Dynamic Backhauling within Converged Networks

Alextian B. Liberato, Magnos Martinello
Software Defined Networks Research Group - NERDS
Federal University of Espírito Santo, Vitoria/ES, Brazil
alextian@ifes.edu.br, magnos@inf.ufes.br

Moises R. N. Ribeiro
Telecommunications Laboratory
Federal University of Espírito Santo, Vitoria/ES, Brazil
moises@ele.ufes.br

Johann M. Marquez-Barja, Nicholas Kaminski, Luiz A. DaSilva
CONNECT, Centre for Future Networks and Communications
Trinity College Dublin, Ireland
{marquejm,kaminskn,dasilval}@tcd.ie

ABSTRACT

To fulfill the emerging demands of mobile users, future networks need to support seamless connectivity. Flexible resource allocation that responds to dynamic trends of user demand and mobility can be achieved through dynamic backhauling that supports elastic routing to setup and teardown paths rapidly as users move and networks adapt. Fortunately, methods such as OpenFlow offer the necessary capabilities to enable such a paradigm. This paper investigates the suitability of such methods in supporting seamless connectivity.

CCS Concepts

•Networks → Network architectures; Network design principles; Network resources allocation; *Network dynamics; Network manageability;*

1. INTRODUCTION

End-to-end network deployments require solutions that encompass both access and backhauling technologies. Currently, wireless access technologies and optical networks are deployed in isolation, considering only those requirements that directly apply to either domain and ignoring the features of the other. When designing an end-to-end network, orchestration of resources in a manner that seamlessly combines both wireless and optical features offers the only means to fully utilize available elements. In that regard we envision a dynamic solution that deploys end-to-end network solutions considering both wireless and optical resources in an integrated manner, combining heterogeneous resources, through virtualization, to enable a converged optical and wireless network. We call this solution Ephemeral Converged Networks (ECNs) [3], ephemeral in the sense of time, since the network is dynamically deployed for the lenght of time that the service requires connectivity, then resources are released.

Enabling an ECN demands a dynamic backhauling mechanism to allow the rapid reconfiguration of network connectivity. In pursuit of such a mechanism, we study the suitability of OpenFlow (OF), on the basis of reconfiguration latency, in supporting seamless connectivity. Moreover, we demonstrate the advantages of dynamic backhauling by exploring initial strategies supporting dynamic end-to-end networks.

LANCOMM, August 22-26, 2016, Florianopolis, Brazil
© 2016 ACM. ISBN 978-1-4503-4426-5/16/08…$15.00
DOI: http://dx.doi.org/10.1145/2940116.2940122

2. EPHEMERAL CONVERGED NETWORKS

The ECN paradigm is based on the concept of Networks without Borders (NwoB) in general, and Ephemeral Wireless Networks (EWN)[1] in particular, which considers a pool of resources (spectrum, radio access networks, storage, etc.) from which a virtual network can be architected and deployed, where user services are decoupled from the underlying network control mechanisms or infrastructure. The components of an ECN, where an end-user subscribes to one or more services offered by the service provider; such services will be tied to connectivity access, which will be provided and controlled by a Broker. This Broker interacts with one or more infrastructure/network provider to determine the appropriate set of resources in the network that will yield the desired coverage and capacity for the network, and that will best meet the needs of the end-user and the service. These resources may be of heterogeneous nature and ownership, including assets in the current mobile and wireless networks, as well as optical networks and their different elements for backhauling. In this paper, our focus is on the feasibility of dynamic backhauling approaches.

3. BACKHAULING EVALUATION

[1]http://www.slideshare.net/JohannMMarquezBarja/ephemeral-wireless-networks

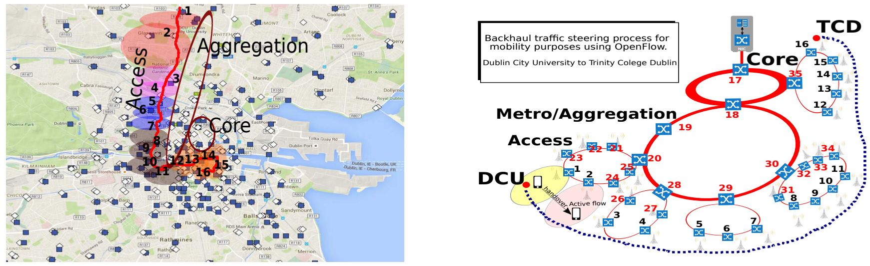

(a) (b)

Figure 1: Backhauling: (a) Trajectory bus and base station location, (b) Network topology in Dublin city with reconfigurability points of attachment.

