

Article

Anticipation of Traffic Demands to Guarantee QoS in IP/Optical Networks

Carolina Pinart

Centre Tecnològic de Telecomunicacions de Catalunya, Av. Carl Friedrich Gauss 7, 08860, Castelldefels, Spain; E-Mail: carolina.pinart@cttc.cat; Tel.: +34 93 645 29 00; Fax: +34 93 645 29 01

Received: 21 August 2010; in revised form: 2 September 2010 / Accepted: 2 September 2010 /

Published: 21 September 2010

Abstract: Traffic in the Internet backbone is expected to grow above a few Tbit/s in 2020. To cope with this, operators are moving to IP/optical network architectures, where IP is the convergence layer for all services. On the other hand, the quality of service (QoS) requirements of future applications encompasses the individualization of services and the assurance of stricter quality parameters such as latency, jitter or capacity. In other words, future optical networks will not only transport more IP data, but they will also have to offer differentiated QoS requirements to services. Finally, some emerging applications, e.g., grid computing, need greater flexibility in the usage of network resources, which involves establishing and releasing connections as if they were virtualized resources controlled by other elements or layers. In this context, traffic-driven lightpath provisioning and service-plane approaches arise as very interesting candidate solutions to solve the main challenges described above. This work reviews the concepts of service-oriented and self-managed networks and relates them to propose an integrated approach to assure QoS by offering flow-aware networking in the sense that traffic demands will be anticipated in a suitable way, lightpaths will be established taking into account QoS information (*i.e.*, impairments) and complex services will be decomposed into optical connections so that the above techniques can be employed to assure QoS for any service.

Keywords: IP convergence; QoS; optical networking; service plane

1. Introduction

The backbone of the Internet is composed of electrical core routers. These routers perform data switching and forwarding digitally, e.g., using synchronized digital hierarchy (SDH) technology, and only make use of fiber optics as the transport medium between adjacent routers. To do so, switches are combined with wavelength-multiplexing systems to increase raw bandwidth at the optical layer, while the IP-layer core electrical routers provide the intelligence needed to route packets between any source and destination on the Internet. The networks forming the Internet backbone consist of multiple service networks on top of a common transport network, which relies on fiber. In this context, the traditional dominant service is voice, but there are also data-oriented services, such as Internet service providers using IP networks, or virtual private networks (VPN) using multi-protocol label switching (MPLS).

The complementary electrical and optical technologies of the Internet backbone are converging into IP/optical network architectures. This means that network technologies are evolving from voice- to data-optimized, where IP is the convergence layer. In the first step of this convergence, IP/optical networks are circuit-switched (or wavelength-routed) networks composed of nodes interconnected with fiber optics that act as optical-layer switches and IP routers simultaneously. These IP/optical routers are capable of transporting data at the optical layer by creating optical paths (lightpaths) that bypass the IP routing infrastructure of the network, *i.e.*, by offloading data to the optical layer. This is known as wavelength routing and results in low-delay, quality-enabled data delivery at the wavelength level. Typically, IP/optical networks are enabled with control and management infrastructures to establish, modify and release lightpaths, e.g., by using the generalized MPLS (GMPLS) technology.

This way, data flows can traverse IP/optical networks through either an optical path or a chain of routing decisions at the IP layer, where IP routers are interconnected with fiber links [1]. This is especially interesting for large flows for two reasons. First, offloading them to the optical layer results in these flows experiencing faster and more reliable transmission with optical switching than with traditional IP routing. Second, remaining flows at the IP level also experience higher quality of service (QoS) because the IP layer is less congested after offloading the large flows to the optical layer. Nowadays, several high-broadband research and education networks are migrating towards this scheme of IP/optical networking. Examples of this migration are Europe's GÉANT (www.geant.net), US' Internet2 (www.internet2.edu), The Netherland's SURFnet (www.surfnet.nl) or Canada's CANARIE (www.canarie.ca) networks. These networks transport data of their users and data of applications demanding very high bandwidth and stringent quality levels, such as grid applications, distributed radio telescopes or high-definition television broadcast.

Typically, human operators (network managers) are in charge of selecting flows to be placed in the optical layer, in general by establishing lightpaths in advance, either manually or by using automated technologies such as GMPLS. Knowing the heavy-hitter behavior of high-bandwidth demanding applications helps network managers establish lightpaths in advance, because heavy-hitter flows are easily identified as flows with a very large number of packets, as defined by Mori *et al.* in [2]. However, there are other "big" IP flows that do not follow heavy-hitter behavior and could also benefit from being moved to lightpaths. These flows are called "elephants" and have the following characteristics: they are persistent in time and few in number but account for most of the Internet traffic volume [3]. This is, though, a very complex decision to be made by human operators, because it

involves analyzing several network parameters for all IP flows traversing the network and then selecting the flows in a timely manner to balance the network properly [1]. In parallel to traffic-driven lightpath provisioning, much effort is being put into investigating the best approaches to establishing lightpaths and then monitoring them once they are in service so that QoS is assured with minimum cost and sufficient efficiency. This is especially the case in all-optical networks where the IP layer (or in general any digital information) is only present at the endpoints of a lightpath, which results in impairment accumulation and in major challenges to traditional performance and fault management.

