Combining Edge Computing and Blockchains for Flexibility and Performance in Industrial Automation

Mauro Isaja
Research & Development
Engineering Ingegneria Informatica SpA (ENG)
Rome, Italy
e-mail: mauro.isaja@eng.it

John Soldatos
IoT Group
Athens Information Technology (AIT)
Maroussi, Greece
e-mail: jsol@ait.gr

Volkan Gezer
Innovative Factory Systems (IFS)
German Research Center for Artificial Intelligence (DFKI)
Kaiserslautern, Germany
e-mail: Volkan.Gezer@dfki.de

***Abstract* — The advent of Industry 4.0 has given rise to the introduction of new industrial automation architectures that emphasize the use of digital technologies. In this paper, we introduce a novel reference architecture (RA) for industrial automation, which leverages the benefits of edge computing, while using blockchain technologies for flexible, scalable and reliable configuration and orchestration of automation workflows and distributed data analytics. The presented RA is unique in blending the merits of blockchains and edge computing, while being compliant with emerging standards for industrial automation, such as RAMI4.0 and the RA of the Industrial Internet-Consortium.**

Keywords-Factory automation; edge computing; blockchain; RAMI4.0; IIRA; Industry4.0.

I. INTRODUCTION

The vision of future manufacturing foresees flexible and hyper-efficient plants that will enable manufacturers to support the transition from conventional "made-to-stock" production models, to the emerging customized ones such as "made-to-order", "configure-to-order" and "engineering-to-order". Flexibility in automation is a key prerequisite to supporting the latter production models, as it facilitates manufacturers to change automation configurations and rapidly adopt new automation technologies, as a means of supporting variation in production without any essential increase in production costs.

In order to support flexibility in automation, the industrial automation community has been exploring options for the virtualization of the automation pyramid, as part of the transformation of mainstream centralized automation models (like ISA-95) to more distributed ones. Several research and development initiatives have introduced decentralized factory automation solutions based on technologies like intelligent agents [1] [2] and Service Oriented Architectures (SOA) [3] [4]. These initiatives produced proof-of-concept implementations that highlighted the benefits of decentralized automation in terms of flexibility. However, they are still not being widely deployed in manufacturing plants, mainly due to that the cost-benefit ratio of such solutions is perceived as unfavourable. Nevertheless, the vision of decentralizing the factory automation pyramid is still alive, as this virtualization can potentially make production systems more flexible and agile, increase product quality and reduce cost.

With the advent of the fourth industrial revolution (Industry 4.0) and the Industrial Internet of Things (IIoT), decentralization is being revisited in the light of the integration of Cyber-Physical Systems (CPS) with cloud computing infrastructures. Therefore, several cloud-based applications are deployed and used in factories, which leverage the capacity and scalability of the cloud while fostering supply chain collaboration and virtual manufacturing chains. Early implementations have also revealed the limitations of the cloud in terms of efficient bandwidth usage and its ability to support real-time operations, including operations close to the field.

More recently, the edge computing paradigm has been explored in order to alleviate the limitations of cloud-centric architectures. Edge computing architectures move some part of the system's overall computing power from the cloud to its edge nodes, i.e., on the field or in close proximity to it –as a means of [5], [6]:

- Saving bandwidth and storage, as edge nodes can filter data streams from the field in order to get rid of information without value for industrial automation.
- Enabling low-latency and proximity processing, since information can be processed close to the field.
- Providing enhanced scalability, through supporting decentralized storage and processing that scales better than cloud processing.
- Supporting shopfloor isolation and privacy-friendliness, since edge nodes at the shopfloor are isolated from the rest of the network.

These benefits make edge computing suitable for specific classes of use cases in factories, including:

- Large scale distributed applications, typically applications that involve multiple plants or factories, which process streams from numerous devices at scale.
- Near-real-time applications, which analyse data close to the field or even control Cyber-Physical Systems such as smart machines and industrial robots.

