uc3m | Universidad Carlos III de Madrid

e-Archivo

This is a postprint version of the following published document:

A 5G Mobile Network Architecture to Support Vertical Industries

Albert Banchs, David M. Gutierrez-Estevez, Manuel Fuentes, Mauro Boldi, Silvia Provvedi

Abstract—The telecom industry is moving from a "horizontal" service delivery model, where services are defined independently from their consumers, towards a "vertical" delivery model, where the provided services are tailored to specific industry sectors and verticals. In order to enable this transition, an end-to-end comprehensive 5G architecture is needed, with capabilities to support the use cases of the different vertical industries. A key feature of this architecture is the implementation of network slicing over a single infrastructure to provision highly heterogeneous vertical services, as well as a network slicing management system capable of handling simultaneous slices. On top of the network slicing technology, functionality needs to be devised to deploy the slices required by the different vertical players and provide them with a suitable interface to manage their slice. In this paper, we design a 5G mobile network architecture to support vertical industries. The proposed architecture builds on ongoing standardization efforts at 3GPP and ETSI and incorporates additional modules to provide enhanced MANO and control functionality as well as artificial intelligence-based data analytics; on top of these modules, a service layer is provided to offer vertical players an easy-to-use interface to manage their services.

Introduction

5G technology provides a new mobile network architecture that, beyond a simple increase in network speed, represents a new paradigm in the end-to-end design of the network and the deployment of new services. The 5G system (5GS) is addressed globally, covering new frequency bands, new radio access technologies, new core capabilities as well as the overall optimization from the mobile radio to the backhaul and core network to run innovative services [1].

Enabled by 5G, the telecom industry is moving beyond voice and mobile broadband, the two core services that have powered the industry until today. Two main opportunities for telecom services involve the "massive" connection of machine-type objects and the "ultra-reliable and low latency" services, in addition to an enhanced mobile broadband service with even higher throughout than 4G. This new service offering entails a transition from a "horizontal" service delivery model, where services are defined independently from the consumers' needs, towards a "vertical" delivery model, where the provided services are tailored to specific industry sectors and verticals.

The transition to the "vertical" delivery model is expected to have a strong impact on the underlying business model of mobile networks. The use cases enabled by 5G address several industry segments that were not effectively covered by 4G and can potentially provide mobile network operators

Albert Banchs is with the University Carlos III of Madrid and with the IMDEA Networks institute; David M. Gutierrez-Estevez is with Samsung Electronics R&D UK; Manuel Fuentes is with the Universitat Politecnica de Valencia; Mauro Boldi is with Telecom Italia; Silvia Provvedi is with Ericsson Italy.


Fig. 1: 5G revenues provided by the different industry segments. Source: [2].

TABLE I: Key Performance Indicators (KPIs) resulting from different vertical use cases and applications.

Metric	Most demanding value	Some relevant verticals
Latency	$<= 5$ms	Media & entertainment
Reliability	99.999%	Manufacturing, health
Density	1M terminals/km2	Energy and utilities
Mobility	500 Km/h	Automotive
Slice deployment time	$<= 90$ min	All
Data rate per user	$>= 50$ Mb/s	Media & entertainment
Location accuracy	$<= 1$ m	Public safety

with new revenue streams, moving away from the traditional Business to Consumer (B2C) model, where revenues come from the end-users, to a Business to Business (B2B) revenue opportunity, with particular emphasis on the industry verticals. According to [2], the overall increase in revenues generated by the different vertical use cases enabled by 5G will be of around \$1,300 billion by 2026, which corresponds to a growth of around 36%, with the distribution across different industry sectors illustrated in Figure 1.

Vertical industries often generate different traffic types that impose very diverse and extreme requirements on the network, illustrated in Table I [3]. This means that the same network needs to efficiently support services with very diverse requirements. To address this, network operators need 5G networks to offer a flexible solution that allows tailoring the network deployment and configuration to serve a customized vertical service.

