Citation for the original published paper (version of record):

Fiorani, M., Rostami, A., Wosinska, L., Monti, P. (2016)
Abstraction Models for Optical 5G Transport Networks.
Journal of Optical Communications and Networking, 8(9): 656-665
http://dx.doi.org/10.1364/JOCN.8.000656

Abstraction Models for Optical 5G Transport Networks

Matteo Fiorani, Ahmad Rostami, Lena Wosinska and Paolo Monti

***Abstract*—The orchestration of radio, transport and cloud resources is a key enabler for efficient service delivery in 5G networks. Orchestration can be achieved with a hierarchical software defined networking (SDN) control architecture, in which a global orchestrator operates above the domain controllers. In such architecture, the abstraction of resources between the controllers and the orchestrator plays a fundamental role for the system performance. In order to reduce the orchestrator complexity the controllers should hide as much details as possible from the orchestrator. On the other hand, the more details are available to the orchestrator the more optimal resource orchestration strategy can be obtained. In order to assess this trade-off, we recently proposed two transport abstraction models, namely big switch (BiS) and virtual link (VL), for centralized radio access networks (C-RANs) with orchestration of radio and transport resources. We observed that VL can provide a more efficient resource orchestration than BiS at the expense of an increased implementation complexity. The contribution of this paper is twofold. We extend the BiS and VL models to make them applicable to any orchestration scenario. Then, we propose a new transport abstraction model, referred to as optical transport transformation (OTT), which aims at achieving efficient resources orchestration with a reduced implementation complexity. We compare the performance of these new abstraction models in a C-RAN as use case in which backhaul and fronthaul traffic is carried over a dense wavelength division multiplexing (DWDM) network. Our results prove that in a C-RAN the best choice for the transport abstraction model depends on the availability and the reachability of the radio resources. If radio resources are scarce compared to transport resources, complex transport abstraction models are not needed and a BiS abstraction is the best choice. On the other hand, if radio resources are widely available and reachable an OTT model guarantees the best overall performance.**

***Index Terms*—Software defined networking (SDN); Wavelength division multiplexing (WDM); Resources abstraction; Centralized RAN (C-RAN); Fronthaul; Backhaul; 5G transport.**

I. Introduction

THE orchestration of radio, transport and cloud resources is seen as a key enabler for efficient service delivery in the 5th generation of mobile networks (5G). Orchestration can reduce significantly the cost and the time required for provisioning new services with respect to today's standards, by allowing the joint allocation of different types of resources [1], [2]. An example is the dynamic provisioning of virtual content delivery network (vCDN) services, which involves the joint allocation of transport (e.g., broadband access) and cloud (e.g., processing and storage) resources [3]. Another example is the dynamic delivery of mobile access services, which involves the joint allocation of radio (e.g., baseband processing), transport (e.g., backhaul), and cloud (e.g., evolved packet core) resources [4], [5], [6].

An efficient orchestration of radio, transport, and cloud resources can be obtained using a hierarchical software defined networking (SDN) control plane [6], [7] where different logically centralized domain controllers are responsible for managing different types of resources (see Fig. 1). Meanwhile, an overarching orchestrator operates on top of the domain controllers and performs the joint allocation of the resources across the different domains. The domain controllers are responsible for providing to the orchestrator an abstract view of the resources in each domain with an adequate level of details to allow for an efficient orchestration of resources, while preserving the scalability of the control infrastructure. As a consequence, a trade-off exists between the level of abstraction (i.e., the amount of information that each domain controller provides to the orchestrator) and the network performance. On the one hand, the controllers should minimize the amount of information shared with the orchestrator in order to limit its complexity and make the control architecture more scalable. On the other hand, the more information is available at the orchestrator the more effective can be the outcome of the resources orchestration work.

In this paper we focus on the abstraction of transport resources, i.e., the information that the transport controller shares with the orchestrator. Specifically, we study the optimal level of abstraction for optical transport networks, which will play a major role in 5G networks, thanks to their capability of providing high capacity in a cost- and energy-efficient way. Several abstraction models for optical transport networks have been recently proposed [8], [9], [10], [11]. In [8], [9] an abstraction model for optical transport networks is presented which offers low blocking probability for connection requests, while guaranteeing high scalability for the control infrastructure. The authors in [10] propose three abstraction models for optical transport networks and compare them in terms of potential for efficient management of network resources and implementation complexity. In addition, abstraction models for optical transport networks are currently under investigation in standardization bodies, e.g., in [11]. However, these abstraction models are not optimized for scenarios in which the optical transport resources are orchestrated together with radio and cloud resources. These scenarios differ from conventional network scenarios because some transport nodes connected to specific radio and/or cloud resources can play a major role in the overall system performance e.g., the service

Manuscript received February 2016.

M. Fiorani, L. Wosinska and P. Monti are with KTH Royal Institute of Technology, Communication Systems Department (CoS), Kistagången 16, SE-164 40 Stockholm, Sweden (e-mail: fiorani@kth.se).

A. Rostami is with Ericsson Research, Färögatan 6, SE-164 83 Stockholm, Sweden (e-mail: ahmad.rostami@ericsson.com).

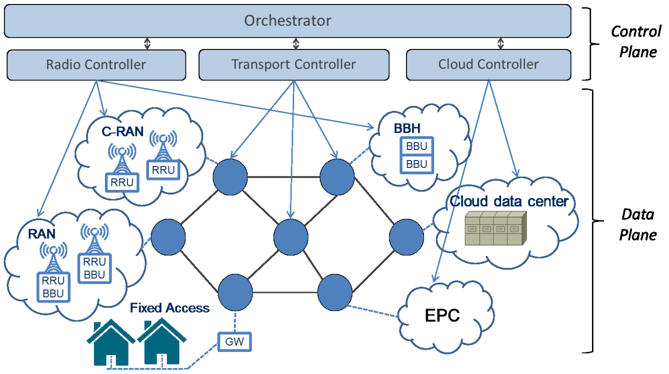

Fig. 1. Reference network architecture. C-RAN: Centralized radio access network, GW: Gateway, BBU: Baseband unit, EPC: Evolved packet core.

blocking probability. In this paper we define these nodes as transport hubs. In order to achieve good blocking performance the transport controller should provide the orchestrator with more detailed information about the transport hubs, while specific information about other transport resources might not be necessary.

