

Towards 6G: Architectural Innovations and Challenges in the ORIGAMI Framework

Livia Elena Chatzieftheriou, Marco Gramaglia, Andres Garcia-Saavedra, Steffen Gebert, Gines Garcia-Aviles, Stefan Geissler, Marco Fiore, Paul Patras, Andra Lutu, Dimitris Tsolkas, Md Arifur Rahman

Abstract—As research in mobile networks is already transitioning from 5G to 6G, we identify a set of fundamental barriers in the current 5G architecture that limit efficient and global operations. We propose innovative architectural solutions that can remove such barriers and lay the foundation for 6G systems. Specifically, we introduce three novel architectural components: the Global Service-Based Architecture (GSBA), the Compute Continuum Layer (CCL), and the Zero-Trust Layer (ZTL). These components collectively aim to enhance network efficiency, security, and scalability, addressing future mobile networks’ dynamic and demanding needs. Furthermore, we discuss the integration of Network Intelligence (NI) that exploits the aforementioned architectural innovations to ensure global operations and services. Ultimately, our proposed vision entails a more adaptive, secure, and intelligent network architecture, setting the groundwork for the next generation of mobile networks.

Index Terms—6G, mobile network architectures, network intelligence, global mobile services, compute continuum

I. INTRODUCTION

The introduction of the 5G technology in the overall mobile network ecosystem has undoubtedly brought advantages in terms of end-user experience and mobile network operator Total Cost of Ownership (TCO) with respect to the legacy generations [1]. However, the expected revolution that the 5G ecosystem should have brought has oftentimes stayed halfway to the envisioned target. Items such as the deep integration of verticals in the network operation or the flexible operation of components fell short in achieving the entirety of their objectives. Therefore, to pave the way towards 6G, some specific actions should be taken:

- Perform a diagnosis of the current shortcomings of the 5G system, carrying out a thorough diagnosis of the current obstacles that are slowing down the achievement of a global service and the efficient usage of the available infrastructure. In §II, we will provide an initial analysis, identifying 8 major barriers in different domains of the network. This aspect should be deeply investigated to ensure the success of 6G, considering both design and implementation perspectives and ensuring the alignment of the network with the future needs of the 6G ecosystem.

L.E. Chatzieftheriou and M. Fiore are with the IMDEA Networks Institute. M. Gramaglia is with University Carlos III de Madrid (UC3M). A. Garcia-Saavedra is with NEC Laboratories Europe GmbH. S. Gebert is with EMnify GmbH. Gines Garcia-Aviles is with i2CAT Foundation. S. Geissler is with University of Würzburg. P. Patras is with Net AI Tech Ltd. A. Lutu is with Telefonica Research. D. Tsolkas is with National and Kapodistrian University of Athens as well as with Fogus Innovations and services P.C. Md Arifur Rahman is with IS-Wireless, Poland.

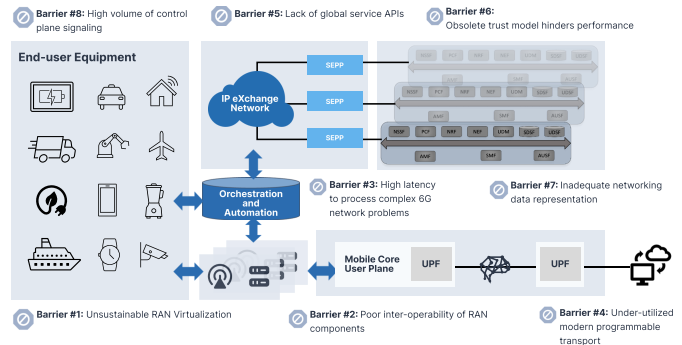


Fig. 1. Limitations of the current 5G Architecture. The architecture innovations ORIGAMI puts forward aim to remove these eight barriers towards 6G.

- Design novel architectural models to overcome the identified barriers. In §III, we introduce architectural components that complement and extend the current 5G architecture: the Global SBA (GSBA), the Compute Continuum Layer (CCL), and the Zero-Trust Layer (ZTL).
- Introduction of new control loops. By introducing the aforementioned components, new feedback loops can be introduced. The 6G Network design shall exploit them by introducing new Network Intelligence (NI) solutions and new Services (as the ones described in §IV) that will contribute to achieve the 6G envisioned KPIs.

