



<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *43rd European Conference and Exhibition on Optical Communication (ECOC), September 2017*.

Citation for the original published paper:

Raza, M R., Rostami, A., Vidal, A., Santos, M., Wosinska, L. et al. (2017)
Priority-Aware Service Orchestration Using Big Data Analytics for Dynamic Slicing in
5G Transport Networks
In:

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-227881>

Priority-Aware Service Orchestration Using Big Data Analytics for Dynamic Slicing in 5G Transport Networks

M. R. Raza⁽¹⁾, A. Rostami⁽²⁾, A. Vidal⁽³⁾, M. A. S. Santos⁽³⁾, L. Wosinska⁽¹⁾, P. Monti⁽¹⁾

⁽¹⁾ KTH Royal Institute of Technology, Electrum 229, SE-164 40 Kista, Sweden, mrraza@kth.se

⁽²⁾ Ericsson Research, Kista, Sweden, ⁽³⁾ Ericsson Research, Indaiatuba, Brazil

Abstract We demonstrate how to efficiently scale up/down resource slices allocated to tenants with different service priorities. Experimental results show that our proposed strategy—based on big data analytics—lowers service degradation by more than 51%, compared to priority unaware approaches.

Introduction

The 5th generation of mobile networks (5G) is expected to support a variety of services, spanning over multiple technology and/or administrative domains, and with diverse requirements. Hence, the network providers are looking into multi-purpose programmable transport network solutions, where the same infrastructure can be used to accommodate multiple tenants simultaneously. Software defined networking (SDN) is the technology of choice to enable such a vision. Using SDN, it is possible to allocate a portion of the end-to-end network resources to each tenant, i.e., a concept also referred to as *slicing*. In order to improve the network resource utilization, the slice assigned to a tenant can be dynamically adjusted based on the temporal and/or spatial variation of the tenant's service requirements¹. We recently presented a proof of concept of *dynamic slicing* in a multipurpose transport network with a simple first-come-first-served (FCFS) orchestration policy². The FCFS policy cannot distinguish among different service requirements when tenants compete for the access of the shared infrastructure resources. Consequently, an infrastructure provider might not be able to allocate its shared resources in the most efficient way. Therefore, in this paper, we propose a priority-aware resource sharing (PARS) strategy for the efficient allocation of shared resources, and experimentally demonstrate how it significantly lowers the service degradation level experienced by the tenants. Results show that the PARS strategy lowers service degradation by more than 51% compared to FCFS policy.

Experimental Set-up and Use Case Definition

Fig. 1a presents the system architecture of the end-to-end programmable platform for dynamic service creation², used to demonstrate the concepts presented in this work. We assume three technology domains. The *optical domain* is based on the dense wavelength division multiplexing (DWDM) technology and provides connectivity services to the other two domains. The transport controller is based on

OpenDaylight (ODL) which has been extended for controlling DWDM switches and tunable transceivers. In this work, an optical network emulator (ONE) is used to mimic the behaviour of a real DWDM network³. The *radio domain* provides mobile broadband services using a number of Long Term Evolution (LTE) Macro and Small cells (i.e., MCs and SCs). They are deployed according to the centralized radio access network (C-RAN) concept, where remote radio units (RRUs) and baseband processing units (BBUs) are interconnected through a high-speed common public radio interface (CPRI). The radio domain relies on the optical domain for the provisioning of CPRI flows, where each CPRI flow requires the full capacity of a wavelength channel. The RAN control functions are dynamically programmed through a centralized command line interface (CLI)-based manager. These functions include activation/configuration of cells and assignment of BBU resources to RRUs. The *cloud domain* comprises two (optically interconnected) datacenters (DCs), which provide compute/storage services to the cloud tenant. The cloud controller is responsible for the management of the DC resources and for the operations of the intra-DC network in each DC. On top of the three domains, an *orchestrator* (i) performs a cross-domain management of the radio, transport, and cloud resources, and (ii) provides a virtualized representation of end-to-end resources to the network applications. The orchestrator is also equipped with a Big Data Analytics (BDA) module, used to predict the behaviour of the tenant applications (e.g., temporal variation of the tenants' traffic levels).