As a result, the application of edge computing to factory automation is extremely promising, since it empowers decentralization in a way that still supports real-time interactions and scalable analytics. It's therefore no accident that there are ongoing efforts to provide edge computing implementations for industrial automation in general and factory automation in particular. Furthermore, reference architectures for IIoT and industrial automation exist, which highlight the importance of edge computing for compliant implementations. In this article, we present a reference architecture (RA) for factory automation based on edge computing, which has been specified as part of the H2020 FAR-EDGE project [11]. The FAR-EDGE RA and associated compliant implementations comprise some unique features and capabilities, which differentiate them from other on-going implementations of edge computing for factory automation. Most of these unique features concern the exploitation of Distributed Ledger Technology (DLT, today commonly referred to as "blockchain") as a means of representing automation and data analytics processes based on Smart Contracts. These can be dynamically configured, stored securely and executed in a distributed way, enabling flexibility and scalability in factory automation processes.

The paper is structured as follows: Section 2, following this introduction, presents state-of-the-art specifications and implementations of the edge computing paradigm for factory automation. It also positions FAR-EDGE against them. Section 3 introduces the FAR-EDGE RA, from a functional and structural perspective. Section 4 illustrates a number of automation use cases and the way in which they can be supported by FAR-EDGE compliant systems. Finally, Section 5 concludes the paper.

II. Related Work

Acknowledging the benefits of edge computing for industrial automation, standards development organizations (SDOs) have specified relevant reference architectures, while industrial organizations are already working towards providing tangible edge computing implementations.

SDOs such as the OpenFog Consortium and the Industrial Internet Consortium (IIC) have produced Reference Architectures. The RA of the OpenFog Consortium prescribes a high-level architecture for internet of things systems, which covers industrial IoT use cases. On the other hand, the RA of the IIC [7] outlines the structuring principles of systems for industrial applications. The IIC RA is not limited to edge computing, but rather based on edge computing principles in terms of its implementation. It addresses a wide range of industrial use cases in multiple sectors, including factory automation. These RAs have been recently released and their reference implementations are still in their early stages.

A reference implementation of the IIC RA's edge computing functionalities for factory automation is provided as part of IIC's edge intelligence testbed [8]. This testbed provides a proof-of-concept implementation of edge computing functionalities on the shopfloor. The focus of the testbed is on configurable edge computing environments, which enable the development and testing of leading edge systems and algorithms for edge analytics. Moreover, Dell-EMC has recently announced the EdgeX Foundry framework [9], which is a vendor-neutral open source project hosted by the Linux Foundation that builds a common open framework for IIoT edge computing. The framework is influenced by the above-listed reference architectures and is expected to be released in 2017. Other vendors are also incorporating support for edge devices and edge gateways in their cloud platforms.

FAR-EDGE is uniquely positioned in the landscape of edge computing solutions for factory automation. In particular, the FAR-EDGE architecture is aligned to the IIC RA, while exploiting concepts from other RAs and standards such as the OpenFog RA and RAMI 4.0 (Reference Architecture Model Industry 4.0) [10]. However, FAR-EDGE explores pathways and offers functionalities that are not addressed by other specification and reference implementations. In particular, it researches the applicability of disruptive key enabling technologies like DLT and Smart Contracts in factory automation. DLT, while being well understood and thoroughly tested in mission-critical areas like digital currencies (e.g., Bitcoin), have never been applied before to industrial systems. FAR-EDGE aims at demonstrating how a pool of specific Ledger Services built on a generic DLT platform can enable decentralized factory automation in an effective, reliable, scalable and secure way. Ledger Services will be responsible for sharing process state and enforcing business rules across the computing nodes of a distributed system, thus permitting virtual automation and analytics processes that span multiple nodes – or, from a bottom-up perspective, autonomous nodes that cooperate to a common goal. This is the project's unique contribution, which sets it apart from similar efforts worldwide.

III. FAR-EDGE RA Oveview

The FAR-EDGE RA is the conceptual framework that drives the design and the implementation of the project's automation platform based on edge computing and DLT technologies. As an RA, its first goal is communication, i.e. providing a terse representation of concepts, roles, structure and behaviour of the system under analysis for the sake of dissemination and ecosystem-building. Its second goal concerns reuse: exploiting best practices and lessons learned in similar contexts by the global community of system architects.