In [12] we proposed two simple transport abstraction models for centralized radio access networks (C-RANs) with orchestration of radio and transport resources. These models, referred to as big switch (BiS) and virtual link (VL), are an adaptation of well-known transport abstraction models [13] and are designed specifically for a C-RAN scenario. BiS is very simple to implement and maintain, but it does not allow for an efficient resource orchestration. On the other hand, VL offers lower blocking probability, but at the expense of a higher complexity. The contribution of this paper is twofold. First we extend the BiS and VL models to make them applicable to any scenario in which radio, transport, and cloud resources are jointly orchestrated. We also propose a new abstraction model, i.e., optical transport transformation (OTT), which aims at providing good service blocking performance at low complexity. We analyze the proposed abstraction models in the case of a centralized radio access network (C-RAN) in which the fronthaul and backhaul traffic is carried over an optical dense wavelength division multiplexing (DWDM) network. We selected this specific use case because, despite the large number of studies on cloud orchestration, there are only a few works targeting the joint provisioning of RAN and transport resources. We compare the performance of all the abstraction models presented in the paper in terms of service blocking probability and update complexity. The results prove that in a C-RAN the best choice of a transport abstraction model depends on the availability and the reachability of the radio resources. If radio resources are scarce compared to transport resources, complex transport abstraction models are not needed and BiS is the best choice. On the other hand, if radio resources are available and reachable an OTT model guarantees the best trade-off between service blocking probability and update complexity.

II. Abstraction Models

In Fig. 1 the reference 5G network architecture is presented [14]. It is divided in data plane and control plane. In the data plane a converged optical transport network interconnects different types of resources, such as fixed access gateways (GW), radio access networks (RANs), baseband processing hotels (BBHs), evolved packet core (EPC) and cloud data centers. Meanwhile, the control plane is organized in a hierarchical SDN architecture. It comprises three separate domain controllers, i.e., radio, transport and cloud. The radio controller is responsible for controlling radio resources (e.g., RRUs and BBHs). On the other hand, the transport controller manages the optical network. Finally, the cloud controller is in charge of controlling the cloud resources, which include the EPC (which we suppose to be virtualized) and the cloud data centers. The controllers are supposed to have full knowledge of the resources in their respective domains and are responsible for providing an abstract view of these resources to the overarching orchestrator using an abstraction model. The orchestrator is in charge of accommodating new service requests (e.g., instantiation of a new vCDN or mobile broadband access service) by allocating the required radio, transport and cloud resources. The orchestrator bases its decisions on the abstract view of the resources provided by the domain controllers.

In this paper we focus on the abstraction of the optical transport network, i.e., the abstraction provided by the transport controller to the orchestrator. We introduce the following definitions. *Transport I/O*: a node in the transport network that is connected to radio and/or cloud resources; *Transport hub*: a specific transport I/O that is connected to important radio and/or cloud resources and plays a major role on the system performance. The transport hubs are a subset of the transport I/Os and depend on the specific use case. *Abstract topology*: view of the transport network possessed by the orchestrator; *Physical topology*: view of the transport network possessed by the transport controller. In the following we describe the considered transport abstraction models, namely BiS, VL and OTT.

In the BiS model the transport controller presents the optical network to the orchestrator as a single node (switch). The input ports of the switch correspond to the transport I/O and the output ports of the switch correspond to the transport hubs. In Fig. 2(a) an example of the abstract network view provided to the orchestrator is shown. The orchestrator does not have any information about the internal network connectivity and resources availability. On the other hand, the transport controller assigns a set of weights to the input ports, where each weight represents the cost for establishing a connection towards an output port. The weights can represent for instance the distance between the corresponding transport I/O and transport hub in the physical topology. The weights are calculated during the generation of the abstract topology and are never updated. The BiS model is very simple because it does not require any update messages to be exchanged between the transport controller and the orchestrator in order to maintain the abstract network view. On the other hand, it provides to the orchestrator only a very coarse knowledge of the transport resources which limits the efficiency of resources allocation.

In the VL model the transport controller presents the optical network to the orchestrator as a set of virtual links interconnecting the transport I/Os and the transport hubs. In

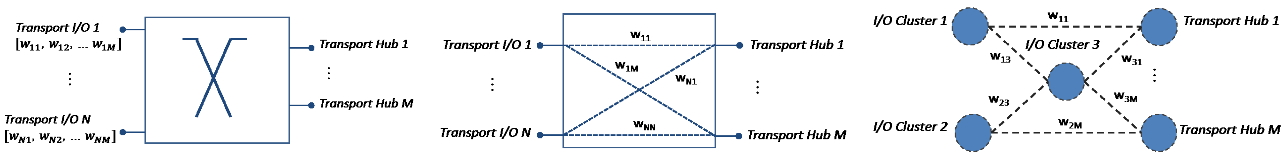

(a) Big switch (BiS). (b) Virtual link (VL). (c) Optical transport transformation (OTT).

Fig. 2. Example of the network views possessed by the orchestrator employing the different transport abstraction models.

Fig. 2(b) an example of the abstract network view generated using the VL model is shown. Each virtual link is assigned one or more weights reflecting the cost for establishing a path between the corresponding transport I/O and transport hub in the physical topology. The weights can be either constant (e.g., representing the distance in the physical topology) or variable (e.g., reflecting the current network resources availability in the physical topology). Every time the connectivity between a transport I/O and a transport hub is lost (i.e., it is not possible to set up a path between them anymore) the corresponding virtual link is deleted from the abstract topology. As a result, the orchestrator is always aware of which transport hub can be reached by any transport I/O. In addition, the orchestrator can exploit the information regarding the transport resources availability to perform traffic engineering decisions (e.g., load balancing). The VL is more complex to implement with respect to the BiS. It requires constant update messages to be exchanged between transport controller and orchestrator in order to maintain the abstract topology. The complexity is a function of the number and the type (i.e., constant or variable) of the weights of the virtual links. On the other hand, VL enables the implementation of more sophisticated orchestration strategies.

