Blockchain as a Key Enabling Technology for Decentralized Cyber-Physical Production Systems*

M. Isaja and A. Calà

Abstract— **Over the last decade, several innovative factory automation architectures have been developed that leverage the Cloud Computing paradigm to improve the flexibility and adaptiveness of the shopfloor. However, their actual adoption by the manufacturing industry is still limited, due mainly to performance and security constraints. To address these issues, the FAR-EDGE project is proposing a reference architecture based on Edge Computing for Industry 4.0 solutions, focusing on a more efficient use of distributed computing power and network bandwidth. This paper reports on the FAR-EDGE experimentation of Blockchain and Smart Contracts as key enabling technologies for Edge Computing in the specific context of decentralized Cyber-Physical Production Systems.**

I. Introduction

The Bitcoin network is undoubtedly the best known use case of the "distributed ledger" concept: a totally decentralized information system where no individual node is ever in control, while the overall consistency of data and logic is maintained by means of a "consensus" algorithm – i.e., nodes must continuously reach an agreement over the network on what the current state of the system is. Despite Bitcoin being a very specialized application that implements a digital currency, the foundation it builds on is not: Blockchain technology, since Bitcoin's beginnings in 2009, has been used to power a great number of alternative platforms, thanks to a new generation of software that added even more innovative capabilities to an already disruptive architecture. In particular, "Smart Contracts" allow the users to define their own custom business logic, turning the task-specific distributed ledger into a distributed computing environment that can support any kind of application.

Lately, these features have attracted a lot of interest from the corporate world. For most of these observers, the value of Blockchain is being an "enabler of trust": a technology that can make existing business processes simpler and new business models possible by removing the need of a trusted intermediary between parties that do not trust or even know each other. Finance, utilities and supply chain management are the most obvious sectors in which such decentralization can be applied with profit, but new use case proposals are emerging nearly every day – sometimes with some merit. "*Blockchain [...] is often perceived as the catalyst of an IT revolution to come, likened by some to the advent of the World Wide Web in the nineteen-nineties*" [1]. There are still significant challenges to overcome in order to fulfil this promise, tough, the most prominent being system scalability and regulatory frameworks (e.g., Blockchain records to be globally accepted as evidence in litigation). According to a 2017 study from Gartner, Blockchain technology is currently at the "*peak of inflated expectations*" and the time frame of full "*mainstream adoption*" will probably start from year 2022 [2].

That said, the FAR-EDGE[1] research project (funded by the Horizon 2020 programme of the European Commission) is pursuing a less far-fetched and more short-term goal, focusing on those unique capabilities of the Blockchain that impact on ICT systems rather than on business ecosystems, enabling decentralized and highly-available solutions rather than trust. In this paper we report about the FAR-EDGE exploration of this topic in the context of factory automation, more specifically for the enablement of machine-to-machine (M2M) collaboration in Cyber-Physical Production Systems (CPPS).

II. The FAR-EDGE Project

The FAR-EDGE baseline value proposition is a Reference Architecture (RA) that applies a novel kind of Edge Computing approach to Industry 4.0 solutions [3], with the purpose of improving the overall flexibility and adaptiveness of the shopfloor. The market driving forces are the mass-customization and re-shoring trends: there's an increasing need for agile production lines that can support "lot-size-one" scenarios and cope with frequent changes in requirements, environment and workload. In this context, FAR-EDGE aims at providing a new design approach and new software tools for CPPS, while at the same time allowing an easy migration path from legacy factory systems. Such migration path can be seen from two different but complementary perspectives, facing each other: a top-down one that considers how to bring computing power near to where it's actually needed, for increased efficiency (the classic Edge Computing approach); and a bottom-up one where the main concern is to decompose the monolithic production line into a number of independent modules that can be dynamically rearranged.