The goal of this paper is to propose a mobile network architecture that meets the requirements of the different vertical industries targeted by 5G, incorporating the capabilities needed to support their advanced use cases. To meet the highly diverse service requirements, we leverage the network slicing technology together with some novel functionalities.
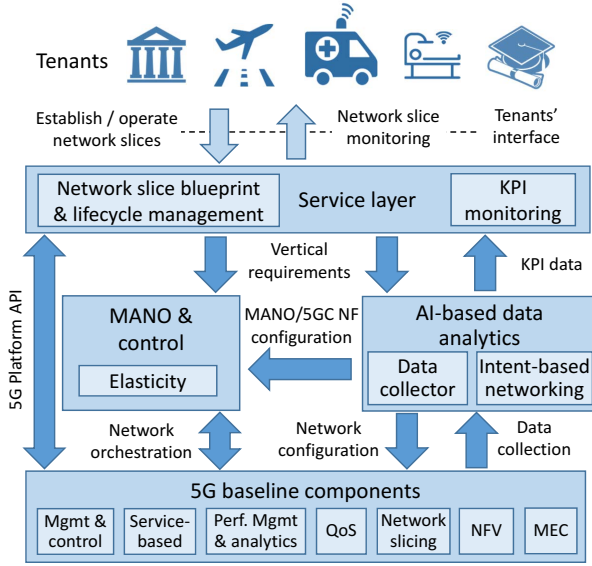
Fig. 2: Proposed 5G architecture for vertical industries.

Overall Architecture

Figure 2 depicts the overall architecture envisaged in this paper. The architecture lies on top of the 5G platform, which consists of baseline components compliant with the latest release of 3GPP. On top of the baseline components, we include some modules that provide functionality required for the different vertical use cases that is not available within the underlying platform. These modules employ standard interfaces to configure and manage the underlying functions.

The architecture envisions two different modules on top of the 5G baseline components:

- The 'MANO (Management and Orchestration) & Control' module is responsible for orchestrating and controlling the different network functions (NFs) and resources, based on the vertical requirements [4].
- The 'Artificial Intelligence (AI)' module takes care of the overall management, control and configuration of the network, relying on AI techniques, including the configuration of the NFs and the MANO algorithms.

The upper layer within the proposed architecture is the so-called service layer, which provides the vertical customers and other network slice tenants with a friendly interface to request and operate network services. This service layer maps the user requests coming from the verticals to primitives in the interfaces with the underlying modules.

In addition to the functions described above, broadcast is also an essential component of the proposed architecture. Indeed, broadcast is necessary for those vertical use cases that involve the delivery of a given content to many users simultaneously. This is the case, for instance, of applications such as ultra-HD immersive video or software updates for IoT devices. Since the use of point-to-point (PTP) unicast transmissions would be highly inefficient to this end, the support for broadcast and multicast technologies providing point-to-multipoint (PTM) transmissions is needed. To provide such functionality, the proposed architecture leverages the

findings of the 5GPPP 5G-Xcast project [5], implementing PTM by means of advanced algorithms in the 5G Core and the 5G Radio Access Network (RAN); specifically (*i*) two new NFs are added into the 5G Core, the 5G-Xcast User Plane Network Function (XUF) and the 5G-Xcast Control Plane Network Function (XCF), responsible for the delivery of multimedia content and for the control of the active broadcast sessions, and (*ii*) a 5G NR Broadcast solution is provided in the RAN relying on the concept of RAN Multicast Area (RMA) to broadcast a given service to a group of cells.

While there have been some individual efforts in the literature and standards around orchestration [6], data analytics [7] and AI-based network operation [8] to design similar components to those of our architecture, the key novelties of this paper are (*i*) the proposal of an architecture that brings all of these components together, and (*ii*) the design of the modules themselves, which contains novel elements; note that some of these novel elements are described in further detail in prior publications focusing on the individual components [4], [5], [9]–[11].

5G Baseline Components

The proposed architecture builds on the 5G functionality currently specified in the standards, adding new functions on top of this functionality. In the following, we briefly describe the standard 5G functionality that we take as baseline; the detailed definition of such baseline components can be found within the corresponding 3GPP specifications [6], [12].

Management and Control Architecture: A service-oriented management and control architecture is adopted for the interaction between service consumers and service producers. This involves operations on fault supervision service, performance management service, provisioning service and notification service, among others. A management service offers management capabilities, which are accessed by management service consumers via standardized service interfaces.

Service-based Architecture: The service-based architecture allows for the flexible addition and extension of functions including proprietary ones. It is centered on services that can register themselves and subscribe to others. This enables a more flexible development of new services, as it becomes possible to connect to other components without introducing specific new interfaces [13].

Performance Management and Analytics: The performance data of the mobile network and the various NFs are used for network monitoring, assessment, analysis, optimization and assurance. The analytics applications collecting these data can thus detect network problems and predict potential issues, allowing to take appropriate actions quickly or even in advance. To allow for such a quick reaction, the data needs to be collected in real time.

