The Optical Internet: Architectures and Protocols for the Global Infrastructure of Tomorrow

Antonio Rodríguez Moral, Paul Bonenfant, and Murali Krishnaswamy, Photuris, Inc.

ABSTRACT

We present a high-level overview of the current state of the network architectures, protocols, and technologies that will serve as the seed for the optical Internet. We further propose that a two-way ripple effect of technologies penetrating from the edge to the core and vice versa, and the associated transformations that result, represent the keys to unlock the full potential of an optical Internet.

INTRODUCTION

The term *optical Internet* has recently joined the long list of overused and omnipresent terms in the telecommunications and networking industry. It is undoubtedly a term that immediately evokes a certain irresistible aura of futuristic networks and next-generation technologies.

The Internet itself, with its many protocols and supported applications, is becoming the reference point for the multiservice network infrastructure of tomorrow. The research community and the networking industry are working at a phenomenal rate — indeed, on "Internet time" — to design and deploy the technologies, protocols, networks, and services that will make that goal a reality.

While many advertise that the optical Internet is already here, we contend that there is still a long way to go to realize such a vision. We submit that the disruptive technology in the march toward the optical Internet has been IP itself. Admittedly, innovations in optical technology (dense wavelength-division multiplexing, DWDM, and optical amplifier technology in particular) have sustained traditional networking applications and offered remarkable capacity gains in recent years; the blistering pace of technology innovation in specialty fibers, forward error correcting codes and optical layer management, Raman amplification and gain equalization, dispersion management, and optical switching technologies has ushered in the age of ultra-long-haul DWDM systems and changed the economy of long-haul optical networking. But it is the IP suite of protocols — a technology with its roots in the desktop and enterprise environment once deemed unsuitable for the transport network "mainstream" — that represents the disruptive force which is slowly but surely creeping "upmarket" [1] toward the core network, with the introduction of multiprotocol label switching (MPLS) and its extensions to multiprotocol lambda switching (MPλS)/generalized MPLS (GMPLS).

At the same time we propose that the realization of the optical Internet will take place only when the capabilities offered by optical technology, available today primarily at the core of the network, spread toward the edges, extending the optical reach as close as possible to the final user. It is this two-way ripple effect of technologies penetrating from the edge to the core and vice versa, and the associated transformations that result from such a combination, which represent the key that will unlock the optical Internet.

Another issue that will be critical in the evolution toward the optical Internet is the design of the control plane of next-generation optical internetworks. We introduce some of the concepts and protocols under consideration to enable rich services such as real-time "point-and-click" optical channel provisioning, optical layer protection and restoration, optical layer network topology autodiscovery, optical layer traffic engineering, and optical bandwidth services management. Some of these features are common in a networking arena that seemed, until now, widely disparate from optical networks: IP networks and their upcoming MPLS extensions. For this reason, the proposal for an IP-centric control plane for next-generation optical networks and optical internetworks based on MPλS, now classified under the generic umbrella name of GMPLS, has emerged as a natural next step.

The remainder of this article is organized as follows:

 0163-6804/01/$10.00

- We outline some of the current major transformations in the network infrastructure that will form the seed for the evolution to the optical Internet.
- We introduce our definition of the optical Internet, some of its distinctive features, and an argument in favor of the possibility that such a global infrastructure will be deployed sometime in the (near) future.
- We give a high-level view of some of the network architectures that will be at the heart of the optical Internet.
- We outline the service models and interconnection models that will be possible for the optical Internet.
- We dedicate some time to what we believe will be the major challenge for the optical Internet: the design and deployment of a feature-rich unifying control plane.
- We summarize, present some conclusions, and glance briefly beyond our current era.

Major Transformations in Network Infrastructure

At this point in history it seems almost obligatory to mention the explosion of traffic driven by the Internet and associated applications, such as the Web. In this section we focus on some of the major transformations this growth is forcing on the network infrastructures deployed by carriers and service providers.

Optical Transport and Switching Capacity and Router Packet Processing Capacity

From an architectural perspective, the growth in data traffic demand is dramatically increasing the rate at which bandwidth is managed in the network core. High-capacity data network elements, such as multigigabit and terabit IP routers, are now available that provide statistical multiplexing and traffic aggregation at the edge and the core network, thus reducing the need for synchronous optical network/digital hierarchy (SONET/SDH) layer multiplexing [2] within the core. As a consequence, service providers are now deploying higher-granularity digital switches (e.g., switching at SONET STS-48 or SDH STM-16 rates). However, as the capacity requirement and wavelength bit rate further increase, network elements and backbone architectures capable of managing bandwidth at the optical channel level (i.e., independent of the wavelength bit rate) will soon be needed to ensure network scalability.

For the same reasons, it is becoming increasingly efficient and economical to perform protection and restoration in the optical layer [3, 4]. In fact, a major network failure, such as a fiber cut or node failure, would impact an extremely large number of routers and IP flows, rendering traditional IP restoration or even MPLS-based restoration either intractable or too cumbersome. Optical layer restoration can be carried out at the optical channel (i.e., wavelength) level or the optical multiplex section (i.e., fiber) level. Different protection/restoration approaches are possible, and therefore different optical network architectures can be implemented (see [3] for a more detailed account).

