The Role of SDN in Application Centric IP and Optical Networks

Victor Lopez, Jose Manuel Gran, Juan Pedro Fernandez-Palacios
Telefónica I+D/GCTO, Madrid, Spain

Domenico Siracusa, Federico Pederzolli
CREATE-NET, Trento, Italy

Ori Gerstel, Yona Shikhmanter
Sedona Systems, Raanana, Israel

Jonas Mårtensson, Pontus Sköldström
ACREO, Kista, Sweden

Thomas Szyrkowiec, Mohit Chamania, Achim Autenrieth
ADVA, Germany

Ioannis Tomkos, Dimitrios Klonidis
Athens Information Technology (AIT), Athens, Greece

Abstract— **Transport IP/optical networks are evolving in capacity and dynamicity configuration. This evolution gives little to no attention to the specific needs of applications, beyond raw capacity. The ACINO concept is based on facilitating applications to explicitly specify requirements for requested services in terms of high-level (technology agnostic) requirements such as maximum latency or reliability. These requirements are described using intents and certain primitives which facilitate translation to technology specific configuration within the ACINO infrastructure. To support this application centric approach, SDN must have a key role in this evolution. There are representative case studies where SDN gives an added value when considering not only the network but also the application layer.**

Keywords—Application-Centric; formatting; style; styling; insert (key words)

I. INTRODUCTION

The Internet has evolved over time into a three tier infrastructure: the top tier constitutes of applications driving the traffic and ultimately the requirements of the lower layers. These applications can be consumer applications such as video, audio, gaming, file-sharing, communication, social networking, consumer cloud access etc. [1], or business applications such as backup or inter-site connectivity or various datacenter-to-datacenter interactions, such as distributed search or VM migration. The application's traffic usually passes through a grooming tier, typically IP/MPLS, which aggregates multiple small flows into larger "pipes" that can be cost-effectively supported by the underlying optical transport tier. Additional grooming of IP traffic can also be performed on OTN before offloading traffic onto optical networks.

Traffic grooming is effective in maximizing capacity utilization and reducing management complexity. However, mapping a large number of small flows, belonging to different applications, into a small number of very large and static lightpaths means that specific application requirements, such as latency or protection constraints, are seldom guaranteed after the grooming process. While some of the requirements may be satisfied implicitly by the configuration of the infrastructure, application agnostic grooming is an obstacle to effective service fulfilment.

The ACINO (Application Centric IP/Optical Networks Orchestration) project [2] proposes a novel application-centric network concept, which differentiates the service offered to each application all the way down to the optical layer, thereby overcoming the disconnect that the grooming layer introduces between service requirements and their fulfilment in the optical layer. This allows catering to the needs of emerging medium-large applications, such as database migration in datacenters. To realize this vision, ACINO will develop an orchestrator, which will expose the capability for applications to define service requirements using a set of high level primitives and intents, and then performs multi-layer (IP and optical) planning and optimization to map these onto a multi-layer service on the network infrastructure. The orchestrator also targets the re-optimization of the allocated resources, by means of an application-aware in-operation planning module.

This work presents case studies where SDN gives an added value when considering not only the network but also the application layer. The rest of the paper is structured as follows: section II presents a high level explanation of the application requirements. Section III identifies the requirements for a packet/optical SDN network orchestrator. Section IV details the case studies, where a SDN orchestrator provides an added-value to applications. Finally, Section V concludes this paper.

II. APPLICATION REQUIREMENTS

Applications requirements can vary depending on the service nature. Typical services are satisfied when there is enough bandwidth for the communication. Nevertheless, there are specific applications that may require further parameters, such as: the maximum latency, the service duration, the level of protection needed, the maximum downtime, encryption, multiple connections for the same application or even diverse routes.

The packet layer transports the traffic of multiple applications by aggregating their traffic onto optical connections. This mapping is done based on the destination address, but it is coarse by nature, as the traffic is not treated according to the requirements of the application that generates it. The main reason to use the network in this way is that the granularity (and cost) of the optical connections is in the order of tens or hundreds of Gigabits per second, but the actual traffic generated by applications is typically one or more orders of magnitude smaller. However, two major trends are driving the change towards a different approach: on the one hand the required bandwidth for the applications is dramatically increasing year over year and there are business applications, like datacenter to datacenter [3], which are not anymore of a magnitude order smaller than the optical connections. On the other hand, the optical layer has created mechanisms to adapt its granularity and offer to the services a more accurate bandwidth allocation [4].

