Towards Machine-Learning-Based 5G and Beyond Intelligent Networks: The MARSAL Project Vision

John S. Vardakas*, Kostas Ramantas*, Eftychia Datsika*, Miquel Payaró†, Sofie Pollin‡,
Evgenii Vinogradov‡, Manos Varvarigos§, Panagiotis Kokkinos §
Roberto González-Sánchez¶, Juan Jose Vegas Olmos‖, Ioannis Chochliouros**, Philippe Chanclou††,
Pierangela Samarati‡‡, Adam Flizikowski x, Md Arifur Rahman x, and Christos Verikoukis†
*Iquadrat Informatica S. L., Barcelona, Spain, Email: {jvardakas,kramantas, edatsika}@iquadrat.com
†Telecommunications Technological Center of Catalonia (CTTC/CERCA), Castelldefels, Spain,
Email: {miquel.payaro, cveri}@cttc.es
‡Department of Electrical Engineering, KU Leuven, Belgium, Email: {sofie.pollin, evgenii.vinogradov}@esat.kuleuven.be
§School of Electrical & Computer Engineering, NTUA, Athens, Greece, Email:vmanos@mail.ntua.gr, kokkinop@central.ntua.gr
¶NEC Laboratories Europe, Madrid, Spain, Email: roberto.gonzalez@neclab.eu
‖NVidia Mellanox, Roskilde, Denmark, Email: juanj@mellanox.com
**OTE Research, Omilos OTE, Athens, Greece, Email: ichochliouros@oteresearch.gr
††Fixed Access Networks Group, Orange Labs, Lannion, France, Email:philippe.chanclou@orange.com
‡‡Universita degli Studi di Milano, Italy, Email:pierangela.samarati@unimi.it
xIS-Wireless, Piaseczno, Mazowieckie, Poland, Email:{a.flizikowski, a.rahman}@is-wireless.com

***Abstract*—5G mobile networks will be soon available to handle all types of applications and to provide service to massive numbers of users. In this complex and dynamic network ecosystem, end-to-end performance analysis and optimization will be key features to effectively manage the diverse requirements imposed by multiple vertical industries over the same shared infrastructure. To enable such a vision, the MARSAL project targets the development and evaluation of a complete framework for the management and orchestration of network resources in 5G and beyond, by utilizing a converged optical-wireless network infrastructure in the access and fronthaul/midhaul segments. At the network design domain, MARSAL targets the development of novel cell-free based solutions by exploiting the application of the distributed cell-free concept and of the serial fronthaul approach, while contributing innovative functionalities to the O-RAN project. In parallel, in the fronthaul/midhaul segments MARSAL aims to radically increase the flexibility of optical access architectures via different levels of fixed-mobile convergence. At the network and service management domain, the design philosophy of MARSAL is to exploit novel ML-based algorithms of both edge and midhaul DCs, by incorporating the Virtual Elastic DataCenters/Infrastructures paradigm. Finally, at the network security domain, MARSAL aims to introduce mechanisms that provide privacy and security to application workload and data, targeting to allow applications and users to maintain control over their data, while AI and and Blockchain technologies will be developed in order to guarantee a secured multi-tenant slicing environment.**

I. Introduction

The number of people living in megacities, that is cities with population greater than 10 million, has increased from 69.5 million in 1975 to a staggering 472.8 million in 2015 [1]. The percentage of network traffic originating and terminating in a city is increasing [2], through the use of popular on-line services and smart city applications. It is expected that Smart Mega Cities will become the main source of data, characterized by massive data growth and processing requirements. Today these requirements are served by a variety of optical and wireless networking and edge/fog/cloud computing technologies and infrastructures deployed in cities and belonging to different providers, while smaller business owners (e.g. stadium operators) are also deploying their own infrastructures [3].

In such an environment, 5G networks are set to address the demands of a fully connected and mobile society, enabling a wide variety of applications over the same infrastructure, while carrying 45% of the total mobile traffic and serving up to 65% of the world's population [4]. These numbers are expected to increase due to the urbanization of global population and the increase of the size and volume of mega-cities that is cities with population greater than 10 million [1].

