Network Slice Mobility in Next Generation Mobile Systems: Challenges and Potential Solutions

Rami A. Addad, Tarik Taleb, Hannu Flinck, Miloud Bagaa, and Diego Dutra

ABSTRACT

Network slicing offers numerous benefits, particularly the ability to deliver highly customizable services to new industry sectors that have been unserved or inadequately served by current mobile network operators. Among new industry use cases that are targeted by the fifth generation (5G) mobile systems, there exist scenarios that go beyond what the current device-centric mobility approaches can support. The mobility of low latency communication services, shared by a group of moving devices, e.g., autonomous vehicles that share sensor data, is a prime example of these cases. These use cases' demands for ultra-low latency can be addressed by leveraging the Multi-Access Edge Computing (MEC) concept, techniques for live migration of virtual resources, Software Defined Networking (SDN), and network slicing. In this article, we define different slice mobility patterns, different methods for grouping users, and different triggers for network slice mobility. Furthermore, we evaluate the mobility of services and network slices based on the simultaneous migrations of multiple containers.

INTRODUCTION

The 5th generation mobile networks (5G) will go beyond providing only high data rates for mobile users, as it will be a platform for a wider communication ecosystem for the Internet of Things and machine-type communications applications [1]. The services of new vertical industries, e.g., automotive, e-health, public safety, and smart grids, impose unique requirements that will push the envelope for high performance, scalability, and availability. To achieve such ambitious goals, 5G has been re-architected from the ground up in comparison to the previous mobile network generations. Decoupling logical network functions from the physical infrastructure forms the basis of deployment of self-contained, programmable and customizable networks [2].

Network Function Virtualization (NFV) and Software Defined Networking (SDN) are the fundamental technologies to implement the separation of logical network functions from the infrastructure leading to network softwarization [3]. SDN offers programmability of the connectivity between the network functions while NFV provides the means to define, instantiate and manage virtualized network functions that are required to create multiple logical (virtual) networks. The sharing of the same underlying infrastructure among isolated and self-contained networks leads to the key concept of network slicing [4].

Network slicing is an active research field within the academic community and among the different standards development organizations (SDOs), such as the Next Generation Mobile Network Alliance (NGMN), Third Generation Partnership Project (3GPP), and International Telecommunication Union — Telecommunication Standardization Sector (ITU-T). In the scope of this article, we define a network slice as an end-to-end (E2E) logical network running on top of a common underlying (physical or virtual) network, fully isolated, with independent control and management, and flexibly programmable to meet service level agreements (SLA) of a specific service. A network slice consists of computing and storage resources, associated with virtual networks, possibly composed of multiple virtual sub-network segments; and may span across multiple technological as well as administrative domains. Mobile users can connect to multiple slices simultaneously depending on the type of services they are using [5]. Furthermore, for latency-sensitive services, MEC [6] provides powerful service delivery with minimal delay. In a MEC environment, MEC servers, as well as network access points, can be sliced to serve multiple slices with very different service characteristics.

However, as users move with their mobile devices from one domain to another, their ongoing mobile communication and service sessions, running on at least a network slice, may suffer a drop in QoS if not a total disconnect. This may happen if their ongoing network service has to be served by instantiating a totally new network slice of the same type at the destination network (new radio resources, backhaul, computing, and storage). This leads to a network slice that has to support moving its computation, following its users' mobility patterns. Additionally, the slice itself may need to adjust its resource allocation, adding more resources or freeing unused ones. These requirements lead us to the definition of the concept of slice mobility.

Each instantiated network slice should ensure service continuity for its end-users. Intuitively, a simple service replication would solve the issue, yet not all services are stateless (i.e., services that do not save user session/context). It is worth not-

ing that the migration of stateless services is less demanding than the migration of stateful (i.e., services that save user session/context) services [7]. To support migration of stateful services, separate mobility mechanisms must be designed and implemented to ensure consistency of the service context. In addition, the resources of a slice could become saturated and over-consumed due to a sudden increase in the service demand, thus actions for slice scaling or breathing should be supported. Fortunately, the same actions that are needed for slice mobility are well suited for slice scaling. Based on these observations, the contributions of this article are:

- The introduction of a number of slice mobility patterns to optimally manage and use slices with their allocated resources, while leveraging the ideas of the Follow me Edge concept introduced earlier in [8]–[10].
- The examination of service migration capabilities needed for slice mobility inside one Infrastructure as a Service (IaaS) cloud domain, that we then extend to cover slice mobility across multiple IaaS domains.
- The definition of several key slice mobility triggers and user equipment (UE) grouping methods to enable efficient slice mobility needed for the described slice mobility use cases.
- A preliminary evaluation of two key enabling technologies for slice mobility patterns: system virtualization with containers and SDN to allow a fast and seamless allocation of resources to a network slice.

The remainder of this paper is structured as follows. The next section summarizes the fundamentally related research work. The following section gives an overview of different slice mobility use cases, followed by an introduction of the key slice mobility triggers and the different methods for grouping users besides slice mobility patterns. We then focus more on the enabling technologies for the proposed solution and the primary evaluation setup for the proof of concept. Finally, the paper concludes in the final section.

RELATED WORK

The authors of [11] focused on extensions for NFV orchestration that provide tailored support for mobility and QoS/QoE for network slices while ensuring efficient utilization of the substrate network resources. They presented several ways to classify and select mobility management (MM) schemes based on the context of the service or type of a network slice. Terminal speed, session continuation requirements, and stability of the endpoint addresses are the main criteria for selecting a mobility scheme. The authors proposed slice-specific MM schemes through the creation of context-dependent configurations of the instantiated network functions. Our envisioned approach differs from this one by focusing on moving slices along with their services based on mobility triggers and grouping of respective UEs.

As previously discussed, the migration process is one of the key enablers of the proposed slice mobility concept, whereby a running virtual instance (i.e., Virtual Machines (VMs) or containers) is migrated across different host machines (i.e., within the same IaaS platform or across multiple IaaS) without disrupting the service. Clark *et al.* [12] introduced the writable working set concept and used it to design a pre-copy based migration procedure that enables the live migration of VMs for the Xen Virtual Machine Manager (VMM). Mann *et al.* [13] presented a network architecture that provides a layer-agnostic and seamless live and off-line VM mobility across multiple data centers. They leveraged SDN and used the principle of location independence in order to handle the inter-data center limits. They obtained better results compared to the default layer 2 networks; in some tests, their solution outperformed the default approach by up to 30 percent. Tay *et al.* [14] evaluated the migration performance of containers and VMs showing that the use of system containers can achieve the same, or even better, capabilities than a VM without its high overhead.