The proposed OTT model is a hybrid solution aimed at combining the advantages of BiS and VL. It is designed to maximize the amount of information provided to the orchestrator regarding the reachability of the transport hubs, while hiding as much as possible the details regarding the reachability of the transport I/Os. Basing on these considerations, the transport I/Os are organized in clusters (e.g., according to the distance to the transport hubs in the physical topology) and each cluster is advertised as a single node in the abstract topology. We refer to these abstract nodes as *cluster I/Os*. On the other hand, each transport hub is advertised to the orchestrator as an individual node. The transport controller presents the optical network to the orchestrator as a set of virtual links interconnecting the cluster I/Os and the transport hubs. Similarly as for the VL model, the virtual links can be characterized by one or more weights and these weights may be either constant or variable. OTT presents potentially lower complexity than VL, while still allowing for the implementation of advanced orchestration strategies.

In the following section we explain how the optical transport abstraction models can be applied to a C-RAN scenario. The considered C-RAN scenario is similar to the one analyzed in [12]. However, in this paper we include a new constraint, i.e., the maximum distance between remote radio units (RRUs) and baseband processing units (BBUs), which plays an important role in terms of overall system performance. For this reason we also modify the way in which we apply the BiS and VL models to the C-RAN scenario with respect to [12]. In addition, in this paper we provide for the first time a detailed description of how the abstract topologies are created and of the algorithms used for the resources allocation.

III. Use Case: Centralized RAN

The C-RAN is an architecture for mobile networks in which the BBUs are decoupled from the base stations and centralized in one or more pools, i.e., the baseband processing hotels (BBHs) [15] (see Fig. 3). The pooled BBU resources can be dynamically shared among a large number of RRUs enabling the efficient utilization of the radio resources as well as allowing the implementation of advanced radio coordination schemes. The C-RAN architecture divides the transport network in two parts, i.e., fronthaul and backhaul. The fronthaul segment is responsible for the traffic between the RRUs and the BBHs, while the backhaul part provides connectivity between the BBHs and the EPC. The fronthaul segment is characterized by high capacity and strict latency requirements [16], which imposes a distance limit between the RRUs and the BBHs, which we denote by the parameter D. In other words, a BBH is *reachable* from a RRU if their physical distance is equal or lower than D.

We assume that both fronthaul and backhaul traffic are carried over an optical transport network based on the DWDM technology. We further assume that a hierarchical SDN control plane is employed to control the C-RAN (see Fig. 3). The radio controller manages RRUs and BBHs, the transport controller controls the optical nodes, and the cloud controller manages the EPC. The orchestrator can dynamically deliver mobile access services which span across the three domains. As an example, the orchestrator can decide to activate (or deactivate) the RRUs and connect (disconnect) them to the EPC dynamically according to the actual mobile traffic demand by assigning dynamically the corresponding radio, transport and cloud resources. In the following, we provide the detailed description of the procedure applied by the orchestrator for connecting a newly activated RRU to the EPC.

In the considered C-RAN scenario the optical nodes connected to the RRUs act as transport I/Os, while the optical nodes connected to the BBHs or to the EPC act as the transport hubs. We distinguish the transport hubs as BBH hubs or EPC hubs. We consider that each BBH is composed of a fixed number of BBU ports, where each BBU port provides the baseband processing for one RRU. Every time a new RRU is activated, the orchestrator needs to provision a fronthaul

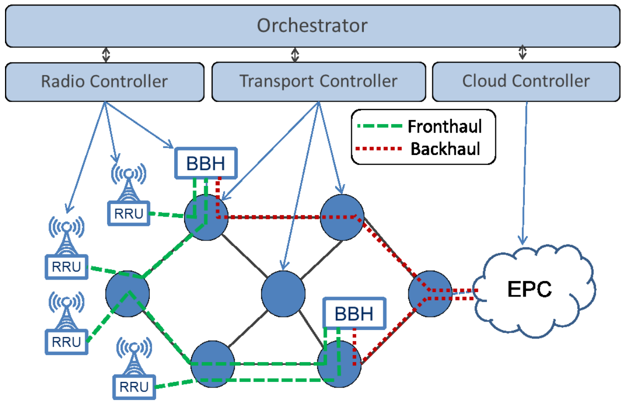

Fig. 3. Reference C-RAN architecture.

connection between the corresponding transport I/O and a BBH hub. The chosen BBH hub must be connected to a BBH with at least one free BBU port to provide the baseband processing functions required by the RRU. In addition, the distance between the BBH and the RRU should be equal or lower than D, i.e., the chosen BBH should be reachable from the RRU. Due to the high capacity and strict latency requirements of fronthaul traffic, a fronthaul connection is set up by reserving a dedicated lightpath between the transport I/O and the chosen BBH hub. In addition, the orchestrator should also provision a backhaul connection between the selected BBH hub and an EPC hub. 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 single lightpath between the selected BBH hub and an EPC hub.

The orchestrator selects the required resources according to the abstraction models provided by the controllers. As a consequence, the employed abstraction model might have a significant impact on the efficiency of the process for activating new RRUs and consequently on the performance of the C-RAN. Since the scope of this work is to compare the performance of different transport abstraction models, we assume a fixed abstraction model for the radio resources (out of the scope of this work) where the detailed use of the BBU ports inside each BBH is not advertised. On the other hand, the orchestrator is informed by the radio controller whenever a specific BBH runs out of free BBU ports. The abstraction model employed by the cloud controller is also out of the scope of this work.

The following terminology will be used in the remaining part of the paper:

- $t \in T$, where T is a set of transport I/Os;
- $b \in B$, where B is a set of BBH hubs;
- $e \in E$, where E is a set of EPC hubs;
- $c \in C$, where C is a set of I/O clusters;
- d_{tb}: distance between transport I/O t and BBH hub b. Corresponds to the length of the shortest lightpath that can be established between t and b in the physical topology;
- $d_{cb} = \max_{t \in c} d_{tb}$: distance between cluster I/O c and BBH hub b;
- (t,b): virtual link between t and b;
- w_{tb}: weight of the virtual link between t and b;
- $\mathbf{w_{be}}$: weight of the virtual link between b and e;
- h_i: boolean variable with value 0 if BBH i has no available BBU ports and value 1 otherwise;
- M: large number.