Based on these premises the idea behind ACINO is to overcome this coarse mapping by placing application-specific traffic flows directly into dedicated optical services, or, at the very least, to groom together a number of application flows with similar requirements into a specific optical service. In this manner, each application would benefit from having a transport service tailored to its specific requirements.

From a high-level perspective, this approach requires control solutions for transport networks that (a) enable applications to express their specific requirements and (b) are able to configure and reserve network resources to create a service that treates properly the application based on its requirements.

III. SUPPORT FOR PACKET-OPTICAL NETWORKS

There are two main challenges when deploying services in packet-optical networks: (1) heterogeneous control planes and (2) different transport technologies inside each layer. A multi-layer SDN approach was proposed in [5] to address the same issue (Fig. 1). The *orchestrator* is in charge of end-to-end connectivity provisioning, using an abstracted view of the network and also covers inter-layer aspects. Each layer has a separate *controller* that is responsible for the configuration of its own technology. Every controller knows the vendor-specific details of its own underlying products and technologies, so each vendor can optimize transmission performance across the optical layer. Furthermore, the optical layer technology does not have to be the same across different optical domains. One domain can have integrated OTN switching capabilities while another domain may use WDM or even flexgrid optical switching. The only important fact to the orchestrator is that the controller offers four key services: (i) Provisioning, (ii) Topology Discovery, (iii) Monitoring and (iv) Path Computation or can easily integrate to an external application that provide such services using a simple API.

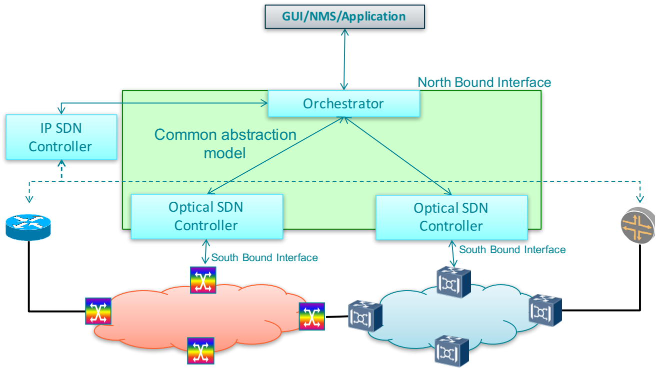

Fig. 1. Multi-layer SDN architecture

Provisioning capability enables the creation, deletion and update of connections in the network. However, to cope with the application requirements, the capability must support explicit routes, route restrictions, service resilience and traffic engineering parameters such as bandwidth and latency. Topology Discovery must export the topology information as well as the resource occupation to verify that a new service can support the application needs. The discovery of the routers and optical devices is also part of this topology discovery function. Monitoring capabilities are important so the multi-layer orchestrator can perform resilience actions that can not be solved by a controller locally. For instance, after a failure in one domain, the orchestrator may request another connection using a second domain. Path Computation is a fundamental characteristic that allows the orchestrator to analyze candidate paths and carry out "what if" analysis. An orchestrator with its global view can optimize end-to-end connections that individual controllers cannot configure.

Using these interfaces, the multi-layer orchestrator can perform the same operations as single or mono vendor controllers but in a multivendor fashion. A differentiation point of the SDN approach in comparison with the management approach [6] is the use of standard interfaces that provides the orchestrator with a vendor agnostic view of the network resources. This would allow to carry out multi-layer restoration operations like multi-layer re-route, which allows sharing a free IP port in each node to recover from any interface failure in the router, or Multi-layer Shared Backup Router (MLSBR), which consists of having extra-shared backup routers to restore the traffic in case of a failure of an IP router [7].

IV. CASE STUDIES

The ACINO consortium has selected some applications and relevant network operations to illustrate, with case studies, the value proposition of SDN in application centric IP/optical networks.

A. *Application-based DataCenter Interconnection*

An initial case study for the ACINO approach is that of a large network-facing application requiring a service with specific characteristics. For example, one business application may need to migrate VMs according to a "follow the sun" approach, which entails regular schedulable large-sized, short lived network flows for which latency is not paramount. Another application, consisting of a distributed, synchronously updated DB, may require a constant low bit rate connection with minimal latency. The owner of a third application may suddenly decide to move all related VMs to a different DC, for example because the current ones lacks infrastructure to support future expansion. This is a one-time, possibly schedulable event involving the bulk transfer of a large amount of inter-dependent VMs, and would therefore need a trade-off between bandwidth and duration, with an emphasis on the former.