According to the FAR-EDGE RA, the topmost tier is the result of splitting some of the central IT systems of the enterprise/factory into a number of smaller, locally-scoped computing units, called Edge Gateways (EG). Each EG is responsible for the direct monitoring and control of "passive" Edge Nodes (EN) that are deployed in close proximity to it – i.e., legacy sensors and actuators. This allows better use of

*Research supported by European Commission's Horizon 2020 program.

M. Isaja is with the Research & Development Unit in Engineering Ingegneria Informatica (ENG), Rome, Italy (e-mail: mauro.isajy@eng.it).

A. Calà is with the Otto-von-Guericke University, Magdeburg, Germany, and Siemens AG, Erlangen, Germany (e-mail: ambra.cala@ovgu.de).

[1] https://www.faredge.eu

network resources and reduced response times, with little or no disruption of the existing shopfloor: an element of paramount importance in real-world manufacturing business, where risk reduction is a priority and green-field implementation is an exception.

The bottom side of the FAR-EDGE RA is the Field Tier, populated by Smart Objects (SO) – i.e., tools or machines with on-board intelligence. SOs have some degree of autonomous behaviour but still require factory-level coordination. The Field Tier is addressing more modern manufacturing systems that do not (entirely) follow the hierarchical ISA-95 architectural standard, also known as the "*automation pyramid*" [4].

The big challenge of the FAR-EDGE approach is then twofold: firstly, to combine older and newer factory architectures in a seamless way, without disrupting the former or hindering the latter; then, to obtain the best from extensive decentralization but still retain control over the manufacturing process as a whole. The key technology for this task is, not surprisingly, the Blockchain, which in the FAR-EDGE RA is positioned on a tier on its own: the Ledger. From a conceptual point of view, the Ledger Tier acts as an intelligent communication channel connecting EG and SO nodes into a peer-to-peer (P2P) network. More concretely, this communication is implemented as shared process state that is maintained on the Blockchain and managed by Smart Contracts, that in the FAR-EDGE RA are named Ledger Services (LS) for consistency. This integrated vision is represented in Fig. 1 below.

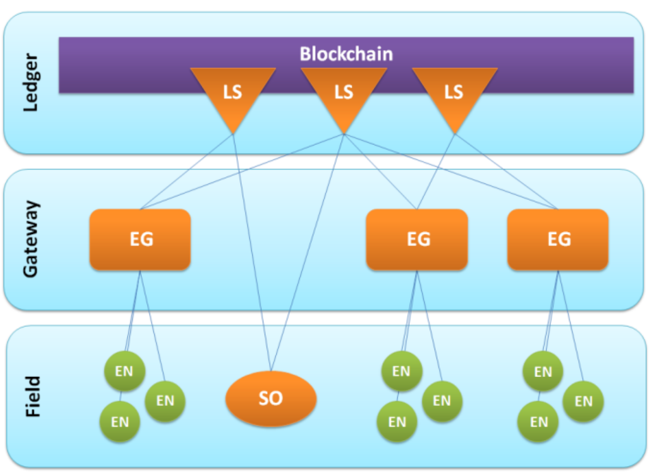

Figure 1 - FAR-EDGE RA: Field, Gateway and Ledger tiers

The FAR-EDGE project is not just about architecture, though: its second and arguably more important value proposition is the Reference Implementation (RI) of a platform which supports RA-compliant solutions. The FAR-EDGE RI provides basic services and tools for building Edge Computing factory systems, covering automation, analytics and simulation scenarios. The high-level design of the FAR-EDGE RI, with all its individual components in place, is shown in Fig. 2 (adapted from the original picture in the FAR-EDGE RA specifications [5]).

