PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Backbone network architectures for IP optical networking

Stefano Baroni

Phil. Trans. R. Soc. Lond. A 2000 **358**, 2233-2254
doi: 10.1098/rsta.2000.0644

10.1098/rsta.2000.0644

Backbone network architectures for IP optical networking

By Stefano Baroni†

Bell Laboratories, Lucent Technologies, 101 Crawfords Corner Road,
Holmdel, NJ 07733, USA

The unprecedented growth of Internet Protocol (IP) traffic is leading Internet Service Providers and network operators worldwide to investigate architectural alternatives for cost effective, reliable, scalable and flexible multi-terabit IP backbones. In this paper, several overlay, service and transport layer networking architectures, which employ IP, multiprotocol label switching, synchronous optical network/synchronous digital hierarchy and dense wavelength division multiplexing technologies, are proposed and analysed. Multiple parameters, such as network capacity, cost, restoration strategy, reconfigurability and accommodation of pre-emptable traffic, are considered for the architectural comparison. Detailed network design and economic analysis are provided for the different alternatives considering a typical nationwide US backbone with projected IP traffic in *ca.* 3 years. Several sensitivity analysis results are also shown, to evaluate the effect of cost changes in some of the critical technological factors in these architectures, such as 10 Gb s^{-1} optics cost or IP router cost. The results show the value of transport layer networking architectures for multi-terabit IP backbones, and how, when compared with service layer architectures, they provide additional desirable features such as wavelength reconfigurability and restoration scalability.

Keywords: optical networking; optical data networking; wavelength division multiplexing; wavelength routing; Internet Protocol; multiprotocol label switching

1. Introduction

With the unprecedented growth in data (and, in particular, Internet Protocol (IP)) traffic, the need for larger and more scalable IP backbone networks is seemingly limitless. However, several important factors have to be taken into consideration when designing an IP backbone, which contrast with traditional voice networks. First is the traffic demand distribution. Voice traffic derives from human-to-human interaction and, hence, it mainly translates into relatively short-distance connections (e.g. between neighbouring cities). Conversely, the global reach of the Internet eliminates the concept of distance, leading to traffic demand biased towards long-reach connections (OECD 1998). Moreover, it changes the traditional traffic demand fluctuation as a function of the time of the day. Second and most important is the high uncertainty in the forecast of IP traffic distribution. While voice traffic demand can be

---

† Present address: Corvis Corporation, 7015 Albert Einstein Drive, Columbia, MD 21046, USA (sbaroni@corvis.com).

Phil. Trans. R. Soc. Lond. A (2000) **358**, 2233–2254

easily derived by looking at changes in the population distribution, IP traffic is dominated by Web applications, i.e. human-to-Web server connections. Since these non-traditional traffic sources (the Web servers) are mobile and can be located anywhere, IP traffic distribution becomes very unpredictable. Other inherent characteristics of IP traffic are its burstiness and asymmetry. These critical requirements are leading Internet Service Providers (ISPs) and network operators worldwide to investigate architectural alternatives for scalable and flexible multi-terabit IP backbones. The latter must support the current traffic demands and cost-effectively absorb traffic growth.

Backbone networks currently designed and deployed for the transport of voice circuits are becoming less and less attractive for IP dominated traffic environments. In fact, synchronous optical networking (SONET)/synchronous digital hierarchy (SDH) ring architectures (Loyd Jones *et al.* 1999; Wauters *et al.* 1999) are much less efficient for the transport of long distance IP demands, given the reduced potential for sharing the ring restoration capacity. Moreover, it is important to note that high capacity data network elements (such as multi-terabit IP routers and asynchronous transfer mode (ATM) switches) are now becoming available, and provide statistical multiplexing and traffic aggregation at the edge of the backbone network, minimizing the need for SONET/SDH layer multiplexing (Doverspike *et al.* 1999). Finally, a complex provisioning mechanism may be required when multiple rings participate in establishing high capacity (e.g. wavelength) connections. Therefore, new technologies and architectures (including different restoration strategies) must be investigated to provide network operators with efficient alternatives for the design and deployment of their IP backbones.

