Compute and network virtualization at the edge for 5G smart cities neutral host infrastructures

Michele Paolino[+], Gino Carrozzo[++], August Betzler[x], Carlos Colman-Meixner[xx], Hamzeh Khalili[x], Shuaib Siddiqui[x], Teodora Sechkova[+], and Dimitra Simeonidou[xx],
[+]Virtual Open Systems, Grenoble, France; [++]Nextworks, Pisa, Italy; [xx]University of Bristol, Bristol, UK;
[x]i2CAT Foundation, Barcelona, Spain
Corresponding author: Michele Paolino, m.paolino@virtualopensystems.com

***Abstract*—Smart cities are one of the most important 5G verticals due to their impact on people's life. The neutral host model is key for this vertical, specially to promote the infrastructure sharing between operators for a pervasive infrastructure deployment. However, it demands enhancements in today's virtualization technologies to support geographically scattered and resource constrained computing and networking elements (i.e., core data centers, edge computing point of presences, and far-edge servers at lampposts, etc.). This paper presents the virtualization enhancements developed within the H2020 5GCity project to support the neutral host model for smart city deployments. The envisioned virtualization technology demands security and technology extensions with a unified view of sliceable and heterogeneous devices and radio technologies (e.g., LTE, 5G, and Wi-Fi). In our solution, security is enhanced by deploying a robust compute node authentication, monitoring, and geo-tagging. In addition, wireless connectivity is extended via an innovative multiple RAN controller approach for the management and control of heterogeneous radio resources. In order to be validated, our virtualization approach is being deployed and demonstrated in the cities of Bristol, Barcelona, and Lucca.**

I. Introduction

The smart city paradigm can provide multiple valuable services to citizens with strict performance, security, and sustainability requirements for the urban infrastructure and local business. The 5G infrastructure promises high performance and resiliency at the cost of higher density in network access points deployments. In urban contexts, 5G will require multiple macro-cells coupled with an even more pervasive deployment of small cells connected to multiple edge Point of Presences (PoPs) (e.g., multi-access edge computing nodes).

Various techno-economic analyses forecast prohibitive costs of pervasive 5G deployments for mobile network operators in case they need to build their own infrastructure [1]. In this market scenario, municipalities and local governments owning public assets (e.g., streets and districts) can take a key role as neutral hosts to offer infrastructure-as-a-service (IaaS) for Mobile Virtual Network Operators (MVNO) and verticals. The neutral host model provides a framework for a infrastructure owner to slice its compute and networking resources according to the needs of verticals providing services for end users. However, the neutral host model requires higher levels of security and new virtualization schemes to deal with pervasive 5G infrastructures and heterogeneous smart cities verticals. Hence, in this work we introduce the security and virtualization enhancements developed by 5GCity project for edge/far-edge, wireless, and multiples PoPs.

978-1-7281-3627-1/19/$31.00 © 2019 IEEE

This paper is organized as follows: Section II describes the objectives and architecture of the 5GCity project. Section III presents the virtualization enhancements introduced, while Section IV provides an overview of related works. Finally, Section V concludes the paper.

II. The virtualization components of 5GCity

H2020 5GCity project aims to validate the benefits of 5G technologies in smart cities by adding neutral host functionalities on a distributed cloud and radio platform [2]. The 5GCity neutral host platform leverages on the integration between Network Functions Virtualization (NFV), Software Defined Networks (SDN) and MEC on distributed cloud and radio platforms. This way, it will allow infrastructure owners to monetize their investment and service providers to deploy collaborative and innovative applications and finally improve the end users Quality of Experience (QoE).

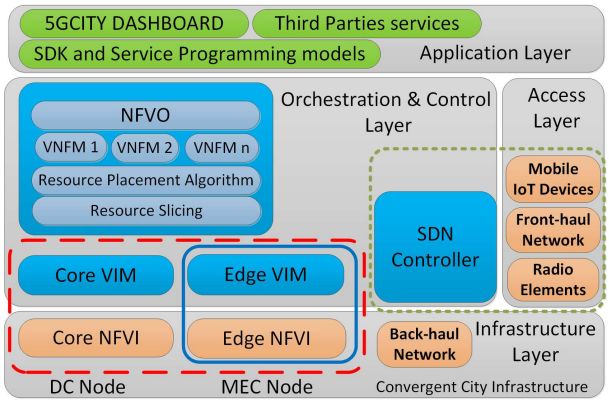

Fig. 1. 5GCity Architecture highlighting the contributions of this paper: Multi-VIM (red dashed box), EdgeVIM/EdgeNFVI (blue box) and wireless virtualization (green dotted box)