With respect to the first cited work, in this study, we introduce the slice mobility patterns, and jointly emphasize the key enabling triggers and the grouping methods of the end users, while the remaining listed related works are used as a basis for enabling the desired slice mobility. As the entrance of new 5G use-cases will beget a highly mobile environment, increase the core network traffic and reduce the latency, this work is a must for achieving the 1 ms latency target for the upcoming 5G mobile systems and beyond.

SLICE MOBILITY USE CASES

Network slicing enables customization of network services and resources for the needs of different use cases. The same underlying infrastructure is shared among the different slices that add their own customized network functions and services to implement a dedicated service network. The services in a slice will have their own dedicated share of e.g. RAN resources, edge computing resources, and control plane resources that will be concatenated to form an end-to-end slice. Such a dedicated slice can be customized for a number of use cases, such as to support autonomously moving equipment (e.g. drones, robotic vehicles, or a fast-moving train carrying mobile users enjoying infotainment services) with mobility patterns and latency requirements that differ from a regular mobile phone. To ensure that the user mobility patterns match with the coverage and availability of the resources allocated to the slice, we introduce the concept of slice mobility that is implemented by carefully coordinated and organized live migrations of Virtual Network Functions (VNFs). In this section, we discuss use cases that clearly show the need for the proposed slice mobility paradigm that 5G and beyond networks should adopt. The first two use cases form a new mobility pattern not perceived before in the previous generation. They also express the perfect need for a mobile slice to follow the users (cars, and UAVs), ensure specific resource availability and continuous delivery during the mobility time.

Network slicing enables customization of network services and resources for the needs of different use cases. The same underlying infrastructure is shared among the different slices that add their own customized network functions and services to implement a dedicated service network.

Autonomous vehicles are thought to be one of the key verticals that would benefit from the upcoming 5G systems. Autonomous vehicles are expected to be served by mobile operators or car manufacturer-operated network slices. Safety of the passengers and the other public on the road mandates high reliability and low-latency connectivity to be an integral part of a slice for autonomous vehicles.

The last use case is proposed to illustrate another aspect of the slice mobility patterns related to the partial mobility which will be detailed in the next section.

DRONE TRAFFIC CONTROL

Unmanned aerial vehicles (UAVs) define the perfect example for predictable paths and high mobility, showcasing the slice mobility use case. As an initial scenario, let us assume that Jonathan is a young researcher working on UAVs, and his mission is to perform a self-swarming control test over two university campuses. Jonathan has ordered a customized slice for his drone experiment. In addition to offering the connectivity to the drones, the slice is hosting virtual flight controllers for all the involved drones. The virtual flight controller is a software application that controls a corresponding drone and receives location information from the drone and radio usage information from the network. The virtual flight controller application is instantiated in a container. Each drone has its own instance of the controller and the controllers need to cooperate to group the drones into swarms in an orderly manner so they move as a group to the same direction. In his experiments, Jonathan noticed two possibilities: the first one is that the swarms of drones can have a synchronized mobility pattern where the swarm moves as a tight group outside the current service area, or a second one where the drones move in smaller groups or one by one out from the initial service area depending on how tight the swarm is. In both situations, the containerized flight controller must follow its drone in order to provide the required low latency and accurate control of the drones. This use case leads to the notions of migrating the slice as a whole (i.e. full slice mobility) or the slice gets split over two service areas.

AUTONOMOUS VEHICLE SUPPORT

Autonomous vehicles are thought to be one of the key verticals that would benefit from the upcoming 5G systems. Autonomous vehicles are expected to be served by mobile operators or car manufacturer-operated network slices. Safety of the passengers and the other public on the road mandates high reliability and low-latency connectivity to be an integral part of a slice for autonomous vehicles. A missed signal from the virtual data and processing analytics of a connected car (vDPACC) can cause an imminent dangerous situation that could result in physical damage or even a death. In addition to the support of low latency and high reliability, the slice must also have a redundancy that can be implemented at the VNF-and connection-levels inside a slice or by a backup slice. The latter approach would lead to simpler overall service orchestration and configuration. When such a slice is relocated, the backup slice also needs to be modified or even moved as well. The slice mobility management mechanism must take into account the availability of backup connections, as well as the redundancy of the VNFs of the slice when migrating the slice and its services. The backup slice should not be relocated at the same time as the primary slice to ensure availability of the service.

RAPIDLY CHANGING VIDEO STREAMING NEED

In this use case, a city hall meeting is taking place in the form of a webinar. All the residents of the city have been invited due to the importance of the decisions to be made. The community has arranged its own slice for this mass meeting. The webinar is also attended by an audience on a train departing from the railway station of the city. The train leaves according to its schedule and enters soon to a neighboring area, or a tunnel, that has much less capacity to offer. The city hall videos' streaming slice needs to be complemented with a new and temporal slice in the new service area that the train is passing by. This new slice solves the bottleneck by accommodating a special Content Delivery Network (CDN) capacity with a video streamer that aggregates multiple separate unicast video streams into a limited number of shared streams that are distributed to the traveling audience of the webinar. Once the train enters into an area with better infrastructure, the purpose-built video CDN slice can be released and the original slice can continue to serve the audience. This use case shows the need for slice splitting and slice merging leading to the concept of slice breathing discussed in the following sections.

KEY ENABLING TRIGGERS, GROUPING ATTRIBUTES, AND SLICE MOBILITY PATTERNS

As was discussed in the use case section, slice mobility comes in a number of different variants or mobility patterns depending on how the slice is to be modified due to the changes in its service consumption. Thus, categorizing slice mobility patterns is a challenging issue. Based on what was introduced earlier, we identify the following slice mobility patterns: full slice mobility, slice splitting, slice merging, slice shrinking and slice breathing. Slice mobility actions are based on a number of triggers characterizing the slice dynamics and its service consumption. We start by introducing the slice mobility triggers that initiate slice mobility actions. The mobility triggers, as well as slice mobility actions, operate on a set of different groupings depending on a service or its resources. We elaborate the key grouping attributes offered by the 5G network specifications that provided the basis for slice mobility patterns and mobility triggers.

SLICE MOBILITY TRIGGERS

In this step, we will identify, discuss and present several key triggers to enable different slice mobility patterns. Triggers broadly relate to the users' mobility, the availability of physical resources and network resources at the hosting edge cloud or federated cloud, resource efficiency utilization, service reliability, and security.