Flow-based QoS: 5G provides QoS with a much higher level of granularity than LTE. The QoS Flow Identifier (QFI) uniquely identifies QoS flows and consists of the following parameters: 5G QoS Identifier (5QI), which comprises the priority level, the delay budget, the error rate, the Allocation Retention Priority (ARP) and the Guaranteed and Max Flow Bitrate. The network treats packets based on their 5QI tags.

Network Slicing: The mobile network can be deployed to support different types of services: eMBB (enhanced Mobile Broadband), URLLC (Ultra-Reliable and Low Latency Communications) and mMTC (massive Machine Type Communications), depending on the customers' needs. This allows transforming the mobile network infrastructure from a single network to a network with logical partitions, each of them providing different service requirements with appropriate network isolation.

Network Virtualization: Most of the NFs are virtualized, becoming Virtual Network Functions (VNFs) and running as software components rather than using dedicated hardware components. In this context, mobile networks become operators' telco-cloud systems. Beyond the virtualization for the Core Network, the RAN NFs are also virtualized. In particular, with the functional split of the RAN between the central unit and the distributed unit, the central unit can be virtualized. The management system for the VNFs is aligned with ETSI Network Function Virtualization (NFV) and MANO.

Multi-access Edge Computing (MEC): The platform provides a virtualized computing environment to enable applications to run as VNFs in network nodes. This offers storage and computational resources at the edge, accessible by application/ service providers and third parties, which can experience reduced latencies and higher speeds.

Enhanced Management, Orchestration and Control

In order to accommodate the demands of the vertical use cases, the architecture comprises an enhanced network slice MANO framework. This function builds on the MANO solution developed by the 5G-MoNArch project [9]. It extends the underlying standard 5G platform with (i) algorithms that go beyond the standards specifications, (ii) the concept of elasticity, and (iii) an agile control plane, leveraging the provided standardized interfaces.

In order to orchestrate a network slice to meet the requirements of vertical customers, the orchestration algorithms devised in [4] are employed. These algorithms select the most appropriate location for the different VNFs instantiated for a network slice, including the ones corresponding to application layer functions that can be located within the network relying on the MEC capabilities of the infrastructure. The algorithms take as input from the underlying infrastructure the resource availability in the different nodes of the network, accounting both for computing and communications resources, and interact with the AI-based data analytics module (presented later) to take the optimal decision.

Based on the output of the orchestration algorithm, the NFV Orchestrator (NFVO) instantiates the network slice as follows: (i) it issues the corresponding requests to the Software-Defined Networking (SDN) controller to instantiate connections between the different network nodes, (ii) it requests the Virtualized Infrastructure Manager (VIM) to reserve the virtual resources at the different network nodes, (iii) it commands the Virtual Network Function Manager (VNFM) to instantiate the VNFs, and (iv) it configures the VNFs and PNFs (Physical Network Functions).

The enhanced MANO framework proposed is further combined with the resource elasticity concept of 5G-MoNArch [4]. Elasticity refers to the ability of a system to gracefully adapt to changes in the amount of available resources while minimizing the disruption perceived by the users. An elastic system thus adapts to load changes in an automatic manner to closely match the available resources while making a use as efficient as possible of the underlying resources. This implies that, when reducing the amount of available resources, the operation of the service should not be significantly degraded. Such a feature is highly desirable for a number of the envisaged use cases based on eMBB where service guarantees are not very stringent. For such use cases, elasticity can help to minimize the cost of the required network deployment.

Elasticity is addressed by means of (i) orchestration-driven elasticity, involving the flexible placement of VNFs, and (ii) slice-aware elasticity, involving cross-slice resource provisioning mechanisms. Orchestration-driven elasticity is provided by taking VNF placement decisions that maximize the efficiency of the system, mitigating available the impact of outages in the available resources. Slice-aware elasticity is achieved by performing end-to-end cross-slice optimization, jointly orchestrating and controlling multiple network slices, maximizing the overall efficient while providing isolation across slices to guarantee their respective Service Level Agreements (SLAs).

In addition to the management functionality described above, the architecture also comprises a highly agile control plane that assists with the overall optimization of the slicing performance, enabled by the service-based architecture and the availability of data analytics. Such control plane, implemented by the set of 5GC NFs, differs from current solutions in that it operates at a shorter timescale as compared to the relatively long timescales of the control architecture of the standard. It assists the overall management of slices by reconfiguring themselves or performing short-term decisions following elasticity principles when the SLAs of the different slices are at risk.