The approach discussed in this paper is on the basis of the SNS JU Project ORIGAMI¹, which will employ this methodology to ensure the successful design of the 6G networks.

II. THE BARRIERS LANDSCAPE

The current technological solutions and architectural models considered for 5G and its future evolutions have brought undeniable advantages to the telecom arena, helping the transition towards comprehensive adoption of the Network Function Virtualization (NFV) paradigm, and contributing to more effective management of cloudified resources. The fundamental shift from a hardware-based approach to a software-prominent one, based on the adoption of best practices from enterprise software development and large open-source communities, is paving the road to more flexible and increasingly efficient networks, as well as new business models, mitigating vendor lock-in, and introducing software tiering. Despite the giant

¹<https://sns-origami.eu>

TABLE I
BARRIERS THAT ORIGAMI AIMS TO REMOVE.

Barrier	In what it consists
#1	Unsustainable RAN virtualization
#2	Poor inter-operability of RAN components
#3	High latency and unreliable Network Intelligence to process complex 6G network problems.
#4	Under-utilized modern programmable transport
#5	Lack of global service APIs
#6	Obsolete trust model hinders performance
#7	Inadequate networking data representation
#8	High volume of control plane signaling

leap forward in the past decade, there are still several barriers that, once overcome, will pave the way toward 6G.

In particular, while the NFV concept was proposed more than a decade ago, only very recently the network architectures developed by main Standardization Development Organizations (SDOs) started to natively embrace software as the main design driver, e.g., O-RAN (RAN domain) and 5G SBA (Core). Moreover, many important functions are very difficult to accommodate with the current virtualization technologies that stem from the cloud domain and with the present closed approach that characterizes MNO operations.

The architectural innovations we propose put forward the aim to remove eight barriers toward 6G in the state of the art that are inherent to the current 5G architecture and technology. We summarize them in Fig. 1 and Table I. Removing them requires fundamental changes in terms of the mobile network architecture, which will go well beyond the current state of the art and are imperative to pave the road towards the development of 6G systems. We now discuss eight prominent barriers in the current 5G architecture that, if not tackled, risk hampering the capabilities of future 6G systems.

1) *Efficiently Managing 6G Network Infrastructure*: The capability to efficiently control and expose the computing resources available in different infrastructure providers, across the virtualized radio access, transport, and core domains. **Barriers #1, #2, #3 and #4** refer to the need to address complexity, interoperability, operating and control latency, and under-utilization of resources to enable the efficient deployment and operation of 6G networks. One important challenge involves the lack of a unified framework for managing the increasingly complex and heterogeneous landscape of 6G infrastructure (**Barrier #1**) [2], which is further exacerbated by the poor interoperability between open RAN components (**Barrier #2**) [3]. Solving these interoperability issues is crucial for enabling real-time control and coordination across various network elements. Fast and efficient Control Plane (CP) functions are essential for maximizing the benefits of high-performance User Plane (UP) functions, and overcoming the currently high latency introduced there (**Barrier #3**). To address this issue, ORIGAMI will explore asynchronous functions, lightweight computation approaches, and cutting-edge computing technologies such as quantum computing. Overcoming **Barrier #4** is crucial for leveraging the full potential of programmable user planes and unlocking innovative 6G VNFs and services.

2) *Open Tenancy and MNO Collaboration*: The possibility of realizing actual global network operations that are natively

open to tenants as well as external MNOs, obstructed by **Barriers #5, #6, #7, and #8**. With these, we show that the current cellular architecture is fundamentally at odds with empowering smaller-scale providers and global service (**Barrier #5**), thus limiting competition and innovation. We question the status quo, where service providers expand their coverage globally only via pre-established agreements between two entities (**Barrier #6**), based on which they cooperate in authenticating and billing users. This relies on the assumption that users rarely cross provider boundaries, which in ORIGAMI we argue is obsolete [4]. Internet players such as Mobile Network Aggregators (MNAs) are already pushing for a new operational model, where devices enjoy global seamless mobility, currently relying only on international roaming, but suffering in terms of quality of service and meeting stringent performance requirements. In this context, the lack of a common network data representation makes troubleshooting service issues a slow and mostly human-driven process (**Barrier #7**). At the same time, the coordination of operations of one end-user across multiple domains puts a high load in terms of control traffic volume (**Barrier #8**) on an infrastructure that was initially designed only to support global mobility as a niche service (i.e., the rare international travel of users and their devices. [5]. Ultimately, the existing scenario significantly constrains the practicality of achieving localized breakouts, thus impeding the comprehensive utilization of the flexibility and performance enhancements envisioned in 5G.

