2020 European Control Conference (ECC)
May 12-15, 2020. Saint Petersburg, Russia

From IoT to Cloud: An End-to-End Virtualization Approach

Aris Leivadeas

Abstract— **The integration of the Internet of Things (IoT) with the Cloud and Edge Computing infrastructures offers unprecedented opportunities for applications and services. Realizing such a holistic and large-scale platform requires a mixture of different technologies to enable the interconnection of the various functional layers of this new end-to-end (E2E) communication paradigm. Virtualization and specifically Network Function Virtualization (NFV), is the key enabling technology for creating and supporting an IoT-aware Edge-Cloud Computing infrastructure. In this tutorial session paper we study how to establish a scalable softwarized E2E solution that will gracefully combine different technologies to provide networking services with high Quality of Service (QoS) and low communication cost.**

I. Introduction

During the last decades, the Internet has been a key factor in building the global information society and leading economy growth. The mix of technologies currently offered in the Internet, such as Cloud/Edge Computing, Network Function Virtualization (NFV), Data Analytics, and Internet of Things (IoT), provide the necessary tools to build a service-centric Future Internet Architecture. Future Internet is expected to rely on virtualization technologies, moving away from dedicated and expensive hardware middleboxes, while IoT will be one of the main sources of data generation in this communication paradigm.

However, the main bottleneck in such an environment is the performance degradation observed when moving from dedicated hardware middleboxes to softwarized solutions. General purpose servers must now accommodate a wide range of Operating Systems (OS) and software packages in order to implement the large range of network services available. At the same time, Internet's resounding success has also turned out to be the biggest obstacle leading to in inefficient utilization of the underlying infrastructure. The excessive amount of the generated traffic from the IoT layer and the need for novel and more targeted networking services has led to a depletion of the available physical resources.

Furthermore, the management and orchestration of such a large-scale platform necessitates efficient mechanisms for controlling the infrastructure that spans from the physical objects / "things" till the Cloud. A tool to facilitate the design of such controllers can be found in the principles of Data Analytics that can predict the network traffic, its resource requirements, and appropriately guide the control decisions.

Thus, the goal of this tutorial session paper is to establish the basis of a scalable softwarized end-to-end solution that will gracefully combine all the aforementioned technologies to provide scalable E2E solutions with high QoS and low communication cost.

To succeed that, first we present the optimal software and hardware configuration of the underlying generic hardware infrastructure (e.g. typical x86 servers, edge nodes, etc.) to optimize the attained performance in terms of throughput and delay providing a cost efficient and attractive solution. Secondly, taking into consideration the limited resources available, especially at the network edge, appropriate control decisions in forms of resource allocation algorithms are introduced. The goal of these algorithms is to minimize the deployment cost (e.g. in terms of operational and capital costs) of the services offered and further ameliorate the throughput and delay performance. Finally, data analytics mechanisms are investigated to remove redundant data, unravel traffic patterns, and predict the network conditions to better orchestrate the network design and resource planning of the E2E platform.

The rest of the paper is organized as follows. Section II presents a brief background of the academic efforts to resolve similar problems. Section III emphasizes on the performance analysis of Cloud and Edge based NFV solutions. Section IV presents possible resource allocation methods for resource orchestration and management in such E2E platforms. Section V, provides the basis and the importance of data analytics in IoT-aware Edge/Cloud communication services. Section VI provides the challenges of controlling these services. Finally, Section VII concludes the paper.

II. Background

A. Network Function Virtualization

With NFV, network functions are now decoupled from dedicated and expensive hardware appliances and they are offered as software entities called Virtualized Network Functions (VNFs). VNFs, such as routers, firewalls, or intrusion detection systems are instantiated on generic servers in the Cloud or Edge leading to a significant capital and cost reduction. Despite the numerous advantages of the NFV technology, two of the main challenges that needs to be addressed are [1]: a) the performance degradation, when using generic servers instead of dedicated hardware and b) the optimal placement of the VNFs to efficiently utilize the cloud/edge resources.