5G changes the landscape of mobile networks in a profound way, with an evolved architecture supporting unprecedented capacity, spectral efficiency, and increased flexibility. Moreover, 5G adopts Edge computing as a key paradigm, evolving from centralized architectures, e.g., based on C-RAN([5]), towards multiple tiers of Edge nodes and a virtualized RAN (vRAN). Open RAN initiatives such as O-RAN have a key role in this evolution, complementing the 3GPP 5G standards with a foundation of vRAN network elements [6]. However, these technologies have been in large developed in isolation between them, making difficult to fully exploit their capabilities in an integrated, end-to-end and secure manner. Algorithms do not only run in the cloud, and optical and wireless links cannot be abstracted in the same way. When going to cell-free networking concepts, more nodes and links will be interconnected, serving local and global secure applications, and thus it is essential to rethink the architecture and algorithms running elastically at the scale of a city or building level.

In general, application traffic flows from and towards end-users and end-devices, served by multiple levels of storage and computing entities from the edge to the cloud, while utilizing a diverse set of wireless and optical technologies in the fronthaul, midhaul and backhaul network segments. These infrastructure resources belong to different administrative domains, operate in parallel in the same network areas and are usually shared between competing flows, computations and data in static and/or statically multiplexed manner. Thus, it is clear that targeted activities need to take place to fully exploit key developments, towards a disaggregated infrastructure model, where a technological infrastructure blocks can be transparently and flexibly replaced by others, while offering similar networking and/or computing offerings and control and monitoring capabilities. Specifically, key advances are required both in the network design and network/service orchestration levels:

• The network infrastructure should be able to support multiple distributed edge nodes and a huge number of access points, which are coordinated and orchestrated by entities in a low-cost and near-zero latency manner;

• A unified and hierarchical infrastructure is essential in order to provide an intelligent management of communication, computation and storage resources, which can be further enhanced by incorporating efficient Machine-Learning (ML) algorithms;

• The support of multiple tenants should be followed by the application of mechanisms that are able to guarantee data and information security and integrity especially in multi-tenant environments, which would play a vital role in enabling various use-cases and industry verticals targeted in 5G and Beyond (B5G) systems.

To face the aforementioned challenges, the EU research project MARSAL proposes a new paradigm of elastic virtual infrastructures that integrate in a transparent manner a variety of novel radio access, networking, management and security technologies, which will be developed under the MARSAL framework in order to deliver end-to-end transfer, processing and storage services in an efficient and secured way. This will be achieved by developing innovations that are based on three pillars to enable a new generation of ultra-dense, cost-efficient, flexible and secure networks: the *network design pillar*, the *virtual elastic infrastructure pillar*, and the *network security pillar*. For the network design pillar, MARSAL pushes cell-free networking towards the distributed processing cell-free concept, and enables wireless mmWave solutions, which will be implemented and integrated with existing vRAN elements, while being in-line with the O-RAN Alliance. MARSAL's second pillar is built based on the Elastic Edge Computing notion, targeting to optimize the functionality of the Mobile Edge Computing (MEC) and the network slicing management systems via a hierarchy of analytic and decision engines. Finally, under its third pillar, MARSAL will develop novel ML-based mechanisms that guarantee privacy and security in multi-tenancy environments, targeting both end users and tenants. In the rest of the paper, we present the MARSAL concept, together with the research advancements that will be implemented throughout the duration of the project.

II. The MARSAL Concept and Architecture

MARSAL aims to provide an evolved architecture towards B5G, offering unprecedented degrees of flexibility and closed-loop autonomy at all tiers of the infrastructure, and significantly improved spectral efficiency via cell-free networking.

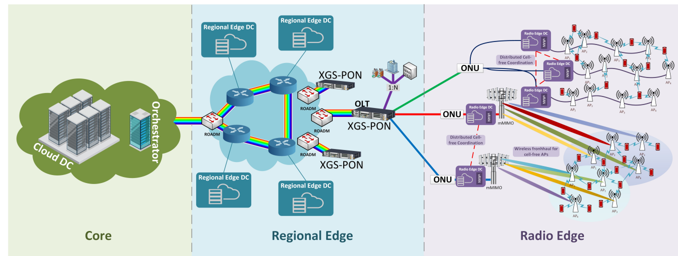

Fig. 1. Main concept of MARSAL project describing a complete 5GB system set-up from access to the core networks

The overall architecture and the structure of the envisioned B5G MARSAL is depicted in the generic schematic of Fig. 1, and includes all the main infrastructure elements that are deployed within MARSAL project. MARSAL adopts an evolved 3GPP NG-RAN which is extended with emerging cell-free technologies for network densification. Moreover, MARSAL contributes innovations at the optical transport domain and significantly evolves the MEC system towards fully elastic Edge Computing. MARSAL will deploy a distributed Edge infrastructure with Data Centres (DCs) structured in 2 tiers, featuring Regional Edge and Radio Edge nodes. Radio Edge DCs will host the Network Functions of the (virtualized) RAN, which fully aligned with the O-RAN specifications .