Historically, it has proven very difficult to accurately predict emerging or growing traffic patterns in the Internet, and we believe this difficulty will persist. Network planners need some flexibility in the network to absorb this uncertainty without the need for costly redesign or locking into architectures that imply some amount of stranded capacity. In the case of optical networks, this requirement translates into the need for optical channel switching and rapid end-to-end optical channel provisioning, enabled by programmable optical crossconnects (OXCs) and optical add/drop multiplexers (OADMs). These next-generation optical networks must also be versatile because some service providers may provide generic optical layer services (e.g., wavelength leasing services) that may not be specific to any particular digital clients.

Finally, it is important to note that the inherent long distance characteristic of IP traffic [5] leads to increasing values of the ratio of pass-through vs. add/drop traffic at each node (larger than 80/20) in the inner portion of the network core. This has a direct impact on backbone network efficiency, and makes the possibility of deploying innovative core network architectures very attractive.

The Optical Internet

Next-Generation Internet: DiffServ, VPNs

After more than 20 years of tremendous success and growth, the time has come for the Internet service model to evolve.

The IP networks of the next decade will have to accommodate a global number of users that increases at unprecedented rates, and who connect at faster data rates due to the widespread use of DSL, cable modems, and future 3G mobile services. Each of these technologies has the potential to multiply by 40 the current typical access rate of 56 kb/s of analog modems. In the enterprise segment the upcoming Gigabit and 10 Gigabit Ethernet revolution promises a similar increase in traffic load, with several emergent service providers building their business models around the possibility of offering 100 Mb/s of Internet connectivity at prices comparable to those paid today for a T1 line (1.5 Mb/s).

But the next-generation Internet is not only about more bandwidth; it is about new service models and richer services. In a commercial context, the flat best-effort service provided by the network may not be the most appropriate any longer: customers requiring the assurance of better quality of service (QoS), and willing to pay more, should be able to get better service than customers paying for basic service [6].

A relatively new service architecture has been proposed to extend the set of services provided by the Internet, which is currently largely limited to best-effort service.

The differentiated services (DiffServ) architecture proposal [7] approaches the problem of QoS support from the point of view of allowing for controlled unfairness in the use of network resources. The DiffServ architecture aims at pro-

As the capacity requirement and wavelength bit rate further increase, network elements and backbone architectures capable of managing bandwidth at the optical channel level (i.e., independent of the wavelength bit rate) will soon be needed to ensure network scalability.

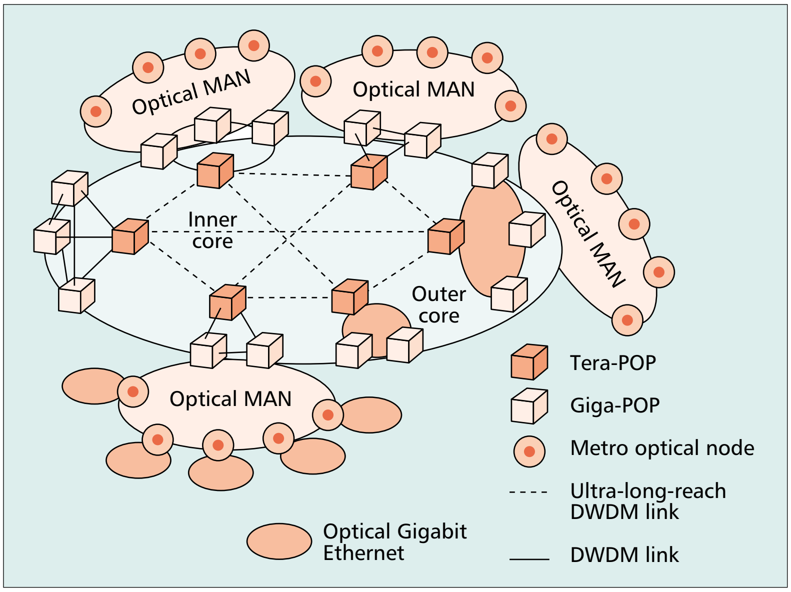

Figure 1. *Possible future optical Internet architectures.*

viding simple and scalable service differentiation by recognizing that most data flows generated by different applications can ultimately be classified into a few general categories (i.e., traffic classes). It does this by discriminating and treating the data flows according to their traffic class, thus providing a logical separation of the traffic into the different classes [7].

It seems likely that a combination of DiffServ and MPLS mechanisms will be deployed to ensure QoS across intradomain and interdomain networks.

Another major transformation in the upcoming Internet will be the widespread use of virtual private network (VPN) services, enabled by a very scalable solution that combines MPLS and Border Gateway Protocol v. 4 (BGP4) [8]. This solution places no constraints on addressing plans used by VPNs, and provides basic security comparable to that provided by frame relay or asynchronous transfer mode (ATM)-based VPNs, but without incurring the overhead and complexity associated with the use of overlay models or IPSec. In addition, this solution provides flexible and scalable support for QoS, with a great deal of flexibility with respect to the control policies that can be assigned to a particular class of service. This solution will also enable the provision of end-to-end VPN services across multiple service providers and carriers.

