The Age of Multi-Layer Networking

Ori Gerstel
Principal Engineer; High-End Routing and Optical Group; Cisco
ogerstel@cisco.com

Abstract: Service provider networks are increasingly challenged by capacity per fiber and cost per bit, yet there is no transmission related innovation in sight that promises to alleviate these bottlenecks. We believe that the solution lies in tight collaboration between the optical layer and IP layer, allowing for much more efficient use of the available resources at both layers.

OCIS codes: (060.4250) Networks; (060.6718) Switching, circuit; (060.6719) Switching, packet

1. Challenges to the evolution of the Service Provider network

Service provider networks are in the midst of several major transitions. On one hand, traffic continues to grow at an exponential rate – around 30-50% per year globally and at a faster rate in some networks [1]. On the other hand, a growing percent of the revenues from the services are going to "over the top" (OTT) service providers, such as Google and Netflix, leaving SPs with almost flat revenues, while being forced to grow their infrastructure. This trend strains the business model of SPs, as the gap between the cost of the network and the revenues from it shrinks.

In the past, the network cost was kept reasonably low, thanks to continuous innovation in optical modulation formats, which enabling cramming more information into a GHz of spectrum on the fiber: thus, over the last 10 years, the capacity per 50GHz spectrum slice increased tenfold, from 10Gbps to 100Gbps. However, there is mounting evidence that the ability to increase the capacity per fiber will become harder as we approach Shannon's (non-linear) limit, as can be seen in the following figure.

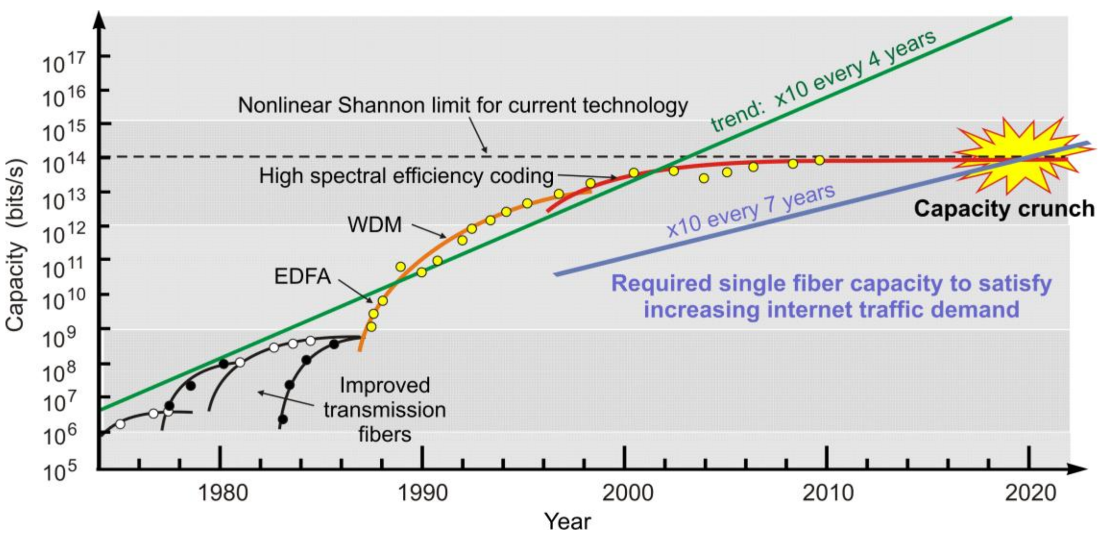

Figure 1: Historic perspective on fiber capacity [2]

The figure shows how several innovations helped keep the pace of capacity growth, but it is not clear what will be the engine of the next wave of growth. Today the upper bound is around 200Gbps per 50GHz using 16QAM modulation – and in a couple of years, as DSP processing power increases – we will achieve 400Gbps using the same modulation format in such a spectrum slice. However this seems to be the end of this wave of innovation: higher order modulation formats, such as 64QAM, will have very limited reach (less than a few hundred Km) and therefore limited applicability. As a result, the industry is turning to parallel solutions, such as superchannels (parallel, tightly spaced channels), or SDM (parallel fiber cores). While both approaches will increase the capacity per fiber (in particular SDM), they are not expected to provide significant cost reductions since the implementation complexity remain linear to the bitrate (more transceivers, instead of a single higher rate transceiver).

