Workspace-based Virtual Networks: A Clean Slate Approach to Slicing Cloud Networks

Romerson Deiny Oliveira[a], Diego Nunes Molinos[b], Marcelo Silva Freitas[c], Pedro Frosi Rosa[d] and Flavio de Oliveira Silva[e]
Faculty of Computing, Federal University of Uberlandia, Uberlandia, Brazil

Keywords: Cloud Computing, Edge Computing, Virtual Network, ETArch, Workspace, Software Defined Networking.

Abstract: Cloud Computing has brought a new vision about what applications might expect from underlying networks. Software Defined Networking, in this context, provides techniques to make networks more flexible to application requirements as well as to virtual networks, which help implement the provisioning of physical resources to support tenants. Cloud-hosted applications assume that networks have high throughput, but do not consider competition through the medium, handled by the switches. This paper aims to present the concept of Workspace as a new paradigm to manage the underlying resources in terms of meeting the communication requirements. We introduce a clean-slate approach based on horizontal addressing that has Workspaces as logical links capable of being parameterized at the level of medium access control. In addition, an overview of a new network element designed to allow the parameterization of Workspaces directly on the hardware is given.

1 INTRODUCTION

Cloud computing has become an expressive paradigm over the Internet in such a way that the number of applications developed and hosted in the Cloud is constantly increasing (Zhang et al., 2010). Virtualization in this respect enables the deployment of virtual networks as well as virtual services, and makes the details of the underlying (physical) network to applications transparent (Son and Buyya, 2018) (Blenk et al., 2015) (Jain and Paul, 2013).

Nevertheless, virtualized resource management is one of the challenges that Cloud environments still have to deal with. For instance, how to quickly deploy new services to customers and how to adjust bandwidth on demand. These problems are directly related to network infrastructures and not just to the virtualized distributed system. In this sense, the concept of programmable networks has been recovered to make network management more flexible.

As stated by (Jain and Paul, 2013) and (Son and Buyya, 2018), several researches explore the concept of Software Defined Networking (SDN) to provide Cloud infrastructure. Software Defined Networks separate control and forward planes as well as enable higher level software to configure the control plane on demand, allowing the infrastructure to quickly adapt to new application requirements (ONF, 2012).

Following the principles of SDN, Entity Title Architecture (ETArch), a clean slate network architecture, is based on a horizontal addressing scheme in which entities communicate with each other through a Workspace. The Workspace is a logical bus, a multicast domain, created according to the requirements of the application (de Oliveira Silva et al., 2012).

Regarding SDN implementations, this is achieved with a new layer of control between applications and infrastructure. This control layer receives instructions from applications and issues configuration actions for the equipment at the infrastructure layer (Son and Buyya, 2018). Nowadays, the OpenFlow protocol can be considered the 'de facto' standard for the interaction between the centralized logic controller and the forwarding devices, e.g. SDN switches.

Despite the flow separation allowed by the OpenFlow switches, there is still no link-level programming, i.e. all streams are performed following the same static medium access control.

[a] https://orcid.org/0000-0003-0083-469X
[b] https://orcid.org/0000-0002-7455-7624
[c] https://orcid.org/0000-0002-2313-1068
[d] https://orcid.org/0000-0001-8820-9113
[e] https://orcid.org/0000-0001-7051-7396

Oliveira, R., Molinos, D., Freitas, M., Rosa, P. and Silva, F.
Workspace-based Virtual Networks: A Clean Slate Approach to Slicing Cloud Networks.
DOI: 10.5220/0007753104640470
In *Proceedings of the 9th International Conference on Cloud Computing and Services Science (CLOSER 2019)*, pages 464-470
ISBN: 978-989-758-365-0

This paper aims to present the concept of Workspace as a new paradigm to manage the underlying communications in terms of meeting the communication requirements of applications such as QoS, mobility, security, and so on. Debates on cloud-hosted applications assume that networks are increasingly larger, but do not consider competition for the physical medium manipulated by switches. We introduce a clean-slate approach based on horizontal addressing that has Workspaces as logical links capable of being parameterized at the level of access control to the medium.