The Optical Internet as the Infrastructure for Next-Generation Internet Backbones

As happens with any overused term, the term optical Internet has multiple definitions. Our definition revolves around a basic premise: the optical Internet will be built with IP/MPLS and the optical layer as the predominant networking layers. This implies that the bulk of transport, switching, multiplexing, and routing functions will be performed only at those layers. The success of IP technology is founded, among other factors, on the wide diversity of link layer and lower layer protocols supported by the IP protocol. We believe that this diversity will continue, but it is very likely that Ethernet and SONET/SDH will be the predominant framing layers for the optical Internet. MPLS will provide the capabilities required to deploy feature-rich data services, to an arguable extent displacing ATM. Our definition of the optical Internet also critically depends on the design and deployment of a unifying control plane, built around IP and MPLS control protocols.

The optical Internet will also provide a new set of capabilities, not present in current IP backbone networks, for the dynamic provisioning of optical bandwidth upon request of IP routers and label switched routers (LSRs). This feature, also known as *optical bandwidth on demand*, will stretch the possibilities of the optical Internet in terms of dynamic behavior and service flexibility. New control interfaces, between routers and optical networks, need to be developed to support this capability. These data and control interfaces, known as the *optical user–network interface* (O-UNI), are briefly described later.

Network Architectures: A New Breed of Networks

The optical Internet will be built around innovative network architectures that become possible only after incorporating several recent technological developments in the areas of optical components and networks and IP/MPLS protocols and systems.

Inner Core/Outer Core: Lattice Networks for the Inner Core

Optical core networks, with several tens (or even hundreds) of terabits per second aggregated point-to-point capacity will be deployed in the next few years, mainly driven by data and, in particular, IP/MPLS applications. These core networks will also exploit the benefits, in scalability and cost efficiency, provided by optical switches.

From an infrastructure and management viewpoint, the feasibility of such networks strongly depends on the following key network components:

- Very high capacity optical line systems (OLSs) — with both high wavelength bit rate (up to 40 Gb/s) and large number of optical channels per fiber (up to a few hundred)
- Various reconfigurable optical switching network elements (e.g., OXCs and OADMs)
- A feature-rich optical layer control plane composed of protocols and algorithms that enable near-real-time optical channel provisioning, optical layer restoration, and distributed network intelligence
- Interworking of the optical layer control plane with IP core router protocols

Several recent technological improvements have led to the availability of OLSs capable of carrying tens of optical channels for several thousands of kilometers without electronic regeneration. These systems (referred to hereafter as *ultra-long-reach OLSs*) may use a combination of distributed Raman amplification, out-of-band or "strong" forward error correcting

(FEC) codes, amplifier gain control, and dynamic dispersion management. Typical systems will transport more than 100 wavelengths operating at 2.5 or 10 Gb/s up to 4000 km without electronic regeneration [9].

The availability of ultra-long-reach optical transport systems, together with some of the characteristics of Internet traffic patterns, where flows traveling distances longer than 1500 km between source and destination predominate (e.g., the so-called bicoastal traffic patterns in the United States), will open the possibility of new backbone network architectures. For example, these new architectures might center around an inner core composed of a reduced number of nodes (e.g., 10 percent of the overall number of core POPs in an ISP), sometimes dubbed Tera-POPs, interconnected through point-to-point ultra-long-reach DWDM links (Fig. 1) in an optical mesh architecture. Each of these Tera-POPs might consist of one or more terabit LSRs, optical switches, and the corresponding ultra-long-reach OLSs (Fig. 2).[1] This configuration optimizes the required optical bypass ratio, allowing for most of the traffic to be switched at the optical layer. At the same time, terabit LSRs, enabling traffic engineering and DiffServ support, handle the IP traffic that needs to be groomed and routed at each Tera-POP. The inner core is very likely to provide 1+1 optical path protection or optical-channel shared mesh restoration schemes (see [9] for a more detailed account).

It is worth noticing that, due to the capabilities of ultra-long-reach optical transport systems, inner core architectures might resemble lattice networks; two nodes (e.g., Los Angeles and New York) which before were six hops away are now separated by a single hop. Distance between nodes in the inner core may no longer be a metric of concern, at least from the IP perspective. This fact has some interesting implications in the optimization of routing algorithms and the inherent network design problems.

This inner core would be surrounded by an outer core consisting of networks of nodes called GigaPOPs, with optical mesh and optical shared protection rings as typical architectures for these regional backbones.

Metropolitan Optical Rings, Optical Gigabit Ethernet, and "Optical Ethernet"

We expect that metropolitan area networks (MANs) supporting Internet connectivity will continue to be characterized by hub-and-spoke or *backhauling* applications served by optical dedicated protection rings, or optical shared protection and optical mesh restoration schemes where more distributed traffic patterns dictate.

