Latency-aware Optimization of Service Chain Allocation with joint VNF instantiation and SDN metro network control

F.J. Moreno-Muro[1], C. San-Nicolás-Martínez[1], M. Garrich[1], P. Pavon-Marino[1,2], O. González de Dios[3], R. Lopez Da Silva [3]

[1] Universidad Politécnica de Cartagena, 30202 Cartagena, Spain, javier.moreno@upct.es
[2] E-lighthouse Networks Solutions, 30203 Cartagena, Spain
[3] Telefonica GCTO, Ronda de la Comunicación, Madrid, Spain

Abstract *This demonstration showcases the open-source Net2Plan tool optimally computing and provisioning Service Chain requests with latency considerations by interfacing to an ONOS-controlled metro network and an ETSI-OSM instance orchestrating a set of OpenStack VIMs.*

1. Introduction

The dramatic forecasted growth of the IP traffic and the incoming 5G era endangers traditional approaches of network control and management. Indeed, Cisco predictions announce an annual increasing of 27% in the IP traffic mainly due to the popularity of the high-bandwidth video content [1], and metro network traffic increased twofold in comparison with backbone traffic in 2017 [2]. Additionally, key performance indicators for 5G networks enforce, among other needs, (i) a scalable management framework permitting agile deployment of applications, (ii) a reduction of the network management OPEX by at least 20% and (iii) the accomplishment of the End-to-End latency less than 1 ms for some of the traffic types. This picture critically challenges optical networking in the metro.

Software-defined networking (SDN) and network function virtualization (NFV) are key technologies for the aforementioned challenging scenario. SDN enables fully-programmable network management decoupling forwarding actions from control decisions whilst NFV provides flexible placement of network functions relying on commodity hardware appliances. SDN and NFV are complementary tools to assess a wide variety of user traffic requests. In this context, the Metro-Haul Project [3] proposes scalable, dynamic and efficient metro networks to efficiently interface 5G access and high-capacity core/backbone networks. The core of the Metro Haul proposal is the dynamic interconnectivity of two different types of nodes, Access Metro Edge Node (AMEN) and Metro Core Edge Node (MCEN) with computational capabilities to permit SDN/NFV functionalities to provision Virtual Network Functions (VNFs) close to the end users. Recent works show that locating content near to end users in an optimal way may partially alleviate the impact in core/backbone networks [4].

In our previous work [5], we demonstrated the specialized open-source planning tool Net2Plan [6] assisting an NFV-Orchestrator Open-Source MANO (OSM) [7] instance in the optimal VNF instantiation and Service Chain allocation directly in an emulated transport network. Two major limitations were present in [5]: the lack of a SDN controller to manage the network resources and the simplified data-plane emulation. The present work expands [5] by (i) including an ONOS [8] instance that controls the optical transport data plane, that is now emulated via Mininet, and controlled through the OpenFlow protocol and (ii) presenting an SDN/NFV allocation algorithm where flow allocation is optimized considering end-to-end latency requirements [5].

2. Demonstration Architecture Overview

This demonstration proofs the benefits of Net2Plan's interconnection with an ONOS-controlled emulated data plane to assist OSM's orchestration of OpenStack Virtual Infrastructure Managers (VIMs) so that latency considerations provide optimal VNF computation, service chain (SC) allocation and deployment in physical hardware appliances distributed in a metropolitan network. An architectural schema of the proposed demonstration is shown in Fig. 1, which is composed of the following functional components:

- Operations Support System (OSS). Represents an operator that deploys an application/service. A GUI has been programmed in Net2Plan to emulate the operator behavior.
- NFV-Orchestrator (NFV-O), implemented with an ETSI OSM instance in charge of the virtualization infrastructure that manages and deploys VNFs leveraging in VIMs.