Group Mobility Trigger: This trigger can be considered as the main catalyst for slice mobil-

ity in a real-life environment. The group mobility trigger could be applied to all use cases cited above (i.e., drone, autonomous cars, and video streaming). In this trigger, a group of users simultaneously move from one location to another. For example, passengers on board a metro or train move simultaneously from one location to another (as in the use case of video streaming in a highly moving entity), which requires full slice mobility. The signal strength can be measured by the users and reported back to the access points. These measurements can be correlated and acted upon already on the access points or deeper in the network, such as at the mobility management entity in the evolved packet system, access management function in the 5G core, or at the life cycle manager controlling the slice. The entity in charge of correlating these measurement reports will pull the group mobility trigger that will, in turn, launch the process of slice mobility to follow the mobility of that particular group of users, using the same slice and its services that were reported by the measurements.

Resource Availability Trigger: Edge clouds tend to always have fewer resources, i.e. network, processing, and storage capacity, than the centralized cloud. Due to the limited resources at the edge, the system level resource consumption must be monitored carefully. Once the upper limit of allowed total used capacity is reached, a "limited resources availability" trigger will be generated with a parameter indicating that the highest allowed resource consumption level of a certain type of system resources is reached. This trigger is sent to the slice life-cycle management to initiate migration of services from the highly-loaded edge cloud(s) toward the centralized cloud while keeping at the edge only delay-sensitive services. Alternatively, the services of the highly-loaded edge cloud could be migrated toward neighboring edge clouds if they have the necessary system-level resources available. Once the resource consumption level decreases below a given parametrized threshold, the edge cloud should send a "resource availability" trigger to indicate that there is room to accommodate more users and services for that particular type of resources.

Let us consider that the users are static and the network resources are exhausted by the requested data from a given group of users. In this situation, a scale-out operation needs to be considered. However, the slice is limited by the available physical resources of the edge cloud and the number of different slice migration strategies that need to be considered with potential impacts on the service QoE. One possible trade-off is to migrate some users or services to other edges or to the centralized cloud that is further away even if that may reduce their QoE.

Reliability Trigger: The reliability trigger would be typically generated by the operation and maintenance protocols and supporting operations, administration and maintenance (OAM) systems of the access point or the edge cloud that monitors the healthiness of the connectivity. In case of a major disaster, there would be an abrupt interruption of connectivity toward the users or neighboring edge nodes or even to the centralized cloud. The services should be therefore simultane-

The key mobility triggers for slice mobility are not alone sufficient to efficiently manage the slice mobility patterns. In addition to triggers, we need means to group various relevant objects so the mobility of the slices of resources consumed by those groups can be separately and efficiently managed.

ously evacuated from one location to another as long as there is still some capacity left for that. All unnecessary actions should be deferred until the system snapshot has been replicated. This trigger will cause full slice mobility where all slices, served by this access point or edge cloud node, need to be migrated elsewhere to ensure the integrity of all valuable data.

Security Trigger: Security is of vital importance for the networks since a compromised entity or service may result in considerable damage to the whole infrastructure. For instance, in case a denial of service (DoS) attack occurs in a specific location "A", the services should be then shifted from that location to a more secure one. An intrusion detection system (IDS) could send a security trigger to start a slice mobility for the compromised slice. When the IDS detects an anomaly in a slice, this trigger is sent to the orchestrator that initiates the migration process for that slice.

Request Overload Trigger: Service request trigger is based on the number of simultaneous service requests stemming from groups of users requesting the services available in an existing slice. As a potential number of requests is likely to overload the requested service and eventually to cause a trigger for limited availability of resources, this overload trigger is sent prior to any hard resource limitations and to initiate smooth overload control and potential redistribution of the service across neighboring nodes.

Service Consumption Trigger: The resource consumption of individual services needs to be monitored. Based on the type of service and its resource consumption, a trigger is generated if the service consumption is under or above predetermined levels. Note that the previously-discussed resource availability trigger deals with the aggregate system level resources whereas this trigger copes with the performance of a single service.

Grouping Attributes

The key mobility triggers for slice mobility are not alone sufficient to efficiently manage the slice mobility patterns. In addition to triggers, we need means to group various relevant objects (e.g. services, users, and network functions) so the mobility of the slices of resources consumed by those groups can be separately and efficiently managed. Next, we will investigate various grouping attributes to support slice mobility patterns.

Grouping by User Subscription Type: Grouping users by using their distinctive subscription types are one of the simplest grouping attributes. In 5G, this identifier is called the subscription permanent identifier (SUPI) and it is globally unique throughout the 3GPP system [15]. Another useful identifier in 5G is the generic public subscription identifier (GPSI) that identifies a 3GPP subscription for different data networks. With these subscription identifiers, we can separate, for example, IoT users of a given service provider from mobile
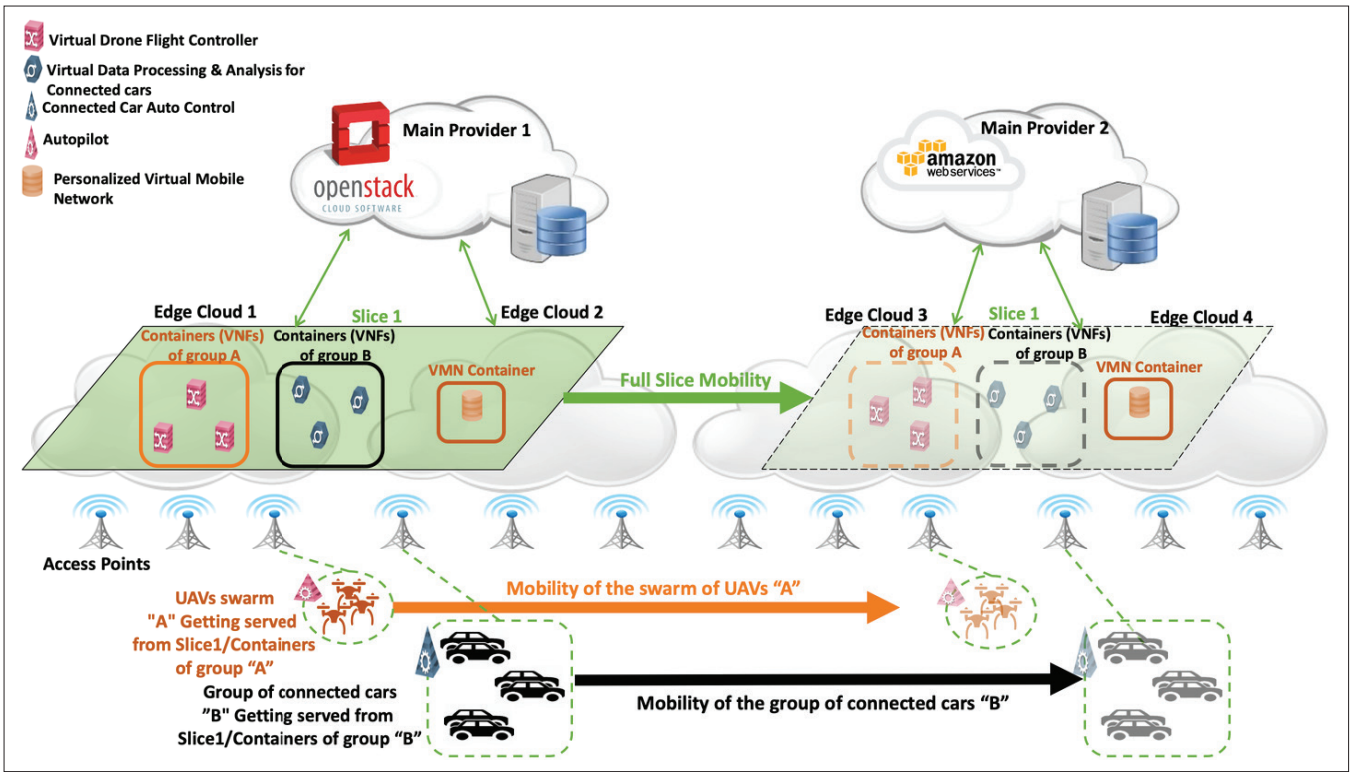
FIGURE 1. Full slice mobility pattern.