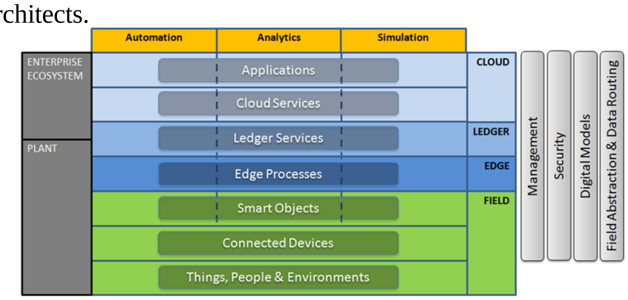

Figure 1. Overview of the FAR-EDGE RA

The FAR-EDGE RA is aligned to IIC's RA concepts and described from two architectural viewpoints: the functional viewpoint and the structural viewpoint, as outlined in following paragraphs.

An overall architecture representation that includes all elements is provided in Figure 1.

Copyright (c) IARIA, 2017. ISBN: 978-1-61208-598-2

A. *Functional Viewpoint*

According to the FAR-EDGE RA, the functionality of a factory automation platform can be decomposed into three high-level Functional Domains - Automation, Analytics and Simulation – and four Crosscutting (XC) Functions – Management, Security, Digital Models and Field Abstraction & Data Routing. To better clarify the scope of such topics, we have tried to map them to similar IIRA concepts. Functional Domains and XC Functions are orthogonal to structural Tiers: the implementation of a given functionality may – but is not required to – span multiple Tiers, so that in the overall architecture representation Functional Domains appear as vertical lanes drawn across horizontal layers. In Figure 2, the relationship between Functional Domains, their users and the factory environment is highlighted by arrows showing the flow of data and of control.

Automation Domain: The FAR-EDGE Automation domain includes functionalities supporting automated control and automated configuration of physical production processes. While the meaning of "control" in this context is straightforward, "configuration" is worth a few additional words. Automated configuration is the enabler of plug-and-play factory equipment (better known as plug-and-produce), which in turn is a key technology for mass-customization, as it allows a faster and less expensive adjustments of the production process. The Automation domain requires a bidirectional monitoring/control communication channel with the Field, typically with low bandwidth but very strict timing requirements (tight control loop). In some advanced scenarios, Automation is controlled – to some extent – by the results of Analytics and/or Simulation. The Automation domain partially maps to the Control domain of the IIRA.

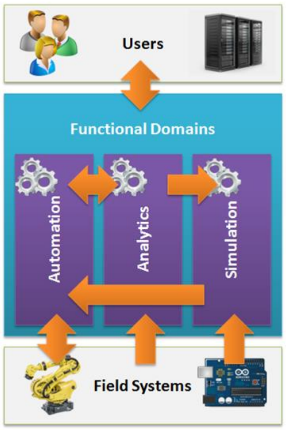

Figure 2. FAR-EDGE RA Functional Domains

Analytics Domain: The FAR-EDGE Analytics domain includes functionalities for gathering and processing Field data for a better understanding of production processes, i.e. a factory-focused business intelligence. This typically requires a high-bandwidth Field communication channel, as the volume of information that needs to be transferred in a given time unit may be substantial. On the other hand, channel latency tends to be less critical than in the Automation scenario. The Analytics domain provides intelligence to its users, but these are not necessarily limited to humans or vertical applications (e.g., a predictive maintenance solution): the Automation and Simulation domains, if properly configured, can both make direct use of the outcome of data analysis algorithms. In the case of Automation, the behaviour of a workflow might change in response to changes detected in the controlled process – e.g., a process drift caused by the progressive wear of machinery or by the quality of assembly components being lower than usual. In the case of Simulation, data analysis can be used to update the parameters of a digital model (as illustrated in the following section). The Analytics domain matches perfectly the Information domain of the IIRA, except that the latter is receiving data from the Field through the mediation of Control functionalities.

Simulation Domain: The FAR-EDGE Simulation domain includes functionalities for simulating the behaviour of physical production processes for the purpose of optimization or of testing what/if scenarios at minimal cost and risk and without any impact of regular shop activities. Simulation requires digital models of plants and processes to be in-sync with the real world objects they represent. As the real world is subject to change, models should reflect those changes. For instance, the model of a machine assumes a given value of electric power / energy consumption, but the actual values will diverge as the real machine wears down. To detect this gap and correct the model accordingly, raw data from the Field (direct) or complex analysis algorithms (from Analytics) can be used.

