5G-INDUCE Platform: Deploying Network Applications in Industry 4.0

Dimitrios Klonidis[1], Franco R. Davoli[2], Nikolaos Pitropakis[3], Georgios P. Katsikas[1], Thanos Xirofotos[1], Roberto Bruschi[2], Chiara Lombardo[2], David Gomez-Barquero[4], Qi Wang[5], Jose Maria Alcaraz Calero[5], Ioannis Tomkos[6], Alexios Birbas[6], Ioannis Giannoulakis[3], Vasileios Samarinas[3], Emmanouil Kafetzakis[3], Dimitrios Nodaros[3], Panagiotis Gouvas[1]

[1]UBITECH LTD, Athens, Greece
[2]University of Genoa – CNIT, Genoa, Italy
[3]Eight Bells LTD, Athens, Greece
[4]Universitat Politecnica de Valencia, Valencia, Spain
[5]University of the West of Scotland, Paisley, United Kingdom
[6]University of Patras, Rio-Patra, Greece
dklonidis@ubitech.eu

Abstract— The successful establishment of 5G technology and the move beyond this requires vertical stakeholders and Network Application developers to be able to smoothly deploy and manage applications in distributed 5G network environments, in a secure fashion and with strict KPI requirements. The innovative solutions within 5G-INDUCE rely on the deployment of an open, ETSI NFV compatible, 5G orchestration platform for the deployment of advanced 5G Network Applications (NetApps). The platform's unique features provide the capability to the NetApp developers to define and modify the application requirements while the underlay intelligent OSS can expose the network capabilities to the end users on the application level without revealing any infrastructure-related information. This process enables an application-oriented network management and optimization approach that is in line with the operator's role as manager of its own facilities, while it offers the operational environment to any developer and service provider through which tailor-made applications can be designed and deployed, for the benefit of vertical industries and without any indirect dependency on a cloud provider. The efforts focus on the Industry 4.0 vertical sector, as one of the fastest growing and most impactful sectors in European economy with high potentials for service development SMEs and with the capability to tackle all diverse cases of service requirements. The separation of concerns between application- and network-level orchestration is in line with the trend toward automated and dynamic slice management that will be required in the evolution toward 6G.

Keywords—5G, Orchestration, Network Applications

I. Introduction

The 4th industrial revolution (Industry 4.0) combines physical and digital technologies under connected system infrastructures, gaining critical operation insights to improve efficiency in manufacturing productivity and quality, but also work safety, environmental protection, and supply chain optimization. The realization of Industry 4.0 does not simply rely on locally installed smart IoT monitoring and automated control technologies but on the combination of emerging technologies over smartly interconnected and collaborating infrastructures [1]. 5G comes as a unique enabler to meet the strict networking requirements set in Industry 4.0, allowing the launching of unique Network Applications (NetApps) tailored to the infrastructure technology capabilities and the requested service provisioning requirements. This is an essential requirement for any extension beyond 5G that envisions the leading role of the vertical stakeholders in the management of their digital processes/services, along with the independent growth of the infrastructure owners, in an environment defined by security.

Towards this direction and in addition to the design and development of 5G platforms and ecosystems that enable infrastructure virtualization and support of edge processing, smooth and costless porting of NetApps in 5G ecosystems needs to be supported too, thus making the Industry 4.0 sector (and in principle any emerging vertical sector) ready to exploit its full potentials, such as production automation, robotics technologies, smart logistics and metering [2]. The successful porting of NetApps requires the attractiveness of vertical service developers and providers (being currently mostly activated in the cloud), and their ability to easily deploy new tailored features, such as performance monitoring and optimization, cognitive decision making and enhanced cyber security capabilities. In turn, the challenging performance and operating requirements of the new class of NetApps should be efficiently supported by the underlying platform, embracing recent well-known technologies, like Network Functions Virtualization (NFV) and Multi-access Edge Computing (MEC) that transform network operators' infrastructures into distributed datacenters with advanced virtualization and software-driven capabilities.