The 5GCity functional architecture (Fig. 1) is composed by four layers; Application, Orchestration and Control, Wireless Access and Infrastructure. The last two layers are organized in three different tiers which integrate both Cloud and Edge. The three-tiers of 5GCity are:

1) A centralized tier where massive computing resource are deployed, i.e., data center (DC) nodes,
2) An edge tier, geographically dispersed, with limited computing resources (e.g., MEC nodes composed by µdata center/street cabinets),

3) A far-edge tier with resource constrained devices at lampposts or other street furniture.

Thanks to the virtualization solutions developed during the project and detailed in Section III, 5GCity provides a neutral host platform by deploying the three tiers in combination with Radio Access Network (RAN) (Fig. 2). The following subsections provide more information about the three conceptual pillars that form 5GCity: the neutral host concept, MEC, and RAN virtualization.

A. The neutral host platform

The term "neutral host" combines two concepts - the aspects of "hosting" and "neutrality". Hosting aspect refers to an entity that provides set of resources available for clients such as mobile network operators to allows them to provide continuous services. Neutrality aspect refers to the host with shared platform to multiple clients or tenants. Neutrality in this context does not implies strict equality between hosted clients, as the resources offered are subject to commercial agreements between the neutral host and the hosted clients that require policy-based management. From a users point of view, the system behavior and services using the resources of a neutral host should be available without user intervention and, ideally, these should be seamless and identical to those provided by the hosted clients dedicated resources.

Technically, the neutral host model allows the building of end-to-end segmented slices, which encompass a wide variety of resources (network, storage and computing). Those slices are leased to service/content providers, which in turns can use the virtual resources they have been assigned and map their services to sets of slices [3].

The 5GCity neutral host platform will allow flexible end-to-end network slicing allocation schemes [4], policies definition to support Service Level Agreements (SLAs) and at demonstrating scale up/down of infrastructure resources, assigned to service providers and/or 5G verticals. Smart cities entities cover pivotal role within the 5G framework and a city municipality is the perfect candidate to cover the role of 5G neutral host. An example 5GCity big picture of neutral host and edge virtualization is summarized on Fig. 2.

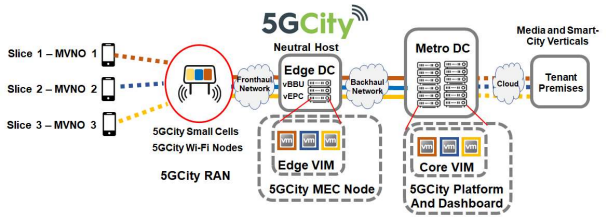

Fig. 2. 5GCity neutral host Vision with multiple MVNOs sharing virtualized functions at the edge

B. Multi-Access Edge Computing (MEC)

The neutral host model, to fully cope with 5G requirements in terms of bandwidth, coverage and latency, should also leverage on distributed edge resources by providing the capability to deploy end-to-end services which span across locally distributed pools of resources. This scenario is well described by MEC network architecture concept which enables cloud computing capabilities and an IT service environment at the edge of the network [5]. This environment is potentially characterized by applications running close to the user equipment thus (i) inherently ensuring ultra-low latency and high bandwidth as well as real-time access to radio network information that can be leveraged by applications, and (ii) offering a technology which could be suitable to be a technical enabler for 5G landscape. MEC architecture poses real challenges in the design of end-to-end services mainly since the resources locally offered can be limited, thus highlighting the need of tight centralized orchestration which can dynamically operate the lifecycle of edge computing applications [6]. MEC nodes will not only provide virtualization of computing at the edge but also the possibility for virtualized radio functions (Fig. 2).