To address the first challenge, a number of hardware and software solutions have been proposed to alleviate the performance deterioration caused by virtualizing the network services. Hardware-based solution usually involves FPGA accelerators to increase the performance of the VNFs [2]. However, such an approach involves additional cost to buy, deploy, and operate the hardware. In contrast, software-based solutions can significantly increase the performance without

Aris Leivadeas is with the Department of Software and Information Technology Engineering, École technologie Supérieure, University of Quebec, Montreal, Canada (e-mail: aris.leivadeas@etsmtl.ca). Aris is supported by the CHIST-ERA 2018 project DRUID-NET "Edge Computing Resource Allocation for Dynamic Networks".

the need to invest more monetary and installation cost [3]. Softwarized-based solutions usually intervene in the hypervisor by proposing appropriate tuning techniques [4], [5]. However, such approaches are difficult to operationalize in a production environment.

Regarding the second challenge, particular emphasis has been placed in allocating a series of interconnected VNFs, called Service Chain (SC), in a pool of available physical resources. Traditionally, the problem was solved assuming the existence of only a Cloud infrastructure [6] [7]. However, recently a lot of attention has been drawn in utilizing the Edge computing infrastructure in conjunction with IoT to benefit from the reduced latency introduced and the high available bandwidth [8] [9].

B. *Data Analytics and Control*

Another significant aspect in an IoT-aware E2E platform is the fact that the IoT data can be very redundant. That is the same data can be repeated without change and without surplus value [10]. IoT data can be proved catastrophic in the subsequent VNF processing, since it can lead to a wastage of resources, while it can also slow down the analysis itself. This can be very restrictive for mission-critical IoT applications that require ultra-low end-to-end communication latency. Thus, appropriate techniques, have been proposed to deal with IoT data redundancy, such as redundancy detection [11], data filtering [12], and data compression [13] among others.

Nonetheless, a missing part of the data analytics and resource allocation solutions proposed in the literature is where and how these mechanisms should be applied. Evidently, this leads to the discussion of the need of controllers and control techniques regarding how, where, and when IoT traffic should be analyzed and scheduled. In such a holistic approach, a combination of different controllers with different goals are required. Possible controllers involve SDN controllers, Cloud controllers, NFV Management and Orchestration Controllers, IoT access controllers, Edge controllers and so on [6] [14].

Hence, in this tutorial paper, we aim to present a holistic approach that takes into consideration the total challenges involved in each technology. The ultimate goal, is to shed light into a very important but still vague research endeavor as the E2E IoT-aware Edge/Cloud platform.

III. Network function virtualization performance

The major advantage of NFV is that it can promise considerable operational and capital cost reductions. Total cost of ownership calculations however, are typically a function of the attainable network throughput, which in a generic server is highly dependent on the underlying server architecture. The number of VNFs running on the server, the Input/Output (I/O) demands, the performance of the underlying hypervisor scheduler, and the packet path from the physical Network Interface (pNIC) of the server into the VNFs are examples of how the server architecture can influence the overall performance. In this tutorial, we will present our results indicating that consolidating many VNFs in the same server can severely deteriorate the throughput performance [15]. Specifically, we show that the VNFs in contrast with traditional cloud VM applications are I/O bound. Thus, the number of associated resources is not always the bottleneck. This relies from the hypervisor and virtual switch inside a generic server that were not designed with the goal to control and schedule traffic intensive applications. Hence, this bottleneck urged in using more efficient I/O techniques such as OVS, OVS-DPDK, FD.io VPP, and SR-IOV [3] [16]. In particular, a hypervisor bypass I/O technique can lead to a significant throughput performance able to saturate the total physical capacity of the NIC. This outcome is of particular importance for IoT and specifically for massive machine connections that can generate huge amount of traffic that needs fast and reliable processing.

Nonetheless, since not all of the I/O architectures can bypass the hypervisor, it is of utmost importance to propose appropriate scheduling algorithms that will efficiently distribute the physical resources of a server to the hosted VNFs. Thus, it becomes evident, especially for handling real-time IoT applications, that the hypevisor should support a real-time scheduler that will guarantee the horizontal scheduling of multiple concurrent hosted VNFs on the same server. This can only be guaranteed by I/O performance optimization techniques, resource scheduling, and a local control co-design that will interpret the resource requirements of the VNFs into control decisions of how the physical resources will be shared among the VNFs.

