

Transport Abstraction Models for an SDN-Controlled Centralized RAN

Matteo Fiorani, Ahmad Rostami, Lena Wosinska, and Paolo Monti

Abstract—In a centralized radio access network (C-RAN) scenario the joint coordination of radio (e.g., remote radio units, baseband units) and transport (e.g., optical cross connects) resources can be achieved via software defined networking (SDN) control plane, where a global orchestrator harmonizes the use of resources across all network segments. The more accurate the information about each domain (i.e., the abstraction of wireless and transport resources) is, the better will be the outcome of the orchestration work. This letter presents three transport resources abstraction models along with their corresponding orchestration policies. Their performance are compared showing that there is not a single best abstraction strategy that fits all the cases. If radio resources are scarce compared to transport resources, complex transport abstraction models are not needed. Contrariwise, if enough radio resources are widely available, more detailed abstraction models are required for achieving good network performance, but at the expense of an increased implementation complexity.

Index Terms—Software defined networking (SDN), centralized RAN (C-RAN), transport abstraction, fronthaul, backhaul, wavelength division multiplexing (WDM).

I. INTRODUCTION

CENTRALIZED RADIO ACCESS NETWORK (C-RAN) is a new concept aimed at a better management of the resources in mobile networks [1]. Unlike traditional RANs, where the radio unit and the *base-band processing unit* (BBU) are co-located at the radio base station site, in a C-RAN architecture BBUs are decoupled from the base stations and centralized in one or more pools, i.e., *BBU Hotel(s)*. In a C-RAN architecture the transport network is divided into two parts: *fronthaul* and *backhaul*. The fronthaul segment is responsible for the traffic between the *remote radio units* (RRUs) and the BBU Hotels, while the backhaul part provides connectivity between the BBU Hotels and the mobile core network, i.e., the *evolved packet core* (EPC). There are stringent requirements (i.e., high capacity, low latency) for the fronthaul network, and the communication between RRUs and BBUs is usually carried out using radio-over-fiber techniques and standard radio interfaces (e.g., the common public radio interface (CPRI) [2]).

The performance of a C-RAN can benefit significantly from a joint coordination among radio, e.g., RRUs and BBUs, and transport resources. A promising option is represented by Software Defined Networking (SDN) [3]. According to the SDN concept, resources in each network segments, i.e., radio and transport, are managed independently by separate logically centralized controllers, which in turn can be organized in a

hierarchical control architecture [4]. An orchestrator sits on top of these controllers with the objective of harmonizing the resource allocation across all the network segments in a scalable manner. The key to the success is exposing to the orchestrator an abstracted view of the resources in each segment with an adequate level of details. On the one hand it is important to hide details of each segment as much as possible to increase the scalability of the control architecture, while on the other hand the orchestrator needs to have enough information to make optimal or quasi-optimal decisions regarding the global resource allocation.

The best level of resources abstraction in SDN-controlled transport networks is currently subject of extensive research. The work in [5] proposes an abstraction model for a packet-over-optical network able to achieve low blocking probability as well as high scalability. The work in [6] proposes and compares three abstraction models for optical networks in terms of maintenance complexity and optimization potential. In addition, abstraction models for SDN-controlled transport networks are currently investigated also by standardization bodies, e.g., in [7]. However, to the best of our knowledge, transport abstraction models for C-RANs have never been investigated.

This letter studies how detailed an abstraction model needs to be in advertising transport resources, focusing specifically on C-RANs with a dense wavelength division multiplexing (DWDM) transport network. More specifically this letter proposes three abstraction models and quantitatively analyzes their blocking probability performance and implementation complexity. From the study presented in the letter it can be concluded that designing an optimal abstraction model for C-RAN transport resources is a delicate task, which requires an adequate information on the availability and reachability of the wireless resources, e.g., position and size of the BBU Hotels. Also, there is a tradeoff between the accuracy of the abstraction level (reflected also in the increased control plane complexity) and the optimality of the decisions regarding resource allocation made by the orchestrator.