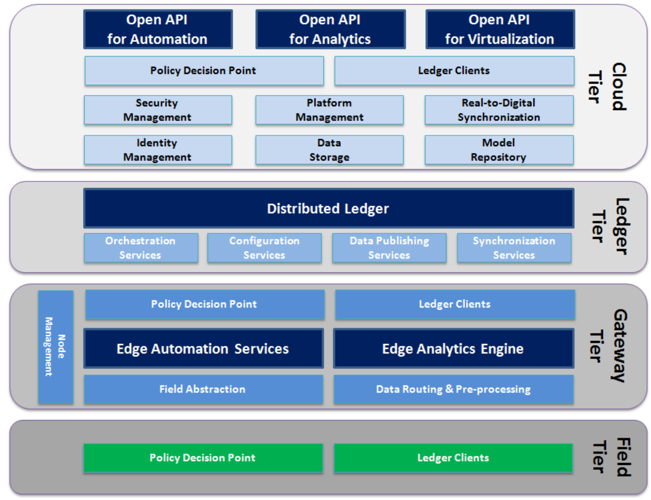

Figure 2 - Components of the FAR-EDGE RI

In this paper, our main concern is the use of the Blockchain and of Smart Contracts as key enabling technologies for decentralized CPPS. So, having provided a general overview of the FAR-EDGE RA/RI, we are now going to focus on the Ledger Tier only. In the actual RI, the Distributed Ledger component depicted above is an Hyperledger Fabric[2] instance. The four boxes right underneath – Orchestration, Configuration, Data Publishing and Synchronization – represent categories of Ledger Services that cover a wide array of use cases. In other words, in FAR-EDGE the Blockchain is an off-the-shelf product (with minor extensions, like the integration with platform-level identity management), which is enriched by application-specific Smart Contract software. The entire Synchronization Services category is dedicated to secure state sharing tasks – i.e., keeping process instances running on EG and SO nodes in-sync, while at the same time protecting shared data from unauthorized access.

III. Blockchain Issues

For those familiar with the technology, the main question raised by the previous chapter is: how can a Blockchain fit into this picture? According to conventional wisdom, Blockchains platform are slow and cumbersome systems with limited scalability and an aversion to data-intensive applications. Nevertheless, while this vision has solid roots in reality (at least at the time of writing, but things are changing), in the context of smart factories these shortcomings are not as relevant as it may seem. In order to substantiate this claim, though, we first need to explain some key points of the technology: those that are both its enablers and its bottlenecks.

First and foremost, the Blockchain is a log of all transactions (i.e., state changes) executed in the system. The log, which is basically a "witness" of past and current system state, is replicated and kept in-sync across multiple nodes. All nodes are peers, so that no "master node" or "master copy" of the log exists anywhere at any time. Internally, the log is a linear sequence of records (i.e., "blocks" containing transactions) that are individually immutable and time-stamped. The sequence itself can only be modified by

[2] https://www.hyperledger.org/projects/fabric

appending new records at the end. The integrity of both records and sequence is protected by means of strong cryptographic algorithms [6]. Moreover, all records must be approved by consensus among peers, using some sort of Byzantine Fault Tolerance (BFT) mechanism as a guarantee that an agreement on effective system state can always be reached, even if some peers are unavailable or misbehaving (in good faith or for malicious purposes) [7][8].

The process described above is all about trust: the consensus protocol guarantees that all approved transactions conform to the business logic that peers have agreed on, while the log provides irrefutable evidence of transactions. For this to work in a zero-trust environment, where peers that do not know – let alone trust – each other and cannot be sanctioned by any higher authority, there is yet another mechanism in place: an economic incentive that rewards "proper" behaviour and makes the cost of cheating much higher than the profit. Given that the whole system must be self-contained and autonomous, such incentive is based on native digital money: a "cryptocurrency". This closes the loop: all public Blockchain networks need a cryptocurrency to fuel their BFT mechanism – a scenario that has been dubbed "cryptoeconomics". For some of them (e.g., Bitcoin), the cryptocurrency itself is the main goal of the system: transactions are only used to exchange value between users. Other systems (e.g., Ethereum) are much more flexible, as we will see further on. That said, cryptocurrencies are problematic for many reasons, including regulatory compliance, and hinder the adoption of the Blockchain in the corporate world.