The experiments conducted were designed to investigate the agility of OpenFlow Protocol for provisioning backhaul paths to support the flexibility required by Ephemeral Converged Networks. Our study considers the network topology illustrated in Fig. 1(b). The network is composed of (i) access networks with base stations distributed throughout the city of Dublin, as denoted by black numbers (from 1 to 16); (ii) OpenFlow enabled switches at the aggregation and at core levels represented by red numbers (from 17 to 35).

Fig. 1 also shows the trajectory of a public bus traveling from Dublin City University (DCU) to Trinity College Dublin (TCD). A bus following this trajectory travels a distance of 6 km in 42 minutes. The delay between the fiber links and the base stations was set to 0.005 ms/km. For the wireless links, we set the delay to 0.003 ms/km. The public transport route is based on real information, publicly available, from Dublin Bus. The deployed fiber layouts are based on an illustrative scenario. In order to evaluate a dynamic backhaul for traffic steering, we have implemented a mobility scenario in the Mininet 2.2 emulation environment, using the Open vSwitch 2.3 software switch. To represent our testbed we used the configuration set of OpenFlow switch HP ProCurve J9451A [1]. In such a scenario, a user inside the bus uses a video streaming application.

The **reconfigurability latency** is defined as the time taken by the process of reconfiguring the backhaul and transferring a data session (video stream) from one network to another [2]. We run a set of experiments measuring that latency assuming that a server connected to the core network (switch 17 in Fig. 1) provides the video streaming. The measure of reconfigurability latency includes: latency communication between the switch and the controller that receives packet_in OF message, the processing of packet_in in the controller, the sending of flow_mod OF message(s) to the switch(es); and the new rule activated in the switch with a new action of output port. Note that reconfigurability can happen at intra access or inter access networks, requiring table updates for a different number of switches. The service connectivity depends on the network agility for route reconfiguration. In our emulation, the route computation is performed by a Ryu SDN controller. The route configuration at the switches is established by OpenFlow, operating in three different modes: **Proactive mode:** the controller inserts all necessary rules proactively into the switches to set the route. In this case, a reconfiguration does not need packet_in OF messages from switches to the controller. **Reactive in group mode:** after receiving a packet_in OF message sent from the *new point of attachment (PoA)*[2] to the controller, the controller is able to set the group of switches that belong to this route. **Reactive in single mode:** the controller updates the rule(s) in switches that are part of the new route, after receiving packet_in OF messages.

The concept of reactive in group mode means that updates for new rules are sent at once to all switches leading to the new PoA serving the mobile user. Only the first switch detecting the new flow sends a packet_in message to the controller. The controller then sends flow_mod messages to the whole group of switches on the shortest path to the destination PoA. In contrast, in the single mode, this communication of packet_in then flow_mod happens in every single hop.

Fig. 2 shows a worst-case analysis (e.g. the longest number of switches in the route) of the reconfigurability latency between PoA-10 to PoA-11. As can be seen, the lower bound (baseline) in this scenario is OpenFlow operating in proactive mode, which exhibits a reconfigurability latency of 400 ms on average. In the OpenFlow reactive mode by group, updates were performed in a group of switches, but all at once after just the first packet_in, resulting in reconfigurability latency of around 600 ms. The difference of 200 ms compared to the proactive mode is essentially due to the communication between PoA and controller. For reactive in single mode, the latency reaches 1900 ms because of communications among switches and the controller.

Fig. 3 presents the reconfigurability latency when the number of switches in the route varies. For the case of reactive in single mode, the reconfigurability latency increases linearly with the number of switches in the route. Their flow tables need to be updated to maintain user connectivity. For 5 switches, the reconfigurability latency increases by a factor of approximately 5

[2] New point of attachment is the base station that receives the new connection.

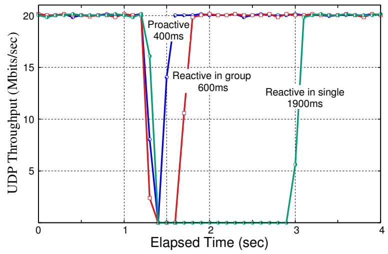

Figure 2: Video streaming flow and its reconfigurability latency between PoA-10 to PoA-11.