II. SDN-CONTROLLED C-RAN ARCHITECTURE

The C-RAN architecture considered in this letter is presented in Fig. 1. In the data plane, RRUs provide wireless mobile access and are connected to optical nodes acting as the (I/O) input/output ports towards the DWDM transport network. The optical nodes consists of reconfigurable optical add-drop multiplexers (ROADMs) performing wavelength switching in the optical domain [8]. The DWDM transport network is organized in a mesh topology, where some network nodes are connected to BBU Hotels and some act as hubs toward the EPC. A transport controller manages the optical elements and sets-up/tears-down lightpaths in the DWDM transport network. RRUs, BBU Hotels, and the EPC are managed by a separate radio controller. The transport and the radio controllers are responsible for abstracting their intra-domain resources and for providing a simplified view of their underlying network to the orchestrator.

Manuscript received February 25, 2015; accepted June 15, 2015. Date of publication June 17, 2015; date of current version August 10, 2015. This work was supported by the Kista 5G Transport Lab (K5) project funded by VINNOVA and Ericsson AB. The associate editor coordinating the review of this paper and approving it for publication was W. Fawaz.

M. Fiorani, L. Wosinska, and P. Monti are with Communications Systems (CoS) Department, Royal Institute of Technology (KTH), 164 40 Stockholm, Sweden (e-mail: fiorani@kth.se).

A. Rostami is with Ericsson Research, 164 83 Stockholm, Sweden (e-mail: ahmad.rostami@ericsson.com).

Digital Object Identifier 10.1109/LCOMM.2015.2446480

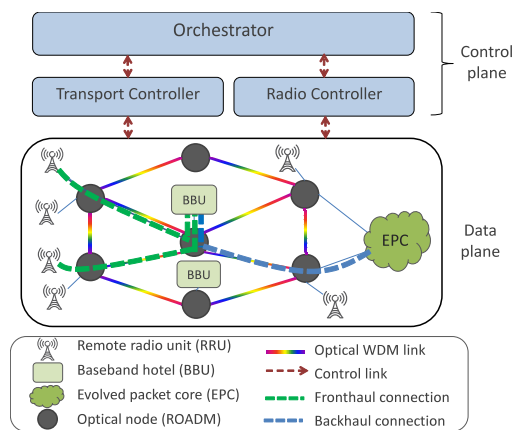


Fig. 1. Reference C-RAN architecture.

We assume that every time a RRU requests a connection to the EPC (e.g., after being activated) the orchestrator needs to provision one fronthaul and one backhaul connection in the DWDM transport network. A fronthaul connection is set up by establishing a dedicated lightpath between a RRU and a BBU Hotel with at least one free BBU port (see Fig. 1). Backhaul traffic is packet-based, thus one backhaul connection can be multiplexed (i.e., using a packet-optical integration technique) with other backhaul traffic over a lightpath between a BBU Hotel and the EPC (see Fig. 1). Based on the resource abstraction information provided by the radio and the transport controllers, the orchestrator selects the I/O ports to be connected to serve the fronthaul/backhaul traffic. If the selected I/O ports can not be connected with a lightpath (i.e., the transport controller realizes that there are not enough wavelength resources) or if there are no free BBU ports in the selected BBU Hotel, the RRU can not be connected to the EPC and the connection request is blocked.

III. TRANSPORT ABSTRACTION MODELS

This section presents three transport abstraction models with their respective orchestration policies for the C-RAN architecture described in Section II. We assume that the orchestrator knows the location of the transport I/O ports, i.e., how RRUs, BBU Hotels, and the EPC are connected to the DWDM transport network. Since the focus of this work is on abstraction models for the optical resources, the proposed orchestration policies use a fixed abstraction model for the wireless resources (out of the scope of this work) where the detailed use of the BBU ports inside each BBU Hotel is not advertised. On the other hand, the orchestrator is informed by the radio controller whenever a specific BBU Hotel runs out of free BBU ports.