III. ORIGAMI APPROACH: LAYERS AND LOOPS

To overcome the barriers above, we propose three architectural innovations that enable three global control loops that will support different Network Intelligence (NI) functionalities and global services. The three envisioned architectural enablers will evolve the architecture beyond the current 5G SBA approach, effectively tearing down these barriers. This re-design, while maintaining the essence of the 5G architectural work, transitions into 6G by integrating the proposed architectural evolutions, as depicted in Fig. 2 and discussed next.

A. Compute Continuum Layer (CCL)

The CCL is an innovative architectural component for 6G systems, envisioned to optimize network operations according to the characteristics of the underlying infrastructure. It will enable efficient resource sharing and will unlock the potential of heterogeneous computing infrastructures across multiple domains. It will support a variety of computing resources, including GPUs, TPUs, FPGAs, ASICs, NPUs, smartNICs, and even novel quantum computing paradigms, accelerating VNFs and complementing conventional CPUs.

The CCL shall pave the way for compute-aware network operations while maintaining the abstractions of a pure virtualization layer. This involves rethinking Network Functions (NFs) at both levels, enhancing efficiency, and reducing resource usage by matching NFs to the CCL. Consequently, the CCL manages available resources, allowing the highest utilization factor in the edge-to-cloud continuum.

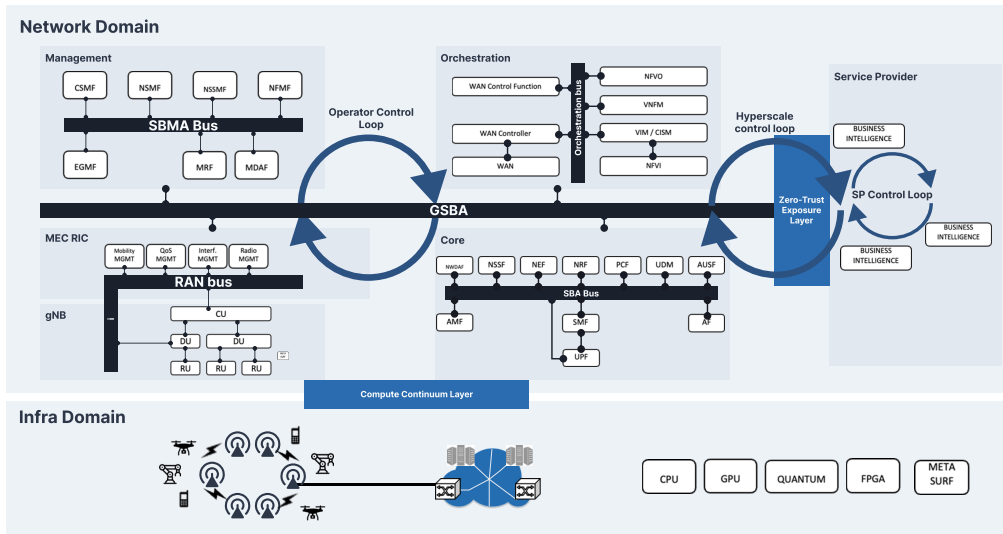


Fig. 2. The ORIGAMI architecture innovations that enable next-generation global services and Network Intelligence (NI) functionalities.




Fig. 3. The ORIGAMI Compute Continuum Layer (CCL)

The CCL will allow a centralized and coordinated abstraction between NFs and specific hardware components within the underlying hardware substrate. This approach simplifies resource management and allocation, promoting efficient utilization and scalability. Furthermore, the CCL accommodates new sensing functionalities within the network, catering to tenants interested in metadata rather than data transport.

The computational resources assigned to specific NFs might impact both the overall performance and their operational behavior, especially on virtualized RANs, where time constraints are tight. Therefore, the CCL will enable the regulation of VNFs by auditing and supporting decisions taken during their operation, ensuring their correct operation and performance in real-time under a large variety of scenarios.

