, this step corresponds to the BPMN messages *msg_01* and *msg_14*, sent respectively to the Coolstore and Dairy Plant. These messages contain the structured order data and trigger the downstream choreography tasks. In the process, the Coolstore prepares shipments according to its seasonal cycle, while the Dairy Plant responds on its weekly production schedule.

The GUI therefore serves as the operational embodiment of the Retailer's initiating role in the choreography. It translates business inputs into formal SDS messages, ensuring that orders are captured, auditable, and compliant with the workflow logic. From this moment on, the negotiation, transport coordination, and delivery follow the paths already modeled in the BPMN, but all start with this simple yet pivotal interface.

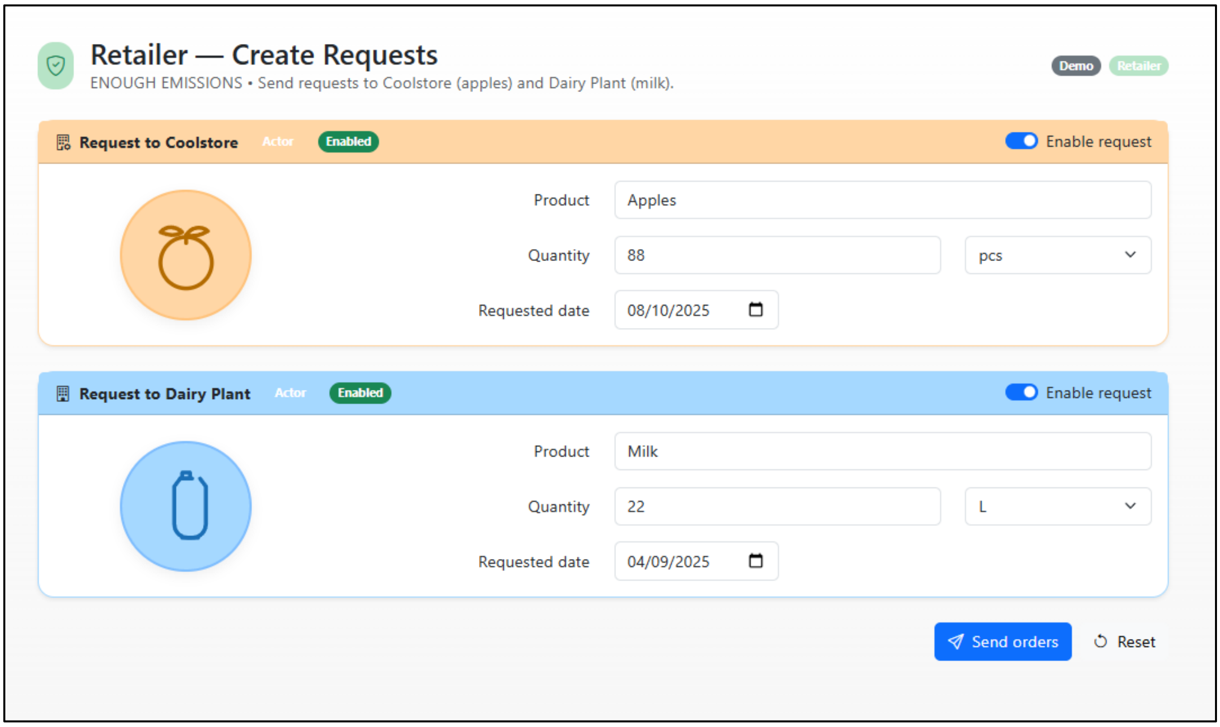

Figure 4.4.3: Retailer's GUI for orders placement. User interface for Retailer – Create Request, with product, quantity, unit, and date settings, and option to send orders.

The *Coolstore – Negotiate & Dispatch* interface (Figure 4.4.4) represents the second stage of the choreography, where the supplier responds to the Retailer's order. At the top, the *Incoming Retailer Request* section shows the original order details—product, quantity, unit, and requested date—received from the Retailer via the SDS system. This ensures that the Coolstore has full visibility of the demand before acting.

Below, the *Propose Alternatives* section allows the Coolstore to suggest modifications, such as an adjusted delivery date or quantity, often motivated by capacity limits or strategies to minimize GHG emissions. For example, the Coolstore may consolidate shipments to reduce empty transport runs or align deliveries with optimal storage opening periods. The proposal is then sent back to the Retailer, corresponding to the BPMN choreography where negotiation messages are exchanged between the two actors.

In addition, the *Next Improvements* field enables the Coolstore to send recommendations for longer-term sustainability improvements. These nudges—such as adopting reusable pallets or optimizing delivery time windows—extend beyond the immediate transaction, supporting continuous process improvement across the SCE.

This GUI therefore serves as the *negotiation and feedback node* of the choreography: it balances operational constraints with sustainability optimization while keeping the Retailer engaged in a constructive dialogue. The integration of proposals and improvement messages reflects how the choreography allows not only transactional exchanges but also the embedding of systemic, collaborative learning.

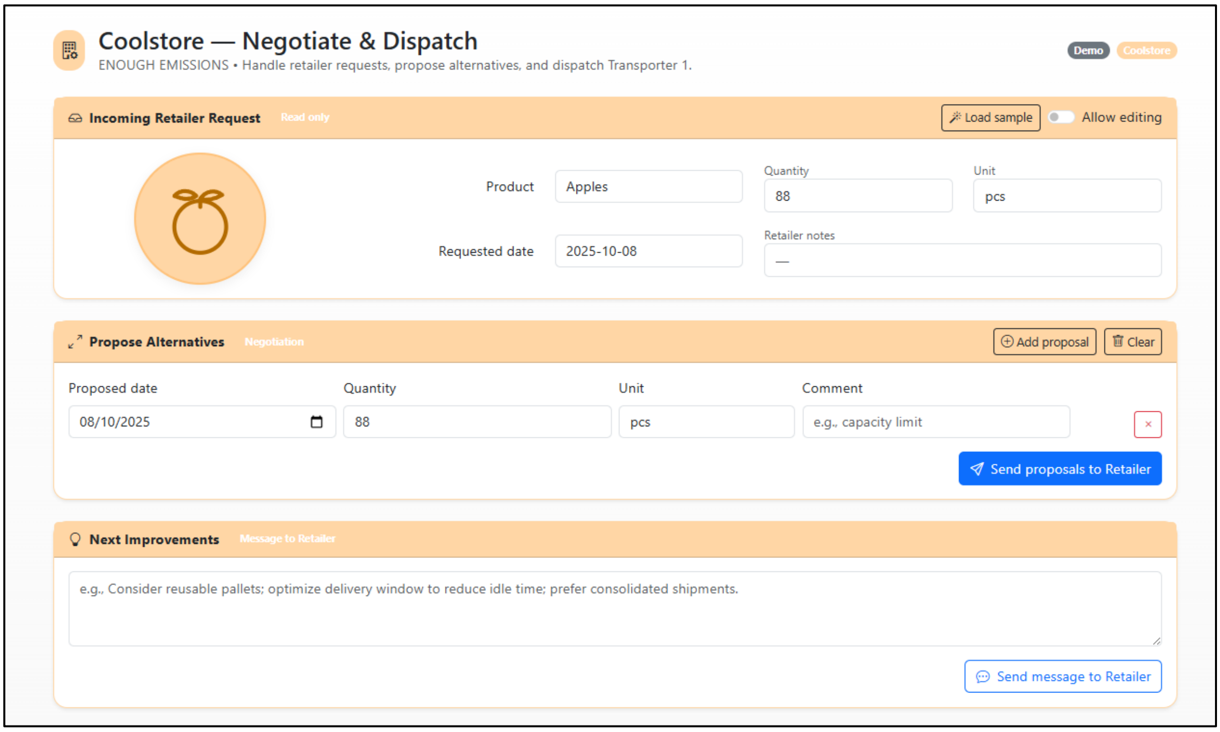

Figure 4.4.4: User interface for Coolstore – Negotiate, proposing alternative quantity/date for improved GHG savings.

The *Negotiate Transport Condition* interface (Figure 4.4.5) illustrates the *Coolstore – Dispatch* phase of the choreography, where the Coolstore finalizes logistics with *Transporter 1*.

At the top, the form captures the *Order reference* and the *Pickup date and time*, ensuring precise scheduling. The *Pickup location* (here, *Coolstore DC, Bay 3*) and the *Delivery location* (e.g., *Dairy Plant – Inbound*) define the physical path of the goods, tying the choreography explicitly to the supply chain geography.

Operational details such as quantity, temperature requirements (e.g., +4°C for cold-chain compliance), and special notes (e.g., “fragile goods: load first, no stacking”) allow the Coolstore to transmit critical conditions to the Transporter. These constraints ensure both product safety and adherence to sustainability objectives, such as minimizing spoilage and optimizing refrigerated loads.

The *Send* button formalizes the negotiation: once both parties agree on these parameters, the dispatch order becomes effective, and the transport phase is activated in the choreography.

This interface, therefore, represents the execution bridge between the Retailer's demand, the Coolstore's negotiation, and the Transporter's operational role. By standardizing how logistics conditions are communicated and accepted, it ensures traceability and reduces the risks of misalignment, while embedding sustainable transport practices into the SCE workflow.

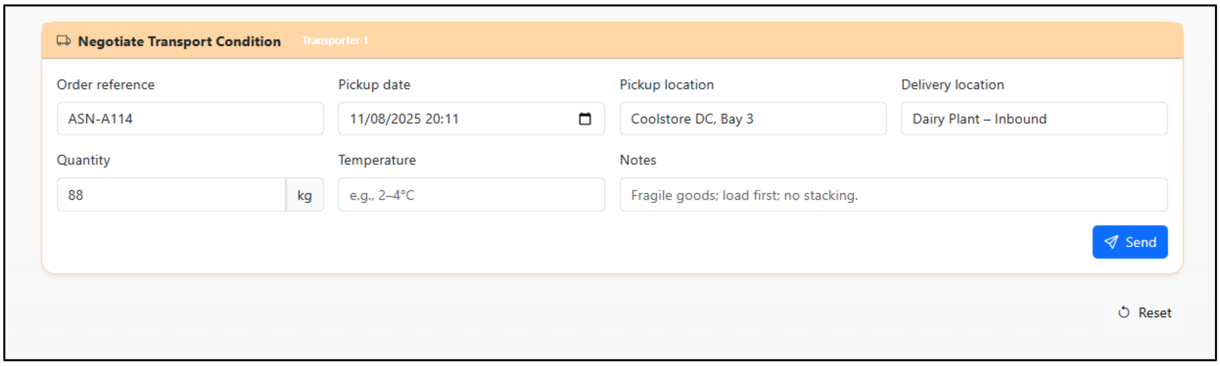

Figure 4.4.5: User interface for Coolstore – Dispatch, showing shipment negotiation with Transporter 1.

The two screenshots together illustrate the *Dairy Plant – Negotiate & Dispatch* interface (Figure 4.4.6) within the SDS platform.

At the top, the Dairy Plant receives the incoming order request from the Retailer, here concerning milk in the amount of 22 lots, with a requested delivery date. This section is read-only, providing a clear view of what the Retailer asked.

In the middle, the *Propose Alternatives* area allows the Dairy Plant to negotiate: it can adjust the proposed date, the quantity, or add explanatory comments (e.g., cooling capacity limits, batching requirements). These proposals are then sent back to the Retailer, enabling iterative negotiation to optimize logistics and CO_2 savings.

The *Next Improvements* section allows the Dairy Plant to suggest more structural changes for the future, such as shifting to returnable crates, consolidating outbound deliveries with other orders, or adjusting order lead-times. This creates a feedback loop that enhances sustainability and efficiency across iterations.

Once the Dairy Plant agrees on conditions with the Retailer, the Dispatch panel (first image) is used to set up the transport details with Transporter 2. Here, it specifies the order reference, pickup date, pickup and delivery locations (from the Dairy Plant outbound gate to the Retailer receiving dock), vehicle type (e.g., a refrigerated van), and temperature requirements. A checkbox allows optimization for lower CO_2 emissions if possible. Submitting this form sends the confirmed shipment request to Transporter 2, completing the dispatch cycle.

In summary, this interface mirrors the Coolstore application but is adapted to the Dairy Plant's role: negotiating with the Retailer, confirming feasible alternatives, and finally coordinating shipment via Transporter 2.

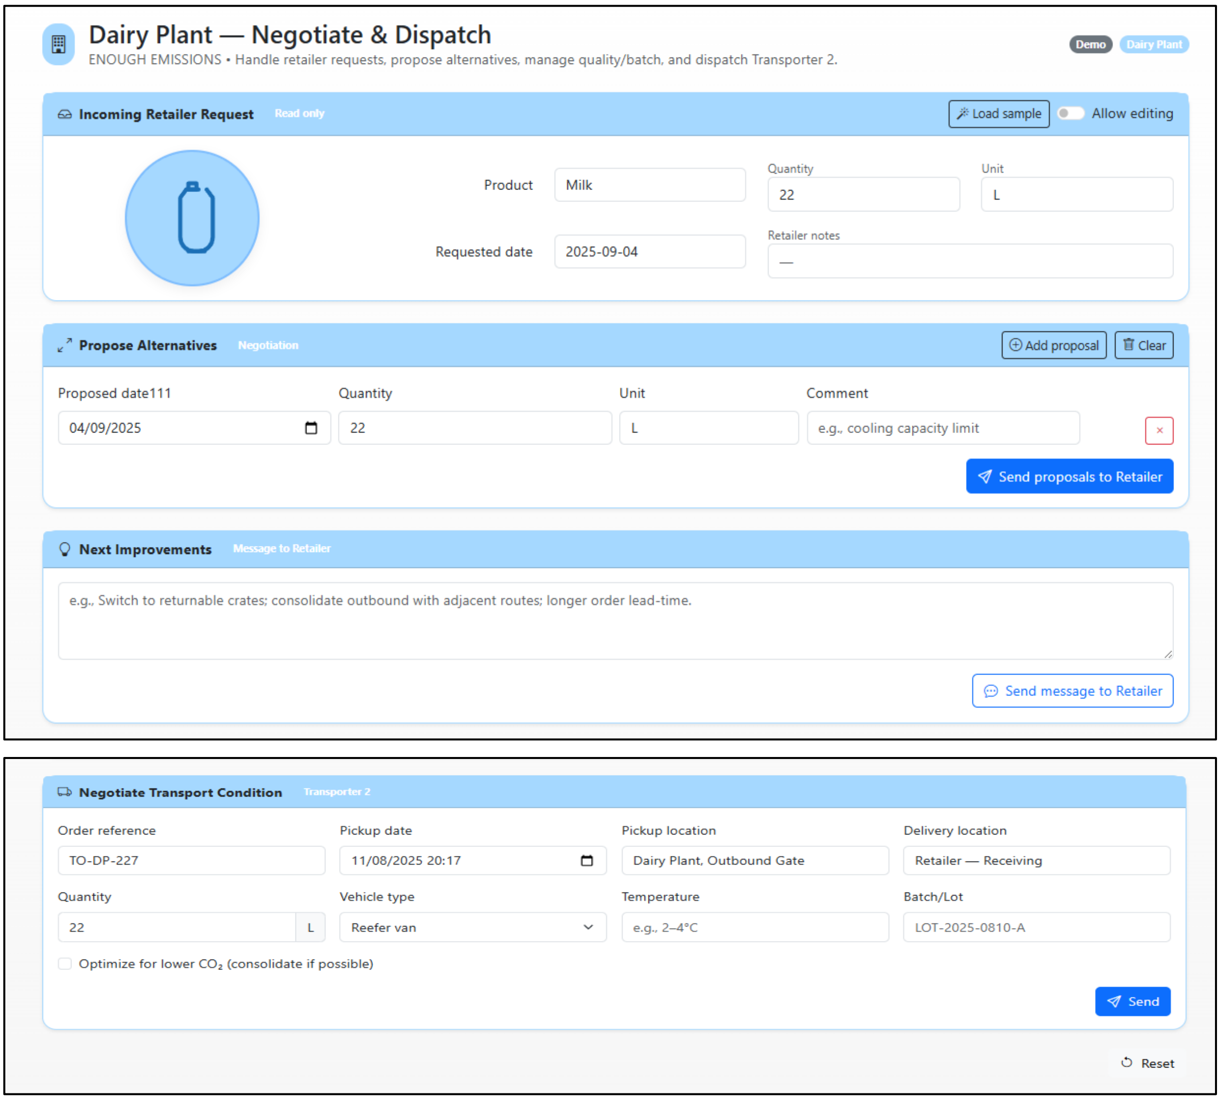

Figure 4.4.6: Dairy Plant's GUIs for negotiation (top) and transport (bottom).

The following screenshot shows the *Transport Orders for Coolstore* interface (Figure 4.4.7), which illustrates the role of *Transporter 1* in the SDS choreography.

At the top, the interface lists the active transport orders with key details: the shipper, pickup and destination locations, estimated delivery date, number of items, and current status (here marked as *ACQUIRED*). Below, the *Order Details* section expands the selected order. It specifies the pickup location (Avenida da Liberdade, Lisbon, Portugal), the destination (Rue de Rivoli, Paris, France), the requested and estimated shipping dates, the shipping method (TIR in this case), and the temperature range required for the shipment. A *Products list* table provides item-level details such as product ID, description, quantity, packaging unit, and total weight.

This corresponds to the point in the choreography where the Coolstore has dispatched its goods, and Transporter 1 formally accepts the order, which then appears in its application. The transporter can now begin planning the shipment, including vehicle assignment, routing, and ensuring compliance with conditions like refrigeration.

Currently, Transporter 1 and Transporter 2 are virtual participants, implemented as automated software agents by ELET (Eletica S.r.l.) partner. They serve a dual purpose: first, to demonstrate how transport management systems can connect to the SDS via SDS Connectors (here Transporter 1 uses RMAS technology, while Transporter 2 uses GraphQL), and second, to simulate a generic transport management process that shows how decision-making contributes to the SCE’s performance.

Like all other actors in the SDS ecosystem, the transporters are not passive executors: they optimize logistics decisions with the explicit goal of maximizing GHG savings. They can evaluate alternative routes, consolidate shipments, or apply other operational improvements.

In the case shown, Transporter 1 manages the delivery of apples from Lisbon to Paris, reflecting the completion of the choreography message *msg_08*. Once the order is acquired, the system provides the transporter with all necessary details to initiate shipment planning, selecting the optimal route and vehicle type for efficiency and sustainability.

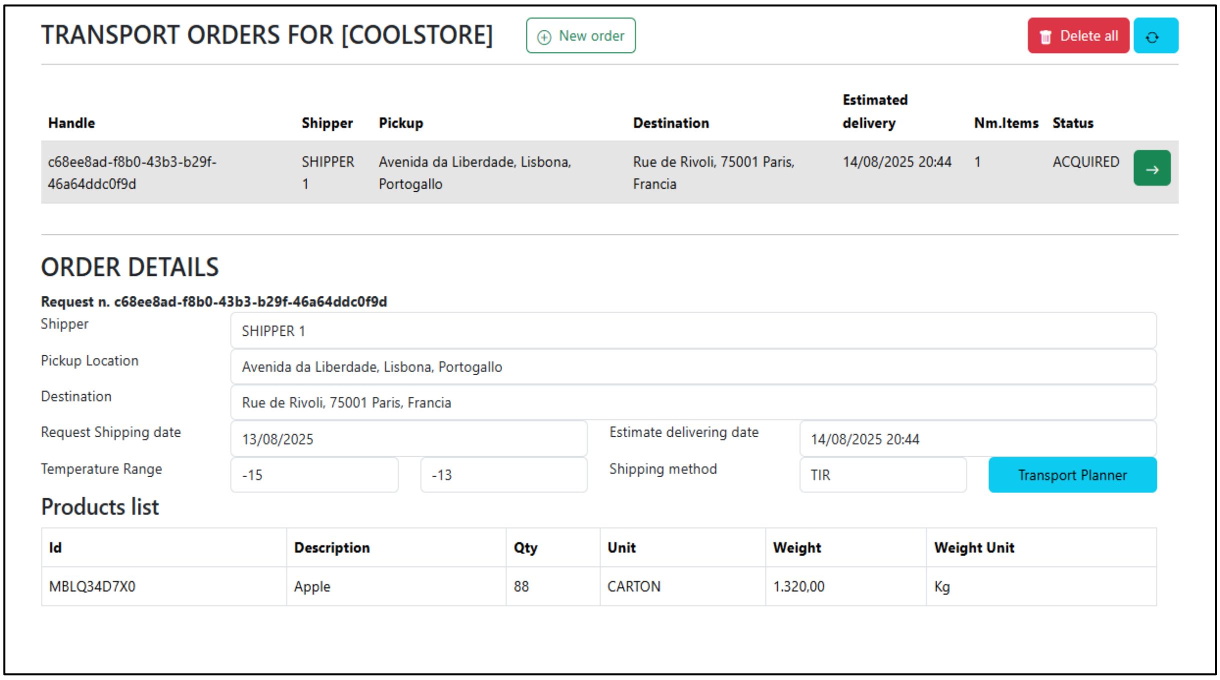

Figure 4.4.7: Entry point GUI of the Transporter 1 Participant.

The following screenshots illustrate the *Transport Planner interface*, which is the step where Transporter 1 (Coolstore) or Transporter 2 (Dairy Plant) begins planning the shipment after receiving the confirmed order.

The first panel (Figure 4.4.8) shows how the *Transport Planner* proposes different possible routes between the pickup and delivery points. In this simulation, the system leverages Google Directions API, which considers distance, duration, and potential traffic or interruptions. Additional data services, such as Copernicus API, could be layered to display environmental conditions (e.g., pollution levels, adverse weather forecasts), offering a more comprehensive decision basis.

Once a route is chosen, the Transport Planner highlights the emission estimates for that route and overlays them on the map. Here, the planner computes the carbon footprint by combining the distance with the chosen vehicle’s fuel and technical characteristics.

In this view, the user can also see a list of available vehicles. Each option is evaluated against the shipment requirements (weight, refrigeration, availability). Vehicles are flagged as *unsuitable* or *unavailable* when they cannot meet constraints—for instance, because of insufficient payload, lack of refrigeration, or already being allocated elsewhere. For each suitable vehicle, the system displays parameters such as fuel type, maximum load, and CO_2 emissions per kilometer and per freight unit, alongside refrigeration-related CO_2 if applicable. This ensures that the selection integrates both logistical feasibility and sustainability performance.

The second screen (Figure 4.4.9) exposes in detail the vehicle data sheet, making transparent all technical specifications (fuel type, refrigeration, weight capacity) and their contribution to emission estimates. This supports an informed selection that balances efficiency and sustainability.

The Transport Planner provides a decision-support environment where shipment planning is not only about logistics but also about minimizing environmental impact. The transporter participant, by interacting with this interface, actively contributes to the Smart Data System (SDS) objective of optimizing supply chains for lower GHG emissions while maintaining operational feasibility.

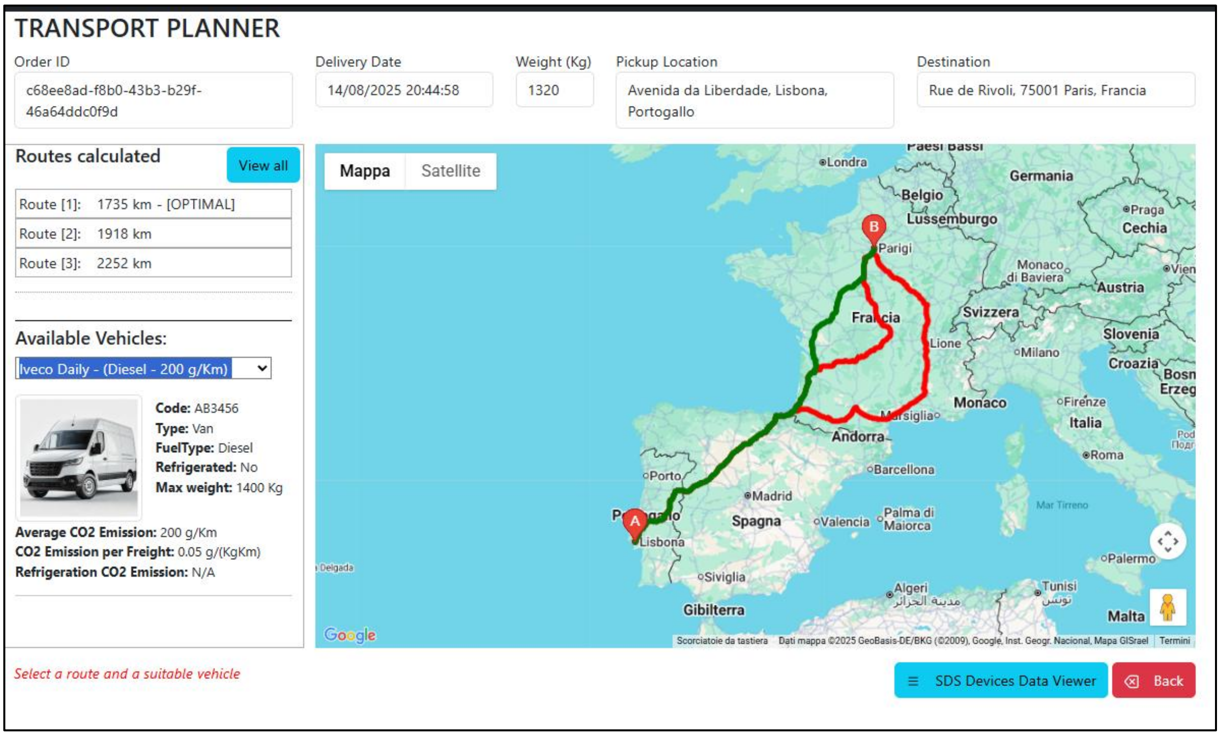

Figure 4.4.8: GUI of the Transporter 1 Participant for travel planning.

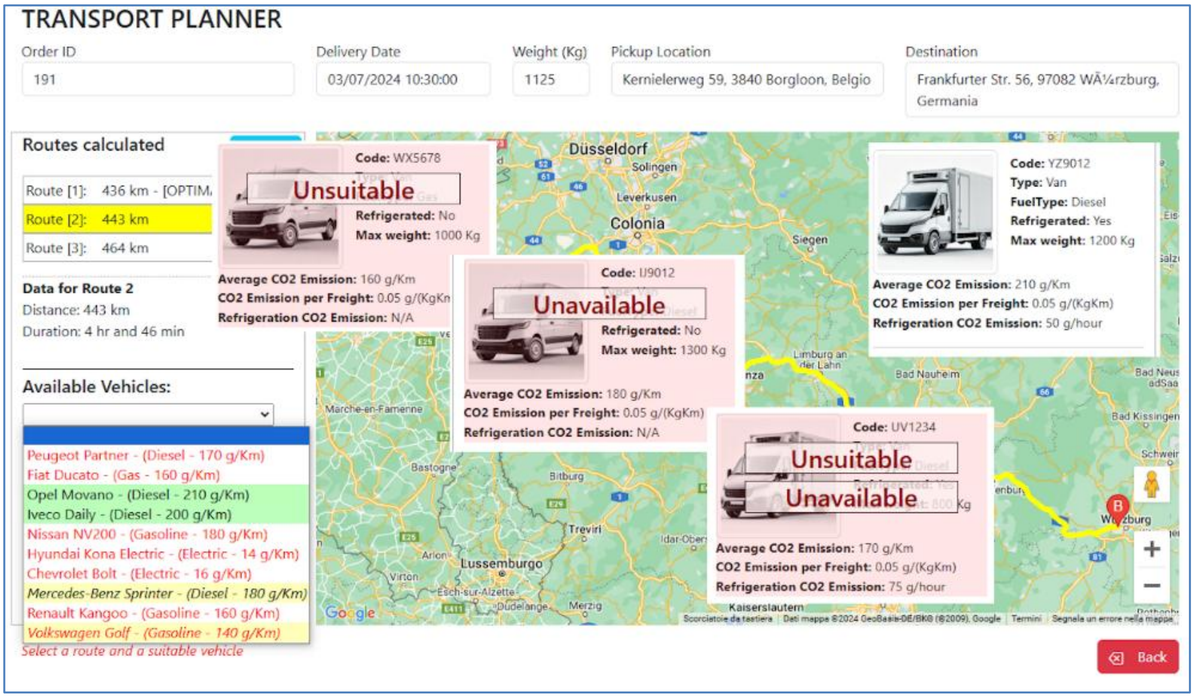

Figure 4.4.9: GUI of the Transporter 1 Participant for travel planning optimization and fleet management.

The Figure 4.4.10 represents the final stage of Transporter 1's workflow in the Coolstore scenario: the delivery confirmation and approval process.

At the top, the order summary is displayed with key identifiers: the shipper, pickup and destination locations, estimated delivery time, number of items, and the current status of the order, which in this case is "PROPOSAL APPROVED". This indicates that the delivery plan has already been validated by the Coolstore.

The central section shows detailed order data: pickup and delivery addresses, requested and actual delivery dates, shipping method, and temperature requirements. Here, the delivery is marked as "ACCEPTED", confirming that the transport plan is aligned with the operational and environmental requirements.

On the right, the Transport Proposal panel consolidates the transport execution plan:

- The vehicle type used (in this case an Iveco Daily, diesel-powered).
- The driver's profile.
- The distance of the selected route (1735 km).
- The estimated cost and duration of the delivery.
- The calculated CO_2 emissions for the journey, which are key to measuring the environmental impact of this transport decision.

At the bottom, the Delivery Confirmation button finalizes the operation. Once triggered, it validates the transport execution, logs the actual CO_2 performance, and provides feedback to the SDS system.