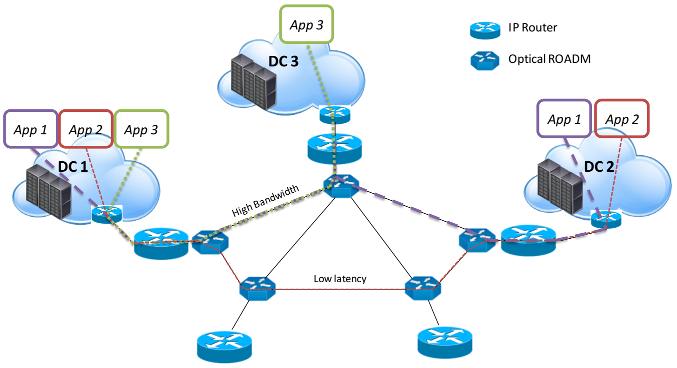

Fig. 2. Application centric Intent-based IP optical orchestration.

Assuming that these applications are housed in the same DC and the first two need connectivity to the same external DC, while the third to a different DC, currently the first two would be mapped to the same service (with shared characteristics), potentially sharing an IP TE link with the service used for the third application. Using the ACINO approach, these applications (or some human operating on their behalf) could explicitly specify their requirements to the border router(s) of the DC, which would contact the ACINO orchestrator to map them in appropriate services. Let's assume that the two DCs are already connected using a L3 VPN with no special requirements save some guaranteed bandwidth, as shown in Fig. 2.

The first application just requires large amount of bandwidth periodically. The orchestrator could decide to periodically set up another, dedicated, WAN link to satisfy the application, or it could reserve extra bandwidth on the existing VPN link (e.g. if the underlying optical connections are overprovisioned with respect to configured tunnels), or even do nothing at all if the baseline service is deemed sufficient.

The second application requires a constant connection with low latency. If the baseline service is not already configured as an IP adjacency served by one (or more) optical connection(s) on the shortest path, the orchestrator would set up such a service and instruct the border router to direct traffic from that application over it.

Finally, the third application requires a large amount of bandwidth on a short notice. Unless the baseline service is largely overprovisioned, the orchestrator might temporarily assign more capacity to it, by leveraging available bandwidth in the optical layer or possibly instantiating new temporary optical connections. Furthermore, since two datacenters need to be connected, the orchestrator may decide to share a single optical adjacency for both connections up to an IP router near one of the datacenters, and simply re-send the traffic destined for the other DC to the optical layer to be carried on a second optical connection.

B. *Enabling dynamic 5G services*

Dense deployments of small cells, often referred to as ultra-dense networks (UDN), will be a major building block for increasing radio access capacity in 5G. The scenario considered here is a UDN deployed in a location where a large crowd of people is gathered in a relatively small area for limited time periods, where all the small cells are switched off most of the time. One example is a stadium where the UDN is active for the duration of a live event, such as a football match or a concert. A stadium UDN deployment would benefit from the ACINO approach through the ability to dynamically request and reserve backhaul transport capacity for when the UDN is active.

A key vision for 5G is that the network should support a high variety of applications with very diverse requirements. For this reason, the 5G service deployment concept of network slicing has emerged, where one physical infrastructure is shared between all these applications by introducing multiple virtual and logically separated instances spanning all network domains, including the transport network. The ACINO approach would enable extending this differentiated 5G application treatment to the optical layer.

Fig. 3 illustrates how the ACINO transport network could support the stadium scenario. A mobile edge cloud is located at or close to the stadium, running selected distributed 5G user plane (UP) or control plane (CP) network functions depending on network slice. From this edge cloud, multiple transport network connections are set up to support the different network slices and applications.

Fig. 3 depicts three application examples with different requirements. One application, expected to dominate the traffic volume in the stadium scenario, is upload/download of videos and photos from members in the audience to/from a central cloud (e.g. YouTube). This application would have high requirements on bandwidth but less focus on latency and reliability. Another application is the remote control of cameras filming the event and transmission of their video streams to a central production studio. This application would have requirements on low latency and high reliability. The third application is connecting a number of different environmental sensors at the stadium to a central location with low requirements on bandwidth and no strict latency requirements.