1.1. Related Work

In the following, research works dealing with service management and QoS assurance in IP/optical networks are outlined. These works will be the basis for the approach proposed in this paper. Fioreze *et al.* developed the concept of self-management [4] to achieve traffic-driven lightpath provisioning. Their goal is to select the IP flows more suitable to be offloaded to lightpaths in IP/optical networks, and also to establish and release these lightpaths in an automated manner. The core to this approach consists in monitoring network data to detect elephant flows based on a set of parameters to decide on transporting these flows at the optical layer in a rapid, automated manner, while the other flows remain at the IP layer. The concept of impairment-aware routing and wavelength assignment (IA-RWA) was introduced in the late 90s by Ramamurthy *et al.* to assure QoS in the establishment process of new lightpaths by means of including information of the optical layer into the routing and resource assignment decisions [5]. That is, the estimation of the impairments experienced by the optical signal transporting data over a lightpath is included in the lightpath establishment process to help assure QoS by anticipating the degradation that data would suffer if the lightpath was established.

This is important because lightpaths transport high amounts of data, which means that establishing too impaired lightpaths may result in enormous data loss, and hence in severe QoS violations. Related with that, non-intrusive monitoring schemes are also being investigated under the framework of optical performance monitoring [6]. This is especially relevant in all-optical (transparent) networks, where data is transported in the optical domain from the network's ingress to egress without undergoing electro-optical conversions in the intermediate nodes, and QoS needs to be assured without having access to digital information about the optical signals being transported in the network.

The application of a service plane to optical networks was first proposed by the Service Oriented Optical Network (SOON) project in 2005 [7]. This proposal improved the ITU-T standard automatic switched transport network (ASTN) architecture [8] by introducing an additional functional layer called service plane, which contained the intelligence for service provisioning and provided abstraction from transport-related implementation details. Further details are provided in Section 2.

This paper reviews the concepts of service-oriented and self-managed networks. Then, it relates them to propose an integrated approach to assure QoS by offering flow-aware networking in the sense that traffic demands will be anticipated in a suitable way, lightpaths will be established with the cooperation between the IP and optical layers to achieve impairment-awareness, and complex services will be decomposed into optical connections so that the above techniques can be employed to assure QoS for any service. The remainder of the paper is organized as follows. Section 2 overviews the

architectures of service-oriented IP/optical networks and self-managed IP/optical networks. In Section 3, we review the main concepts of QoS in IP/optical networks in terms of the relation between QoS and the previous architectures (service-oriented and self-managed networks). Section 4 proposes an integrated approach for QoS assurance in IP/optical networks involving anticipation of traffic demands to offload elephant flows to the optical layer, also decomposing complex services into a series of connection requests, and then establishing lightpaths incorporating IA-RWA. Finally, in Section 5 we draw conclusions and outline future work.

2. Service-oriented and Self-managed Optical Networks

Assured QoS in IP/optical networks should allow network customers to create, modify and release services in a dynamic manner and according to pre-established service level agreements (SLA). To achieve this, two main architectures have been recently proposed: service-oriented optical networks and self-managed networks.

2.1. Service-oriented Optical Network Architectures

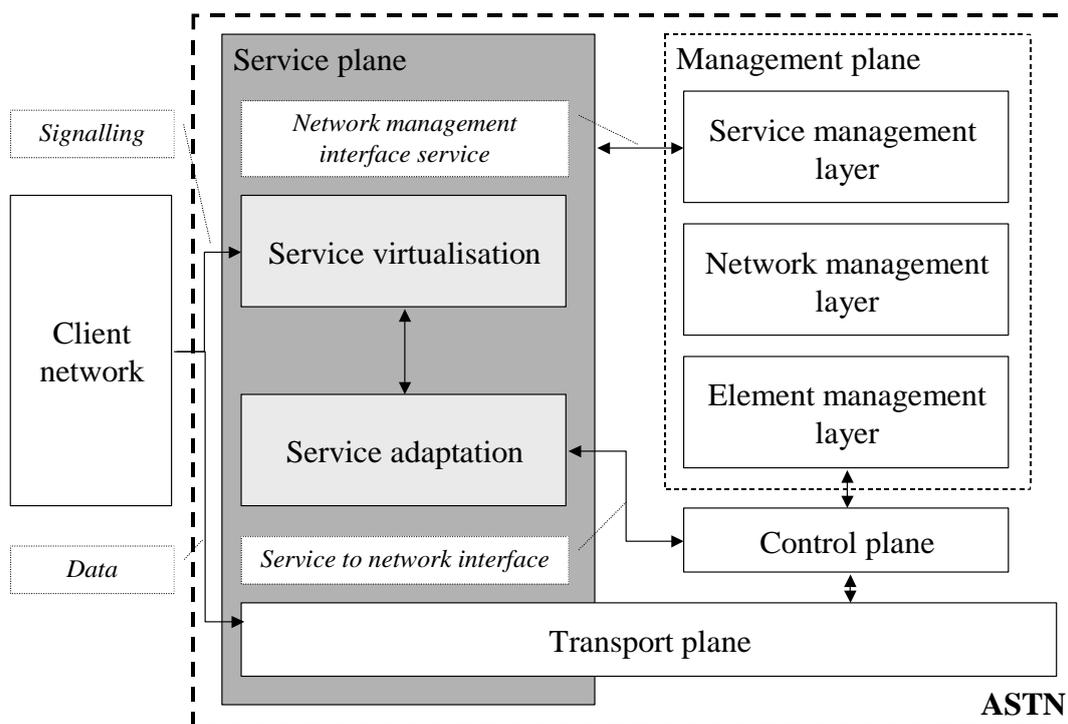
Figure 1 depicts the service plane (highlighted in gray) introduced by the SOON project in [7], and relates it with the elements and interfaces of the ASTN planes of management, control and transport, described in [8]. The figure also shows the two main building blocks of the service plane:

- Service virtualization, which allows client applications to refer network resources in abstract terms, such as session or topology.
- Service (connectivity) adaptation, which translates connectivity requests into a set of messages that the control plane can understand. The format and contents of these messages depend both on the nature of connectivity requests and on the technology in which the control plane is implemented.