As new services emerge in the MAN environment, the nature of traffic in the network may change quite radically. Time-of-day-sensitive applications such as storage area networks lend themselves well to a more flexible "pay-per-use" model afforded by an MPλS/GMPLS-enabled control plane. At the same time, a new and perhaps unexpected player will soon enter the metropolitan and wide area networking (WAN) arenas. Gigabit Ethernet (GbE) and the emerging 10 Gigabit Ethernet (10 GbE) will, for the first time, reach beyond the local area networking (LAN) environment, creating an unprecedented transformation in end-network services, pricing structures, and technology integration.

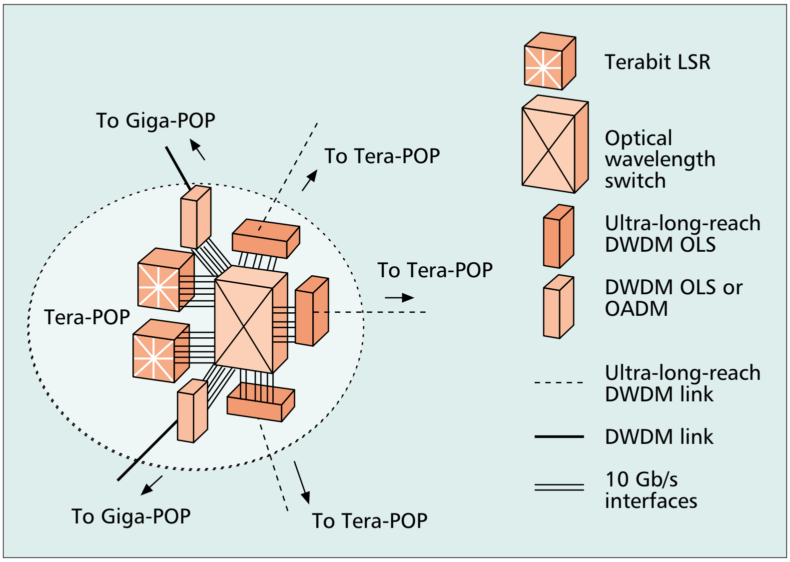

Figure 2. *An example tera-POP functional internal architecture.*

As a result, we envision that the optical Internet will have two facets: classical packet over SONET/SDH (POS) in the core network, plus GbE and 10 GbE signals transported over WDM networks in MAN and very likely WAN environments. This so-called *optical Ethernet* will likely become the de facto standard for gigabit Internet connectivity. And, surprisingly, this may be the case not only for enterprise customers, but for residential services as well. Recent efforts along with emerging carriers in Europe, Canada, and the United States [13, references therein], for example, provide a glimpse into a possible future when GbE and 10 GbE services may be as common as dialup, ISDN, cable modem, or DSL services today.

Service Models and Interconnection Models

The Optical Internet Network Model: UNI and NNI

The network model considered in several of the ongoing efforts for standardization of optical internetworks (e.g., [10]) consists of IP routers or LSRs attached to an optical network, and connected to their peers over dynamically established switched lightpaths (optical channel trails). In many cases an ISP owns and deploys IP routers and LSRs and interfaces with another carrier, who in turn owns and operates the optical network elements. In some other cases it is the ISP who owns all the IP routers, LSRs, and optical network elements. The optical network is assumed to consist of multiple optical subnetworks interconnected by optical links in a general topology. This network may be multivendor. In the near term, it may be expected that each subnetwork will consist of a single vendor's opti-

[1] *Tera-POPs will consist of these types of functions, the elements of which may or may not be integrated into multifunctional/multiservice equipment.*

This network model provides a clear separation of logical control interfaces at the client-optical network interface, referred to as the Optical User-Network Interface, and at the optical sub-network interface, referred to as the Optical Network-Network Interface.