At the network design level, the emphasis is on innovations at the RAN and fronthaul domains, that will unlock the potential of cell-free networking in future 5GB networks. MARSAL targets the development of novel cell-free networking mechanisms that will allow the significant scaling up of AP deployment in a cost-effective manner, by exploiting the distributed processing cell-free concept. MARSAL proposes to disaggregate the traditional cell-free Central Processing Unit (CPU) in Distributed Units (DUs) and a Central Unit (CU) in line with the 3GPP NG-RAN architecture. Regarding the wireless fronthaul links, MARSAL aims to design and prototype an innovative mmWave Hybrid MIMO solution. MARSAL's cell-free innovations will be implemented and integrated with existing vRAN elements for the first time, and will be contributed back to the O-RAN project.

At the network service and management level, MARSAL will offer a joint optimization framework for the efficient, self-driven orchestration and management of communication, storage and computational resources in an integrated way. To this end, MARSAL aims to optimize the functionality of the MEC system and the Network Slicing Management system via a hierarchy of analytic and decision engines. Moreover, MARSAL envisages a new approach based on the Elastic Edge Computing notion, aiming to provide proposing an evolved MEC system which allows application functions to by disaggregated and dynamically migrated across the edge and/or from the RAN to the core cloud and vice-versa.

At the network security level, MARSAL will introduce mechanisms that guarantee privacy and security in multi-tenancy environments, targeting both end users and tenants. Thus, OTT application providers and end users can maintain control over their data when relying on shared infrastructures, without sacrificing on functionality. MARSAL's developed

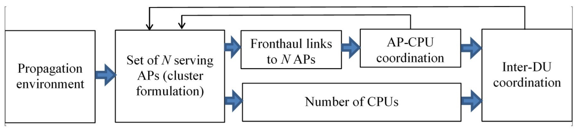

Fig. 2. MARSAL's proposed Cell-Free optimization flow, involving inter-AP and inter-DU coordination

security mechanisms and policies will be developed explicitly for the MARSAL network architecture, providing trustworthiness and resiliency against security failures or breaches, thus guaranteeing data and computation integrity, and generating privacy preserving data representations.

III. MARSAL's Network Design Pillar

To enable fast and flexible network densification, novel cell-free networking and mmWave dense fronthaul solutions will be proposed in MARSAL. Relying on the already existing distributed cell-free KUL testbed, channel data will be collected, distributed processing algorithms and architectures will be designed, which will be integrated in a B5G architecture.

A. Data-driven Cell-Free in 5G and Beyond Architectures

MARSAL envisions cell-free networking as a key component of B5G RANs, that will offer unprecedented spectral efficiency and performance which is not constrained by inter-cell interference. Cell-free networking is a user-centric paradigm, where UEs can share the same resources and are concurrently served by multiple APs. The Spectral Efficiency achieved is a function of the level of coordination between APs (e.g., it is almost doubled when Channel State Information (CSI) is shared with the CPU for performing joint combining/precoding [8]). However, a practical cell-free RAN will require the significant scaling-up of AP deployments, which is not possible with current approaches where APs are served by a single CPU. To overcome this limitation, MARSAL will disaggregate the traditional CPU in multiple DUs, in line with 3GPP's 5G architecture. MARSAL's innovations will focus on distributed processing, with clusters of APs and DUs coordinating via fronthaul links as shown in Fig. 2, and addressing two main pillars: firstly, MARSAL will address cluster formation, dynamically allocating a sub-set of APs to each UE. Secondly, MARSAL will consider clusters of APs served by multiple DUs, where inter-DU coordination is required for signal decoding. MARSAL will explore inter-DU coordination requirements and their effect in Spectral Efficiency, and propose dynamic adaptability of the coordination levels jointly addressing AP-DU and DU-DU coordination. MARSAL will leverage a cell-free testbed offered by KU Leuven testbed to deploy various cell-free topologies e.g. tree-based Fronthaul [9] and Serial Fronthaul [10].