Crosscutting Functions: Crosscutting Functions address common specific concerns. Their implementation affects several Functional Domains and Tiers. They include.

- **Management**: Low-level functions for monitoring and commissioning/decommissioning of individual system modules..
- **Security**: Functions securing the system against the unruly behaviour of its user and of connected systems. These include digital identity management and authentication, access control policy management and enforcement, communication and data encryption.
- **Digital Models**: Functions for the management of digital models and their synchronization with the real-world entities they represent. Digital modes are a shared asset, as they may be used as the basis for automated configuration, simulation and field abstraction – e.g., semantic interoperability of heterogeneous field systems.
- **Field Abstraction & Data Routing**: Functions that ensure the connectivity of business logic (FAR-EDGE RA Functional Domains) to the Field, abstracting away the technical details – like device discovery and communication protocols. Data routing refers to the capability of establishing direct producer-consumer channels on demand, optimized for unidirectional massive data streaming – e.g., for feeding Analytics.

 ISBN: 978-1-61208-598-2

B. *Structural Viewpoint*

The FAR-EDGE RA uses two classes of concepts for describing the structure of a system: Scopes and Tiers.

Scopes are very simple and straightforward: they define a coarse mapping of system elements to either the factory - Plant Scope - or the broader world of corporate IT - Enterprise Ecosystem Scope. Examples of elements in Plant Scope are machinery, Field devices, workstations, SCADA and MES systems, and any software running in the factory data centre. The Enterprise Ecosystem Scope comprises ERP and PLM systems and any application or service shared across multiple factories or even companies – e.g., supply chain members.

Tiers are a more detailed and technical-oriented classification of deployment concerns. They can be easily mapped to scopes, but they provide more insight into the relationship between system components. This kind of classification is quite similar to OpenFog RA deployment viewpoint, except for the fact that FAR-EDGE Tiers are industry-oriented while OpenFog ones are not. FAR-EDGE Tiers are one of the most innovative traits of its RA, and are described in following paragraphs.

The Field Tier is the bottom layer of the FAR-EDGE RA and is populated by Edge Nodes (EN), i.e. any kind of device that is connected to the digital world on one side and to the real world to the other. ENs can have embedded intelligence (e.g., a smart machine) or not (e.g., a sensor or actuator). The FAR-EDGE RA honours this difference: Smart Objects are ENs with on board computing capabilities, Connected Devices are those without. The Smart Object is where local control logic runs: it's a semi-autonomous entity that does not need to interact frequently with the upper layers of the system. As shown in Figure 3. ENs is actually located over field devices.

The Field is also populated by entities of the real world, i.e., those physical elements of production processes that are not directly connected to the network, and as such are not considered as ENs: Things, People and Environments. These are represented in the digital world by some kind of EN wrapper. For instance, room temperature (Environment) is measured by an IoT sensor (Connected Device), the proximity of a worker (People) to a physical checkpoint location is published by an RFID wearable and detected by an RFID Gate (Connected Device), while a conveyor belt (Thing) is operated by a PLC (Smart Object).

The Field Tier is in Plant Scope. Individual ENs are connected to the digital world in the upper Tiers either directly by means of the shopfloor's LAN, or indirectly through some special-purpose local network (e.g., WSN) that is bridged to the former. From the RAMI 4.0 perspective, the FAR-EDGE Field Tier corresponds to the Field Device and Control Device levels on the Hierarchy axis (IEC-62264/IEC-61512), while the entities there contained are positioned across the Asset and Integration Layers.

The Edge Tier is the core of the FAR-EDGE RA. It hosts those parts of Functional Domains and XC Functions that can leverage the edge computing model, i.e., software designed to run on multiple, distributed computing nodes placed close to the field, which may include resource constrained nodes. The Edge Tier is populated by Edge Gateways (EG): computing devices that act as a digital world gateway to the real world of the Field. These machines are typically more powerful than the average intelligent EN (e.g., blade servers) and are connected to a fast LAN. Strategically positioned close to physical systems, the EG can execute Edge Processes: time- and bandwidth-critical functionality having local scope. For instance, the orchestration of a complex physical process that is monitored and operated by a number of sensors, actuators (Connected Devices) and embedded controllers (Smart Objects); or the real-time analysis of a huge volume of live data that is streamed from a nearby Field source.