broadband users of the same or different providers, or prepaid users from enterprise users. Because this grouping is very coarse, it will be performed in combination with other types of groupings.

Grouping by Access Type: Users can be grouped by the access type they use. Access types could be 3GGP (e.g. 4G, 5G, 5G small cell) or non-3GPP (e.g. WiFi and Wi-Max) accesses or multi-access. This grouping attribute is coarse and often needs to be complemented with other grouping attributes.

Grouping by Network Slice Type: 3GPP defines three standard types of slices:
• eMBB (enhanced Mobile Broadband) slice
• URLLC (ultra-reliable low latency communications) slice
•MIoT (massive IoT) slice.
A UE can be connected at the same time to multiple slices. A slice is associated with a slice identifier, called network slice selection assistance information (NSSAI) that contains information about the slice type [15]. When a UE connects to the 5G network, it will use this identifier to express which slice it wants to join. NSSAI is also used in binding services and 5G network functions to a particular slice. It is very critical for slice mobility to identify that particular slice that will be modified or moved as it impacts multiple bindings between multiple entities. Unfortunately, NSSAI is unique only within one operator domain which introduces additional complexity when a slice is moving across operator boundaries.

Grouping by Service Area: Grouping of the consumed resources only is not enough because service area and service availability limitations may exist. A trivial example would be the case of a user "A" that is consuming an ultra-low latency service. If user "A" is a static user or his mobility pattern is restricted to a single service area, this user cannot be grouped with other users that are allowed to use the low latency service in other service areas.

Grouping by Access Characteristics: Grouping users by their experienced access characteristics constitutes one of the essential groupings attributes to identify a slice mobility pattern. Access characteristics take into account the radio metrics of the access, throughput, and frequency of handovers. For example, in our third use case where passengers are traveling in a high-speed train and receive video service over a network slice, one useful grouping of the users would be based on users experiencing same radio characteristics, same radio frequencies, same handover frequency and using same access points in addition to receiving the same service.

Grouping by Geographical Location: Geographical location is an obvious grouping attribute that is applicable to all above-mentioned use cases. Geographical location-based grouping is often combined with the other grouping methods, particularly with the mobility pattern grouping. In Geographical location-based grouping, users, services or slices are classified based on their current location.

SLICE MOBILITY PATTERNS

We classify slice mobility patterns into the following categories: i) full slice mobility; ii) partial slice mobility which includes slice breathing, slice splitting, and slice merging; and iii) slice mobility optimizer which contains slice shrinking pattern. In the remainder of this section, each slice mobility pattern is described, highlighting its respective triggers. As discussed earlier, slice mobility
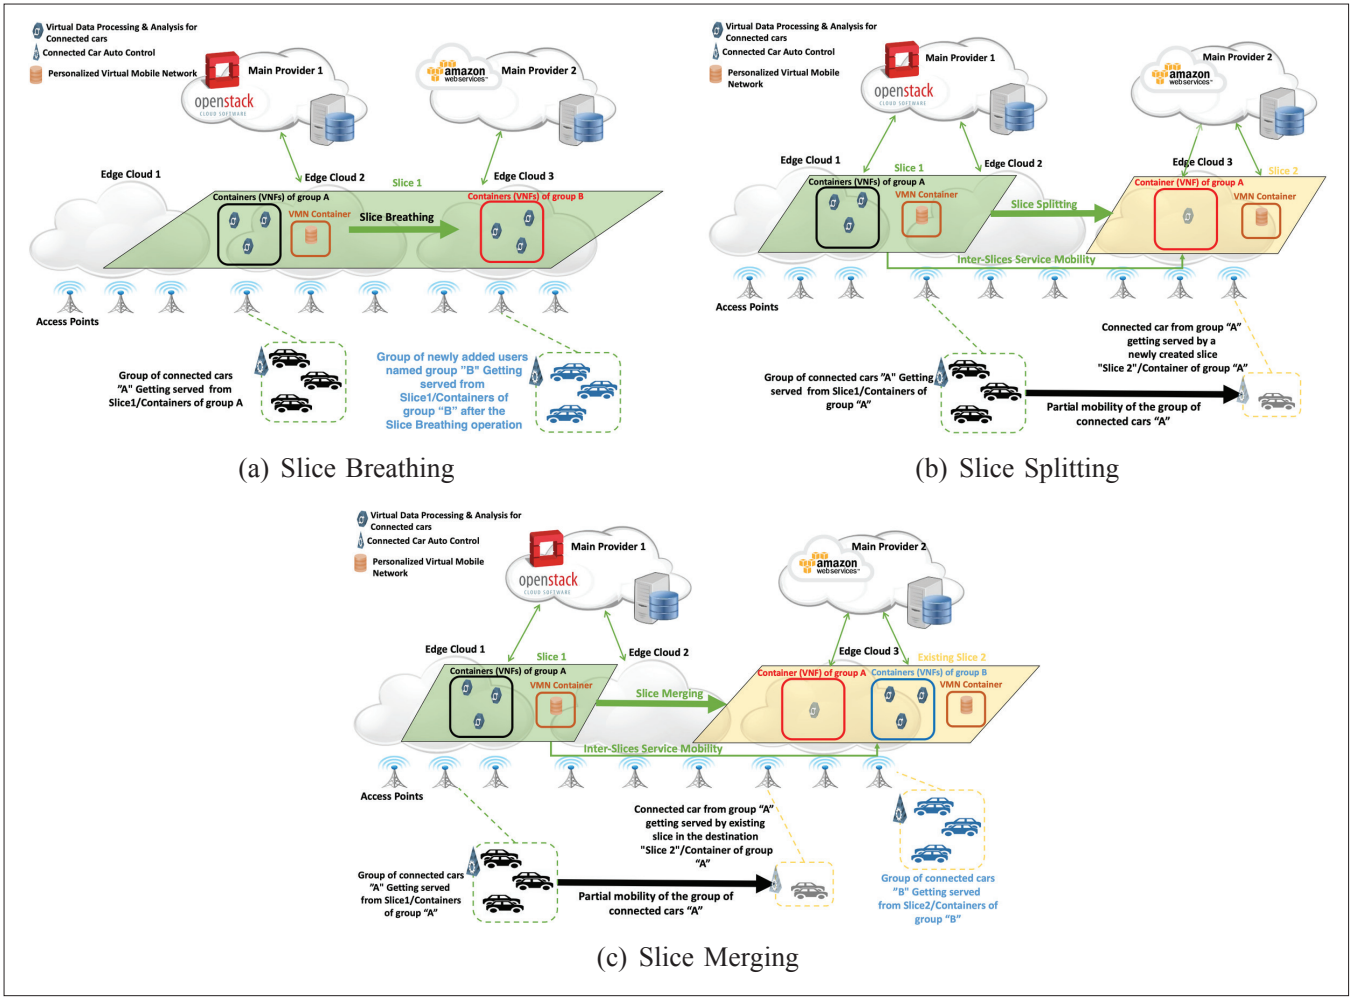
FIGURE 2. Partial slice mobility patterns.