B. mmWave Hybrid Fronthaul for Cell-Free networks

Massive MIMO fronthauling has been very recently proposed for interconnecting APs and CPUs in the cell-free networking paradigm [11]. MARSAL will propose a new Hybrid MIMO fronthauling approach, specifically targeting cell-free networks, with advanced beamforming and beam-sharing capabilities. Thus, a new AP topology adaptation in cell-free networks, and advanced scenarios can be supported, with APs reassigned to different DUs on demand, as shown in Fig. 3. Moreover, point-to-multipoint connectivity will be

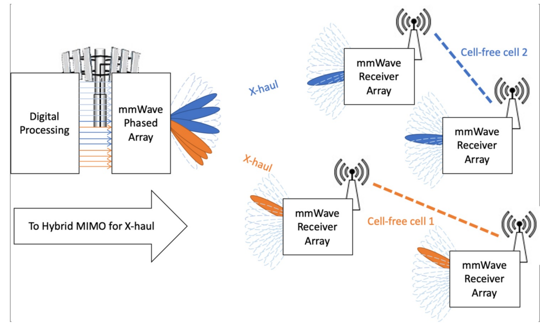

Fig. 3. MARSAL's Hybrid MIMO beam-steering and beam-sharing concept

supported via beam-sharing from multiple APs. To this end, MARSAL will design and build prototypes of mmWave Hybrid MIMO nodes, leveraging PT's RFIC mmWave beamforming transceiver and phased-array antenna module. MARSAL will propose many ground-breaking innovations, related to optimal AP deployment strategies, new formulations of dynamic AP clustering and adaptive coordination problems, a new mmWave interface similar to 3GPP's Xn directly interconnecting DUs, and new functional splits in the context of the cell-free fronthauling compatible with 3GPP's Options 1-8.

C. MARSAL's Cell-Free vRAN

The MARSAL project will be aligned with the ORAN Alliance architecture, which represents an evolution of Cloud RAN (C-RAN), further disaggregating the Base Band Unit (BBU) and complementing the 3GPP 5G standards with a foundation of virtualized RAN (vRAN) network elements and packet-based interfaces. Specifically, ORAN disaggregates the BBU in a DU with the real-time functions, and a CU with the non-real time functions. The latter is further disaggregated into the CU-User Plane (CU-UP) and the CU-Control Plane (CU-CP) as shown in Fig. 4. MARSAL's CU User Plane function (i.e., CU-UP) and DU will be deployed at MARSAL's Radio Edge, and the CU-CP Near-RT RIC at the Regional Edge. The cell-free high-PHY functions will be integrated with the vDU and MAC scheduler via a 5G FAPI-like interface, adding cell-free support in ORAN's architecture for the first time.

The MARSAL project will contribute with many cell-free related innovations to the ORAN architecture. Specifically, distributed cell-free networking support will be implemented as part of the O-DU module, appropriately modifying the PHY, MAC, and RLC sublayers. The role of orchestration in the midhaul (NFVO, MEC) will also be considered. Moreover, cell-free Radio Resource Management (RRM) will be implemented as part of the Near-RT RIC, which will be derived from ACC's dRAX platform. The dRAX plug-in framework will be leveraged to implement the data-driven RRM functionality for the cell-free architecture in line with the O-RAN specifications for training and deploying ML inference within the Near-RT RIC [12] . Extensions of existing interfaces and control plane protocols will be applied at various reference points (e.g. fronthaul, E2/O1 interfaces, etc.) to incorporate support for cell-free networking, after appropriate requirement analysis.

IV. MARSAL's Virtual Elastic Infrastructure Pillar

MARSAL envisages a new paradigm of Edge infrastructures based on the notion of Elastic Edge Computing, aiming

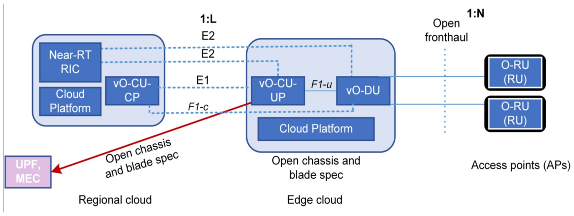

Fig. 4. MARSAL's vRAN components and interfaces, aligned with O-RAN

to overcome the isolation and underutilization of resources deployed at Edge nodes, and offer zero perceived latency to smart connectivity applications.