Another key pointy of Blockchain technology that is worth mentioning is the problem of "transaction finality". Most BFT implementations rely on "forks" to resolve conflicts between peer nodes: when two incompatible opinions on the validity of some transaction exist, the log is split in two branches, each corresponding to one alternate vision of reality – i.e., of system state. The other nodes of the network will then have to choose which branch is the valid one, and will do this by appending their new blocks to the "right" branch only. Over time, consensus will coalesce on one branch (the one having more new blocks appended), and the losing branch will be abandoned. While this scheme is indeed effective for achieving BFT in public networks, it has one important consequence: there is no absolute guarantee that a committed transaction will stay so, because it may be deemed invalid *after* it is written to the log. In other words, it may appear only on the "bad" branch of a fork and be reverted when the conflict is resolved. Clearly enough, this behaviour of the Blockchain is not acceptable in scenarios where a committed transaction has side effects on other systems.

This is how first-generation Blockchains work. For all these reasons, public Blockchains are, at least to date, extremely inefficient for common online transaction processing (OLTP) tasks. This is most unfortunate, because second-generation platforms like Ethereum have introduced the Smart Contract concept. Smart Contracts were initially conceived as a way for users to define their custom business logic for transaction – i.e., making the Blockchain "smarter" by extending or even replacing the built-in logic of the platform. It then became clear that Smart Contracts, if properly leveraged, could also turn a Blockchain into a distributed computing platform with unlimited potential. However, distributed applications would still have to deal with the scalability, responsiveness and transaction finality of the underlying BFT engine, which significantly limits the range of possible use cases.

To tackle this problem, the developer community is currently treading two separate paths: upgrading the BFT architecture on the one side, relax functional requirements on the other. The former approach is ambitious but slow and difficult: it is followed by a third generation of Blockchain platforms that are proposing some innovative solution, although transaction finality still appears to be an open point nearly everywhere. The latter is much easier: if we can assume some limited degree of trust between parties, we can radically simplify the BFT architecture and thus remove the worst bottlenecks. From this reasoning, an entirely new species was born in recent years: "permissioned" Blockchains. Given their simpler architecture, commercial-grade permissioned Blockchains are already available today (e.g., Hyperledger, Corda), as opposed to third-generation ones (e.g., EOS, NEO) which are still experimental.

IV. Permissioned Blockchains

Permissioned Blockchains are second-generation architectures that do not support anonymous nodes and do not rely on cryptoeconomics. Basically, they are meant to make the power of Blockchain and Smart Contracts available to the enterprise, at least to some extent. Their BFT is still a decentralized process executed by peer nodes; however, the process runs under the supervision of a central authority. This means that all nodes must have a strong digital identity (no anonymous parties) and be trusted by the authority in order to join the system. Trust, and thus access to the Blockchain, can be revoked at any time. The BFT protocol can then rely on some basic assumptions and perform much faster, narrowing the distance from OLTP standards in terms of both responsiveness and throughput. Some BFT implementation (e.g., Hyperledger Fabric) also support final transactions, as consensus on transaction validity can be reached in near-real-time *before* anything is written to the log.

The key point of permissioned Blockchains is that they are only partially decentralized, leaving governance and administration roles in the hands of a leading entity – be it a single organization or a consortium. This aspect is a boon for enterprise adoption, for obvious reasons. Typically, these networks are also much smaller than public ones, with the positive side effect of limiting the inefficiency of data storage caused by massive data replication across peer nodes. Overall, we can argue that permissioned Blockchains are a viable compromise between the original concept and legacy OLTP systems. But then, to what extent? Can we identify some use cases that a state-of-the-art permissioned Blockchain can effectively support? This is exactly what the FAR-EDGE project aims at, with the added goal of validating claims on the field, by means of pilot applications deployed in real-world industrial environments.