The use case under exam assumes three tenants: two mobile service providers (MSPA and MSPB) and one cloud service provider (CSP). MSPA and MSPB get a slice of radio and transport resources, which consists of a collection of dedicated (i.e., MCs, BBUs and the wavelength resources to connect them) and shared (i.e., SCs, BBUs, and wavelengths) resources. CSP is assigned a slice of cloud (i.e., compute and storage in both DC1 and DC2)

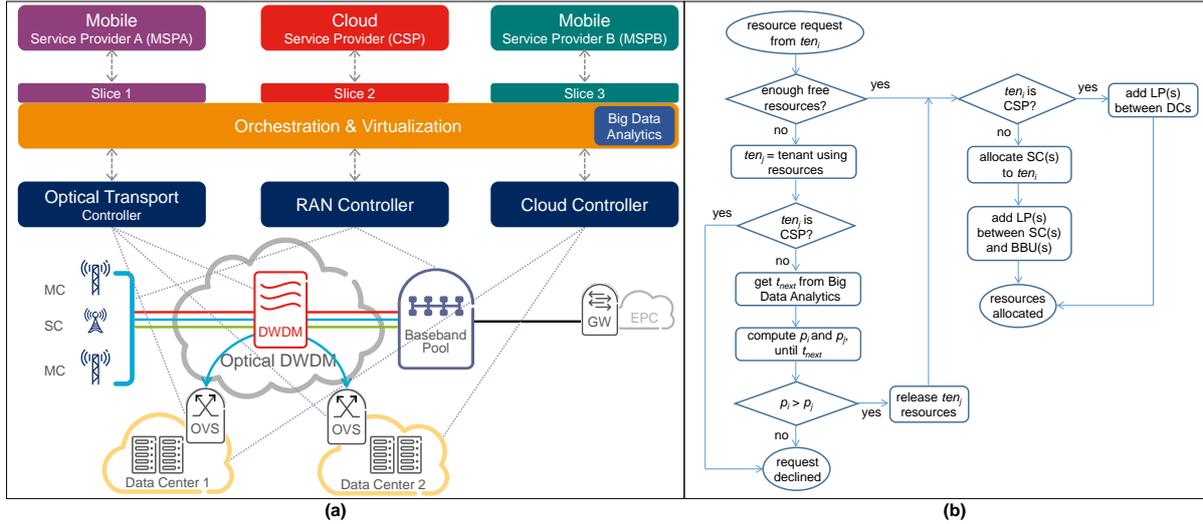


Fig. 1: (a) System architecture (b) Flow chart of the Priority Aware Resource Sharing (PARS) strategy

resources, as well as the transport (i.e., wavelength) resources to connect the DCs when data needs to be backed up from DC1 to DC2. The cloud resources are dedicated for the CSP, whereas the corresponding transport resources are taken from the transport resources shared among all the three tenants. Each tenant monitors the use of its resources and decides when it should scale up/down its own slice. For example, when MSPA realizes that the demand of its users exceeds a certain threshold, it asks the orchestrator to scale up the slice (i.e., activate a number of shared SCs and connect them to BBUs using wavelength resources that are in the shared pool). Similarly, when it is time for a backup, the CSP asks the orchestrator to assign wavelength resources (from the shared pool) to connect DC1 to DC2.

The orchestrator is in charge of resolving possible conflicts among tenants for accessing shared resources. If a tenant cannot access shared resources when needed, it experiences services degradation and the infrastructure provider will have to pay a penalty proportional to the degradation level. This penalty varies depending on the priority of a given service. To minimize the penalty, the infrastructure provider needs to adopt a smart orchestration policy. For example, applying a simple FCFS orchestration policy² might lead to high penalties. To minimize the overall penalty, we propose a priority-aware resource sharing (PARS) policy for assigning the shared resources to tenants' slices. PARS utilizes the BDA module to predict when tenants might need to access shared resources.