This service plane concept can be implemented in many ways. For example, Martini *et al.* propose and evaluate a service platform for optical transport networks based on the session initiation protocol (SIP) that enables applications to exchange messages with the network to generate service requests for the reservation of the resources they need. The proposed service platform subsequently translates such requests into technology-dependent directives for the IP/optical network nodes [9]. The service platform is composed of three functional layers:

- An application-oriented network layer, which manages application signaling and generates the service requests. This layer handles the communication among end hosts and maps generic communication instances into network service sessions that are managed using the SIP protocol.
- A service plane, which manages the mapping between service requests and technology-specific directives. This plane follows the architecture depicted in Figure 1.
- A GMPLS-enabled transport network infrastructure, which performs data forwarding. Note that this infrastructure includes the control and transport planes depicted in Figure 1.

Figure 1. Service Oriented Optical Network (SOON)’s service plane building blocks and interfaces. Adapted from [7].



Another implementation example is the European PHOSPHOROUS project, in which Harmony is used as the basis for a service-oriented IP/optical network [10]. Harmony is a multi-domain path provisioning system where users and grid applications can reserve end-to-end paths in advance. In other words, Harmony is a network resource brokering system that allows heterogeneous domain interoperability by performing an inter-domain resource brokering over network resource provisioning systems (NRPS) such as GMPLS. Harmony is based on a web service architecture and contains a network service plane (NSP). In PHOSPHOROUS, the communication interfaces of the NSP and the NRPS adapters implement common web service operations for topology and reservation management. The optical network domains publish only their border endpoints (interfaces) to the NSP. Intra-domain connections are hidden to the NSP and maintained by the corresponding NRPS.

More recently, Takeda *et al.* [11] have proposed a thin layer for seamless interconnection of multi-technology transport networks, which implicitly involves a service-oriented optical network. The rationale for this proposal is as follows: the evolution from voice- to data-oriented networks results in multi-technology transport networks, which shall be interconnected seamlessly to provide end-to-end services. In this context, the thin layer is an enabling technology to achieve this goal through providing point-to-point links between endpoints with the required QoS. To do so, the thin layer encompasses the necessary control and management functionalities.

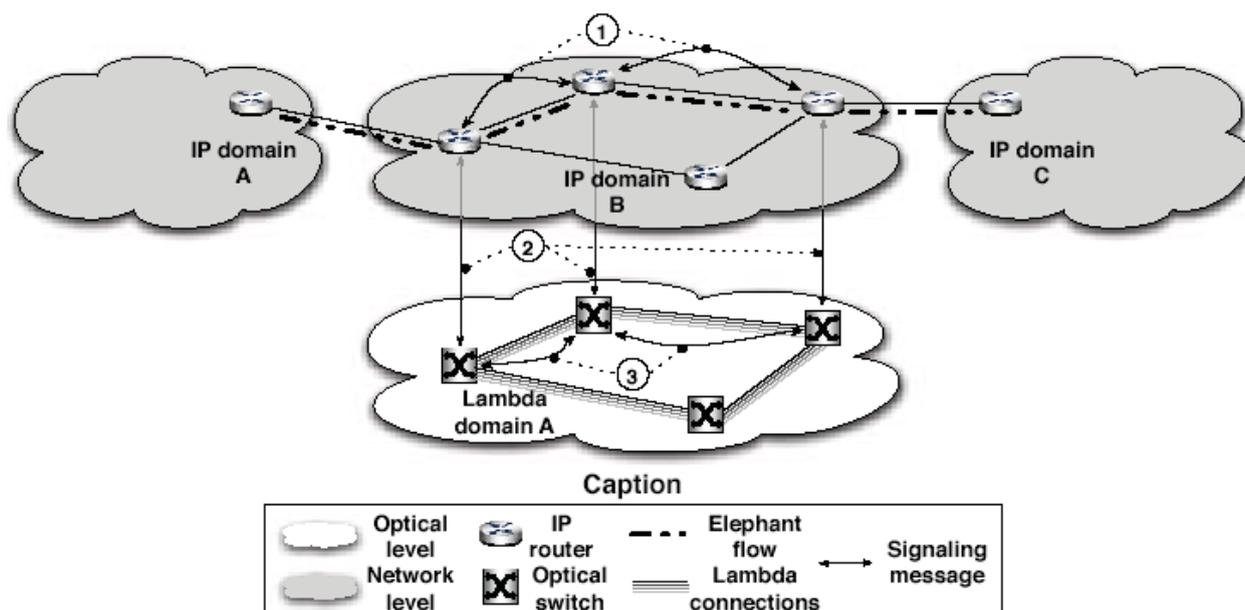
2.2. Self-managed Optical Networks to Handle Elephant Flows

In the above-described service-oriented optical networks, either users or applications are in charge of requesting the establishment of new lightpaths when they need to transmit data, if this is not

performed by the network managers as in currently deployed optical networks. In self-managed networks, it is the optical network that infers when new lightpaths are needed by detecting elephant flows and by subsequently requesting the establishment of lightpaths to offload these flows to the optical layer. Figure 2 illustrates the architecture of a self-managed optical network.