AI-based Data Analytics

In our design, AI plays a very prominent role, featuring an AI-based Data Analytics framework that allows for autonomous and more efficient control, management and orchestration of the network. The framework is particularly well aligned with the concepts defined by 3GPP SA2 for the Network Data Analytics Function (NWDAF) of the 5G Core (5GC) [7], and it incorporates the basics of the functionality defined by 3GPP SA5 for the Management Data Analytics Function (MDAF) [6] and the Intent Driven Management Service [14] of the management plane.

Precisely because the above standardization efforts are still ongoing, there is no current full-blown data analytics-assisted architecture ready for deployment. The proposed architecture utilizes the most mature components and fills out the missing gaps with novel modules that intend to influence the standards discussions in control, management and orchestration planes. Our AI-based data analytics framework is depicted in Figure 3 and builds on the following main guidelines:
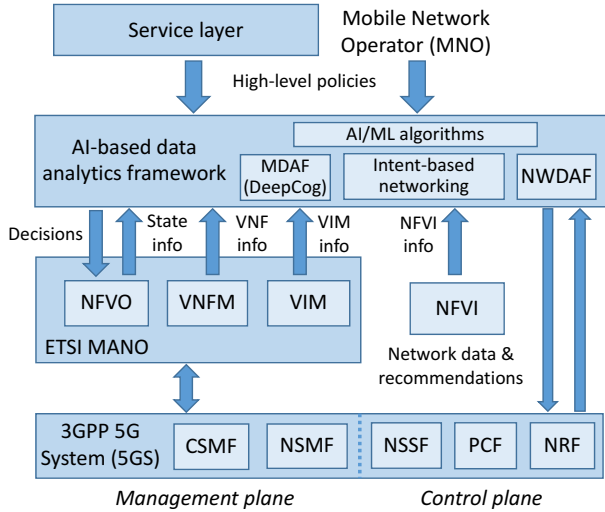
Fig. 3: Proposed AI-based data analytics framework.

- We rely on control plane data analytics specified within Rel-16 (i.e., NWDAF) for the optimization of the 5GC NFs, proposing novel methods that leverage the available analytics to optimize performance for the following 5GC NFs: (i) Networks Slice Selection Function (NSSF), (ii) the Policy Control Function (PCF), and (iii) Network Repository Function (NRF).

- We utilize an orchestrator and management system that relies also on analytics for making its decisions, enhancing the Communication Service Management Function (CSMF) and the Network Slice Management Function (NSMF) and proposing an AI-based capacity forecasting tool named DeepCog to assist orchestration decisions.

- On top of the above functions, the AI-based Data Analytics layer follows the modularized system architecture of ETSI ENI [8], comprising the following modules which are responsible for gathering, representing and understanding the needed information and leveraging such information to take the respective decisions while meeting the end-to-end goals: (i) Policy Management, (ii) Context Awareness, (iii) Situational Awareness, (iv) Cognition Management, and (v) Knowledge Management.

In the following we describe in further detail the data analytics framework proposed for the control and the management/orchestration planes, respectively.

Data analytics for the control plane: In the control plane, analytics allow NFs to optimize their behavior at run-time, typically at a much faster speed than what network management and orchestration systems allow. Control plane decisions are usually constrained because of not having visibility on the underlying infrastructure, leaving all scaling decisions out of its scope. However, slice-level load balancing and congestion control as well as QoS management can benefit from the NWDAF analytics defined in [7] on slice load, service experience and Quality of Experience (QoE). To realize this, our design includes the following enhancements to 5GC NFs:

- NSSF: this NF, in charge of selecting the set of Network Slice instances serving a UE, optimizes real-time

slice load decisions by being constantly informed by the NWDAF on both load status and service experience statistics and predictions. We propose a novel slice selection and load control functionality for NSSF that, based on slice load and service experience analytics, decides at the session establishment phase which slice optimally serves each of the new UEs arriving in the network.

- PCF: this NF supports a unified policy framework to govern network behavior, including the determination of QoS parameters relevant to each application and the service provided by the network slices. We propose an extension of PCF functionality that, informed by NWDAF analytics on UE and application service experience, adapts service QoS parameters across all UEs on a slice in such a way that the slice SLA is satisfied.