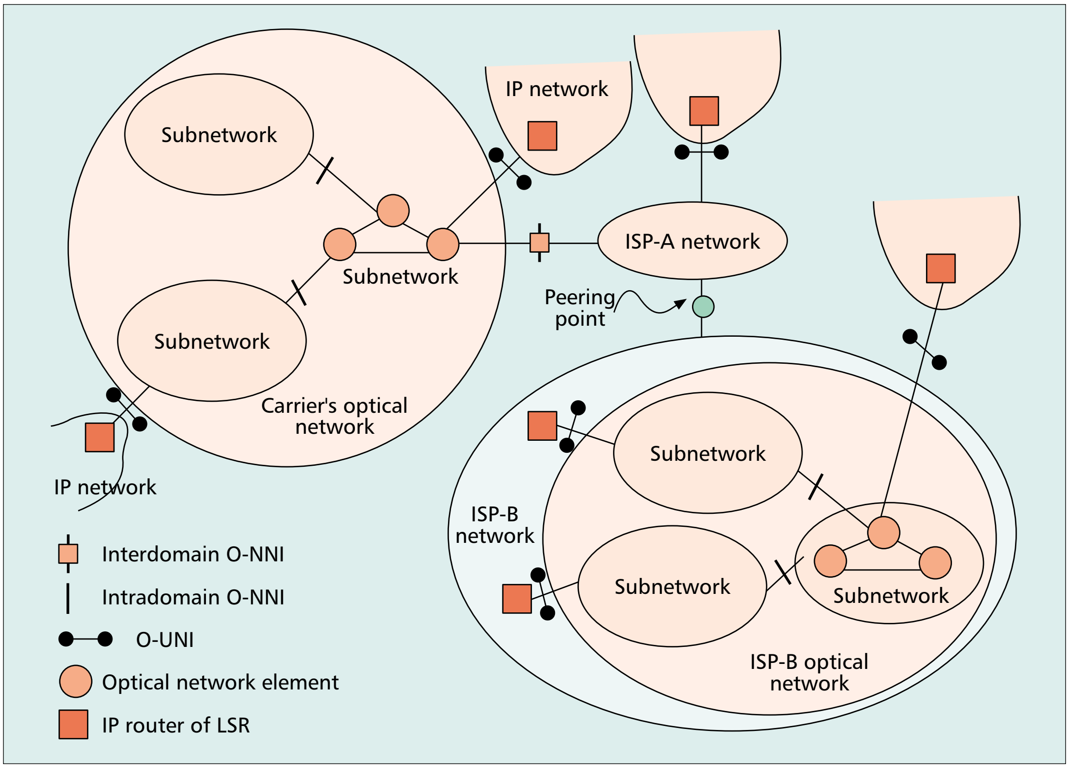

Figure 3. *A high-level optical Internet architecture model.*

cal network elements. This network model is shown in Fig. 3.

This network model provides a clear separation of logical control interfaces at the client optical network interface, referred to as the *optical UNI* (O-UNI), and at the optical subnetwork interface, referred to as the *optical network–network interface* (O-NNI).

The services defined at these interfaces determine the type and amount of control information that flows across them. It is, in principle, possible to have a unified service definition across both these interfaces such that there is virtually no difference in the type of control information that flows across them. On the other hand, it may be required by the several service providers and carriers involved in the network architecture to minimize the flow of control information, especially routing-related information, over the O-UNI. In this case, the O-UNI and O-NNI may look different in some respects.

Each of these interfaces can be categorized as public or private, depending on network context and service models. Routing information (i.e., topology state information) can be exchanged across a private interface, without any restriction; on the other hand, the amount and type of routing information that can be exchanged across a public interface may be limited by the use of explicit restrictions (abstraction, filtration, etc.) [10]. This distinction in the relationships that may occur across private or public control interfaces, very similar to the intradomain vs. interdomain separation of routing information and protocols in the Internet, is the source for the different interconnection models (overlay, peer, augmented) outlined below. With this distinction taken into account, the network model is general enough to accommodate all the possible deployment scenarios of carrier and service provider interconnection.

Interconnection Models: Overlay Model, Peer Model, Augmented Model

The architecture alternatives for the optical Internet model can be better understood if we keep a clear separation between the data and control planes of the O-UNI.

The optical network provides a service to external entities (e.g., IP routers and LSRs) in the form of coarse-granularity fixed-bandwidth transport pipes (optical paths). Such optical paths must necessarily be in place before communication (at the IP layer) between IP routers at the edge of the optical network can begin. Thus, the data plane of such optical internetworks is realized over an overlay network of optical paths. On the other hand, IP routers and optical network elements can have either a client-server or peer relationship on the control plane, especially regarding the implementation of routing protocols that allow dynamic discovery of IP endpoints attached to the optical network. With this distinction in mind, it is clear that the different possible network architectures are defined essentially by the organization of the control plane [10].

The different models for organization of the control plane are referred to as the *overlay model*, the *peer model*, and the *augmented model*; they differ from each other in a number of ways.

First of all, their scope of application varies. The overlay model is of interest to supercarriers, optical backbone providers, and those ISPs which lease optical infrastructure from optical backbone providers. In this case the control model calls for completely independent IP/MPLS and optical layer control planes. The augmented and peer models address the problem of tightly integrated IP/MPLS and optical layers for optimized optical internetworks. They are applicable to ISPs that are also optical backbone providers, and wish to integrate the design and operation of the optical layer and IP layers for an optimized optical internetwork.

Second, the models differ in the relationship between address spaces and the exposure of optical layer topology to the IP/MPLS layer (and vice versa). The overlay model uses independent addressing schemes and mandatory opacity of the optical layer network topology (i.e., the optical layer network topology is not exposed at all to the service layers). The peer and augmented models follow a common addressing scheme (e.g., optical network elements become IP addressable devices, as noted in [11]), and allow for full or partial opacity or transparency of the optical layer network topology under administrative control.

The third parameter for variation is the way in which routing protocols (next section) run on the IP and optical layers. The overlay model is based on two sets of separate and isolated protocol instances for both layers, while the peer model uses a single common set of instances. In between lies the augmented model, which uses separate but coordinated instances (e.g., addressing and routing information may be shared between both instances).