The remainder of the document is structured as follows: Section 2 presents the background and related works about fundamental concepts, highlighting the supporting technologies for Cloud networks and briefly presents the ETArch architecture. Section 3 provides an overview of the use of ETArch for slicing Cloud networks proposal. Section 4 shows the results of work in progress until now. Finally, Section 5 offers some concluding remarks and makes suggestions for future work.

2 BACKGROUND AND RELATED WORK

This section brings some noteworthy concepts about current and proposed cloud support.

Cloud computing represents the transformation in computing nowadays. The virtualization technique employed in cloud computing allows multiple servers to serve multiple applications for several users. Resources such as processing and storage have been greatly improved, but network infrastructure resources have made slow progress in recent years (Son and Buyya, 2018).

2.1 Cloud Virtual Networks Management

The virtualization technique has been widely explored in the field of cloud computing, offering distributed computing in which a pool of virtualized and dynamically scalable computing services are delivered on demand (Duan et al., 2012).

Cloud virtual networks introduce flexibility by decoupling service providers and allow multiple heterogeneous network from different service providers, sharing the same common underlying physical network resources (Chowdhury and Boutaba, 2010).

Virtual networks was conceived to 'slice' a given physical network infrastructure. Basically, it abstracts the underlying physical network and then creates detached virtual networks (slices). In the network field, there are several techniques that can create network slices, such as Wavelength Division Multiplexing (WDM), Virtual Local Area Networks (VLAN) and Multiple Protocol Label Switching (MPLS) (Blenk et al., 2015).

To enable virtualization of networks in cloud computing environments, it is necessary to virtualize nodes, links, switches, and all resources in the physical network infrastructure. The granularity level of the slice directly reflects the effectiveness of managing the network resources (Foukas et al., 2017) (Chowdhury and Boutaba, 2010).

In coarse granulation, each function is responsible for a large portion of the network operations (e.g., LTE EPC), whereas in the fine grained functions, each of the coarse grained functions is divided into several sub-functions. For the same example, LTE EPC can be broken into functions responsible for mobility and traffic forwarding (Foukas et al., 2017).

There are several initiatives that seek mechanisms to minimize the lack of flexibility and dynamism of virtual networks. Although usual, virtualization requires powerful abstraction capabilities to keep the operational complexity hidden (Amarasinghe and Karmouch, 2016).

Software Defined Networking breaks with the traditional distributed control of the real Internet architecture, by logically centering the control of an entire physical SDN network on a single logical SDN controller (Blenk et al., 2015). The programming capability offered by SDN can be used to implement virtual SDN, so an SDN network can be virtualized by inserting a hypervisor between the physical SDN network and the SDN control plane. The SDN hypervisor virtualizes the physical SDN network and creates isolated virtual SDN (vSDN) networks that are controlled by their respective SDN virtual controllers (Son and Buyya, 2018) (Blenk et al., 2015).

Network virtualization in SDN still lacks research and studies. There are obstacles that decrease the performance of vSDN and consequently in the quality of service delivered by the network. A visible problem is the abstraction of the hardware SDN switches, where the elements can present significant variations compromising the efficiency of the network. Recent studies present partial virtualization proposals for networks and physical switches (Kuzniar et al., 2014)(Lazaris et al., 2014)(Blenk et al., 2015).

2.2 Entity Title Architecture - ETArch

Among several new architectures projects (Hasegawa, 2013), the ETArch project emerged (de Oliveira Silva et al., 2012) idealizing a clean slate architecture to meet the requirements of future networks. ETArch is based on a horizontal title-based addressing scheme in which communicating entities request content by subscribing to them, which triggers a process of dynamic network configuration in order to provide content for requesting entities. Content is delivered through a channel, called Workspace, which aggregates various communication elements, by allowing the user's entities to express their requirements and capabilities over time.