- NRF: this modules is in charge of keeping an NF profile of all NFs belonging to a slice, including their instantaneous load; whenever a new instantiation is needed, it provides a set of candidate NF instances for a certain NF type. With NWDAF analytics, we enhance NRF by adding an NF instance pre-selection step so that not only instantaneous NF load is taken into account, but also statistics and predictions. This way, load balancing is embedded in the selection process of the new NF instance among the candidate set.

Data analytics for management and orchestration planes: In the management and orchestration planes, the data used as input by the AI-based analytics framework is provided by the NFV Infrastructure (NFVI) and the MANO system, as shown in Figure 3. The NFVI information provides the knowledge on the computational resources' capabilities (such as the type of CPU and memory, accelerators, etc.) along with their availability (i.e., the status and utilization level). Based on this information and running AI-based algorithms, the framework influences and optimizes the placement decisions made by the VIM, while ensuring that the resulting resource allocation satisfies the respective slice SLAs. AI-based data analytics is also employed to optimize resource provisioning and admission control of network slices [10]. All these algorithms are integrated within the slice management functions, namely CSMF and NSMF.

3GPP's efforts to leverage data analytics for management and orchestration are being channelized through the MDAF, but the specification of this function is still in a very preliminary stage [6]. Our approach proposes a novel optimization to MDAF called DeepCog [11], which leverages demand analytics to provide predictive slice management decisions targeting resource utilization efficiency. DeepCog builds on deep learning to forecast the capacity needed to accommodate future traffic demands within individual network slices based on the collected data traces of each slice. In this way, it anticipates future loads in the various slices and uses this to (i) assist the VIM to take decisions that avoid resource outages, and (ii) timely trigger up/down scaling or in/out scaling of the associated resources. The DeepCog algorithm has been carefully designed to accommodate the offered load at all times thereby avoiding resource underprovisioning, as

Fig. 4: Capacity forecasting with DeepCog. Top: actual demand and forecasted capacity for a slice providing video service. Bottom: level of overprovisioning (blue) and underprovisioning (red) over time. DeepCog almost never incurs in underprovisioning. Source: [11].

otherwise the SLAs with tenants would not be met (typically implying high costs in terms of fees to tenants for violating their SLAs). This is depicted in Figure 4, which shows that DeepCog almost never incurs in underprovisioning. Furthermore, as shown in [11], by tuning the algorithm to provide the desired level of underprovisioning, DeepCog can meet diverse vertical requirements in terms of reliability. This shows that, by combining data analytics and AI-based techniques, our approach is able to effectively optimize resource orchestration while meeting vertical service requirements.

SERVICE LAYER

In order to meet the needs of the vertical customers, the architecture includes a service layer that provides them with a suitable interface. The design of this service layer is aligned with the incipient efforts on Exposure Governance Management Function (EGMF) [6] at 3GPP, yet its scope goes much beyond that of the standard.

The service layer focuses on tenants such as vertical industries that require a customized network slice from the infrastructure to meet their specific requirements of the provided service. To request, establish and maintain a network slice, a tenant requires an interface from the infrastructure that includes the following functionalities:

- *Instantiate a network slice*: When a tenant needs a network slice, it has to issue a request to the infrastructure indicating information such as: (i) the area that needs to be covered by the network slice, (ii) the load that needs to be supported, (iii) the requirements that need to be satisfied in terms of rates, latencies and reliability, (vi) the mobile terminals and devices that belong to the slice, etc.
- *Orchestrate application-layer virtualized functions*: In case the tenant needs to place some of its application-layer functions within the network infrastructure employing MEC technology, it has to provide the virtualized function to the infrastructure indicating the constraints associated to the placement of such function; this could be the case, for instance, of an AR/VR server with very low latency requirements.

- *Monitor the network slice*: After its network slice has been instantiated, the tenant needs to be able to monitor the service provided by the network slice to confirm that its requirements are being satisfied and take corrective actions when needed. For instance, if the slice's load increases, a larger slice may be requested.
- *Operate the network slice*: The tenant needs to be able to perform some operations on a running network slice, such as adding new users to the slice, increasing its coverage, changing the requirements or the load, re-orchestrating application-layer virtualized functions, etc.