In the following sections, we primarily present the conceptualization of the NetApp-aware 5G system platform along with its overall architecture. Finally, our work sheds light on the implementation plans for the NetApp management

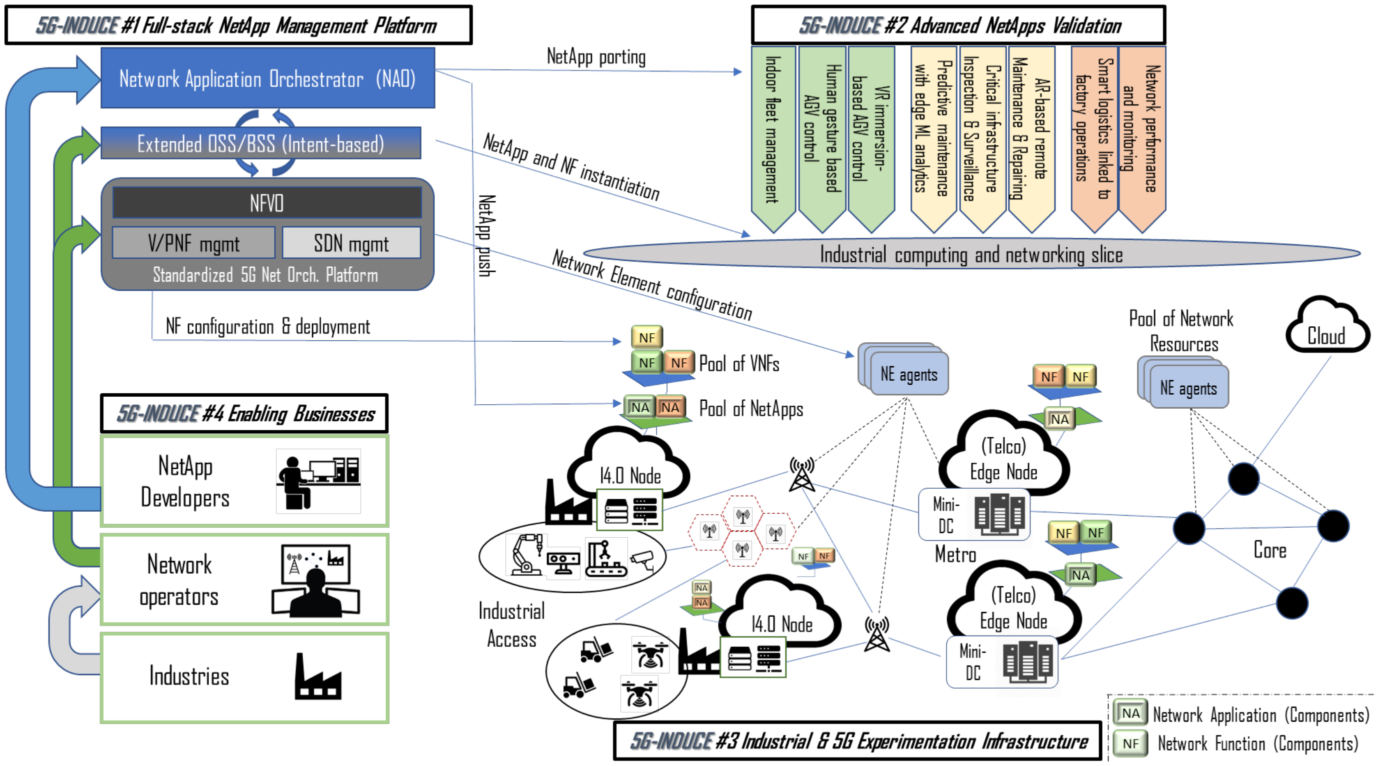

Fig. 1. The 5G-INDUCE vision and the adopted approach for easy (i) porting and/or (ii) development of industry 4.0 NetApps over advanced Experimentation Facilities combining real 5G and private industrial networks

platform which constitutes the core part of the overall system and the functional interfacing block between the NetApps and the deployment and management of network and computational resources, in a vision oriented to the evolution toward enhanced 6G capabilities.