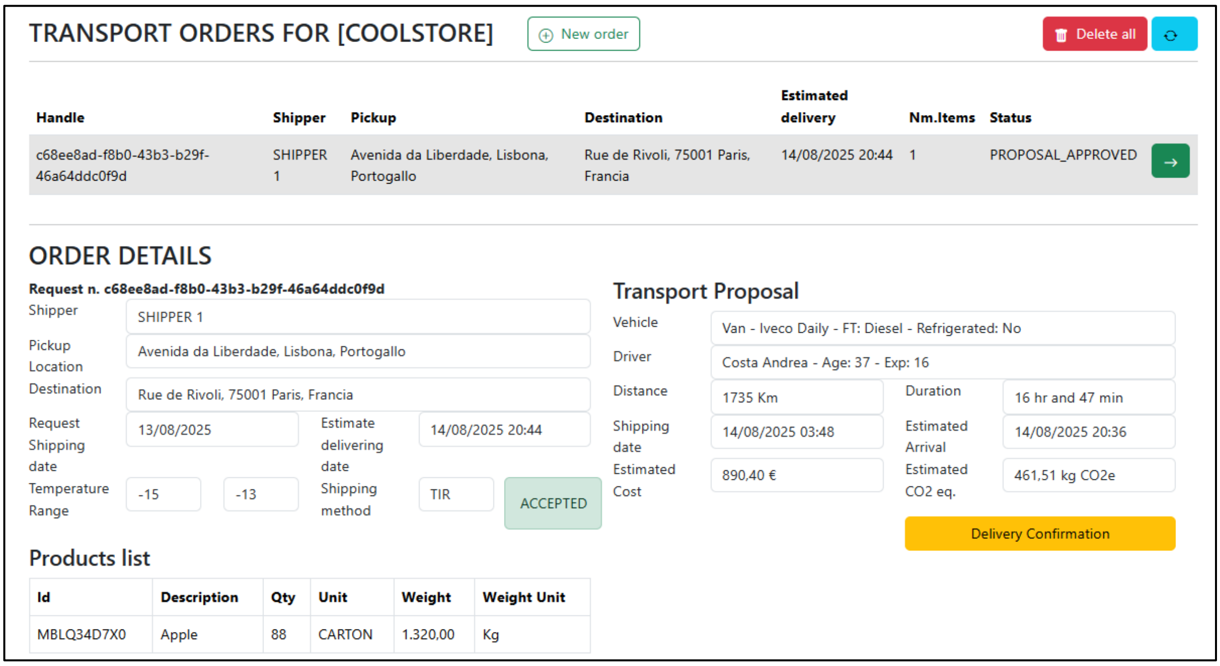

Figure 4.4.10: GUI of the Transporter 1 Participant in the approval stage.

After Transporter 1 completes the delivery creation and planning process—driven by decision-making criteria that prioritize GHG savings—the transporter proceeds to issue the delivery confirmation (Figure 4.4.11). In this stage, the final parameters of the transport, such as distance traveled, journey duration, and vehicle characteristics, are consolidated. These values allow the system to recalculate the effective CO_2 emissions, comparing them to the initial estimates provided at the proposal stage. This recalculation is essential for determining the actual GHG savings or overruns achieved.

In the simplified emulation, only the distance and duration are used for this computation. However, in more advanced use cases, additional factors such as the load factor (truck fill ratio) become decisive for optimizing operations and minimizing emissions.

Once the confirmation is submitted, the Coolstore application automation automatically sends structured feedback to Transporter 1 (via msg_10 in the choreography), reporting on the quality and outcome of the service. This closes the loop in the transport process, ensuring that both logistical and sustainability performance are captured and fed back into the Smart Data System (SDS) for continuous improvement.

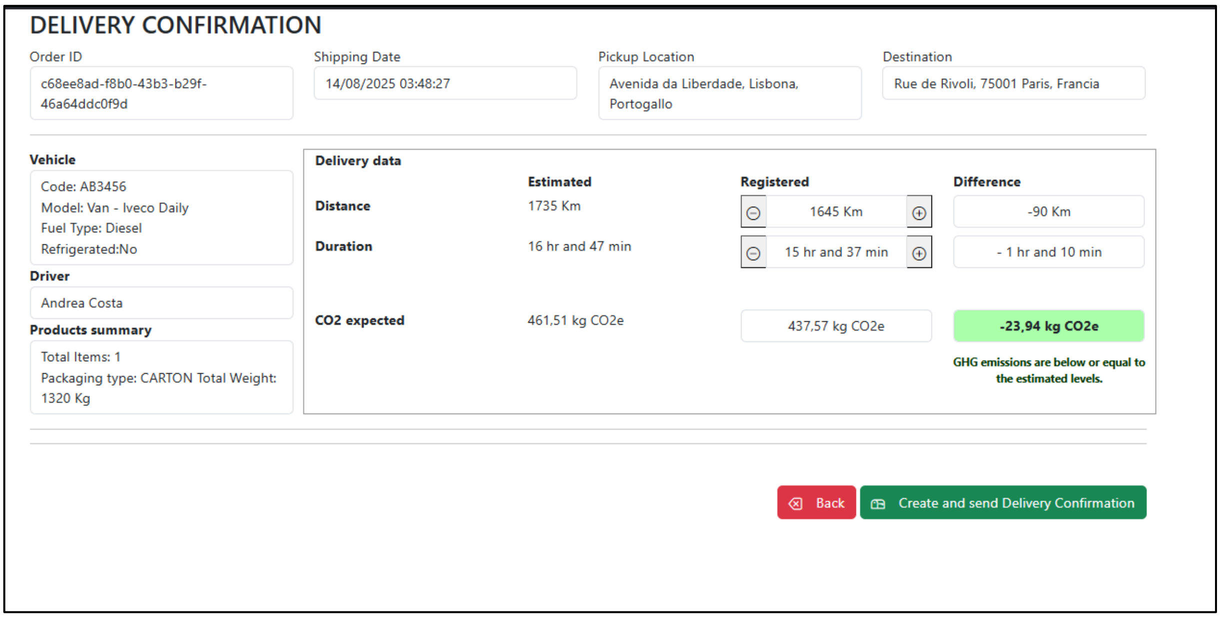

Figure 4.4.11: GUI of the Transporter 1 Participant in the confirmation stage.

The branch for the Transporter 2 and the Dairy Plant will be very similar here, and we avoid more details for the sake of some brevity. Nonetheless, we remark that each of the Participant can use and adapt a specialized release of their management interface: this is a relevant part of the SDS business.

Finally, the shared monitoring dashboard provides a consolidated view of the GHG savings performance achieved across successive iterations of the SCE. By collecting data from all the exchanged messages, the system generates both numerical and graphical representations of the KPIs. In particular, it highlights the contribution of each message type to the overall performance, making the monitoring both transparent and actionable.

As illustrated in the example in Figure 4.4.12, the black cumulative curve represents the primary KPI—total CO_2 savings—while the stacked bars reveal the differentiated impact of each activity/message. Over multiple sessions, this enables stakeholders to observe whether collaboration is driving progressive improvement or, in some cases, deterioration.

Short-term improvements usually result from local, private actions by individual participants within a single session. Medium-term effects emerge when lessons from past sessions are considered in subsequent negotiations. Finally, long-term benefits can be achieved by drawing upon the archived history of multiple SCEs, which allows for systemic learning and more effective strategic planning.

This continuous feedback loop demonstrates how the SDS framework turns operational collaboration into a measurable path of sustainability improvement, closing the cycle from negotiation to monitoring.

Figure 4.4.12: Cumulative graph across several Sessions of the GHG savings achieved in the SCE.

To conclude this section, we emphasize that these preliminary experiments have successfully consolidated the core use of the SDS framework infrastructure and its associated business use case. In particular, this case has demonstrated—also within the activities of Work Package 6—how the SDS platform can create digital entities for different supply chains and foster continuous and collaborative improvement among participants toward the achievement of common goals. Perhaps the most interesting aspect is that this outcome is achieved through a compromise between the private objectives of each participant and, at the same time, the overall objectives of the supply chain and the broader context in which it operates. This, in our opinion, constitutes some strenght and novelty of the SDS approach.

5 BUSINESS USE CASE 2: BUYING FROM THE MARKETPLACE

We recall in Figure 5.1, for convenience of the reader, the business use case 2 picture that is discussed in more details in section 3.

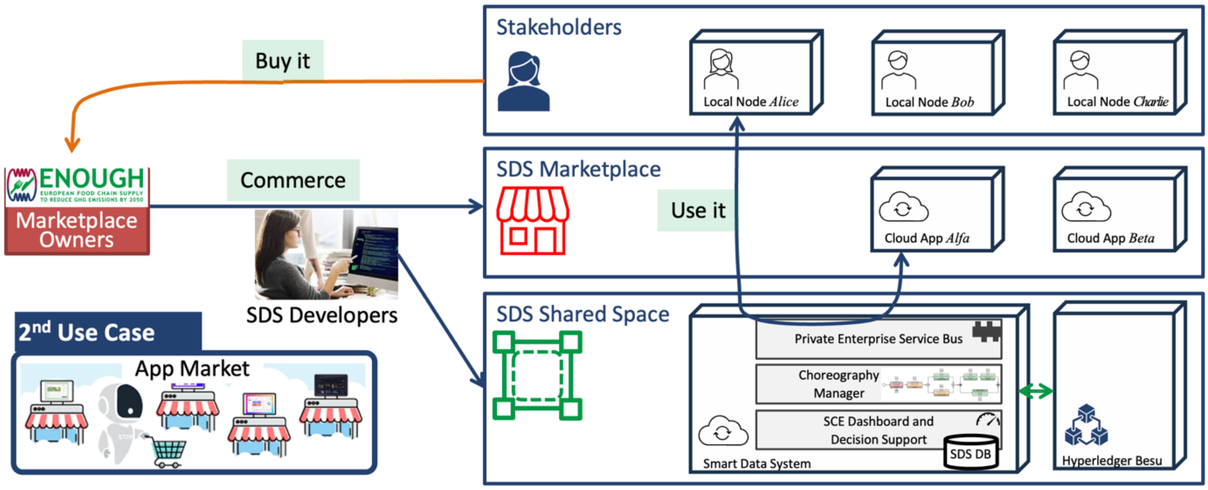

Figure 5.1: Business case 2 – Clients access additional third-party applications and services via the SDS Marketplace (see section 3 for more details).

Clients have the possibility to buy and use a wide range of applications and services available in the SDS Marketplace. These offerings may include advanced dashboards for monitoring and visualization, specialized data repositories or knowledge sources, decision-support systems, artificial intelligence tools for predictive analytics, and other digital services that enhance supply chain operations. Such applications can be provided by third-party vendors, original equipment manufacturers (OEMs), or marketplace owners themselves, thereby enriching the ecosystem with diverse solutions.

To ensure seamless integration, however, every external application must undergo a process of adaptation to the SDS system. This adaptation is generally minor and is typically supported or entirely carried out by the SDS system developers, who provide the necessary connectors, interfaces, and compliance checks. This guarantees interoperability within the SDS infrastructure while lowering the entry barriers for external providers. In this way, the marketplace extends the functional value of the SDS platform, giving participants access to a continuously expanding catalog of digital tools while maintaining coherence, security, and transparency across the system.

To demonstrate this, in this section it is reported a new study that the ELET (Eletica S.r.l.) partner is undergoing in order to demonstrate the capability of the SDS to receive a remote autonomous automation unit as a special kind of Participant or, alternatively, as part of the tools that a Participant (or Actor) can have available. This configuration is mainly used in DCS (Distributed Control System) technology typical of industrial automation — also called today the IIoT (Industrial Internet of Things) context.

We start by recalling briefly some foundational concepts and definitions.

A major concept is the *SDS Connector* system, which is a software and hardware component installed in the Actor's premises and systems. The *SDS Connector* implements an *RMAS transceiver*, which is the part of the software that creates the SSH tunnel, the sockets, and that handles the bi-directional replication of the *Local DB* to synchronize and connect with the *SDS DB*. The *Local DB* is implemented with an SQLite embedded DBMS. There are 4 different types of *SDS Connector* as follows:

- **Type 1**. The *RMAS Transceiver* written in Python language. This solution is the simplest for the generic users. It is cross-platform (in operating systems sense) as the Python portability is leveraged.
- **Type 2**. The *RMAS Transceiver* written in C language. This solution has the advantage of minimum impact on resources, and the code can remain closed to the user. A version has to be compiled for every kind of operating system, which is a minor drawback. In case of SQLIte used as Local DB, the C language is ideal for the performance and the features that can be exploited.
- **Type 3**. This is like Type 2, but here special embedded and small libraries and C implementation are used. It is the version with the smallest footprint suited to IIoT devices like sensors and actuators.
- **Type 4**. This is the alternative obtained with GraphQL technology. In this case no Local DB is strictly necessary, although suggested. This is the case in which the stakeholder's application has to be adapted by themselves in order to connect to the SDS. The only specification that is shared is the data model. Moreover, a corresponding configuration must be obtained on the SDS side by configuring the GraphQL adapter in the SDS node.

The study reported in this section concerns the Type 3 kind of connector (Figure 5.2 and 5.3).

The Type 3 connector application is an embedded system that is designed here to interface with refrigeration control systems via RS485 communication. The system currently utilizes an STM32L4R9ZIJX microcontroller to read and store data from the refrigeration system into an SQLite database. This stored data is then shared with the SDS via a GSM connection, as the device is not connected to a local LAN or Wi-Fi network. Additionally, the node is equipped with a GPS module to track the device's location. The application is in an experimental phase but is designed to be adaptable to the SDS, regardless of the nature of the data collected. This development is part of ELET's ongoing efforts within the project to enhance data management and monitoring solutions.

The increasing need for efficient data management in refrigeration systems, particularly in transport vehicles like refrigerated vans, has led to the development of a specific Type 3 connector application. The Type 3 connector leverages an STM32 microcontroller to continuously monitor and record key parameters from the refrigeration control unit. The application enables reliable data storage and later retrieval for analysis and reporting. Furthermore, it is designed with adaptability in mind, ensuring seamless integration with future developments in the overall system architecture. Importantly, while the current implementation utilizes the STM32L4R9ZIJX, the application's architecture is platform-agnostic, allowing it to be easily ported to other microcontroller families. Thus, this study has been only a first instance of a wide range of Type 3 connectors that depend on the specific requirements.

The Type 3 Actor application is built around the STM32L4R9ZIJX microcontroller, which plays a central role in managing communication, data handling, and GPS tracking.

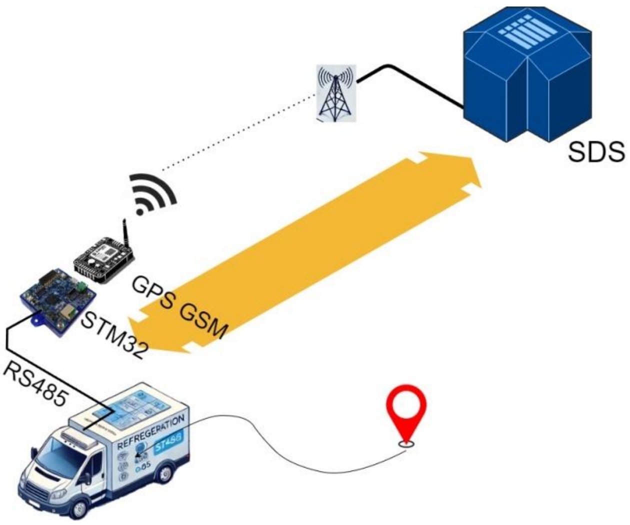

Figure 5.2: The specific Type 3 Actor here studied.

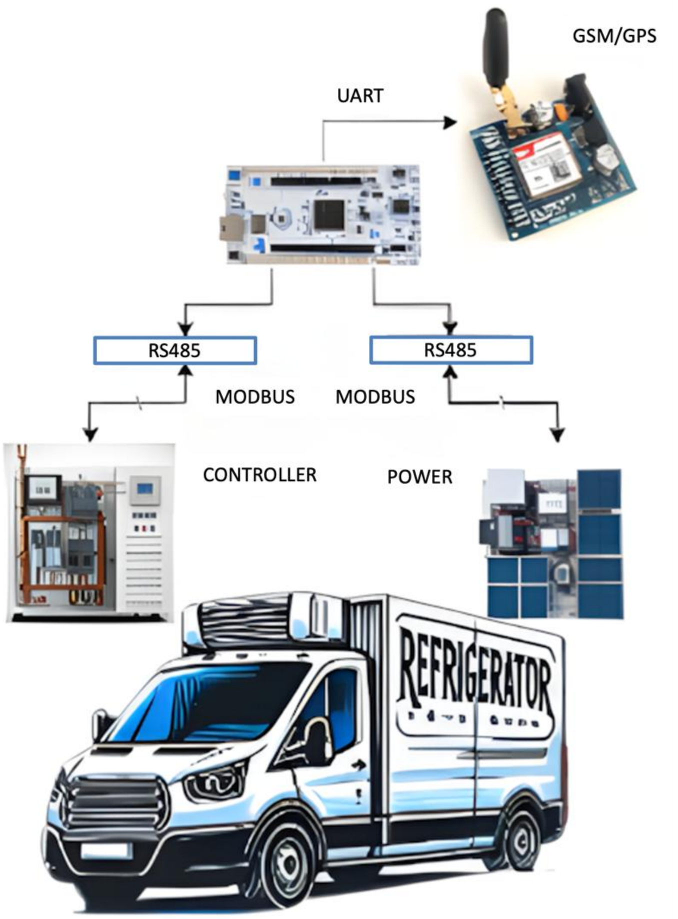

Figure 5.3: Picture of the hardware of the type 3 Node prototype.

The architecture comprises the following components:

- **STM32L4R9ZIJX Microcontroller:** The STM32L4R9ZIJX is the core of the application, responsible for interfacing with the refrigeration control system via RS485 protocol. Known for its high performance, low power consumption, and extensive peripheral support, the STM32L4R9ZIJX ensures reliable data acquisition and processing. However, the application is designed to be platform-independent, allowing for easy migration to other microcontroller units (MCUs) as needed.
- **RS485 Communication:** RS485 is a widely used serial communication standard in industrial environments, appreciated for its robustness and ability to handle long-distance data transmission. The Type 3 Actor uses this protocol to receive real-time data from the refrigeration system.
- **SQLite Database:** The microcontroller stores the received data in an SQLite Database. SQLite is chosen for its lightweight, serverless, and self-contained nature, making it ideal for embedded systems. The database ensures that all recorded data is easily accessible for future use, either for local processing or transmission to other systems.
- **GSM Connectivity:** Since the device is not connected to a local LAN or Wi-Fi network, data sharing with the SDS occurs via a GSM connection. This ensures that the node can communicate with the SDS regardless of its location, providing flexibility in deployment.
- **GPS Module:** The node is also equipped with a GPS module, allowing it to track and log its location. This feature is particularly useful for applications where monitoring the geographical position of the refrigeration unit is crucial, such as in fleet management.
- **Smart Data System (SDS) Integration:** The Type 3 Actor is designed to share the recorded data with the Smart Data System (SDS). Although the SDS is still under development, the Type 3 Actor's design ensures that it can adapt to the SDS regardless of the specific data types it will eventually handle. This adaptability is a key feature, ensuring that the Type 3 Actor remains relevant and useful as the broader system evolves.

The Type 3 Actor application follows a straightforward workflow:

1. **Data Acquisition:** The STM32L4R9ZIJX microcontroller continuously reads data from the refrigeration control system via the RS485 interface.
2. **Data Storage:** The acquired data is immediately stored in the onboard SQLite database. This ensures that data integrity is maintained even if communication with the SDS is temporarily unavailable.
3. **GPS Tracking:** The GPS module continuously tracks the device's location, logging this data alongside the refrigeration system data in the SQLite database.
4. **Data Sharing via GSM:** The microcontroller shares the stored data with the SDS via a GSM connection. This process is designed to be flexible, allowing for various data types to be transmitted without requiring changes to the microcontroller's firmware.

Concerning the business use case 2, any service or application that does not belong to the core management of the SCE Choreography is to be considered something that can be conveyed through the SDS Marketplace. Connecting this to the SCE Choreography of the previous section, this means that a Participant like the Transporter 2, might want, for their decision making, have a tool that acquires information for them (privately) from the fleet of trucks in order to make better assessments.

In particular, here we have tested this architectural provision by means of the integration of the first decision making application that is rendered available and accessible to one or more of the Participants. In this case, and for the first Demo 1 in WP6, we make the case that it is a service purchased only by one Participant.

This application is provided as a Node in the SDS infrastructure, so something more general than Actors and Participants, which concerns mostly a component connected to the SDS Infrastructure rather than a specification of its role. This means that a suitable SDS Connector has been deployed in the machine where the application is run. The SDS Infrastructure is the made ready to accept and transmit the values of all the variables that is needed by the application to acquire data.

The updated system can seamlessly integrate into an SDS choreography by leveraging the capabilities of the connector developed for Type 3 Nodes. This connector enables the system to act as an autonomous participant within the SDS framework, facilitating efficient communication and data exchange.

In practice, the Participant wants to access this feature by a widget on their SCE management GUI as explained in the previous section during planning. The button is shown in Figure 5.4.

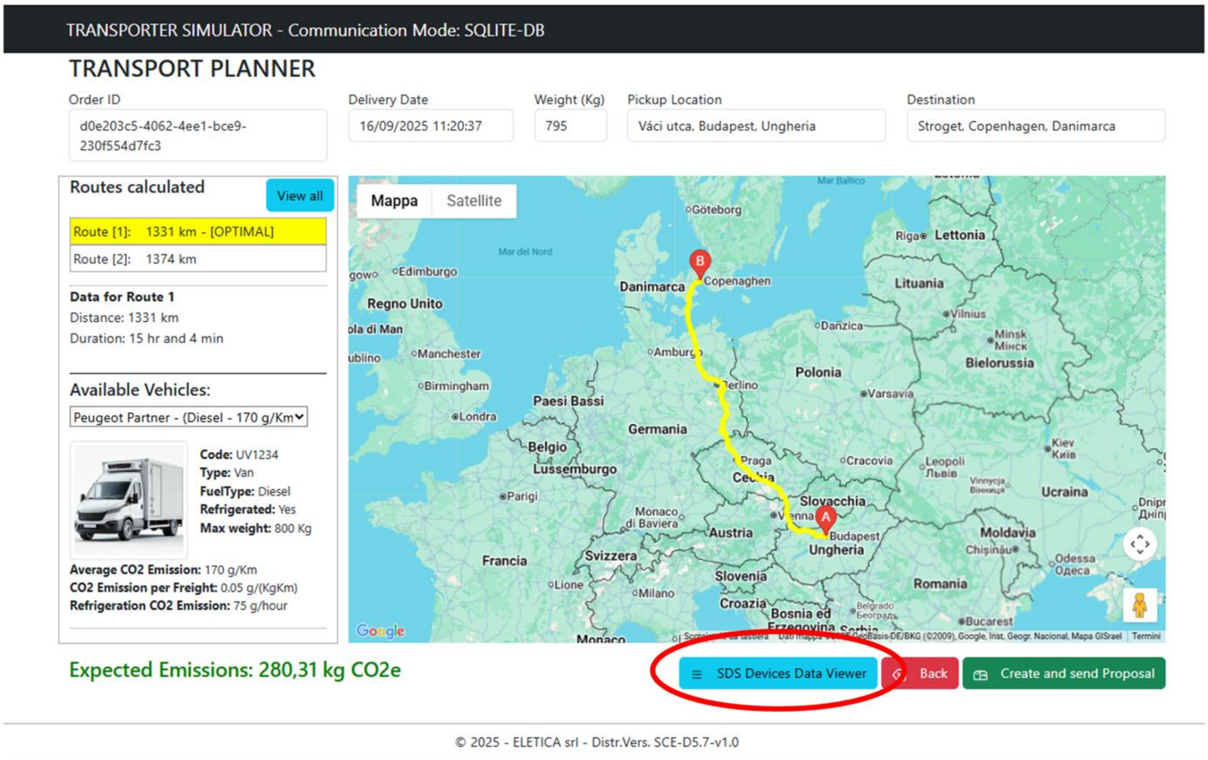

Figure 5.4: Highlighting of the button that calls the tool for the Transporter 2, during decision making for planning.

This widget will simply bring to the Web page displayed in Figure 5.5, which provides access to this tool.

Figure 5.5: The page of the Devices Data Viewer for SDS.

The SDS Devices Data Viewer serves as a demonstration of how the SDS can operate as a comprehensive platform for collecting, transmitting, and visualizing real-time operational data. It illustrates in practice the capability of the SDS to function not only as a coordination and negotiation tool between supply chain actors but also as a monitoring and decision-support environment where operational parameters are continuously captured, processed, and made accessible.

In this example, the focus is placed on a refrigerated van equipped with advanced monitoring systems that track both the refrigeration unit and the vehicle's battery pack. These measurements are critical for ensuring that perishable goods are maintained within the required temperature range, while also keeping track of the energy efficiency and environmental footprint of the transport operation. Data such as cooling unit performance, energy consumption, battery charge cycles, and temperature stability can be streamed into the SDS platform, where they are aggregated and contextualized alongside other supply chain information.

For a transporter participant, acquiring a tool like the SDS Devices Data Viewer provides a clear competitive advantage. It allows them to make data-driven decisions regarding route optimization, vehicle allocation, energy efficiency, and compliance with sustainability targets. In addition, the transparency of the collected data facilitates reporting to clients and stakeholders, while also enabling integration with broader supply chain performance indicators. Ultimately, this demonstration highlights the potential of the SDS platform to act as a real-time digital twin of transport operations, supporting both immediate decision-making and long-term strategic improvements in sustainability.

The user interface in Figure 5.6, represents one view of the Devices Data Viewer for SDS.

The panel lists the most recent 25 readings transmitted by a monitored device—in this case, a refrigerated van equipped with sensors for the refrigeration unit, the battery pack, and a GPS system for geolocation. For each record, the interface provides:

- Sent date (sampled data) → the timestamp when the measurement was taken.
- Timestamp → the precise time of acquisition in synchronized system format.
- GPS coordinates → latitude, longitude, and altitude for accurate tracking of vehicle position.

- Values → the raw telemetry package including detailed parameters from refrigeration monitoring, energy status, and potentially other environmental variables.

The presence of long hexadecimal or encoded strings in the Values field demonstrates the raw, unfiltered device communication, ensuring tamper-proof traceability before further decoding and visualization. This emphasizes the SDS principle of secure data transfer, where information flows through SDS Connectors and nodes with end-to-end integrity.

The interface highlights how SDS can serve as a backbone for transparent, auditable logistics, ensuring that operational data—such as temperature stability, battery charge, or GPS traces—can be captured and later validated. For a transporter, this enables both real-time operational awareness and compliance reporting, while for the SDS ecosystem it proves the integration of IoT telemetry into broader supply chain management.

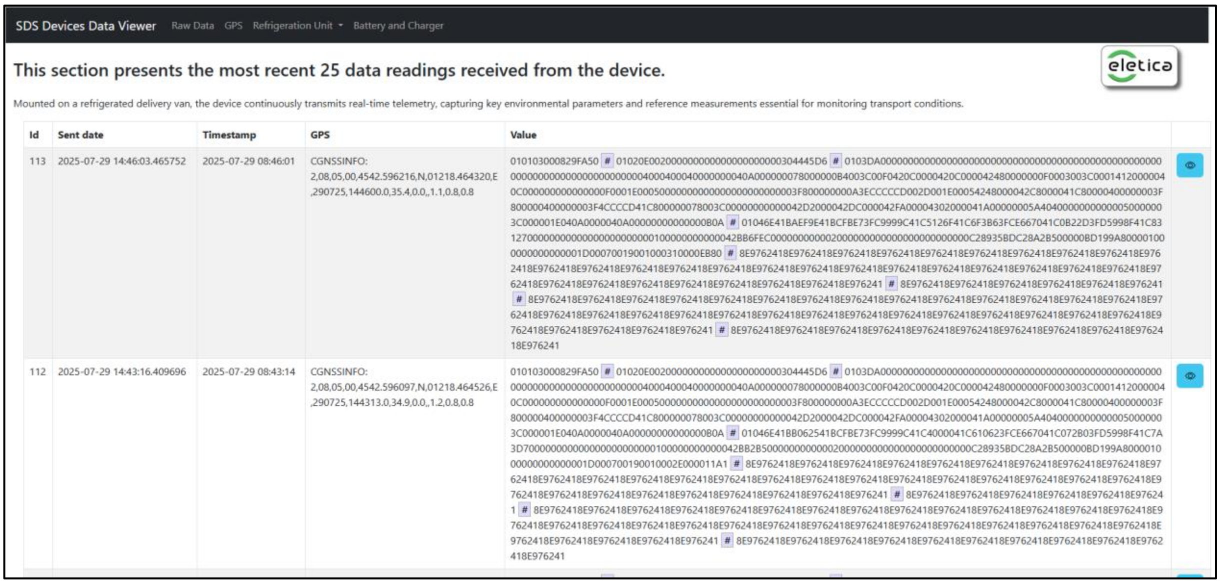

Figure 5.6: The first raw data view page of the Devices Data Viewer for SDS.

The view in Figure 5.7 user interface represents the *Decoded Variables Overview* of the *SDS Devices Data Viewer*, where raw telemetry transmitted by a refrigerated van is decoded into meaningful operational parameters. The goal is to provide a clear, structured, and auditable view of the vehicle's critical subsystems for monitoring and decision-making.

- Refrigeration unit variables (left panel): This section reports the working conditions of the cold-chain equipment. It includes duty cycles (when compressors are active), key temperature readings within the storage compartment, setpoints configured by the operator, and defrost cycle events. It also highlights anomalies, such as deviations from expected cooling performance, which could compromise product safety.
- Battery and charger variables (right panel): This area details the electric system, listing the state of charge (SoC), battery voltage, current flows, and depth of discharge. Additional entries show the health index, which indicates long-term battery performance and degradation. These values are essential for sustainable operations since they directly influence both reliability and CO_2 efficiency of electric or hybrid vehicles.