Algorithm 1 Big switch (BiS)

Orchestrator:
$s \leftarrow 0$; $w_{ts} \leftarrow D$;
for $b = 1$ **to** $|B|$ **do**
  if $h_b = 1$ **and** $w_{tb} \leq w_{ts}$ **then**
    $s \leftarrow b$;
  end if
end for
if $s = 0$ **then**
  Request blocked; Exit;
else
  Send to transport controller request for lightpath between t and s;
end if
Transport Controller:
Run RWA on the physical network;
if $\exists$ lightpath between t and s with length $\leq D$ **then**
  Establish lightpath; Exit;
else
  Request blocked; Exit;
end if

Now we explain how the BiS, VL and OTT abstraction models can be employed by the orchestrator to dynamically establish fronthaul connections between newly activated RRUs and BBHs. The procedure for establishing the corresponding backhaul connections between the BBHs and the EPC is very similar and is not described in the paper because of the space limitation. In addition, it is worth noting that backhaul connections for multiple RRUs can be statistically multiplexed over the same lightpath making their impact on the C-RAN performance negligible compared to the impact of the fronthaul connections [12].

A. BiS model for C-RAN

The abstract topology presented to the orchestrator using a BiS abstraction in the C-RAN is shown in Fig. 4(a). The input ports of the switch correspond to the transport I/Os, while the output ports correspond to the BBH and EPC hubs. The BiS model for C-RAN has been extended with respect to the one presented in [12] in order to take into account the distance constraint between RRUs and BBHs. In this new model each transport I/O t is assigned a set of $|B|$ weights, which correspond to the distance in the physical topology between t and each of the the BBH hubs (i.e., $w_{tb} = d_{tb}$). Similarly, each BBH hub is assigned a set of $|E|$ weights that correspond to the distance to each of the EPC hubs (not shown in Fig. 4(a)). All these weights are calculated by the transport controller when the abstract topology is created for the first time and are never updated.

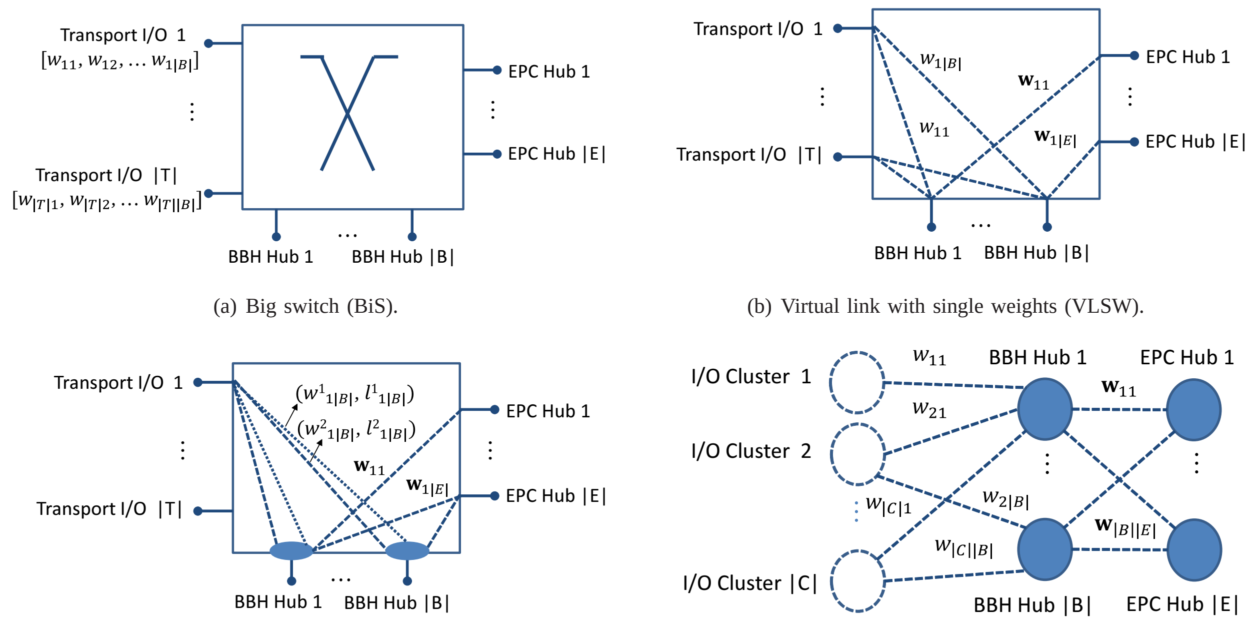

(a) Big switch (BiS). (b) Virtual link with single weights (VLSW).

(c) Virtual link with multiple weights (VLMW). (d) Optical transport transformation (OTT).

Fig. 4. Example of abstract views possessed by the orchestrator in the C-RAN scenario using different transport abstraction models.

The fronthaul connection for a newly activated RRU attached to transport I/O t is provisioned by establishing a new lightpath between t and b, where b satisfies the following requirements: (i) is connected to a BBH with at least one available BBU port (i.e., $h_b = 1$), (ii) is reachable from t (i.e., $w_{tb} < D$), (iii) has the lowest weight among those that satisfy (i) and (ii). If a BBH hub b satisfying these requirements is not identified, then the orchestrator blocks the RRU activation. Otherwise, the orchestrator sends the request for establishing a lightpath between t and b to the transport controller. Note that the decisions of the orchestrator are based on the static weights provided by the transport controller at the time the abstract topology was created. As a consequence, the weights might not reflect correctly the status of the transport resources at every point in time. This means that the orchestrator may take a wrong decision, i.e., it may choose to interconnect a transport I/O t and a BBH hub b which cannot be connected in the physical topology (e.g., because of the lack of wavelength resources). In this case, the request for establishing the lightpath between t and b is blocked by the transport controller which in turn triggers the blocking of the RRU activation. The complete pseudo-code for establishing a fronthaul connection between a newly activated RRU and a BBH using the BiS model is reported in Algorithm 1.

B. VL model for C-RAN

In this subsection we present two possible implementations of the VL abstraction for the C-RAN scenario, i.e., the VL with single weight (VLSW) and the VL with multiple weights (VLMW).