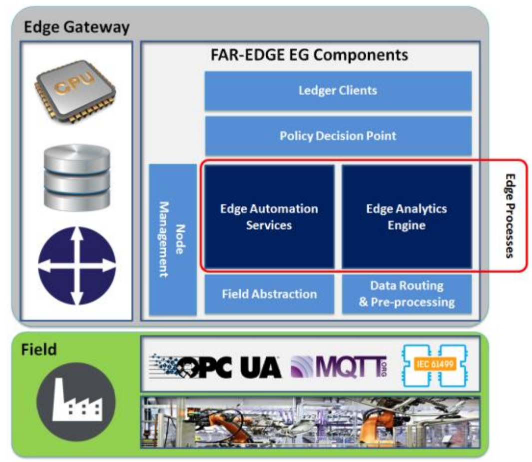

Figure 3. Edge Tier in the FAR-EDGE RA

Deploying computing power and data storage in close proximity to where it is actually used is a standard best practice in the industry. However, this technique basically requires that the scope of individual subsystems is narrow (e.g., a single work station). If instead the critical functionality applies to a wider scenario (e.g., an entire plant or enterprise), it must be either deployed at a higher level (e.g., the Cloud) – thus losing all benefits of proximity – or run as multiple parallel instances, each focused on its own narrow scope. In the latter case, new problems may arise: keeping global variables in-sync across all local instances of a given process, reaching a consensus among local instances on a global truth, collecting aggregated results from independent copies of a data analytics algorithm, etc. The need for peer nodes of a distributed system to mutually exchange information is recognized by the OpenFog RA. The innovative approach in FAR-EDGE is to define a specific system layer – the Ledger Tier – that is responsible for the implementation of such mechanisms and to guarantee an appropriate Quality of Service level.

The Edge Tier is in Plant Scope, located above the Field Tier and below the Cloud Tier. Individual EGs are connected with each other and with the north side of the system, i.e., the globally-scoped digital world in the Cloud Tier – by means of the factory LAN, and to the south side through the shopfloor LAN. From the RAMI 4.0 perspective, the FAR-

EDGE Edge Tier corresponds to the Station and Work Centre levels on the Hierarchy axis (IEC-62264/IEC-61512), while the EGs there contained are positioned across the Asset, Integration and Communication Layers. Edge Processes running on EGs, however, map to the Information and Functional Layers.

The Ledger Tier is a complete abstraction: it does not correspond to any physical deployment environment, and even the entities that it "contains" are abstract. Such entities are Ledger Services, which implement decentralized business logic as smart contracts on top of a distributed ledger. Ledger Services are transaction-oriented: each service call that needs to modify the shared state of a system must be evaluated and approved by Peer Nodes before taking effect. Similarly to "regular" services, Ledger Services are implemented as executable code; however, they are not actually executed on any specific computing node: each service call is executed in parallel by all Peer Nodes that happen to be online at the moment, which then need to reach a consensus on its validity. Most importantly, even the executable code of Ledger Services can be deployed and updated online by means of a distributed ledger transaction.

Ledger Services implement the part of Functional Domains and/or XC Functions that enable the edge computing model, through providing support for their Edge Service counterpart. For example, the Analytics Functional Domain may define a local analytics function (Edge Service) that must be executed in parallel on several EGs, and also a corresponding service call (Ledger Service) that will be invoked from the former each time new or updated local results become available, so that all results can converge into an aggregated data set. In this case, aggregation logic is included in the Ledger Service. Another use case may come from the Automation Functional Domain, demonstrating how the Ledger Tier can also be leveraged from the Field: a smart machine with embedded plug-and-produce functionality can ask permission to join the system by making a service call and then, having received green light, can dynamically deploy its own specific Ledger Service for publishing its state and external high-level commands.

The Ledger Tier lays across the Plant and the Enterprise Ecosystem Scopes, as it can provide support to any Tier. The physical location of Peer Nodes, which implement smart contracts and the distributed ledger, is not defined by the FAR-EDGE RA as it depends on implementation choices.