We depict the envisioned CCL in Fig. 3. It should provide a feature-rich API that fulfills a twofold objective: *i*) provide an abstraction of the heterogeneous and disaggregated (in terms of technology and execution environments, respectively) computing infrastructure that can be used by the 6G Network Userland, including all software components being executed in the network; and *ii*) expose 6G Userland-specific tools to exploit the underlying infrastructure components.

In this way, the underlying infrastructure can be efficiently

pooled and effectively used just when required by 6G Userland apps. This aspect is of particular importance to avoid vendor lock-in while guaranteeing performance in challenging environments such as next-generation 6G deployments.

In summary, ORIGAMI’s CCL aims to be a groundbreaking architectural component for 6G systems that tailors network operations to the underlying infrastructure, supports diverse computing resources, and preserves virtualization layer abstractions. By enabling compute-aware operations and accommodating sensing functionalities, the CCL paves the way for the next generation of communication networks.

B. Zero-Trust Exposure Layer (ZTL)

The ZTL enables the matching between the internal operation of the service provider’s business logic and the network operator’s continuous optimization. This cooperative control loop can unleash the full potential of the network deployment by *i*) having a more direct influence on the operation internals and *ii*) unlocking beyond data-transfer functionality such as remote sensing or digital twinning directly to the providers. Thanks to this view, the ZTL can provide verticals with network hyper scalers in a similar way computing hyper scalers (such as Amazon AWS, Microsoft Azure, and Google Cloud Services) are doing already for other purposes.

The ZTL enables both vertical and horizontal exposure:

1) *Vertical Exposure – Network Application Analytics Fusion*: The data analytics framework proposed by 3GPP in TS 23.288 introduced a giant leap forward in how analytics are produced and managed within the network. The requirement for network automation has influenced the design of the 3GPP system standardized in R15. Before this, the data and analytics generation involved communication between network elements and their managers through proprietary interfaces. However, with the subsequent consolidation in R16 and R17, the system architecture was re-designed to natively support the collection




Fig. 4. Business agreements to enable international roaming

of analytics towards automation loops. The Network Data Analytics Function (NWDAF) is the cornerstone of this system, which gathers data from other NFs, computes analytics, and shares them with other consumer functions in the network. A fundamental part of this exchange is the one with the management system, where the Management Data Analytics Function (MDAF) acts as a bridge towards other selected functions, such as the radio-related one. In this framework, 3GPP defined analytics that are usually backed by relevant use cases involving interactions with other network functions in the 5G Core Domain and with the Management system: for R18 we have 14 categories ranging from UE-related analytics to NF performance. While the analytics system provides a framework to effectively perform self-optimization tasks in the 5G Core and Management domains, including federated learning, all these metrics are prominently related to network management. In a scenario like the one described in §II (Barrier #6), where the service provider is capable of directly customizing the network behavior in a zero-trust fashion, also the NWDAF-provided analytics shall reflect this view.

This problem is related to the traditional problem of QoS to QoE mapping, which has been studied a lot in the past, also in the context of 5G systems [6]. In general, the analytics framework proposed by 3GPP deals with QoS metrics that are related to the different Network Functions in the core. Still, Service Providers need to optimize their own metrics that are indeed related to the QoS perceived by users but may substantially differ because of, e.g., business-related aspects that are unknown to the network operator.

Thus, the ZTL enables specific APIs that can be leveraged to provide customization of the analytics provided by the NWDAF, to also take into account feedback coming to the overarching service provider, as exemplified below.

2) *Horizontal Exposure – Enabling the Global MNO Model.*: Support for *things* operating globally has become critical for IoT verticals, from connected cars to logistics and wearables, and explains the commercial success of IoT platforms. All IoT device manufacturers need a global connectivity solution, and this requirement is bound to become urgent with the rapid acceleration of IoT deployments. IoT verticals rely on providers who can ensure reliable data connectivity across the globe, such as M2M platforms that depend on the cellular ecosystem. International mobile roaming is thus an essential service for IoT verticals. Depending on the use case (e.g., automotive, logistic tracking, smart meters), roaming may be required occasionally or permanently. Different IoT verticals come with potentially different requirements: while

logistics services, for instance, may prioritize international roaming to track assets in flux, payment services depend on signal reliability, where terminals always connect, and select an alternative network in the event the first one fails.