V. THE FAR-EDGE LEDGER TIER

The first problem that FAR-EDGE had to face was to define the "performance envelope" of current Blockchain implementations, so that validation cases could be shaped according to the sustainable workload. The idea was to set the benchmark for a "Blockchain comfort zone" in terms of a few objective and measurable Key Performance Indicators (KPI), targeting the known weak points of the technology:

- Transaction Average Latency (TrxAL) – The average waiting time for a client to get confirmation of a transaction, expressed in seconds.
- Transaction Maximum Sustained Throughput (TrxMST) – The maximum number of transactions that can be processed in a second, on average.

The benchmark was set by stress-testing, in a lab environment, actual Blockchain platforms. These were selected after a preliminary analysis of the permissioned Blockchains available from open source communities, using criteria like code maturity and, most importantly, finality of transactions. The only two platforms that passed the selection were Hyperledger Fabric (HLF) and NEO. The stress test was then conducted using BlockBench, a specialized testing framework [9], and a simple configuration of eight nodes on commodity hardware.

HLF emerged from tests as the only viable platform for CPPS applications, given that NEO is penalized by a significant latency (~7s.), which is independent from workload (the expected result for a "classical" Blockchain architecture that aggregates transactions into blocks and defines a fixed delay for processing each block). On the contrary, HLF was able to accept a workload of up to 160 transactions per second with relatively low latency (0.1-1s.). On heavier workloads, up to 1000 transactions per second, NEO is instead the clear winner, thanks to its constant latency, while HLF's performance progressively degrades (>50s.). This workload profile however, while appealing for high-throughput scenarios (e.g., B2C payment networks), is not compatible with basic CPPS requirements. Consequently, the Blockchain performance benchmark was set as follows:

- **0.1 <= TrxAL <= 1.0**
- **0 <= TrxMST <= 160**

This is also considered the performance envelope of the FAR-EDGE Ledger Tier, as the HLF platform has been adopted as its baseline Blockchain implementation.

VI. CPPS BLOCKCHAIN USE CASES

Having marked some boundaries, the FAR-EDGE project then proceeded with the identification of some pilot applications for the validation phase. The starting point was a set of candidate use cases proposed by our potential users, who were eager to tackle some concrete problems and experiment with some new ideas. The general framework of this exercise is described here.

As already explained in Chapter II, the main objective in FAR-EDGE is to achieve flexibility in the factory through the decentralization of production systems. The catalyst of this transformation is the Blockchain, which – if used as a computing platform rather than a distributed ledger – allows the virtualization of the automation pyramid. The Blockchain provides a common "virtual space" where data can be securely shared and business logic can be consistently run. That said, users can leverage this opportunity in two ways: one easier but somewhat limited, the other more difficult and more ambitious.

The easiest approach is of the brown-field type: just migrate (some of) the factory's centralized monitoring and control functionality to Ledger Services on the Ledger Tier. Thanks to the Gateway Tier, legacy centralized services can be "impersonated" on a local scale by Edge Gateways: the shopfloor – that hardest environment to tamper with in a production facility – is left untouched. The main advantages of this configuration are the mitigation of performance bottlenecks (heavy network traffic is confined locally, workload is spread across multiple computing nodes) and added resiliency (segments of the shopfloor can still be functional when temporarily disconnected from the main network). Flexibility is also enhanced, but on a coarse-grained scale: modularity is achieved by grouping a number of shopfloor Edge Nodes under the umbrella of one Edge Gateway, so that they all together become a single "module" with some degree of self-contained intelligence and autonomy. Advanced Industry 4.0 scenarios, like "plug-and-produce", are out of reach.

The more ambitious approach is also a much more difficult and risky endeavour in real-world business, being of the green-field type. It's about delegating responsibility to Smart Objects on the shopfloor, which communicate with each other through the mediation of the Ledger Tier. The business logic in Ledger Services is higher-level with respect to the previous scenario: more about governance and orchestration than direct control. The Gateway Tier has a marginal role, mostly confined to big data analytics. In this configuration, central bottlenecks are totally removed and the degree of flexibility is extreme. The price to pay is that a complete overhaul of the shopfloor of existing factories is required, replacing PLC-based automation with intelligent machines.