Technological choices include wavelength division multiplexing (WDM) (see Loyd Jones *et al.* 1999; Wauters *et al.* 1999), ATM (see Wright & Yarlagadda 1999), IP (see Kaufman *et al.* 1999), and soon multiprotocol label switching (MPLS) (see Rosen *et al.* 1999). Network architecture alternatives include IP routers (or label-switching routers (LSRs), if MPLS is used) interconnected by point-to-point optical links, or via SONET/SDH or optical (WDM) ring, and optical (WDM) mesh infrastructures.

One of the key factors to reduce the overall network cost is the simplification of the protocol stack, as it enables network operators to minimize the number of network elements. The work in Anderson *et al.* (1999) addressed different possible choices for the protocol stack and provided a high-level architecture analysis. Similarly, the restoration mechanism is of primary importance for the network performance, operation and management, and, therefore, cost. Considerable work has been carried out to evaluate and compare different protection and restoration alternatives at the transport/optical (SONET/SDH and WDM) layer (see, for example, Gerstel 1998; Manchester *et al.* 1999; Johnson *et al.* 1999; Sato & Okamoto 1999), whereas the survivability problem in multi-layer networks is only at an initial stage (Demester 1999). However, in all these studies, no network design and quantitative evaluation of the different alternatives were provided.

The combined problem of network design and restoration has been carried out for WDM optical networks in Van Caenegem *et al.* (1998), Doshi *et al.* (1999) and Baroni *et al.* (1999). However, the considered traffic was given in node-pair wavelength-demand, and architectural requirements specific of IP services were not addressed. Furthermore, when IP traffic has been considered for the design and cost analysis of long-distance transport network architectures, such as in Doverspike *et al.* (1999) and

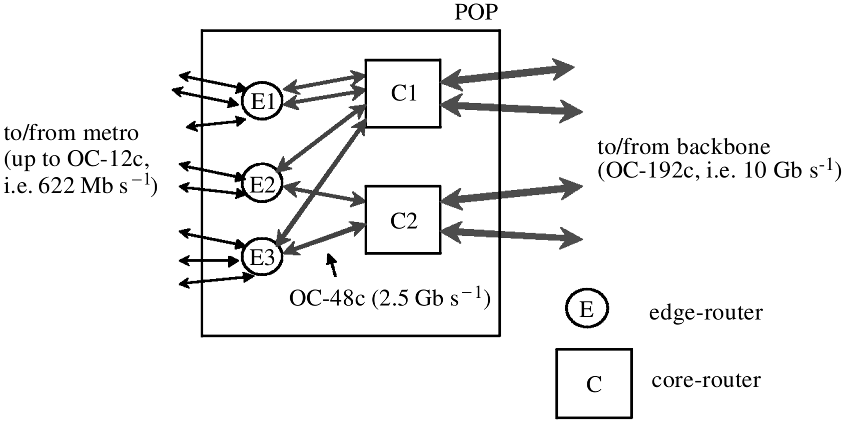

Figure 1. POP configuration (logical connectivity).

Doshi *et al.* (1998), this has always been only one component of the total network traffic (consisting also of circuit-switched services).

Therefore, fundamental analysis is required to evaluate and compare network architectures that are optimized for transport of IP traffic only. There are two main reasons for this, as follows.

(1) The bandwidth explosion driven by the popularity of the Internet has led to a paradigm shift in the telecommunications industry: from voice-optimized circuit-switched services to data-optimized packet-switched services. Many traffic forecast reports foresee that, around the year 2003, more than 90% of all the traffic will be data traffic (with IP as the predominant protocol).

(2) Network operators and bandwidth providers are deploying 'green-field' IP backbone networks.