Note that in Fig. 3, even though all three applications enter and exit the ACINO transport network at the same nodes, they

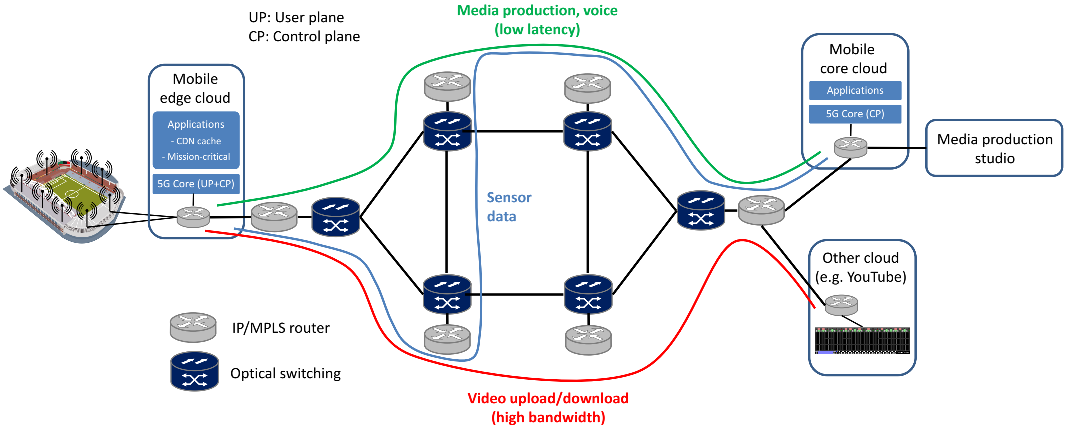

take different paths through the network. Of course, all traffic does not have to pass through the same exit node; more latency sensitive applications may e.g. be directed to a more local datacenter, and some applications may even run on servers in the edge cloud located at the stadium.

C. *Application-specific protection strategies*

Optical layer restoration and multi-layer restoration have been researched extensively and it can account for very substantial savings in the total number of required router interfaces and transponders, on the order of 40–60 percent in the core [9]. In both cases, it is assumed that some of the responsibility for restoring from failures is moved to the optical layer, since it is much more cost effective to build in spare capacity in the optical layer than to do so in the IP layer. The former approach assumes that the optical layer alone is responsible for restoring from a failure – and therefore the selection of the restoration path is insensitive to the needs of the IP layer. The latter approach does take IP layer needs into account, but for the aggregate traffic that traverses the failed IP links.

An example for the behaviour of a latency-sensitive traffic can be found next. In this example, we assume all optical links have the same length and that a service can tolerate a latency of 4 optical links. Fig. 4a shows a service routed over this network in green (on the left), and its routing over the network after the optical recovery from a failure, assuming it takes the same IP layer path. This is acceptable for non-latency-sensitive traffic, however if the max latency is 4 hops, then the IP layer should route the service over a different IP path – as shown in Fig. 4b. If the protection would be only done at the optical layer the service would be rerouted as shown in Fig. 4c. However, there is an increment on the delay as the service has to cross the whole ring. Therefore, the ACINO orchestrator will perform an application-specific restoration to provide each application with the required resilience mechanism.

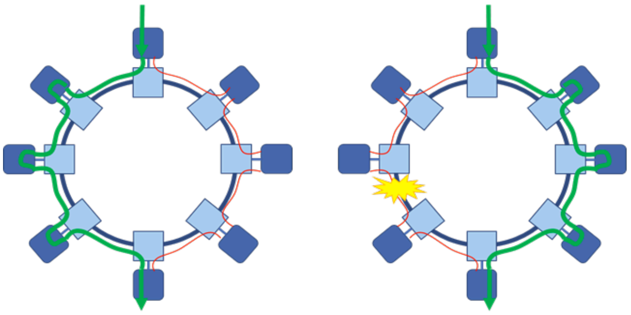

(a) Normal model (b) IP restoration

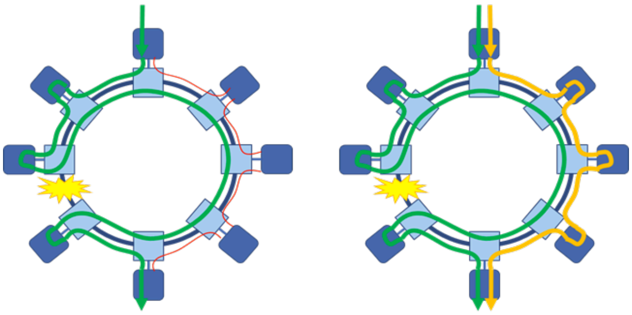

(c) Optical restoration (d) App specific restoration

Fig. 4. Example for restoration under application latency constraints.