C. Radio Access Network (RAN)

The RAN is an emerging architecture for the 5G framework and strictly coupled with MEC architecture [7], [6]. It is characterized by the virtualization of two main functions, Baseband Unit (BBU) and the Evolved Packet Core (EPC) that are moved away from the base station and placed in Central Offices (COs) or in the edge or cloud. This empowers the neutral host scenario with the capability of sharing the radio access part, by slicing its radio access network in multiple tenants operated by MVNOs [4]. Moreover, given the high flexibility of the neutral host framework the share models of Multi Operator Radio Access Network (MORAN) and Multi Operator Core Network (MOCN) can be deployed by providing a wide range of deployment solutions [3] An example of MORAN and MOCN is displayed on Fig. 2, where multiple MVNOs share virtualized functions at the edge through virtualization techniques described in this article.

III. 5GCity virtualization extensions

Virtualization is a key neutral host enabler because it abstracts computing and networking infrastructure resources to provide customers' applications with usable logically partitioned instances (e.g., virtual machines, unikernels, network slices, etc.). However, the high number of heterogeneous solutions interconnected to build the city infrastructure represents a challenge for the virtualization layer that is requested to support different tenants with multiple VIM and orchestration solutions, non-homogeneous and geographically scattered devices (smart gateways, lampposts, smart devices, etc.) and competing wireless technologies (Wi-Fi, LTE, etc.).

This section details the Multi-PoP, security and wireless virtualization extensions developed in 5GCity to support the concept of neutral host.

A. Multi-Points of Presence (Multi-PoP) virtualization

As mentioned in Section II, the 5GCity three-tier infrastructure is composed of different NFVI sites geographically distributed across the city, from the data center down to the edge PoP and the far edge (i.e., lamppost). Multiple NFVI-PoPs are generally deployed to allow the appropriate geographical distribution and availability of resources. These NFVIs can be under control of different Virtualized Infrastructure

Manager (VIM) and consequently implement diverse resource and service management approaches.

The orchestration of inter-NFVI connectivity related to inter-PoP Network Services is a critical factor in such deployments. Depending on the different types of VIMs in use, the implementation of the inter-NFVI network connectivity may vary. In 5GCity, we distinguish between: a) PoPs under the control of different types of VIM (e.g., OpenStack, fog05, etc.); b) single VIM (e.g., OpenStack) to control all the PoPs possibly in a multi-region deployment.

In this context, some important design factors to be taken into account when planning and deploying a NFV infrastructure include:

- *Single-VIM vs. Multi-VIM configurations*, e.g., depending on a mix of strategic policies and best practices to keep resources and infrastructures separated, e.g., for different application workflows;
- *Traffic steering capabilities* when virtual services are deployed across different PoP (e.g., at data center and at edge);
- *Ownership of the underlying network infrastructure*. e.g. single vs. multiple operators (i.e., NFVI owned and operated by single or by multiple operators).
- *Specific geographical constraints of the underlying physical infrastructure*, e.g. in terms of available connectivity across the three-tiers (core, edge, far-edge).

Infrastructure heterogeneity combined with the aforementioned design factors create different scenarios of deployment (Fig. 3) where a set of specific network features will be enabled to fully support 5GCity architecture and its intended Use Cases.

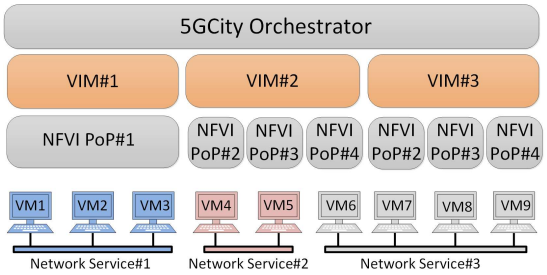

Fig. 3. Multi-PoP and Multi-VIM 5GCity (Scenario 1, 2 and 3)

Scenario #1. This scenario consists in standard single domain OpenStack deployment with single controller node co-located with a set of compute nodes (see Fig. 3 left for VIM#1 and NFVI-PoP#1). This scenario is the most simple scenario, and refers to an usual OpenStack installation. This scenario is recommended for LAB deployments or for early stage city pilot deployments where no edge nodes are available.

Scenario #2. It refers to a standard single domain OpenStack deployment with single controller Node and compute nodes deployed at data center and at edge level (see Fig. 3 center for VIM#2 and NFVI-PoP#2-4). To ensure that Layer 3 (L3) traffic coming from external networks (e.g., traffic from the user or traffic from IoT sensors) is at first handled at edge and then passed to core, the Neutron Distributed Virtual Routing (DVR) [8] service must be enabled. Neutron DVR allows to distribute L3 agents over the various compute nodes of the OpenStack Cloud, thus overcoming the limitation for L3 traffic (i.e., IP packets) to be all processed by the single L3 agent deployed in the OpenStack controller node, as in standard installations. In fact, the L3 agent deployed in the controller node is able to instantiate virtual network elements (i.e., virtual routers), and allows to a) route traffic between internal subnets; b) route traffic from internal subnet to external subnet (SNAT); c) route traffic from external subnet to internal subnet (DNAT) with floating IP.