In this work, several overlay, service-layer and transport-layer network architectures that are candidates for multi-terabit IP backbone applications are proposed and analysed. The architecture alternatives rely on IP, MPLS, SONET/SDH, and WDM technologies. A typical nationwide US backbone with projected IP traffic in *ca.* 3 years is considered for the network design and economic analysis. The architectures are compared according to multiple parameters, such as network capacity, cost, restoration strategy, reconfigurability and accommodation of pre-emptable traffic (Baroni *et al.* 2000*a*).

The rest of this paper is organized as follows. Section 2 introduces some preliminary considerations and concepts that are the basis of the architectural analysis. Sections 3, 4 and 5 provide detailed description of the proposed overlay, service and transport-layer networking architectures, respectively. In § 6, the results of network design and economic analysis are presented for a typical nationwide US backbone. The main conclusions of this work are summarized in § 7.

2. Architectural considerations

The network nodes, also referred to as points of presence (POPs), consist of multiple *edge* and *core* routers (figure 1). They are assumed to be collocated in the

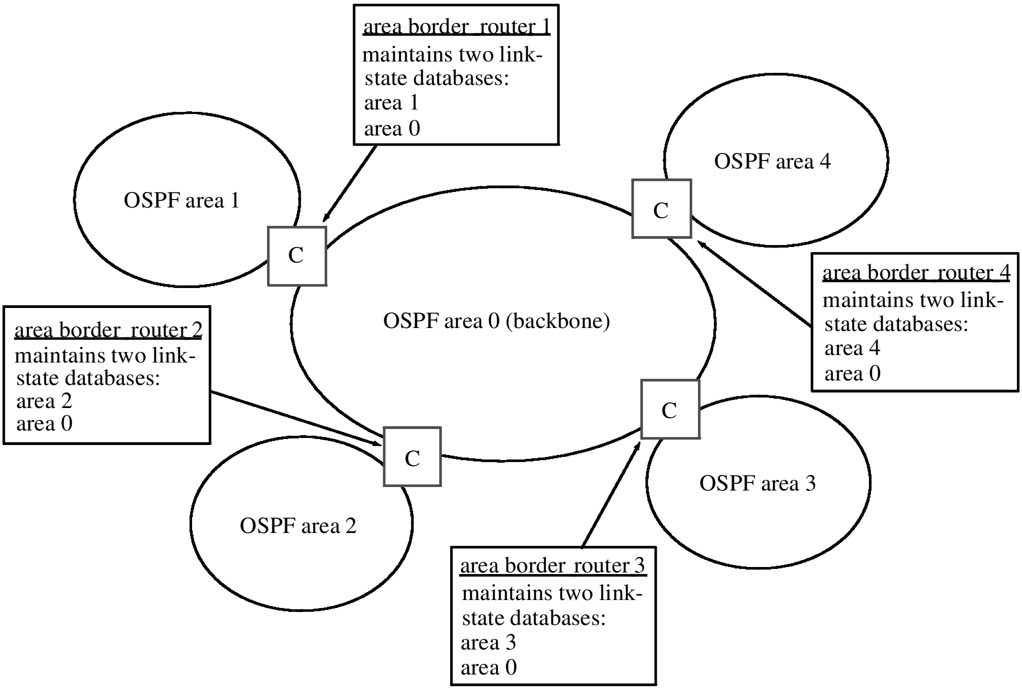

Figure 2. OSPF areas and area border routers.

same office, to avoid the large cost associated with traffic backhaul from edge to core routers. In this study, large IP backbone networks are considered with the traffic demand requiring multiple OC-192c (10 Gb s^{-1}) connections per POP pair.† Therefore the interconnections between core routers are assumed to be at OC-192c rate (see figure 1).

The edge routers face the metro and access portions of the network and collect low bit-rate tributaries (up to OC-12c, i.e. 622 Mb s^{-1}). The IP packets are then aggregated up to OC-48c (2.5 Gb s^{-1}) and forwarded to the core routers. Therefore, the minimum bandwidth from the edge routers into the backbone is OC-48c.