A. Software-Defined control plane towards Fixed Mobile Convergence

MARSAL will propose and implement a novel hierarchical control plane solution, federating the SDN controllers of the fixed and mobile segments of the network under a common orchestration subsystem. While the O-RAN Orchestration and Automation layer includes a Non-RT SDN Control & Policy function, it is deployed at the Core tier and therefore operates at a relatively coarse time-scale of multiple seconds, while only targeting the RAN segment. For the control plane of the mobile segment, MARSAL proposes the disaggregation of the Non-RT into Near-RT SDN Control function that will be hosted by the Near-RT RIC at the Regional Edge nodes. Thus, near real-time reaction to workload variations will be supported, at sub-second timescales. Moreover, MARSAL proposes the deployment of Software-Defined Transport Network Controllers (i.e., SDTNs) at the Regional Edge to control the fixed segment. Both domains will be federated under MARSAL's Core Tier NFVO, based on ETSI OSM, which will provide Network Slicing As a Service (NSaaS) functionality as per 3GPP TR 28.801 specifications. Thus, end-to-end slicing with centralized orchestration is supported, while still allowing innovative closed-loop (or ML-driven) control of each individual domain.

MARSAL also proposes a novel Fixed Mobile Convergence (FMC) solution, to facilitate integrated connectivity of mobile and fixed (i.e., FTTH) services. MARSAL's solution involves two transmission approaches seamlessly integrated at the Regional Edge node, including a standard Point-to-Point (PtP) connection with or without WDM and a very disruptive point-to-multipoint (PtMP) approach based on XGS-PON modules, as detailed in Fig. 5. An energy efficiency scenario will also be considered, with traffic aggregation on a limited set of wavelengths which will allow shutting-down individual SFP+ transceivers. Moreover, predictive dynamic slicing approaches will also be explored, leveraging an ML-driven control loop to trigger slice reconfiguration proactively, based on convolutional neural networks that predict traffic fluctuations.

B. MARSAL's Elastic Edge Cloud Infrastructure

The 5G Service-Based Architecture has adopted Edge Computing as a key paradigm, and has defined the necessary interfaces (e.g., Mp2 reference point) and enabling technologies that allow Edge deployments to be fully integrated with the 5G Core. Moreover, TR 23.758 specifications define a new Application Architecture for verticals via the EDGEAPP initiative ([13]), allowing OTT applications to be deployed at

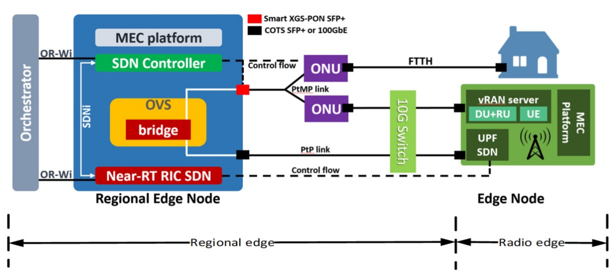

Fig. 5. MARSAL's FMC architecture and distributed control plane

Local Area Data Networks at the network Edge. MARSAL will leverage an existing MEC platform (e.g. StarlingX) that participates in ETSI's plug-tests, but it will be fully integrated with the 5G NFVI , supporting coordinated resource allocation for MEC applications and 5G Network functions. This is accomplished by coordinating two different management and orchestration sub-systems (i.e., the NFVO and MEC Orchestrator) which interact through the Mm1 reference point, defined as part of the ETSI MEC 017 specifications.

While previous approaches adopted a common VM-based technology stack for MEC and NFV, MARSAL approach will be based on Cloud-Native technologies (i.e., Docker Containers, Kubernetes Virtual Infrastructure Managers (VIMs)) which are widely regarded as the future of vertical application development [14]. While support for Kubernetes VIMs is gradually emerging in MEC platforms (e.g., in StarlingX), there is currently a gap in supporting disaggregated Cloud-Native apps. To fill this gap, MARSAL proposes extensions to the MEO to support the disaggregation of application functions, that will be defined as collections of helm charts, both horizontally (i.e., across Edge sites) and vertically (i.e., from the cell site towards the core cloud). MARSAL will consider the deployment of the aforementioned functions either at the "bare metal"of the MEC hosts' NFVIs or within VNFs, as proposed in the NFV-IFA 029, thus compatible with the NSaaS sub-system. Moreover, MARSAL will extend the Mobile Edge platform at the host level, to allow MEC apps to be accessed by any UE, irrespective of physical location. Dynamic Virtual Network Embedding algorithms will be explored, to determine the optimal disaggregation of application functions at any Edge DC, considering Compute, Networking, and Storage constraints, thus achiving increased resource utilization.