II. Main Concept

In 5G-INDUCE, the vision mentioned above is realized with the deployment of two linked frameworks:

- a NetApp Orchestration (NAO) framework, interfacing with the end users and serving their Application requirements and
- a NFV Orchestration (NFVO) framework that manages the distributed resources over the virtualized infrastructure.

In this architecture, the network management part is typically owned by the telecom operator, which is reluctant to expose any infrastructure details to end users; however, end users require such information for optimized application deployment. Therefore, an intelligent Operations Support System (OSS) is typically required in the middle and provides the interface that translates Application requests into Network connectivity and resource allocation requirements. The NAO is integrated with the OSS extending its capabilities to the interfacing with the vertical end users and the application developers, while operators gain the ability to apply polices and any intelligent data analytics methodology at the service layer. Therefore, such a scheme maintains the edge resource management capabilities of the standardized ETSI NFVO framework, while enabling an application-oriented network management and optimization approach.

We note explicitly that this separation of concerns between application- and network-level orchestration is in line with the vision of future directions toward 6G (as outlined, e.g., in [3], [4]). In particular, this architectural vision is matched by Fig. 6 in reference [3], advocating an increasing presence of autonomic network management and automated (dynamic) network slicing, which will be also supported by Artificial Intelligence / Machine Learning (AI/ML) algorithms [5], [6].

III. Architecture

The overall 5G-INDUCE system architecture follows a four-axis structure as depicted in Fig. 1. These four axes highlight the core activities of the project:

Axis #1 – Full-stack NetApp Management Platform: Takes state-of-the-art research technologies, developed and tested in the previous 5G PPP research project initiatives, to a substantially higher technology-readiness level required by the Industry 4.0 standards, while at the same time offering the first ecosystem for advanced overlay 5G NetApps. Specifically, the 5G-INDUCE platform integrates state-of-the-art control and data plane developments ranging from industrial IoT devices and 5G radio access nodes to a state-of-the-art OSS and a scalable microservices-based cloud orchestration platform, each with a clear evolution roadmap towards an integrated full-stack 5G NetApp management platform.

Axis #2 - Advanced Industry 4.0 NetApps: Showcasing the support of a variety of innovative Industry 4.0 market verticals through the demonstration of advanced use cases, meeting demanding Industry 4.0 and 5G KPIs, such as ultra-low latency, fast service deployment, high service reliability, while accommodating industrial-grade scalability demands.

Axis #3 – Industrial & 5G Experimentation Infrastructure: Heterogeneous real-world industrial fields at scale provided by both global (with active presence in Europe) and European industrial players, on key industrial sectors, such as automotive, energy, and home appliances. 5G-INDUCE partners the aforementioned industrial players together with large European mobile network operators and vendors, aiming to build a unique fabric of tightly interconnected 5G infrastructures with industrial fields, which will lead the Industry 4.0 activities in Europe. Prior to the deployment of 5G-INDUCE NetApps on the real experimentation infrastructure, 5G-INDUCE utilizes a state-of-the-art 5G platform integrator and NetApp DevOps testbed for pre-deployment testing and validation.

Axis #4 – End Users (Enabling Businesses): Creates huge impact for key market stakeholders through a win-win-win (win3) business model that generates mutual benefits for: (i) industrial players, (ii) network operators, and (iii) a large pool of heterogeneous SMEs/start-ups acting as NetApp developers, while maintaining discrete and distinct roles for each of these stakeholders.

By putting all axes together, 5G-INDUCE, enables a full-stack NetApp Management Platform (axis #1), which efficiently facilitates the development of advanced Industry 4.0 NetApps (axis #2) over smart large-scale 5G and Industry 4.0 experimentation facilities (axes #3), while establishing a dramatically better and sustainable business model (axis #4) that involves all the key stakeholders, while maintaining discrete and distinct roles for each of them.