The first transport abstraction model is called *Big Switch* (BS). With this model the transport network is presented to the orchestrator as a single node (switch) with no information about the internal transport interconnections, i.e., the connection links between the transport I/O ports and their level of occupancy. As a result, the orchestrator does not know if two transport I/O ports can be connected with a lightpath in the transport physical topology. A schematic representation of how the transport network is abstracted at the orchestrator using BS is shown in Fig. 2(a). This model does not need any update from the transport controller, making it very simple to be maintained. Using BS it is possible to define the following orchestration

policy: each newly activated RRU is connected to the EPC via a default BBU Hotel. The pseudo-code for selecting the BBU Hotel and for establishing a connection between the RRU and the BBU Hotel is shown in the lower part of Fig. 2(a). It can be seen that the activation process of a new RRU is blocked by the transport controller if no lightpaths can be established in the physical topology between the RRU and the BBU Hotel selected by the orchestrator.

The second abstraction model is called *Virtual Link with Constant Weights* (VLCW). With this model the transport network is presented to the orchestrator as a set of virtual links interconnecting the transport I/O ports. Each virtual link is assigned a constant weight (i.e., representing the physical shortest path distance between the I/O ports). When the connectivity between two I/O ports is lost (i.e., no lightpath can be setup between them anymore) the corresponding virtual link is deleted from the abstracted topology. As a result, the orchestrator always knows which are the transport I/O ports that can be connected using a lightpath. A schematic representation of how the transport network is abstracted at the orchestrator using VLVW is shown in Fig. 2(b). This abstraction model has a higher complexity compared to VLCW. In fact, anytime a new RRU is successfully activated, the transport controller is required to report to the orchestrator the connectivity status of all I/O ports to update the transport topology abstraction. Using VLCW it is possible to define the following orchestration policy: each newly activated RRU is connected to the EPC using the closest reachable BBU Hotel with available BBU resources. The pseudo-code for selecting the BBU Hotel and for establishing a connection between the RRU and the BBU Hotel is shown in the lower part of Fig. 2(b).

The third abstraction model is called *Virtual Link with Variable Weights* (VLVW). This model is a different version of the virtual link model, where each virtual link is assigned a weight reflecting the actual number of lightpaths that can be established between the corresponding transport I/O ports. As a result, the orchestrator has a notion of all the possible lightpaths that can be used to interconnect RRUs, BBU Hotels, and the EPC. This abstraction policy can be used to apply more sophisticated traffic engineering solutions, i.e., load balancing in the case of this specific work. A schematic representation of how the transport network is abstracted at the orchestrator using VLVW is shown in Fig. 2(c). VLVW has a higher complexity than BS and VLCW. In fact, each time a new RRU is successfully activated, the transport controller is required to recalculate the weights of all the virtual links and to send the updates to the orchestrator. This requires a relatively large number of updates in order to maintain an up-to-date abstracted topology. Using VLVW it is possible to define the following orchestration policy: each newly activated RRU is connected to the EPC via the BBU Hotel (with available resources) that can be reached using the virtual link with the lowest weight. The pseudo-code for selecting the BBU Hotel and for establishing a connection between the RRU and the BBU Hotel is shown in the lower part of Fig. 2(c).

The pseudo-code describing the orchestration policy used for connecting BBU Hotels and the EPC (i.e., backhaul connections) is not shown in Fig. 2 due to space limits. Notice that backhaul connections for multiple RRUs can be statistically multiplexed over the same lightpath making their impact on the RRU activation blocking probability negligible compared to the impact of fronthaul connections (i.e., at least one order of magnitude higher according to our simulations).