The centralization of the routing functionalities in the controller node makes the standard OpenStack networking less efficient in a typical 5GCity scenario where the controller node is located in data center premises, compute nodes are located at the edge and IP traffic coming from external network must be consumed directly by VNF residing at the edge nodes. Through Neutron/DVR all the SNAT/DNAT operations can be performed locally at compute node level.

Scenario #3. It addresses a multi-domain OpenStack deployment, each domain consisting in at least a controller node and a set of compute nodes (see Fig. 3 right for VIM#2-3 and NFVI-PoP#5-7). OpenStack Tricircle [9] can ensure that Layer 2 (L2) cross-domain networking automation is enabled. Tricircle is a plugin for networking automation that works across Neutron instances in multi-region deployments. In order to allow the VNFs inter-connection across regions, Tricircle enables the L2 networking among distributed Neutron instances (e.g., IP address space management, IP allocation and L2 network segment global management). In terms of resource orchestration, Tricircle allows Neutron to work as one cluster in multi-region OpenStack clouds, actuating the orchestration of virtualized networking resources across multiple OpenStack clouds. All VNFs, in an OpenStack tenant domain and provisioned in different clouds, can be interconnected via the global virtualized networking resources.

It could be also possible to envision a design scenario for multi-domain hybrid deployments where far-edge VIM is a non-OpenStack deployment. In this scenario it is recommended to handle cross-VIM communications through an overlay tunnel. Overlay mechanisms available for OpenStack are VxLAN, GRE, IPSEC. Specific tunnel technology to be used depends on far-edge VIM capabilities (e.g., fog05).

Among the various configuration options, the most applicable to the specific physical infrastructures available in 5GCity [2] is the Scenario #2 with Neutron/DVR. This has been used in the 5GCity pilot at the City of Lucca (Italy) to implement direct routing towards the edge of the traffic generated by CCTV security cameras used to implement the Unauthorized Waste Dumping Detection use case. With DVR, we managed to route the IP streams from external CCTV cameras on dedicated networks directly into the edge computing instances where video analytics functions are deployed. Through this configuration, CCTV traffic can bypass the Neutron agent and router in OpenStack controller (deployed at 5GCity core), reducing latency for the connections and thus increasing TCP throughput over the link.

B. EdgeVIM and EdgeNFVI

Security and trust are particularly important in smart cities environments because of their distributed architecture and the potential privacy issues related to the the data they use. In fact, citizens data (coming from cameras, mobility services, health, etc.) needs to be well protected to avoid data leakages that can be sold or used for retaliations by attackers. The 5GCity Edge VIM and Edge NFVI provide a virtualization-based security and trust infrastructure for Arm-based edge devices that enable enhanced security, authenticated devices, geo/asset tagging and secure storage. This infrastructure includes VNF, NFVI and VIM-extensions, setting the ground of security and trust features at the lower level of the software architecture.

At the base of the 5GCity Edge VIM and Edge NFVI extensions there is VOSYSmonitor[10], a system partitioner for Arm devices that leverages Arm TrustZone to enable a Trusted Execution Environment (TEE)[11] (i.e., a secure area of the main processor that provides an isolated and trusted environment). This TEE is used for the implementation of a virtualized Trusted Platform Module (TPM)[12], a set of security features standardized by the Trusted Computing Group. The virtualized TPM (vTPM) functions are made available to the VNFs as well as to the hypervisor. As for the VNFs, vTPMs are used to enhance the security of the network functions with secure storage and cryptographic algorithms. Standardized and open APIs are used to call secure services, in a way that provides portability and legacy application support. With regard to the hypervisor, the vTPM functions are used to expose trusted computing features to the Edge VIM.