events are generally triggered by the mobility of a group of end-users, the availability of the needed resources in the cloud or by the security aspects of the requested services.

Full Slice Mobility: In this use case, we consider a group of users moving to a different location; e.g., a swarm of UAVs moving from the service area of an edge cloud 1 to the service area of another edge cloud 3 (Fig. 1). This group mobility will trigger a service migration process of the entire services and associated resources used by this group. This leads to the notion of full slice mobility as all resources and the services of a slice are impacted by the mobility and are migrated from the original resource location to a new one. Several triggers may be used to indicate the need for this type of mobility. The most important ones are group mobility, limited resource availability, and reliability triggers.

Partial Slice Mobility: Full slice mobility is an expensive operation but some of the mobility patterns can be supported with less tedious operations. In partial mobility, only some identified resources of a slice are to be migrated. Thus, there are two possible scenarios: either the network slice will be extended to allow more coverage (Slice Breathing) or a different network slice will be considered in the destination. For the latter, we consider two cases: i) either existing slices; or ii) newly-created slices. Both of these divergent

cases yield two types of slice mobility, namely Slice Splitting and Slice Merging.

Slice Breathing: Another use case, depicted in Fig. 2(a), refers to when the group of users "C" attached to the existing slice (slice 1) causing skewed resource consumption, over-consuming a subset of the offered services of the slice 1. A slice breathing operation will be triggered, causing replication of the content of highly-loaded micro-services (e.g. containers), followed immediately by user redirection to this newly-created (sub-)slice to guarantee seamless continuity of the services.

The slice breathing operation can be based on several triggers. Clearly, the group mobility trigger, the resource availability trigger, the service overload trigger, and the service consumption trigger can be used to initiate a slice breathing mobility pattern. To summarize, in the slice breathing operation, a slice is temporarily expanded by combining service replication and service migration processes with redirection of users to another slice to match a sudden need for scaling a slice. In the case of unfeasibility of slice breathing operations, slice splitting and slice merging are adopted depending on the situations.

Slice Splitting: In the slice splitting case, we opt to create a new slice for the upcoming group of users as shown in Fig. 2(b). This action will be permitted by doing an inter-slice service mobility to
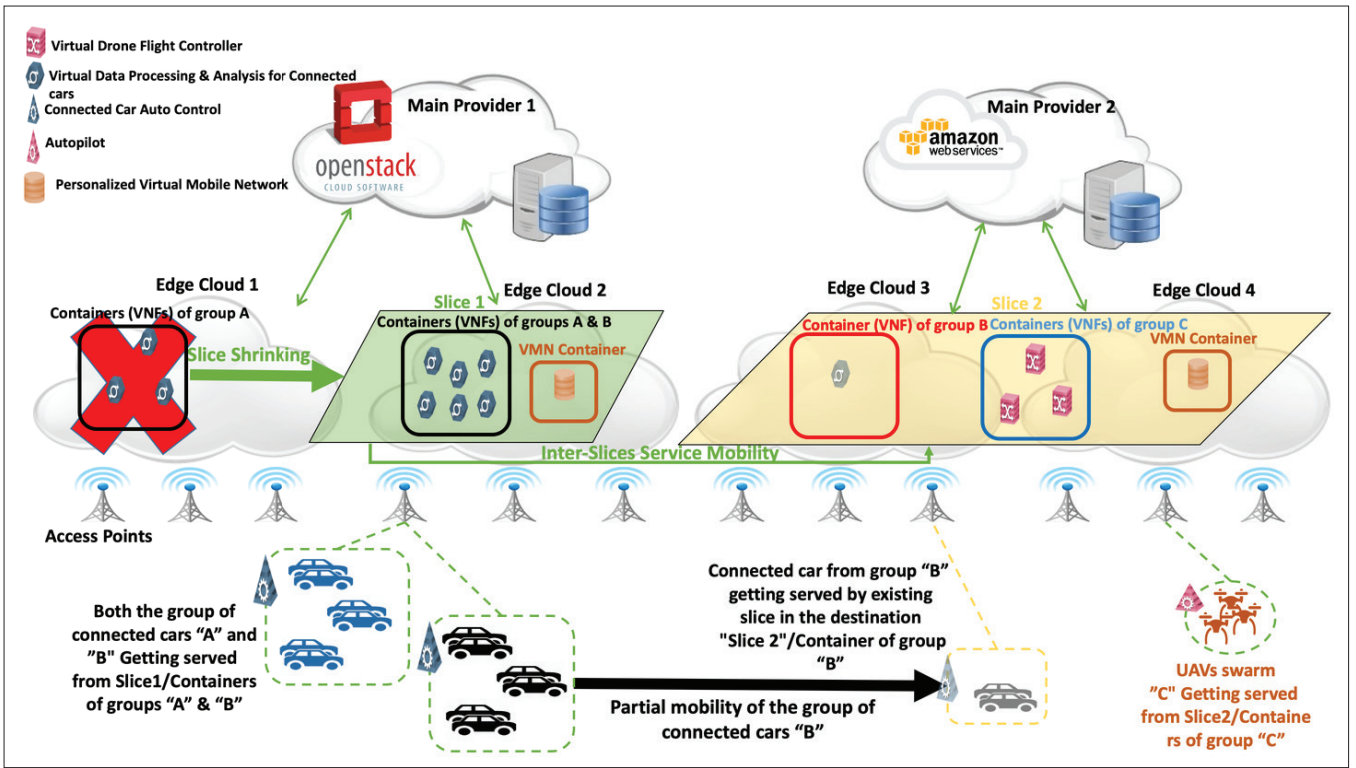
FIGURE 3. Slice shrinking.

ensure the availability of the service in the newly-added slices by exploiting the ability of inter-data-center multiple migrations.