It is important to note that many of the tenants employing a network slice are players outside the traditional mobile network ecosystem that do not have the skills and expertise to manage mobile network services. Therefore, in order to allow such players to enter the network slicing market without imposing an unacceptable burden on them, it is of utmost importance to provide an intuitive and easy-to-manage interface to establish and operate a slice. For the implementation of the service layer, two alternative interfaces are provided to the tenants: (i) an API with different functions to automatically instantiate and operate slices, and (ii) a web interface that allows to instantiate and operate slices manually. The first one is useful for those verticals that have management software that is running their service and wish to automate the network slicing operations. The second one is more useful to those tenants that run their service manually.

In order to cope with the complexity of the system while providing a simple interface to manage and operate slices, we adopt an intent-based approach where the policies coming from vertical customers simply indicate their 'business intent', declaring high-level service policies rather than specifying detailed networking configuration. By leveraging AI techniques, the network understands such policies and aligns to them, using context and analytics to learn and adapt to changing needs and conditions. The powerful abstraction level offered by such an intent-based interface allows to specify service policies in a way that is close to the customer's natural language.

MEETING VERTICAL REQUIREMENTS

In the following, we explain how the proposed architecture components contribute to meet the vertical requirements identified in the KPIs presented in Table I, providing some experimental results for the individual components.

Table II summarizes the relationship between the various architecture components and the KPIs. There are four relevant KPIs for the architecture components presented in this paper: latency, reliability, slice deployment and data rate. Network slicing is the technology underlying the architecture and is essential towards the achievement of all of them. Results provided in [9] show that our network slicing technology can effectively deploy isolated slices for different services.

AI-based data analytics is also central to achieving many of the KPIs, as it allows to take intelligent decisions when placing VNFs and allocating resources across the network. This ultimately results in network slices that are efficiently

TABLE II: Contribution of the proposed components to the relevant KPIs.

KPI	Network Function	Contribution to the KPI
Latency	Network slicing	Network slicing allows to customize the delay provided by each slice
	MANO	By placing the VNFs impacting latency close to the end-users, low latencies can be provided
	AI-based data analytics	By leveraging data on resource usage, intelligent orchestration decisions can be taken to maximize overall efficiency while meeting latency requirements
Reliability	Network slicing	Network slicing guarantees isolation between slices, thus allowing to provide a reliable service to selected slices
	AI-based data analytics	By means of AI-based algorithms, we can determine the necessary resources that need be allocated to a slice to guarantee its requirements
Slice deployment time	Network slicing	Network slicing allows to deploy different slices for different services
	MANO	MANO takes care of the instantiation the NFs required for a given slice, automating the process of deploying a slice and reducing the times involved
	Service layer	The service layer provides an easy-to-use interface to tenants for the deployment of new slices
Data rate per user	Network slicing	Network slicing ensures isolation between slices, protecting the data rates provided to each slice
	AI-based data analytics	AI-based data analytics can be applied to many user plane functions, substantially improving the system efficiency and the delivered data rates

orchestrated to provide the desired services in terms of latency, reliability and data rates. It has been shown in [15] that, by dynamically allocating resources in network slicing, a much more efficient usage of the (radio and compute) resources can be achieved, with gains of up to 10 in the radio access and up to 2 in the central cloud. Furthermore, [11] shows that AI-based algorithms can effectively realise such gains.

The orchestration capabilities provided by MANO are essential to deploy slices tailored to the desired services. Experimental results from [4] confirm that latencies as low as 5 ms can be achieved by placing critical VNFs at the edge of the network. In terms of slice deployment, results from two different testbeds [4] corroborate that MANO technologies can deliver slice deployment times in the order of minutes.

Note that the above results cover only a subset of the KPIs identified in Table I. As a matter of fact, this paper has focused on some specific components that contribute to some essential KPIs but not all of them. The remaining KPIs are satisfied by the baseline 5G components. For instance, aspects such as mobility and density are addressed by physical layer components which are beyond the scope of this paper.

CONCLUSIONS

5G enables a new business model for mobile networks where network customers are no longer end-users but vertical industries, which demand very diverse services from the mobile network with extreme requirements in many cases. To meet the requirements of such vertical players, a new network paradigm is required around network slicing, where independent slices are instantiated for different tenants and customized to meet their specific requirements. Standard bodies such as 3GPP and ETSI have already defined the baseline 5G platform for network slicing. However, the design of a comprehensive solution to provide 5G mobile network services to vertical industries requires of additional components to deploy and manage the slices providing the vertical services. In this paper, we have presented a 5G mobile network architecture for this purpose. The architecture comprises enhanced MANO functions to orchestrate network slices, an agile control plane

to accommodate changing demands on the fly, and AI-based data analytics modules to support orchestration and control decisions. On top of this, we have designed a service layer adapted to the needs of vertical players. The proposed architecture combines state-of-the-art technology from standards with innovative modules designed to effectively support vertical industries' services.