- GPS variables (in metadata at the top): Latitude, longitude, and timestamp provide real-time geolocation tracking of the van. This allows correlation of refrigeration and battery events with transport routes, environmental conditions, and delivery schedules.

By decoding these variables, the SDS platform transforms encrypted IoT telemetry into operationally relevant insights. This enables participants in the supply chain to monitor compliance, optimize performance, and ensure transparency. The decoded data also feeds into higher-level KPIs (like CO_2 savings dashboards), supporting both immediate corrective actions and long-term sustainability assessments.

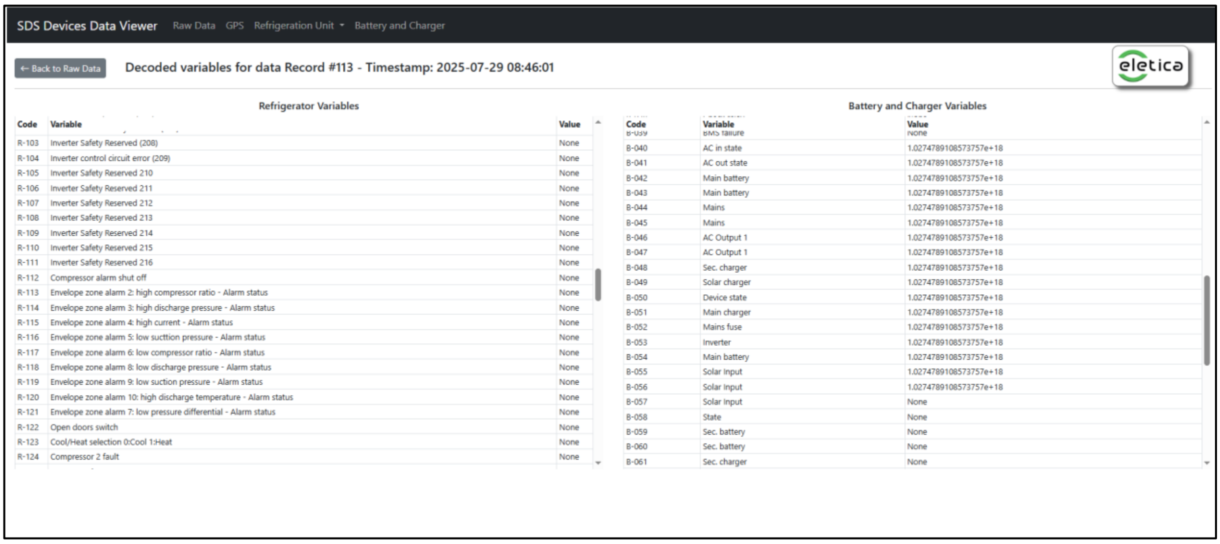

Figure 5.7: Decoded Variables Overview of the SDS Devices Data Viewer.

Figure 5.8 illustrates the *GPS Geolocation Features* of the *SDS Devices Data Viewer*, providing a spatial and temporal view of transport operations. It is designed to give stakeholders full transparency into where a vehicle is, where it has been, and how its movement correlates with other operational data.

- Real-time vehicle position on interactive map (right side): The interface shows the current location of the refrigerated van as a marker on an interactive map. This visualization updates continuously, allowing operators to follow the vehicle live as it moves along its delivery route.
- Supports route analysis and event correlation: The accumulated GPS traces provide input for route optimization and emissions analysis, since distances traveled and vehicle operating conditions directly affect CO_2 calculations. When combined with environmental KPIs, the map helps visualize how operational decisions—such as detours or idle periods—impact sustainability performance.

This interface demonstrates how the SDS integrates location intelligence with operational telemetry, turning raw geospatial data into actionable insights for efficiency, safety, and GHG savings.

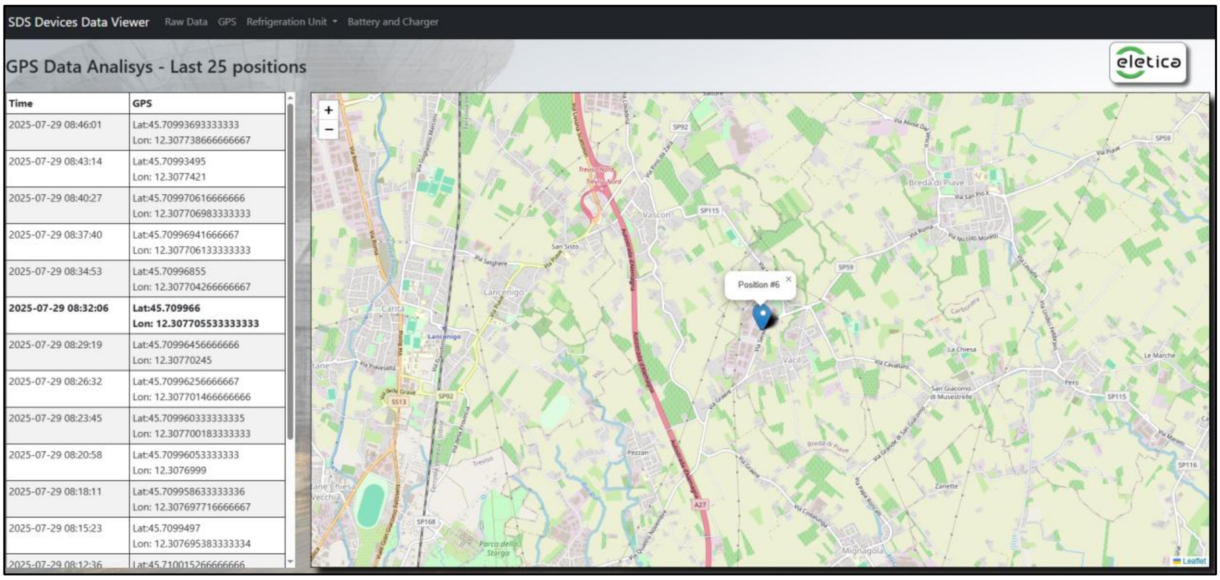

Figure 5.8: GPS Geolocation Features of the SDS Devices Data Viewer.

Figure 5.9 shows sample graphs from the *Refrigeration Unit Data Analysis* module within the SDS *Devices Data Viewer*. The purpose of these charts is to provide continuous monitoring of refrigeration system performance, a critical element for transport of perishable goods where product quality and energy efficiency are tightly coupled.

- Cooling Performance Overview (top chart):

This time-series view displays multiple sensor readings: cold room temperature, discharge and suction temperatures, suction pressure, and external temperature. By comparing setpoints with actual sensor values, operators can detect deviations, inefficiencies, or malfunctioning cooling cycles. Peaks and troughs indicate cooling cycles and possible defrost operations.

- Compressor Demand and Rotor Speed (middle chart):

This panel focuses on the compressor's workload and motor performance. It shows the compressor request signal against the actual rotor speed, highlighting whether the system is meeting demand efficiently. A mismatch or lag could reveal mechanical inefficiencies, abnormal load, or energy waste.

- Superheat and Valve Operation (bottom chart):

Here the focus is on thermodynamic balance: the superheat value (a critical safety and efficiency parameter) is plotted against the expansion valve opening. Stable superheat control ensures both energy efficiency and protection of the compressor. Variations can reveal leaks, refrigerant charge issues, or valve anomalies.

Together, these three dashboards allow stakeholders to:

- Monitor setpoint vs. actual performance in real time.
- Detect defrost cycles and other recurring behaviors.
- Identify early anomalies before they affect product safety.
- Quantify how refrigeration contributes to GHG emissions, since compressor duty cycles and inefficiencies directly drive energy consumption.

This visualization demonstrates how the SDS framework can integrate detailed refrigeration telemetry into the broader supply chain monitoring system, enabling evidence-based optimization of transport operations.

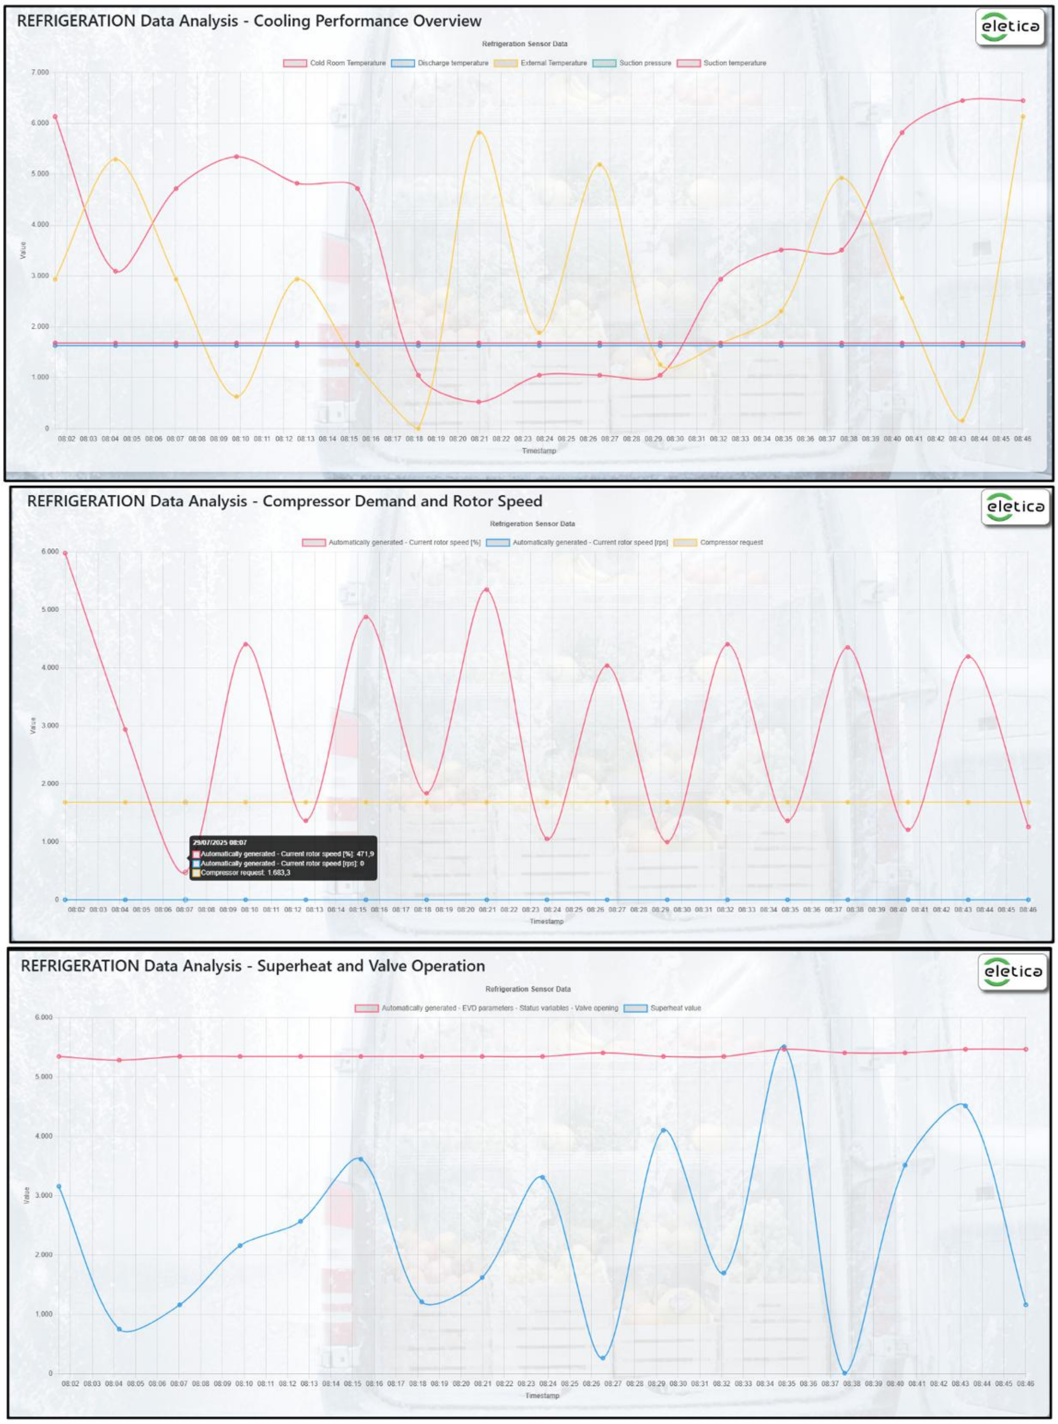

Figure 5.9: Refrigeration Unit Data Analysis.

Lastly, Figure 5.10 presents the *Battery and Charger Monitoring* dashboard within the *SDS Devices Data Viewer*. Its purpose is to provide continuous supervision of the energy storage system in a refrigerated transport vehicle, ensuring reliability, efficiency, and sustainability of operations. The view is organized into four time-series graphs:

- State of Charge (SOC) (top-left): Displays the battery's remaining capacity as a percentage over time. Stable SOC trends indicate balanced consumption and charging cycles, while drops or irregularities could signal excessive discharge or abnormal loads.
- Voltage (top-right): Tracks the electrical potential across the battery terminals. Consistent voltage within nominal limits reflects healthy operation, while sudden dips or spikes may reveal faults, cell imbalances, or charging irregularities.
- Current (bottom-left): Shows the flow of electrical current (both discharge and recharge). Monitoring current helps identify energy demand patterns, peak loads, and charging intensity. Prolonged high current could indicate stress on the system.
- Temperature (bottom-right): Captures the thermal conditions of the battery pack. Stable temperature ensures safety and efficiency, whereas abnormal rises may point to cooling issues, overcharging, or risks of thermal runaway.

In combination, these metrics support:

- Depth of discharge analysis, to prolong battery life.
- Identification of charging cycles and efficiency.
- Monitoring of a health index (linking SOC, current, and temperature) for predictive maintenance.
- Assessment of how battery performance contributes to the overall GHG footprint of the transport operation.

By integrating these parameters into the SDS framework, operators can make informed decisions about vehicle scheduling, charging policies, and equipment maintenance, directly improving sustainability and operational reliability.

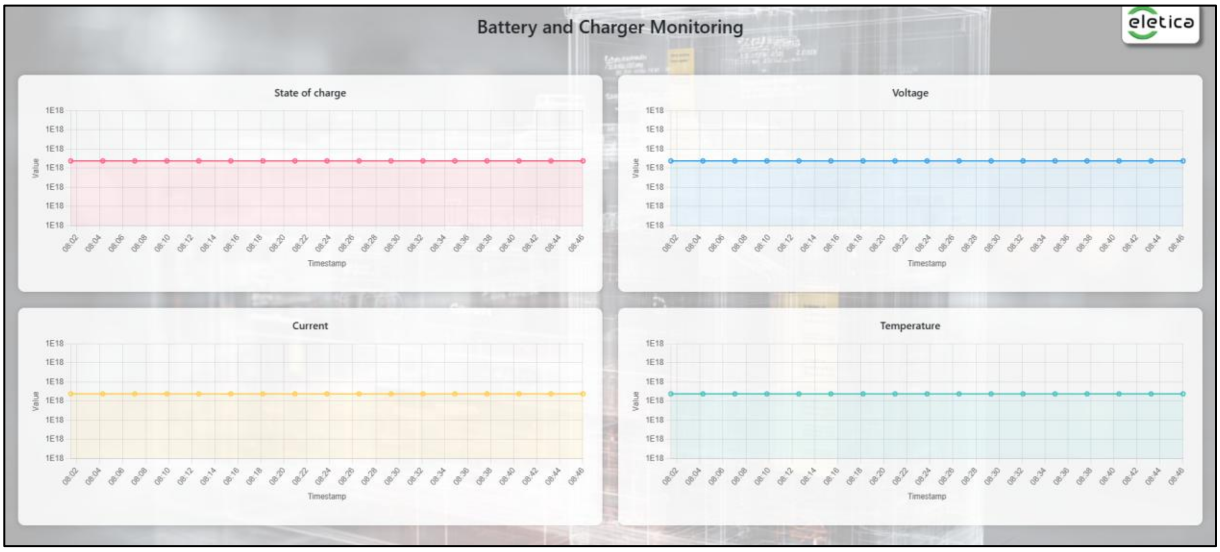

Figure 5.10: Battery & Charger – Sample Graphs

6 BUSINESS USE CASE 3: SELLING TO THE MARKETPLACE

Differently from the previous business use case, here the provision of additional tools in the SDS Marketplace is pushed by third parties (with respect to the SDS Business Unit) that make available their products to the SDS clients. We recall the business case scheme in Figure 6.1. In this case, the SDS Marketplace operates not only as a hub where clients can purchase or access existing applications, but also as a platform that allows external stakeholders—such as technology providers, software developers, or service companies—to introduce their own solutions and make them available to the wider ecosystem. These solutions may range from specialized data analytics modules, advanced decision-support tools, optimization engines for supply chains, blockchain-based certification services, or even domain-specific applications designed for sustainability monitoring, logistics, or energy management.

To enable this integration, the SDS Developers play a key role by supplying the SDS Connector suite, which ensures interoperability and secure communication with the SDS Infrastructure. This connector standardizes data formats and interaction protocols, so that even heterogeneous applications can seamlessly interact with other components of the system. In addition to the technical connector, developers may also provide a dedicated Web access point and user interface (as in the previous business use cases) within the SDS Shared Space. This allows the newly integrated solution to be visible, accessible, and usable by clients in the same environment where they already manage their SCEs and supply chain processes.

The inclusion of third-party applications under this framework also requires minimal adaptation steps to align with SDS governance rules, security standards, and message-exchange protocols. Once validated, these applications enrich the ecosystem, creating a broader marketplace of interoperable tools that foster innovation, lower entry barriers for specialized providers, and offer clients a wide choice of services tailored to their operational and sustainability needs.

In this way, the SDS Marketplace evolves into a collaborative digital ecosystem, where value is generated both by the platform owners and by external contributors, and where participants benefit from a continuously expanding portfolio of solutions that enhance efficiency, sustainability, and competitiveness.

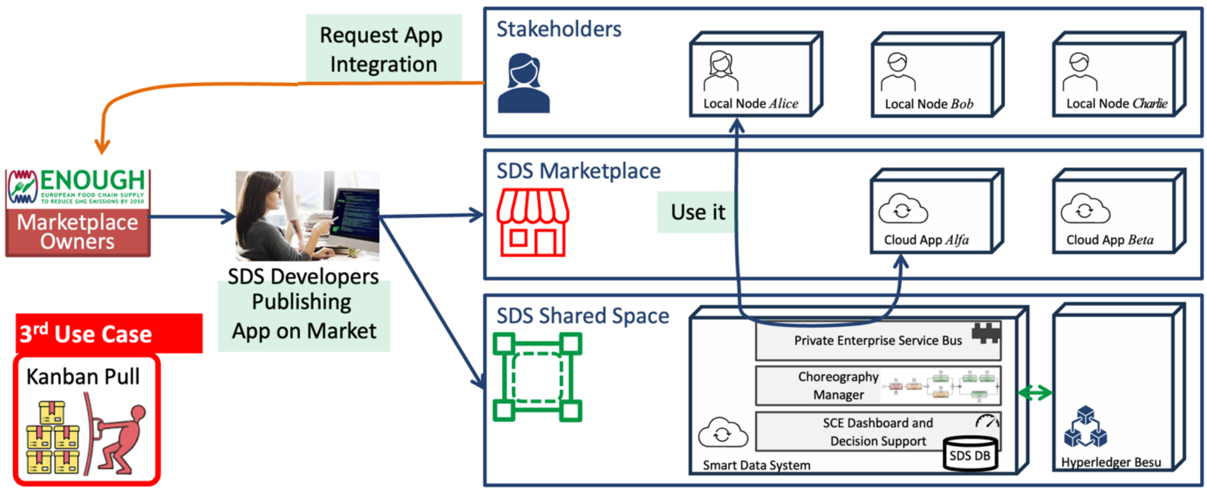

Figure 6.1: Business case 3 – Stakeholders integrate and sell their solutions in the SDS Marketplace through SDS connectors (see section 3 for more details).

To demonstrate a practical instance of this business use case, an experiment was carried out in collaboration with a group of advanced cybersecurity students of UNIVPM, who were tasked with assuming the role of a software provider. Their assignment was to design and implement an application that could be layered on top of the SDS infrastructure and subsequently offered in the SDS Marketplace as a value-added service. This exercise aimed to showcase the openness and extensibility of the SDS framework, and to highlight how external actors, even those not originally part of the consortium, can contribute innovative solutions that enrich the ecosystem.

The application developed in this context is particularly noteworthy, as it introduces the concept of a **private blockchain currency** and an associated **credits market** within the SCE (Supply Chain Entity). The central idea is to leverage tokenized incentives as a mechanism to nudge and reward virtuous behavior among participants.

For example, companies that adopt more sustainable practices or improve their operational efficiency could be rewarded with digital credits, which in turn could be exchanged, accumulated, or used to signal reputational gains within the marketplace. This represents a concrete instantiation of how behavioral economics and distributed ledger technologies can be combined within the SDS to foster both collaboration and competition toward shared sustainability goals.

Perhaps the most important feature demonstrated through this experiment is the independence of third-party applications from the SDS core infrastructure. Once connected via the SDS Connector and given access to the SDS database schema, applications can operate in parallel, without interfering with the underlying orchestration and data exchange mechanisms that drive the SCE. The SDS data layer, structured around a publish/subscribe paradigm, ensures that all relevant data flows are made available in a standardized and secure way. This allows external applications to subscribe to the streams they require, process them according to their own logic, and provide complementary services without disrupting the stability or scalability of the core system.

This demonstration clearly showed that the SDS is not merely a closed platform but a **flexible and extensible digital ecosystem**, capable of hosting diverse and innovative applications that may originate from academia, industry, or independent developers. Such openness is a critical enabler for long-term growth, as it allows the marketplace to evolve dynamically and to respond to the emerging needs of supply chains and sustainability-oriented enterprises.

One major task in the ENOUGH project's Work Package 5 (WP5) was to explore how blockchain technologies can be harnessed to support sustainability within food supply chains. The blockchain layer has been introduced as one of the core mechanisms of the SDS Infrastructure. Nevertheless, the TRL level of this research remains lower than the other parts of the SDS Infrastructure. Thus, we decided to keep it as a research laboratory part to be developed with suitable follow ups, while not blocking the SDS business which is already viable for the market otherwise. The research conducted has demonstrated that many parts of the blockchain framework is still not mature and deserves more time to become part of the core SDS. This research is already available (or soon when under press at the moment of this writing)[26,27,28,29,30,31,32,33,34].

With this demonstration conducted with the SCE_D5.3 instance (see section 4), we can instead demonstrate a mature application, already viable for the SDS Marketplace, with minimum fixes, although here the developers were aware only of the concerned parts of the SDS DB schema, and not other technicalities known by the SDS Business Unit, to simulate perfectly what a third party would be allowed to know to sell their product in the SDS Marketplace.

The central idea was to tokenize sustainability, creating a transparent and traceable mechanism for tracking greenhouse gas (GHG) reductions, compensatory actions, and company reputations. By assigning CO_2 tokens and environmental reputation scores, the system incentivizes companies to adopt greener practices and establishe accountability.

Two main dashboards have been developed to manage and visualize sustainability data. The Global Dashboard provides a real-time overview of all companies participating in the system. It highlights key metrics, comparative analytics, overall GHG reduction trends, and collective sustainability performance, serving stakeholders who need to monitor the broader impact. The Company Dashboard, by contrast, is a personalized tool where each company can track its own performance. It displays direct emissions, CO_2 credits earned, compensatory actions, and reputation.

---

[26] Spegni, F., Fratini, L., Pirani, M., & Spalazzi, L. (2023, March). ChoEn: A smart contract based choreography enforcer. In *2023 IEEE international conference on pervasive computing and communications workshops and other affiliated events (PERCOM workshops)* (pp. 86-91). IEEE.

[27] Cacopardo, A., Cucchiarelli, A., & Spalazzi, L. (2023). A Soulbound Token-based Reputation System in Sustainable Supply Chains. In *INTERNATIONAL CONFERENCE ON EMBEDDED WIRELESS SYSTEMS AND NETWORKS (EWSN)*.... Junction Publishing.

[28] Pirani, M, Cucchiarelli, A., Spalazzi, L. (2024), "Blockchain and the unexpected: taming complexity in sustainable supply chains", at WOSC 2024 – World Organisation of Systems and Cybernetics, WOSC 19th Congress 2024, Lady Margaret Hall, Oxford, UK, 11-13. September 2024.

[29] Pirani, M., Cucchiarelli, A., & Spalazzi, L. (2024, November). A Role of RMAS, Blockchain, and Zero-Knowledge Proof in Sustainable Supply Chains. In *IECON 2024-50th Annual Conference of the IEEE Industrial Electronics Society* (pp. 1-4). IEEE.

[30] Pirani, M., Cucchiarelli, A., Naeem, T., & Spalazzi, L. (2025, May). Verifiable Actor Model Systems Through Relational-Model Multi-Agent System and Zero-Knowledge Proofs. In *2025 IEEE 8th International Conference on Industrial Cyber-Physical Systems (ICPS)* (pp. 01-06). IEEE.

[31] Pirani, M., Cucchiarelli, A., Naeem, T., & Spalazzi, L. (2025). A Blockchain-Driven Cyber-Systemic Approach to Hybrid Reality. *Systems*, *13*(4), 294.

[32] Naeem, T., Pirani, M., & Spalazzi, L. (2025), "Evidence-based Oracles Using Bayesian Network", (in press) in 2025 21st International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)

[33] Pirani, M., Bonifazi, G., Cucchiarelli, A., Naeem, T., & Spalazzi, L. (2025), "Holonic Oracle Constructivism in Cyber-Physical Systems", (in press) in 2025 IEEE International Conference on Systems, Man, and Cybernetics (SMC), October 5–8, 2025, Vienna, Austria.

[34] Naeem, T., Pirani, M., & Spalazzi, L. (2025), "Universal Wallet for Trustless Cross-Chain Interoperability via Merkle Proofs", (in press) in 23rd International Conference on Pervasive Intelligence and Computing (2025) – The 4th International Workshop on Hybrid Internet of Everything Models for Industry 5 (HIEMI 2025)

Central to the system is the ability to monitor CO_2 credits, GHG savings, and compensatory actions. Companies accumulate CO_2 credits through direct emissions reductions or through verified compensatory initiatives. Graphical interfaces display savings over time, either on a monthly or yearly basis, allowing both short- and long-term analysis. Comparative dashboards also benchmark companies against one another, ranking them by reputation scores and cumulative savings.

The system implements five main types of transactions: rewards for verified GHG savings, penalties for excessive emissions, token transfers between companies, compensation rewards for offsetting initiatives, and token requests to obtain credits from peers. Every transaction is recorded on the blockchain, ensuring immutability, transparency, and prevention of double-counting. Each company's transaction history is fully traceable, and its CO_2 credits and sustainability actions are always verifiable.

Companies are able to register compensatory initiatives by providing descriptions, saved CO_2 amounts, and supporting certificates. These actions are archived in a public history, ensuring verifiability. Blockchain technology guarantees integrity by making all such records immutable, auditable, and resistant to manipulation.

The platform employs two main token types. Fungible Tokens (FTs), representing CO_2 credits, are interchangeable and divisible, functioning much like currency. They are earned for each ton of CO_2 saved. Soulbound Tokens (SBTs), representing company reputation, are unique and non-transferable, reflecting a firm's long-term commitment to sustainability. Reputation tokens cannot be bought or traded; they must be earned through direct actions and consistent performance.

The reputation mechanism is based on a weighted formula that incorporates five factors: direct GHG reductions, streak bonuses for consistent performance, compensatory actions, donations of credits to others, and penalties for excess emissions. The raw score is normalized into a 0–100 ranking, allowing fair comparison across the network. This transparent system promotes continuous improvement, rewarding proactive behavior while penalizing unsustainable practices.

The integration of blockchain into sustainability management offers several key advantages: immutable trust, as every transaction is permanent and tamper-proof; complete transparency, as all activities are traceable and auditable; no double counting, since credits are unique digital assets; and automated efficiency, where smart contracts regulate rewards, penalties, and transfers, minimizing administrative burdens. Together, these benefits establish a reliable ecosystem that fosters accountability and collaboration.

This sustainable blockchain framework introduces a novel mechanism for aligning corporate sustainability with financial and reputational incentives. By tokenizing CO_2 reductions and embedding them into a transparent blockchain environment, companies are motivated to adopt greener practices, engage in compensatory actions, and collaborate with peers. The dual-token system of fungible credits and soulbound reputational assets strengthens accountability and supports a systemic culture of continuous improvement.

We will describe this application with the help of some screenshots in the following.

Figure 6.2 introduces the central theme: "Sustainable Blockchain – The Future of Food Supply". The core message is that blockchain technology can act as an enabler of sustainability within food supply chains, offering innovative tools such as the tokenization of emissions and the creation of transparent reputation mechanisms. The idea is that each ton of CO_2 saved or offset can be transformed into a digital token, and that virtuous behaviors are rewarded not only with environmental benefits but also with economic and reputational gains. In this way, companies are incentivized to adopt more responsible practices.