From the RAMI 4.0 perspective, the FAR-EDGE Ledger Tier corresponds to the Work Centre, Enterprise and Connected World levels on the Hierarchy axis (IEC-62264/IEC-61512), while the Ledger Services are positioned across the Information and Functional Layers.

The Cloud Tier is the top layer of the FAR-EDGE RA, and also the simplest and more "traditional" one. It is populated by Cloud Servers (CS): powerful computing machines, sometimes configured as clusters, which are connected to a fast LAN internally to their hosting data centre, and made accessible from the outside world by means of a corporate LAN or the Internet. On CSs runs that part of the business logic of Functional Domains and XC Functions that benefits from having the widest of scopes over production processes, and can deal with the downside of being physically deployed far away from them. This includes the planning, monitoring and management of entire factories, enterprises and supply chains (e.g., ERP and SCM systems). The Cloud Tier is populated by Cloud Services and Applications. Cloud Services implement specialized functions that are provided as individual API calls to Applications, which instead "package" a wider set of related operations that are relevant to some higher-level goal and often expose an interactive human interface.

The Cloud Tier is in Enterprise Ecosystem scope. The "Cloud" term in this context implies that Cloud Services and Applications are visible from all Tiers, wherever located. It does not imply that CSs should be actually hosted on some commercial cloud. In large enterprises, the Cloud Tier corresponds to one or more corporate data centres (private cloud), ensuring that the entire system is fully under the control of its owner.

In terms of RAMI 4.0, the FAR-EDGE Cloud Tier corresponds to the Work Centre, Enterprise and Connected World levels on the Hierarchy axis (IEC-62264/IEC-61512), while the Cloud Services and Applications are positioned across the Information, Functional and Business Layers.

IV. REFERENCE USE CASES

In following paragraphs we present some indicative use cases that will be supported by FAR-EDGE.

A. Wheel Alignment Smart Station

This scenario is centred around the concept of an autonomous cyber-physical system (CPS): a self-contained plant module (workstation) comprising smart machines/tools and locally-scoped monitoring/control logic. Such module operates as a block-box: internally, it implements automated machine/tool configuration and workflows; externally, it integrates with the factory's IT backbone (e.g., MES/ERP) by means of a "public" interface that provides the required functionality while hiding the module's internals.

The concrete use case that the FAR-EDGE project is developing in this scenario targets a wheel alignment workstation for the manufacturing of industrial vehicles. The use case is complex, as it builds on a production process that is currently in place: its full description would go beyond the scope of this paper. To summarise, the added value of introducing the FAR-EDGE platform is this context is twofold. Firstly, it enables smart tools, i.e., an IoT-ready nut driver – to be dynamically deployed on any physical workstation and to be timely reconfigured (torque adjustment) to match fast-changing requirements, as a wide array of truck models is processed along the same production line. Secondly, it allows the entire workstation to be easily relocated to other plants that share the same IT backbone. As a positive side effect, the workstation, being mostly autonomous, is also able to operate with little or no disruption when temporarily disconnected from the network.

According to the FAR-EDGE RA, locally-scoped automation and analytics are Edge Processes belonging to the Edge Tier. From the implementation perspective, Edge

Processes are hosted by an Edge Gateway, which is an integral part of the wheel alignment workstation. In particular, smart tool deployment is in charge of the Edge Automation Services (EAS) component, which communicates with local field devices through the Field Abstraction layer. EAS also interacts with a digital model of the plant in order to retrieve and update information.

B. *Plug-and-Produce Conveyor Belt*

As in the previous example, the foundation of the FAR-EDGE use case is an existing logistic process in a real-world factory. The scenario is that of a large production plant where finished products, stacked on pallets, are moved by a single conveyor belt to a warehouse. Pallets only contain product items of the same type, but each pallet can be different as the conveyor is the outlet of multiple assembly lines working in parallel; the exact product type sequence on the conveyor at any given time is not predictable. When the pallets reach the warehouse, they are dispatched to a number of "exit bays" for immediate shipping, temporary storage or other destinations (e.g., defective products). The dispatching logic should take into account product type on the one hand, bay configuration, capacity and status on the other. In its current implementation, a PLC-based dispatching system does its best to match the input stream (product type ID scanned on pallets) with the output channels (static configuration of exit bays), taking into account the daily production schedule. However, this approach does not allow for any significant schedule change and/or "hot" reconfiguration of the exit bays.