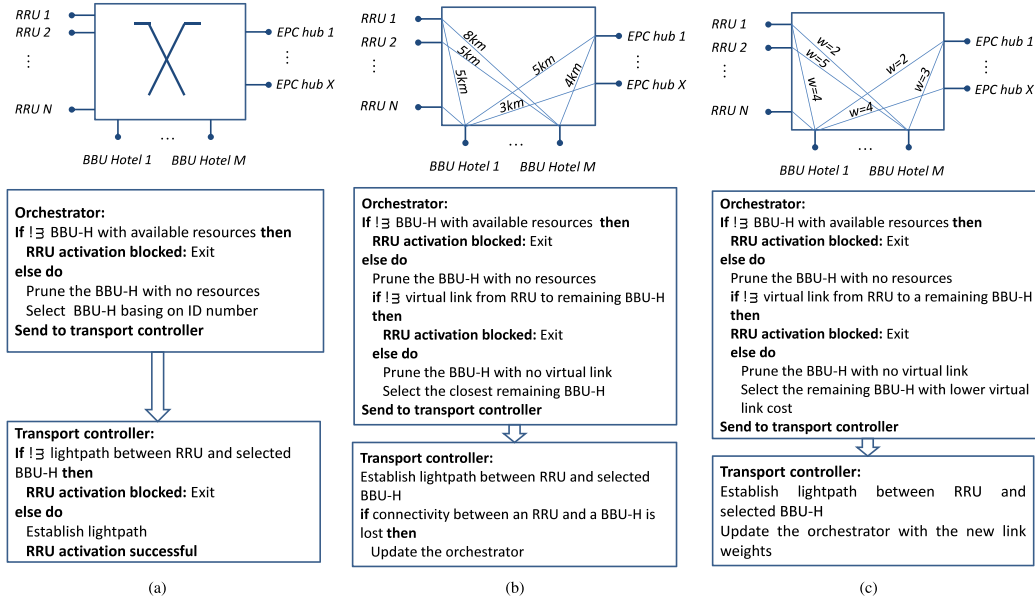


Fig. 2. Representation of the proposed transport abstraction models and pseudo-code of the corresponding orchestration policies. BBU-H: Baseband Hotel. (a) Big switch (BS). (b) Virtual link constant weight (VLCW). (c) Virtual link variable weight (VLVW).

IV. NUMERICAL RESULTS

In order to evaluate the performance of the proposed transport abstraction models we implemented an event driven C++ simulator. The reference topology (Fig. 3(a)) is a metro/aggregation network with 38 nodes, 59 bidirectional fiber-links, and 256 wavelengths per fiber (W) [8]. Each network node is assumed to be connected to several RRUs requesting connections to the EPC for a duration, i.e., holding time (h), that is exponentially distributed with an average of 50 time units. The inter-arrival time of the connection requests at each node is exponentially distributed with an average that has been varied to mimic different traffic load values. Two network nodes (with nodal degree 5 and 6) host a BBU Hotel, while two other nodes act as hubs toward the EPC (Fig. 3(a)).

One important factor to consider while analyzing the performance of the proposed abstraction models is the availability of BBU ports with respect to the amount of transport resources that can be used to reach them. This aspect can be modeled with the parameter η defined as the ratio between the total number of BBU ports available in the network and the maximum number of BBU ports that can be potentially reached via the available wavelength resources:

$$\eta = \frac{(1+b) \cdot \sum_{i=1}^{N_H} N_{BBU}^i}{W \cdot \sum_{i=1}^{N_H} D_{BBU}^i}. \quad (1)$$

b represents the ratio between the backhaul and the fronthaul traffic associated to a single RRU, where the fronthaul traffic is the one generated by the RRU and the backhaul is its counterpart after being processed at a BBU Hotel. N_H is the total number of BBU Hotels, W is the number of wavelengths per fiber, N_{BBU}^i is the number of BBU ports in the i^{th} BBU Hotel, and D_{BBU}^i is the degree of the optical node connected to the i^{th} BBU Hotel. When $\eta < 1$ there are less BBU resources deployed with respect to the available wavelength resources, i.e., the optical transport network is over-dimensioned. In this case, a connection request from a RRU to the EPC is more likely to be blocked because of the lack of available BBU ports. On

the other hand, if $\eta > 1$ there are more BBU ports deployed than the ones that can be potentially reached, i.e., the wireless network resources are over-dimensioned with respect to the wavelength resources. In this case, a connection request from a RRU to the EPC is more likely to be blocked because of the lack of wavelength resources in the DWDM transport network.