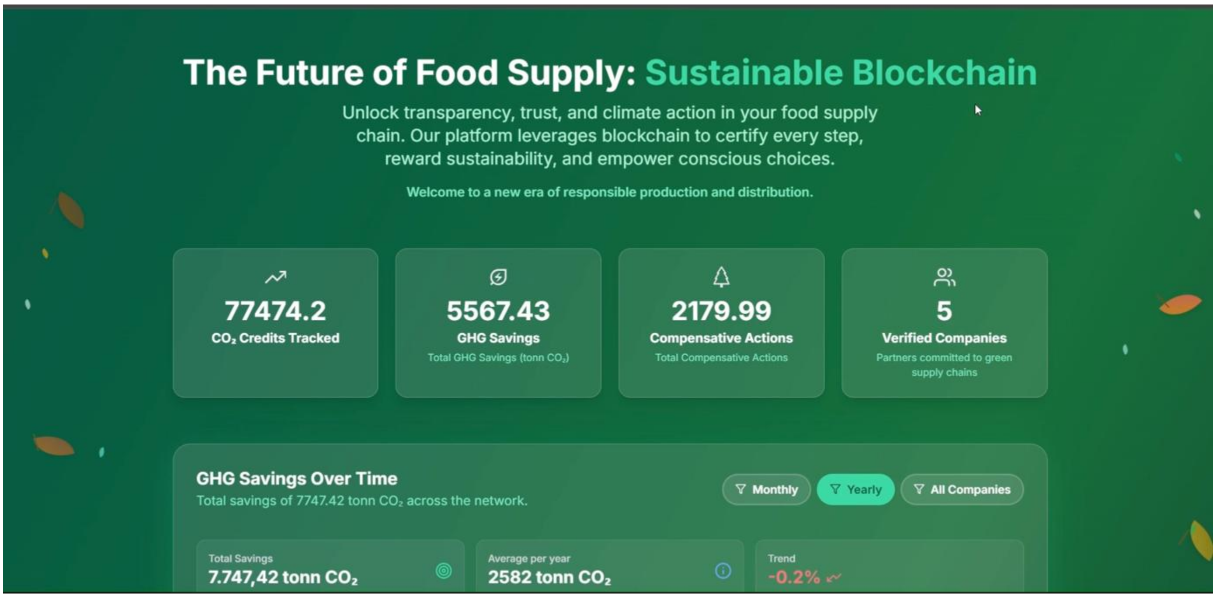

Figure 6.2: Initial page of the application[35].

Figure 6.2, introduces the concept of the Sustainable Blockchain, highlighting how blockchain technology can be applied to monitor, certify, and incentivize sustainability in food supply chains. The system is built around four main indicators that together provide a transparent and auditable picture of environmental performance. The first of these is the amount of CO_2 credits tracked, which represents the total credits owned by the company as a direct result of its recognized efforts to reduce emissions or to compensate for them. Closely linked to this, the indicator for GHG savings expresses in tonnes of CO_2 the actual greenhouse gas reductions achieved through operational improvements, efficiency measures, or sustainable practices. Complementing these, the system also accounts for compensatory actions, which quantify the amount of CO_2 offset through external measures such as tree planting, renewable energy projects, or the purchase of certified carbon offsets. Finally, the number of verified companies provides an additional layer of credibility, showing how many partners have been authenticated within the system and are actively monitored for their sustainability commitments.

The dashboard in the figure provides a concrete snapshot of how these metrics are reported: more than 77,000 CO_2 credits are tracked, over 5,500 tonnes of CO_2 have been saved, more than 2,100 compensatory actions have been carried out, and five companies have been verified in the system. These values exemplify how the platform aggregates performance data and makes it visible to all participants.

At the center of the view, the banner frames the value proposition: blockchain is used to unlock transparency, trust, and climate action. By certifying every step of the process, rewarding sustainable behavior, and empowering conscious decision-making, the system ensures that improvements are not only pursued but also demonstrated in a way that is tamper-proof and easily auditable. In this sense, the Sustainable Blockchain represents a step forward toward responsible production and distribution, where private company actions align with broader environmental goals and where data integrity is guaranteed through the immutable structure of blockchain technology.

[35] All the figures provided in the images and in the discussion are given only for demonstration purposes. This application is a simulation and might not have a realistic cgrounding and ounterpart in some of its parts.

Figure 6.3, presents the project dashboards, which serve as the visual and operational interface for the different stakeholders involved. Two levels of dashboards are highlighted:

- the Global Dashboard, which provides a real-time overview of all participating companies, enabling comparative analysis, trend visualization of CO_2 reduction, and overall performance assessment;
- the Company Dashboard, which is tailored to each company, offering targeted indicators, the balance of CO_2 tokens, the reputation score, and operational suggestions for improving sustainable practices.

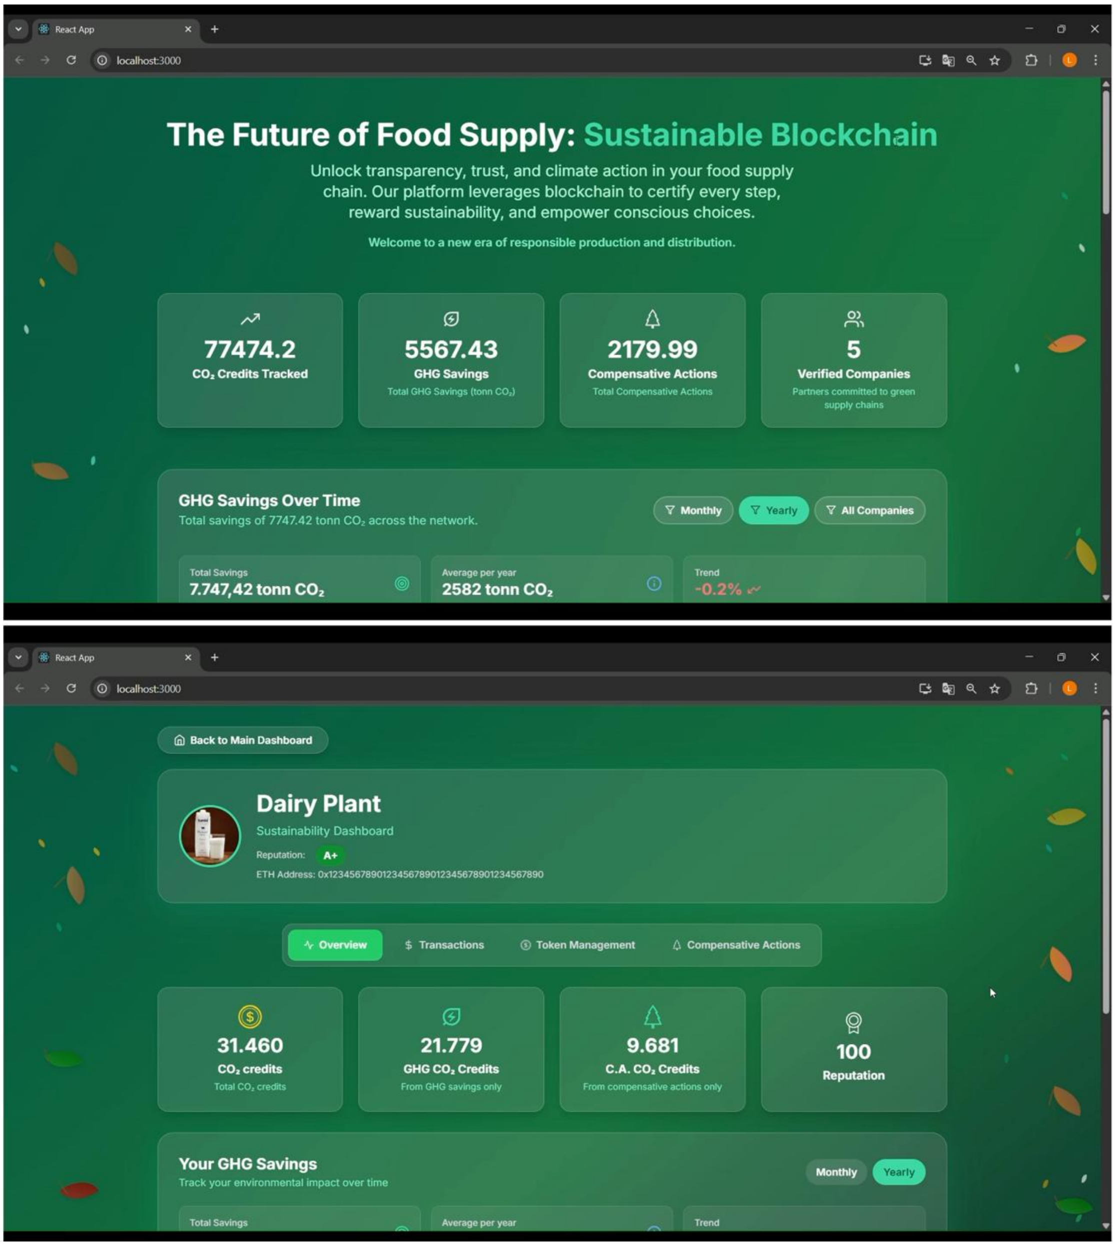

Figure 6.3: Global (top) and Company (bottom) dashboards.

The key concepts monitored by the Global system are: CO_2 Credits, representing the total amount of credits held by a company; GHG Savings, referring to the direct reduction of greenhouse gas emissions;

Compensatory Actions, which include certified offsetting initiatives; and finally, the number of Verified Companies, i.e., those monitored within the network.

This context is supported by extensive data visualization through time-series graphs. One example is the chart of GHG Savings per month (Figure 6.4), which enables the analysis of monthly CO_2 savings per company or per reference year. The graph highlights not only total and average values but also short-term trends, which are useful for promptly identifying potential performance declines.

The view of GHG Savings per year (Figure 6.5) expands the perspective to a multi-annual horizon. Companies can observe their trajectory of emission reductions on a yearly scale, benchmark themselves against sector averages, and identify consolidated trends that signal long-term sustainability progress.

Figure 6.4: GHG Savings per month.

Figure 6.5: GHG Savings per year.

Figure 6.6, introduces the **Sustainable Companies dashboard**, a comparative tool designed to evaluate and benchmark the performance of different companies within the system. At the heart of the dashboard is the **Reputation Score**, a metric expressed on a 100-point scale that reflects each company's sustainability performance. The score itself is derived from multiple factors, including verified greenhouse gas (GHG) reductions, compensatory actions, and collaborative behavior. To make the interpretation more intuitive, the score is also mapped into a **Reputation Category**, with labels such as A+, B, D, or E, similar to grading systems commonly used in sustainability reporting.

Alongside this qualitative evaluation, the dashboard also quantifies the **Total GHG Savings** in tonnes of CO_2 equivalent, offering a direct measure of the cumulative environmental impact achieved by each participant. This column captures both reductions from operational improvements and offsets from compensatory measures, thus giving a holistic view of climate performance.

The example displayed in the figure shows how companies compare to one another. The Dairy Plant leads with a perfect reputation score of 100, corresponding to category A+, and a total saving of more than 3,100 tonnes of CO_2. The Coolstore follows with a score of 72, placing it in category B, supported by over 2,000 tonnes of savings. The Transporter has a lower score of 38, falling into category D despite achieving nearly 1,000 tonnes of savings, which suggests gaps in consistency or behavior. The Retailer performs worse, with a score of only 22 in category E, even though its cumulative savings exceed 1,400 tonnes. Finally, the Display-cabinet manager shows minimal contribution with a score of 0 and just over 120 tonnes saved, also ranking in category E.

Through this comparative view, the dashboard highlights not only the absolute environmental contributions of each company but also their relative reputation and trustworthiness within the ecosystem. It underscores the principle that sustainability is not measured solely by raw savings but also by continuous improvement, reliability, and alignment with shared goals across the supply chain. In this way, the dashboard becomes a transparent accountability tool, encouraging participants to both celebrate high performance and address shortcomings in order to move collectively toward greater climate responsibility.

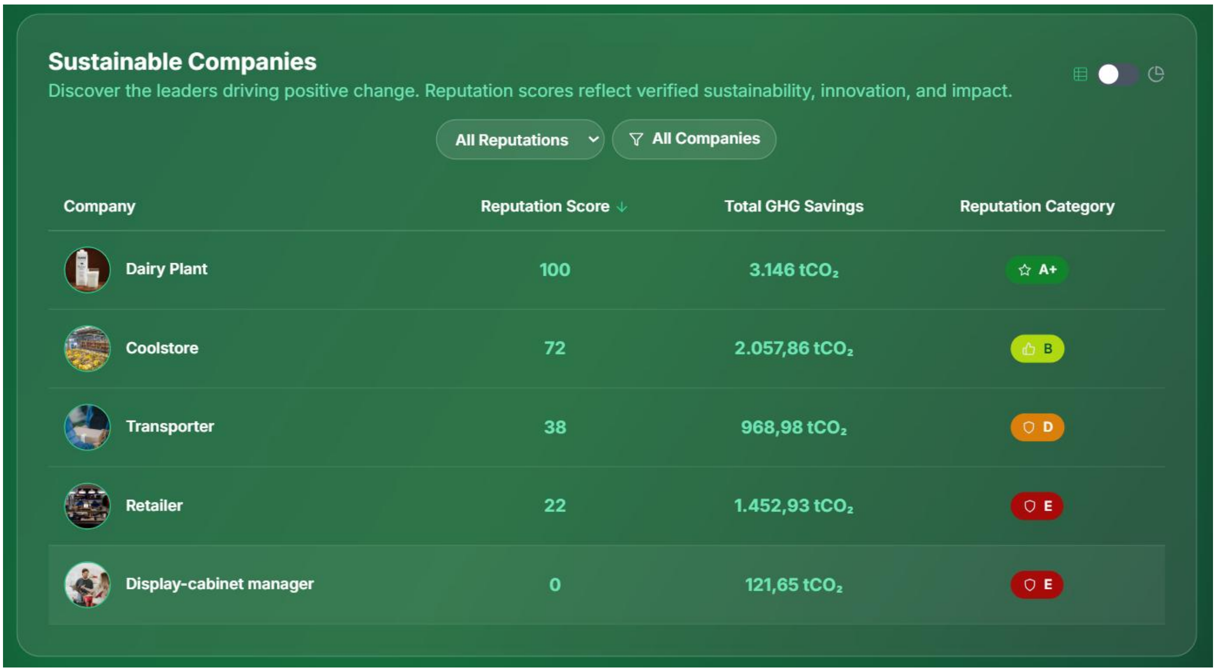

Figure 6.6: The Sustainable Companies dashboard.

Figure 6.7 expands the view of sustainable companies from individual performance to an aggregate perspective. Instead of focusing on single entities, it groups companies according to their Reputation Category, showing the relative distribution across the network. The central element is a pie chart that divides participants into performance classes, making visible at a glance how many companies are excelling, how many are performing moderately, and how many are lagging behind.

The reputation scale on the right provides the reference: companies scoring between 95% and 100% fall into A+, representing leaders with outstanding sustainability practices. Scores between 70% and 84% correspond to category B, showing good but not perfect performance. On the lower end, D covers scores between 30% and 49%, signaling critical weaknesses, while E (0–29%) represents companies with very poor sustainability results.

In the current snapshot, the distribution is heterogeneous: 20% of companies are in A+, another 20% in B, and 20% in D. The largest share, however, is in E with 40%, highlighting that many participants still fall far short of the expected standards. This distribution suggests that while a minority of companies are setting ambitious examples, the majority still require substantial improvement.

This view is powerful because it highlights systemic strengths and weaknesses rather than isolated performances. It enables decision-makers, regulators, or consortium coordinators to identify where to focus collaborative improvement efforts. For instance, the prevalence of E-category companies could trigger targeted programs to assist low performers in adopting best practices. Conversely, A+ and B companies can be showcased as benchmarks and role models within the ecosystem, sharing methods and innovations that others can adopt.

In essence, this slide transforms raw sustainability scores into a **collective diagnostic tool**, providing an overview of the network’s health and a foundation for strategic interventions to lift the overall performance of the supply chain.

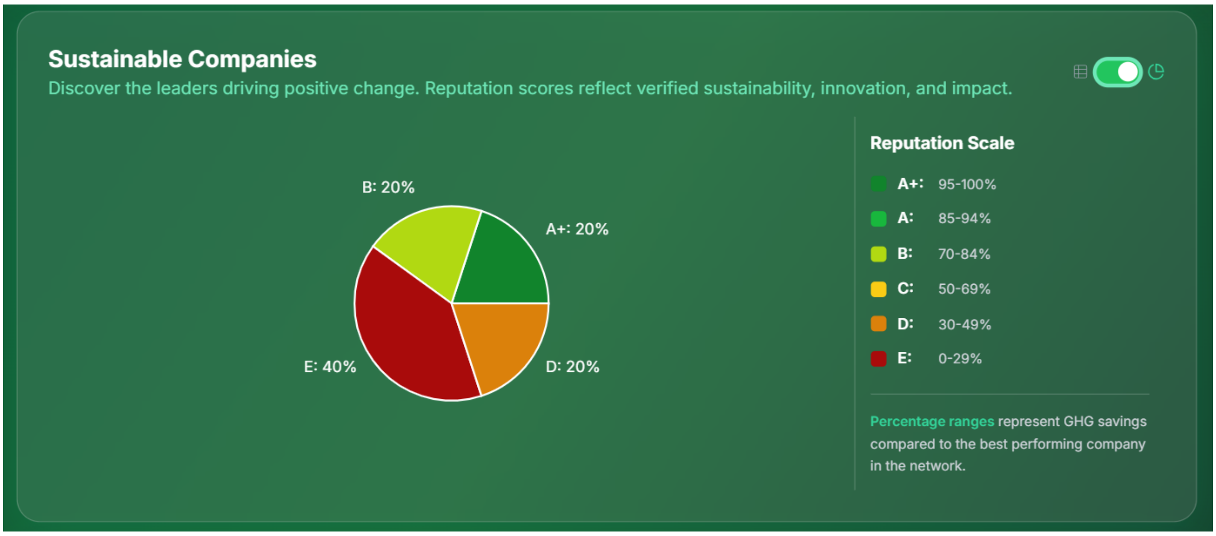

Figure 6.7: The view of sustainable companies from individual performance to an aggregate perspective.

Figure 6.8, provides a comprehensive overview of sustainability metrics by comparing how different companies contribute to reducing greenhouse gas emissions, both through their own operations and through compensatory activities. The table shown at the center of the slide distinguishes clearly between two types of contributions: GHG Savings, which measure direct reductions achieved internally, and Compensative GHG Savings, which reflect emissions offset via external initiatives such as carbon credits or approved compensatory projects.

The data highlight how performance varies across companies. The Dairy Plant emerges as the strongest contributor, with over 2,177 tons of CO_2 saved internally and an additional 968 tons offset through compensation, demonstrating both operational efficiency and proactive engagement in offsetting measures. The Coolstore also shows a solid performance, with around 1,452 tons saved directly and over 600 tons offset, reflecting a balanced approach between internal improvements and compensations. The Transporter, though operating at a smaller scale, still reports measurable savings of 726 tons directly and 242 tons through compensation, proving that even logistics players can make a meaningful impact.

The Retailer presents a mixed profile: while it records 1,088 tons of direct savings, its compensatory actions are relatively modest at 364 tons, suggesting room to expand offsetting activities to complement its operational gains. Finally, the Display-cabinet manager stands out negatively, with only 121 tons of direct savings and no compensations at all, highlighting a critical gap in engagement with sustainability practices.

By presenting both categories of impact side by side, the slide underlines the importance of not only reducing emissions at the source but also investing in offsetting strategies to reach broader sustainability goals. This comparative view allows stakeholders to identify leaders, laggards, and opportunities for improvement, while at the same time reinforcing the value of a dual strategy: operational efficiency plus compensatory responsibility.

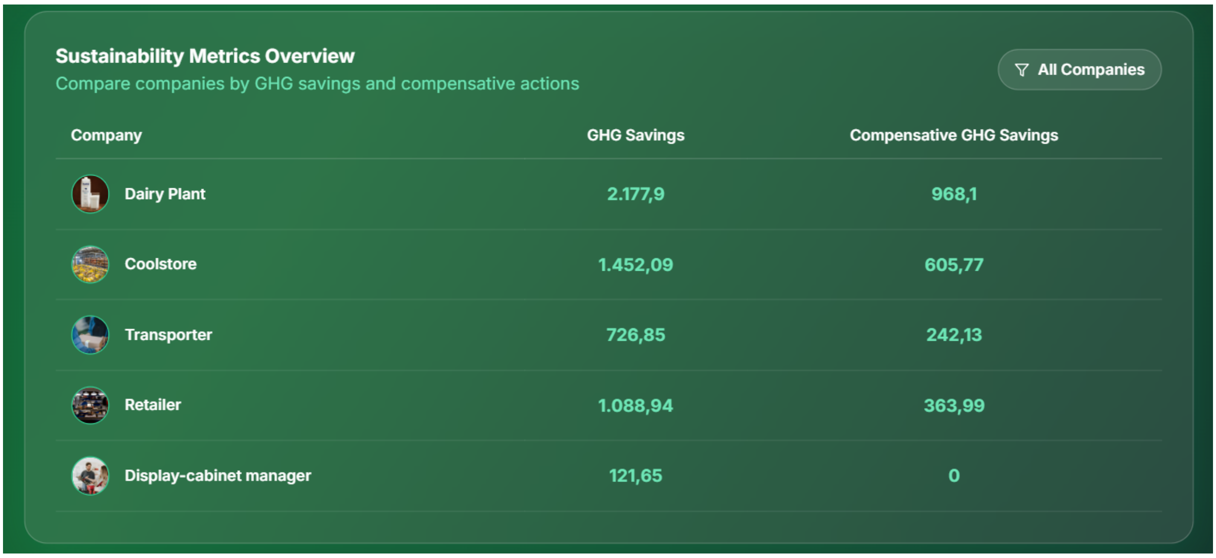

Sustainability Metrics Overview
Compare companies by GHG savings and compensative actions
All Companies

Company	GHG Savings	Compensative GHG Savings
Dairy Plant	2.177,9	968,1
Coolstore	1.452,09	605,77
Transporter	726,85	242,13
Retailer	1.088,94	363,99
Display-cabinet manager	121,65	0

Figure 6.8: A comprehensive overview of sustainability metrics.

Figure 6.9, introduces the **Company Dashboard**, a central tool that illustrates the sustainability profile of an individual company within the SDS ecosystem. In this example, the focus is on the **Dairy Plant**, which is presented with a clear, structured view of its environmental performance and accountability.

At the top, the dashboard displays the company's reputation category, rated here as A+, alongside its public Ethereum address, ensuring full transparency of its transactions on the blockchain. This reinforces the principle that sustainability claims are not only self-reported but also verifiable through immutable records.

The main section of the dashboard highlights four key metrics. First, the total CO_2 credits owned by the company, here shown as 31,460, representing the cumulative environmental assets under its control. Second, the share of those credits that come from direct GHG savings, amounting to 21,779, which reflects actual emission reductions achieved through operational improvements. Third, the portion earned via compensatory actions, here 9,681, documenting the offsets the company has implemented, such as reforestation or renewable energy investments. Finally, the company's overall reputation score, rated at 100, serves as a composite indicator of its trustworthiness and consistency in sustainable behavior.

By combining on-chain transparency, direct performance indicators, and a synthesized reputation score, the Company Dashboard offers both a snapshot and a deep dive into how a single actor contributes to the broader sustainability goals of the supply chain. It allows stakeholders—whether partners, regulators, or consumers—to quickly assess not just the stock of credits, but also the quality of the company's actions and the credibility of its claims. This makes it both a monitoring tool and a reputational instrument, aligning business performance with accountability in the green transition.

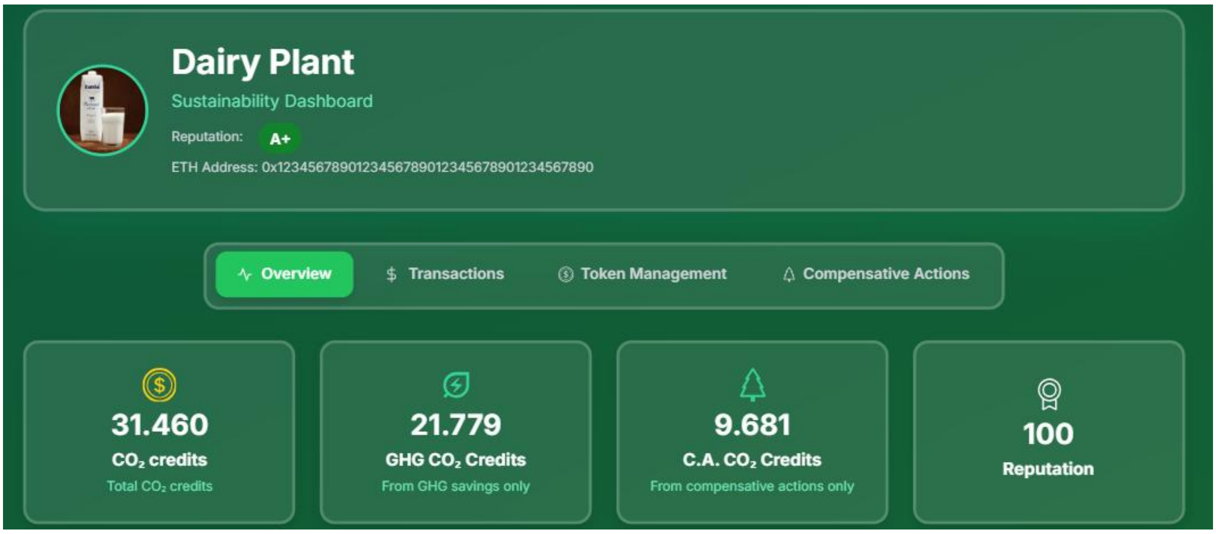

Figure 6.9: A detail of the Company Dashboard (for the Dairy Plant Participant of the SCE).

There are the **five fundamental types of transactions** that structure this SDS system's carbon credit economy application, ensuring that every environmental action—positive or negative—is consistently tracked, recorded, and auditable. Together, they provide the rules that govern how greenhouse gas (GHG) savings, penalties, and transfers are accounted for across the network.

The first category is the **Reward Received**, which is triggered whenever a company achieves verifiable GHG savings. This type of transaction translates actual operational improvements into digital credits, rewarding efficiency and sustainable practices. On the opposite side is the **Penalty Received**, applied when a company emits beyond its permitted threshold. By deducting credits, this mechanism introduces accountability and discourages environmentally harmful behavior.

A third category, the **Token Transfer**, comes into play when one company chooses to transfer credits to another, reflecting peer-to-peer exchanges. This could occur, for instance, when a surplus holder supports a partner in need, enabling flexibility within the ecosystem. Closely related is the **Token**

Request, which occurs when a company formally asks for credits and, upon approval, receives them from another participant. This emphasizes collaboration and ensures that the credits flow where they are most needed.

Finally, the **Compensation Reward** recognizes companies that offset emissions through certified compensatory actions—such as reforestation projects or renewable energy investments. These activities generate additional credits, validating not only reductions but also proactive contributions to climate goals.

By standardizing these five transaction types—rewards, penalties, transfers, requests, and compensations—the SDS platform creates a transparent, traceable, and balanced token market. Every credit can be tracked back to its origin, preventing double counting and ensuring that sustainability claims are both measurable and credible.

Figure 6.10 presents the **Company Transaction History**, a central feature that ensures transparency and accountability within the SDS framework. From the *Transactions* tab, users can review the full, time-ordered ledger of all activities related to a company's CO_2 credits. Each entry is recorded with details such as the type of transaction, its description, the amount of credits involved, the date, and the corresponding blockchain hash for verification.

The history includes all five categories of transactions previously introduced: **Rewards Received** for verified GHG savings, **Penalties** for excess emissions, **Token Transfers** between companies, **Token Requests** when credits are obtained from others, and **Compensation Rewards** generated through certified offset projects. For example, in the screenshot shown, we see records of a token request fulfilled, multiple rewards for GHG savings, a transfer to another participant, and a reward linked to an "Ocean Cleanup Project."

This setup provides end-to-end traceability: every credit earned, spent, or exchanged is documented and anchored on-chain, meaning it cannot be tampered with or lost. Companies gain a reliable tool for internal reporting, auditors have verifiable evidence, and partners can trust that credits reflect genuine actions. In practice, the transaction history functions as both an operational logbook and a compliance report, bridging sustainability efforts with financial and reputational accountability.

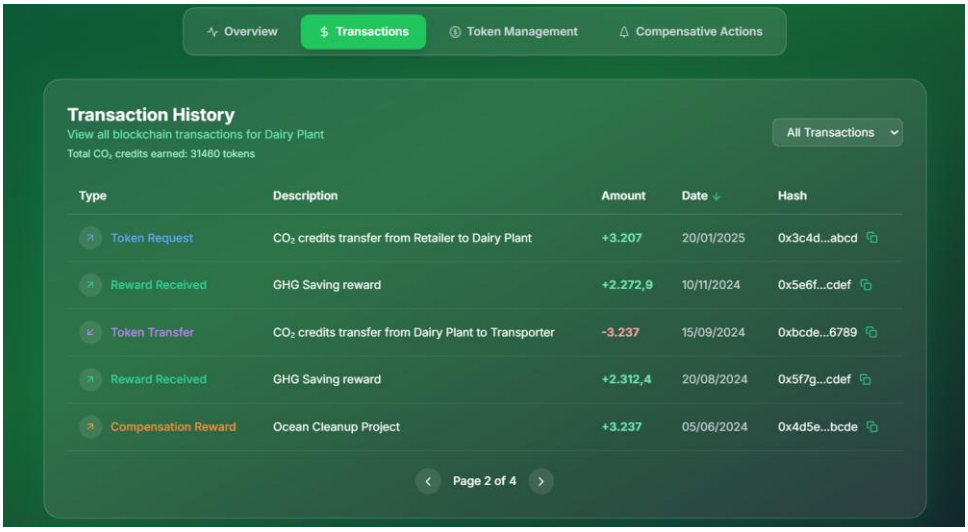

Figure 6.10: The Company Transaction History.

Figure 6.11 shows the **Request & Transfer CO_2 Credits Interface**, which is one of the most practical features of the application. It shows how companies can directly manage their carbon credits by either asking for new ones or transferring them to other participants in the network.

On the top, we see the **Request CO_2 Credits** panel. Here, a company can enter the amount of credits it needs and specify the purpose of the request, ensuring that all credit demands are contextualized and auditable. This prevents arbitrary allocation and keeps the system transparent.

