# Multi-layer orchestration for application-centric networking


Ori Gerstel
Sedona Systems, Raanana, Israel
orig@sedonasys.com

Victor Lopez
Telefónica I+D/GCTO, Madrid, Spain
victor.lopezalvarez@telefonica.com

Domenico Siracusa
CREATE-NET, Trento, Italy
domenico.siracusa@create-net.org



***Abstract*—We argue that the implementation of services in an IP-optical network should be driven by the needs of the specific applications, and explain why this requires a centralized orchestration architecture.**



***Keywords—multi-layer networks, SDN, application-centric, IP-optical netowrks, centralized control.***


## I. INTRODUCTION

Transport IP/optical networks are evolving in two main directions: achieving very large capacity for transmitting unprecedented amounts of data between optical devices, and providing cost-effective high-speed services, dynamically created in response to requests from high-bandwidth users via novel control plane solutions. To enable the former, technologies like superchannels and FlexGrid are used. To pursue the latter direction, new mechanisms are required to abstract and partition network resources, in order to provide them as a service. The need for such an on-demand service is driven to a large extent by increasingly complex data center systems, and the desire of service providers to participate in the value chain for this new world and move away from becoming dumb pipe providers.

However, this evolution gives little to no attention to the specific needs of applications, beyond raw capacity. Some applications – like large scale offline backup – can tolerate high latency and even service disruption, while other applications – like synchronous data replication are extremely sensitive to both latency and service disruption. Currently, the service creation process multiplexes heterogeneous traffic from applications with different requirements onto a unified IP layer; consequently, all applications are treated in the same way by the optical layer, without any differentiation. In fact, their requirements are actually hidden by the grooming process of the packet layer, which aggregates traffic received from applications into flows that are blindly transported by the optical layer through big optical pipes. To the best of our knowledge, no solution has been proposed so far to manage transport network resources on a per-application basis, providing a thoughtful differentiation mechanism.

In this paper, we introduce a novel approach for the creation of application-centric network services, by orchestrating the IP and the optical layers in a joint fashion, yielding a network-wide optimized configuration that matches the needs of each application all the way down to the optical layer. To this end, we exploit the flexibility provided by FlexGrid technology, which allows such services to carve out a small portion of the optical capacity and manage it independently. We initially discuss the proposed approach and then we focus on one of its key enablers: the joint orchestration of IP/MPLS and optical resources.


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


## II. APPLICATION REQUIREMENTS

Applications requirements can be very different in nature. Most often, the bandwidth required for satisfactory communication is the main parameter; however, specific applications may also have further requirements, including:

- The maximum latency the application can tolerate
- The duration of the required service
- The need for network-level protection against failures
- The maximum acceptable downtime in case of a failure
- The need for encrypted communication
- Multiple concurrent connections among two or more sites
- Diversely routed services that will not fail together.

As mentioned earlier, the IP layer usually aggregates the traffic generated by multiple applications by their destination address, thus mapping many heterogeneous services into a single transport link, served by means of an optical connection. Such inaccurate mapping has been hitherto tolerated essentially because there was no better way of mapping the IP traffic into optical lightpaths without severely affecting the efficiency of the underlying optical layer. This has been dictated by the large gap between the minimum granularity handled by the optical layer (tens or hundreds of Gigabits per second) and the actual traffic generated by applications, which in most cases is one or more orders of magnitude smaller. However, there are two major trends that justify the pursuit of a different approach: on one side the bandwidth required by applications is increasing year over year (it will suffice to think about the capacity needed for high-quality videos). This is particularly relevant for business applications, such as data center to data center communications, e.g. virtual machine migration, replication or indexing [1], which are typically in the order of many Terabytes. On the other side, the granularity of services that can be offered by the optical network can now be finer, due to recent advances in physical layer technologies such as sliceable transponders [2].

Given these premises, a novel approach called Application-Centric IP/Optical Network Orchestration (ACINO) [3] proposes 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 specifics requirements and (b) are able to configure and reserve network resources to create a service that exactly matches those requirements.

## III. Multi-layer network operations

The rest of the paper focuses on the second requirement. Once the specific needs of the applications are somehow known, the most appropriate configuration for both the IP and the optical layer must be chosen. Such a multi-layer approach is also needed because even the most basic operation can yield different network behaviors depending on which layer it is applied to. For example, it would stress different resources, or result in different costs or energy consumption. Thus, multi-layer visibility enables the provisioning of specialized per-application treatment discussed thus far. For instance, while a simple capacity requirement could be satisfied by an appropriate combination of optical circuits and logical adjacencies at the IP layer, the survivability requirement of an application may be fulfilled by means of either IP protection or optical restoration, or even a combination of the two. This choice, in turn, influences the latency perceived by the application in case of failures, which, if expressed as a requirement, can be used to guide the selection of the most appropriate alternative. In simple words: the ability to assign different operations to different layers is paramount for realizing highly optimized application-centric services.