In our study b is assumed to be constant for all RRUs and it is equal to 0.1 [2]. The number of BBU ports in the BBU Hotels (N_{BBU}^i) has been varied in order to consider different values of η . We assume that to establish lightpaths the transport controller employs shortest path routing and first fit wavelength assignment (with wavelength continuity constraint). The blocking probability values of the three abstraction models using the orchestration policies presented in Section III are presented in Fig. 3 as a function of the traffic load and for different values of η . The presented results have a confidence interval not exceeding 6%, with 95% confidence level. Similar simulation results were also obtained for different network topologies, but are not presented here because of space constraints.

When $\eta = 0.5$ (Fig. 3(b)) all the transport abstraction models perform the same. This is because the blocking probability is dominated by the scarcity of BBU ports. As a consequence, employing a virtual link abstraction model to allow for a better use of the transport resources does not lead to any benefit. On the other hand, when $\eta = 0.9$ (Fig. 3(c)) the unavailability of wavelength resources becomes noticeable, resulting in a clear difference between the blocking probability obtained with BS and with the virtual link-based abstraction models (i.e., around half an order of magnitude at moderate and high loads). This is due to the fact that with BS the orchestrator has no information about the actual availability of wavelength resources and may choose to connect I/O ports which cannot be interconnected in the physical topology. Meanwhile, the two virtual link-based models have almost the same performance, because the blocking probability due to the unavailability of BBU ports is still significant and limits the benefit of having a very accurate view of the transport resources. When $\eta = 1$ (Fig. 3(d)) the unavailability of wavelength resources starts to play a more prominent role and BS shows very poor performance with

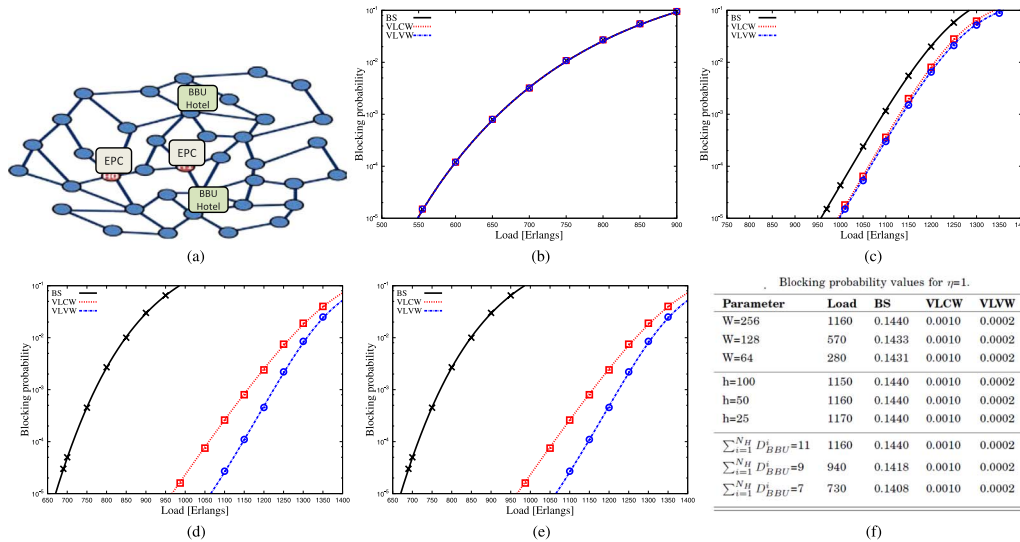


Fig. 3. Reference topology and RRU activation blocking probability as a function of the traffic load and η . (a) Metro topology. (b) $\eta = 0.5$. (c) $\eta = 0.9$. (d) $\eta = 1.0$. (e) $\eta = 1.1$. (f) Sensitivity analysis.

respect to the virtual link-based models. In addition, the difference in blocking performance between VLCW and VLVW becomes evident, i.e., at moderate loads VLVW shows almost one order of magnitude lower blocking probability values than VLCW. This is due to the fact that with VLVW the orchestrator is provided information about all the available lightpaths that can be used to connect two I/O ports, allowing the orchestrator to utilize more efficiently the transport network resources (i.e., load balancing). Finally, with $\eta > 1$ (Fig. 3(e)) the blocking probability is mostly dominated by the unavailability of wavelength resources and the difference in performance among the abstraction models becomes very significant. In particular, BS shows blocking probability values more than three orders of magnitude higher than the one of VLVW. This is because of the lack of information on the availability of transport resources, i.e., all the connection requests are sent to a default BBU Hotel even if it might not be reachable through the transport network.