On the bottom, the **Transfer CO_2 Credits** panel enables a company to send credits it already owns to another participant. To do this, the sender specifies the recipient's blockchain (Ethereum) address and the number of credits to be transferred. Since all transactions are logged on-chain, each transfer is verifiable, immutable, and linked to the identity of both parties.

Together, these two functions create the backbone of **peer-to-peer carbon credit balancing** within the SDS ecosystem. Companies that generate a surplus through strong performance or compensatory actions can redistribute credits, while those falling short can cover their emissions responsibly[36]. Importantly, the process remains simple and accessible through the user-friendly interface, while the blockchain layer guarantees trust and accountability.

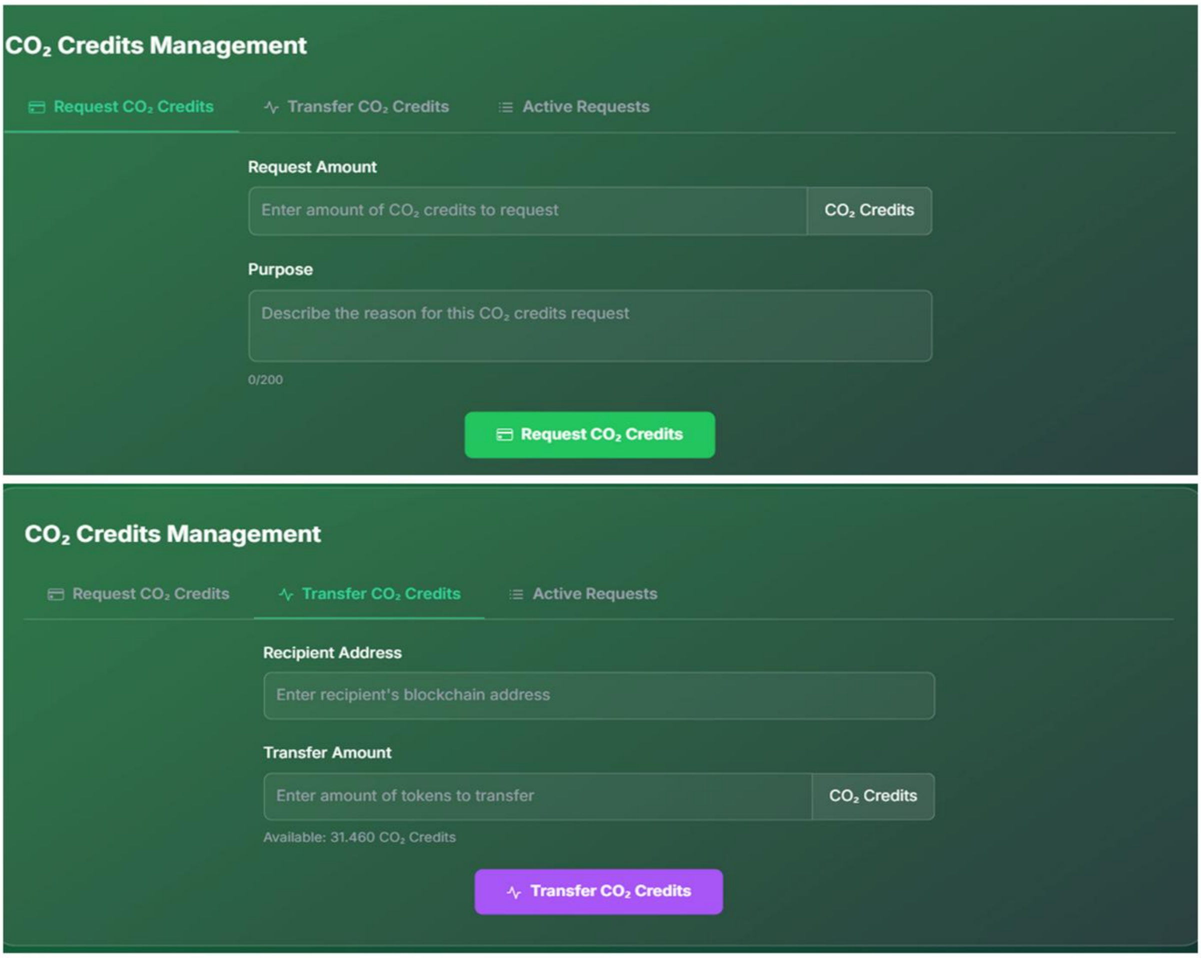

Figure 6.11: Transfer (bottom) and request (top) of CO_2 credits.

[36] In the case the tokens are bought of course there is always a debt that this actor has to recover in the future for an actual responsible conduction. Temporary debt is allowed to give the chance to an actor to initiate a transformation and an improvement. A mechanism that is useful and incentivizing in the case of certain processes or small-medium sized companies that experience a lag in the participation in an otherwise virtuous collective.

Figure 6.12 shows the **Active Request Interface**, which is a central point for managing all pending CO_2 credit requests across the network. The interface provides visibility into who is requesting credits, how many they are asking for, and the specific purpose behind the request. For example, in the illustration, the Coolstore requests 3,000 credits to offset emissions, the Transporter requests 2,500 credits to meet its annual sustainability goal, and the Retailer requests 1,500 credits for event-related compensation.

Each request is accompanied by a *Details* button, which, when selected, reveals the full request information. This includes the requester's name, blockchain address, the exact number of credits needed, and the reason for the request. From this view, the responding company can decide whether to fulfill the request and directly transfer the required credits, making the system both operationally efficient and transparent.

The Active Request Interface plays a crucial role in facilitating peer-to-peer collaboration within the SDS ecosystem. It enables companies with surplus credits to identify opportunities to support others in need, reinforcing the marketplace dynamics of balancing carbon responsibilities. By combining traceability, purpose-driven requests, and blockchain-based execution, the platform ensures that credit distribution is fair, auditable, and aligned with both individual company needs and collective sustainability goals.

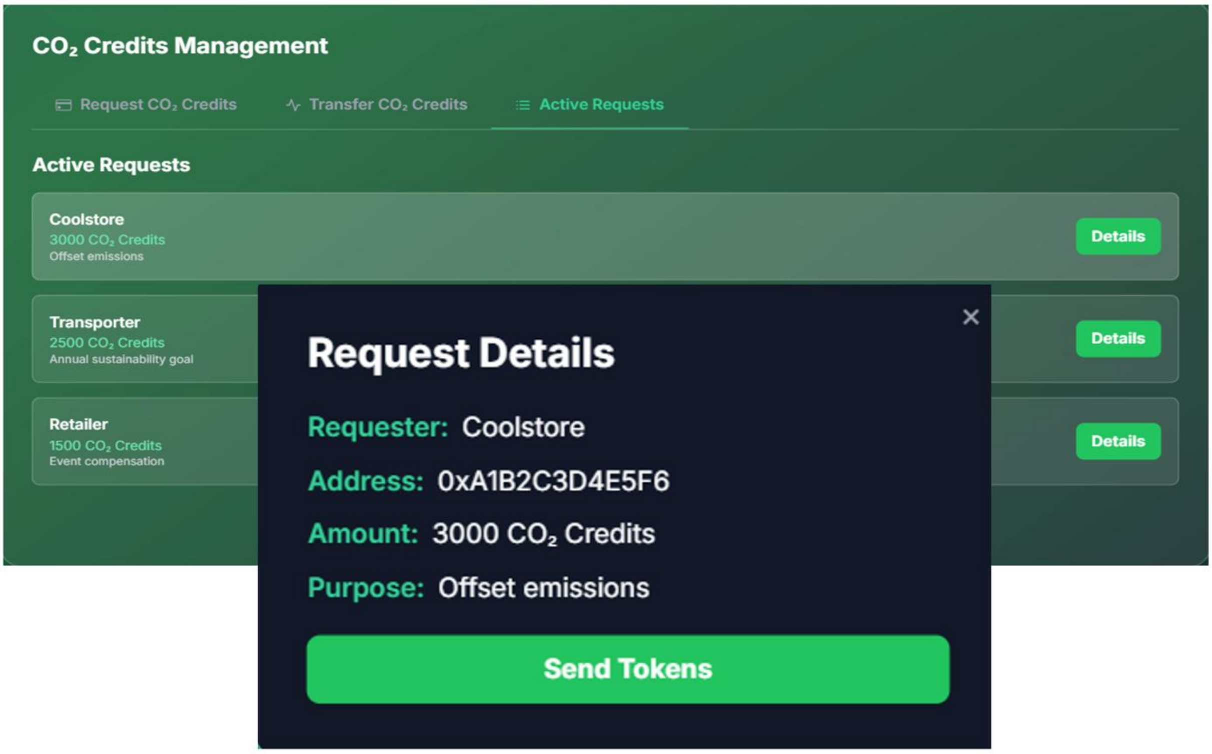

Figure 6.12: The Active Request Interface.

Figure 6.13 presents the New Compensative Action interface, a key feature that allows companies to formally register initiatives aimed at offsetting their carbon footprint. The process is designed to ensure transparency and credibility in how compensations are introduced into the system. To submit a new action, the company must first specify an object, essentially a concise title that captures the initiative's essence. They then provide a detailed description, explaining what the compensative action entails, its scope, and its expected benefits.

An essential part of this submission is quantification: the company must declare the **total tonnes of CO_2 saved** as a direct result of the action. This figure is not taken at face value but must be supported by documentation. For this reason, the interface includes a field to upload a supporting PDF certificate, which acts as verifiable evidence—such as third-party audits, official approvals, or certified reports. Only after these details are provided can the action be submitted for validation within the SDS platform.

This mechanism ensures that compensative actions are not just declarations but documented and traceable commitments. It strengthens accountability by linking each initiative to verifiable evidence and integrates seamlessly with the SDS crediting system, where approved actions can generate additional CO_2 credits.

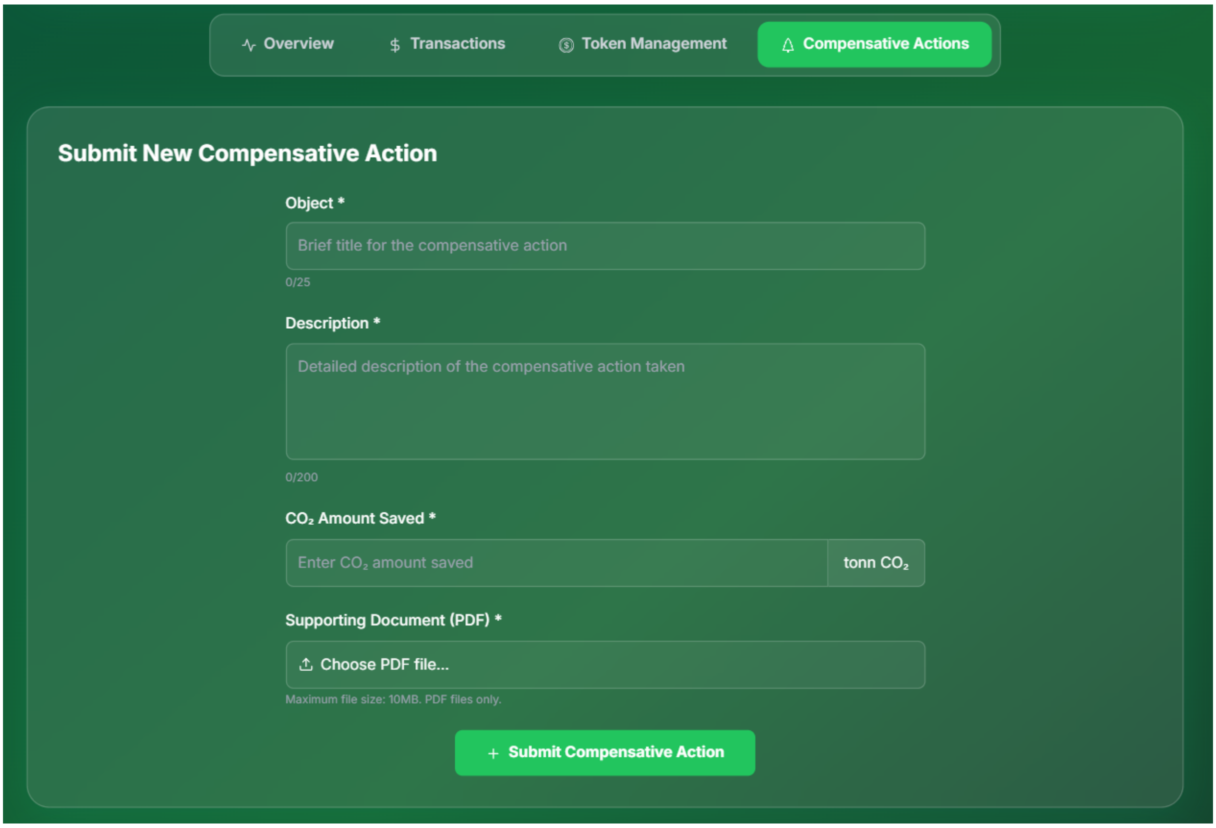

Figure 6.13: The compensative action GUI.

Figure 6.14 shows the **Compensative Action History** interface, which provides a full record of all the compensatory measures that a company has undertaken and had approved. The aim is to ensure transparency and accountability in how organizations offset their emissions. Each entry in the history lists the project title, its approval status, a short description of the initiative, the amount of CO_2 saved or offset, and the date of registration. For added verification, a direct download link is available, giving access to the official supporting documents—such as certificates or signed agreements—that validate the action.

For example, projects like *Forest Restoration*, *Ocean Cleanup Project*, and *Wetland Conservation* are included in the sample view. Each contributes in a different way: planting trees to capture CO_2, removing plastic waste to protect ecosystems, or conserving wetlands as natural carbon sinks. These projects are not just numerical contributions; they highlight the diverse strategies companies can adopt to balance emissions with restorative actions.

By consolidating this information, the history page serves two purposes. Internally, it allows companies to keep track of their long-term environmental commitments. Externally, it provides stakeholders, regulators, and customers with a clear, auditable trail of compensations, ensuring that claims of sustainability are backed by verifiable evidence. This reinforces the credibility of the SDS system and promotes trust across the supply chain ecosystem.

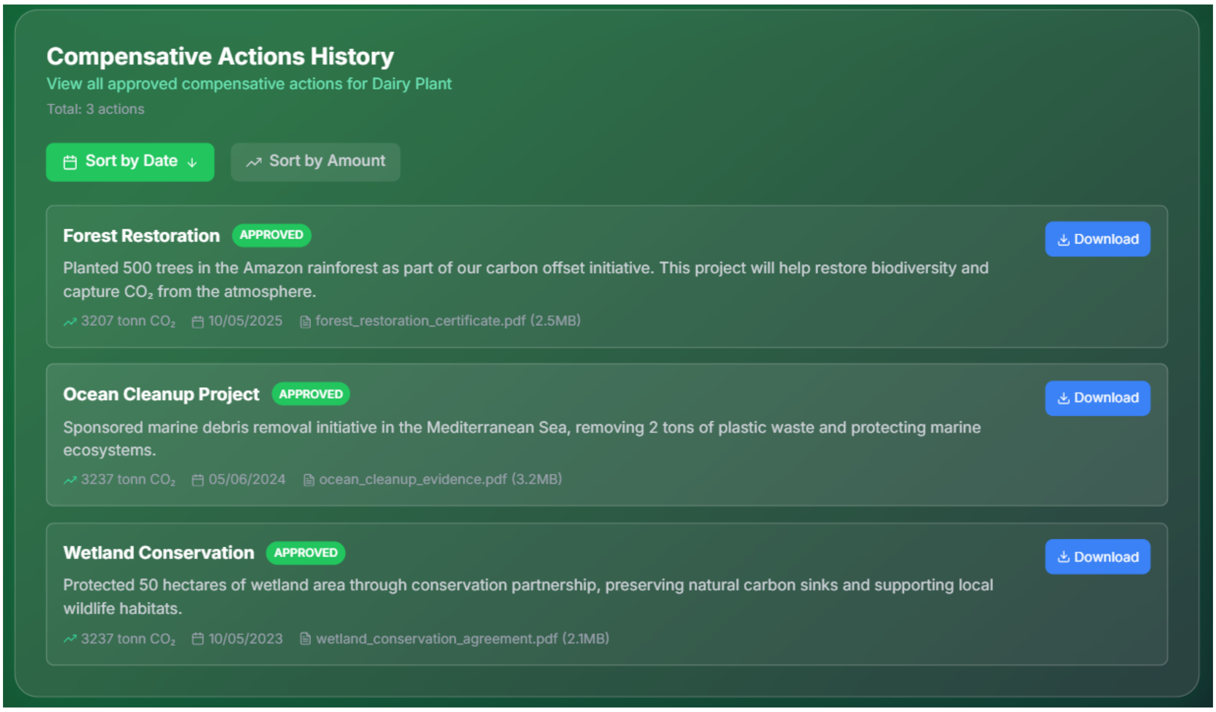

Figure 6.14: The compensative action history GUI.

All in all, the proposed system leverages blockchain technology to create a sustainable ecosystem where companies are rewarded for their environmental performance and held accountable for their excesses. At the foundation of this model are the intrinsic benefits of blockchain itself. The immutability of transactions ensures that once data is recorded, it cannot be tampered with, thereby eliminating risks of greenwashing and preserving trust in the reported figures. Transparency is guaranteed because all GHG savings and compensatory actions are fully traceable and auditable on the distributed ledger. Moreover, by tokenizing CO_2 credits as unique digital assets, the system prevents double counting, ensuring that credits cannot be fraudulently claimed more than once. Efficiency is introduced through smart contracts, which automate the processes of rewarding, transferring, and verifying transactions, creating an ecosystem that is both secure and streamlined.

To operationalize this, the platform introduces two categories of tokens. **Fungible Tokens (FTs)** represent CO_2 credits. Like currency, these tokens are interchangeable and divisible, allowing companies to earn, stack, or transfer them based on their performance in reducing emissions.

Alongside these, **Soul Bound Tokens (SBTs)** capture reputation. Unlike FTs, these tokens are unique, non-transferable, and cannot be bought or sold; they must be earned through genuine action. The Reputation SBT reflects a company's ongoing commitment to sustainability, thereby turning intangible behavior into a measurable and visible digital asset. This dual-token system creates a straightforward cause-and-effect logic, offering financial incentives for sustainability while embedding trust and accountability into corporate (Participant's) practices.

CO_2 credits themselves are generated through automated rules. For every verified ton of CO_2 saved, a company earns a fixed number of credits (X). Conversely, for every ton of excess emissions, a deduction (Y) occurs. These parameters are configurable, meaning the system can be fine-tuned to emphasize certain sustainability goals or adapt to sector-specific requirements. Thus, credits act as both a reward and a deterrent, steering companies toward continuous improvement.

Beyond credits, the system integrates a dynamic reputation mechanism. Reputation is not static but shaped by four pillars. The most significant factor is direct GHG reductions: companies that consistently lower emissions are heavily rewarded, while those with excesses are penalized. Consistency is another pillar, with a "streak bonus" rewarding sustained positive performance, although even a single negative event can reset it. Proactive offsetting through compensatory actions—such as verified reforestation or ocean cleanup projects—provides an additional boost. Finally, ecosystem collaboration encourages companies to donate credits to peers in need, fostering solidarity across the network.

The reputation score today is calculated through a very simple weighted formula, which has to be considered as a mere placeholder for a more specific and well-designed one:

Reputation = (GHG Savings × X) + (Streak Bonus × Y) + (Compensatory Actions × Z) + (Donations × K) − (Excess Emissions × V).

This raw score is then normalized to a 0–100 scale, ensuring comparability across all participants. By benchmarking against both the highest and lowest performers, the system creates a fair and transparent ranking.

Together, credits and reputation form a dual system of performance indicators. Credits represent the quantitative stock of emissions reduced or compensated, while reputation represents the qualitative trust and credibility a company builds over time. These elements, supported by blockchain's integrity and smart contract automation, enable a marketplace of sustainable behavior where companies can improve individually and contribute to collective goals.

In essence, the platform goes beyond simply monitoring emissions: it transforms environmental responsibility into a transparent, incentivized, and collaborative framework, aligning private actions with global sustainability objectives.

7 BUSINESS USE CASE 4: CHANNELLING INNOVATION THROUGH THE MARKETPLACE

The final business use case is dedicated to actions that go beyond the development of products for the SDS Marketplace as in previous business use case 2 and 3 respectively treated in chapter 5 and 6.

Here the goal is to create a persistent innovation channel for conveying research through a technology transfer process. This means that the academic or research entities are more involved in this case rather than the in the others.

Again, this channel is created by the peculiar inclusivity and flexibility of the SDS concept and its general purpose infrastructure.

In Figure 7.1 we recall the scheme of this business use case for reader's convenience.

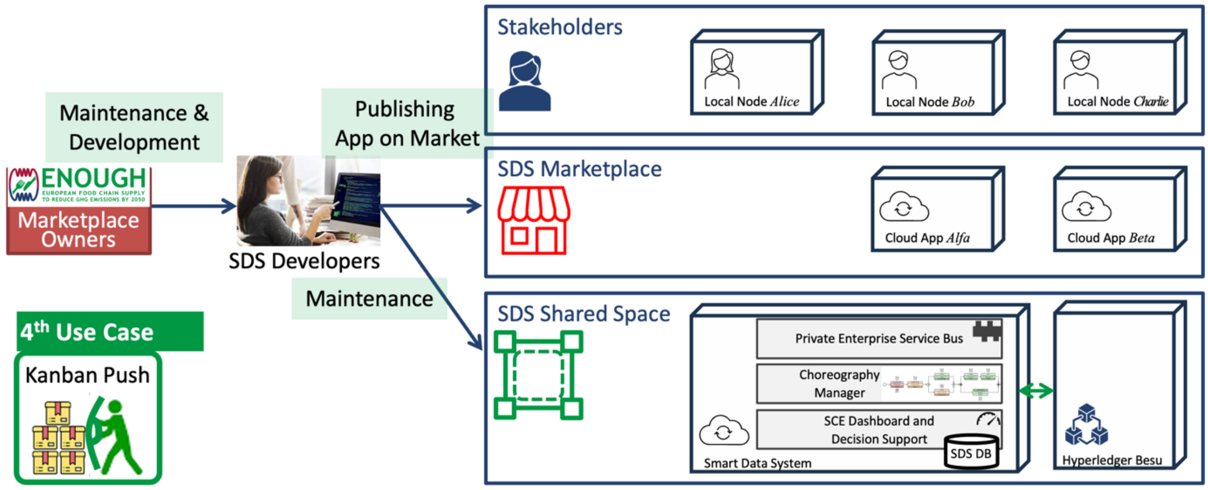

Figure 7.1: Business case 4 – SDS Owners introduce new modules into the marketplace, ensuring continuous upgrades (see section 3 for more details).

In this fourth business use case, the SDS Owners themselves take an active role in introducing new modules and functionalities into the SDS Marketplace. Unlike third-party providers, who require a more extensive phase of requirement collection and integration support, the Owners of the system have intrinsic knowledge of the SDS architecture, its internal processes, and the interoperability layers already in place. This privileged position allows them to focus on rapid deployment and seamless integration, as the technical and organizational prerequisites are already well understood within the consortium. It is in this similar to business use case 2, but it is more dedicated on the one hand to research and innovation of the SDS framework itself and, on the other hand, to use it at best for challenging and disruptive innovations.

The addition of new modules can serve multiple purposes. On one side, it enables the continuous enrichment of the SDS ecosystem with functionalities that respond to emerging needs identified during real-world testing and stakeholder feedback. On the other, it supports the alignment of the platform with evolving policy frameworks, sustainability targets, and market demands that interdisciplinary research allows. Since these modules are developed by the Owners, their integration requires only minimal adaptation and validation, streamlining the process considerably compared to external applications.

Equally important, this business case emphasizes the role of ongoing system maintenance and upgrades, which are coordinated and driven by the ENOUGH business unit. Regular updates ensure that the SDS remains secure, resilient, and capable of accommodating new non-functional requirements such as performance, scalability, or compliance with evolving standards. Maintenance activities include bug fixing, optimization of existing services, and refinement of interfaces to guarantee a smooth user experience across different participant categories.

In practice, this means that the SDS Marketplace is not a static repository of tools but a dynamic environment in which core functionalities are continuously improved, expanded, and aligned with both technological innovation and business objectives. By coupling the strategic oversight of the ENOUGH business unit with the technical expertise of the SDS Owners, this model ensures that the platform remains adaptive, future-proof, and able to deliver sustained value to its users.

A first concrete materialization of this fourth business use case has been realized through the conveyance and valorization of research results previously developed in the framework of European projects, recontextualized into a completely new ground that has opened promising routes for medium- and long-term research. This step demonstrates how the SDS Marketplace can act as a catalyst for innovation by enabling the translation of past scientific achievements into practical applications embedded in the SDS ecosystem.

In this specific instance, the UNIVPM partner undertook the task of improving and integrating two strands of research that originated in earlier H2020 projects, reshaping them and combining them in order to generate novel outcomes. The first of these research strands concerned a methodology for applying mixed reality to the management of facilities, an approach that leverages immersive visualization and interaction technologies to support operators and decision-makers in navigating complex infrastructures. The second strand, by contrast, was rooted in the management of complexity in systems of systems, and in particular addressed the challenges of coordination, interoperability, and adaptive governance across heterogeneous subsystems.

The innovative step has been to merge these two methodologies into a unified decision-making application and methodological framework that blends the Holonic Management Tree (HMT)[37] approach with the affordances of Mixed Reality (MR)[38,39]. This integrated methodology, named **HMT+MR**, provides a powerful tool for visualizing, analyzing, and coordinating complex systems of systems, enabling human decision-makers to navigate multiple layers of interdependencies with enhanced clarity and situational awareness. By embedding the holonic paradigm within an immersive MR environment, it becomes possible not only to abstract complexity into hierarchical decision trees, but also to interact with these abstractions in a more intuitive and collaborative manner.

The connection of this new application with a Supply Chain Entity (SCE) arises naturally within the SDS platform. As demonstrated in Demo 1, the SDS provides the data backbone and integration mechanisms that allow the HMT+MR methodology to be instantiated within a living supply chain scenario. In this context, the support and contribution of TUGraz have been particularly significant, as their case study (Dairy Plant Participant in the SCE) offered a concrete, real-world environment where

---

[37] Pirani, M., Carbonari, A., Cucchiarelli, A., Giretti, A., & Spalazzi, L. (2024). The Meta Holonic Management Tree: review, steps, and roadmap to industrial Cybernetics 5.0. *Journal of Intelligent Manufacturing*, 1-42.

[38] Vaccarini, M., Spegni, F., Giretti, A., Pirani, M., & Carbonari, A. (2024). Interoperable mixed reality for facility management: a cyber-physical perspective. *Journal of Information Technology in Construction*, *29*, 573-595.

[39] Carbonari, A., Pirani, M., & Giretti, A. (2023). Leveraging BIM and mixed reality to actualize lean construction. In *Proceedings of the 31st Annual Conference of the International Group for Lean Construction (IGLC31)* (pp. 116-127). International Group for Lean Construction.

the methodology could be tested, verticalized, and refined. By applying HMT+MR to their owned supply chain scenario, it was possible to demonstrate the full sequence of operations, from data acquisition through SDS Connectors, to decision-making enhanced by the MR layer, and finally to collaborative optimization among participants.

This effort represents not only a technical demonstration but also a conceptual advancement. It shows how the SDS platform can serve as a bridge between theoretical research and practical implementation, turning methodologies that were originally conceived as stand-alone academic contributions into operational tools that create value for supply chain participants. Furthermore, it exemplifies how the SDS Marketplace can host applications that are both domain-specific and methodologically innovative, providing stakeholders with solutions that address real industrial needs while also advancing scientific inquiry.

In broader terms, the integration of HMT+MR within the SDS framework confirms the role of the platform as a fertile ground for the co-evolution of business applications and research methodologies. It highlights, not the least, how past European investments in research can be amplified and redirected towards new horizons, where sustainability, digitalization, and systemic management converge. The implications are particularly relevant for future supply chains, which will increasingly rely on such advanced decision-making aids to cope with complexity, uncertainty, and the pressing demand for environmental responsibility.

As the details and components of this methodology are profound and extensive, we restrict ourselves here to a concise exposition focused on the main topics. For a more in-depth treatment, the reader is referred to the scientific publications already available from the authors of this document, as well as those currently under development at the time of writing.

7.1 The Holonic perspective in supply chain Choreographies

The supply chain entity (SCE), as formalized in the choreography view (see Figure 4.4.1), represents the network of participants, tasks, and exchanges that sustain the flow of goods, information, and value across a collaborative ecosystem. While BPMN choreographies typically emphasize the sequencing and synchronization of message exchanges, they can be enriched by a systemic interpretation that places them in a recursive structure. This is achieved through their association with a holarchy tree, an organizational paradigm inspired by Arthur Koestler's concept of *holons*: entities that are simultaneously wholes in themselves and parts of larger systems. The structure of the self-similarity concept here applied is displayed in Figure 7.1.1.

The figure illustrates the structure of the **Holonic Management Tree (HMT)**, which provides a recursive framework for modeling complex systems of systems such as supply chain entities (SCEs). The tree is organized into multiple **levels of granularity** (levels 0 to 3 in this case), where the lowest level consists of **ground cells**, or tree leaves, representing the most elementary units of action, such as a transport operation, a machine cycle, or a sensor measurement. These leaves are progressively aggregated into higher-level holons, which embody larger subsystems (for example, a logistics hub, a production unit, or a retailer's operations), up to the top level, which represents the global behavior of the SCE. The self-similarity arises from imposing only four types of causal relationships between the parts: series, parallel, assembly, and expansion, which are typical layouts in production cells.