The 5GCity Edge VIM is based on OpenStack and leverages the Edge NFVI to support trusted computing functions. Asset tagging and geo-tagging are supported thanks to specific extensions to the OpenStack scheduler that have been developed in order to use an attestation service, coupled with an attestation agent that runs on each trusted compute node. Figure 4 shows how the attestation service is linked to the agents to certify the trustworthiness of a specific compute node for a given request. In fact, in order to enforce security and enable multi-tenancy, the attestation procedure is repeated for each request.

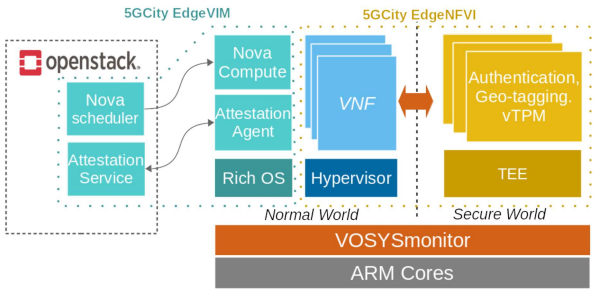

Fig. 4. 5GCity EdgeVIM and EdgeNFVI solution for trusted computing

The added benefits of security come at the price of performance penalty. The computational overhead of the attestation service and the secure infrastructure is measured in an experiment comparing it with a vanilla OpenStack deployment [13]. The results show a ∼2% increase of the average VM creation time, an acceptable result that can be further improved in future implementations.

C. Wireless Virtualization

5G will integrate different types of radio technologies, such as evolution of LTE, the 5G New Radio (NR) and Wi-Fi based technologies, and it will enable slicing to allow third parties to instantiate RAN connectivity on demand to provide services for smart cities [14]. In 5GCity, the targeted dense edge deployments are composed of a potentially large number of wireless LTE and Wi-Fi links, requiring the instantiation of multiple virtual networks over a single, shared physical infrastructure. Following the paradigms applied for compute virtualization (e.g., VIMs interacting with the NFVI to instantiate services), 5GCity introduces the concept of one or multiple underlying SDN-based RAN controllers that similarly handle the radio resources: a RAN controller exposes a set of resources to the 5GCity platform that can be allocated to different slices. As such, a slice can include a single or a group of wireless interfaces out of the entirety that is deployed in the city. In the same way a bare metal compute node can be virtualized to host multiple tenants and VMs, RAN elements can be virtualized and be associated to different services. Since RAN solutions can be heterogeneous, e.g. there can be custom devices deployed by the city using the 5GCity RAN controller, or commercial solutions with proprietary control software that come with their own RAN controllers, 5GCity implements an infrastructure abstraction. It enables the platform to support different types of RAN controllers and to integrate the underlying RAN technologies.

In wireless network mediums, virtualization can be performed in different ways. When speaking in general terms of wireless virtualization, techniques like time-hopping or TDMA enable slicing in the wireless medium. For Wi-Fi, SDN-based solutions have been proposed that enable the management of virtual access points, with initial work carried out in the definition of algorithms and protocols to enable resource allocation. In 5GCity we look at other ways to implement network virtualization with the same goals: to provide network slicing and isolation to support the neutral host case for 5G smart cities. The 5GCity wireless virtualization basically consists in sharing a physical wireless interface among a set of tenants or services and defining a configuration and management plane between the physical devices and the 5GCity platform. For example, in LTE, for each tenant a Public Land Mobile Network (PLMN) ID can be instantiated on the same carrier to differentiate between the tenants, each PLMN ID associated to a different service. In Wi-Fi, this corresponds to instantiating virtual access points on top of physical access points and attaching them to services hosted in the 5GCity infrastructure. In order to set up the RAN devices and to assign them to wireless slices, the 5GCity platform talks over a REST API with the deployed RAN controllers (Fig. 5). The configuration and management of the RAN devices is performed via the OpenFlow and OvSDB (data plane configuration) and NETCONF (device configuration) protocols. The wireless slices created in this way can be attached to any other services hosted by the 5GCity platform, being under control of its orchestration layer. In the city deployments we validate that both small cells and Wi-Fi nodes can be configured in any combination via the

platform and that up to 6 slices can be enabled per device. Further, using a traffic rate administration mechanism inherent to the 5GCity platform, QoS can be applied in the RAN by assigning different lower bound airtime ratios across the entirety of devices belonging to a slice via global optimization schemes. With the NS3 simulator, we validate this solution, e.g. using 5 Wi-Fi access points, running 5 slices on top, where each slice is assigned 20% of the airtime and where 6 terminals per slice and AP are distributed equally across the APs that generate On/Off traffic following a pareto distribution. In average, with the 5GCity solution we observe a deviation from the ideal slice airtime of only 1.8%, compared to an error of 6.7% when only using local optimization per Wi-Fi node.