Finally, an additional variable is the way in which label switching protocols (next section) run on the IP/MPLS and optical layers. A single monolithic label switching protocol would be very interesting architecturally and administratively because of its potential simplicity, conceptual integrity, and ease of management, especially from the perspective of network operations control. But the semantics of label switching, and the establishment and maintenance of label-based optical paths in an optical network may be different from LSPs in MPLS networks.

Each model has a clear rationale. The rationale for the overlay model is based on the fact that network providers and backbone builders leasing optical channels to ISPs will not want to disclose any information about the internal details of the optical network infrastructure, such as topology or capacity sharing due to shared restoration. A client/server model is then the only option for the relationship between the control planes of the IP/MPLS layers and the optical layer. On the other hand, the rationale for the integrated and peer models is based on the fact that service providers (ISPs, backbone builders) that deploy and control IP routers, LSRs, and optical networking equipment can optimize the design and runtime control of their optical internetworks. Moreover, past experiences with overlay models (IP over ATM) provide some useful insights regarding the problems associated with these models (e.g., mapping between UNI and NNI, routing scalability). As the optical layer becomes more analogous to ATM (in the sense of providing *virtual circuits*, VCs), the same problems can be expected for optical internetworks.

For example, the classical IP over ATM overlay model presents some scalability issues due to the amount of routing information that must be exchanged in response to a topology change. In the case of IP routers connected by a full mesh of ATM VCs, the number of immediate neighbors to any router equals the number of routers around the ATM cloud minus 1 (itself). The number of VCs required to form such full mesh equals $n(n-1)/2$, where n is the number of routers. And the number of adjacencies any router keeps equals $n - 1$. Adjacencies enable a router to keep track of which other routers are directly connected to it, including whether any links between them are operational or not, and are used to exchange routing information with those neighbors. It can be shown that the amount of routing information that needs to be exchanged in such a network in response to a topology change can be as much as on the order of n^4 [14]. This can lead to a serious scalability problem when n is large.

The Next Hop Resolution Protocol (NHRP) can be used to solve this problem because it allows routers to establish VCs over which they can send data without having to establish a routing adjacency between them. On the other hand, this approach presents some problems, including the need to deploy and run several NHRP servers, much like ATM Address Resolution Protocol (ATMARP) servers; the possibility of introducing persistent forwarding loops; and the lack of multicast support.

In the case of interworking between the IP/MPLS and optical layers, the lightpaths are analogous to ATM VCs, and the optical network elements are analogous to ATM switches. Clearly, a client/server model will suffer from the same set of problems already encountered in the IP over ATM overlay model.

While the debate rages as to which will emerge as the dominant model, from a broad perspective, the initiatives underway in the IETF and elsewhere, when taken together, provide a framework from which the appropriate service model can emerge. Depending on the carrier or service provider's application, and whether they own, manage, and maintain either the IP routers, the optical transport networking equipment, or both, the framework will allow for either separate and distinct or converged control planes.

A Unifying Control Plane for the Optical Internet

We expect optical component and systems technology will continue to evolve very rapidly, with additional disruptive technologies coming into place in the next few years. An overall network architecture for the optical Internet that allows gradual and seamless introduction of these disruptive technologies into the network without time-consuming and costly changes is therefore fundamental. There is a certain consensus around the idea that a flexible control plane capable of providing rich functionality in terms of routing and signaling, while at the same time hiding vendor-specific implementations or administrative policies, will be one of the most critical elements for the realization of the optical Internet.

In the case of interworking between the IP/MPLS and optical layers, the lightpaths are analogous to ATM VCs, and the optical network elements are analogous to ATM switches. Clearly, a client/server model will suffer from the same set of problems already encountered in the IP over ATM overlay model.

The challenge for the IP Optical research and industry communities is how to build a control plane for the optical layer capable of accommodating, within a single encompassing framework and set of protocols, all the possible service and interconnection models described earlier.