IV. IoT-Aware edge/cloud Resource allocation

IoT devices are in general constrained in terms of resources (e.g. CPU, memory, and storage) and energy (power is supplied by batteries). Hence, to support the requirements of the IoT applications in terms of processing and networking, Cloud Computing has risen as the leading paradigm. Cloud can offer an abundant of computing and communication resources that can scale up and down with varying user demands. This makes it very suitable for IoT applications due to their dynamic and ad-hoc nature. However, Cloud resources are usually accessible through the Internet, which is characterized by a best-effort connection with varying quality. To this end, application tasks and networking functions that require stringent resource guarantees such as real-time processing, local storage, or local context-awareness cannot be always satisfied using traditional Cloud Services [17]. Edge Computing can fill this gap by bringing Cloud Computing capabilities closer to the “things”.

One major challenge of this hybrid Edge/Cloud infrastructure is how to efficiently share the available resources in the Edge and Cloud to satisfy the requirements of the offloaded network services and application tasks (instantiated in the form of VNFs). These tasks and services are expressed as SCs that are associated with specific IoT applications. The major consideration with theses SCs is how to allocate them by minimizing the communication delays and the deployment costs [18] [19]. However, such a resource allocation technique when integrating the Edge with the Cloud is not straightforward. The reason is that we need to find the appropriate edge or cloud sites to host the VNFs/tasks and find the routing paths to interconnect them. The creation of a routing path that traverses multiple different domains from things, to gateways, to backhaul networks, edge networks, transport networks, and data center networks is far from easy. Especially, when we have to consider a set of different and possible conflicting goals.

Thus, in this part of the tutorial, we will demonstrate how to efficiently utilize the resources between the Edge and Cloud, in order to reconcile the conflicting goals of latency performance and deployment cost. Minimizing the deployment cost essentially amounts to minimize the amount of used servers. However, this approach can lead to an over-utilization of the server's resources that can have a serious effect on the overall network delay. Furthermore, there are additional goals to improve the resource allocation technique followed such as load balancing within the Edge or Cloud infrastructures and load balancing between the Edge and Cloud infrastructures. Hence, in the particular tutorial we will present a novel multi-objective resource allocation approach satisfying different goals and requirements generated from the IoT devices, Edge, and Cloud infrastructures in order to provide an end-to-end IoT aware Edge/Cloud solution [20].

At the same time, we will show the importance of including the performance analysis of the VNFs during the resource allocation decisions. In particular, the VNF performance analysis should be used as a feedback to the resource allocation, since both problems are closely correlated and can greatly affect the resulted performance. This knowledge is the main input element of the controller part of the hypervisor in order to dynamically and in real-time schedule the available resources at the hosts VNFs.

V. DATA ANALYTICS IN IOT

IoT analytics is a general term that describes how to process data generated from the IoT devices in order to produce information that makes sense. IoT analytics can be classified in three different categories. The first category is based on the fact that the IoT data can be very redundant. Hence, appropriate data cleansing and filtering techniques can be applied. The second category is the most wildly explored and is based on how data analysis can be applied by the IoT applications to provide an intelligent view related to the IoT system. For example, how collected information from thousands temperature, pressure, wind, and humidify sensors can determine the likelihood of a storm and its possible paths [21]. The third category is based on the network analytics. This is the least explored area, but nonetheless the most interesting aspect when offloading application's tasks and networking services as VNFs at the Edge and the Cloud.

During IoT network analytics the focus is to analyze data to determine loss or degradation in connectivity. This is particularly important for two reasons: i) first IoT access network technologies are lossy and unreliable with direct impact in the overall E2E performance, ii) without this kind of analysis there will not be any data and thus no other category of data analytics.

The network analytics mechanisms need a number of inputs in order to predict the access network conditions. These inputs are needed from the very beginning of the design of an IoT system and can be gathered by answering the following questions. How heterogeneous are the IoT devices? Are they battery-powered? With what frequency they send data? How rich are the data? Are the devices moving? What is the distance of the devices from the gateway? Which access technology is used?