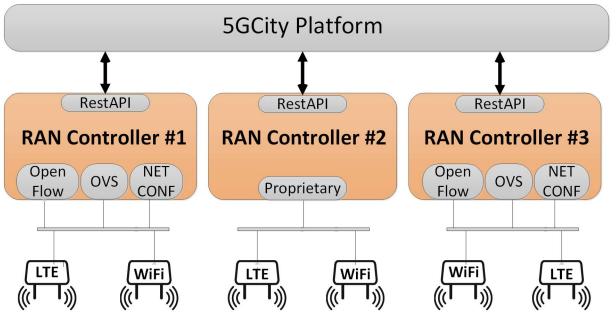

Fig. 5. 5GCity RAN Controller solution used to apply slicing to Wi-Fi, LTE and other radio access technologies.

IV. Related Works

Solutions presented in this paper relate to works tackling smart cities evolution towards virtualized infrastructures from various perspectives. For instance, authors of [6] apply the concepts of MEC with a focus on user mobility. They suggest QoS improvement by ensuring that the content is moving following the location of the user. On the other hand, the H2020 5GINFIRE project consortium [15] had a similar goal of driving forward the virtualization technologies convergence. The technical objective of the research is to build a 5G NFV-based reference ecosystem of experimental facilities for early tests and trials accelerating the verticals deployment. Still with the same focus is Mosaic5G [16] which provide as-a-service platform for 5G research. This is a collaboration of components on different levels. Among them are the FlexRAN, enabling a software-defined RAN and LL-MEC, a MEC platform aligned with ETSI. The project does not include the infrastructure and its management.

For what concerns the presented 5GCity virtualization extensions and in particular the Multi-PoP aspects, NFV infrastructures for 5G are generally distributed across multiple geographical locations to better cope with the actual availability of computing and network resources in various areas of the city/network. All the Cloud Management Systems (e.g., OpenStack, VMware vCloud, AWS, etc.) offer solutions for distributing the infrastructure PoPs across multiple interconnected sites, of different sizes and mixed/heterogeneous solutions. In this area, two major types of solutions emerged in state of the art to address the multi-site orchestration problem from different viewpoints. *Multi-site operations via multi-VIM* refer to cases in which the NFV orchestrator can handle various - possibly heterogeneous - Virtualized Infrastructure Managers (VIM). Relevant examples of this solutions are the ETSI Open Source MANO stack [17], which is capable to support multiple different VIMs including OpenStack, VMware vCloud Director, Amazon Web Services (AWS) and OpenVIM, and Tacker [18], which supports a multi-site OpenStack architecture. Instead, *Multi-site operations via multi-region single VIM* address cases in which the NFV orchestrator connects to a single VIM which in turns manages a NFV infrastructures further split in regions. Relevant solutions in this area are the OpenStack Tricircle [9] and Kingbird [19] projects which, respectively, provide network automation in Neutron in multi-region OpenStack deployments, centralized quota management and views of distributed virtual resources, synchronization of ssh keys, images, flavors, security groups, etc. across regions.

As for the security and trusted computing features of 5GCity, there are two main hardware technologies that form the basis of a trusted infrastructure. In particular, *Intel SGX* or Intel Software Guard Extensions are architecture extensions to provide secure memory regions (enclaves) in which dedicated parts of the applications can run protected and isolated [20]. On the other hand, *ARM TrustZone* allows the execution of a secure environment in parallel with a rich operating systems [21]. While the EdgeVIM and EdgeNFVI are based on Arm TrustZone and target the edge of the network, other works are mostly focusing on data centers and on Intel based trusted computing techniques[22]. Examples are given by [23], that is proposing a complete solution for a secure cloud infrastructure including platform integrity, external attestation, VM integrity together with a reference architecture. Similarly, [24] shows how to secure containers in a Kubernetes environment using Intel SGX support. On the other hand, the authors of [25] are integrating Arm and Intel Trusted computing technologies to solve the security problems of the infrastructure. The focus of this work is more at the IoT/Edge level and does not address the problems of NFV and smart cities.