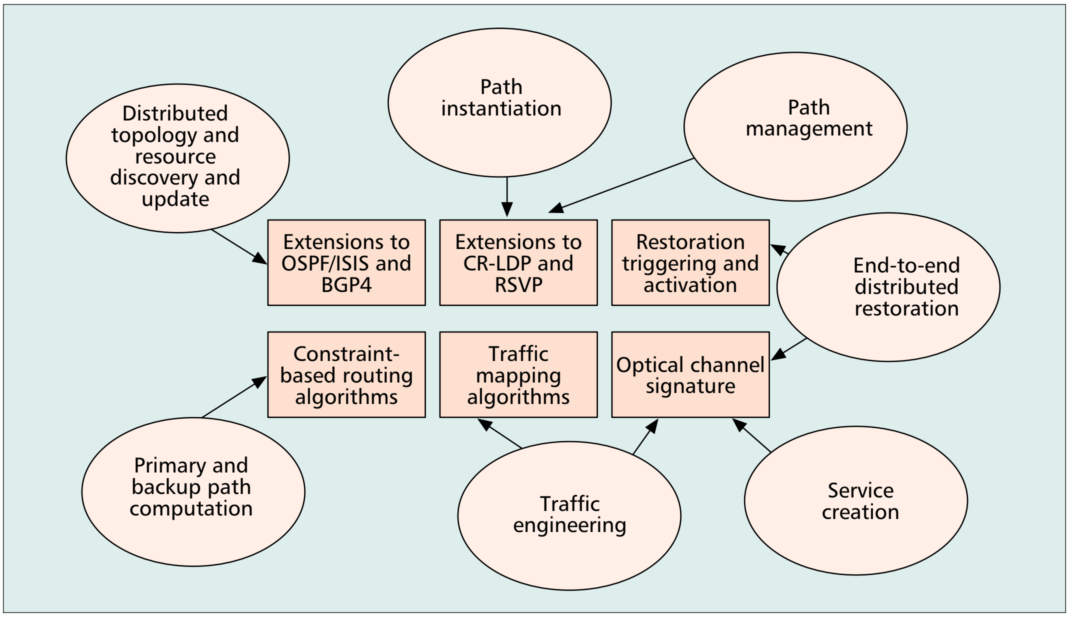

Figure 4. *Protocols and algorithms for the control plane of the optical Internet.*

The challenge for the IP optical research and industry communities is how to build a control plane for the optical layer (and an interlayer control plane between the optical and IP/MPLS layers) capable of accommodating, within a single encompassing framework and set of protocols, all the possible service and interconnection models described earlier. The possibility of designing such a control plane is at hand, thanks to the flexibility introduced by several control protocols already present in the IP and MPLS layers.

The basic philosophy for the control plane advocated by the MPλS/GMPLS concept can be summarized as follows: the control plane is composed of a set of IP/MPLS-centric algorithms and distributed protocols running in all the nodes of the optical internetwork. A set of routing protocols and algorithms —based on appropriate "optical" extensions to Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS), or BGP — maintains a synchronized network topology database, and advertises topology state information to maintain and refresh that database. A constraint-based routing algorithm on each node may then use the information in the topology database and other relevant details to compute appropriate paths (for primary and restoration optical paths) through the optical domain. Once a path is computed, a signaling protocol (actually an extended or tailored version thereof) such as Resource Reservation Protocol (RSVP) or constraint-based routing Label Distribution Protocol (CR-LDP) can be used to instantiate the path. Optical paths can then be maintained (rearranged or terminated) as LSPs.

In general, a whole range of IP-based service applications and powerful resource optimization algorithms together with new and extended IP protocols need to be developed for efficient traffic engineering of emerging OXC and OADM-based IP/data networks (see [11, 15] for more detailed descriptions). Figure 4 illustrates how these various entities interwork to achieve this goal.

A few issues particular to the nature of optical internetworks must be addressed. The most relevant are the support for bidirectionality in lightpaths, failure recovery, and signaling for restoration [12].

Optical channel trails are bidirectional. That is, the output port selected at an OXC for the forward direction is also the input port for the reverse direction of the path. This contrasts with MPLS environments where LSPs are of a unidirectional nature. Adequate extensions to RSVP Traffic Engineering (RSVP-TE) and CR-LDP are needed to handle this requirement, avoiding collisions and race conditions that might appear during the optical channel trail setup phase.

The issue of how to incorporate protection and restoration at the optical layer — in its multiple possibilities, ranging from 1+1 optical path protection, optical channel shared protection rings (OCh/SPRINGs), or optical mesh-based shared restoration — and the associated algorithms to compute backup paths and signaling required to initiate and propagate restoration events possibly remains the biggest challenge at this moment. This issue is even more acute when considering the multilayer implications of peer models where both MPLS fast rerouting and optical layer protection/restoration mechanisms may coexist and interwork. Some initial approaches to this issue can be found in [10].

Conclusions

The concept of "Internet time" as shorthand for the rapid rate of technological innovation in the Internet is by now well established. We could argue that in recent years the rate of technological innovation in optical components and optical networks has been even faster than that of the Internet itself, and that it was a matter of time until both phenomena converged, giving rise to a networking revolution that will even further transform the network infrastructure of the Internet as we know it.

The optical Internet, as the network infrastructure that will enable the backbones of the next-generation Internet, is a concept ripe with potential and promise. It will result from a powerful combination of:

- Optical innovations (ultra-long-reach systems, flexibly reconfigurable OADMs, and large-scale OXCs) enabling new network architectures and unprecedented levels of bandwidth management
- A distributed control plane (MPλS/GMPLS) based on "reused and extended" protocols, algorithms, and software artifacts from IP/MPLS networks
- Industry groundswell activities continually pushing "downmarket" technologies (e.g., Ethernet) "upmarket"

The two-way ripple effect of technologies penetrating from the edge to the core and vice versa, and the associated transformations that result, represent the keys to unlock the potential for an optical Internet.