ETArch's Title addressing eliminates the ambiguity between Identifier and Locator, which still exists in the IP addressing hierarchy, and replaces multiple addresses, namely those used in multiple layers of the TCP/IP stack, for a single address for the entity (de Souza Pereira et al., 2010).

ETArch does not have a fixed layer structure as in the TCP/IP architecture. In place of transport and network layers, ETArch defines the 'Communication' layer. As shown in Figure 1, this layer differs from traditional ones by being flexible and shaping according to application-specific communications requirements such as, for example, QoS, low power consumption and bandwidth.

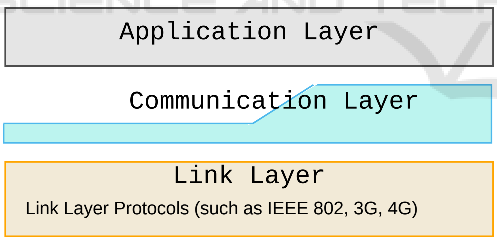

Figure 1: ETArch layer architecture pattern (de Oliveira Silva et al., 2012).

2.2.1 Entity and Title

An ETArch entity can be anything that can communicate through the network. Thus, it can be a host, a server, a switch, a user, a process, a Workspace, among others. An entity has at least one title, a set of requirements and capabilities. Examples of requirements are throughput, security and jitter.

Before start communicating, an entity expresses its needs by sending its requirements, in addition to its capabilities, to the control plane. Using such information, the network will be configured to meet the requirements to support communication.

The Title is an identifier, unambiguous and independent of the type and topology of the underlying network. An entity may have one or more titles, but a title identifies only one entity. The title is independent of the location of the entity, which facilitates the management of entity mobility (Guimarães et al., 2014).

2.2.2 Domain Title Service - DTS

DTS represents the control plane of an ETArch system. It is a distributed system composed of interconnected DTS Agents (DTSA), each one controlling a subset of the forwarding elements of the network.

The DTSA is a local domain. To allow inter-domain communication, ETArch has the Master DTSA (MDTSA), which is an agent that has a superset of DTSA functions, responsible for implementing the inter-domain communication interface.

2.2.3 Workspace

A Workspace can be defined as a logical bus, regardless of the topology and type of the underlying network, which interconnects multiple entities that need to communicate according to specific requirements. A Workspace is naturally multicast, meaning that one primitive transmitted by an entity reaches, without retransmissions, all entities participating in that Workspace.

Workspaces does not exist in the initial network configuration. If an entity wants to offer content, it needs to create a Workspace and register it in DTS through a DTSA. When another entity is willing for this content, it must ask DTS to register it and to attach it to the Workspace that offers that content.

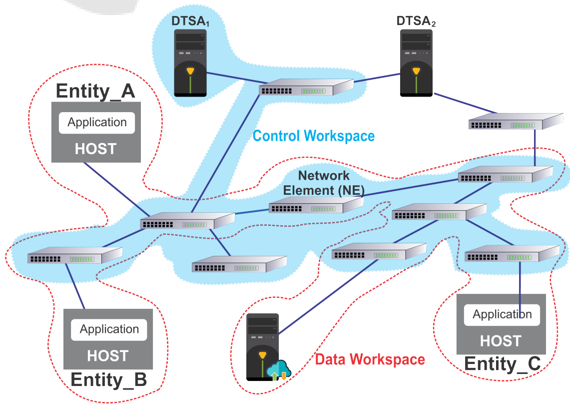

Figure 2: Ecosystem ETArch (de Oliveira Silva et al., 2012).

ETArch specifies two types of Workspaces, namely, Data and Control. The Figure 2 shows an example of Data Workspace of the ETArch architecture. In this figure, the DTS domain consists of two agents, DTSA1 and DTSA2, each controlling a subset of the Network Elements (NEs). To this end, each agent abstracts the topology of the NEs it controls.

Note that traffic, in the illustrated Data Workspace, is only duplicated in NEs closest to consumer entities. Problems related to discovering Workspaces and determining the best routes to expand their boundaries are discussed in (Souza Neto et al., 2015).