To demonstrate how different application needs drive to different solutions at both the packet layer and optical layer, consider the simple ring network in Fig. 1(a). The network consists of 8 sites, each containing an optical node and an IP node and connected in a ring as shown by the thin red lines. Let's assume the distances between all nodes are the same, and focus on a latency sensitive service (in green), which can tolerate 4 fiber spans maximum. Please note that in large scale networks the latency is dominated by light propagation and not by router latency, which is actually very small in core routers. Moreover, let us assume that this service is protected by the IP layer. In normal conditions, the service is routed over the "left" side of the ring, as shown by the green line in Fig. 1(a). When a failure occurs, the service is routed over the other side of the ring as shown in Fig. 1(b) – still meeting its latency constraint. However, let's assume that there is a desire to use optical restoration in this network due to its considerably lower cost. This would cause the service to be routed as in Fig. 1(c), violating its latency constraint. This does not imply that the service provider must give up the cost savings promised by optical restoration since other services, which are not as latency sensitive, could still use such restoration. The ACINO solution: in this case, let the latency sensitive applications use IP protection (the orange line in Fig. 1(d)) while services for less sensitive applications can use optical restoration – as the green line in Fig. 1(d).

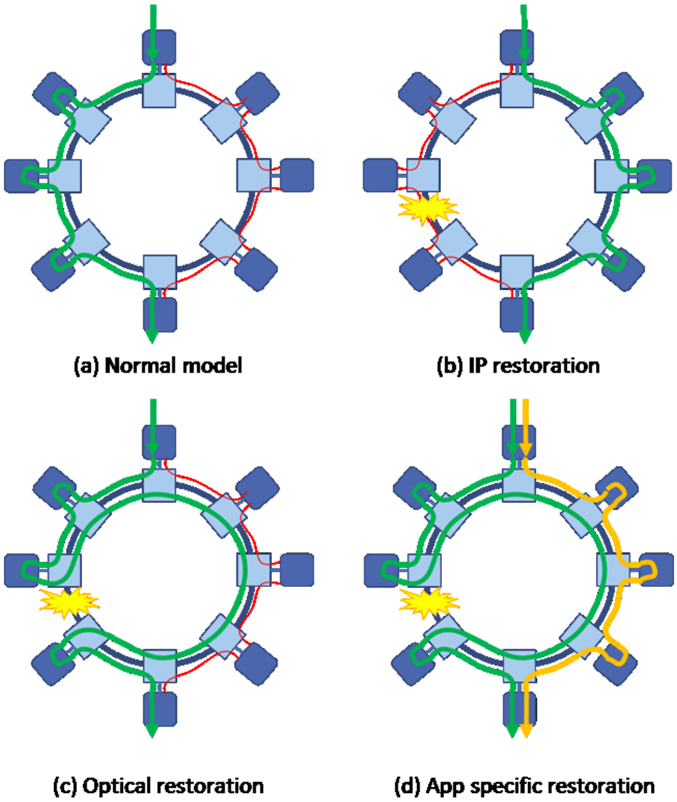


Fig. 1. Example for restoration under application latency constraints.

Clearly, real networks are more complex and the decision of which services should use which restoration approach depends on the specific routing of the service. This is essentially a new variant of the known joint IP/optical network optimization problem for which there are many known techniques in the literature. The key challenges that are currently taken into account relate to the development of efficient multi-layer path computation schemes [4]. The multi-layer optimization problem is still under active investigation, however, considering the high complexity involved, typically the model for the IP layer is not very accurate. Among the works proposing multi-layer solutions, [5] tries to optimize the explicit routing of demands in the IP/MPLS and optical layers, while [6] discuss about enhanced resilience mechanisms. Authors of this paper have already been active in this field: for instance, [7] investigates the impact of IP layer routing policies on multi-layer network design. Nevertheless, to the best of our knowledge, there is no work that proposes an application-centric multi-layer optimization approach, which is part of the ongoing investigation of the ACINO project.

Once a new multi-layer network configuration is computed, it must be implemented in the actual network. This requires orchestrating the underlying resources. A possible framework for orchestrating packet and optical networks is described by the ABNO architecture [8], which combines a number of technologies already under standardization within IETF and demonstrated in a multi-layer scenario [9].

## IV. The Need for a Multi-Layer Orchestrator

In ACINO, we take the realistic assumption that there will be two or more distinct controllers for IP and optical devices, typically (but not necessarily) built by equipment vendors, leveraging different control platforms (OpenDaylight and ONOS are two promising examples). In accordance to SDN principles, these controllers should expose their northbound interface (NBI) to a common orchestrator, which would receive requests directly from applications and would provision the required services and carry out multi-layer optimization. This is demonstrated in Fig. 2.