Raw bandwidth growth is just one of several trends that may be as challenging to the evolution of the network:

- Consumer traffic is now much larger than business traffic – skewing the required technologies towards more dynamic IP based technologies
- The number of major sources of bandwidth for this traffic is becoming much smaller. For example, in the US, Netflix traffic represents a third of the overall peak traffic. This implies larger traffic fluctuations due to failures of peering points or changes in content provider policies
- Cloud computing will make it easy to mobilize an application from one server to another based on power savings, proximity to the users and commercial considerations, further increasing traffic dynamism
- The Internet of Everything – which will turn billions of devices into active users of the internet, will have an unknown – yet surely dramatic impact – on the network.

These trends imply that it will be increasingly hard to predict traffic behavior. As a result, the network planner will have to plan for the unknown, but how does one do this without significantly over-provisioning the network and further straining the SP business model? The solution lies in a much more agile network.

2. Agile optical networking as a first step to address the challenge

To address the aforementioned trends, the optical layer will have to be:

- Streamlined: the increased pressure on SP margins implies that the future network must be as efficient as possible, and this implies removing as many network layers and the interfaces between layers as possible – in fact most core networks are already evolving to a two layers model: a transport layer and a service layer.
- Flexible: the lack of ability to forecast how traffic will evolve implies that the network will have to be as flexible as possible in providing the right amount of capacity where it is needed using the most effective modulation format. Purpose built transponders that only support one format will be ineffective.
- Reconfigurable: since traffic patterns will change more frequently, the network will have to support graceful release, redeployment, and reoptimization of resources. Without these capabilities, resource will sit idle and the cost of the network will grow well beyond the required cost.

These characteristics are a necessary but insufficient condition to achieving a true low cost and ultra-efficient solution. What good is a high degree of network agility, if at the end it is limited by existing planning practices that assume a slow and manually-intensive process to implement a change in the network? The only way to reap the benefits of this agility is to invent new planning practices that involve the service layer that drives the need for optical agility (which is typically an IP network).

3. The role of multi-layer collaboration

The IP layer must closely interact with the optical network to optimize how optical resources and IP resources are used. This will allow the network to quickly move optical capacity to where it is needed by the IP layer, instead of today's approach, of over-provisioning static IP links, to address different possible changes in traffic patterns over a static optical layer.

To better understand the difference between today's approach and the desired future approach, consider the example in Figure 2, which depicts two traffic scenarios (on the left). Each scenario implies different capacity needs from one router to three other routers (in Gbps units). These scenarios could result from different failures in the network, from unexpected traffic growth, or changes in peering arrangements. Either way, the network design must accommodate both scenarios without having to redeploy gear. Figure 2(a) show how this is done using today's static IP network over static optical network. The planner must provision each link to an adjacent router to account for the maximum capacity needed for both scenarios. In this example, this means 2x100G wavelengths to the first and second router and 3x100G wavelengths to the third router. A design with an agile network in mind is shown in Figure 2(b). Here the planner must only consider the total number of wavelengths needed to accommodate both scenarios – 5x100G wavelengths in this case, but how they are distributed amongst the neighbors is not important, since they can be redistributed quickly by the network as the traffic pattern changes. In a nutshell, the planning process is changing from sizing each link to the maximum capacity required to sizing each node to the maximum total capacity

The application that best captures the value of this new approach is multi-layer restoration. In this case, the IP layer relies on the optical layer to restore failed links in the event of an optical layer failure, using the same router interfaces and transponders [3]. This reuse of interfaces is the key reason for the significant savings achieved by the scheme – in the order of 40% of router interfaces and transponders on several real European core network models – not to mention the associated rack-space and power.