In FAR-EDGE, both paths are going to be explored with different use cases combining on automation, analytics and simulation. We here give one full example of each type. The information is preliminary, as at the time of writing they are still work in progress: according to the project timeline, the complete pilot use cases will be ready in the third quarter of 2019.

The first use case follows the *lightweight* brown-field approach. The legacy environment is an assembly facility for industrial vehicles. The pilot is called "mass-customization": the name refers to capability of the factory assembly line to handle individually customized products having a high level of variety. If implemented successfully, mass-customization can give a strategic advantage to target niche markets and meet diverse customer needs in a timely fashion. In particular, the pilot factory produces highly customized trucks. The product specification is defined by up to 800 unique variants, and the final assembly includes approximately 7000 manufacturing operations and handles a very high degree of geometrical variety (axle configurations, fuel tank positions etc.). Despite the high level of variety in the standard product, at some production sites 60% of the produced trucks have unique customer adaption.

In the pilot factory, the main assembly line is sequential but feeds a number of finishing lines that work in parallel. In particular, the wheel alignment verification is done on the finishing assembly line and is one of the last active checks done on trucks before they leave the plant. This opens up an opportunity to optimize the workload. In the as-is scenario, wheel alignment stations are statically configured to accommodate specific truck model ranges: products must be routed to a matching station on arrival, creating a potential bottleneck if model variety is not optimal. As part of the configuration, a handheld nut runner tool needs to be instructed as to the torque force to apply.

In the to-be solution, according to the FAR-EDGE architectural blueprint, each wheel alignment station is represented at the Edge Tier level by a dedicated Edge Gateway box. The EG runs some simple ad-hoc automation software that integrates the Field systems attached to the station (e.g., a barcode reader, the smart nut runner) using standard IoT protocols like MQTT. The EG also runs a peer node that is a member of the logical Ledger Tier. A custom Ledger Service deployed on the Ledger Tier implements the business logic of the use case. The instruction set for the products to be processed is sent in JSON format to the Ledger Service, once per day, by the central ERP-MES systems: from that point and until a new production plan is published, the Ledger and Edge Tiers are autonomous.

When a new truck reaches the end of the main line it is dispatched to the first finishing line available, achieving the desired result of product flow optimization. Then, when it reaches the *wheel alignment station*, the chassis ID is scanned by a barcode reader and a request for instructions is sent, through the automation layer on the EG, to the Ledger Service. The Ledger Service will retrieve the instruction set from the production plan – which is saved on the Ledger itself – by matching the chassis ID. When the automation layer receives the instructions set, it parses the specific configuration parameters of interest and sends them to the nut runner, which adjusts itself. The wheel alignment operations will then proceed as usual. A record of the actual operations performed, which may differ from those in the instruction set, is finally set back to the Ledger and used to update the production plan.

An overall view of the use case is given in Fig. 3.

While the product flow optimization mentioned above is the immediate result of the pilot, there are some additional benefits to be gained either as a by-product or as planned extensions.

Firstly, the wheel alignment station – together with its EG box – becomes an autonomous module that can be easily added / removed and even relocated in a different environment. This scenario is not as far-fetched as it may seem, because it actually comes from a business requirement: the company has a number of production sites in different locations all over the world, each with their own unique MES maps. The deployment of a new module with different MES maps is currently a difficult and costly process.

Secondly, in the future the truck itself may become a Smart Object that communicates directly with the Ledger Tier. Truck-Ledger interactions will happen throughout the entire lifecycle of the truck – from manufacturing to operation and until decommissioning – with the Ledger maintaining a digital twin of the truck.