To enable (current) global operation of mobile devices (both smartphones and things), the cellular ecosystem integrates many entities and procedures, including specific infrastructure to connect MNOs, business partnerships or the use of third-party Data Clearing Houses (DCHs) for billing. Many of these rely on specifications rooted in dated and arcane practices, involving long waiting periods for financial clearing, complex billing models, and disparate mechanisms for dealing with inter-MNO disputes. Specifically, to support customers of a Home MNO (HMNO) roaming in the network of a Visited MNO (VMNO) both networks must have a commercial agreement. With a technical solution in place, commercial roaming is then possible, and MNOs’ customers can use their respective partners’ networks to extend coverage. MNOs generate roaming revenue by charging their roaming partners as a function of the data/voice/SMS the partner’s users (inbound roamers) generate on the visited network. The roaming partners must each account the activity of roaming clients in a given VMNO. Then, by exchanging and comparing these records, the VMNO can claim revenue from the partner HMNO. Adding to this complex scenario is the challenge of the relationship between MNOs and other players in the cellular ecosystem, such as the IPX Providers. Fig. 4 shows a schematic view of the business interactions between entities in the cellular ecosystem to ensure that devices can currently enjoy global service.

Despite the fast evolution of radio technologies over the past decades, this logic of interconnection has remained largely unchanged. Even more, the associated platforms and systems are opaque, and with only little innovation. The advance of the inter-provider charging system is especially challenging, as it requires standardization, and subsequent joint evolution and deployment in the networks involved. This is becoming an important problem, as the current inter-provider charging system for global services imposes penalties in terms of performance, operational costs, and business revenues. At the same time, the vision for next-generation cellular systems (6G and beyond) sets an extraordinarily high bar for cellular networks, which are expected to provide smart and global connectivity to this massive volume of heterogeneous terminals, operating in different environments around the world. Meeting this ambitious goal requires growing the already substantial complexity of the cellular ecosystem to instantly orchestrate physical resources and functions across different network domains, in sync with time-varying user demand and multi-tenancy requirements, while offering a global service.

Creating an architecture based on a new trust model marks a significant change, requiring a shift to a paradigm where IoT device identities are not tied to a single operator. This approach, grounded in decentralized identity, separates user authentication from operator services, paving the way for innovative business models and dynamic customer charging strategies.

C. Global Service-Based Architecture (GSBA)

Both CCL and ZTL (together with the other legacy domain buses, such as 3GPP SBA) rely upon the Global (GSBA), which is meant to enable the declaration and management of services between different domains (e.g., network operators, and infrastructure providers). In our vision, a “domain” maps to an entity that owns different sub-domains, such as the radio access network, the core network, and the international carrier network. Within this ecosystem, one of the major challenges we aim to tackle is the inherent lack of trust between the different entities, which makes resource sharing and the deployment of novel business models difficult. While legacy domain buses such as 3GPP SBA shall assist the GSBA, there exist domains where new buses need to be developed.

1) *The RAN bus*: It operates within the RAN Intelligent Controller (RIC) platform, specifically utilizing a RAN bus that facilitates key performance measurements (KPM) collection from RAN nodes and dispenses RAN control (RC) decisions to manage the RAN infrastructure. The RIC platform hosts various xApps that leverage the RAN bus for several functions, including collecting RAN node performance data via E2 service model KPMs (E2SM-KPM), a service offered by the O-RAN Alliance. Additionally, xApps can modify Information Elements (IEs) in different signaling messages, facilitated by the RAN node using E2SM RAN Control (E2SM-RC), without needing to decode entire network messages.

Conflicts can arise within the RIC when xApps, especially third-party ones, alter IEs in ways that clash with each other. These conflicts are categorized as direct, indirect, and implicit, each requiring specific mitigation strategies. ORIGAMI’s approach includes designing conflict mitigation techniques to ensure smooth interoperability of these third-party xApps. The RIC’s Conflict Mitigation component plays a crucial role in resolving these issues, with strategies like post-action verification and tailored management for indirect and implicit conflicts. By effectively managing these conflicts, third-party xApp providers can optimize RAN node performance, with each xApp targeting distinct metrics for improvement.