The abstract topology related to the VLSW implementation is shown in Fig. 4(b). The VLSW was already introduced in [12], but here we provide more details on how the topology is built and the resource allocation algorithm is used during the orchestration process. The transport I/O t is connected with a virtual link to BBH hub b only if $d_{tb} < D$. The virtual link (t,b) is characterized by a single weight that corresponds to the distance in the physical topology between t and b, i.e., $w_{tb} = d_{tb}$. The BBH hubs and the EPC hubs are fully connected with virtual links, which are also characterized by weights corresponding to their distance in the physical topology.

The fronthaul connection for a newly activated RRU attached to transport I/O t is provisioned using the same orchestration strategy as for the BiS. However, in VLSW the transport controller after setting up a new lightpath re-computes the distances between transport I/Os and BBH hubs. In case that some distances are changed, the transport controller updates the abstract topology accordingly. In addition, the transport controller may remove virtual links from the abstract topology (e.g., in case all wavelength resources are occupied). As a consequence, the orchestrator is always aware if transport I/O t can be connected to BBH hub b in the physical topology. The complete pseudo-code for establishing a fronthaul connection between a newly activated RRU and a BBH using the VLSW model is reported in Algorithm 2.

The VLMV implementation is shown in Fig. 4(c) and is an extension of the model already presented in [12]. In the new VLMW model, transport I/O t is connected to any reachable BBH hub b with two virtual links, i.e., $(t,b)^1$ and $(t,b)^2$. The virtual link $(t,b)^1$ (dotted line in Fig. 4(c)) represents the lightpath between t and b that employs the fiber links with more available wavelength resources. It is characterized by a weight w^1_{tb} proportional to the wavelength resources and a length l^1_{tb}. The virtual link $(t,b)^2$ (dashed line in Fig. 4(c))

Algorithm 2 Virtual Link with Single Weight (VLSW)

```
Orchestrator:
s ← 0; w_ts ← D;
for b = 1 to |B| do
    if h_b = 1 and ∃(t,b) | w_tb ≤ w_ts then
        s ← b;
    end if
end for
if s = 0 then
    Request blocked; Exit;
else
    Send to transport controller request for lightpath between
    t and s;
end if
Transport Controller:
Run RWA on the physical network;
Establish lightpath between t and s;
for t = 1 to |T| do
    for b = 1 to |B| do
        Compute w_tb;
        if w_tb ≤ D then
            if w_tb changed then
                Update w_tb in the abstract topology;
            end if
        else
            Remove (t,b) from the abstract topology;
        end if
    end for
end for
Exit;
```

Algorithm 3 Virtual Link with Multiple Weights (VLMW)

```
Orchestrator:
s ← 0; x ← 0; w^1_ts ← M; w^2_ts ← M; N_vl ← 2;
for b = 1 to |B| do
    for i = 1 to N_vl do
        if h_b = 1 and (t,b)^i | (l^i_tb ≤ D and w^i_tb < w^i_ts) then
            s ← b; x ← i;
        end if
    end for
end for
if s = 0 then
    Request blocked; Exit;
else
    Send to transport controller request for lightpath between
    t and s on virtual link no. x;
end if
Transport Controller:
Run RWA on the physical network;
Establish lightpath between t and s using (t,s)^x;
for t = 1 to |T| do
    for b = 1 to |B| do
        Compute l^1_tb, l^2_tb, w^1_tb, w^2_tb;
        if (l^1_tb or l^2_tb or w^1_tb or w^2_tb) changed then
            Update the abstract topology;
        end if
    end for
end for
Exit;
```

represents the lightpath between t and b with shortest distance. It is also characterized by a weight w^2_{tb} proportional to the wavelength resources and a length l^2_{tb}. The BBH hubs and the EPC hubs are fully connected with virtual links, which are characterized by single weights corresponding to their distance in the physical topology.

The fronthaul connection for a newly activated RRU attached to transport I/O t is provisioned by establishing a new lightpath between t and b, where b satisfies the following requirements: (i) is connected to a BBH with at least one available BBU port (i.e., $h_b = 1$), (ii) is reachable from t (i.e., $l_{tb} < D$), (iii) can be reached using the lightpath that employs the fiber links with more available wavelength resources among those that satisfy (i) and (ii). If a BBH hub b satisfying these requirements is not identified, then the orchestrator blocks the RRU activation. Otherwise, the orchestrator sends the request for establishing a lightpath between t and b to the transport controller. Note that by always choosing the lightpath that utilizes the fiber links with more available wavelength resources the orchestrator performs efficient load balancing over the transport resources. This is done with the objective of minimizing the probability of blocking an RRU activation due to the lack of wavelength resources. After setting up a new lightpath, the transport controller is responsible for updating the weights and the lengths of all the virtual links in the abstract topology. The pseudo-code for provisioning a fronthaul request using the VLMW model is shown in Algorithm 3.

C. OTT model for C-RAN

The abstract topology presented to the orchestrator using the OTT model in the C-RAN is shown in Fig. 4(d). The transport I/Os are organized in $|C|$ I/O clusters according to their distance to the BBHs. Each I/O cluster is presented as a single node in the abstract topology. Moreover, each BBH hub and EPC hub is advertised as a separate node. Cluster I/O c has a virtual link connecting to BBH hub b only if $d_{cb} < D$. A virtual link (c,b) is characterized by a weight w_{cb} that is proportional to the amount of wavelength resources available in the physical topology between t and b, where t is a transport I/O that belongs to c and has the longest distance to b (i.e., $d_{cb} = d_{tb}$). The BBH hubs and the EPC hubs are fully connected with virtual links characterized by weights corresponding to their distance in the physical topology.

The fronthaul connection for a newly activated RRU attached to transport I/O t is provisioned using the same orchestration strategy as for the VLMW. As a consequence, using OTT the orchestrator performs also load balancing over the transport resources. However, this load balancing is not as efficient as in VLMW because less information about the actual usage of the transport resources at the transport I/Os are available at the orchestrator. Still, the orchestrator has

Algorithm 4 Optical Transport Transformation (OTT)

Orchestrator:
$s \leftarrow 0$; $w_{cs} \leftarrow M$;
for $b = 1$ **to** $|B|$ **do**
  if $h_b = 1$ **and** $\exists\,(c, b) \mid w_{cb} < w_{cs}$ **then**
    $s \leftarrow b$;
  end if
end for
if $s = 0$ **then**
  Request blocked; Exit;
else
  Send to transport controller request for lightpath between t and s;
end if
Transport Controller:
Run RWA on the physical network;
Establish lightpath between t and s;
for $c = 1$ **to** $|C|$ **do**
  for $b = 1$ **to** $|B|$ **do**
    Compute d_{cb};
    if $d_{cb} \leq D$ **then**
      Compute w_{cb};
      if w_{cb} changed **then**
        Update w_{cb} in the abstract topology;
      end if
    else
      Remove (c,b) from the abstract topology;
    end if
  end for
end for
Exit;

a detailed view of the usage of the transport resources at the BBH hubs and EPC hubs that uses for performing load balancing. The pseudo-code for serving a new connection request from an RRU connected to transport I/O t belonging to cluster I/O c is shown in Algorithm 4.