We performed a sensitivity analysis by varying the value of W , h and the placement of the BBU Hotels, i.e., which results in a different degree of the nodes to which BBU hotels are connected ($\sum_{i=1}^{N_H} D_{BBU}^i$). Fig. 3(f) presents blocking probability values for BS and VLVW obtained with $\eta = 1$ and for an input load such that the blocking probability of VLCW is 0.1%. It can be observed that, in each experiment the BS and VLVW models show almost the same blocking probability. As a consequence, we can conclude that changing the selected parameters leads to shifting the blocking probability curves of Fig. 3(d) along the x-axis (i.e., the load) without affecting their relative performance. Similar results were also obtained for different values of η , but are not reported in the letter because of the space limit. We can conclude that the relative performance of the abstraction models is not a function of the specific simulation parameters.

To evaluate the complexity of the three abstraction models, we measured the number of updates (between the transport controller and the orchestrator) necessary to maintain a correct abstracted topology. The number updates has a major impact on the scalability of the control architecture (e.g., maximum number of network nodes that can be controlled without the orchestrator taking wrong decisions because of outdated information). BS does not require any update and thus does not introduce any scalability limitation. On the other hand, VLCW

and VLVW require a non negligible number of updates to maintain their abstract topology. In our simulations it was found that depending on the load VLCW needs between 15% and 35% less updates than VLVW, which shows always the highest complexity.

V. CONCLUSION

This letter presents three transport abstraction models along with their respective orchestration policies suited for an SDN-controlled C-RAN architecture. It is found that, when BBU resources are scarce (i.e., $\eta < 0.9$), sophisticated transport abstraction strategies are not needed, and a simple big switch model is the best solution. On the other hand, when wavelength resources start to become a bottleneck for the system (i.e., $\eta \geq 0.9$), more advanced transport abstraction models and orchestration policies are required to achieve good blocking probability performance. In particular, VLCW shows a relatively good tradeoff between blocking performance and complexity, while VLVW has the best blocking performance at the cost of relatively high complexity.

REFERENCES

- [1] I. Chih-Lin *et al.*, "Toward green and soft: A 5G perspective," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 66–73, Feb. 2014.
- [2] CPRI Specification V6.0, Aug. 2013.
- [3] K. H. Kim and N. Feamster, "Improving network management with software defined networking," *IEEE Commun. Mag.*, vol. 51, no. 2, pp. 114–119, Feb. 2013.
- [4] P. Ohlen, B. Skubic, Z. Ghebretensae, W. John, and M. Shirazipour, "Software-defined networking in a multi-purpose DWDM-centric metro/aggregation network," in *Proc. IEEE GLOBECOM*, Dec. 2013, pp. 1233–1238.
- [5] P. Iovanna, F. Ubaldi, F. Di Michele, J. P. Fernandez-Palacios Gimenez, and V. Lopez, "E2E traffic engineering routing for transport SDN," in *Proc. IEEE OFC*, Mar. 2014, pp. 1–3.
- [6] A. Autenrieth *et al.*, "Evaluation of virtualization models for optical connectivity service providers," in *Proc. IEEE ONDM*, May 2014, pp. 264–268.
- [7] D. Cecarelli and Y. Lee, "Framework for abstraction and control of transport networks," Internet Eng. Task Force (IETF), Fremont, CA, USA, IETF draft, Mar. 6, 2015.
- [8] S. Zhang, M. Xia, and S. Dahlfort, "Fiber routing, wavelength assignment and multiplexing for DWDM-centric converged metro/aggregation networks," in *Proc. IEEE ECOC*, Sep. 2013, pp. 1–3.