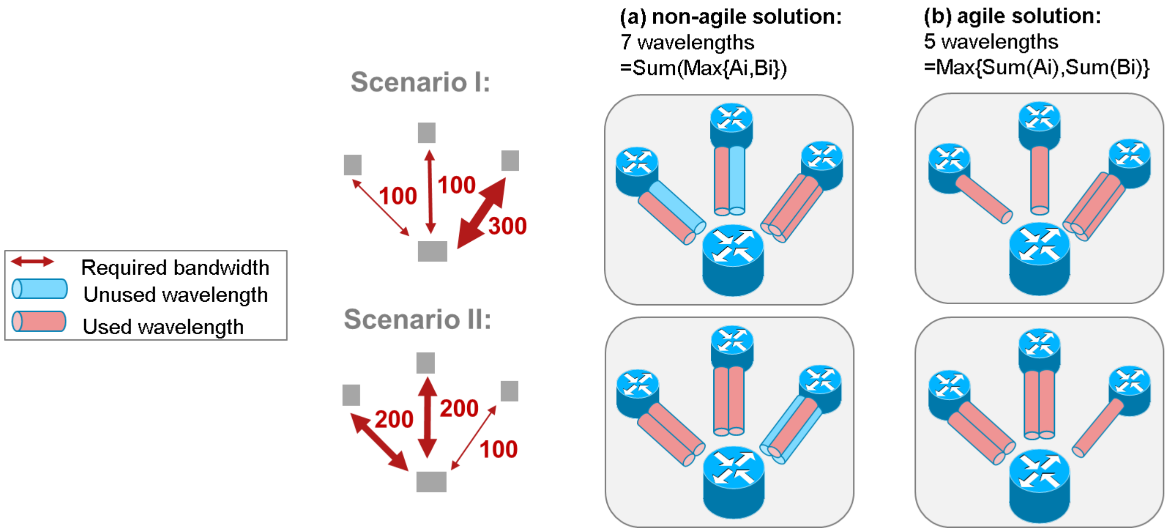

Figure 2: How network agility reduces cost

Many other multi-layer applications promise additional savings: from optical layer optimization, to automated optical bypass of routers, to disaster recovery. Some of these applications (including restoration) require fast reaction and a high degree of availability – pointing to the need for a distributed control plane between layers. Others require a high degree of sophistication in understanding how a proposed change in the network will impact the routing of traffic flows in the IP layer and the resulting impact of its service-level agreement (SLA) – pointing to the need for centralized control with a high degree of user interaction. We believe that the best architecture is a hybrid one, with both distributed and centralized control elements. In addition, a new type of network management is needed to provide the operator with sufficient information and control over such an automated network. This architecture is shown in Figure 3. Once these innovations are implemented, we believe that the pressure of service provider network will be significantly reduced.

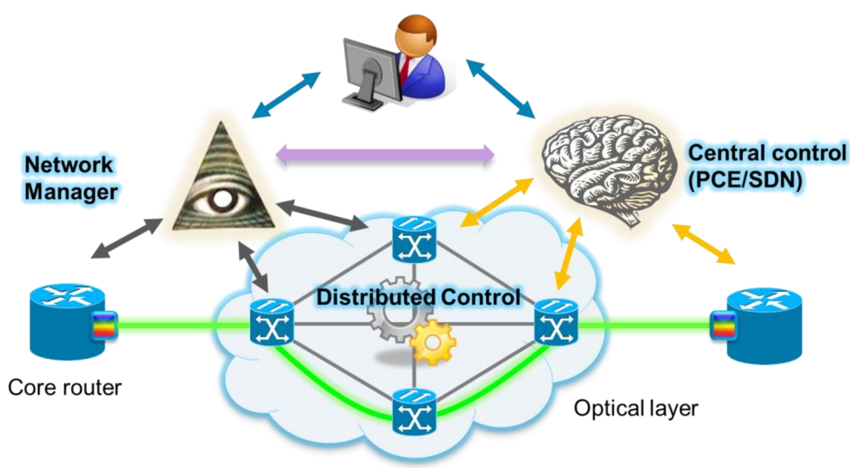

Figure 3: Overall agile network architecture

4. References

[1] Cisco VNI http://www.cisco.com/en/US/netsol/ns827/networking_solutions_sub_solution.html
[2] Courtesy of David J. Richardson, adapted from "Filling the Light Pipe", Science 15 October 2010: Vol. 330 no. 6002 pp. 327-328
[3] O. Gerstel et al., "IP-Optical Interaction during Traffic Restoration", In proc. of OFC'13, 2013.