C. Self-Driven Virtual Elastic Infrastructures via multi-objective optimization

Beyond-5G networks will be characterized by Self-Driven infrastructures, with pervasive ML and closed-loop autonomy at all layers. However, while the concept of a"global brain"that controls the infrastructure top-down is attractive, it can easily become the bottleneck due to the massive number of data generated by complex 5G infrastructures. MARSAL will propose a novel, distributed approach (Fig. 6) that involves Analytic Engines at all tiers of the Edge infrastructure, and Decision Engines at the two Core-Tier orchestration subsystems. Analytic Engines, the first pillar of automation, analyse and federate measurements to achieve Context Awareness, and Decision Engines, as the second pillar, Plan and React to Context changes, delegating data-driven local control decisions to the lower tiers of the hierarchy.

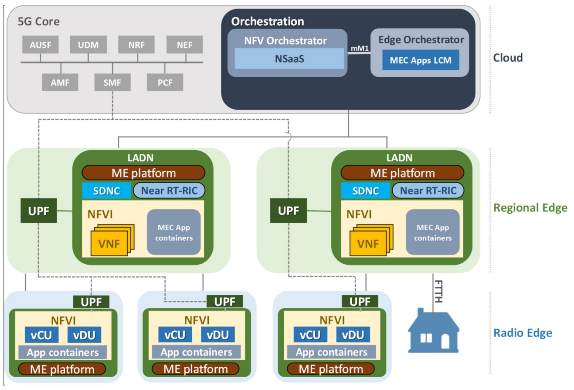

Fig. 6. MARSAL 3-tier architecture

For the first pillar of automation, MARSAL will design and implement an innovative, decentralized approach to achieve global Context Awareness, using for the first time Representation Learning and Embedding Propagation (EP) algorithms . State-of-the-art context representation methods for 1D and 2D data-sets are not appropriate for MARSAL's diverse network and application data, that can be best represented with a graph-like abstraction. To this end, network slices and MEC applications will be represented as the nodes of a knowledge graph, along with their defining variables and parameters (e.g., SLA requirements, latency budgets, cost considerations, energy efficiency goals). MARSAL proposes to apply EP mechanisms to build the node representations (or embeddings) of the knowledge graph, iterating over the data and minimizing the differences among neighbouring embeddings in the graph.

For the second pillar of automation, the resulting embeddings that represent the current state (or Context) of the MARSAL infrastructure in a highly compressed form (i.e. encoded as multidimensional normalized arrays) can be transmitted to the Core Tier Decision Engines. The embeddings are fed to downstream ML algorithms implemented by the Decision Engines that jointly orchestrate Network Slices, Network Services and MEC applications continuously and automatically evaluating current context under required policy. Due to the high number of (potentially conflicting) parameters and policy requirements involved, MARSAL will consider multi-objective optimization techniques that achieve different trade-offs between optimality and complexity.

V. MARSAL's Network Security Pillar

Multi-tenancy is foreseen as a pillar of future 5G and beyond networks, and will help recoup the big investments required by telecom operators. In multi-tenant infrastructures, there is a distinction among Mobile Network Operators (MNOs), Infrastructure providers, and Mobile Virtual Network Operators (MVNOs). MVNOs and OTT application providers generally lack spectrum licenses, and hence rely on MNOs to lease Network Slices. MARSAL will propose mechanisms which allow the collaboration of all these entities without assuming trust among participants, or requiring participants to share their business or operational data, as well as ML-driven slice security mechanisms to defend from a rogue tenant.

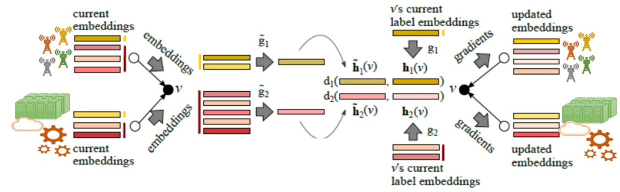

Fig. 7. MARSAL's EP Operation Schema.