The Control Workspace is responsible for transmitting the control plane primitives. It is a common Workspace, however, its function is to create a logical bus for communication of DTSAs. Thus, the entities that compose it are the DTSAs of a domain (DTS) and the switches that interconnect them.

Control Workspaces are not dynamic, such as Data Workspaces. The latter are created and removed as required by network applications. The control workspaces are created in the network bootstrap and maintained during its operation. This type of Workspace is used for inter-domain routing, since communications between DTSAs is the reason for its existence (Souza Neto et al., 2015).

In order to implement the concept of Workspaces, two protocols were proposed for the control plane: ETCP (Entity Title Control Protocol) and DTSCP (Domain Title Service Control Protocol), the first one for communication between entities and DTSAs and the second for communication between two or more DTSAs. The services of each protocol, as well as their respective primitives can be found in (de Oliveira Silva et al., 2012).

3 ETARCH TO SLICING CLOUD NETWORKS

This section presents an approach that considers the use of Workspaces to implement future virtual networks.

3.1 Workspace-based Slices

Applying virtualization concepts to SDN networks offers numerous benefits. For example, slicing the physical network into multiple virtual networks provides the lessee with only the resources requested by the virtual network.

In turn, the ETArch architecture was designed based on the assumption that current network architectures, based on the "one size fits all" approach, are no longer adequate for future networks. That is, networks of the future must be able to meet the diverse demands of applications with diverse requirements.

As presented in Section 2.2, Workspace is a fundamental concept to communications in ETArch. In addition to its intrinsic multicast nature, it can also be emphasized that the concept of QoS are also intrinsic to its nature, just because a Workspace is deployed by DTSAs on network switches according to the requirements and capabilities of the entity that create the Workspace. Other entities that may be attached, must meet the Workspace requirements.

Note that a single virtual SDN network must support a variety of applications. But, since each application may have specific requirements, it is appropriate to think that one slice should have be associated to one application. Then, if the workspaces can be seen as slices, they can offer finer-grained resource management whether compared to virtual networks, such as SDN virtual slices.

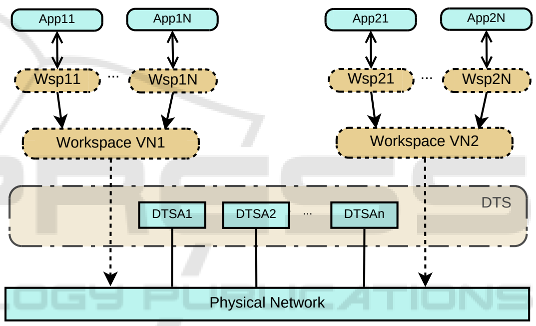

Figure 3: Workspace-based virtual networks.

The Figure 3 shows how a DTS could instantiate, through DTSAs, two Workspaces, namely (`Workspace VN1` and `Workspace VN2`), that would be two slices of the physical network. Then, they would create support for instantiating a Workspace by the application within each slice.

Considering the possibilities and benefits that the implementation of this idea could bring to the management of future network resources, this position paper defends the potential applicability of Workspaces to the provisioning and control of virtual networks, including future cloud environments. However, to make it real, a whole new network element and its interface must be designed.

3.2 Workspace Interfaces

A complete overview of the Workspace-based architecture is represented in figure 4. At the top, the DTSA and End User entities are using the Workspaces as a communication link. There are no

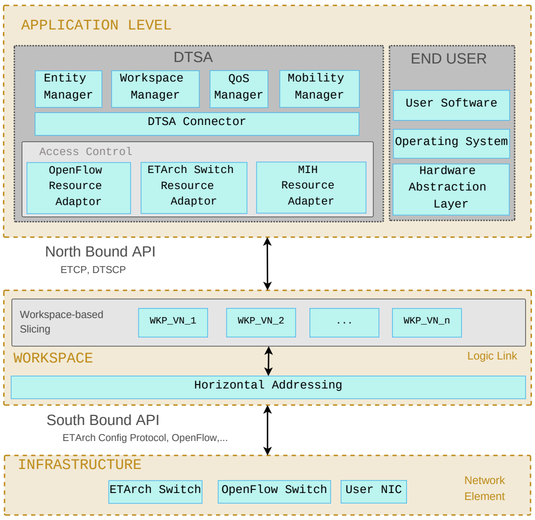

Figure 4: Workspace-based Architecture Overview.