Two types of relations are represented in the model. **Tree relations** (inter-level), shown as vertical connections, capture the hierarchical dependency between levels, highlighting how local actions and performance indicators propagate upward to inform the behavior of the whole system. **Structural relations** (intra-level), represented by solid arrows, describe horizontal interactions among entities at the same level, such as coordination between two transporters or collaboration between plants.

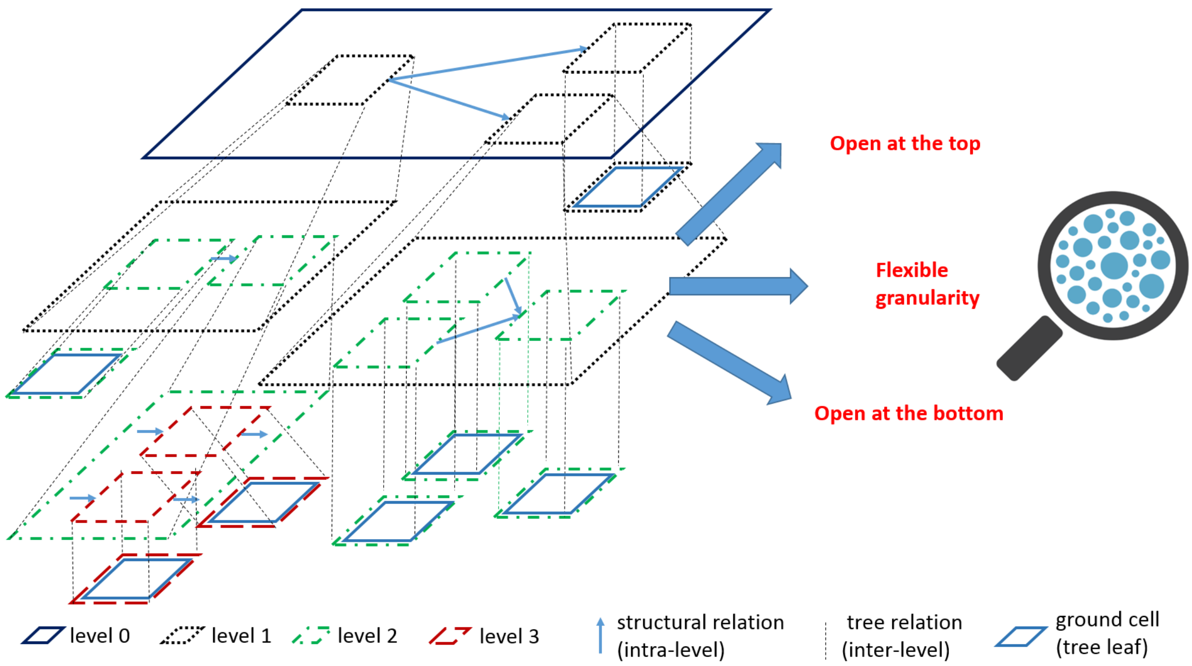

Figure 7.1.1: The self-similarity concept of HMT structures for systems of systems (Pirani et al., 2024)[40].

The model is characterized by three essential properties. It is open at the bottom, allowing the inclusion of new ground cells, such as new participants or processes. It exhibits flexible granularity, meaning that the system can be analyzed either at fine detail (leaf level) or at aggregated views (higher levels) depending on managerial needs. Finally, it is open at the top, enabling the integration of higher-order layers, such as multi-supply-chain ecosystems or strategic networks.

The magnifying glass on the right emphasizes the importance of the ground cells as the fundamental measurement points where performance indicators, such as CO_2 savings, efficiency, or service levels, are collected. These values feed the recursive structure of the tree, ensuring that local improvements or degradations can be assessed in terms of their impact on the global system. The figure conveys how the Holonic Management Tree enables a scalable and adaptive representation of supply chains, **linking the micro-level of operational data with the macro-level of system-wide decision-making**.

In the holonic representation, each node in the choreography—from a producer sending data, to a transporter negotiating shipments, to a retailer confirming deliveries—can be mapped as a holon. These holons are arranged hierarchically in a tree that reflects the recursive structure of the supply chain. At the lower levels, the leaves of the tree represent atomic tasks or measurable events, such as

[40] Pirani, M., Carbonari, A., Cucchiarelli, A., Giretti, A., & Spalazzi, L. (2024). The Meta Holonic Management Tree: review, steps, and roadmap to industrial Cybernetics 5.0. *Journal of Intelligent Manufacturing*, 1-42.

the completion of a transport operation, the maintenance of a refrigeration unit, or the execution of a quality inspection. At higher levels, intermediate nodes aggregate these results into broader process indicators, culminating in the root, which embodies the overall performance and sustainability posture of the entire SCE.

This dual representation offers a powerful cognitive shift. Instead of interpreting the choreography as a flat diagram of sequential steps, the holarchy allows stakeholders to appreciate its recursive depth. Each element is simultaneously accountable for its own performance and for its contribution to the system above it. This framing is particularly valuable in contexts such as sustainability-driven supply chains, where local improvements must be assessed not only in isolation but also in their cumulative systemic impact.

The structure of the HMT can be directly connected to the way we represent a Supply Chain Entity (SCE) through BPMN choreographies. In the choreography, we see participants such as Retailer, Coolstore, Dairy Plant, and Transporters exchanging structured messages to coordinate orders, shipments, and negotiations. Each of these participants and their tasks corresponds to leaves or holons in the tree structure.

At the ground cell level, every task in the choreography (for example, “Retailer sends order,” “Coolstore proposes alternative,” or “Transporter confirms delivery”) produces measurable outputs, such as quantities delivered, CO_2 savings, or penalties for inefficiencies. These are the fundamental data points — the “leaves” of the tree.

Moving up the hierarchy, these leaves are aggregated into higher-level holons that reflect the performance of an entire participant. For instance, all actions performed by the Dairy Plant (order preparation, negotiation, dispatching) are grouped to form the Dairy Plant’s contribution to sustainability and efficiency. Similarly, the Coolstore’s operations are aggregated to represent its systemic impact.

The tree relations (inter-level) illustrate how local actions in the choreography propagate upwards. A small efficiency improvement in transport planning (leaf level) contributes to better performance of the Transporter holon (level 2), which then affects the overall coordination of the SCE (level 1), and eventually impacts the global sustainability metrics of the supply chain (level 0).

The structural relations (intra-level) mirror what happens inside the choreography when multiple actors interact at the same level. For example, negotiations between Coolstore and Retailer, or parallel operations of two Transporters, are captured as horizontal interactions among holons that need to coordinate to avoid conflicts and optimize outcomes.

Crucially, the properties of the HMT — open at the bottom, flexible granularity, and open at the top — match the dynamics of SCE choreographies. The openness at the bottom allows new participants (e.g., another supplier or transporter) to be added seamlessly into the choreography but also a reurrence to more details (increase level of detail) with the deepening of one level. Flexible granularity supports decision-making at different levels: the system manager may zoom into a single transaction between Retailer and Coolstore or zoom out to see cumulative performance across the entire chain. Openness at the top provides the possibility to integrate one SCE choreography into a broader ecosystem of supply chains, enabling comparisons or cooperation across industries.

In practice, this connection means that the decision-making tool based on HMT+MR (Holonic Management Tree plus Mixed Reality) can measure, predict, and assess the effects of each local action in the choreography, not in isolation but as part of the recursive system structure. Thus, when the Retailer negotiates delivery dates or when a Transporter optimizes its route, the impact of these actions is immediately visible not only locally but also on the global indicators of the supply chain, such as total GHG savings.

The recursive nature of the holarchy ensures that every improvement at the leaves of the tree has a cascading effect on the whole SCE. For instance, consider a logistics operator who optimizes the routing of deliveries, reducing fuel consumption by 10%. At the local level, this translates into a measurable gain in CO_2 savings, which can be recorded as a KPI in the decision-making tool. However, the benefit does not remain confined to the transport holon. It propagates upward: the distribution process node shows enhanced performance, which in turn improves the aggregated environmental footprint of the supply chain segment, ultimately contributing to the global sustainability metrics of the SCE.

This ripple effect works both positively and negatively. A successful improvement in the leaves can lift the entire structure, while inefficiencies or failures at the local level can weaken the system as a whole. For example, if the dairy plant in the choreography fails to meet refrigeration standards, the deterioration not only affects product quality locally but also impacts the retailer's ability to deliver fresh goods, undermining customer satisfaction and eroding the SCE's global reputation. The holarchy makes explicit this interdependence, reminding participants that systemic health depends on the alignment and continuous improvement of all its components.

Importantly, this perspective supports both short-term and long-term improvement strategies. Short-term gains can be achieved through immediate local actions—such as adjusting shipment frequency, reducing idle times, or optimizing machine duty cycles—while medium- and long-term gains emerge from the accumulation of iterative improvements across multiple sessions of the SCE. Over time, as the history of interactions is recorded and archived, patterns of successful practices can be identified and generalized, enabling the design of better processes and even new business models.

7.2 Decision-making as a cybernetic process

Cybernetics is the interdisciplinary study of systems of control and communication in machines, living beings, and organizations, focusing on feedback mechanisms that regulate behavior and achieve goals. We refer the reader to some literature for a more profound definition and its implications[41].

To operationalize this recursive view, the HMT+MR decision-making tool developed in the project integrates measurement, prediction, and assessment capabilities. It continuously monitors the performance of leaves, calculating local KPIs such as GHG savings, energy efficiency, or delivery accuracy. These results are then propagated upward, aggregated, and normalized along the holarchy, ultimately producing a synthetic assessment of the entire SCE.

The novelty of the tool lies in its ability to not only report but also **predict** the systemic effects of local decisions. When a participant takes a local action—say, adopting a new compensatory measure, introducing a different vehicle for transport, or adjusting production schedules—the tool can estimate its impact at higher levels of the tree. This predictive function is crucial in avoiding sub-optimizations: improvements that look beneficial locally but create hidden trade-offs globally. For example, increasing production throughput at a plant might reduce unit costs but overload the transport system, leading to delays and increased emissions. By simulating the propagation of actions through the

[41] Pirani, M., Carbonari, A., Cucchiarelli, A., Giretti, A., & Spalazzi, L. (2024). The Meta Holonic Management Tree: review, steps, and roadmap to industrial Cybernetics 5.0. *Journal of Intelligent Manufacturing*, 1-42.

holarchy, the tool reveals these trade-offs in advance, allowing participants to recalibrate their strategies.

From a theoretical standpoint, this capability resonates with Ashby's *law of requisite variety*[42], a cornerstone of cybernetics. Complex systems can only be effectively regulated if the regulatory mechanisms embody a level of variety commensurate with that of the system itself. The holonic decision-making tool embodies this principle: by modeling the recursive complexity of the supply chain, it provides the necessary variety to match the dynamic interactions among participants. This transforms the SCE into a cybernetic system where feedback loops are made explicit, measurable, and actionable.

The integration of the holarchy with the decision-making tool thus provides both a conceptual and practical framework for continuous improvement. On the conceptual side, it shifts the mental model of supply chains from linear sequences to recursive systems of systems (confront Figure 2.2.1 and Figure 2.2.2), where each participant is simultaneously autonomous and interdependent. On the practical side, it delivers a digital instrument that collects, processes, and visualizes performance data, enabling stakeholders to test, validate, and refine their strategies in an evidence-based manner.

In practice, this means that Participants in the SCE Choreography—such as the Retailer, Dairy Plant, Coolstore, and transporters—can see how their local behaviors contribute to or detract from the global objectives of the SCE. They can experiment with alternative actions, observe their predicted and actual effects, and engage in negotiations with other participants to align their strategies. Over time, the tool supports the emergence of cooperative behaviors and shared norms, fostering a culture of collaborative problem-solving.

This approach also aligns with the broader goals of sustainability and Industry 5.0. By embedding systemic awareness into everyday decision-making, it ensures that improvements are not pursued at the expense of others but are harmonized with collective goals such as CO_2 reduction, energy efficiency, and resilience. It also creates an infrastructure for learning across sessions: each iteration of the SCE produces data that can be mined for lessons learned, supporting long-term innovation.

The mapping of SCE Choreography into a holarchy tree, coupled with the decision-making tool, represents a breakthrough in supply chain management. It operationalizes the principle that every local action has systemic consequences, providing stakeholders with the means to visualize, measure, and predict these effects. By doing so, it fosters continuous improvement both at the local and global levels, ensuring that private goals of participants remain aligned with the collective objectives of the supply chain.

Perhaps the most significant contribution of this approach is its capacity to reconcile autonomy and collaboration. Participants retain their local decision-making power but operate within a framework that makes explicit their systemic responsibilities. This **balance between private optimization and global coordination** is precisely what is needed to drive sustainable, resilient, and adaptive supply chains in an era where complexity and uncertainty are the norm.

[42] Pirani, M., Cucchiarelli, A., Naeem, T., & Spalazzi, L. (2025). A Blockchain-Driven Cyber-Systemic Approach to Hybrid Reality. *Systems*, *13*(4), 294.

7.3 HMT+MR applied to Dairy Plant

In this experiment, we had the opportunity to explore the potential of the HMT+MR methodology in the context of one of the case studies made available by the partners involved in the work of WP5. In particular, the partner TUGraz kindly provided us with access to their Dairy Plant Participant within the demonstrative SCE. This case acted as a strawman example through which we could verify, in practice, both the adaptability and the prospective usefulness of the methodology in real supply chain contexts.

The narrative scenario we constructed can be described as follows. A decision maker at the Dairy Plant—whether an individual manager or a dedicated team—takes an active role within the framework of continuous improvement of the SCE_D5.3, as introduced in Section 4. This decision maker, while engaging in negotiations with the Coolstore Participant, does not act passively, but instead produces internal forecasts that are informed by a dual modeling perspective: on the one hand, a representation of the internal dynamics of the Dairy Plant's own processes, and on the other hand, an understanding of the activities of the other participants in the SCE.

This modeling effort is not abstract but grounded in concrete informational flows. The decision maker considers all messages exchanged within the choreography, as well as the historical record of such exchanges, which is readily accessible through the dashboards of the SDS system. This holistic information basis enables the decision maker to build a systemic vision of the supply chain dynamics in which they are embedded.

On this foundation, the decision maker proceeds to create the HMT model. The SCE is decomposed into a holonic tree structure, in which each node represents a meaningful subsystem or activity, while the leaves capture more granular operational processes. Alongside this decomposition, key performance indicators (KPIs) are defined—in this case with particular emphasis on GHG savings, given their centrality for the ENOUGH project. Through this process, the decision maker constructs a model that reflects both the complexity and the interdependence of the supply chain reality.

It is important to emphasize that this process of association, decomposition, and modeling of the SCE through the HMT approach does not follow a single "correct" path. On the contrary, it can be performed in a variety of ways depending on the perspective, objectives, and domain knowledge of the decision maker. This flexibility embodies a principle of what may be described as "pragmatic constructivism"[43]: the system's representation is not a fixed truth but a purposeful construction, designed to serve the improvement of both local and global performance. This holds whathever the agents that construct the models of reality are humans or machines[44].

What is particularly significant is that, regardless of the specific shape that the HMT model takes, its inherent mechanisms guarantee that the resulting framework will guide toward overall improvement of the supply chain. Local improvements achieved at the leaves of the holarchy cascade upwards to higher levels, ultimately producing systemic gains across the SCE as a whole. The synergy between local optimization and global coherence is precisely the added value of the HMT methodology, especially when supported by the immersion and visualization possibilities offered by Mixed Reality.

Taken together, this case illustrates how HMT+MR can enable decision makers to move beyond fragmented or short-term reasoning. By combining cybernetic principles, systemic decomposition, and

[43] Lissack, M. (2019). Yes, we can, and do, design our understanding: The roles of ascribed coherence and ascribed realism in our sense-making. *She Ji: The Journal of Design, Economics, and Innovation*, *5*(4), 401-419.

[44] Raikov, A. N., & Pirani, M. (2022). Human-machine duality: What's next in cognitive aspects of artificial intelligence?. *IEEE Access*, *10*, 56296-56315.

immersive interaction with data and models, the methodology fosters both a pragmatic orientation to action and a scientific rigor in assessing the consequences of local decisions on the broader choreography of the SCE. This is why we argue that HMT+MR holds the potential to become a foundational decision-making tool in future digitalized and sustainable supply chains.

In Figure 7.3.1 we show the start of a decomposition of the SCE_D5.3 Choreography after two decompostion steps have been performed by the decision maker. As a notation, the elliptical elements are a node of the HMT, and the letter besides the name indicates the kind of structure chosen (p=parallel, s=series, a=assembly, e=expansion). Thus in this figure, we have decomposed the SCE into a series, then downwards into another series P1s and a leaf P2. The P1s has been decomposed into two leaves P1.1 and P1.2. Red rectangles denote leaves of the HMT.

By continuing with this decomposition we arrive to a point in which there is the switch between the public level and the private level for the Dairy Plant.

We should emphasize, before proceeding further, that the decomposition process within the construction of the HMR is not carried out arbitrarily, but rather is triggered whenever a performance bottleneck is identified at a given step of the model. The rationale is that the model should remain as compact and parsimonious as possible, while still retaining explanatory and operational power. Only when strictly necessary does a leaf node undergo decomposition, transforming into a refined branch.

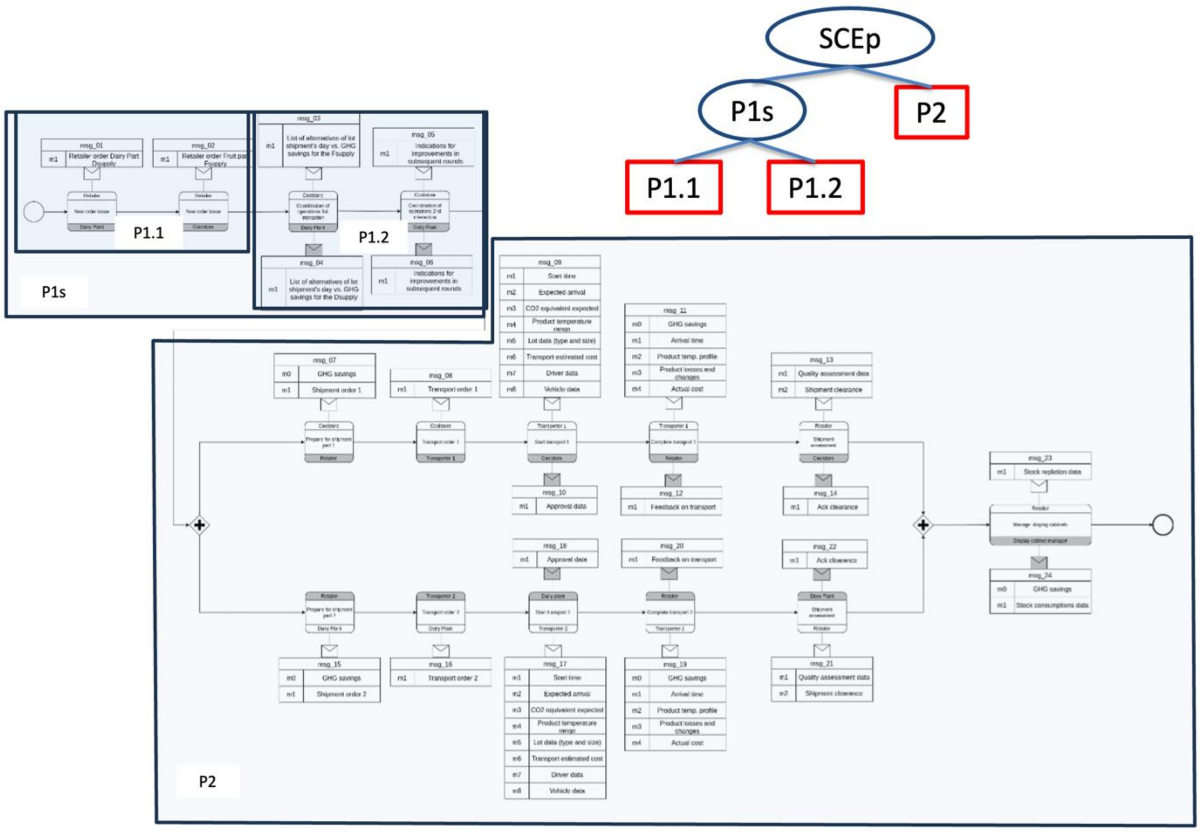

Figure 7.3.1: First two steps of the HMT decomposition.

This refinement enables the decision-maker to capture greater detail in the representation of the system, with the explicit purpose of identifying a feasible and effective action to resolve the observed bottleneck. In this sense, decomposition is both a diagnostic and prescriptive mechanism: diagnostic, because it helps localize where the constraint is occurring, and prescriptive, because it offers a structured pathway toward designing interventions that can improve the overall performance of the system[45].

After the system's bottleneck drive into a focus for the Dairy Plant Participant, there is a switch between the public and the private level, as from that point on, the modelling will depend only on private information and knowledge the decision maker has of it internal processes and plants. we could arrive into such situation as shown in Figure 7.3.2.

Before going deeper into this detail, note that up to this point the decision making could have been a collective process, as we stayed at the Choreography, and so public, level of the SCE, where message information is shared between all the interested Participants. So, decision making is distributed or collective up to here. From this moment on, we will dive into the local internal process of the Diary Plant Participant. The HMT can prosecute, but the decision making is now proprietary and hidden for the Participant. Nevertheless, any improvement performed in the inside, can have a good effect on the whole. Moreover, the predictions that are enabled by the HMT methodology, allows the Participant to assess the impact of its internal choices to the whole SCE.

Being the next step so an important discontinuity between the public choreography level and the private level, we have marked it in red in Figure 7.3.2.

At this point the ball is in the hands of the manager of the Dairy Plant. She has to find a way to improve their internal processes in order to minimize their m0(msg_04) or, in other words, their GHG savings amount.

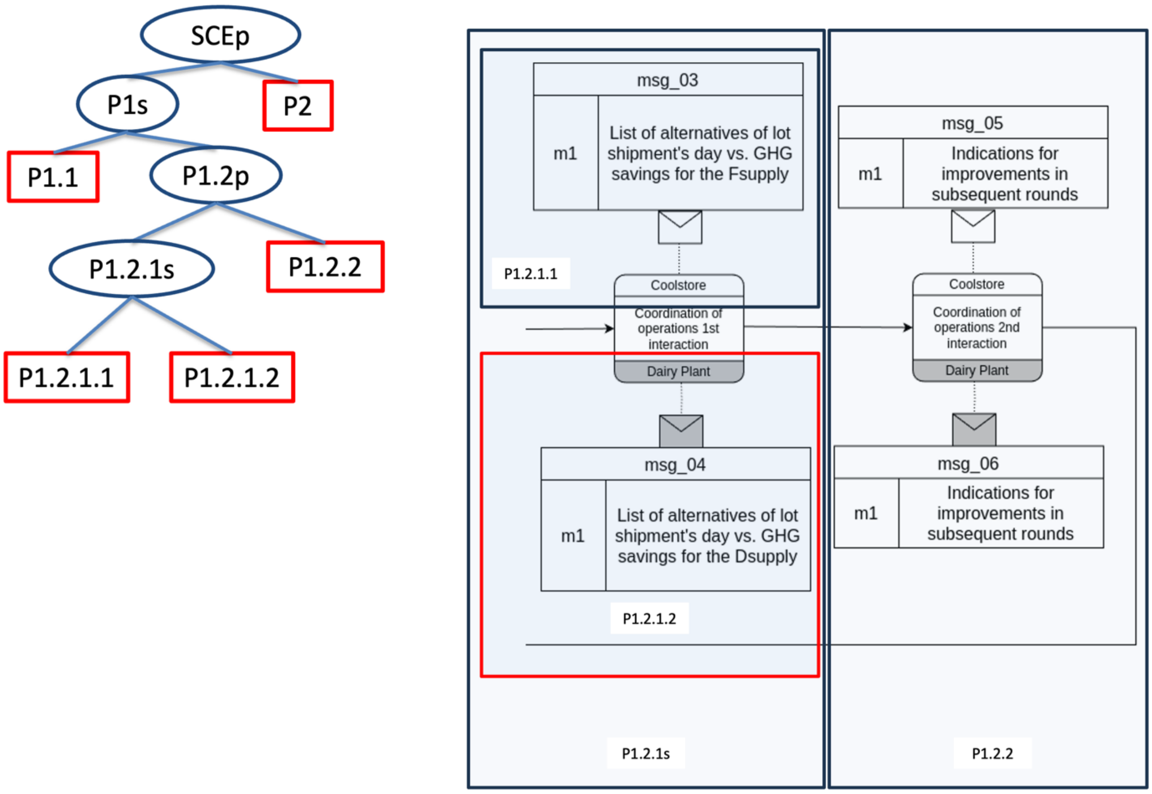

Figure 7.3.2: The HMT starts focusing the Dairy Plant.

[45] Pirani, M., Carbonari, A., Cucchiarelli, A., Giretti, A., & Spalazzi, L. (2024). The Meta Holonic Management Tree: review, steps, and roadmap to industrial Cybernetics 5.0. *Journal of Intelligent Manufacturing*, 1-42.

Nonetheless, an interpretation problem is soon raised by the Dairy Plant manager. The very nature of Activity 03 and 04 is to provide a "plan of the GHG savings" in order to be the most aligned. The narration was (for SCE_D5.3) that they had to provide each other with a profile of their own GHG savings prediction during the next week. They use this prediction to improve their and joint performance by trying to match temporally the two GHG savings profiles: the more they are aligned, the best for the SCE. Thus, we have here two dimensions: one is the absolute levels of GHG savings, the other is the temporal alignment. This is the beauty (in the eye of the beholder) of the HMT, as it can enter a refinement phase in order to substitute this model with a more suitable one. Actually, the managers of both Participants (Dairy Plant and Coolstore) think that the assembly structure could fit better than a series. The assembly structure has the "head" cell that is a new KPI that usually has the semantics of how effectively an assembly of two improvement actions is done.

In this case that new KPI could measure and target the alignment. If we name this cell as P1.2.1h, the KPI for P1.2.1h could be expressed as a function of performance depending on a distance—to keep this discussion short and readable, we refer the reader to Pirani et al. (2024)[46] and the references therein for the exact and complete definitions of OEE and Peff.:

$$OEE(P1.2.1h) = Peff(P1.2.1h) = -\frac{1-\varepsilon}{d_{max}}|d| + 1 \qquad (1)$$

where d is the distance between the maximum value of the profiles of GHG saving, and d_{max} is the maximum distance in days between these peaks. When the maxima are aligned d=0 and OEE(P1.2.1h)=1, and when d=d_{max} OEE(P1.2.1h)=ε.

This means a nudge to collaboration to let the process be tailored in order to make d=0. Of course, this can be done by the action of one or the other Participant or both. In our example we imagine that the Coolstore finds no easy way to do that, and so the Diary Plant strives in order to reach this shared result, or at least they try first.

We remind here that, to do this joint negotiation, the two Participants have request to the SDS platform to produce a Web page where this negotiation can happen quickly (because they do not want to mess with this task). In Figure 7.3.3 we report a snapshot of that provision from the SDS framework. From the knowledge of the best day, and knowing their current predictions, the Diary Plant manager tries to give a look inside their process to see if they can do something.

In Figure 7.3.4 we show the assembly structure pictorially, while in the Figure 7.3.5 the refinements with respect to previous Figure 7.3.2, are made. In particular, the green element stands there for the head cell of the assembly, which is a new focus and goal for the Diary Plant manager.

At this point we will begin the descent into the branch of the HMT that will relate exclusively to the Dairy Plant process manager.

[46] Pirani, M., Carbonari, A., Cucchiarelli, A., Giretti, A., & Spalazzi, L. (2024). The Meta Holonic Management Tree: review, steps, and roadmap to industrial Cybernetics 5.0. *Journal of Intelligent Manufacturing*, 1-42.

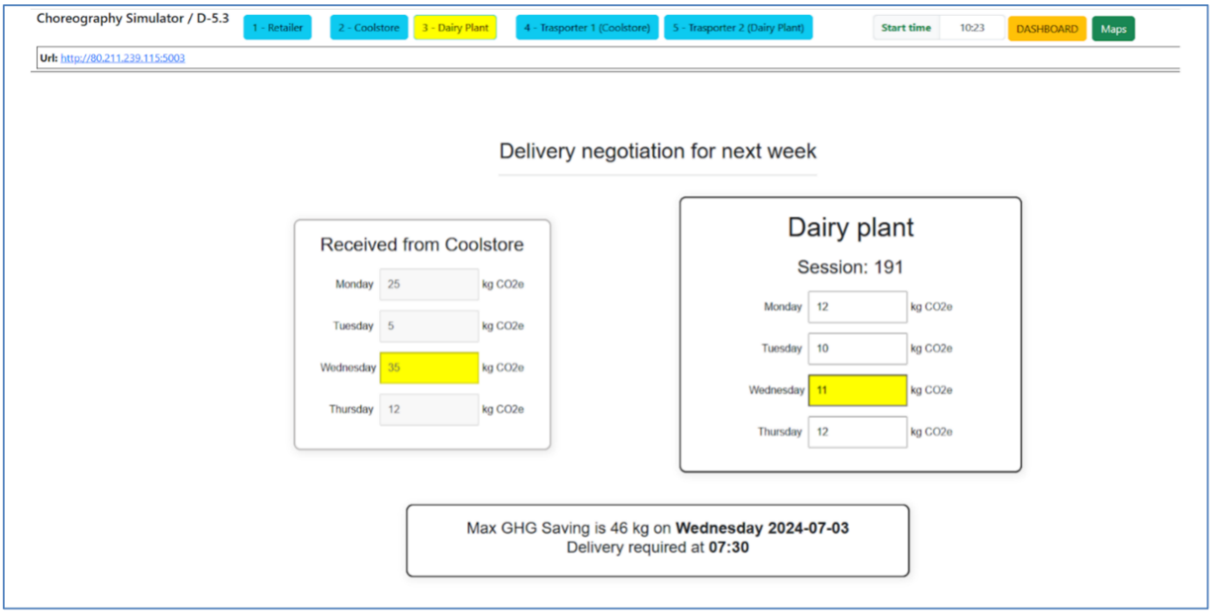

Figure 7.3.3: Joint decision on best day for maximum GHS savings.