Wireless virtualization for Wi-Fi has been discussed in a variety of works. The approach followed by the authors of [26] introduces slicing by using lightweight APs in form of virtual interfaces that are assigned on a per-client base and can be shifted between physical devices to support mobility. Yet, the proposal is missing the mechanisms to isolate the wireless resources consumed by each virtual interface as it is done in 5GCity. Another solution for the scheduling over IEEE 802.11 mesh networks is proposed by the authors of [27]. However, practical aspects are not discussed and the proposed solution does not provide orchestration of resources among virtual interfaces operating on specific channels. For a broader overview and the application of SDN in wireless networks we refer the interested reader to [28], a work that explains the benefits, but also highlighting drawbacks it introduces.

V. Conclusion and Future Work

In 5G smart cities environments virtualization technologies are challenged by the heterogeneity of wireless, hardware, and software deployments, as well as the security threats inherited by the neutrality of city infrastructure. To deal with those issues, the H2020 5GCity project is working to provide virtualization extensions for city wide infrastructures for neutral host model as enablers for 5G and smart cities deployments. In

this work we presented solutions to combine multiple VIMs and deploy them in multiple PoPs, to secure scattered edge devices based on ARM architecture and to manage Wi-Fi and LTE radio resources through an SDN-based RAN controller.

In more detail, as for Multi-PoP, a first deployment of Neutron/DVR has been completed in a real world scenario in the city of Lucca. Future work will go in the direction of supporting Scenario #2 (i.e., Neutron/DVR in case of L3 traffic) or Scenario #3 (i.e., Tricircle for L2 traffic) in all of the three targeted cities.

In terms of security extensions for the edge, a first prototype of the EdgeVIM is available with compute nodes authentication, system security monitoring, asset, and geotagging. The deployment of EdgeVIM is ongoing for the city of Bristol. The next development steps are to enable the virtualization of the vTPM at the VNF level, in a way that both VMs, containers and unikernels can be allocated on the city NFVI with trusted computing features for secure processing.

Finally, as for wireless virtualization, the RAN controller design will evolve to support additional wireless technologies, to provide enhanced features for mobility, as well as self-healing mechanisms, with the intention to contribute to 5G standards and specifications.

As part of the 5GCity schedule we are deploying our proposed approach in Bristol, Barcelona, and Lucca. A full validation will be conducted in the month of October 2019 by including the integration of multiple RAN controllers from different vendors, Multi-PoP scenarios and a demonstration of EdgeVIM capabilities against edge devices tampering and attacks.

ACKNOWLEDGMENT

This work was partly funded by the European Commission through 5GCity project (grant no. 761508) part of the 5G PPP in Horizon 2020 programme. The paper solely reflects the views of the authors. EC is not responsible for the contents of this paper or any use made thereof.

REFERENCES

[1] G. Smail and J. Weijia, "Techno-economic analysis and prediction for the deployment of 5g mobile network," in *2017 20th Conference on innovations in clouds, internet and networks (ICIN)*. IEEE, 2017, pp. 9–16.

[2] "H2020 5gcity project," https://5gcity.eu/, accessed: 2019-05-15.

[3] K. Samdanis, X. Costa-Perez, and V. Sciancalepore, "From network sharing to multi-tenancy: The 5g network slice broker," *IEEE Communications Magazine*, vol. 54, no. 7, pp. 32–39, July 2016.

[4] H. Khalili, A. Papageorgioui, S. Siddiqui, C. Colman Meixner, G. Carrozzo, R. Nejabat, and D. Simeonidou, "Network Slicing-aware NFV Orchestration for 5G Service Platforms," *EuCNC*, 2019.

[5] G. Baldoni, P. Cruschelli, M Paolino, C. Colman Meixner, A. Albanese, A. Papageorgiou, H. Khalili, S. Siddiqui, and D. Simeonidou, "Edge Computing Enhancements in an NFV-based Ecosystem for 5G Neutral Hosts," *IEEE NFV-SDN*, 2018.

[6] T. Taleb, S. Dutta, A. Ksentini, M. Iqbal, and H. Flinck, "Mobile edge computing potential in making cities smarter," *IEEE Communications Magazine*, vol. 55, no. 3, pp. 38–43, 2017.