IV. Numerical Results

In this Section we analyze and compare the performance of the transport abstraction models in the C-RAN scenario. Firstly, we introduce the considered performance metrics and then we present and discuss the numerical results.

A. Performance Metrics

The performance metrics that we use for evaluation of the proposed transport abstraction models are the blocking probability and the update complexity. The blocking probability is defined as the probability that a new connection request from an RRU is blocked due to the lack of available resources. An RRU connection request can be blocked either (i) because no available BBU ports are found in any reachable BBH or (ii) because no transport (i.e., wavelength) resources are found to accommodate the corresponding fronthaul and/or backhaul traffic. The update complexity is defined as the number of control messages that need to be exchanged between the transport controller and the orchestrator in order to maintain the abstract topology.

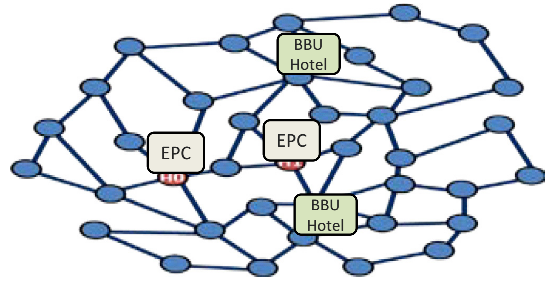

Fig. 5. Reference metro topology.

One important factor to consider while analyzing the performance of the proposed abstraction models in the C-RAN scenario 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 dimensioning 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:

$$\eta = \frac{\lceil (1+r) \cdot \sum_{i=1}^{|B|} N_{BBU}^{i} \rceil}{W \cdot \sum_{i=1}^{|B|} D_{BBU}^{i}}. \quad (1)$$

Here, r 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 BBH. In our study r is assumed to be constant for all RRUs and it is equal to 0.1. W is the number of wavelengths per fiber (in our study we assumed W=256), N_{BBU}^{i} is the number of BBU ports in the i^{th} BBH and D_{BBU}^{i} is the nodal degree of the i^{th} BBH hub. 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 from a newly activated RRU to the EPC is more likely to be blocked because of the lack of an available BBU port in a reachable BBH. On the other hand, if $\eta > 1$ there are more BBU ports deployed than the ones that can be potentially reached, i.e., the radio network resources are over-dimensioned with respect to the wavelength resources. In this case, a connection from a newly activated RRU to the EPC is more likely to be blocked because of the lack of wavelength resources for accommodating the fronthaul and/or backhaul traffic. 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).

B. Simulation Results

In order to evaluate the performance of the proposed transport abstraction models we implemented an event driven C++ simulator. The reference transport topology (Figure 5) is a

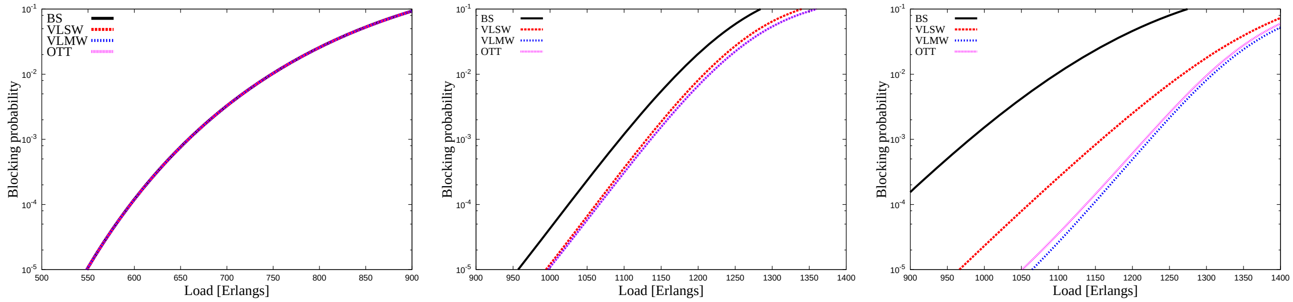

(a) $\eta = 0.5$. (b) $\eta = 0.9$. (c) $\eta = 1$.

Fig. 6. Blocking probability of the proposed abstraction models as a function of the load and for different values of η.

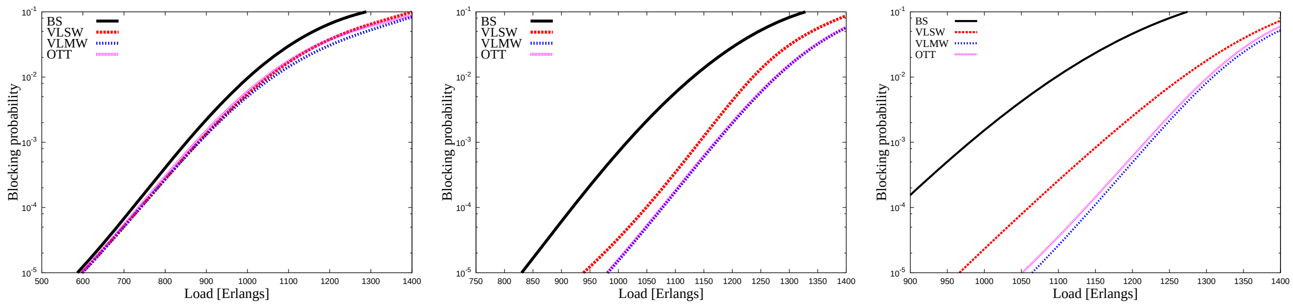

(a) $D = 20\,km$. (b) $D = 30\,km$. (c) $D = 40\,km$.