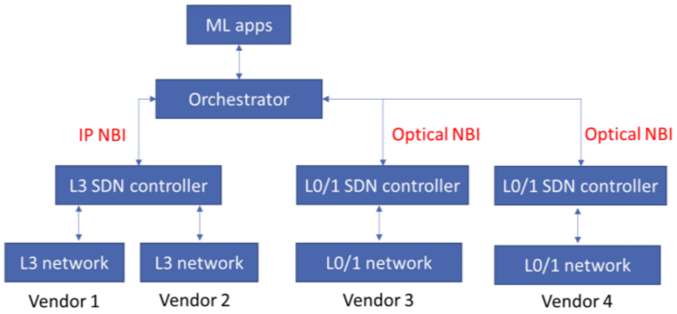


Fig. 2. Control architecture.

It should be noted that we do not see a likely convergence of controllers for optical networks since each vendor has their own specific way to account for impairments. At the IP layer, however, we believe that a single controller will suffice to control all routers. Such multi-vendor support has been the goal of most IP layer controllers and is in line with the standards-based collaboration of equipment at this layer.

To achieve application-centric control, the orchestrator is expected to interact with network controllers operating at different layers, as well as with the users of the service. To this end, a Control Orchestration Protocol (COP) that abstracts a common set of control plane functions used by various SDN controllers, allowing the interworking of heterogeneous control plane paradigms, has been presented in [10].

An early demonstration of the above hierarchical architecture has been discussed in [11, 12]. The demonstration uses Cisco and Juniper commercial core routers and optical gear from Infinera, Cisco, co-located with the routers at the Telefonica lab. Each optical vendor provided their SDN controller, and a single controller was used for the IP layer as per the above figure. We used Sedona Systems' Multi-Layer application platform (MAP) [13], to interact with the various controllers, each with its own NBI. MAP had a separate driver to adapt the vendor specific details to a common network model and towards several control applications, such as Sedona's ML visibility application (MLV) and its optimization application (MLO) to optimize the IP network.

The demonstration showed how the network reacts to an increasing load on IP links (simulated by increasing the bandwidth reservation of an MPLS tunnel), by automatically adding a new IP link between routers – instead of the trivial (but sub-optimal) solution of upgrading the capacity of the congested link. The demonstration did not yet involve users driving services in an application-centric manner, and this is left for further investigation of ACINO.

## V. Conclusions And Future Work

In this paper we motivated the need for application-centric control over both the IP and the optical layers in service provider networks, namely, the ACINO approach. We also discussed a realistic architecture for such a multi-layer control system and referred to an early demonstration of this architecture. Much work is still needed in developing the interface that allows applications to specify their specific needs and in orchestration algorithms that can translate these needs into an implementation that optimally fits the network. In particular, we envision that there will be a need to group lower bandwidth services into higher bandwidth connections in the optical layer, but in a way that respects their needs and constraints. Significant work is also needed in developing optimization algorithms that can take a set of applications with disparate needs and map them onto network functions – such as the IP layer and optical layer – in a cost-effective manner.

## References

[1] 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.

[2] O. Gerstel, "Flexible Use of Spectrum and Photonic Grooming", Photonics in Switching 2010.

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

[4] J.E. Gabeiras et al., "Is Multi-layer Networking Feasible?", in Elsevier Journal on Optical Switching and Networking, April 2009, Vol. 6, Issue 2, Pages 129-140.

[5] P. Iovanna et al., "Integration of Planning and Control Plane in Packet Optical Multilayer Network," OFC/NFOEC 2013.

[6] M. Pickavet, et al.,"Recovery in multilayer optical networks," JLT, Jan. 2006.

[7] E. Palkopoulou et al. "Impact of IP Layer Routing Policy on Multi-Layer Design", OFC/NFOEC 2014, paper W1K.4.

[8] D. King, A. Farrel, "A PCE-Based Architecture for Application-Based Network Operations", IETF RFC 7491.

[9] A. Aguado et al., "ABNO: a feasible SDN approach for multi-vendor IP and optical networks," OFC/NFOEC 2014, paper Th3I.5

[10] R. Vilalta et al., "The Need for a Control Orchestration Protocol in Research Projects on Optical Networking", European Conference on Networks and Communications (EUCNC), 2015.

[11] O. Gerstel et al. "The Need for SDN in Orchestration of IP over Optical Multi-Vendor Networks", ECOC 2015.

[12] Multi-layer/multi-vendor control demonstration, OFC15: https://www.youtube.com/watch?v=pmzdsf-oHCw

[13] Sedona Systems, Multi-Layer application platform (MAP) http://sedonasys.com/technology/