The answers to these questions will help train the analysis but will not solve the ambiguity of the communication. The reason is that access networks are usually wireless, lossy and unreliable. Thus, the purpose of the IoT network analytic algorithms should not only be the traffic analysis but also traffic prediction in order to make the communication as deterministic as possible.

Obviously, the output of the algorithm, will be the driving force for designing the resource allocation algorithms. Based on the prediction of the traffic gathered at the gateways of the access network, we will be able to plan and schedule the necessary amount of resources at the Edge and the Cloud. Furthermore, the network analysis will dictate when the application tasks and VNFs should be applied locally in the IoT devices and when remotely at the Edge and/or Cloud.

VI. CONTROL CO-DESIGN IN AN IOT E2E PLATFORM

The combination of the VNF performance analysis, resource allocation algorithms, and network analytics can create the basis for a viable E2E IoT platform. To understand how these three ostensibly different theoretical approaches can be combined, we need to answer the question of who is responsible for executing them?

A possible answer is to have a centralized controller that will have a global view of the E2E system. However, such an approach can introduce tremendous delays and scalability issues. Thus, we need to have a hierarchical deployment of controllers that jointly manage the different layers of the network.

We identify three main controllers in this hierarchical deployment. In the level of VNF performance, the server's hypervisor acts as horizontal controller that is responsible to instantiate, monitor, and guarantee the performance of the VNFs according to the decision of the resource allocation algorithm. For the resource allocation algorithm, we need a vertical controller that will have an overall view of the whole physical infrastructure. Furthermore, considering an Edge/Cloud interplay, different controllers should exist at each infrastructure. Finally, the network analysis can be applied by a controller at the gateway level, which is responsible to take the decision of when to interconnect the "things" with the rest of the system.

There is a clear interdependency of the three types of controllers and a necessity for a coordination between them. Usually, at the Edge and Cloud the connection between the vertical and horizontal controllers are well-established (e.g. OpenStack, CloudStack and KVM/VMware hypervisors). However, the real challenge is how to interconnect the vertical controllers of the Edge and Cloud and both vertical controllers with the gateway controller. During the tutorial we will present the possible solutions and open challenges towards establishing this control communication and co-design.

VII. CONCLUSION

IoT is a recent technology gaining considerable momentum with hundreds of billion devices expecting to be connected to the Internet in the following decade. This extraordinary amount of IoT traffic will generate new network service needs, while requiring stringent Service Level Agreements. Thus the integration of all recent technologies, such as NFV, Cloud and Edge computing, and data analytics is deemed necessary in order to sustain an E2E IoT-aware communication platform. Hence, this tutorial paper provides a network paradigm where the main Future Internet technologies are intermixed to provide an end-to-end agile and robust architecture providing low-cost but high efficient networking solutions.

REFERENCES

[1] J. Halpern and C. Pignataro, "Service Function Chaining (SFC) Architecture", Tech. Report, IETF, Oct. 2015.

[2] C. Kachris, G. Sirakoulis, and D. Soudris, "Network Functions Virtualization based on FPGAs: A Framework for all-Programmable Network Devices", Technical report, Democritus University of Thrace, Greece, Jun. 2014.

[3] N. Pitaev, M. Falkner, A. Leivadeas, I. Lambadaris, "Multi-VNF Performance Characterization for Virtualized Network Functions", IEEE Int. Conf. on Network Softwarization (NetSoft), pp. 1-5, Jul. 2017.

[4] J. Martins, M. Ahmed, C. Raiciu, V. Olteanu, M. Honda, R. Bifulco, F. Huici, "Click-OS and the Art of Network Function Virtualization", In the 11th USENIX Symposium on Networked Systems Design and Implementation, pp. 459-473, Apr. 2014.

[5] L. Nobach, D. Hausheer, "Open, Elastic Provisioning of Hardware Acceleration in NFV Environments", IEEE Int. Conf. on Networked Systems (NetSYs), pp. 1-5, Mar. 2015.

[6] A. Leivadeas, M. Falkner, I. Lambadaris, and G. Kesidis, "Optimal Virtualized Network Function Allocation for an SDN enabled Cloud", Elsevier J. on Computer Standards & Interfaces, vol. 54, no. 4, pp. 266-278, Jan. 2017.