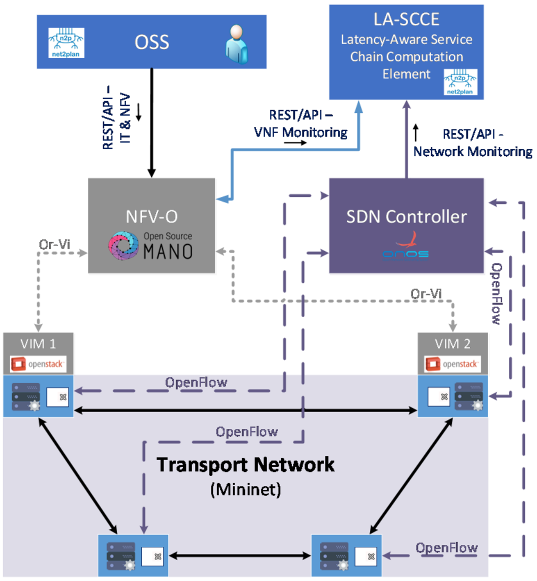

Fig. 1: Demonstration Architecture

- Virtual Infrastructure Managers (VIMs): are responsible for the instantiation and hosting the virtual machines (VMs) of the VNFs. OpenStack provides this actions.
- SDN-control of the metropolitan network is performed by ONOS via OpenFlow.
- Emulated transport is provided by Mininet so that realistic packet flows through a metropolitan-size network are established between the VNFs instantiated in commodity hardware appliances.
- Latency-Aware Service Chain Computation Element (LA-SCCE) represents an evolution of the classical Path Computation Element (PCE) tuned for service chain allocations, where the path is constrained to traverse a sequence of VNFs while satisfying latency requirements. The LA-SCCE is implemented as an algorithm in the Net2Plan planning tool.

3. Demonstration Testbed Configuration

Fig. 2 shows a schema of the physical implementation of this demonstration which is composed of a personal laptop, three high-performance mini-PCs and two regular auto-configured switches. In particular, the Net2Plan planning tool and OSM (in a VM) run in the personal laptop. The two mini-PCs emulate AMEN/MCEN nodes in a metropolitan network with a VIM OpenStack instance in each of them. One mini-PC runs the ONOS SDN controller and the Mininet emulated metropolitan transport network.

Fig. 2 also clearly differentiates the control-plane network (in blue) from the data-plane network (in green). In the control plane, one regular switch interconnects OSM, the two OpenStack instances, ONOS and Net2Plan. The data-plane network provides data-plane layer-3 connectivity using the other switch with specific physical interfaces to communicate the emulated hosts in Mininet with the VNFs instantiated by the VIMs in the mini-PCs. The OpenStack VIMs are set up with an internal private network to place the VMs of the VNFs. The VMs have external connectivity given by floating IPs of a public network with access to the data plane network.

4. Demonstration workflow

The demonstration follows the following steps:

a. Firstly, in the load process, Net2plan receives the entire information of the NVF and IT resources via REST/API from OSM and the VIMs and the emulated transport network data from ONOS.

b. Fig. 3 (a) shows how the user defines the service chain request (origin and destination nodes, sorted sequence of VNFs, bandwidth and latency) from the Net2plan GUI.

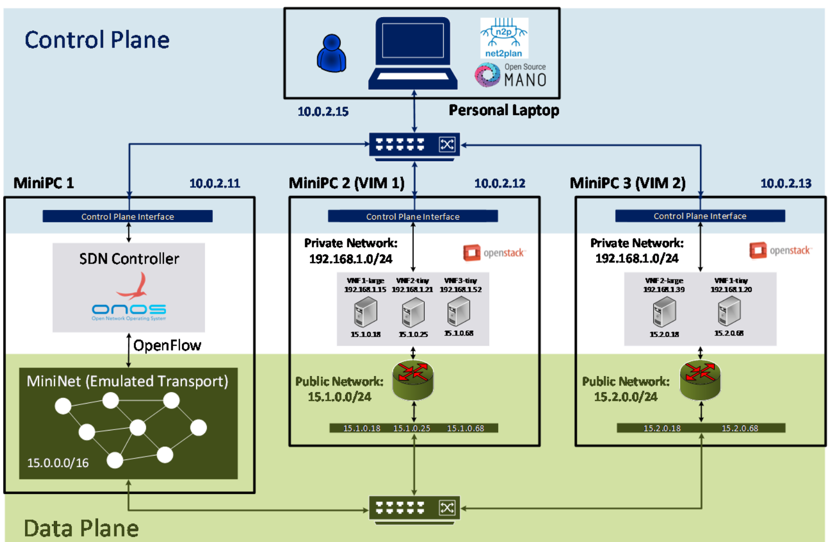

Fig.2 – Demonstration testbed configuration.