The key to this approach is flow characterization, not service characterization as in service-oriented networks. There are several research works on flow characterization, most of which are based on the analysis of traces and on the comparison of observation points. These works, e.g., [13–15], are mainly focused on estimating precise flow size distribution or packet distribution. Note that data for flow characterization can be collected at one or more observation points or can be obtained by measurement techniques such as packet sampling. The metrics considered for analyzing flows are typically packets, bytes and port distribution, and some works also consider effects on the flow size distribution, e.g., packet counts, SYN information and sequence number information.

Figure 2. Architecture of a self-managed optical network. Extracted from [12].



Other relevant works are those of Papagiannaki *et al.* [16] and Lan *et al.* [17]. In [16], the authors focus on the analysis of the volume (size) of big IP flows as well as on their historical behavior. The authors argue that the flow volume may be very volatile, which results in different load values captured at different time intervals. In [17], the authors characterize big flows by observing their size, duration, throughput (rate), and burstiness, as well as by examining their correlation considering the traditional 5-tuple flow definition: group packets with the same source and destination IP addresses, same transport protocol, and same transport protocol source and destination port numbers. Alternatively, [1] proposes a novel set of parameters to predict the traffic volume generated by elephant flows, which are summarized in Table 1. Note that volume prediction is employed in the self-management approach to reduce the cost of high-volume flows staying at the IP level, in terms of consumption of network resources, by detecting them in a timely manner so that they can be transferred to the optical layer.

Table 1. Network parameters to detect elephant flows.

Flow identifier parameters	Flow behavior parameters
TCP/UDP port numbers	Duration
IP addresses	Number of packets in the flow
Network segments (subnets and autonomous system)	Throughput (Bps, bits per second)
Protocols	Throughput (Pps, packets per second)
Type of service (ToS)	

3. Service Level Agreements in IP/Optical Networks

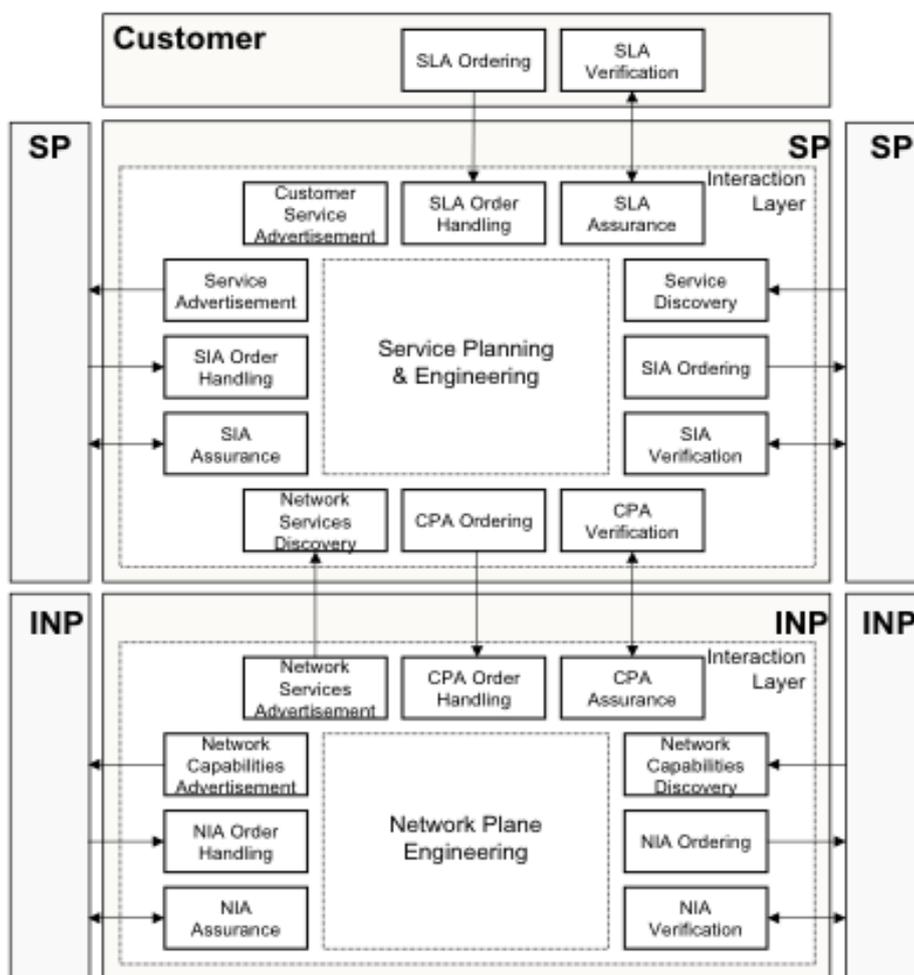
Connectivity between two geographically dispersed locations is usually provided by an independent third party called a service provider. The SLA is then a contract agreed upon between the service provider and the service consumer, and is the basis for guaranteeing QoS because it specifies in measurable terms the services furnished by the provider and the penalties assessed if the provider cannot meet the goals established. In other words, the SLA details the attributes of the service provided to the customer. These parameters span bandwidth, fault tolerance, recovery time, reliability, availability, response time, packet/burst loss and bit error rate, to name a few. In circuit-switched networks, QoS also includes maximum connection establishment delays and blocking probability. In the Internet, the reality is that end-to-end SLAs are scarce, and SLAs among Internet service providers are limited, which gives an idea of the complexity of managing QoS in deployed optical networks. Besides, the choice of ‘optical QoS’ parameters to be included in SLAs and the ways in which these parameters shall be measured or interpreted is still an ongoing debate.