PARS: Priority Aware Resource Sharing

The intuition behind PARS is to use the prediction information from BDA to decide if it is advantageous (from a penalty minimization point of view) to pre-empt some of the shared

resources already assigned to a tenant running a low-priority service, in order to accommodate a tenant with high-priority. The assumption made while designing PARS is that transport resources assigned to CSP cannot be pre-empted once provisioned, in order to avoid interrupting the DC back-up process. Fig. 1b presents the flowchart of PARS strategy. When a tenant (i.e., ten_i) requests to use shared resources at time t_o , the orchestrator checks their availability. If free, the orchestrator allocates the resources to ten_i . If ten_i is the CSP, the orchestrator only needs to establish one or more lightpaths (LPs) between the two DCs. If ten_i is a mobile tenant, the orchestrator assigns as many SCs as needed and connects them to the BBU pool. On the other hand, if there are no free shared resources, the orchestrator checks if the tenant that is using the resources (i.e., ten_j) can be pre-empted. In this case, the orchestrator needs to understand if the penalty (p_i) incurred by not accommodating ten_i is higher or lower than the penalty (p_j) incurred by pre-empting ten_j . To do this, the orchestrator identifies t_{next} as the time instance in which the needed resources can be expected to become available (i.e., thanks to the BDA output). The penalties for ten_i and ten_j are computed by:

$$p = \alpha \cdot \left[\int_{t_o}^{t_{next}} C_{req}(t) - \int_{t_o}^{t_{next}} C_{prov}(t) \right], \quad (1)$$

where α is the penalty coefficient of tenant, $C_{req}(t)$ is the total capacity required, and $C_{prov}(t)$ is the total capacity provided during the time interval (t_o, t_{next}) . If $p_i > p_j$, ten_j will be pre-empted, otherwise ten_i will not be provided the resources that it needs.

Performance Evaluation

PARS has been implemented in the demo setting described in Fig. 1a. For the experimental evaluation, we consider a simple case where

each of the MSPA and MSPB gets one dedicated MC, one BBU, and one wavelength. MSPA and MSPB also share one SC, one BBU, and one wavelength. The DWDM network has three wavelengths, where two are dedicated for connecting MC of each mobile tenant to BBU. The third wavelength is shared among MSPA and MSPB (for connecting SC to BBU), and CSP (for backup between the two DCs). The BDA is trained for 30 days using random variations of the traffic profiles taken from³ (for MSPA and MSPB) and from⁴ (for CSP). The traffic predictions for three tenants are shown in Fig. 2. Note that, the CSP profile is a step function that describes when and for how long a DC backup process takes place. The figure also shows the value of the threshold (i.e., 0.6) used by the mobile tenants to trigger the slice scale up procedure.

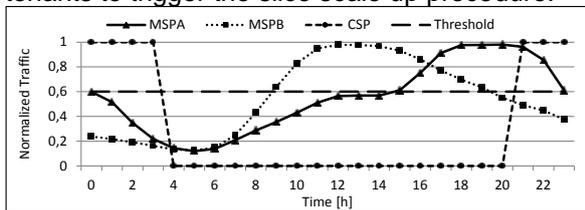


Fig. 2: Traffic profiles learned from Big Data Analytics.