IV. THE NETAPP MANAGEMENT PLATFORM

This section provides technical insights about the design and implementation of the 5G-INDUCE axis #1, i.e., the full-stack 5G-INDUCE NetApp Management platform, which consists of two main parts: (i) an underlying network platform realized as an advanced 5G-INDUCE OSS and (ii) the overlay 5G-INDUCE NAO. Both parts are illustrated in Fig. 2 and described in turn below. An additional objective of this section is to highlight which parts of the full-stack 5G-INDUCE NetApp Management platform derive from previous 5G initiatives and/or existing standards (blue boxes in Figure 2), what extensions are required (yellow boxes in Figure 2), and which additional components are added (green boxes in Figure 2), in order to satisfy the 5G-INDUCE requirements.

A. The 5G-INDUCE OSS

This section presents the design and technical characteristics of the 5G-INDUCE OSS, which acts as a glue between the vertical stakeholder domain (i.e., the NAO) and the well-known building blocks present in production telco platforms (e.g., in accordance with ETSI NFVO standards [7]).

Previous related OSS solutions, as for example the one from the EC H2020 project MATILDA [8], have demonstrated rather simple network slices associated with static compute resources (i.e., for fixed application graphs). These features comprise a useful basis to begin with, but they fall short to satisfy the advanced requirements imposed by the envisioned NetApps and their broad range of networking features. The main implementation goals for the 5G-INDUCE OSS are:

- To develop an OSS Northbound (NB) API to manage the lifecycle of network slices and (edge) compute/storage resources. In the case of MATILDA this information is "hidden" by the OSS, as both the allocation and release of network slices and resources are static.
- To design a new OSS NB API that will enable advanced fully automated low-level operations of the OSS over 5G (and Beyond) network slices. This new API, combined with the revised lifecycle management above, will permit on-the-fly modification of a network slice's resources, including dynamic traffic steering and on-demand (edge)

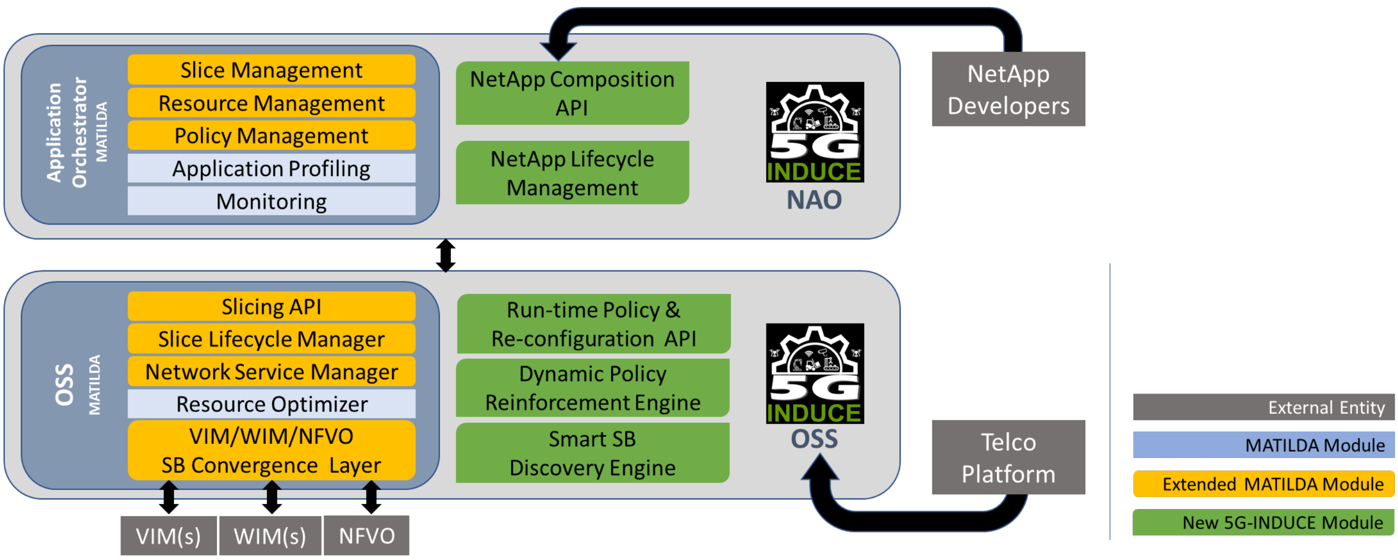

Fig. 2. Structure of the NAO and OSS and extensions/additions of components in support of the NetApp developers.