It goes without saying that guaranteeing QoS is a must for IP/optical networks. The applications being deployed across the public Internet today are increasingly mission-critical, whereby business success can be jeopardized by poor network performance. In other words, IP/optical networks will not be much promising unless they can guarantee a predictable performance as specified by the QoS parameters established in the SLAs. Transmission performance is a key factor regarding QoS in optical networks because optical signals carry huge amounts of data, which means that poor transmission quality has a dramatic impact on QoS. For this reason, there are many works in the literature on impairment-aware lightpath provisioning (IA-RWA), e.g., [5], and fault management schemes adapted to IP/optical networks, especially for all-optical transport networks, e.g., [18].

The QoS requirements of future applications encompass the individualization of services and stricter quality parameters such as latency, jitter or capacity. In other words, IP/optical networks will not only transport more IP data, but they will also have to offer differentiated QoS requirements and traffic patterns. Finally, some emerging applications, e.g., grid computing, need greater flexibility in the usage of network resources, which involves establishing and releasing connections as if they were virtualized resources that can be controlled by other elements or layers. In this context, traffic-driven lightpath provisioning (described in Section 2.2) arises as a very interesting candidate solution to solve the major challenges described above. On the other hand, the European AGAVE project introduced the concept of parallel Internets to enable end-to-end service differentiation across multiple administrative domains. These are coexisting parallel networks composed of interconnected, per-domain network planes and allow for differentiation in terms of QoS, availability and resilience [19].

The rationale behind AGAVE is that SLAs are established between customers and service providers (SP), while connectivity-provisioning agreements are established between service providers and IP network providers (INP). Internally, SP have service-interconnection agreements. AGAVE defines a network plane as a logical partition of the network resources internal to an INP that is designed to transport traffic flows from services with common connectivity requirements for differentiated traffic treatment in terms of routing, forwarding (queuing and scheduling mechanisms) and resource management (admission control, traffic shaping and policies). Then, parallel Internets are inter-domain extensions of these network planes from the perspective of a single INP, and based on agreements with peer and/or remote INPs (INP interconnection agreements). Figure 3 depicts AGAVE’s functional architecture, which allows achieving network support for service differentiation without being service-aware.

Figure 3. AGAVE’s functional architecture. Extracted from [19].



4. Integrated Approach for QoS Assurance in IP/Optical Networks

In this section, we propose an approach to handle connectivity (transport) and service-level QoS in an integrated way. As described in Sections 1.1 and 2, research works in the literature handle service management and QoS assurance separately. Therefore, our aim is to tackle both issues at the same time and in an aligned way in the context of IP/optical networks, taking into account QoS from the

application layer to the physical (optical) layer. Transport-level QoS is assured by combining techniques to anticipate to traffic demands and impairment-aware lightpath provisioning (IA-RWA). Application-level QoS is assured with tailored service management, which is achieved by an adaptation of the service-oriented optical network architecture described in Section 2.1, coupled tightly with traffic-driven lightpath provisioning and IA-RWA.

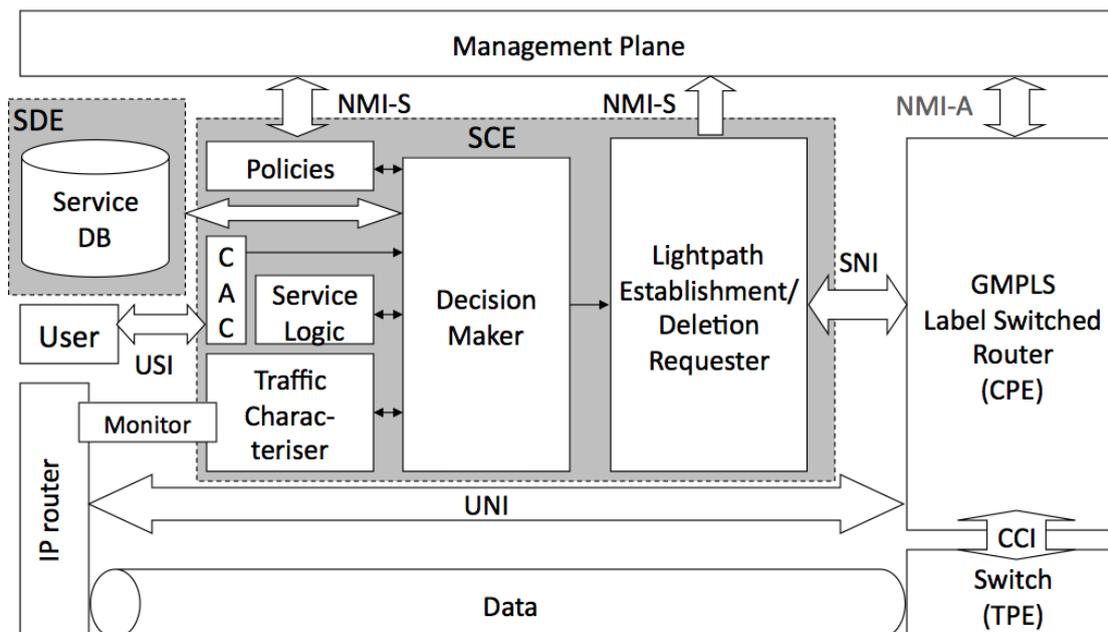
Besides assuring QoS in an integrated way, our proposed approach will enable new business models for IP/optical networks, such as provisioning as a service. That is, the offer of efficient, quality-enabled transport of user-driven services over IP/optical networks 'as a service' by enabling an interface for adaptation of these smart services to connectivity requests on backbone optical networks, especially GMPLS-managed (see USI in Figure 4). Such a model would decouple the innovation of services and service logic from protocols and network elements, support service portability between systems, with relocation and replication in the lifecycle, and support QoS assurance both in the setup and during the life of the connectivity services associated to the user-driven or smart services.