Finally, we glance into a future governed by "network Darwinism," wherein only the fittest networks and network elements survive. Optical networking systems have evolved in an age of "Net heads vs. Bell heads" — an era characterized by the phenomenal growth of the Internet and IP, coupled with the more measured growth of traditional telecommunications services; disparate applications, running over separate networks, managed by different organizations. Current events signal the end of this era. Distinctions have blurred as networks have converged, and only organizations that marry the two — IP/data and optical/transport — will survive.

The two-way ripple effect of technologies penetrating from the edge to the core and vice versa, and the associated transformations that result, represent the keys to unlock the potential for an optical Internet.

References

IETF drafts and RFCs are available at http://www.ietf.org

[1] C. M. Christensen, *The Innovator's Dilemma*, Harvard Business School Press, 1997.

[2] R. D. Doverspike, S. Philips, and J. R. Westbrook, "Future Transport Network Architectures," *IEEE Commun. Mag.*, vol. 37, no. 8, Aug. 1999, pp. 96–101.

[3] P. Bonenfant, "Optical Layer Protection and Restoration," *Optical Networking*, A. Bononi, Ed., Springer-Verlag, 1999, pp. 77–88.

[4] O. Gerstel and R. Ramaswami, "Optical Layer Survivability: A Services Perspective," *IEEE Commun. Mag.*, vol. 38, no. 3, Mar. 2000, pp. 104–13.

[5] OECD rep., "Internet Infrastructure Indicators," Oct. 28, 1998; available at http://www.oecd.org

[6] L. Mathy, C. Edwards, and D. Hutchison, "The Internet: A Global Telecommunications Solution?" *IEEE Network*, July/Aug. 2000, pp. 46–57.

[7] D. Black *et al.*, "An Architecture for Differentiated Services," IETF RFC 2475, Dec. 1998.

[8] E. Rosen and Y. Rekhter, "BGP/MPLS VPNs," IETF RFC 2547, Mar. 1999.

[9] A. Rodriguez-Moral *et al.*, "Optical Data Networking: Protocols, Technologies and Architectures for Next Generation Optical Transport Networks and Optical Internetworks," *IEEE J. Lightwave Tech.*, vol. 18, no. 12, Dec. 2000, pp. 1855–70.

[10] B. Rajagopalan *et al.*, "IP over Optical Networks: A Framework," IETF Internet draft, draft-many-ip-optical-framework, work in progress.

[11] D. Awduche *et al.*, "Multi-Protocol Lambda Switching: Combining MPLS Traffic Engineering Control With Optical Crossconnects," IETF draft, draft-awduche-mpls-te-optical, work in progress.

[12] B. Rajagopalan *et al.*, "Signaling for Fast Restoration in Optical Mesh Networks," IETF draft, draft-bala-restoration-signaling, work in progress.

[13] A. K. Bjerring and B. St. Arnaud, "The Coming Revolution in Dark Fibre Networks," available at http://www.canarie.ca/advnet/canet3/fibre.html

[14] B. Davie and Y. Rekhter, *MPLS Technology and Applications*, Academic Press, 2000.

[15] A. Banerjee *et al.*, "Generalized Multiprotocol Label Switching: An Overview of Routing and Management Enhancements," *IEEE Commun. Mag.*, vol. 39, no. 1, Jan. 2001, pp. 144–50.

Biographies

ANTONIO RODRIGUEZ-MORAL (arodmor@photuris.com) leads the definition of network architectures and next-generation systems at Photuris, Inc. Before joining Photuris he was at Bell Laboratories, Lucent Technologies, where he was involved in the definition and analysis of network architectures, product evolution planning, and strategy for next-generation DWDM networks, and in research and development of network management systems and high-speed IP networks. Prior to joining Bell Laboratories he was with AT&T Network Systems in Europe, where he led several research and development projects for SDH and passive optical networks. He received his M.S. degree in electrical engineering from the Technical University of Madrid, Spain.

PAUL BONENFANT (pbonenfant@photuris.com) serves as chief architect at Photuris, Inc. His experience spans SONET/SDH, WDM, and optical networking transport architecture, product evolution planning, network survivability, and associated global standards development. Before joining Photuris he served as a business development manager for Mergers and Acquisitions in Lucent's Optical Networking Group, and led a group responsible for optical network architecture evolution. Prior to joining Lucent, he led requirements and standards development for SONET/SDH self-healing rings and dense WDM systems at Bell Communications Research (Bellcore, now Telcordia). He holds a B.S. degree in engineering and applied science, and an M.S. degree in electrical engineering, both from the California Institute of Technology in Pasadena, California.

MURALI KRISHNASWAMY (murali@photuris.com) is responsible for defining the GMPLS control plane architecture for next-generation systems at Photuris, Inc. His other responsibilities include prototyping and testing control layer protocols for routing, signaling, and restoration. Earlier he was with Lucent's Optical Networking Group where he was involved in defining and prototyping new networking architectures for core IP routers and SONET/SDH and DWDM systems. Prior to that, at AT&T he successfully designed and implemented the industry's first PSTN-Internet internetworking architecture for IP-based call control. He received his Master's degree from the Stevens Institute of Technology, New Jersey.