compute/storage resource re-allocation at run-time. This is a key requirement to permit the overlay 5G-INDUCE NAO to dynamically modify NetApp application graphs (e.g., in reaction to specific events originating from IoT devices).

- To redesign a flexible Southbound (SB) convergence layer to interface not only with the available 5G facilities, but also smart Industry 4.0 infrastructures and control blocks in an adaptive and transparent fashion. To further reinforce the 3GPP approach on 5G capability exposure, the 5G-INDUCE OSS will allow policy-driven hiding/anonymization of sensible data towards overlay vertical stakeholders.

The 5G-INDUCE OSS leverages a standardized 5G Network Orchestration Platform as shown at the bottom part in Fig. 2. This platform interacts with the afore-described SB convergence layer of the extended OSS through a number of standard direct interfaces towards well-known and widespread building blocks, such as the Open-Source MANO (adopting the standard ETSI NFV Os-Ma-nfvo reference point [7]), Kubernetes as Virtual Infrastructure Manager (VIM) for both the NFV Orchestrator and the edge computing support [9], and ONF-based Wide-area Infrastructure Manager (WIM).

B. The 5G-INDUCE NetApp Orchestrator (NAO)

The most radical contribution of the 5G-INDUCE project is the NAO, which interacts with the enhanced 5G-INDUCE OSS layer as exposed by the telco provider. In a nutshell, the scope of the NAO is to undertake the deployment, real-time management, and graceful un-deployment of Industry 4.0 NetApps, while inherently providing elasticity and compliance with certain high-level NetApp policies.

Modern vertical applications (essentially NetApps in the 5G domain) consist of a chain of cloud-native components that can be managed independently, as far as their scaling aspect is concerned. Each application component is bundled in an orchestration-friendly way, i.e., as a VM image, a container, or even a unikernel. An earlier version of the NAO [10] has been developed under the MATILDA project and was backwards compatible with all three industry-leading approaches mentioned above, offering also important telco-interplay capabilities such as Bi-directional interaction with a slice-enabled OSS, Resource-constrained slices, Application profiling and Policy enforcement.

One of the most crucial benefits of the MATILDA application orchestrator is its clear architectural separation from the telco. In other words, the MATILDA application orchestrator is completely unaware of the internal mechanisms that are used in order to satisfy the slice intent constraints. On the one hand, this is a desirable behavior, but on the other hand if the MATILDA orchestrator wants to utilize a pure layer-3 network service (e.g., layer-3 firewalling on a specific link), it has to “a priori” (i.e., prior to the Slice Intent request) request the telco provider for a discrete list of network services that are exposed along with non-functional details of their realization (e.g., what type of firewall, what is the maximum throughput that can be accommodated, etc.). Moreover, the MATILDA orchestrator does not provide fine-grained separation between application and network-level services. For instance, to restrict the traffic that goes through a database and a web server application component, the MATILDA orchestrator would duplicate the necessary firewall rules in each VM, which is far from ideal. Instead, an elegant solution would be to request the OSS to deploy a single firewall NF which could be dynamically associated with the application-level graph, forming an extended graph of both V/PNF and application components (i.e., a NetApp graph). However, this solution is not supported by the MATILDA platform, as neither the MATILDA application orchestrator can dynamically modify an application graph, nor the MATILDA OSS can modify its underlying network slice on-the-fly.