[7] ETSI, "ETSI GS MEC 003 V1.1.1, Mobile Edge Computing (MEC) Framework and Reference Architecture," Tech. rep., GlobalPlatform Inc, Tech. Rep., 2016.

[8] "Openstack project - neutron distributed virtual routing," https://wiki.openstack.org/wiki/Neutron/DVR, accessed: 2019-05-15.

[9] "Openstack project - tricircle," https://wiki.openstack.org/wiki/Tricircle, accessed: 2019-05-15.

[10] P. Lucas, K. Chappuis, M. Paolino, N. Dagieu, and D. Raho, "Vosys-monitor, a low latency monitor layer for mixed-criticality systems on armv8-a," in *LIPIcs-Leibniz International Proceedings in Informatics*, vol. 76. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[11] GlobalPlatform, "TEE System Architecture v1.1, GPD_SPE_009," Tech. rep., GlobalPlatform Inc, Tech. Rep., 2017.

[12] N. Sumrall and M. Novoa, "Trusted computing group (TCG) and the TPM 1.2 specification," in *Intel Developer Forum*, vol. 32, 2003.

[13] T. Sechkova, E. Barberis, and M. Paolino, "Cloud & edge trusted virtualized infrastructure manager (vim) - security and trust in openstack," in *2019 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2019.

[14] J. S. Walia, H. Hmminen, and M. Matinmikko, "5g micro-operators for the future campus: A techno-economic study," in *2017 Internet of Things Business Models, Users, and Networks*, Nov 2017, pp. 1–8.

[15] A. Gavras, S. Denazis, C. Tranoris, H. Hrasnica, and M. B. Weiss, "Requirements and design of 5g experimental environments for vertical industry innovations," in *Wireless Summit (GWS), 2017 Global*. IEEE, 2017, pp. 165–169.

[16] N. Nikaein, C.-Y. Chang, and K. Alexandris, "Mosaic5g: Agile and flexible service platforms for 5g research," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 3, pp. 29–34, 2018.

[17] "Etsi open source mano - etsi osm," https://https://osm.etsi.org/, accessed: 2019-05-15.

[18] "Openstack project - tacker," https://wiki.openstack.org/wiki/Tacker, accessed: 2019-05-15.

[19] "Openstack project - kingbird," https://wiki.openstack.org/wiki/kingbird, accessed: 2019-05-15.

[20] V. Costan and S. Devadas, "Intel sgx explained." *IACR Cryptology ePrint Archive*, vol. 2016, no. 086, pp. 1–118, 2016.

[21] S. Pinto and N. Santos, "Demystifying arm trustzone: A comprehensive survey," *ACM Computing Surveys (CSUR)*, vol. 51, no. 6, p. 130, 2019.

[22] J. Greene, "Intel trusted execution technology (white paper)," *Online: http://www. intel. com/txt*, 2012.

[23] R. Yeluri and E. Castro-Leon, *Building the Infrastructure for Cloud Security: A Solutions View*. Apress, 2014.

[24] S. Vaucher, R. Pires, P. Felber, M. Pasin, V. Schiavoni, and C. Fetzer, "Sgx-aware container orchestration for heterogeneous clusters," *arXiv preprint arXiv:1805.05847*, 2018.

[25] R. Pettersen, H. D. Johansen, and D. Johansen, "Secure edge computing with arm trustzone," in *Proc. 2nd Int. Conf,. Internet Things, Big Data Secur.(IoTBDS)*, 2017, pp. 102–109.

[26] J. Schulz-Zander, L. Suresh, N. Sarrar, A. Feldmann, T. Hühn, and R. Merz, "Programmatic orchestration of wifi networks," in *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, ser. USENIX ATC'14. Berkeley, CA, USA: USENIX Association, 2014, pp. 347–358. [Online]. Available: http://dl.acm.org/citation.cfm?id=2643634.2643670

[27] R. Riggio, D. Miorandi, and I. Chlamtac, "Airtime deficit round robin (adrr) packet scheduling algorithm," in *2008 5th IEEE International Conference on Mobile Ad Hoc and Sensor Systems*, Sep. 2008, pp. 647–652.

[28] B. Dezfouli, V. Esmaeelzadeh, J. Sheth, and M. Radi, "A review of software-defined wlans: Architectures and central control mechanisms," *CoRR*, vol. abs/1809.00121, 2018. [Online]. Available: http://arxiv.org/abs/1809.00121

View publication stats