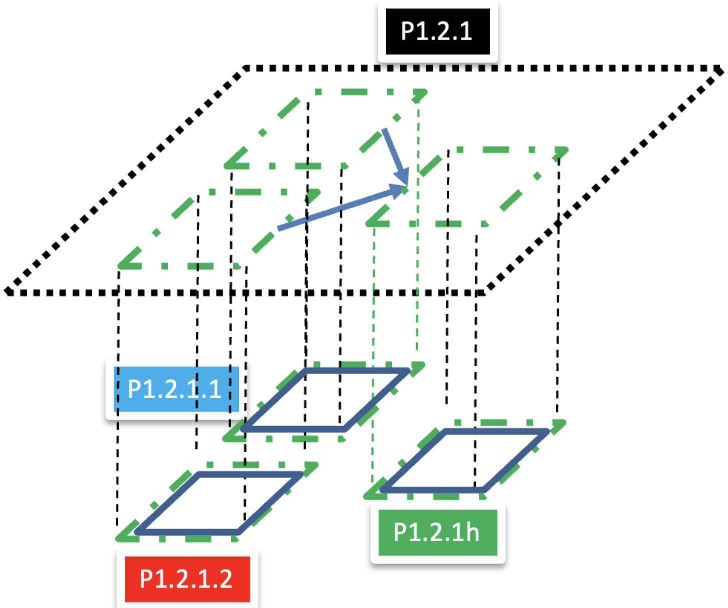

Figure 7.3.4: Pictorial description of the assembly structure decomposition of P1.2.1. P1.2.1h is called the "head" of the assembly.

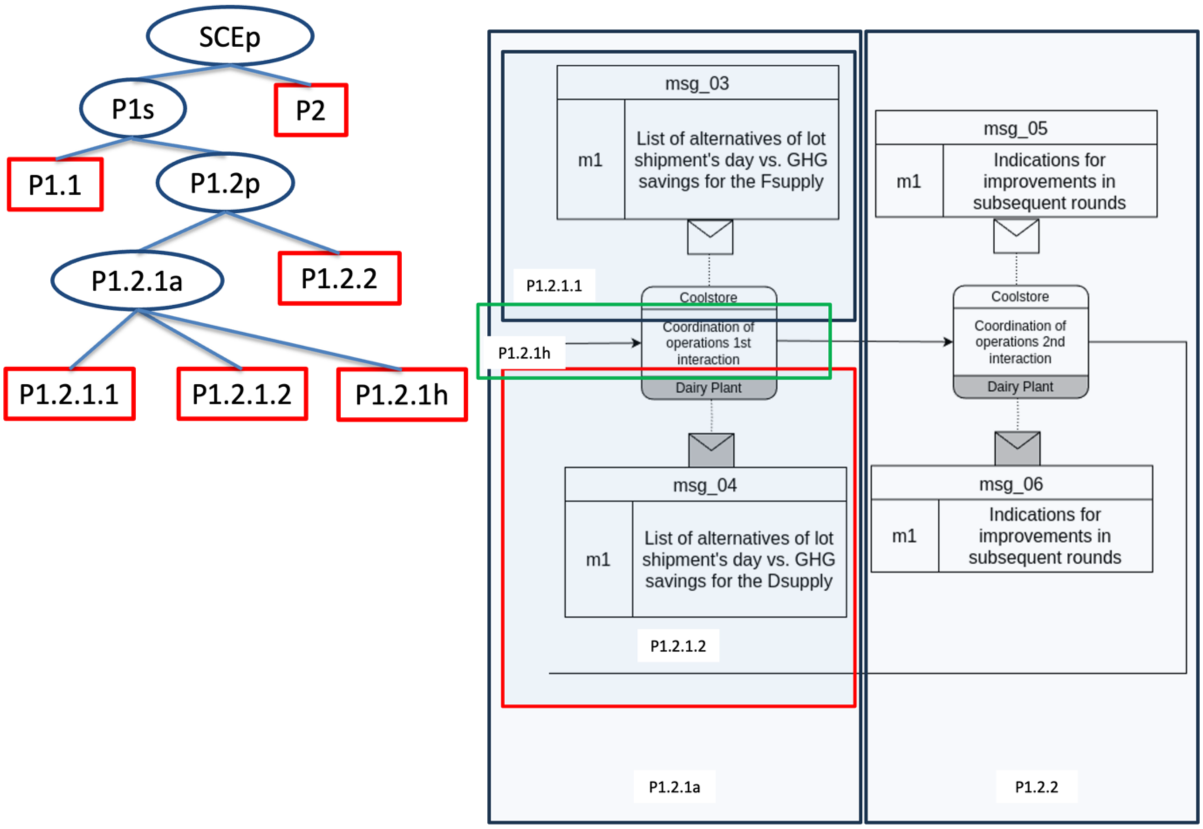

Figure 7.3.5: Last refinement in the HMT of the SCE Choreography, at the public level.

This part, thanks to the information and available facilities and environments (conveyed by TUGraz), will allow us to better express the MR part of the experiment as well. While HMT is applicable in general even without the use of MR, the latter component acquires fundamental value for HMT because it involves the physical presence of the manager in the environment in which the actions take place.

The Diary Plant manager has now to face an improvement on its process. She committed to the SCE and sees a whole picture now, to which she must abide for the collective and individual good at the same time.

The problem at hand is to act in two performance indicators. The first is related to P1.2.1.2, where any increment in GHG savings would provide a plus to the whole supply chain system. The second is related to P1.2.1h, which means being aligned with the temporal profile of GHG saving of another Participant, in order to achieve a joint action that maximizes GHG savings, at least for the products that will be produced and sold in the context of the SCE, and leave others for other sales channels (in the short term).

To do this, the manager might want to use the HMT+MR methodology as we will envision in this discussion.

In Figure 7.3.6, it is shown “the knowledge” of the manager on the Plant. The details on the goals and the technology of this plant have been provided in the context of the activities of Demo 3 in WP6.

Here, for the current purpose, we only recall briefly some of the notions.

The most important focus for this plant is to use the HTHP technology to achieve GHG savings. The HTHP system raises waste heat of the cooling system of the dairy production process to temperature levels suitable for process heat supply. This energy is then accumulated for its reuse. The primary mode here available is a storage and stock of hot water. The hot water storage will be used to heat up cleaning fluids for the CIP (cleaning in place) processes of e.g. milk transport trucks and production equipment. Usually this heating consumes natural gas. In this case, the use of natural gas is avoided, partially or completely, because of the hot water available in the storage.

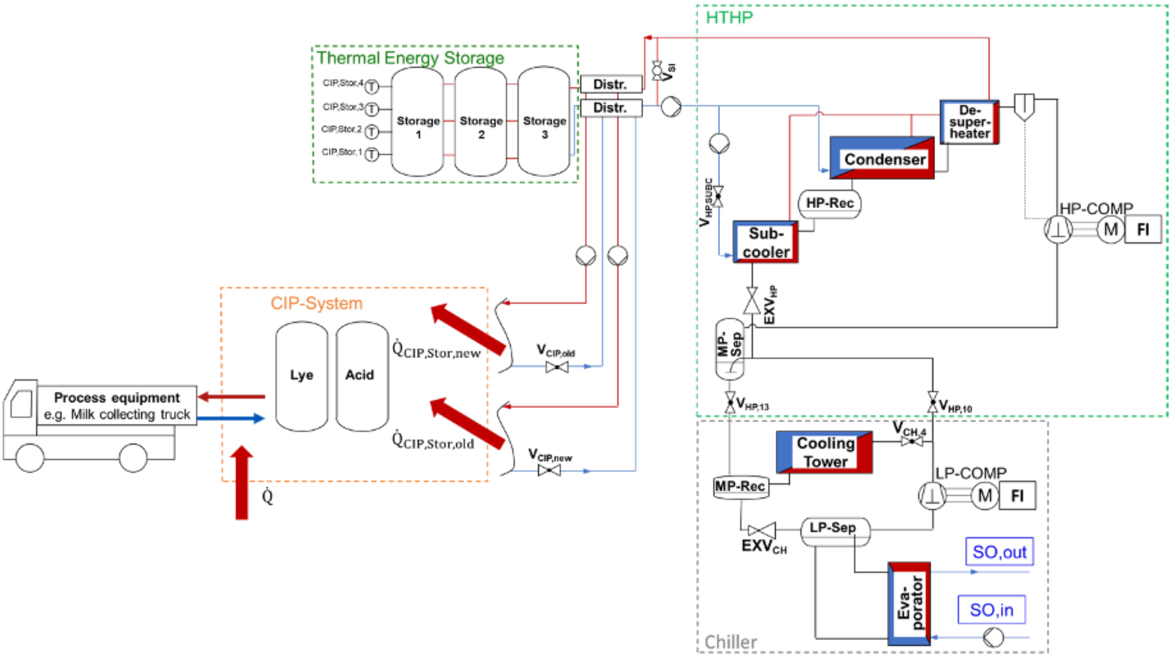

Figure 7.3.6: A scheme of the Diary Plant main components.

In this particular and current scenario, the manager knowns that:

1. Nothing can be done to control the amount of waste heat and periodicity of the production process at the moment. No control to the production is foreseen in the short-term neither possible.
2. The HTHP is quite state-of-the-art system too, and a very few tweaks are still possible to optimize, at least in the short term. The only degree of freedom for improvement at the moment is checking that the HTHP systems works as expected with mostly an improvement made by means of maintenance and calibration operations.
3. In the short term, the parameters of the thermal energy storage rig can be tweaked a bit, but no new tanks can be provided, so capacity is fixed. Anyway, this part is the actual driver of the HTHP system. The GHG saving is active only when the HTHP is on and activated by the sensors and the control in the tanks.
4. The heating demand of the CIP-system is determined by the acutual production processes and depends on different aspects, such as used equipment and specific production processes, which makes it very complex. The consumption of hot water depends on the arrivals of the vehicles and production steps carried out, and at the moment data of arrivals do not provide much but some regularity for some prediction. In any case, the CIP activity is the ultimate

driver of the GHG savings[47]. And the temporal profile of the arrivals and production steps demanding CIP is the only GHG savings profile prediction in first approximation.

Having the former scenario, the manager starts to apply the HMT methodology at the private Dairy Plant level.

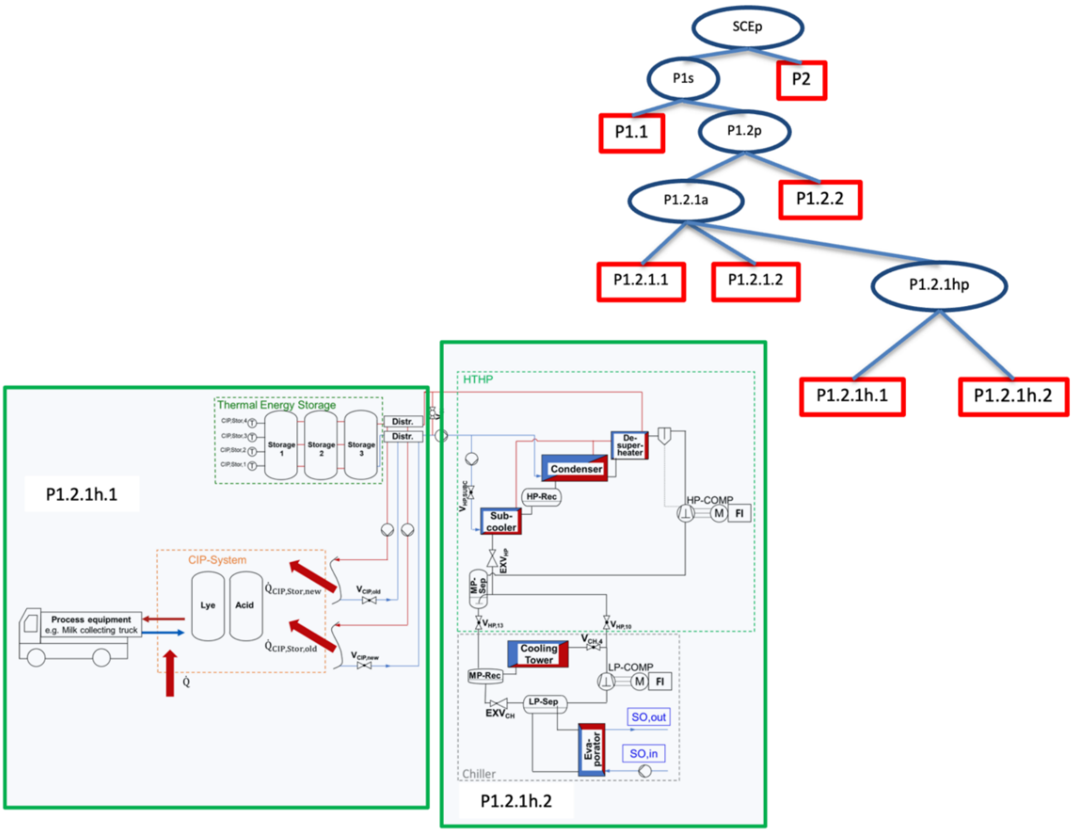

Figure 7.3.7: First partitioning of the Diary Plant, under the temporal profile alignment objective.

A first partitioning is made as in Figure 7.3.7. The manager focuses first the actions that can benefit the temporal alignment of the GHG savings between Coolstore and Diary Plant during the incoming week (based on predictions). This was accounted by the KPI defined for P1.2.1h.

To achieve something that can model the way in which this objective will be reached, the manager thinks that there will be a parallel relationship between the Partition that contains the production process and the HTHP on one side (P1.2.1h.2), and the energy storage and the CIP on the other (P1.2.1h.1). This is motivated by the fact that there is no clear trade-off between the HTHP part and the storage part for what is concerning the alignment of the GHG saving. So, simply the P1.2.1h.2 will be considered not so relevant for this objective (up to current knowledge), while the focus clearly goes to P1.2.1h.1 where some actions can be made to plan the demand of renewable energy during the

[47] This statement requires contextualization in order to avoid misleading interpretations. If there were no CIP processes, there would be no heat demand and thus no emissions. Of course, covering heat demand with the HTHP saves emissions and the temperature in the Storage tanks control the HTHP. However, we need to take care that heat delivered by the HTHP is not automatically seen as savings, only if gas demand can be reduced (otherwise, maximising heat demand and covering it with the HTHP would lead to maximum savings).

week. Analogously, the manager thinks that another partition will focus more on CIP rather that the storge part and then achieve P1.2.1h.1.1 for the storage and P1.2.1h.1.2 for the CIP, as in Figure 7.3.8.

Considering this parallel structural relation, manager might want to act primarily where something is possible to be performed to improve the situation in the short term.

Currently simulations and predictions about the intensity of CIP activity are modelled and available from the Dairy Plant. This information can be used. It will be processed internally by the Participant without the need to upload these data or models in the HMT+MR node. In case an automatic routine can produce the prediction of the peak this data is transmitted to the HMT+MR as an acquisition for the update of the KPI for P1.2.1h.1.2. It is done with the usual SDS connector system.

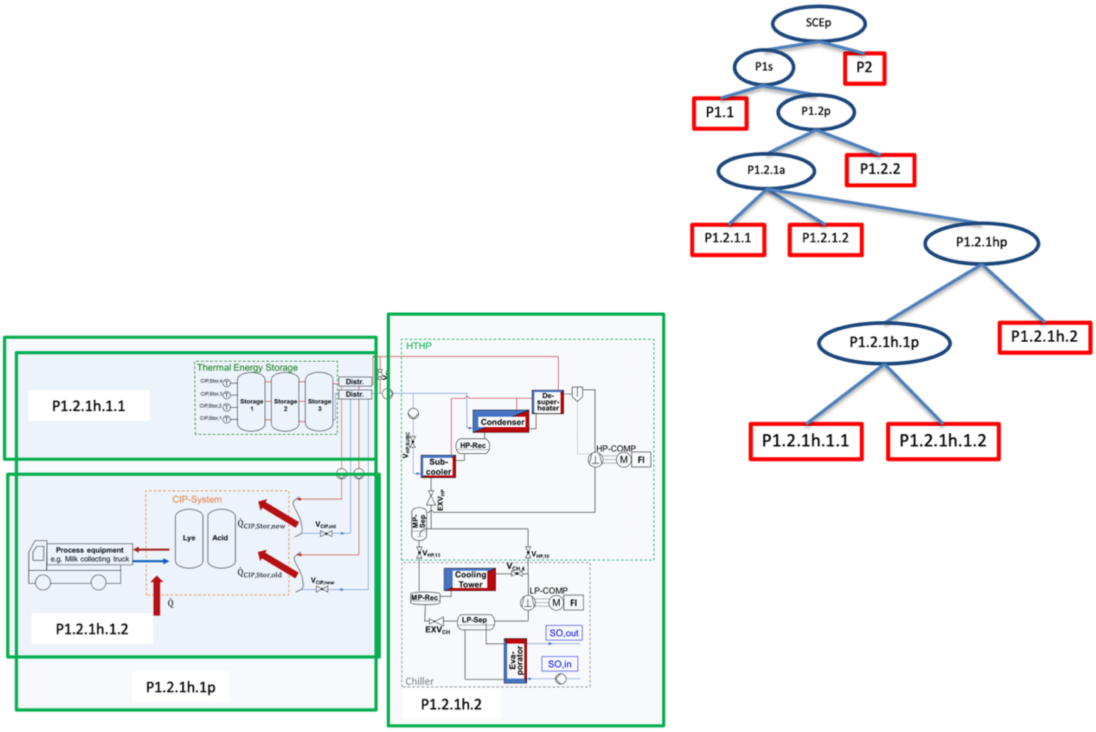

Figure 7.3.8: Second step of partitioning of the Diary Plant, under the temporal profile alignment objective.

Having coped with the partition P1.2.1hp, now it is the time to deal with P1.2.1.2. The goal semantics of this Partition is: to raise as much as possible GHG savings towards their targeted values.

For this step the manager sees a series structure P1.2.1.2s as in Figure 7.3.9, which will be divided again in P1.2.1.2.1 and P1.2.1.2.2 as before.

This is motivated by the fact that if the HTHP part is not working for some reason, the whole storage of energy fails to provide GHG savings, and so the sustainable CIP. Also, in the opposite direction, if the performance of the CIP and storage partition is low, there will be no use to act on the HTHP partition to improve the savings as the HTHP is driven by the CIP demand frequency: the higher the frequency of truck arrivals the more the HTHP will be activated to recover the heat waste of the production plant.

This time, the next refinement would involve the partition of the HTHP and the production, as not much can be done to improve the rate of GHG savings in the CIP and energy storage part. The focus is now on the efficiency of the machineries that can be affected by some malfunctioning or loss of efficiency during times. The new partitioning shown in Figure 7.3.10.

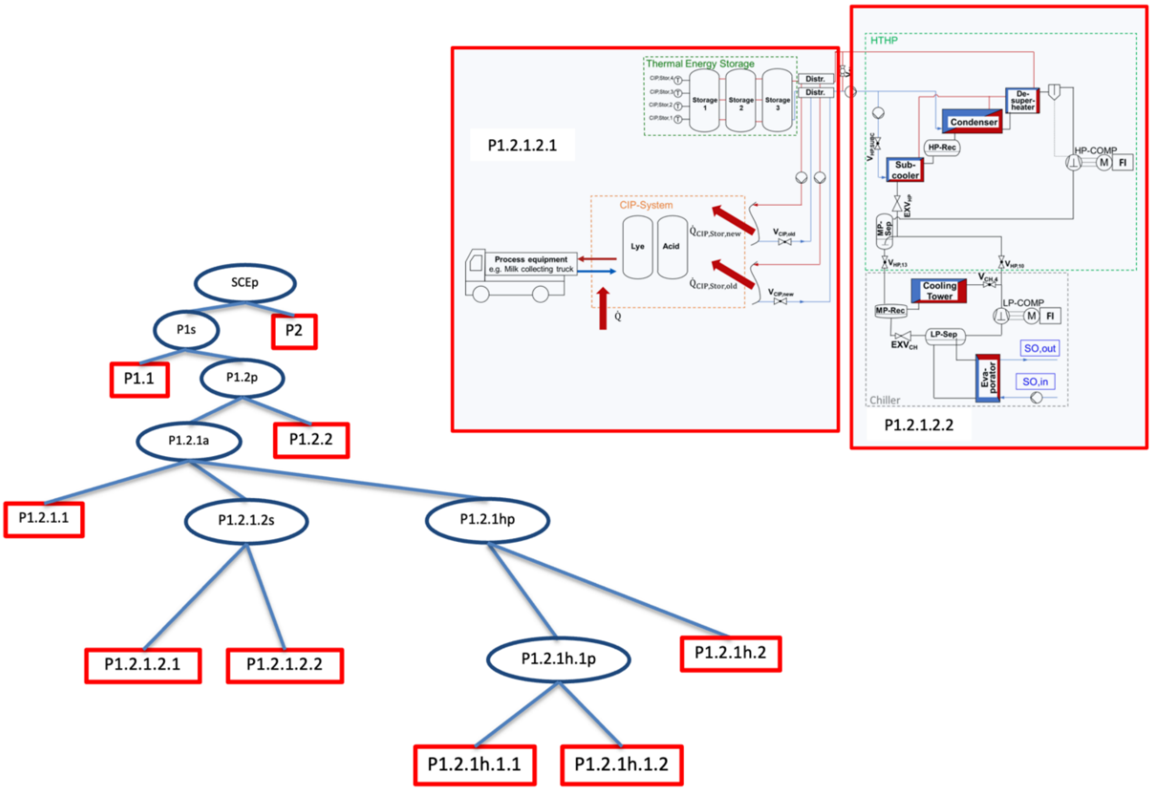

Figure 7.3.9: First partitioning of the Diary Plant, under the objective of absolute GHG savings.

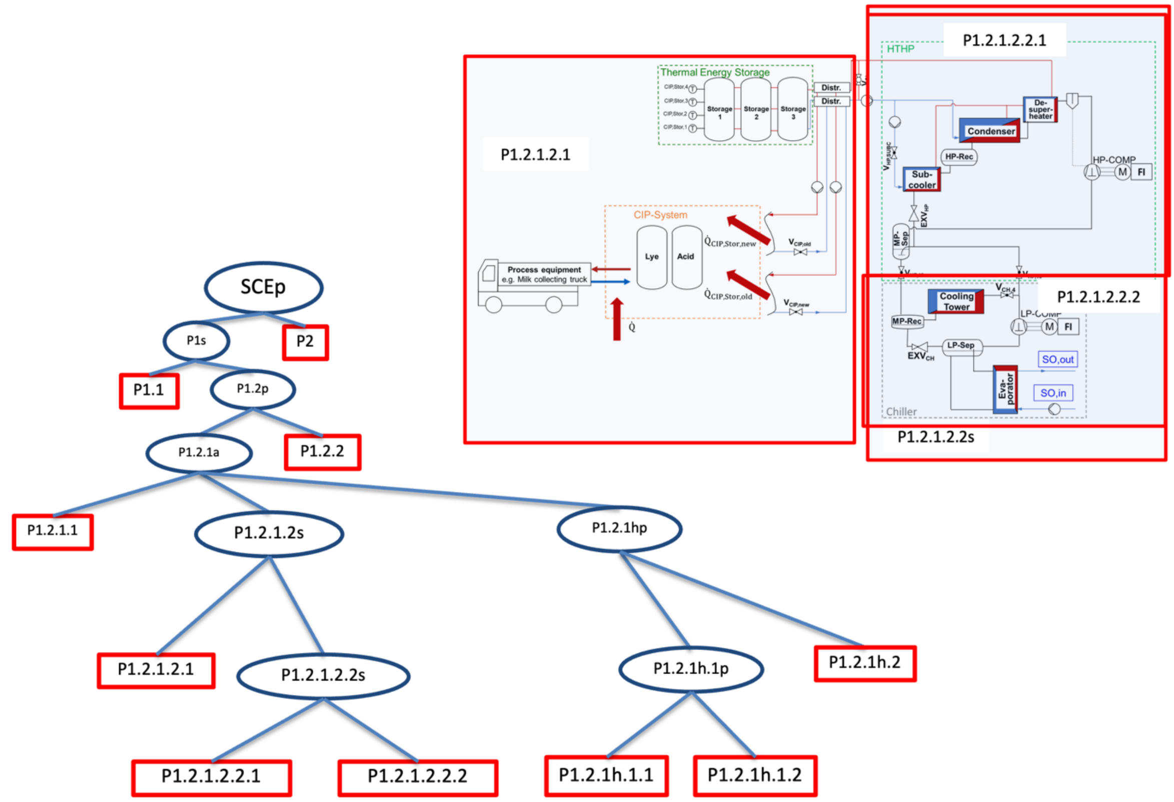

Figure 7.3.10: Second step of partitioning of the Diary Plant, under the objective of absolute GHG savings.

To conclude the discussion on the experimental setting for HMT+MR in this case, we note that even if the real measurements are not available, and no real long session can be performed during the project in the real plant with data acquisition, the whole procedure can be tested and validated by the use of the reduced and detailed models available. TUGraz provides models of the significant parts of the plant with different levels of detail . The detailed model (e.g. a Dymola model of the HTHP process) will take the role of a reality and the simplified one of an assessment by measurements or simulations.

All these models are being calibrated with real data coming from some experimental campaign. Data set is already available.

These models feed the KPIs by means of the SDS connectors and are used by the HMT+MR node that is acting and connected as any other Actor in the SDS infrastructure.

There is also to note that the HMT+MR method gains more and more strength and usefulness as the detail increases. If we continued to detail the HTHP partition, for example, we could isolate specific parts of the plant and their components to identify the part that is less efficient or in higher maintenance demand.

In fact, the utility and true value of MR comes with the presence of the operator on site. Presence is essential to build the most appropriate HMT model at the time, but more importantly to build only the minimally necessary. Model refinements are driven on the needs of the moment, based on the operator's assessments of the context in which the problem is located.

For example, in continuing the refinement of partition P1.2.1.2.2.1 (HTHP), the operator in attendance might request a schematic of the plant and compare it with what is in front of him. Having the schematic, and possibly also the 3-D model of the plant, with HMT the operator would then be guided to deal with those parts that are underperforming at a certain time. Having solved the problem, the operator's attention then can be guided (with the virtual floating arrow) to the part that needs attention (shown in Figure 7.3.11). Having arrived there, the operator can decide to refine the model as needed.

The Figure 7.3.11 illustrates the concept and practical application of HMT+MR (Holonic Management Tree + Mixed Reality) within the SDS framework, showing how decision-making can be directly linked to both digital and physical environments.

On the left, we see a web-based rendering of the Holonic Management Tree (HMT). This digital representation allows decision-makers to structure processes, identify nodes (such as processes, sub-processes, or bottlenecks), and associate performance indicators. Each node can be expanded or decomposed to reflect the real complexity of the system, as discussed earlier with the selective decomposition principle.

At the bottom and right of the figure, we see how this digital HMT model is synchronized with a Mixed Reality environment. Through augmented or mixed reality devices, the user can visualize the same decision nodes directly overlaid on physical objects in the facility. For example, machinery or rooms in a Dairy Plant can be highlighted with KPIs and performance values derived from the HMT.

This integration has several consequences:

- **Synchronization with physical reality**: actions taken in the virtual HMT are reflected in the physical plant, ensuring that decision-making is not abstract but operational. Every action performed has to correpond to an actual action performed manually on the system (in this case just recorded by the HMT) or directly actioned by a virtual widget connected to an automation system in the best cases.

- **Physical actions and feedback**: operators can enact changes in the real system and immediately see their impact on the decision model.
- **Direct guidance**: the system highlights where the problem or bottleneck is, helping the user focus attention and interventions efficiently.
- **Decision-effect coupling**: because the model is interactive, every decision's consequences can be traced and assessed in real time, improving accuracy in continuous improvement cycles.

In the context of the SCE choreography, this means that when the Dairy Plant or Coolstore participates in negotiations and operations, their decision-makers can use HMT+MR to understand not only their own processes but also the impact on the global supply chain entity (SCE). Thus, local actions, such as adjusting a cooling cycle or shipment preparation, are visualized in terms of global effects on sustainability, efficiency, and GHG savings.

In our experimental session prepared for the Demo 1 in the context of WP6, we have constructed a 3D model for being prepared to feed the HMT+MR with a proper set of models that will be used to localize physically the theatre of the actions for the operator. In Figure 7.3.12, 7.3.13, and 7.3.14 we showcase three such views. In Figure 7.3.15, the Hololens 2 has been endowed with the 3D models aligned to the physical place, and is able to activate the HMT+MR for the operator situated in that environment. This image shows a virtual overlay of the Holonic Management Tree directly inside the Dairy Plant environment. By using a Mixed Reality headset, decision-makers can visualize both the facility infrastructure (pipes, machines, layouts) and, at the same time, the holonic decomposition of processes represented by the HMT.