These are important gaps that will be addressed by 5G-INDUCE as depicted in Fig. 2. Specifically, the 5G-INDUCE NAO will leverage the powerful advancements of the 5G-INDUCE OSS, such as the refined SB convergence layer and the advanced functions (i.e., 5G Network Exposure, Network Slice Selection, and Policy Control functions) exposed by the “Network Service Manager”, to (i) realize a substantially enhanced slicing model, suitable for accommodating advanced slices for NetApps and (ii) virtualize and elegantly expose these resources to the overlay NetApps, forming a dramatically richer NetApp composition API, without disclosing sensitive information (e.g., UE or topology formation data) about the underlying network. Moreover, the 5G-INDUCE NAO will support run-time management of the NetApp lifecycle and graph formation, in order to dynamically (i) enable, (ii) scale, (iii) modify. or (iv) disable NetApp capabilities according to relevant events (e.g., the number and type of connected devices/things, their bandwidth on the radio link, their positioning, etc.). This feature is the key for advanced Industry 4.0 NetApps which pose excessive adaptation and dynamicity requirements at scale. Finally, all these advancements will be made possible while maintaining backward compatibility with the MATILDA application orchestrator with zero administrative cost.

V. RELATED NETAPP DEVELOPMENT AND ORCHESTRATION

At present, there are three distinct ways for an application to become cloud native: (i) to be hosted within VMs; (ii) unikernels, or (iii) containers. VMs have been widely used during the last ten years. They are extremely mature and well isolated, since hypervisors offer a strong sandboxing layer. Their main disadvantage is the high virtualization overhead, since each VM carries its own operating system. Unikernels [11] exhibit much lower virtualization overhead using highly-streamlined OS with only the necessary libraries required by the application. In return, they offer significant reduction in image sizes, improved efficiency, and acceptable security. The most dominant ones include OSV.io [12] and Mirage [13]. While Unikernels are extremely small in terms of footprint, their disadvantage is the lack of generalization during their creation. Finally, Containers provide isolation of an application from the OS and the physical infrastructure in which it resides using kernel-level virtualization (i.e., namespaces and cgroups). Such isolation provides means for the safe execution of applications; hence applications cannot contaminate the OS and furthermore they cannot contaminate resources used by other applications. The advantage of containers is their small footprint and their disadvantage is the weak (i.e. kernel-based) isolation scheme. The effort to maximize the utilization of cloud resources provided significant momentum to containers.

When it comes to orchestrating cloud native components there is no silver bullet for orchestrating containers and VMs simultaneously. It could be argued that Docker Swarm [13], and Kubernetes [9] are well-proven, production grade systems for deployment, scaling, and management of containerized applications, while OpenStack Heat [14] is a VM-based orchestrator. These orchestrators offer, out of the box, cluster management, static scaling, desired state reconciliation, load balancing and rolling updates. The MATILDA orchestrator [8] is able to deploy containerized apps on top of a modified telco OSS, orchestrating both VMs and containers, and therefore offering the advanced feature of automatic instrumentation and dynamic policy management. A set of intelligent orchestration mechanisms have been introduced and developed related to optimal deployment of cloud applications over programmable infrastructure, runtime policies enforcement, as well as mechanisms interacting with services provided by telecom operators through Northbound APIs. On a pre-commercial level, only a few cloud orchestrators have started moving towards bridging the gap between the cloud and telco world, by making relevant extensions in order to support extended networking functionalities (e.g., the open source Calico project). Such extensions, though, are still pre-mature and the ones that have managed to offer sufficient result are fiddly, demanding, and non-user friendly.

5G-INDUCE aims to incorporate the design of an interface capable of chaining and managing the VNFs in an abstract and user-friendly way, by providing a set of extensions. The goal is to equip the NetApp developers with a NetApp onboarding tool that allows automated porting, chaining, and reconfiguration of their solutions over the 5G platform stack. Moreover, the envisioned NetApp component repository (populated for example with already deployed component chains and specialized component features for security, processing, etc.), is a useful instrument for the design and extension of the NetApps over the platform. Through the support of the provided Northbound APIs, in combination with the provided advanced network functions, the telecom operators may claim a new service enabling role over their infrastructures providing the development environment to SMEs; at the same time, they can also offer the interface to service providers, in order to deploy their solutions in an application-oriented fashion over the infrastructure and with a view to the needs of the vertical stakeholders.