Fig. 7. Blocking probability of the proposed abstraction models as a function of the load and for different values of D.

metro/aggregation network with 38 nodes and 59 bidirectional fiber-links [17]. For simplicity and without loss of generality we assume that each fiber-link has a length of $5\,km$ and we considered different possible values for the parameter D, i.e., the maximum distance between RRUs and BBHs. Two network nodes (with nodal degree 5 and 6) act as BBH hubs (i.e., $|B| = 2$), while two other nodes act as EPC hubs (i.e., $|E| = 2$). The other network nodes act as transport I/Os (i.e., $|T| = 34$) and are assumed to be connected to RRUs equipped with tunable optical transceivers. For the implementation of the OTT abstraction model we assume that the transport I/Os are organized in four clusters (i.e., $|C| = 4$) according to their distances to the BBH hubs. We assume that each RRU is connected 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. All 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 the space limitation.

Figure 6 shows the blocking probability of the four abstraction models as a function of the traffic load and for different values of η. The maximum distance between RRUs and BBHs has been set to $40\,km$, i.e., $D = 40\,km$. It can be observed that when $\eta = 0.5$ (Fig. 6(a)) all the transport abstraction models achieve the same blocking probability. This is because the blocking probability is dominated by the scarcity of BBU ports, i.e., there is high probability that an RRU activation is blocked because no available BBU ports are found in a reachable BBH. As a consequence, employing a complex abstraction models to provide the orchestrator with a detailed view of the transport resources does not lead to any benefit. On the other hand, when $\eta = 0.9$ (Fig. 6(b)) the unavailability of wavelength resources becomes noticeable, resulting in a clear difference between the blocking probability obtained with BiS and with the other models (i.e., around five times higher at moderate and high loads). This is due to the fact that with BiS the orchestrator has no information about the actual availability of the wavelength resources and may try to connect a transport I/O with a BBH hub which is not reachable in the physical topology. Meanwhile, the VLSW, VLMW and OTT achieve 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. 6(c)) the unavailability of wavelength resources plays a prominent role and BiS shows very poor performance with respect to the other models. In addition, the VLSW also shows clearly higher blocking probability than VLMW and OTT especially at moderate loads where the difference is up to one order of magnitude. This is due to the fact that with VLMW and OTT the orchestrator is provided a more detailed information about the available wavelengths resources at the BBH and EPC hubs, allowing the orchestrator to utilize the transport network resources more efficiently (i.e., by load balancing). Finally, it can be observed that the OTT achieves almost the same performance as VLMW, showing that an accurate view of the wavelength resources usage at the transport I/Os is not necessary.

Figure 7 presents the blocking probability of the four

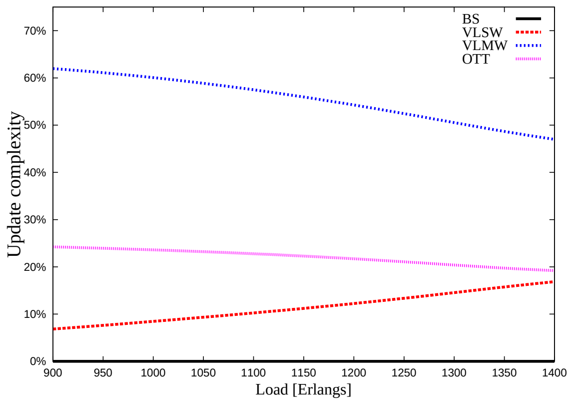

Fig. 8. Percentage of updates with respect to not using abstraction.

abstraction models as a function of the traffic load and for different values of D. The dimensioning parameter η has been set to 1, so that the blocking probability is dominated by the unavailability of wavelength resources. It can be observed that when $D = 20\,km$ (Fig. 7(a)) the abstraction models present almost the same performance, with BiS showing slightly higher blocking probability at high loads. The reason is that with $D = 20\,km$ several transport I/Os can reach only one BBH. As a consequence, the orchestrator is not able to use the information exposed by complex transport abstraction models to perform load balancing and achieve high performance. With $D = 30\,km$ (Fig. 7(b)) the majority of the transport I/Os can reach two BBHs and the difference in the performance of the abstraction models becomes more evident. Specifically, BiS shows a blocking probability up to one order of magnitude higher than VLCW, which in turn has up to five times higher blocking probability than VLMW and OTT. Finally, with $D = 40\,km$ (Fig. 7(c)) all the transport I/Os can reach two BBHs, which significantly increases the difference in the performance between the abstraction models.

Figure 8 shows the update complexity of the abstraction models as a function of the load for $D = 40\,km$ and $\eta = 1$. The number of updates of the proposed abstraction models is normalized against the maximum number of updates, which would be required in case the orchestrator is given full knowledge of all the resources in the optical transport network. Similar results as in Fig. 8 were obtained with different network configurations. It can be observed that BiS shows by far the lowest update complexity as it does not require any update messages to be exchanged between the transport controller and the orchestrator. The VLSW updates the abstract topology only when the distance (i.e., the length of the shortest available lightpath) between a transport I/O and a BBH has changed. As a consequence, its update complexity is relatively low, i.e., less than 20%. Consequently, VLSW can be employed in relatively large network scenarios without bringing any significant scalability problems. The update complexity of the VLSW increases with the load because it is more probable that at high loads the distances between transport I/Os and BBH hubs changes after the establishment of a new connection. On the other hand, the VLMW updates the abstract topology very frequently in order to maintain the information about the current resources utilization in the network. As a consequence, its update complexity can be as high as 65%. This might lead to serious scalability problems in large network scenarios. The update complexity of VLMW is higher at low loads because very frequent updates are required to maintain the correct weights for all the virtual links. The complexity decreases while increasing the load because some virtual links are deleted from the abstract topology (i.e., once the corresponding resources are expired) and thus they do not require updates. Similarly to VLMW, the OTT needs to update the abstract topology frequently in order to maintain the information about the current resource utilization in the network. However, the number of virtual links in OTT is much lower than in VLMW thanks to the use of I/O clusters. Hence, OTT ensures a much lower update complexity than VLMW, which is below 25%. This means that OTT can be applied in relatively large network scenarios.