[7] M. Mechtri, C. Ghribi, and D. Zeghlace, "Scalable Algorithm for the Placement of Service Function Chains", IEEE Trans. On Network and Service Management, vol. 13, no. 3, pp. 533-546, Aug. 2016.

[8] W. He, S. Guo, Y. Liang, and X. Qiu, "Markov approximation method for optimal service orchestration in IoT network", IEEE Access, vol. 7, pp. 49538-49548, Apr. 2019.

[9] R. Munoz, R. Vilalta, N. Yoshikane, R. Casellas, R. Martinez, T. Tsuritani, and I. Morita, "Integration of IoT, Transport SDN and Edge/Cloud Computing for Dynamic Distribution of IoT Analytics and Efficient Use of Network Resources", IEEE J. on Lightwave Technology, vol. 36, no. 7, pp. 1420-1428, Jan. 2018.

[10] K. Papadokostaki, G. Mastorakis, S. Panagiotakis, C. X. Mavromoustakakis, C. Dobre, and J. M. Batalla, "Handling Big Data in the Era of Internet of Things", Springer book series on Advances in Mobile Cloud Computing and Big Data in the 5G Era, vol. 22, pp. 3-22, 2017.

[11] T. Yu, X. Wang, and A. Shami, "Recursive Principal Component Analysis-Based Data Outlier Detection and Sensor Data Aggregation in IoT Systems", IEEE Int. of Things, vol. 4, no. 6, pp. 2207-2216, Dec. 2017.

[12] D.-Y Kim, Y.-S. Jeong, and S. Kim' "Data-Filtering System to Avoid Total Data Distortion in IoT Networking", MDPI Symmetry, vol. 9, no. 1, pp. 1-13, Jan. 2017.

[13] B. R. Stojkoska, and Z. Nikolovski, "Data Compression for Energy Efficient IoT Solutions", IEEE Int. Telecom. Forum (TELFOR), pp. 1-4, Nov. 2017

[14] R. Munoz, R. Vilalta, R. Casellas, A. Mayoral, and R. Martinez, "Integrating Optical Transport Network Testbeds and Platforms to Enable End-to-End 5G and IoT services", IEEE Int. Conf. on Optical Networks (ICTON), pp. 1-4, Sep. 2017.

[15] M. Falkner, A. Leivadeas, I. Lambadaris, and G. Kesidis, "Performance Analysis of Virtualized Network Functions on Virtualized System Architectures", IEEE Computer Aided Modeling and Design of Communication Links and Networks (CAMAD), pp. 71-76, Oct. 2016.

[16] N. Pitaev, M. Falkner, A. Leivadeas, and I. Lambadaris, "Characterizing the Performance of Concurrent Virtualized Network Functions with OVS-DPDK, FD.io VPP and SR-IOV", ACM Int. Conf. on Performance Engineering (ICPE), pp. 282-292, Apr. 2018.

[17] R. Cziva, C. Anagnostopoulos, and D. P. Pezaros, "Dynamic Latency-Optimal VNF Placement at the Network Edge", IEEE Int. Conf. on Computer Communications (INFOCOM), pp. 693-701, Apr. 2018.

[18] Y. Nam, S. Song, and J.-M Chung, "Clustered NFV Service Chaining Optimization in Mobile Clouds", IEEE Comm. Letters, vol. 50, no. 1, pp.30-39, Feb. 2017.

[19] L. Zanzi, F. Giust, and V" Sciancalepore, "M2EC: A Multi-tenant Resource Orchestration in Multi-Access Edge Computing Systems", IEEE Int. Conf. on Wireless Communications (WCNC), pp.1-6, Apr. 2018.

[20] A. Leivadeas, G. Kesidis, M. Ibnkahla, and I. Lambadaris "VNF Placement Optimization at the Edge and Cloud", MDPI J. on Future Internet.vol. 11, no. 3, Mar. 2019.

[21] D. Hanes, G. Salgueiro, P. Grossetet, R. Barton, and J. Henry, "IoT Fundamentals: Network Technologies, Protocols, and Use Cases for the Internet of Things", ISBN: 978-1-58714-456-1, Cisco Press, 2017.