The group mobility trigger, the resource availability trigger, the security trigger, the service overload trigger, and the request overload trigger will be part of the envisaged triggers for the slice splitting use case. However, for these slice mobility patterns, the system should consider the service consumption behavior of end-users as a trigger to choose between slice splitting or slice merging. Effectively, if the requested service is not a delay-sensitive service, the system may opt for a slice splitting action, wait until the creation of the new slice, and start the mobility process by performing migration and replication actions.

Slice Merging: On the other hand, considering the availability of interoperability between two slice providers (i.e., two IaaS), the system may optimize the distribution of the new containers forming a set of slices and allowing an inter-slice connectivity. This kind of slice mobility could be very useful in case of no network coverage from the initial slice provider. A multiple migrations process added to a replication process should be envisaged to enable this use case.

The slice merging mechanism presented in Fig. 2(c) is a quite similar mechanism to slice splitting with the exception of the creation phase of the slice because in this case ultra short latency services will be considered, which means that we cannot tolerate the extra time for the slice creation. Due to this constraint, a slice merging mechanism is employed. In case of different slice providers, an interoperability contract is intuitively assumed to be established between these different providers.

Slice Mobility Optimizer: The Slice Mobility Optimizer mechanisms aim for avoiding the waste of unused resources as a result of the mobility of users, a reduced usage of services, and/or the end of the service usage.

Slice Shrinking: In Fig. 3, we illustrate the case of a group of users that moved and are served by a different slice (i.e., either an existing one or newly created) in their new location (i.e. Edge Cloud 3 & 4). The remaining users are served by Edge Cloud 2, and resources originally used in Edge Cloud 1 will be released, resulting in the shrinkage of the original slice. Slice shrinking is a cleaning mechanism that should take place to release all previously used resources that became obsolete during the migration process.

After deciding which slice mobility pattern to execute, the system should run a control algorithm that determines the optimal number of virtualization instances and sets the rules for slice shrinking. For instance, in Fig. 3, we have two containers serving the groups of users A & B. The system can use the number of requests that container 1 or 2 is able to handle. If either of them is able to provide the required service with respect to the SLA terms negotiated before, it can take over the services provided by the other one, that will be ultimately turned off.

Table I summarizes the relationships between the slice mobility patterns and the different mobility triggers. The slice shrinking pattern is not included in the table as it is supposed to be an automatic process activated after each slice mobility pattern to free resources that become unused after resource migration.

ENABLING TECHNOLOGIES AND EVALUATION SETUP

This section presents our preliminary evaluation of two key enabling technologies for Slice Mobility: system virtualization with containers and SDN. These technologies allow a fast and seamless allo-
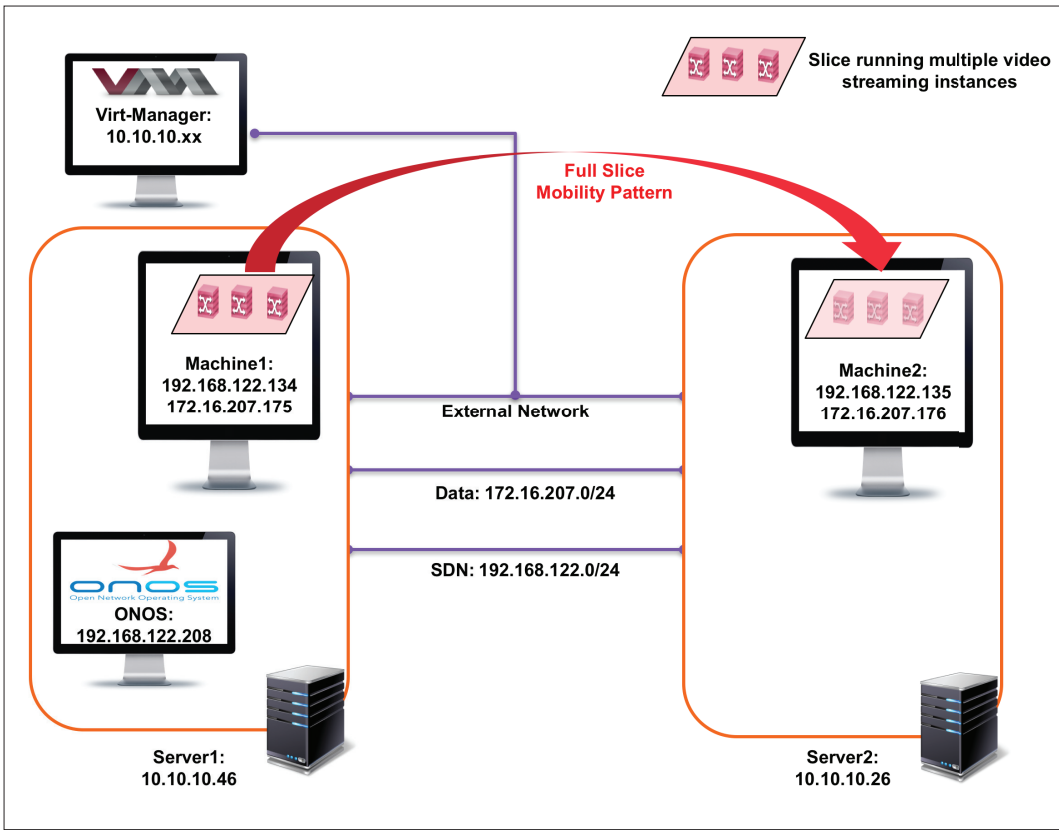
FIGURE 4. Testbed setup.

cation of resources to a network slice, while a complete evaluation of the behavior of the migration procedure under different physical setups and user mobility patterns are outside the scope of this work. In the performance evaluation, we consider parallel migrations to better understand the limits of the envisioned system. Our focus on the parallel migration arises from our perceived need to support multiple simultaneous migrations during a slice reconfiguration caused by user mobility across domains. Figure 4 portrays the testbed environment envisioned to emulate the slice mobility patterns, targeting the third use case discussed previously. In this use case, a slice consists of a set of streaming functions that are used for streaming video services. The testbed consists of two physical servers (i.e., two administrative IaaS domains) as depicted in Fig. 4.