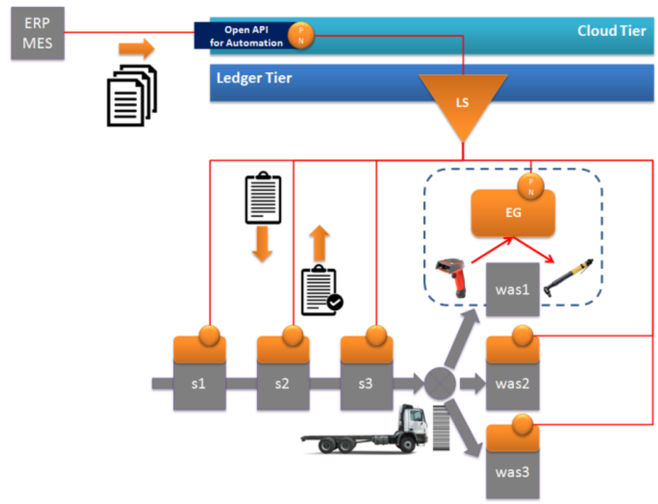

Figure 3 – Truck "mass-customization" use case

The second use case follows instead the *heavyweight* green-field approach. The pilot belongs to a white goods (i.e., domestic appliances) factory. The objective of the pilot is "reshoring", which in the FAR-EDGE context means enabling the company to move production back from offshore locations, thanks to a better support for the rapid deployment of new technologies (i.e., shopfloor Smart Objects) offered by the more advanced domestic plants. In this particular plant, a 1km. long conveyor belt moves pallets of finished products from the factory to a warehouse, where they are either stocked or forwarded for immediate delivery. The factory/warehouse conveyor is not only a physical boundary but also an administrative one, as the two facilities are under the responsibility of two different business units. Moreover, once the pallet is loaded on a delivery vehicle, it comes under the responsibility of a third party who operates the delivery business.

In the as-is scenario, the conveyor feeds 19 shipping bays, or "lanes", in the warehouse. Each lane is simply a dead-end conveyor segment, where pallets are dropped in by the conveyor and retrieved by a manually-operated forklift (basically, a FIFO queue). Simple mechanical actuators do the physical routing of the pallets, controlled by logic that runs on a central "sorter" PLC. The sorting logic is very simple: it is based on a production schedule that is defined once per day and on static mappings of the lanes to product types and/or final destinations. This approach has one big problem: production cannot be dynamically tuned to match business changes, or at least only to a very limited extent, because the fixed dispatching scheme downstream cannot sync with it. The problem is not only in software: the physical layout of the system is fixed.

In the to-be solution, the shipping bays become Smart Objects that can be plugged in and out at need (see Fig. 4). They embed simple sensors that detect the number of pallets currently in their local queue, and a controller board that runs some custom automation logic and connects directly to the Ledger Tier (i.e., without the mediation of an Edge Gateway). A custom Ledger Service acts as a coordination hub: it is responsible for authorizing a new “smart bay” that advertise itself to join the system (plug-and-produce) and, once accepted, to apply the sorting logic. This is based on the current state of the main conveyor belt, where incoming and outgoing pallets are individually identified by an RFID tag, and on “capability update” messages that are sent by smart bays each time they undergo an internal state change (e.g., number of free slots in the local queue, preference for a product type). The production schedule is not required at all, because sorting is only calculated on the actual state.