2) *Network Intelligence bus*: As promoted by recent initiatives, the NI stratum has been proposed as part of the reference architecture for 6G. To define the NI stratum organization and operations, a reference representation of complex NI algorithms has been proposed, enforcing a hierarchy of NI Services (NISs), NI Functions (NIFs), and atomic NIF Components (NIF-Cs) (Fig. 5). This approach presents a unified framework that combines the operational hierarchy of NI components in the Network Intelligence Orchestrator (NIO) and the N-MAPE-K representation of NIF-Cs, advancing the vision of a complete NI stratum. The framework demonstrates how multiple NIF-Cs can create NIFs, which can be combined to form NISs, such as a reliable virtualized RAN (vRAN) service.

3) *Holistic SMO for Cost reduction*: Mobile operators are heavily investing in network performance optimization and Total Cost of Ownership (TCO) reduction. The GSBA will play a critical role in fostering the development of new Service Management and Orchestration (SMO) platforms. By




Fig. 5. Hierarchical taxonomy of NI stratum elements. An NIF corresponds to an individual NI instance that assists a specific functionality.

integrating existing silo-style tools and modules, such as radio design planning, RAN configuration, and monitoring and performance management, the GSBA streamlines the mobile radio access management and orchestration process. Efficient interoperability with O-RAN related modules, such as RIC modules, is achieved through the use of standardized interfaces, services, and open-source software. Additionally, the GSBA ensures the automation of the Life Cycle Management (LCM) for deployed apps and algorithms by adopting CI/CD and MLOps approaches, while enhancing network infrastructure awareness at the service level for increased efficiency.

To further improve network performance, we propose a transition from monolithic Network Functions (NFs)-based core architecture to a microservices-based architecture using a service mesh paradigm. Its theoretical and quantitative assessments can help evaluate its impact on the existing architecture, ultimately driving innovation and optimization within mobile networks. The implementation of an API manager adhering to the standardized 3GPP Common API framework is crucial for ensuring secure and interoperable API provisioning and consumption. This API manager facilitates streamlined interactions between components, improving overall network management. The GSBA’s connection to domain buses, e.g., RAN bus, NI bus, and 3GPP bus (the current 3GPP Service-Based Architecture model), enables more cohesive and efficient management of network resources, promoting collaboration between intelligent algorithms for optimal operation.

These ideas coalesce to create the foundation for novel closed loops that can be harnessed by specific NI algorithms and innovative global services. By effectively connecting scattered and siloed domains into a unified fabric, the GSBA will allow intelligent algorithms to coexist and cooperate according to Network Operator policies. This integration also makes previously unattainable interactions possible, such as the global orchestration of RAN resources involving the RIC.

As a consequence of the above, the GSBA will act as a gateway for more flexible interactions between Service Providers and Network Operators, addressing existing barriers and enabling novel control loops. These control loops present new opportunities for both parties: Service Providers can enjoy a broader range of exposed capabilities, leading to seamless customization, while Network Operators can diversify their service portfolios based on service type, associated KPIs, and customization capabilities. By adopting tiered pricing

models akin to those used by Software-as-a-Service (SaaS) companies, Network Operators can monetize various levels of functionality and appeal to a wider customer base.

IV. NETWORK INTELLIGENCE (NI) AND GLOBAL DIGITAL SERVICES

ORIGAMI aims to collaboratively design and develop innovative network architectural extensions that ensure synergistic operation and address the challenges mentioned earlier, and utilize them by specific NI algorithms that capitalize on the newly created closed loops. On the one hand, ORIGAMI's **Operator Control Loop** extends existing domain-specific control loops, like those in the RAN or core domain, to achieve a holistic network view. The GSBA unifies these domains, allowing intelligent algorithms to cooperate under Network Operator policies. This integration enables operations like global RAN resource orchestration, which standard 5G architecture can't support due to the need for tight inter-domain coordination. On the other hand, the ZTL introduces the **Hyperscale Control Loop**, enhancing Service Providers and Network Operators' collaboration. This enables a broader range of network capabilities and service customization, allowing Operators to offer diverse services and monetization models, similar to SaaS companies. To effectively enable this vision, a number of NI functionalities and services will exploit these closed-loops, on top of ORIGAMI's architectural innovations. We now discuss the solutions proposed by ORIGAMI to overcome the barriers discussed in II.

Data-driven task offloading for reliable vRAN acceleration. To address Barrier #1, ORIGAMI proposes a vRAN approach, focusing on resource pooling and opportunistic offloading to the most efficient computing resources. This approach aims to reduce costs without compromising reliability.