Both servers run Ubuntu version server and have KVM as a hypervisor. Server 1 hosts two VMs; one acts as the source cloud for the slice mobility while the other VM acts as an SDN controller that manages the communications between the different servers and their clients. As the SDN controller, we used ONOS (Open Network Operating System). However, any other SDN controller could be used as well. The second physical server, server 2, hosts only one VM that plays the role of the destination cloud. Both servers and all VMs are running Ubuntu 16.04 LTS with the 4.4.0-64-generic kernel with four CPU cores and 4 GB of memory. The connection between the servers is 1Gbps. An additional host is used for accessing the testbed from an external network.

As previously explained, the container technology offers better overall performance compared

Triggers/slice mobility	Full slice mobility	Slice splitting	Slice merging	Slice breathing
Group mobility	Yes	Yes	Yes	Yes
Resource availability	Yes	Yes	Yes	Yes
Reliability	Yes	No	No	No
Security	No	Yes	Yes	No
Service overload	No	Yes	Yes	Yes
Service consumption	No	Yes	Yes	Yes

TABLE 1. Relationship between key enabler triggers and slice mobility patterns.

to VMs. The use of a system level container, i.e., Linux Container (LXC), technology allows us to support a wider range of applications with the proposed framework. In our environment setup, we therefore use the container technology LXC 2.8 and CRIU 2.6. We created three containers; each of them is running an NGINX server to stream videos to three different clients.

In the performance evaluation, we chose the full slice mobility pattern as it is the most complicated and computation-heavy scenario. The migration process triggers the migration of running instances in parallel from the source cloud to the destination cloud. To automate the orchestration process, we built a framework to handle the parallel container migrations. We considered the resource availability trigger to enable the parallel migrations. We used a mix between the grouping by service area method, and the grouping by geographical location method since our clients are
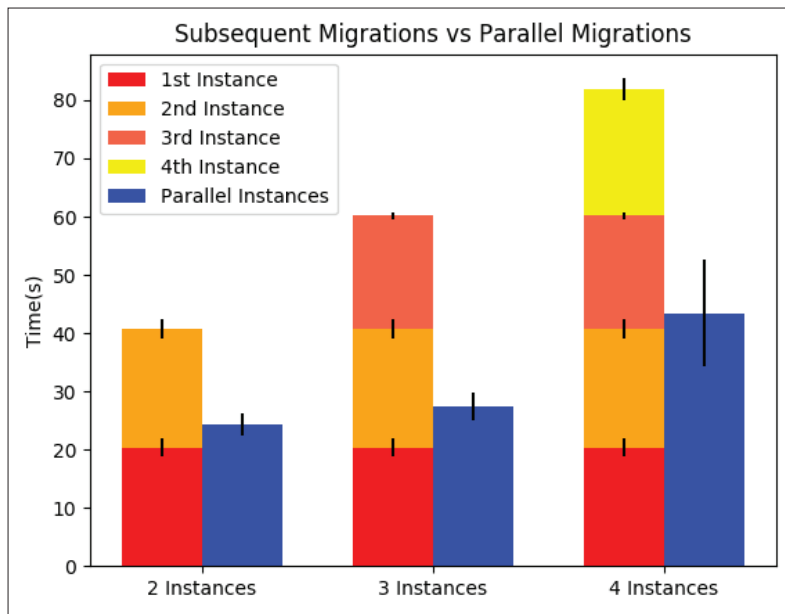
FIGURE 5. Subsequent migrations vs parallel migrations.

static and are served from their respective video streaming servers. By leveraging the SDN paradigm, we created an overlay networking on top of the physical network. Such an action allowed us to design an isolated network topology over multiple physical data centers, and at the same time, carrying out inter-datacenter (i.e., inter-IaaS) parallel migrations.

For enabling the full slice mobility, the test is performed using the iterative migration approach [16]. We adopt the pre-copy logic to perform the whole operations of live migration starting with the copy of the disk until the memory copy. Then, the SDN controller handles the path redirection phase and finally the container is restored to that target node.

Fig. 5 respectively shows the duration of two, three and four subsequent migrations vs the same number of parallel migrations. The test is performed with identical sizes of containers in order to carry out fair tests while comparing the parallel migrations to the subsequent one. In Fig. 5, the X-axis represents the number of LXC migrated containers while the Y-axis represents the migration time in seconds. For each bar (i.e., parallel and sequential migration), we plotted the 95 percent confidence interval of the mean.

The mean total migration time for two subsequent migrations was 40.7123s and the 95 percent confidence interval for this experiment was 2.0007s. For the parallel migrations, we obtained a mean total migration time of 24.3023s and 1.9764s as its 95 percent confidence interval. Regarding the migration of three instances, we achieved a mean total migration time of 60.1735s and 27.3984s, respectively, for both sequential and parallel migrations. Finally, on the subject of four instances migrations scenario, the mean total migration time is 81.9131s for the sequential migrations with 95 percent confidence interval of 2.6310s, while these values are 43.4268s and 9.0867s, respectively, for the parallel migrations. The results obtained through this evaluation reveal that the parallel migrations for enabling a prompt slice mobility are more efficient when compared to the legacy sequential migration strategy.

CONCLUSION AND REMARKS

The network slicing paradigm unquestionably offers a powerful apparatus to support verticals' services. According to this paradigm, a group of users is associated with dedicated computing, storage, and network resources tailored for a given vertical use case. The use of specialized resources in a slice implies that such resources are not available everywhere in the network but require careful resource allocation policies and control. Therefore, there is a strong need to extend the notion of mobility that is traditionally limited to user devices or services without much concern of combined resource availability needed for network slicing. With network slicing, the mobility events of users and services become more correlated leading to new mobility patterns. In this article, we introduced slice mobility patterns with corresponding grouping methods and relevant mobility triggers.

The proposed mobility patterns leverage MEC to offer ultra-short latency communication infrastructure, SDN for fast flow resumption, live migration to ensure the high-availability of services and the Follow me Edge concept to support end-users' mobility. The proposed framework is validated using a realistic testbed. Interesting results were obtained, implying total time migration evaluation depends on the used video streaming delivery content. The obtained results also demonstrate the dominance of the proposed parallel migrations when compared to the legacy sequential migrations strategy. Based on the achieved results and the presented key enabling triggers in addition to the grouping methods, it can be concluded that a mechanism to select the right combination of techniques to be used for efficiently supporting a slice mobility pattern action is indispensable. For the authors, a future research direction would be to investigate Artificial Intelligence techniques for smart and cost-efficient triggering and grouping methods for slice mobility patterns, as well as the migration of several coordinated and dependent services, known as service function chain (SFC).