assumptions about the physical topology of network elements for end users. They talk to each other following the rules of the ETArch protocols.

At the bottom, the network elements act on the forwarding plane. Among them, logical links are established on demand in accordance with DTSA guidelines.

The ETArch switches in this scenario, through the south-bound interface, bring a new hardware-implemented API to provide programmability downto the medium access control method, i.e., the embedded software defines hardware operations using reprogrammability - which is discussed in Section 4.

Especially at the south-bound, the interface between the Workspace and the network element is quite thin. Parameters such as traffic prioritization, flow rate, hardware encryption, and Workspace discovering are configured directly on flexible hardware.

According to (Kalyaev and Melnik, 2015), on some switches, OpenFlow support is considered as enhancement of existing features and adversely affects their cost. Virtually, the switches do not have full hardware support for newer versions of OpenFlow, since hardware is inherited from classic solutions. For economic reasons, manufacturers do not always announce the possibility of using OpenFlow in their decisions. In addition, with special attention to the lower layers of the network, OpenFlow switches do not allow any changes to link-level connections, especially the hardware/software interface.

4 RESULTS AND THE ROAD AHEAD

As stated earlier, programmable networks are useful for enabling the transfer of content in the cloud. Network flexibility depends on features such as programmability level, programmable communications abstractions, and architectural domain or application domain (e.g., signaling, management, transport). This impacts the construction of architectures as well as the services offered, considering that they are strongly related to the programming methodology, adopted middleware and potentially networked node support (Campbell et al., 1999).

One problem with this is that all implementations assume there is some connectivity at the network link level. Note that SDN implementations are most often implemented in OpenFlow-based switches. At this point, and intending to have greater control of the QoS parameters of the transmissions, the Workspace south-bound interface requires some flexible network element to support the programmability, even at the lower levels of the network.

The Link layer is the level at which hardware and software interface in the implementations. Therefore, having a new network element, such as a switch Ethernet and TCP/IP-independent, would be a very versatile solution for establishing distinct Workspaces with different hardware-level features.

In this scenario, we propose a programmable switch prototyped in reconfigurable hardware. The main idea is to offer a switch that allows different methods of medium access control, that differentiate the different communication needs, and that allows the parameterization of each Workspace being created. The notion of separate control and forwarding planes remains unchanged, but the programmability is also added to the forwarding plane as well as the possibility of change hardware configurations on the fly.

We provide an API-like (hardware-implemented) to DTSAs so they can combine parameters and have a subset of possible Workspace configurations strictly on hardware. Figure 5 shows how Workspaces are implemented. Within the ETARch switch, there is a set of Datapaths (and respective PHYs), a Control Unit and a Workspace Table. Each Datapath is connected to all others Datapaths to request and confirm frame retransmissions. In addition, the Control Unit configures each workspace with a subset of Datapath functions/cores, as represented in the colored planes of Figure 5.

Each physical datapath can be used according to the needs of each Workspace. Then, independent traffic can be found on the same switch port belonging to different workspaces. In addition, a Workspace may be flowing on different ports at the same time.