Enhancing Management, Stability, and Security in the 6G Architecture. ORIGAMI targets advanced ML models for network management in Open RAN architectures, addressing inter-operability issues of Barrier #2. It focuses on anomaly detection, traffic forecasting, and dynamic control, leveraging its architectural innovations for enhanced data collection.

Integration of novel compute paradigms and robust Network Intelligence. ORIGAMI aims to develop a comprehensive NI toolbox for 6G networks, focusing on scalable learning algorithms, hybrid learning approaches, and deep-learning-based optimization. It also explores Bayesian learning and Quantum computing-assisted optimization for tackling complex 6G problems. This toolbox will address Barrier #3.

Effective, distributed and streamlined access to u-plane computing capabilities. Addressing Barrier #4 (Underutilized modern programmable transport), ORIGAMI proposes solutions for integrating ML models into u-plane programmable hardware, developing distributed solutions for heterogeneous user planes, and streamlining the development and deployment process for user-plane intelligence.

Enabling the Global Operator Model: ORIGAMI leverages DLT to enable dynamic partnerships in the ecosystem, focusing on real-time billing, automated dispute resolution, trust,

and confidentiality. ORIGAMI proposes a blockchain solution for direct interaction among mobile operators and supports different charging models, tackling Barriers #5 and #6.

Anomaly Detection. ORIGAMI integrates real-time anomaly detection functionalities, especially for M2M platforms, that address Barrier #7. Thus, it proposes the use of IoT Global SIMs and focuses on proactive anomaly management, leveraging federated learning for efficient detection and response.

Network Core traffic analysis and optimization. ORIGAMI explores the benefits of a Service Communication Proxy (SCP)-enabled 6G network core, using ML techniques and graph theory for traffic analysis and optimization. In this way, ORIGAMI aims to reduce control-plane signaling overhead and improve network core operation efficiency.

V. CONCLUSION

The ORIGAMI framework effectively overcomes existing network limitations, crucial for advancing to 6G. We introduce three key innovations: the Global SBA (GSBA), Compute Continuum Layer (CCL), and Zero-Trust Layer (ZTL), each addressing specific barriers in network efficiency, security, and scalability. Additionally, we show how integrating concrete Network Intelligence and Global Services into our framework enhances its capability to handle complex network demands. These elements collectively mark ORIGAMI as a significant step forward in mobile network evolution, offering a comprehensive approach to the next generation of connectivity.

ACKNOWLEDGEMENT

ORIGAMI project has received funding from the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101139270. L.E. Chatzieftheriou is a Juan de la Cierva awardee (JDC2022-050266-I), funded by MCIU/AEI/10.13039/501100011033 and the European Union "NextGenerationEU"/PRTR, and MADQuantum-CM project, funded by the Regional Government of Madrid and the EU "NextGenerationEU"/PRTR.

REFERENCES

- [1] J. Rendon Schneir, J. Bradford, K. Konstantinou, et al., "Cost assessment of elasticity for a 5g broadband network," *Transactions on Emerging Telecommunications Technologies*, vol. 33, no. 4, p. e4425, 2022.
- [2] J. Ayala-Romero, et al., "Mean-Field Multi-Agent Contextual Bandit for Energy-Efficient Resource Allocation in vRANs," in *IEEE INFOCOM 2024-IEEE Conf. on Computer Communications*, pp. 1–10, IEEE, 2024.
- [3] L. Lo Schiavo, et al., "YinYangRAN: Resource Multiplexing in GPU-Accelerated Virtualized RANs," in *IEEE INFOCOM 2024-IEEE Conference on Computer Communications*, pp. 1–10, IEEE, 2024.
- [4] S. Alcalá-Marín, et al., "Global mobile network aggregators: taxonomy, roaming performance and optimization," in *Proc. of the 20th Annual Internat. Conf. on Mobile Systems, Apps and Services*, pp. 183–195, 2022.
- [5] A. Lutu, D. Perino, M. Bagnulo, and F. E. Bustamante, "Insights from operating an ip exchange provider," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pp. 718–730, 2021.
- [6] F.Z. Yousaf et al., "Network slicing with flexible mobility and QoS/QoE support for 5G Networks," in *2017 IEEE International Confer. on Communications Workshops (ICC Wrkshps)*, pp. 1195–1201, IEEE, 2017.