REFERENCES

[1] N. Alliance, "5G white paper," Tech. Rep., Feb. 2015. [Online]. Available: https://www.ngmn.org/uploads/media/ NGMN 5G White Paper V1 0.pdf
[2] K. Ramantas et al., "A C-RAN Based 5G Platform with a Fully Virtualized, SDN Controlled Optical/Wireless Fronthaul," in *2018 20th International Conference on Transparent Optical Networks (ICTON)*, Bucharest, Romania, July 2018.
[3] J. Ordonez-Lucena et al., "Network Slicing for 5G with SDN/ NFV: Concepts, Architectures, and Challenges," IEEE Commun. Mag., vol. 55, no. 5, May 2017, pp. 80–87.
[4] I. Afolabi et al., "Network Slicing; Softwarization: A Survey on Principles, Enabling Technologies; Solutions," IEEE Commun. Surveys Tutorials, vol. PP, no. 99, 2018, pp. 1–1.
[5] T. Taleb et al., "PERMIT: Network Slicing for Personalized 5G Mobile Telecommunications," IEEE Commun. Mag., vol. 55, no. 5, May 2017, pp. 88–93.

[6] T. Taleb *et al.*, "On Multi-Access Edge Computing: A Survey of the Emerging 5G Network Edge Cloud Architecture and Orchestration," IEEE Commun. Surveys Tutorials, vol. 19, no. 3, Third Quarter 2017, pp. 1657–81.

[7] I. Farris *et al.*, "Lightweight Service Replication for Ultra-Short Latency Applications in Mobile Edge Networks," in *2017 IEEE International Conference on Communications (ICC)*, Paris, France, May 2017.

[8] A. Ksentini, T. Taleb, and M. Chen, "A Markov Decision Process-based Service Migration Procedure for Follow me Cloud," in *2014 IEEE International Conference on Communications (ICC)*, Sydney, NSW, Australia, June 2014.

[9] T. Taleb and A. Ksentini, "Follow me Cloud: Interworking Federated Clouds, and Distributed Mobile Networks," *IEEE Network*, vol. 27, no. 5, Sept. 2013, pp. 12–19.

[10] A. Aissioui *et al.*, "On Enabling 5G Automotive Systems Using Follow me Edge-Cloud Concept," *IEEE Trans. Vehicular Technology*, vol. 67, no. 6, June 2018, pp. 5302–16.

[11] F. Z. Yousaf *et al.*, "Network Slicing with Flexible Mobility and QoS/QoE Support for 5G Networks," in *2017 IEEE International Conference on Communications Workshops (ICC Workshops)*, Paris, France, May 2017.

[12] C. Clark *et al.*, "Live Migration of Virtual Machines," in *Proceedings of the 2Nd Conference on Symposium on Networked Systems Design & Implementation -Volume 2, ser. NSDI'05*. Berkeley, CA, USA: USENIX Association, 2005. [Online]. Available: http://dl.acm.org/citation.cfm?id=1251203.1251223

[13] V. Mann *et al.*, "CrossRoads: Seamless VM Mobility across Data Centers through Software Defined Networking," in *2012 IEEE Network Operations and Management Symposium*, Maui, HI, USA, April 2012.

[14] Y. C. Tay, K. Gaurav, and P. Karkun, "A Performance Comparison of Containers and Virtual Machines in Workload Migration Context," in *2017 IEEE 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*, Atlanta, GA, USA, June 2017.

[15] 3GPP, "System Architecture for the 5G System; Stage 2," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 23.501, 03 2018, version 15.1.0. [Online]. Available: https://portal.3gpp.org/desktopmodules/Specifications/ SpecificationDetails.aspx?specificationId=3144

[16] R. A. Addad *et al.*, "Towards a Fast Service Migration in 5G," in *2018 IEEE Conference on Standards for Communications and Networking (CSCN)*, Paris, France, Oct 2018.

Biographies

Rami Akrem Addad is currently pursuing the doctoral degree at Aalto University, Finland, in the Department of Communications and Networking, School of Electrical Engineering. He received the Licentiate degree and the Master degree from the University of Sciences and Technology Houari Boumediene, Algeria in 2015 and 2017 respectively, both degrees with high distinction and honors. He has been involved in the European Project ANASTACIA Horizon 2020 for addressing cyber-security concerns by leveraging SDN, NFV and Cloud architectures. His research interests include 5G network architecture, cloud-native technologies and approaches, network softwarization and slicing mechanisms, MEC, NFV, SDN, and distributed systems.

Tarik Taleb received the B.E. degree (with distinction) in information engineering in 2001, and the M.Sc. and Ph.D. degrees in information sciences from Tohoku University, Sendai, Japan, in 2003, and 2005, respectively. He is currently a professor with the School of Electrical Engineering, Aalto University, Espoo, Finland. He is the founder and Director of the MOSA!C Lab. He is the Guest Editor-in-Chief for the IEEE JSAC series on network Softwarization and enablers.

Hannu Flinck received the M.Sc. and Lic.Tech. degrees in computer science and communication systems from Aalto University (formerly, Helsinki University of Technology) in 1986 and 1993, respectively. He was with the Nokia Research Center and the Technology and Innovation Unit of Nokia Networks in various positions. He is a research manager with Nokia Bell Labs, Espoo, Finland. He has been actively participating in a number of EU research projects. His current research interests include mobile edge computing, SDN, and content delivery in mobile networks, particularly in 5G networks.

Miloud Baga received the bachelors, masters, and Ph.D. degrees from the University of Science and Technology Houari Boumediene Algiers, Algeria, in 2005, 2008, and 2014, respectively. He is currently a senior researcher with the Communications and Networking Department, Aalto University. His research interests include wireless sensor networks, the Internet of Things, 5G wireless communication, security, and networking modeling.

Diego L. C. Dutra is currently working as a professor at Federal University of Rio de Janeiro (UFRJ), Brazil, where he is also a member of the COMPASS Laboratory. He received a B.Sc. in computer science from UFF/Brazil, his M.Sc. and D.Sc. degrees in the Systems Engineering and Computer Science Program at the Federal University of Rio de Janeiro, Brazil, in 2007 and 2015, respectively. He has worked as a postdoctoral researcher in the COMPASS/UFRJ and MOSA!C Lab/Aalto, from 2015 to 2016 and 2016 to 2017, respectively. His research interests include computer architecture, high-performance computing, virtualization, cloud computing, wireless networking, and Software-Defined Systems.