REFERENCES

[1] R. E. Hattachi, J. Erfanian, and B. Daly, "5G White Paper," NGMN Alliance Tech. Rep., Feb. 2015.

[2] Ericsson White Paper, "The 5G business potential: Industry digitalization and the untapped opportunities for operators," 2017.

[3] 5GPPP Programme Management Report, "5G PPP phase II KPIs (Annex)," 2018.

[4] D. Gutierrez-Estevez *et al.*, "Architecture and mechanisms for resource elasticity provisioning," 5G-MoNArch project Deliverable D4.1, June 2018.

[5] T. Tran *et al.*, "Mobile Core Network," 5G-Xcast project Deliverable D4.1, May 2018.

[6] 3GPP TS 28.533 v16.0.0, "Management and orchestration of networks and network slicing; Management and orchestration architecture (Release 16)," June 2019.

[7] 3GPP TS 23.288 v16.1.0, "Architecture enhancements for 5G System (5GS) to support network data analytics services (Release 16)," June 2019.

[8] ETSI White Paper 22, "Improved operator experience through Experiential Network Intelligence (ENI)," Oct. 2017.

[9] G. Garcia-Aviles, M. Gramaglia, P. Serrano, and A. Banchs, "POSENS: A Practical Open Source Solution for End-to-End Network Slicing," *IEEE Wireless Commun.*, vol. 25, no. 5, Oct. 2018, pp. 30-37.

[10] D. Bega, M. Gramaglia, A. Banchs, V. Sciancalepore, and X. Costa-Perez, "A Machine Learning approach to 5G Infrastructure Market optimization," *IEEE Trans. Mobile Comput.*, published online on Feb. 1, 2019.

[11] D. Bega, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Perez, "DeepCog: Cognitive Network Management in Sliced 5G Networks with Deep Learning," in *Proc. of IEEE INFOCOM*, Apr. 2019.

[12] 3GPP TS 23.501 v16.2.0, "System Architecture for the 5G System (Release 16)," Sept. 2019.

[13] A. Minokuchi and S. Isobe, "5G Core Network Standardization Trends," *NTT DOCOMO Technical Journal*, vol. 19, no. 3, Jan. 2018, pp. 5-12.

[14] 3GPP TR 28.812 v0.6.0, "Study on scenarios for Intent driven management services for mobile networks (Release 16)," June 2019.

[15] C. Marquez, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Perez, "How Should I Slice My Network?: A Multi-Service Empirical Evaluation of Resource Sharing Efficiency?" in *Proc. of ACM MOBICOM*, Oct. 2018.

Albert Banchs [SM] (banchs@it.uc3m.es) received his M.Sc. and Ph.D. degrees from the UPC-BarcelonaTech in 1997 and 2002. He has a double affiliation as Professor at the University Carlos III of Madrid and Deputy Director of the IMDEA Networks institute. He was Technical Manager of the 5G-MoNArch European project and is Deputy Technical Manager of 5G-TOURS.

David M. Gutierrez-Estevez [M] (d.estevez@samsung.com) received the B.E. degree from the University of Granada in 2009 and the M.S. and Ph.D. degrees from the Georgia Institute of Technology in 2011 and 2014. He is currently a Principal Research/Standards Engineer at Samsung Electronics R&D UK, where he is contributing to 3GPP standardization as SA2 delegate.

Manuel Fuentes (mafuemue@iteam.upv.es) received his M.Sc. and Ph.D. degree from the Universitat Politecnica de Valencia (UPV) in 2012 and 2017. He worked at Samsung Electronics R&D UK and is currently at UPV working on the 5G-TOURS project.

Mauro Boldi (mafuemue@iteam.upv.es) graduated at Politecnico di Torino in 1998. He is at the TIM research centre, where he is Coordinator of the 5G-EVE European project, Deputy Coordinator of 5G-TOURS, and Vice-Chairman of the EEPS group in the ETSI EE Technical Committee.

Silvia Provvedi (silvia.provvedi@ericsson.com) graduated at the University of Genoa in 1989. She is at Ericsson Italy, where she holds the position of Program Manager & Senior project manager. She is Coordinator of the 5G-TOURS project.