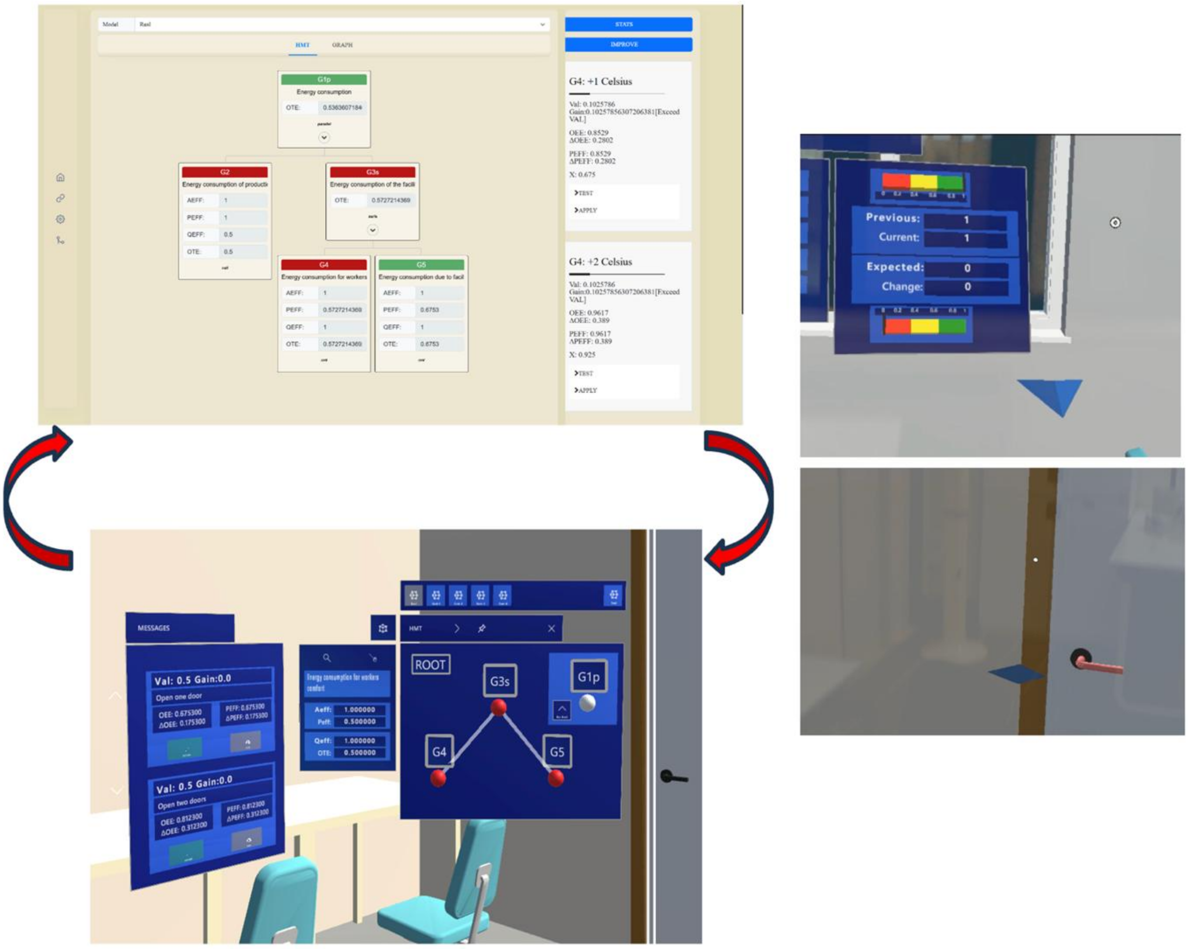

Figure 7.3.11: A gallery of the HMT+MR elements, made of Web-based GUIS, mixed reality superpositions with headsets (Hololens 2 in this case).

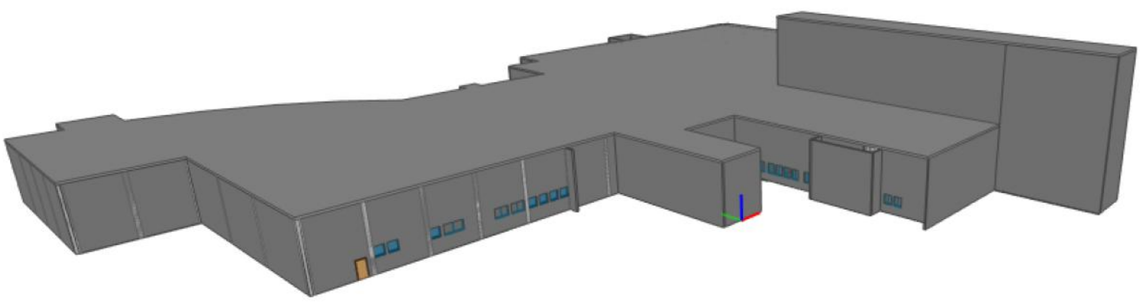

Figure 7.3.12: 3D model of the whole plant of the Diary Plant demonstrative case.

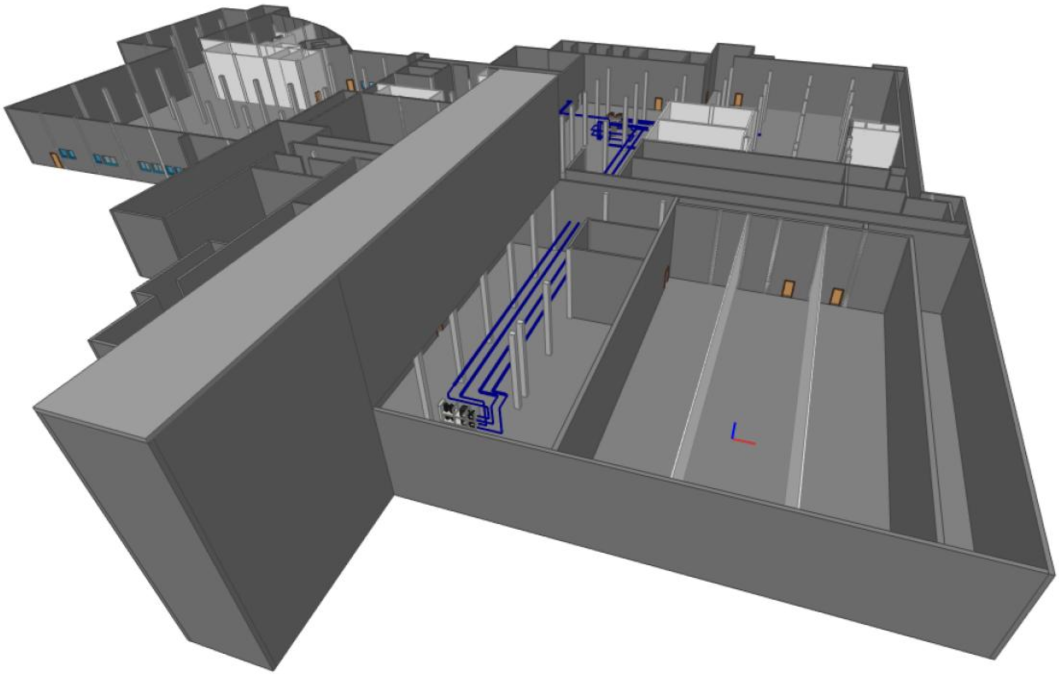

Figure 7.3.13: 3D a detail of the plant of the Diary Plant demonstrative case, HTHP premises.

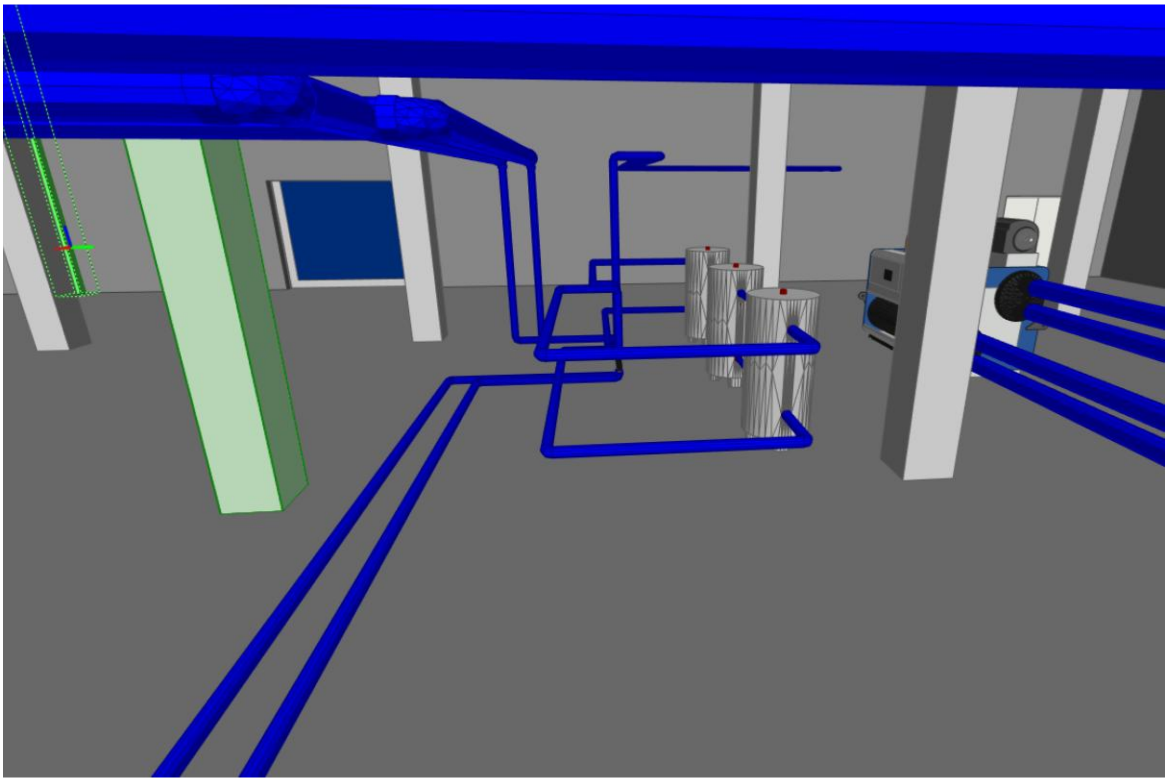

Figure 7.3.14: 3D a detail of the plant of the Diary Plant demonstrative case, CIP energy storage.

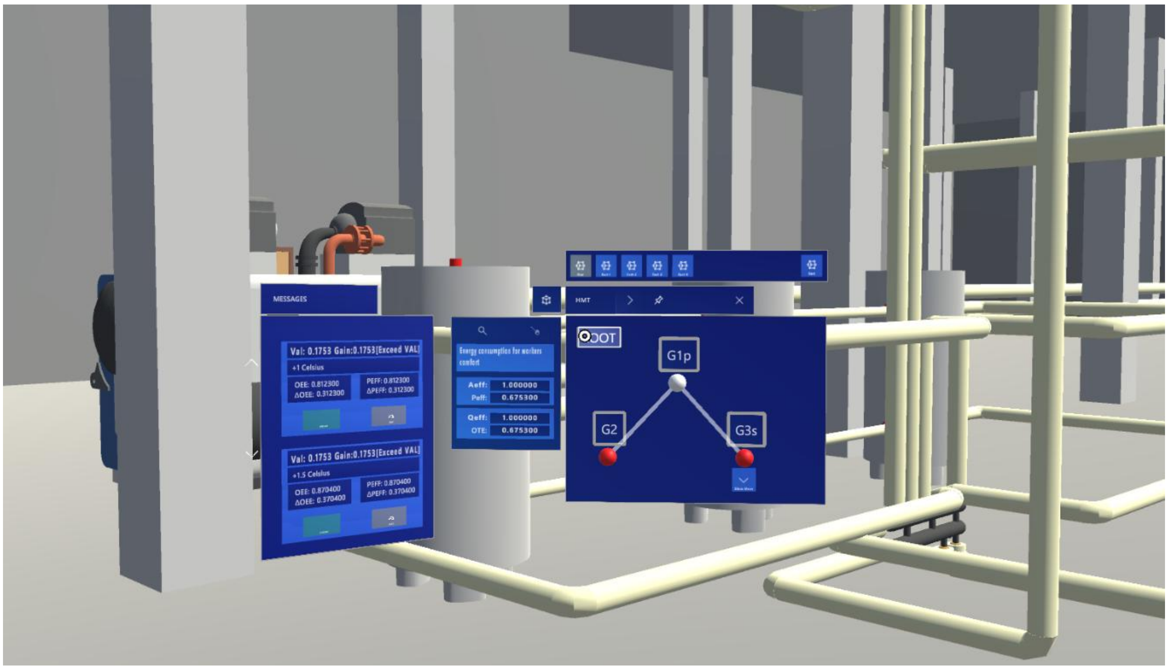

Figure 7.3.15: Mixed Reality (MR) available to the operator in the CIP storage room.

An apparent limitation of the HMT methodology is the need to have to access digital resources, models, simulators, information etc. and create a new branch or node of the HMT with these, along with automatic means of acquisition and implementation.

This means “digitizing on the fly.” But this is both a limitation and a strength of the method.

HMT can be considered an **incremental and continuously evolving digital twin**. One of the most critical limitations in the design and implementation of digital twins lies in the inherent tension between comprehensiveness and feasibility. On the one hand, a digital twin should cover a vast range of real-world features in order to faithfully capture the dynamics of the physical system; on the other hand, such extensive coverage is often unrealistic due to constraints in data availability, computational resources, and modeling complexity. This trade-off typically forces designers to make selective choices, which in turn risks overlooking important aspects of the system under study. The Holonic Management Tree (HMT) methodology provides a distinctive way to address this challenge: by explicitly recording, tracking, and continuously monitoring the design trade-offs and their associated costs, HMT ensures that the process of simplification or abstraction does not become opaque or arbitrary. Instead, it becomes a controlled and transparent element of the overall system design. In this sense, HMT enhances the accountability and sustainability of digital twin development, while also offering decision-makers a structured mechanism to assess whether further refinements or decompositions are warranted.

HMT is a decision-making engine that provide the manager with a continuous awareness about the where and when investments and efforts must be concentrated in order to solve cogent (and real) problems. This tool can be used to support decision making. Otherwise, the decision to invest to apply sensors (or better ones) into a plant part is quite arbitrary and made on gut feelings or by experience of some technical crew, but without a transferrable knowledge about the motivations and context taken for such decisions. HMT renders this process “rational” and “scientific” in one word systematic.

All the decision taken are recorded and discussed basing on what happened during the process. These recordings can be accessed and mined afterwards to constitute a lessons repository and a process mining source. In this methodology, there is an implicit nudge for the manager to well justify her decisions, but also a way to justify some failures in case of wrong decisions. Nonetheless, the HMT mechanisms and inherent monotonicity guarantees that any action is made "at best" in the situation, and no big harm can be made on the plant if the feedback of the HMT are taken into account in the decisions. The HMT warns early in the case of bad modelling of the reality that is under control. It also guides the manager step by step with minor risks.

This safety of the HMT is traded off with the wish of an optimal solution that is a common goal for any plant manager. But we remind that in case of complexity, by definition, no global optimality can be reached by design, just because there are no conditions for optimal modelling. Optimality acceptation can be used only in local, temporary, and well controlled sense.

The famous sentence (a synthesis attributed to) of George E.P. Box (1976)[48] "**All models are wrong, but some are useful"** is a clear guide to HMT. In his paper, Box discusses the role of statistical models in scientific inquiry, highlighting their utility despite their inherent simplifications. It reminds readers that while models cannot capture every detail of reality, they remain crucial tools for learning and decision-making. HMT states that "when and how much a model is useful, we can tell you: the sooner, the better". HMT states with Box: "The experimental design is [...] a movable window looking onto the true state of nature. Its positioning at each stage is motivated by current beliefs, hopes, and fears" of the plant manager[49].
HMT+MR is a special taste of HMT that is strengthened by the possibility for the operator to "filter" some complexity being in the very near condition for a good, situated modelling.

[48] Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, *71*(356), 791-799.
[49] Continuing relevant quotas from Box(1976): " 2.3 Parsimony. Since all models are wrong the scientist cannot obtain a "correct" one by excessive elaboration. On the contrary following William of Occam he should seek an economical description of natural phenomena. Just as the ability to devise simple but evocative models is the signature of the great scientist so overelaboration and overparameterization is often the mark of mediocrity.
2.4 Worrying Selectively. Since all models are wrong the scientist must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad."

8 LESSONS LEARNT

In this final section we propose briefly some discussion points that concern the lessons learnt during the activities and the research conducted in the Work Package 5 of the project.

To evaluate the SDS results more quantitatively, we have to associate Key Performance Indicators (KPIs) to the main dimensions of performance:

– **Latency:** average message transmission time across SCE Participants (measured <200 ms in demo runs). Nonetheless, latency up to some minutes are foreseen dependending on the foreseen introduction of an enforcement control made by Blockchain technologies. Usually, such lags are not critical in real supply chains.

– **Energy efficiency:** percentage reduction in GHG emissions enabled by coordinated decisions (target: 8–12% in demo cases). In general, the energy savings are highly depentent on the processes that are participating in the SCE with high range of variability. The SDS system only guarantees that all the participants are nudged towards a continuous improvement that has to be evaluated both in the short and the long term.

– **Scalability:** number of concurrent Actors and SCEs supported without degradation (benchmarked up to 50 Actors). The early assessment have been done with a standard cloud infrastructure and dockerization. Scalability issues are critical and are a focus in the business plan for the exploitation of the SDS.

– **Compliance:** degree of alignment with EU sustainability standards and RAMI 4.0 (qualitative and quantitative mapping). These sustainability alignments have been qualitatively assessed through the collaboration with experts in the consortium, in particular with close collaboration with work package 7 leaders. RAMI 4.0 is being considered by design of the SDS, and fully compatible.

– **Productivity:** throughput improvements from coordinated scheduling (measured as % increase in OEE/OTE).”. This feature has been directly addressed with the introduction of the HMT decision making, which core definitions of the KPIs are based on Overall Equipment Effectiveness (OEE) and Overall Throughput Effectiveness. The actual meaning of them ranges a vast set of different processes that are being normalized and unified through the HMT paradigm.

These KPIs are the basis of the business plan for the exploitation of the SDS, which is under development since the start of the last semester of project activities. Future work will extend this KPI framework by running additional stress tests, collecting user feedback from real stakeholders, and refining the benchmarks according to domain-specific requirements (e.g., food supply chain, logistics, manufacturing).

8.1 Recommendations for Industry

One of the most important lessons emerging from our work is the need to **keep things simple**. Industrial ecosystems, particularly those dealing with food supply chains (FSC), are inherently complex, involving multiple actors, heterogeneous infrastructures, and diverse interests. If solutions become over-engineered, they risk being rejected or not adopted. Simplicity ensures accessibility, usability, and scalability.

At the same time, we must acknowledge that the food supply chain is complex by nature. This complexity cannot be eliminated, but it can be managed through methods, digital tools, and structured collaboration. Here lies the value of choreography and systemic approaches: they allow industry players to operate within a structured but flexible framework.

A key enabler of progress is **collaboration**. No company can solve sustainability and digitalization challenges in isolation. Supply chains are interdependent by definition, and meaningful results emerge only when participants commit to common goals and share relevant information, while still respecting their competitive positioning.

Looking ahead, we align with the **paradigm of Industry 5.0**, which emphasizes not just automation and efficiency, but the central role of **humans**. Technology is not an end in itself: it is about empowering human decision-making, guiding operators, and augmenting their capabilities through AI, mixed reality, or decision-support tools. This human-centric view helps balance private interests with collective sustainability outcomes.

Sustainability is no longer optional. Regulatory frameworks, consumer expectations, and corporate responsibility all converge toward reducing emissions, improving efficiency, and embedding green practices. It is not just about compliance, but also about competitiveness and resilience in a market increasingly sensitive to environmental impact.

Emerging technologies like blockchain open major opportunities to certify, track, and exchange information in a secure and transparent way. They can build trust across stakeholders, support credits markets for GHG reduction, and unlock new business models where virtuous behavior is recognized and rewarded.

This is part of a broader digital transformation, where data, connectivity, and advanced analytics redefine how decisions are taken and how processes are optimized. Digital twins, IoT, AI, and mixed reality are examples of technologies that, if properly integrated, can radically enhance efficiency and sustainability.

To succeed, however, companies must adopt a **systems engineering perspective**. This means designing processes and technologies not as isolated silos, but as components of a larger interconnected system. Such an approach enables both top-down strategy and bottom-up innovation, ensuring that improvements in one area are aligned with the overall performance of the ecosystem.

Finally, it is essential to **think holistically**. Industry leaders must develop the capacity to see the whole and the parts at once: to grasp how each local action affects the global system, and how system-wide objectives can cascade into concrete, local practices. This dual vision is crucial to reconcile efficiency, competitiveness, and sustainability, and to fully embrace the opportunities of Industry 5.0.

To operationalize collaboration, industry actors should adopt shared digital platforms where KPIs such as latency, carbon footprint reduction, and throughput efficiency are tracked transparently across participants. Adoption of blockchain-based certification should be tied to quantifiable targets, for example reducing transaction verification times by 20% or ensuring traceability for 95% of shipments. Companies are encouraged to join consortia to co-develop digital twin pilots with clear milestones: short-term (process visualization), mid-term (predictive maintenance), and long-term (autonomous optimization). Such roadmaps help translate broad visions into measurable industry-wide actions.

8.2 Recommendations for policy makers

From the experience gained, several recommendations emerge for policy makers, whose role is crucial in shaping the conditions that enable sustainable and digital transformation of the food supply chain (FSC).

First, there is a strong need to **incentivize the digital transformation of the FSC**. The adoption of digital platforms, data-driven tools, and collaborative infrastructures can drastically improve efficiency, transparency, and sustainability across the chain. However, without the right policy support, many actors—especially smaller ones—may lack the resources to make this shift.

This brings us to a key point: the transformation must be **inclusive to SMEs**. Large corporations often have the resources to adopt new technologies, but small and medium enterprises form the backbone of the agri-food sector in Europe and beyond. If left behind, the transformation will remain fragmented, and the benefits will be unevenly distributed. Policies should therefore lower entry barriers for SMEs, simplifying procedures and offering targeted support.

A central task for policy makers is to **lower market barriers and reduce the costs of transformation**. Digital technologies, infrastructures, and expertise can be expensive. Public investment, subsidies, and fiscal incentives should reduce this burden, ensuring that even the smallest actors can participate in collaborative and sustainable digital ecosystems.

In parallel, policies should **support technology transfer in sustainability**. Many solutions already exist—developed in research projects, innovation hubs, and pilot programs—but they rarely reach market adoption at scale. Bridging this gap requires structured initiatives that connect research with industry, fostering uptake of methods and tools that have already been validated.

Another recommendation is to **provide incentives to the exploitation and dissemination of initiatives like ENOUGH**. Too often, European research projects conclude successfully but their results are not fully capitalized. Policymakers should encourage companies, associations, and institutions to adopt and spread these outcomes, so that the knowledge does not remain confined to academic circles but contributes to real industrial impact.

Moreover, there is value in **providing follow-ups to initiatives like the SDS Framework**. Rather than treating research projects as isolated episodes, a policy-driven continuity would ensure that their frameworks evolve, adapt, and expand into practical standards and platforms. This continuity would multiply the long-term impact of public funding.

Policy makers must also be encouraged to **think systemically**. The challenges of the food supply chain—sustainability, resilience, digitalization, and inclusivity—are deeply interconnected. Fragmented policies risk producing contradictions or inefficiencies, while systemic policies can align multiple objectives at once.

For this reason, policies should **promote technologies and approaches that foster systemic thinking**. Tools like supply chain choreography, digital twins, and holonic management methodologies provide frameworks for understanding interdependencies and designing collaborative improvements. Supporting such approaches through regulation and funding would accelerate systemic change.

Finally, **interdisciplinarity must be facilitated and promoted**. The transformation of food supply chains involves agriculture, logistics, energy, digital technologies, and human factors. Policymakers should design programs that encourage collaboration across disciplines, sectors, and expertise, ensuring that solutions are not only technically sound but also socially, economically, and environmentally viable.

Policy makers play a pivotal role in creating an enabling environment where innovation can thrive. By reducing barriers, supporting SMEs, promoting systemic thinking, and valorizing research results like those of ENOUGH, they can ensure that the digital and sustainable transformation of food supply chains becomes a reality, delivering benefits for industry, society, and the environment.

Policy makers should establish targeted fiscal incentives such as tax credits for SMEs adopting certified digital twin solutions, or matching grants for collaborative innovation projects. To reduce regulatory uncertainty, governments could create regulatory sandboxes where novel SDS-based solutions can be tested under controlled conditions. In parallel, coordination with European and international standardization bodies (e.g., CEN, ISO, IEEE) should be promoted to fast-track SDS concepts into interoperable standards. This dual approach—financial support and regulatory clarity—ensures that SMEs not only access technology but can deploy it with reduced risk and long-term sustainability.

8.3 Recommendations for society

The transformation of food supply chains and the broader move toward sustainability cannot be achieved by industry and policymakers alone. **Society has a decisive role**, because ultimately, it is the collective behavior of consumers, communities, and citizens that shapes demand, drives acceptance of new solutions, and determines whether change becomes permanent.

First, it is essential for individuals to **act responsibly as consumers**. Every purchase carries weight, and small daily choices—such as favoring local products, reducing waste, or supporting companies with transparent sustainability practices—aggregate into significant impacts across entire supply chains. Being a responsible consumer is not only about buying less or differently, but about recognizing the influence we have as participants in the global economy.

Equally important is to **actively search for and support sustainability in products**. Labels, certifications, and transparent information are tools that allow consumers to make informed decisions. Choosing products with lower carbon footprints, ethical sourcing, or reduced packaging not only sends a strong signal to producers but also creates economic incentives for companies to adopt greener and fairer practices. By rewarding sustainability in the marketplace, consumers directly accelerate the transition to a more responsible economy.

At the same time, society must **avoid falling into the trap of worshipping digital technologies**. Technology is not an end in itself but a tool—one that must be shaped, critiqued, and guided to ensure it serves human and environmental well-being. Blind adoption can create new inequalities, exclude vulnerable groups, or lead to unintended consequences. Instead, citizens should participate in debates, demand transparency, and engage in co-creation processes that keep technology aligned with public interest.

Finally, there is a moral responsibility to **promote technologies that act for the social good**. This means supporting not only innovations that improve efficiency or profit margins, but also those that address broader societal challenges: reducing greenhouse gas emissions, improving working conditions, strengthening community resilience, or ensuring food security for all. Civil society organizations, local communities, and individuals can amplify these values by raising awareness, advocating for fairness, and ensuring that innovation contributes to collective rather than purely private gains.

In conclusion, the lessons for society highlight that **everyone is a stakeholder in the transformation of food supply chains**. Responsible consumption, support for sustainability, critical engagement with technology, and promotion of social good are not abstract ideals but practical steps that communities can take. If embraced collectively, they provide the cultural and social foundation on which industries and policymakers can build lasting change, ensuring that digital transformation and sustainability truly benefit both people and the planet.

Consumers can actively drive systemic change by supporting certification schemes that verify sustainability claims via blockchain or SDS-based infrastructures. For example, by choosing products carrying a “low-carbon digital twin” label, citizens directly encourage firms to invest in traceability and emission reductions. Community-driven reputation systems, where consumers evaluate companies’ adherence to sustainability commitments, could complement regulatory oversight and framework. Furthermore, civil society organizations should pilot incentive mechanisms—such as local credit systems rewarding waste reduction—that demonstrate how individual choices translate into measurable societal benefits.

9 CONCLUSION

The Smart Data System (SDS) has been developed as an innovative digital backbone for sustainable food supply chains, integrating business processes, digital platforms, and advanced digital industry and decision science paradigms into a coherent and operational framework. Its ability to manage Supply Chain Entities (SCEs) as evolving digital twins, combined with blockchain-based trust mechanisms and holonic decision-making enhanced by hybrid and mixed reality tools, confirms that the SDS is more than a technical artifact. It is a systemic enabler, aligning business transformation with the sustainability objectives of the ENOUGH project and the broader aspirations of Industry 5.0.

The project has shown that collaboration, openness, and adaptability are indispensable in addressing the inherent complexity of modern food supply chains. By adopting agile development principles, engaging diverse stakeholders, and maintaining an open architecture, the SDS has proven capable of continuous evolution. It can incorporate new modules, respond to changing requirements, and align with emerging standards such as RAMI 4.0, the Asset Administration Shell, and Industry 4.0/5.0 reference models. These qualities make it not only a research demonstrator but also a foundation for long-term exploitation well beyond the project's completion.

While section 2 introduces to the concept evolution and development, chapters 3 to 7 of this report illustrated how the SDS vision has been translated into practical applications. The life cycle management of SCEs demonstrated how digital twins can be configured, monitored, and improved across multiple iterations, establishing the basis for continuous systemic learning. The Marketplace scenarios revealed the SDS's inclusivity: enabling clients to purchase services, supporting third-party providers in selling innovative solutions, and allowing the SDS owners to channel research-driven innovations into practice. These use cases highlight the SDS's dual nature—both as infrastructure that guarantees interoperability and as a business ecosystem that fosters value creation, co-innovation, and competitiveness.

Section 8 complemented these demonstrations with key lessons learned and actionable recommendations. For industry, the SDS provides a concrete pathway to digital transformation, transparency, and measurable greenhouse gas (GHG) reduction. For policymakers, it underlines the urgency of systemic and interoperable frameworks that can guide regulation and incentivize sustainable practices. For society at large, it shows how trust, accountability, and human-centric innovation can be embedded in digital transformation processes, ensuring inclusiveness and long-term acceptance.

Crucially, the SDS illustrates the value of systemic thinking: where human decision-making, digital twin monitoring, automation, and policy alignment converge to support sustainability, competitiveness, and resilience. Simplicity in design, openness to stakeholder participation, and holistic approaches have emerged as decisive success factors.

In conclusion, the SDS should not be regarded as an endpoint but rather as a foundation—an evolving framework with the capacity to absorb new technologies, standards, and business models. Its openness, modularity, and adaptability ensure that it will remain relevant in the years to come, serving as a cornerstone for sustainable digital transformation. By bridging the gap between research, business, and societal needs, the SDS paves the way for greener, more resilient, and more inclusive food supply chains, and sets a benchmark for the digital infrastructures of the future.

enough-emissions.eu

Disclaimer: This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101036588.