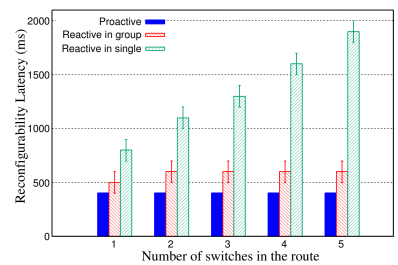

Figure 3: Analysis of the reconfigurability latency for different modes of operation.

(compared to the proactive mode).

Scalability issues should also be considered for the different backhauling operation modes. The number of simultaneous users depends on the switch capacity in terms of its flow table size and dynamics. As can be seen in Figure 1(b), flow tables will be more in demand as we go toward the Central Office as those switches closer to the core will handle more and more flows from different PoA. Evidently, the root element (i.e., switch 17) for all paths used for reaching PoAs will be the scalability bottleneck. Thus, from this element we can roughly infer an estimate for how many flows could be simultaneously served (at least one per mobile user), the average number of active users per PoA, and also the reconfigurability dynamics.

Assuming an OpenFlow switch HP ProCurve J9451A [1], flow tables of 1500 entries are available, which support at most 40 flow_mods per second. As our setup in Figure 1(b) presents 16 switches connected to PoA, on average, over 90 simultaneous flows could be served for proactive mode for the whole network (assuming no route aggregation at the root switch). In addition, in the worst case, 40 (out of 90 active flows) could be under reconfiguration. Although bringing the best reconfigurability latency results, proactive mode will certainly limit mobility dynamics in ECNs due to flow_mod rate limitations in the root switch. Moreover, the 16 routes need to be fully configured (i.e., end-to-end between Central Office and PoAs) through rules with no expiration time limit.

By using reactive modes, on the other hand, we can avoid the end-to-end pre-configuration needed in the proactive mode. Therefore, the number of simultaneous flows, and also the supported set of flows under reconfiguration, can be expanded up to the total flow table capacity of the bottleneck switch. Thus, on average, over 90 active users can be served per PoA. In the worst case, up to 40 active users could be under handover per second across the whole network.

4. CONCLUSION AND FUTURE WORK

This paper has investigated the suitability of OpenFlow in supporting seamless connectivity under dinamically-varying demand profiles that may come from wireless subscribers populations. Our initial results show an analysis of OpenFlow operating the network in three different modes. Tradeoffs between reconfigurability latency and scalability are presented for each mode of operation in order to understand the network agility demanded for the backhauling within converged networks. In the reactive mode by group, the network scales with the number of simultaneous flows to be served, keeping high the number of active users per PoA and reconfigurability latency remains low to reconfigure the route.

Although our preliminary evaluation indicates the feasibility of the use of OpenFlow, for future work we plan to investigate the scalability and agility of the whole network considering routes established by source routing approaches and switches stateless at the core. We also plan to explore some of the proposed solutions through a proof-of-concept experiment using testbeds that span research laboratories in Europe and Brazil.

Acknowledgments

This work has received funding from the European Union's Horizon 2020 for research, technological development, and demonstration under grant agreement no. 688941 (FUTEBOL), as well from the Brazilian Ministry of Science, Technology and Innovation (MCTI) through RNP and CTIC. In addition, in Brazil it has also received funding from CNPq, CAPES and FAPES. We also acknowledge support from the Science Foundation Ireland under grants No. 10/IN.1/I3007 and 14/IA/2527.

5. REFERENCES

[1] D. Y. Huang, K. Yocum, and A. C. Snoeren. High-fidelity switch models for software-defined network emulation. In *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, HotSDN '13, pages 43–48, New York, NY, USA, 2013. ACM.

[2] J. Marquez-Barja, H. Ahmadi, S. Tornell, C. Calafate, J. Cano, P. Manzoni, and L. Da Silva. Breaking the vehicular wireless communications barriers: Vertical handover techniques for heterogeneous networks. *IEEE Transactions on Vehicular Technology*, 64(12):5878–5890, Dec 2015.

[3] J. Marquez-Barja, M. Ruffini, N. Kaminski, N. Marchetti, L. Doyle, and L. DaSilva. Decoupling Resource Ownership From Service Provisioning to Enable Ephemeral Converged Networks (ECNs). In *25th European Conference on Networks and Communications EUCNC [To appear]*, 2016.