D. *Secure Transmission as a Service*

The frequency of cyber-attacks on critical network infrastructure and public/private entities connected to the Internet has been growing significantly and has translated into significant financial costs for end-customers or businesses. As a result, more and more businesses are investing in improving their security infrastructure. Moreover, server-to-server communication, especially datacenter interconnects (DCI) are increasingly becoming an important service offering for network operators.

One of the most common trends for securing the communication has been to push for end-to-end encryption (IPSec, HTTPS). End-to-end encryption is flexible and independent of the underlying network infrastructure, making it relatively easy to deploy. However, the flexibility comes at the cost of increased processing requirements at both server and client endpoints, increased latency and reduced throughput from the network. Shifting the responsibility to the network, reduces the processing complexity at the endpoints (servers) and allows network operators to optimize the $/bit cost for encrypting traffic between two remote sites. Different mechanisms such as IPSec, IEEE MacSec, and custom all-optical encryption differ on the cost of deployment, availability, latency and throughput. The selection of the best encryption mechanism for the applications is a feature that the ACINO orchestrator will support.

E. *Dynamic Virtual CDN deployment*

Content delivery networks (CDN) are used to deliver content from servers located in datacenters (DCs) to end users. In this case, we will consider video distribution as a reference application for the CDNs. Videos are delivered by using the IP layer. Therefore, the choices for an operator’s CDN location are limited to the IP core. Possible deployment sites include regional DCs at the access routers (AR), high density DCs at the transit routers (TR) or even national DCs at the interconnection level (IX). Deployment of video applications only on National CDN has the advantage of high utilization of the video servers, since users access the same datacenters, thus statistically using the resources more often. However, each connection or video between the user and one of those national video servers would pass the whole national network. As the data flows are unicast for video on demand or time shifted services, the (redundant) overhead in the network would be massive. On the other hand, the video servers deployed at the AR locations minimize the network overhead, but increase the video platform’s costs by increasing the number of locations and redundant copies of the same content. Moreover, if the number of customers using these video servers is small, the dimensioning and caching is not efficient.

As the traffic nature is dynamic (events in stadiums, popular contents in regions, unexpected high penetration…), the use of a virtual CDN infrastructure make sense. The virtual CDN infrastructure has the same capabilities as current CDN deployments, but it runs in standard x86 servers. This way the vCDN provider can dynamically activate VMs in locations that require more video servers (Fig. 5). Once the video server is activated, the contents can be transferred from a national/regional datacenter to sync the most popular content. This requires an inter-datacenter transfer, which will require network resources for a one-time sync between the fixed and the virtual CDN site. This approach can reuse the investment in datacenters that operators are doing for virtualizing services or even Virtual Network Functions (VNFs).

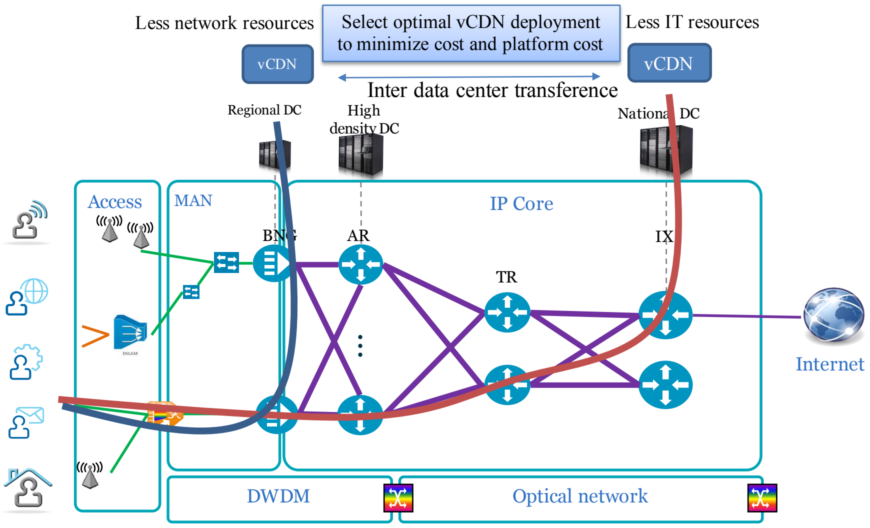

Fig. 5. Virtual CDN scenario.