Figure 4 illustrates the proposed architecture to assure QoS in IP/optical networks that integrates the features of anticipation of traffic demands (in the form of traffic-driven lightpath provisioning) and impairment-awareness in the lightpath provisioning process, along with decomposition of services into lightpaths. This architecture is based on the service plane concept described in Section 2.1 and can be implemented in a centralized or distributed way, that is, with a single or several service plane nodes. If it is centralized, the service plane node contains the SCE and SDE elements highlighted in gray.

The proposed architecture will be implemented in a testbed infrastructure called ADRENALINE Testbed®, which is managed by the Centre Tecnològic de Telecomunicacions de Catalunya (CTTC). ADRENALINE's architecture was conceived to explore cost-efficient solutions to address the ever-growing demands of IP data traffic, together with the need for replacing the SDH nodes in Internet backbone networks. The testbed leverages reconfigurable wavelength-routed network functionalities with connection-oriented Ethernet technologies, which are both controlled by GMPLS-enabled IP/optical nodes. For further information about the testbed, the interested reader is referred to [20].

Generically, the functions of a service plane can be split in two logical elements: one containing information about the services (service data element, SDE) and another managing the translation of IP traffic and services/applications into provisioning requests. Note that these elements are named 'Service virtualization' and 'Service adaptation' in Figure 1. The SDE is essentially a database for service characterization. The other element, called service control element (SCE), is the one performing connectivity adaptation and service virtualization using information from the SDE, as well as QoS policies from the pre-established SLAs. These two elements are highlighted in gray in Figure 4. Note that the network customers ('User' in Figure 4) interact with the SCE by means of a user-service interface (USI), more precisely with the connection administration control (CAC) element of the SCE. The SCE also interworks with the management plane defined in the ASTN standard [8] through the network management interface-service (NMI-S) and with the control plane (labelled as GMPLS label switched router in Figure 4) through the service-network interface (SNI). Note also that, for completeness, Figure 4 also depicts the control and transport plane elements (CPE and TPE, respectively) defined in [8], as well as the interface between IP routers and GMPLS nodes (UNI), also contained in the ASTN standard.

Figure 4. Proposal for QoS-oriented service plane architecture.



4.1. Traffic-driven Lightpath Provisioning Based on the Service Plane

One of the key roles of the service plane is to forecast the behavior of traffic entering the IP/optical network and to automate the generation of provisioning requests by IP routers to the GMPLS control plane. This is based on applying the concept of self-managed networks described in Section 2.2. In a first stage, the service plane will be optimized for a wavelength-routed all-optical network, and then upgraded to Ethernet over this network. Taking this into account, the basic functionalities of an SCE, especially those related to the characterization of IP traffic into flows that can be translated into lightpath services are to be validated in the ADRENALINE Testbed®. In other words, the SCE is acting as the self-management engine of the proposed approach, and performs the functions depicted in Figure 2. The IP flows considered in this work include heavy-hitter flows and elephant IP flows, and are analysed and characterized by the ‘Traffic Characterizer’ component of the SCE. This component is fed with real-time measurement information from the IP router using one or more techniques, such as packet sampling, provided by the ‘Monitor’ element in Figure 4. These techniques are applied to the traffic already being transported in the network, so that proactive actions can be taken in a timely manner if it can be inferred from the measurements that the IP layer should be decongested, *i.e.*, by offloading. Note that most network standards focus on assuring QoS in the set up process, that is, for new services. For example, in Carrier Ethernet networks, Ethernet probing can be applied during provisioning to validate that Ethernet services are turned up according to the SLAs. In MPLS with Traffic Engineering (MPLS-TE) networks, end-to-end routing paths are set with QoS parameters before forwarding data, reserving resources in one aggregated class. In both cases, once the services have been set up, continuous SLA conformance reporting through proactive service and network monitoring is the basis for any detection of SLA violations and subsequent corrective action.

A key aspect for traffic-driven lightpath provisioning is the necessary trade-off between flow duration and decision time, which in the proposed approach will be calculated by the ‘Decision Maker’

component of the SCE taking into account the applicable QoS policies. Note that timeliness in decision making, *i.e.*, selecting a flow to be offloaded to the optical layer, is very important because if the decision is taken too soon there is a high probability that the flow is not going to last long enough, while if it is postponed, the flow will be consuming resources at the IP layer unnecessarily [1]. For example, Fioreze *et al.* propose to choose an elapsed duration of 5 minutes to define a flow as being persistent in time because flows with this or higher duration represent over 70% of the total traffic [1].