A. ML and Blockchain technologies for trust-less multi-tenant slicing

MARSAL aims to deliver a decentralized, blockchain-based platform that supports network slicing transactions via smart contracts, targeting multi-tenant infrastructures for the first time. In this platform, the MNO, MVNOs, and OTT vertical application owners form a decentralized autonomous organization, which can dynamically negotiate network slice contracts, flexibly integrating large and small players without the need for a centralized entity. Smart Contracts facilitate direct contracts among entities that can be dynamically renegotiated based on real-time supply and demand. MARSAL's smart contract platform will be implemented with a private, permissioned blockchain solution where tenants will co-own the validator nodes' network after approval and authentication.

Going a step further, MARSAL will propose and implement new privacy-preserving context representations, which will allow MARSAL's data-driven NSaaS subsystem to operate without exposing tenants' business and operational data. Context awareness requires the exchange of local embeddings (via EP mechanisms) that represent the nodes of a knowledge graph, risking information leakage. These embeddings have inherent anonymization properties, as they represent nodes as compressed, high-dimensional arrays, while the application of EP algorithms iteratively minimizes the differences among neighbouring embeddings, further decreasing the risk of re-identification of the original node. MARSAL will propose innovative techniques to guarantee that embeddings can't be reversed (Fig. 7), and can be shared among competing partners and the NSaaS sub-system, without any risk of disclosing confidential information.

B. Policy-driven data protection and integrity assurance

Modern hyper-converged NFVIs that are gradually being adopted by the telco ecosystem tend to consolidate storage, compute and networking resources under a unified VIM. Storage resources are distributed in clusters of storage pools, and can be deployed at nodes that are part of the Core and Edge NFVIs, forming a Distributed Cloud Storage (DCS) solution [7]. MARSAL will propose a new DCS solution, which offers strong guarantees that the reconstruction of any portion of an original storage resource (e.g., a data file) is impossible, unless an entity has access to all chunks along with the encryption key. This will be accomplished with the application of an All-Or-Nothing-Transform (AONT) to the storage resources and new fragmentation and allocation strategies that also consider the minimum number of nodes required to reconstruct a resource in a multi-tenant infrastructure. While AONT has been previously studied in the context of P2P networks ([15]), it is proposed in MARSAL as an innovative solution in a multi-tenant infrastructure environment with potentially untrusted nodes, ensuring that the unavailability of a single data chunk prevents the reconstruction of any part of the original plaintext.

MARSAL will also propose and implement an innovative NFS Gateway, controlled by the OTT application provider, to serve as the foundation of trust. The DCS gateway will be the intermediary between the (trusted) OTT application, and (untrusted) DCS infrastructure. The NFS gateway will be extended to implement a novel data pipeline for the controlled sharing of data among different parties. The gateway will support for first time a policy language, extending SoTA solutions such as Open Digital Rights Language or JSON-LD, effectively enforceable in the data protection context of different stakeholders that will permit the specification of sharing and processing restrictions over data. Finally, the NFS gateway will implement a novel probabilistic scheme that protects the integrity of computations based on the randomized injection of pre-computed and replicated computational tasks. Thus, a unified solution for data obfuscation and integrity assurance will be implemented for the first time, which varies the probability of randomized injections based on the degree of protection or performance required.

C. Hardware Accelerated, ML-based data plane security and malicious traffic detection

To improve the performance of current signature-based solutions for dealing with zero-day or evolving attacks, MARSAL provide hardware accelerated solutions for a decentralized Threat Detection Engine and a centralized Threat Analysis Engine. ML-based threat detection, that has demonstrated an improved ability to extract complex non-linear relationships in attack data , will be leveraged for the design of MARSAL's Threat Detection Engine (TDE). Moreover, MARSAL will leverage the capabilities of a new generation of MELX programmable SDN Switches, that support the new P4 programming language for stateful processing, header inspection, and packet metadata extraction, operating at wire speeds . The programmable pipeline of the aforementioned SDN Switches will be exploited in combination with the P4 stateful registers to perform cyberattack detection and mitigation via embedded ML algorithms (e.g., for feature extraction) that will be implemented at the Network Switch Operating System. This will allow MARSAL's data-plane to behave as a distributed barrier against threats, securing the entire transport infrastructure and intercepting cyber-attacks at a very early stage. Detected cyber-attacks can be isolated by the data-plane at the level of individual traffic flows, going beyond traditional slice-centric approaches and towards micro-segmentation.