V. Conclusions

In this paper we analyzed the performance of three abstraction models for optical 5G transport networks. Two of them are extended versions of the ones previously proposed in [12] while the third abstraction model is completely new and aims at achieving the best trade-off between blocking performance and complexity. The models are optimized for scenarios in which the optical transport resources are orchestrated with radio and cloud resources. In such scenarios, some transport nodes (referred to as transport hubs) play a fundamental role on the overall system performance and consequently require a special attention in the abstraction process. We explained how the proposed transport abstraction models can be applied to a C-RAN in which fronthaul and backhaul traffic are carried over a DWDM network. We analyzed the performance of the abstraction models in the C-RAN in terms of service blocking probability and update complexity.

Our results show that the best abstraction model for C-RAN depends on the availability and the reachability of the radio resources. If radio resources are scarce compared to transport resources (i.e., $\eta < 0.9$) or difficult to reach (i.e., $D < 30\,km$) complex transport abstraction models are not needed and a simple BiS can be the best solution. On the other hand, if radio resources are available (i.e., $\eta \geq 0.9$) and reachable (i.e., $D \geq 30\,km$) the newly proposed OTT guarantees the best trade-off between the blocking probability and update complexity.

Acknowledgment

The work described in this paper was carried out with the support of the Kista 5G Transport Lab (K5) project funded by VINNOVA and Ericsson AB.

References

[1] P. Ohlen, B. Skubic, A. Rostami, K. Laraqui, F. Cavaliere, B. Varga, N. Lindqvist, ”Flexibility in 5G Transport Networks: the Key for Meeting the Demand for Connectivity”, Ericsson Technology Review, vol. 92, no. 8, pp. 1-8, October 2015.

[2] M. Fiorani, P. Monti, B. Skubic, J. Mårtensson, L. Valcarenghi, P. Castoldi, L. Wosinska, ”Challenges for 5G Transport Networks”, IEEE Proc. of Conference on Advanced Networks and Telecommunication Systems (ANTS), December 2014.

[3] P. Veitch, M.J. Mc Grath, V. Bayon, ”An Instrumentation and Analytics Framework for Optimal and Robust NFV Deployment”, IEEE Communications Magazine, vol. 53, no. 2, pp. 126-133, February 2015.

[4] C.J. Bernardos, A. De La Oliva, P. Serrano, A. Banchs, L.M. Contreras, Hao Jin, J.C. Zuniga, ”An Architecture for Software Defined Wireless Networking,” IEEE Wireless Communications, vol. 21, no. 3, pp. 52-61, June 2014.

[5] P. Ohlen, B. Skubic, A. Rostami, Z. Ghebretensae, J. Mårtensson, K. Wang, M. Fiorani, P. Monti, L. Wosinska, ”Data Plane and Control Architectures for 5G Transport Networks”, IEEE Proc. of European Conference on Optical Communications (ECOC), pp. 1-3, September 2015.

[6] A. Rostami, K. Wang, Z. Ghenretensae, P. Ohlen, B. Skubic, ”First Experimental Demonstration of Orchestration of Optical Transport, RAN and Cloud based on SDN”, IEEE/OSA Proc. of Optical Fiber Communication Conference (OFC), pp. 1-3, March 2015.

[7] R. Munoz, et al., ”Transport Network Orchestration for End-to-End Multilayer Provisioning Across Heterogeneous SDN/OpenFlow and GMPLS/PCE Control Domains”, IEEE/OSA J. of Lightwave Technology, vol. 33, no. 8, pp. 1540-1548, April 2015.

[8] P. Iovanna, F. Ubaldi, F., F. Di Michele, J.P. Fernandez-Palacios Gimenez, V. Lopez, ”E2E Traffic Engineering Routing for Transport SDN”, IEEE/OSA Proc. of Optical Fiber Communications Conference (OFC), pp. 1-3, March 2014.

[9] P. Iovanna, A. Germoni, F. Testa, G. Cossu, V. Lopez, R. Sabella, ”Multilayer Control for Packet-Optical Networks”, IEEE/OSA J. of Optical Communications and Networking, vol. 5, no. 10, pp.86-99, October 2013.

[10] A. Autenrieth, T. Szyrkowiec, K. Grobe, J.P. Elbers, P. Kaczmarek, P. Kostecki, W. Kellerer, ”Evaluation of Virtualization Models for Optical Connectivity Service Providers”, IEEE Proc. of Optical Network Design and Modeling (ONDM), pp. 264-268, May 2014.

[11] D. Cecarelli, Y. Lee, ”Framework for Abstraction and Control of Transport Networks”, Internet Engeneering Task Force (IETF), Fremont, CA, USA, IETF draft, Mar. 2015.

[12] M. Fiorani, A. Rostami, L. Wosinska, P. Monti, ”Transport Abstraction Models for an SDN-Controlled Centralized RAN”, IEEE Communications Letters, vol. 19, no. 8, pp. 1406-1409, August 2015.

[13] D. Drutskoy, E. Keller, J. Rexford, ”Scalable Network Virtualization in Software-Defined Networks,” IEEE Internet Computing, vol. 17, no. 2, pp. 20-27, March-April 2013.

[14] P. Ohlen, B. Skubic, A. Rostami, M. Fiorani, P. Monti, Z. Ghebretensae, J. Mårtensson, K. Wang, L. Wosinska, ”Data Plane and Control Architectures for 5G Transport Networks”, IEEE/OSA J. of Lightwave Technology, available online, 2016.

[15] J. Bartelt, P. Rost, D. Wubben, J. Lessmann, B. Melis, G Fettweis, ”Fronthaul and Backhaul Requirements of Flexibly Centralized Radio Access Networks,” IEEE Wireless Communications, vol. 22, no. 5, pp. 105-111, October 2015.

[16] M. Fiorani, B. Skubic, J. Mårtensson, L. Valcarenghi, P. Castoldi, L. Wosinska, P. Monti, ”On the Design of 5G Transport Networks”, Springer Photonic Network Communications, vol. 30, no. 3, pp. 403-415, December 2015.

[17] S. Zhang, M. Xia, S. Dahlfort, ”Fiber Routing, Wavelength Assignment and Multiplexing for DWDM-centric Converged Metro/Aggregation Networks, IEEE Proc. of European Conference of Optical Communications (ECOC), pp. 13, September 2013.