4.2. Automation of Connectivity Adaptation and Service Virtualization Based on the Service Plane

Another function of the service plane is to automate the connectivity adaptation and service virtualization necessary to handle complex services, such as a mesh of VPNs, prior to requesting the provisioning of connectivity that can be "understood" by the GMPLS control plane, *e.g.*, lightpaths or Ethernet bandwidth. This adaptation and virtualization can be performed in single- and multi-layer optical networks. This will be achieved in the ADRENALINE Testbed® using the reconfigurable wavelength-routed network functionalities and applying or not the connection-oriented Ethernet technologies on top. Here, the key element of our approach is the SDE, which is based on the adaptive multi-service optical network architecture proposed by Zhang *et al.* in [21]. The service database integrated in the SDE includes translation strategies from composite services, *e.g.*, VPN, into provisioning requests to the GMPLS control plane in charge of establishing and deleting connections, either lightpaths or Ethernet bandwidth. These translation strategies contribute to the common goal of achieving adaptive service provisioning through the accomplishment of service customization.

Figure 5. Adaptive decomposition of complex services. Extracted from [21].

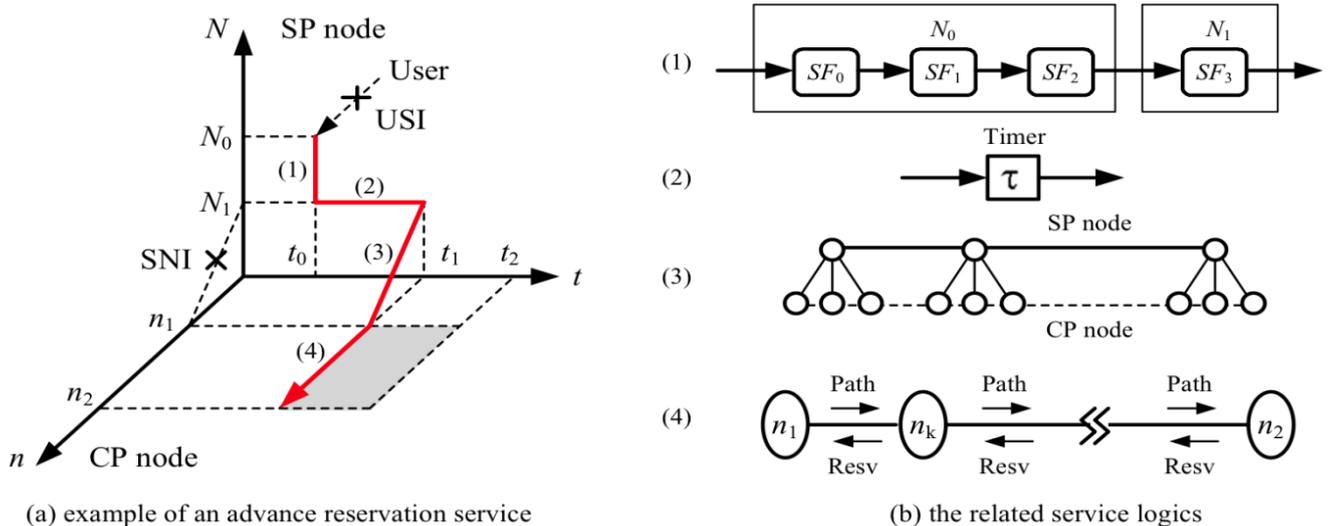


Figure 5 illustrates the rationale behind adaptive decomposition of complex services. Figure 5a depicts the advanced reservation service, which is similar to what human operators can do using suitable interfaces or what IP routers do using the UNI interface in ASTN [8]. Axis N represents nodes in the service plane, in this case in a distributed architecture (*i.e.*, SCE and SDE elements in different service-plane nodes). Axis n represents control plane nodes, while axis t represents the time. A user sends a request to service plane node N0 at t0 to reserve a lightpath between n1 and n2 from t1 to t2.

N0 (SDE element) analyses the request and maps it into four base service functions: the network information service (SF0), the network monitoring predictions service (SF1), the request transfer service (SF2) and the advanced connectivity service (SF3), which are plotted in Figure 5b. These functions match those depicted in the SCE and SDE elements illustrated in Figure 4. Note that the process flow is shown by a red solid arrowed line in Figure 5a, from steps (1) to (4).

Because of the diversity of service logics, the focus of the proposed approach is on decomposing complex services into lower-order services, *i.e.*, basic services that will translate directly into provisioning requests to the GMPLS control plane and will focus on a set of complex services. This involves, besides the SDE, the components ‘Service Logic’ and ‘Lightpath Establishment/Deletion Requester’ of the SCE depicted in Figure 4. Note that the latter component incorporates IA-RWA into the establishment request process implicitly by means of impairment-constrained based routing integrated in the GMPLS control plane, which is an implementation of IA-RWA techniques [22].

5. Conclusions and Future Work

IP/optical networks will not be promising unless they can guarantee a predictable performance as specified by the QoS parameters established in SLAs. This paper has proposed an integrated approach to assure QoS in the service provisioning process of IP/optical networks based on the service plane concept. This approach is fully automatic and can handle traffic-driven provisioning by anticipating traffic demands using IP traffic characterization and self-management concepts, as well as complex service demands through decomposing such services into sets of connection requests. In both cases, a control plane based on GMPLS can understand and perform the establishment and deletion requests issued by the service plane and apply routing-based IA-RWA techniques to the establishment process to enforce QoS assurance. Future work involves the implementation and performance validation of the proposed approach in the ADRENALINE Testbed®, and the evaluation of the benefits of this approach in the new business model of provisioning as a service for optical networks.

Acknowledgements

This work is part of the DORADO project, supported by the Spanish Ministry of Science and Innovation (MICINN) under the contract TEC2009-07995.