The FAR-EDGE platform is redesigning the above described "primitive" CPS with the introduction of Smart Objects (exit bays with embedded computing power and network connectivity) and of a Ledger Tier (a Distributed Ledger exposing Ledger Services) where *decentralized* configuration and orchestration logic resides. The basic use case is Plug-and-Produce: new bays can be added to the working system, and existing bays can be put offline, at any time: the Ledger Tier is responsible for granting permission and for keeping the digital model of the plant in-sync with the real world. Once online, new bays are immediately able to negotiate with the plant controller their services – e.g., ask for more products of a given type when *actual* processing capacity exceeds the incoming flow. The innovative approach in FAR-EDGE, where blockchain is used to implement Ledger Services, avoids potential single-point-of-failure problems and scalability bottlenecks.

V. CONCLUSIONS

The edge computing paradigm provides many compelling advantages for the implementation of digital automation platforms, including the ability to analyze information close to the field, as well as the ability to flexible (re)configure real-time automation workflows. This is the reason why several edge computing platforms for industrial automation are already under implementation. FAR-EDGE takes these implementations to the next level, through enhancing edge computing implementations with the merits of blockchain technologies, notably in terms of representing and implementing automation and analytics operations as scalable and flexibly configurable smart contracts. Blockchain concepts have already been introduced in the FAR-EDGE RA, which serves as a basis for implementing automation, analytics and digital simulation use cases. In addition to providing open source implementation of FAR-EDGE systems, our project will provide tangible research findings regarding the applicability of blockchain for factory automation.

ACKNOWLEDGMENT

This work has been carried out in the scope of the FAR-EDGE project (H2020-703094). The authors acknowledge help and contributions from all partners of the project.

REFERENCES

[1] P. Leitão, "Agent-based distributed manufacturing control: A state-of-the-art survey", Engineering Applications of Artificial Intelligence, vol. 22, no. 7, pp. 979-991, Oct. 2009.

[2] P. Vrba,., Review of Industrial Applications of Multi-agent Technologies," Service Orientation in Holonic and Multi Agent Manufacturing and Robotics, Studies in Computational Intelligence Vol. 472, Springer, pp 327-338, 2013

[3] F. Jammes and H. Smit, "Service-Oriented Paradigms in Industrial Automation Industrial Informatics," IEEE Transactions on., pp. 62 – 70, vol. 1, issue 1, Feb, 2005.

[4] T. Cucinotta and Coll, "A Real-Time Service-Oriented Architecture for Industrial Automation," Industrial Informatics, IEEE Transactions on, vol. 5, issue 3, pp. 267 – 277, Aug. 2009.

[5] W. Shi, J. Cao, Q. Zhang, Y. Li and L. Xu, "Edge Computing: Vision and Challenges," in IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637-646, Oct. 2016. doi: 10.1109/JIOT.2016.2579198

[6] F. Bonomi, R. Milito, J. Zhu, S. Addepalli, «Fog computing and its role in the internet of things», Proceedings of the first edition of the MCC workshop on Mobile cloud computing, MCC '12, pp 13-16.

[7] Industrial Internet Consortium. 2017. The Industrial Internet of Things Volume G1: Reference Architecture, version 1.8. (2017). [Online], Available from: http://www.iiconsortium.org/IIRA.htm

[8] Industrial Internet Consortium. IIC Edge Intelligence Testbed. 2017. [Online], Available from: http://www.iiconsortium.org/edge-intelligence.htm

[9] EdgeX Foundry Framewok 2017. [Online], Available from: https://www.edgexfoundry.org/

[10] K. Schweichhart. "Reference Architectural Model Industrie 4.0 - An Introduction", April 2016, [Online], Available from: https://ec.europa.eu/futurium/en/system/files/ged/a2-schweichhart-reference_architectural_model_industrie_4.0_rami_4.0.pdf

[11] H2020-703094 FAR-EDGE Project 2017. [Onine], Available from: http://www.far-edge.eu