Moreover, MARSAL will propose a centralized Threat Analysis Engine (TAE), that will operate as an ML Fusion Centre, collecting and correlating metadata and features extracted from the P4 pipeline of the decentralized TDE. This will allow complex attacks such as Advanced Persistent Threats that simultaneously target multiple network nodes, to be detected. Furthermore, it will provide system-wide consistency and correlation for events occurring within all the involved P4 pipelines. MARSAL will exploit the flows' destination IPs, and specifically their sequences, since this information is both unencrypted and readily available as part of the SDN Switches' telemetry framework. Thus, the TAE will associate observed flows (with unknown status) with malicious ones based on the sequence of IPs accessed. MARSAL's solution will involve feeding a Deep Neural Network with sequences of IPs to build a vector representation of network flows; intuitively, flows with a small distance from malicious flows should also be flagged as malicious.

VI. Conclusion

This article presents the vision of the MARSAL project proposing a complete framework for the management and orchestration of network resources in 5GB systems. MARSAL's three main pillars are presented, targeting to build innovative advancements related to the network design, management, and security of the network infrastructure. These innovations will establish the main MARSAL vision, which is to become a new paradigm of elastic virtual infrastructures that integrates in a transparent manner a variety of novel radio access, networking, management and security technologies, in order to deliver end-to-end transfer, processing and storage services in an efficient and secured way

Acknowledgment

This work is supported by the EU research project MARSAL (101017171).

References

[1] Statista Megacities - Statistics & Facts, 2020, [Online], Available: https://www.statista.com/topics/4841/megacities/.

[2] CISCO VNI. Cisco Visual Networking Index: Forecast and Trends, 2017–2022 White Paper. [Online], Available: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.htm

[3] G.Kalfas, et al., "Next Generation Fiber-Wireless Fronthaul for 5G mmWave Networks", *IEEE Commun. Mag.*, vol. 57 , no. 3 , pp. 138-144, Mar. 2019.

[4] Ericsson, Ericsson Mobility Report, Nov. 2019. [Online]. Available: https://www.ericsson.com/en/mobilityreport/reports/november-2019

[5] K.Ramantas, et al., "A C-RAN based 5G platform with a fully virtualized, SDN controlled optical/wireless Fronthaul", *in proc. IEEE ICTON*, Bucharest, Romania, 1-5 Jul. 2018.

[6] N. Kazemifard, and V. Shah-Mansouri, "Minimum delay function placement and resource allocation for Open RAN (O-RAN) 5G networks", *Computer Networks*, vol. 188, 107809, Apr. 2021.

[7] R.Muñoz, et al., "The CTTC 5G End-to-End Experimental Platform: Integrating Heterogeneous Wireless/Optical Networks, Distributed Cloud, and IoT Devices", *IEEE Veh. Technol. Mag.*, vol. 11, no. 1, Mar. 2016.

[8] E. Björnson and L. Sanguinetti, "Making Cell-Free Massive MIMO Competitive With MMSE Processing and Centralized Implementation", *IEEE Trans. Wireless Commun.*, vol. 19, no. 1, pp. 77-90, Jan. 2020.

[9] L. Van der Perre, et al., "RadioWeaves for efficient connectivity: analysis and impact of constraints in actual deployments", *in proc. 53rd Asilomar*, Pacific Grove, CA, USA, 3-6 Nov. 2019.

[10] P. Frenger, et al., "Improved antenna arrangement for distributed massive MIMO", Patent application WO2018103897, 2017.

[11] S. Blandino, et al., "Multi-User Hybrid MIMO at 60 GHz Using 16-Antenna Transmitters", *IEEE Trans. Circuits and Systems I: Regular Papers*, vol. 66, no. 2, pp. 848-858, Feb. 2019.

[12] [ORAN-WG2.AIML] O-RAN Working Group 2 AI/ML workflow description and requirements, v01.00 10/2019, [Online], Available: https://www.o-ran.org/s/ORAN-WG2AIMLv0100.pdf

[13] 3GPP SA6 initiatives to enable new vertical applications, [Online], Available: https://www.3gpp.org/news-events/conferences/2082-sa6_verticals

[14] 5G-PPP Software Network Working Group, "Cloud Native and 5G Verticals' services", [Online], Available: https://5g-ppp.eu/wp-content/uploads/2020/02/5G-PPP-SN-WG-5G-and-Cloud-Native.pdf February 2020

[15] E. Bacis, et al., "Securing Resources in Decentralized Cloud Storage", *IEEE Trans. Inf. Forens. and Sec.*, vol. 15, pp. 286-298, 2020.