References and Notes

1. Fioreze, T.; Granville, L.; Sadre, R.; Pras, A. A Statistical Analysis of Network Parameters for the Self-management of Lambda-Connections. *Lect. Note. Comput. Sci.* **2009**, *5637*, 15–27.
2. Mori, T.; Takine, T.; Pan, J.; Kawahara, R.; Uchida, M.; Goto, S. Identifying heavy-hitter flows from sampled flow statistics. *IEICE Trans. Commun.* **2007**, *E90-B*, 3061–3072.
3. Mori, T.; Uchida, M.; Kawahara, R.; Pan, J.; Goto, S. Identifying Elephant Flows through Periodically Sampled Packets. In Proceedings of the ACM SIGCOMM Conference on Internet Measurement, Taormina, Italy, 25–27 October 2004.
4. Fioreze, T.; van de Meent, R.; Pras, A. An Architecture for the Self-management of Lambda-Connections in Hybrid Networks. *Lect. Note. Comput. Sci.* **2007**, *4606*, 141–148.

5. Ramamurthy, B.; Datta, D.; Feng, H.; Heritage, J.P.; Mukherjee, B. Impact of transmission impairments on the teletraffic performance of wavelength-routed optical networks. *IEEE/OSA J. Lightwave Tech.* **1999**, *17*, 1713–1723.
6. Kilper, D.C.; Bach, R.; Blumenthal, D.J.; Einstein, D.; Landolsi, T.; Ostar, L.; Preiss, M.; Willner, A.E. Optical Performance Monitoring. *IEEE/OSA J. Lightwave Tech.* **2004**, *22*, 294–304.
7. Martini, B.; Baroncelli, F.; Castoldi, P.; Muchanga, A.; Wosinska, L. The Service Oriented Optical Network (SOON) Project. In Proceedings of International Conference on Transparent Optical Networks, Barcelona, Spain, 3–7 July 2005.
8. ITU-T. *Requirements for Automatic Switched Transport Networks (ASTN)*; Recommendation G.807/Y.1302; International Telecommunication Union: Geneva, Switzerland, 2001.
9. Martini, B.; Campi, A.; Baroncelli, F.; Martini, V.; Torkman, K.; Zangheri, F.; Cerroni, W.; Castoldi, P.; Callegati, F. SIP-based Service Platform for On-demand Optical Network Services. In Proceedings of Optical Fiber Communication Conference, San Diego, CA, USA, 22–26 March 2009.
10. Figuerola, S.; Garcia-Espin, J.A.; Ferrer, J.; Willner, A. Scalability Analysis and Evaluation of the Multi-domain, Optical Network Service Plane in Harmony. In Proceedings of European Conference on Optical Communications, Vienna, Austria, 20–24 September 2009.
11. Takeda, T.; Inoue, I.; Shiomoto, K. Thin Layer for Seamless Interconnection of Multi-Technology Transport Networks. In Proceedings of International Conference on Optical Networking Design and Modeling, Kyoto, Japan, 31 January–3 February 2010.
12. Fioreze, T.; Oude Wolbers, M.; Van de Meent, R.; Pras, A. Characterization of IP Flows Eligible for Lambda-Connections in Optical Networks. In Proceedings of Network Operations and Management Symposium, Salvador de Bahia, Brazil, 7–11 April 2008.
13. Thompson, K.; Miller, G.J.; Wilder, R. Wide-Area Internet Traffic Patterns and Characteristics. *IEEE Network* **1997**, *11*, 10–23.
14. Kim, M.S.; Won, Y.J.; Hong, J.W. Characteristic Analysis of Internet Traffic from the Perspective of Flows. *Comput. Commun.* **2006**, *29*, 1639–1652.
15. Ribeiro, B.; Towsley, D.; Ye, T.; Bolot, J.C. Fisher information of sampled packets: An application to flow size estimation. In Proceedings of the ACM SIGCOMM Conference on Internet Measurement, Rio de Janeiro, Brazil, 25–27 October 2006.
16. Papagiannaki, K.; Taft, N.; Diot, C. Impact of flow dynamics on traffic engineering design principles. In Proceedings of Annual Joint Conference of the IEEE Computer and Communications Societies, Hong Kong, China, 7–11 March 2004.
17. Lan, K.-C.; Heidemann, J. A measurement study of correlations of internet flow characteristics. *Comput. Netw.* **2006**, *50*, 46–62.
18. Mas, C.; Tomkos, I.; Tonguz, O. Failure Location Algorithm for Transparent Optical Networks. *IEEE J. Sel. Area. Commun.* **2005**, *23*, 1508–1519.
19. Griffin, D. Network Planes and Parallel Internets: A Framework for Lightweight QoS Provisioning. In Proceedings of AGAVE/EMANICS Workshop on the Management of Network Virtualisation, Brussels, Belgium, 6 November 2007.
20. ADRENALINE Testbed®. Available online: http://wikiona.cttc.es/ona/index.php/ADRENALINE_testbed® (accessed on 30 August 2010).

21. Zhang, J.; Wang, L.; Chen, X.; Gu, W. AMSON: An Extended Architecture for Adaptive Service Provisioning in Transport Networks. In Proceedings of Optical Fiber Communication Conference, San Diego, CA, USA, 24–28 February 2008.
22. Martínez, R.; Pinart, C.; Cugini, F.; Andriolli, N.; Valcarenghi, L.; Castoldi, P.; Wosinska, L.; Comellas, J.; Junyent, G. Challenges and requirements for introducing impairment-awareness into the management and control planes of ASON/GMPLS WDM networks. *IEEE Commun. Mag.* **2006**, *44*, 76–85.

© 2010 by the authors; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).