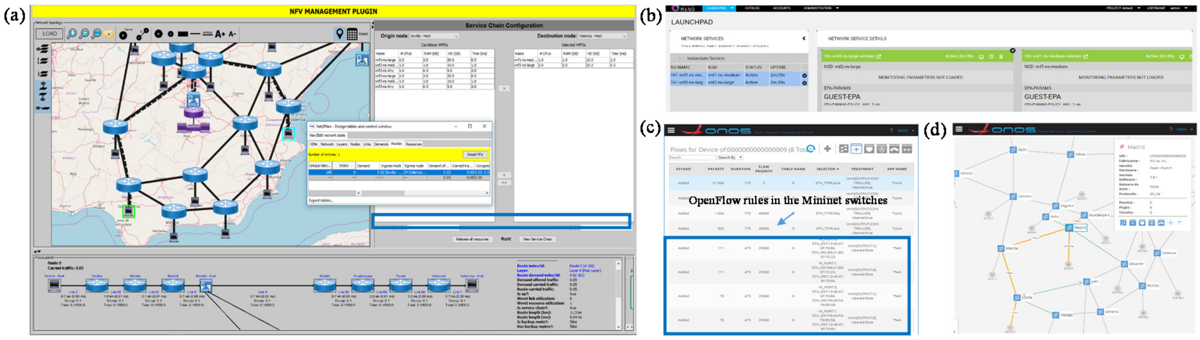

Fig.3 – Snapshots of the open-source modules used in the demonstration. (a) Net2Plan graphical user interface (GUI) showing the (Spanish) and the route implementation passing through an emulated AMEN/MCEN (Madrid) with two instantiated VNFs. (b) OSM GUI showing the two VNFs up and running. ONOS GUI showing (c) OpenFlow rules in a virtual Mininet switch (d) the (Spanish) topology and several flows.

c. Net2plan receives the service chain request and calculates k-shortest service-chain-based paths that need to guarantee the End-to-End latency specifications. The algorithm returns the SC implementation details that meet the specifications, which includes the transport network path and the placement of the VNFs in the VIMs. Note that Net2plan is the only open-source module capable to accomplish the latency needs providing optimal VNF placement and transport paths accounting on realistic information from an SDN-controlled network regarding propagation time and processing time of the VNFs.

d. OSM is notified about the VNFs placement and starts the instantiation of the VNFs in the corresponding VIMs as indicates Fig.3 (b).

e. Once the VNFs are ready, as it can be seen in Fig 3 (c) and (d), ONOS creates the OpenFlow rules following the indications of the Net2plan algorithm providing real connectivity between the origin and destination nodes through the sorted sequence of VNFs.

5. Conclusions

In a dynamic SDN/NFV environment, this demonstration proves (i) Net2Plan's latency-aware SC path computation leveraging on VNFs' processing time and realistic information provided by an ONOS instance that controls an emulated Mininet data plane and (ii) Net2Plan's assistance to OSM so that the chosen SC path is actually instantiated in commodity hardware appliances managed by OpenStack. This demonstration relies on open-source software modules and paves the way to accomplish, among others, the latency requirements of the incoming 5G era.

Acknowledgements

The research leading to these results has received funding from the European Commission for the H2020-ICT-2016-2 METRO-HAUL project (G.A.761727), the Spanish National Project ONOFRE-2 (TEC2017-84423-C3-1-P) and the H2020-MSCA-IF-2016 INSPIRING-SNI project (G.A.750611).

References

[1] Cisco Visual Networking Index: Forecast and Methodology, 2016-2021.

[2] Metro Network Traffic Growth: An Architecture Impact Study, Bells Labs, 2013.

[3] Metro-Haul Project. https://metro-haul.eu/.

[4] P. Pavon-Marino, F.J. Moreno-Muro, N. Skorin-Kapov, "Evolution of Core Traffic for Growing CDNs: Is the growth rate of core network traffic overestimated?", *OFC,* 2017.

[5] F.J. Moreno-Muro, *et. al*, "Joint Optimal Service Chain Allocation, VNF instantiation and Metro Network Resource Management Demonstration", *OFC*, 2018.

[6] Net2Plan [Online] Available: http://www.net2plan.com/

[7] Open Source Mano. [Online]. Available: https://osm.etsi.org/

[8] Open Network Operating System (ONOS) [Online]. Available: https://onosproject.org/