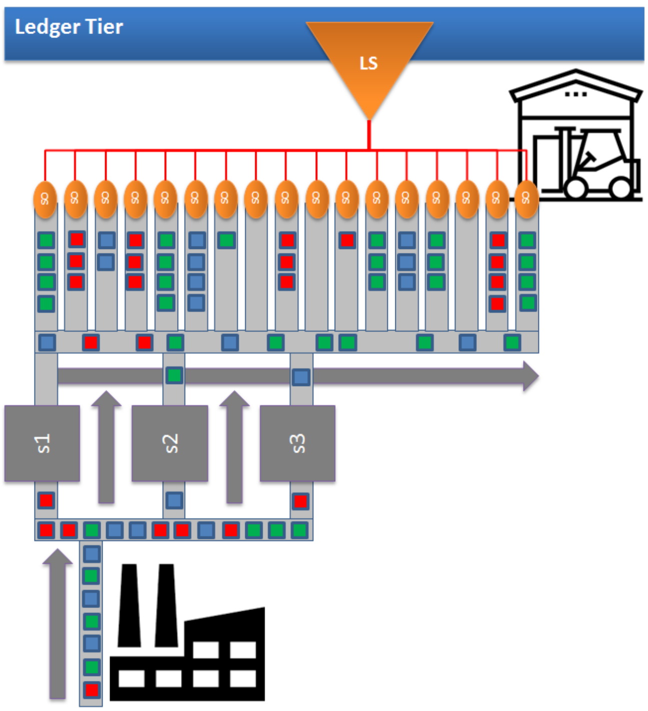

Figure 4 – White goods “reshoring” use case

VII. CONCLUSIONS

Within the last years several decentralized control architectures have been developed, based on cloud computing and web services, highlighting the benefit of decentralized automation in terms of flexibility of heterogeneous devices at the shopfloor. However, they are not yet fully deployed in real manufacturing environments. The FAR-EDGE project aims at supporting industries in their digital transformation towards Cyber-Physical Systems (CPS) for manufacturing and Industrial Internet of Things (IIoT) by providing them with an open platform for factory automation based on edge computing and IoT/CPS technologies, according to the vision of decentralizing factory automation.

The particularity of the FAR-EDGE RA presented in this paper is the introduction of the Ledger, as support of the Edge layer, based on Blockchain and Smart Contracts technologies for truly distributed process logic, i.e. secure state sharing rules, in decentralized CPPS. The paper also presents the general framework of the identified industrial pilot applications of FAR-EDGE RA for the validation phase within the project.

Future work will be focused on the implementation, test and validation of the Ledger, as well as the FAR-EDGE RA, to the project’s industrial use cases, planned by the third quarter of 2019.

ACKNOWLEDGMENT

Most of the work reported in this paper was done in the scope of the FAR-EDGE research and innovation project, funded by the European Commission’s Horizon 2020 program under grant agreement n. 723094.

REFERENCES

[1] M. Isaja, J. Soldatos, “Distributed Ledger Technology for Decentralization of Manufacturing Processes”, proceedings of the 1st IEEE International Conference on Industrial Cyber-Physical Systems (ICPS), 2018.

[2] M. Walker, "2017 Hype Cycles Highlight Enterprise and Ecosystem Digital Disruptions", Gartner, 2017.

[3] M. Isaja, J. Soldatos, “Combining Edge Computing and Blockchains for Flexibility and Performance in Industrial Automation”, proceedings of the 11th International Conference on Mobile Ubiquitous Computing, Services and Technology (UBICOMM), 2017.

[4] T. Sauter, S. Soucek, W. Kastner, D. Dietrich, “The evolution of factory and building automation”, IEEE Industrial Electronics Magazine, September 2011.

[5] M. Isaja, “FAR-EDGE achitecture and components specifications”, FAR-EDGE project public deliverable D2.4, 2017. Retrieved from: https://goo.gl/9wNYGA

[6] H. Halpin, M. Piekarska, "Introduction to Security and Privacy on the Blockchain," IEEE European Symposium on Security and Privacy Workshops (EuroS&PW), Paris, 2017, pp. 1-3.

[7] L. Lamport, R. Shostak, M. Pease, “The Byzantine Generals problem”, ACM Transactions on Programming Languages and Systems, volume 4 n. 3, p. 382-401, 1982.

[8] Z. Zheng, S. Xie, H. Dai, X. Chen, H. Wang, “An overview of Blockchain technology: architecture, consensus, and future trends”, proceedings of IEEE 6th International Congress on Big Data, 2017.

[9] T. Dinh, J. Wang, G. Chen, R. Liu, C. Ooi, K. L. Tan, “BLOCKBENCH: a framework for analyzing private Blockchains”, unpublished, 2017. Retrieved from: https://arxiv.org/pdf/1703.04057.pdf