Our experiment emulates a network in operation for 10 days. For each day, the traffic requirements for the three tenants are randomly generated by varying the profiles retrieved from³ and⁴. The tenants' traffic requirements are sampled every hour. The performance of PARS is benchmarked against FCFS. In our demo experiment, we define α , β and γ as the penalty coefficients for MSPA, MSPB, and CSP respectively. For MSPA and MSPB, the penalty is proportional to the over-the-threshold portion of the traffic that cannot be supported. For CSP, the penalty is proportional to the delay experienced before backup can be started. We consider three cases where the penalty coefficients have different relative values. Fig. 3 shows the average value of the service degradation (scaled by value of penalty coefficients) for the three tenants using FCFS and PARS. **Case I** (Fig. 3a) is for tenants with equal priorities, i.e., $\alpha=\beta=1$, $\gamma=0.4$ penalty units [PUs], which corresponds to the three tenants experiencing the same penalty

value over one hour of maximum degradation. With FCFS, MSPB has very low service degradation. When it is over the threshold, MSPA and CSP usually do not require a slice scale up (Fig. 2). MSPA mostly waits for MSPB to release the shared resources, while CSP needs to wait for MSPA. With PARS, CSP pre-empts MSPA which, in turn, tries to pre-empt MSPB, leading to an increase in MSPB degradation. On average, the total degradation with PARS is 51% better than with FCFS. In **Case II** (Fig. 3b) MSPB has the highest priority (i.e., $\alpha=1$, $\beta=100$, $\gamma=0.4$ [PUs]). In this case, again PARS outperforms FCFS. The tenant with the worst performance is MSPA because it is not allowed to pre-empt MSPB. Similar conclusion can be drawn for **Case III** (Fig. 3c), where CSP has the highest priority and MSPB has higher priority than MSPA ($\alpha=1$, $\beta=10$, $\gamma=4$ [PUs]). In this case, there is no degradation observed for CSP using PARS, as compared to very high value obtained with FCFS.

Conclusions

This paper designed and demonstrated a priority aware strategy (based on big data analytics predictions) for dynamically scaling up and down resource slices in an end-to-end programmable platform for dynamic 5G service creation. The proposed strategy brings significant improvements in terms of lowering service degradation, as compared to the first-come-first-served scaling approach.

Acknowledgements

This work has the support of the Kista 5G Transport Lab (K5) project funded by VINNOVA and Ericsson, and of the H2020-ICT-2014 project 5GEx (Grant Agreement no. 671636).

References

- [1] M. R. Raza, et al., "Benefits of Programmability in 5G Transport Networks," OFC, 2017.
- [2] A. Rostami, et al., "An End-to-End Programmable Platform for Dynamic Service Creation in 5G Networks," OFC, 2017.
- [3] M. R. Raza, et al., "Demonstration of Dynamic Resource Sharing Benefits in an Optical C-RAN," JOCN, 2016.
- [4] F. Morales, et al., "Virtual Network Topology Reconfiguration based on Big Data Analytics for Traffic Prediction," OFC, 2016.

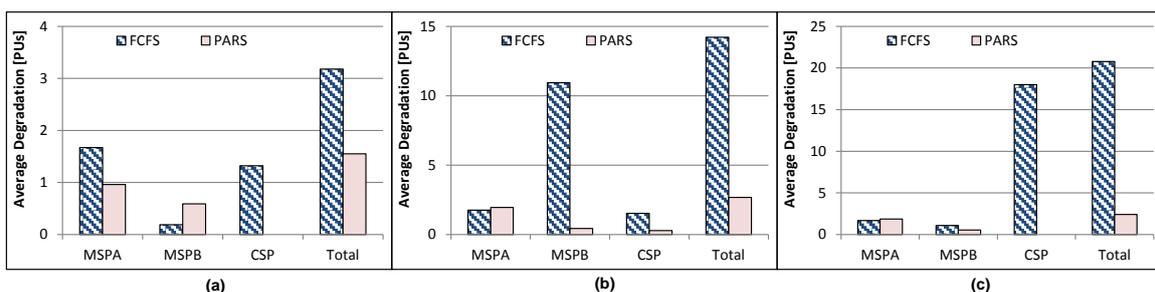


Fig. 3: Average service degradation FCFS vs. PARS: (a) Case I, (b) Case II, and (c) Case III.