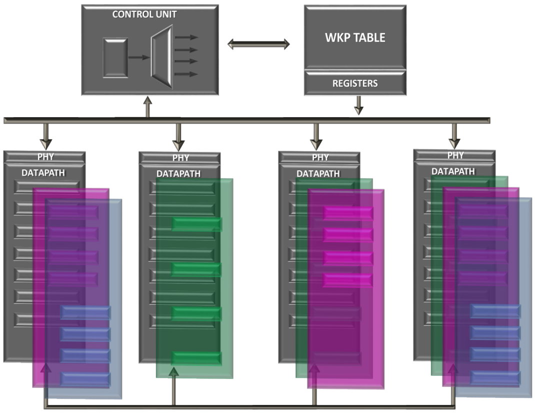

Figure 5: ETArch Switch - Virtual Datapaths.

The granularity of the Workspaces (looking from below) relies on the capabilities and parameters explained in following.

Workspace Prioritizing: traffic sorted according to several levels of Workspace priority on message queues.

Parallel Scheduling: a component of the scheduler that takes advantage of the spatial parallelism of the hardware to implement different threads with different scanning frequencies. This means that frames are forwarded with different flow rates according to Workspace requirements (configured by DTSA). For example, text, audio, and video may have different rates to be forwarded.

Non-linear throughput: thread frequencies can be modified on demand, and at line rate by the DTSA, i.e., the scheduler changes the rate in which each Workspace has its frames being forwarded to the corresponding output ports. This is useful mainly to control the QoS of critical Workspaces.

Workspace Manipulation: Workspaces can be added, edited, extended or deleted at line rate by responding DTSA requests (e.g., due to QoS agreements or security requirements).

Resource Partitioning: the switch itself can inform the DTSA about being (or not) able to accept new Workspaces based on its capacity and the current throughput being used (with direct impact on the QoS).

Multicast Inherent: multicast created at the link level by streams added in a workspace table addressed by titles.

As stated by (Kalyaev and Melnik, 2015), a key challenge for SDN is how to handle high-touch, high-security, and high-performance packet processing streams. There are two essential elements to consider: performance and flexibility. To handle this trade-off, the FPGA is a midterm between GPP and ASIC and it can achieve custom data path processing of over 200 Gb/s per device and still keep the flexibility of a reconfigurable hardware.

Taking this into account and bearing in mind that the cost of Non Recurrent Engineering (NRE) for a FPGA prototyping is lower than for an ASIC, an FPGA was chosen to be the prototyping platform. Up to now, we have the switch prototyped (on Altera's FPGA Cyclone II) and able to add, create, edit and remove Workspaces according to DTSA requests. Performance and stress tests are still being performed. It is possible to realize that the DTS environment has the closest control of the QoS parameter on each link in the network, controlling the scanning frequency of the scheduler and configuring each Workspace parameter on the south-bound interface.

The work in progress at this time is related to managing security in the environment and virtualizing different types of network architectures in terminals, according to the needs of the entities and the aspects of the resident operating system.

5 CONCLUDING REMARKS

This work proposes a network architecture to support the communication requirements of cloud computing, which takes advantage of the clean-slate architecture approach. Network programming capability, scalability, and on-demand provisioning can be enhanced using the Workspace concept.

Implementing Workspaces to support Cloud environment allows more granular management of the infrastructure. Using Workspaces to slice networks means allocate one logical bus per application offering some content rather than using slices, where different applications can share the slice with different requirements.

Nevertheless, implementing a clean-slate approach is something challenging, because it requires changes at the endpoints. These changes may have an impact on the network interface cards as well as in the operating system modules. This is not trivial to overcome.

Some SDN premises are essential for cloud traffic on networks. The ETArch approach, with the new switch design, addresses the separation of hardware and software while maintaining programmability in both. In addition, it provides the availability of programmable open interfaces of the lowest level of the network, the virtualization of the network infrastructure using Workspaces as logical buses (and implemented as virtual data paths), resource partitioning and coexistence of independent network or architectures over the same physical network hardware.

The FPGA-based switch allows to coexist in a single network element, traffic with specific communication requirements. This is useful for distributing content across the network by categorizing it by different levels of QoS, such as real-time video transfers, Web audio conferencing, and document downloads.

ACKNOWLEDGEMENTS

This project was built with the support of MEHAR team. We want to thank all who contributed to our research. Brazilian agency CAPES has partially funded this work.