VI. Conclusions

This concept paper describes the vision, architecture and implementation aspects of the 5G-INDUCE project. The main goal is the development of an intelligent end-to-end orchestration platform for advanced 5G NetApps applicable in the broader Industry 4.0 sector. The core part of the innovative platform consists of: a) a NetApp orchestration layer that includes a user-friendly and automated NetApp porting interface with advanced NetApp component life-cycle management capabilities for dynamic modifications in the required resources, and b) a smart OSS layer. All this allows: i) application offloading for increased performance and security, ii) more generic application support through a policy-driven dynamic and fully automated association of applications with relevant network functions, and iii) discovery and exposure of smart infrastructure capabilities (through the NFVO).

The platform is intended to be deployed over 3 realistic 5G infrastructures comprising innovative industrial sector environments interconnected to the 5G core at operator premises.

Acknowledgment

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 101016941. The content of this article reflects the authors’ views and the Commission is not responsible for any use that may be made of the information it contains.

References

[1] 5G PPP Architecture Working Group, “5G empowering vertical industries”, Feb. 2016.

[2] 5GACIA white paper, “5G for Automation in Industry Primary use cases, functions and service requirements, Jul. 2019.

[3] I. F. Akyildiz, A. Kak, S. Nie, “6G and Beyond: The Future of Wireless Communications Systems,” *IEEE Access*, vol. 8, pp. 133995-134030, July 2020.

[4] *6G Waves Magaz*ine, no. 2, Autumn 2020.

[5] European Technology Platform NetWorld2020, “Smart Networks in the context of NGI,” Strategic Research and Innovation Agenda 2021-27, Sept. 2020.

[6] D. Bega, M. Gramaglia, R. Perez, M. Fiore, A. Banchs, X, Costa-Pérez, “AI-Based Autonomous Control, Management, and Orchestration in 5G: From Standards to Algorithms,” *IEEE Network*, vol. 34, no. 6, pp. 14-20, Nov./Dec. 2020.

[7] S. Kekki, et al., “MEC in 5G networks,” ETSI Whitepaper no. 28, June 2018.

[8] P. Gouvas, A. Zafeiropoulos, E. Fotopoulou, T. Xirofotos, R. Bruschi, F. Davoli, "Separation of concerns among application and network services orchestration in a 5G ecosystem", In *EuCNC Workshop on “From Cloud ready to Cloud Native Transformation: What It Means and Why It Matters”*, Ljubljana, Slovenia, Jun. 2018.

[9] Kubernetes: Production-Grade Container Orchestration, Available online: https://kubernetes.io/.

[10] R. Bruschi, et al., "5G-Ready Vertical Applications Orchestration", MATILDA Whitepaper, Available at: https://private.matilda-5g.eu/Publications/DownloadPublication/50

[11] A. Madhavapeddy et al. “Unikernels: Library operating systems for the cloud,” ACM SIGARCH Computer Architecture News, vol. 41, no. 1, pp. 461-472, Mar. 2013.

[12] S. Wu,, C. Mei, H. Jin, D. Wang, “Android unikernel: Gearing mobile code offloading towards edge computing,” Future Generation Computer Systems, vol. 86, pp. 694-703, Sep. 2018.

[13] N. Naik, “Building a virtual system of systems using docker swarm in multiple clouds,” In 2016 IEEE International Symposium on Systems Engineering (ISSE), Edinburgh, UK, Oct. 2016.

[14] R. Kumar, N. Gupta, S. Charu, K. Jain, S. K. Jangir, “Open source solution for cloud computing platform using OpenStack,” International Journal of Computer Science and Mobile Computing, vol. 3, no 5, pp. 89-98, 2014.