F. *Application-centric in-operation network planning*

Current transport networks are statically configured and managed, because they experience rather limited traffic dynamicity. This leads to long planning cycles to upgrade the network and prepare it for the next planning period. The planning procedure aims that the network can support the forecast traffic and deal with failure scenarios, thus adding extra capacity and increasing network expenditures. The second drawback of current static procedure is that the results from network capacity planning are manually deployed in the network, which limits the network agility. The main reason for this is that current provisioning systems are not always deployed as the number of configurations in IP/optical backbone networks is relatively small. The authors in [9] proposed the term “in-operation network planning”. The main idea of in-operation network planning is to take advantage of novel network reconfiguration capabilities and new network management architectures to perform in-operation planning, aiming at reducing network CAPEX by minimising the over-provisioning required in today’s static network environments.

Within ACINO the in-operation planning concept is extended and includes application awareness. This denotes that overall planning of the resources is not performed solely for the optical infrastructure considering the aggregated data from the upper layer, but on a ‘per application demand’ basis considering both the IP and Optical layer resources. With the ACINO dynamic multi-layer approach, the incoming requests are classified in a sense that they generate different service requirements to the network which translate to different use of the available resources. These requirements are evaluated in real-time providing the optimal routing through IP and optical domain or establishing new IP or/and optical lightpaths if this is required.

V. CONCLUSIONS

This papers details the case studies where SDN can play an important role for Application Centric IP and Optical Networks. It presents an analysis of the application requirements as well as the key elements to support the network operations for packet-optical scenarios. As future work, the authors will analyze each case study to validate the improvement of the application centric approach in packet-optical networks.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Commission within the H2020 Research and Innovation program, ACINO project, Grant Number 645127, www.acino.eu.

REFERENCES

[1] Sandvine Global Internet Phenomena report, 2H 2013, https://www.sandvine.com/downloads/general/global-internet-phenomena/2013/2h-2013-global-internet-phenomena-report.pdf

[2] ACINO H2020 EU Project, http://www.acino.eu/wp-content/uploads/2015/02/ACINO_Public-Project-Presentation.pdf

[3] A. Mahimkar et al., “Bandwidth on demand for inter-data center communication”, Proceedings of the 10th ACM Workshop on Hot Topics in Networks, HotNets-X, 2011.

[4] O. Gerstel, “Flexible Use of Spectrum and Photonic Grooming”, Photonics in Switching 2010.

[5] O. Gerstel, V. López: The Need for SDN in Orchestration of IP over Optical MultiVendor Networks, in European Conference on Optical Communication (ECOC), September, 2015.

[6] A. Martínez, M. Yannuzzi, V. Lopez, D. López, W. Ramírez, X. Masip-Bruin, R. Serral-Gracià, M. Maciejewski, J. Altmann: Network Management Challenges and Trends in Multi-Layer and Multi-Vendor Settings for Carrier-Grade Networks, in IEEE Communications Surveys and Tutorials, June 2014, Issue 99.

[7] A. Mayoral, V. López O. Gerstel, E. Palkopoulou, J. P. Fernández-Palacios and Ó. González de Dios: Minimizing resource protection in IP over WDM networks: Multi-layer Shared Backup Router , in Proc. Optical Fiber Conference (OFC), M3B.1, Mar 2014.

[8] O. Gerstel, C. Filsfils, T. Telkamp, M. Gunkel, M. Horneffer, V. Lopez and A. Mayoral:Multi-Layer Capacity Planning for IP-Optical Networks, in IEEE Communications Magazine Feature Topic "Advances in Network Planning", January 2014, Vol. 52, Issue. 1, pp. 44-51.

[9] L. Velasco, D. King, O. Gerstel, R. Casellas, A. Castro, and V. Lopez: In-operation Network Planning, in IEEE Communications Magazine Feature Topic "Advances in Network Planning", January 2014, Vol. 52, Issue. 1, pp. 52-60.