REFERENCES

Amarasinghe, H. and Karmouch, A. (2016). Sdn-based framework for infrastructure as a service clouds. In *Cloud Computing (CLOUD), 2016 IEEE 9th International Conference on*, pages 782–789. IEEE.

Blenk, A., Basta, A., Reisslein, M., and Kellerer, W. (2015). Survey on network virtualization hypervisors for software defined networking. *arXiv preprint arXiv:1506.07275*.

Campbell, A. T., De Meer, H. G., Kounavis, M. E., Miki, K., Vicente, J. B., and Villela, D. (1999). A survey of programmable networks. *SIGCOMM Comput. Commun. Rev.*, 29(2):7–23.

Chowdhury, N. M. K. and Boutaba, R. (2010). A survey of network virtualization. *Computer Networks*, 54(5):862–876.

de Oliveira Silva, F., Gonçalves, M. A., de Souza Pereira, J. H., Pasquini, R., Rosa, P. F., and Kofuji, S. T. (2012). On the analysis of multicast traffic over the entity title architecture. In *Networks (ICON), 2012 18th IEEE International Conference on*, pages 30–35. IEEE.

de Souza Pereira, J. H., Kofuji, S. T., and Rosa, P. F. (2010). Horizontal addressing by title in a next generation internet. In *Networking and Services (ICNS), 2010 Sixth International Conference on*, pages 7–11. IEEE.

Duan, Q., Yan, Y., and Vasilakos, A. V. (2012). A survey on service-oriented network virtualization toward convergence of networking and cloud computing. *IEEE Transactions on Network and Service Management*, 9(4):373–392.

Foukas, X., Patounas, G., Elmokashfi, A., and Marina, M. K. (2017). Network slicing in 5g: Survey and challenges. *IEEE Communications Magazine*, 55(5):94–100.

Guimarães, C., Corujo, D., Silva, F., Frosi, P., Neto, A., and Aguiar, R. L. (2014). IEEE 802.21-enabled entity title architecture for handover optimization. In *Wireless Communications and Networking Conference (WCNC), 2014 IEEE*, pages 2671–2676. IEEE.

Hasegawa, T. (2013). A survey of the research on future Internet and network architectures. *IEICE transactions on communications*, 96(6):1385–1401.

Jain, R. and Paul, S. (2013). Network virtualization and software defined networking for cloud computing: a survey. *IEEE Communications Magazine*, 51(11):24–31.

Kalyaev, A. and Melnik, E. (2015). Fpga-based approach for organization of sdn switch. In *2015 9th International Conference on Application of Information and Communication Technologies (AICT)*, pages 363–366.

Kuzniar, M., Peresini, P., and Kostic, D. (2014). What you need to know about sdn control and data planes. *EPFL, Lausanne, Switzerland, Tech. Rep. EPFL-REPORT-199497*.

Lazaris, A., Tahara, D., Huang, X., Li, E., Voellmy, A., Yang, Y. R., and Yu, M. (2014). Tango: Simplifying sdn control with automatic switch property inference, abstraction, and optimization. In *Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies*, pages 199–212. ACM.

ONF (2012). Software-defined networking: The new norm for networks. *ONF White Paper*. https://www.opennetworking.org/images/stories/downloads/sdn-resources/white-papers/wp-sdn-newnorm.pdf.

Son, J. and Buyya, R. (2018). A taxonomy of software-defined networking (sdn)-enabled cloud computing. *ACM Computing Surveys (CSUR)*, 51(3):59.

Souza Neto, N. V. d., de Oliveira Silva, F., Rosa, P. F., and de Souza Pereira, J. H. (2015). Control Plane Routing Protocol for the Entity Title Architecture: Design and Specification. In *ICN 2015*, pages 185–190, Barcelona, Spain.

Zhang, Q., Cheng, L., and Boutaba, R. (2010). Cloud computing: state-of-the-art and research challenges. *Journal of internet services and applications*, 1(1):7–18.