An *interior gateway routing protocol* (IGP), such as the *open shortest path first* (OSPF), runs within the ISP network. Core routers are not only critical for the functions of traffic aggregation and packet forwarding, but also to reduce the size of routing tables in edge routers and to reduce OSPF traffic across the backbone. As IP networks become larger, so does the link-state database at each router. As shown in figure 2, the routing table scalability problem is typically solved by implementing a two-level hierarchy, consisting of OSPF *areas* (Black 1999), where groups of (access and edge) routers have an 'area-centric' view of the network; only the core routers belonging to the same area, also called *area border routers* from the OSPF perspective, have knowledge beyond their area. They maintain both local (intra-area) and remote (to the backbone area 0) link-state databases.

---

† In the network design presented in §6, the traffic demand of the network considered was given in OC-48c (2.5 Gb s^{-1}) node-pair connections. However, the aggregated traffic resulted in multiple OC-192c (10 Gb s^{-1}) connections per POP pair.

The edge and core routers are assumed to provide MPLS functionality (Rosen *et al.* 1999), and, thus, are referred to as label-switching routers (LSRs). This feature affects the overall network performance, providing simplified packet forwarding, efficient explicit routing, and enabling traffic engineering and service differentiation. Moreover, from a network design viewpoint, MPLS allows us, for any pair of routers, to split the traffic across diverse (i.e. not necessarily equal cost) paths.

It is important to note that the aggregated traffic between any two edge routers in the backbone can actually consist of several label-switched paths (LSPs) (Rosen *et al.* 1999).† For example, a single OC-48c between two edge routers can actually carry dozens, hundreds, or even thousands of LSPs, depending on the bandwidth granularity of such LSPs.‡ From a network design and architecture analysis perspective, the number of LSPs that are served by a single transport connection is not important. However, it is a key parameter when considering scalability and performance of LSRs, and also with respect to the restoration performance and scalability of MPLS-based architectures.

Given the increasing capacity being processed by the core routers, multi-homing architectures are usually deployed as a solution to overcome their low reliability. This approach provides connectivity from each edge router to (at least) two core routers, ensuring service availability by means of route diversity. In this way, traffic from an edge router can find a second path out of the POP if one of the core routers fails. Dual-homing also allows traffic load balancing and non-disruptive, in-service software upgrade of core routers, not otherwise equipped with this feature.

Two different solutions can be implemented to exploit this dual-connectivity. In the first option, referred to as *dual-feeding*, the traffic from the source edge router is *simultaneously* fed to both core routers. For example, this can be performed at the MPLS layer by setting up two LSPs (i.e. primary and secondary) between source and destination edge routers—instead of only one LSP (primary)—which go through separate core routers at both (source and destination) POPs. At the receiving POP, the destination edge router (the egress LSR) receives and monitors primary and secondary LSPs and input selection is performed locally. In the second case, simply referred to as *dual-homing*, no dual-feeding is implemented, but the secondary LSP is used for the transport of pre-emptable traffic. Similar to the first option, both LSPs have the same ingress and egress edge routers and go through separate core routers at each POP. In a fault-free state, the destination edge router (the egress LSR) receives both primary and secondary LSPs, with the latter carrying pre-emptable traffic, therefore increasing the overall traffic transport efficiency. When the primary LSP fails, signalling is needed to coordinate pre-emption of restoration capacity and switching of service traffic to the secondary LSP, as described in § 4.

Routers with fully duplicated switching fabric and software controllers have recently been developed and become available on the market, referred to as *carrier class* or *ultra-reliable routers* (URRs) (see, for example, Opalka & Soman 1999). These routers provide protection against fabric and interface failures by implementing

---

† A label-switched path (LSP) is a path through one or more LRSs that is followed by IP packets receiving the same forwarding treatment within the MPLS network. At each LSR, packets belonging to a given LSP are routed according to the assigned label, without looking inside the packet header. MPLS enables us to create multiple levels of label hierarchy (Rosen *et al.* 1999).

‡ MPLS allows for a wide range in bandwidth granularity of LSPs, from a single application flow, which can require a few Mb s^{-1}, to a whole IP address prefix, which can require several Gb s^{-1}.

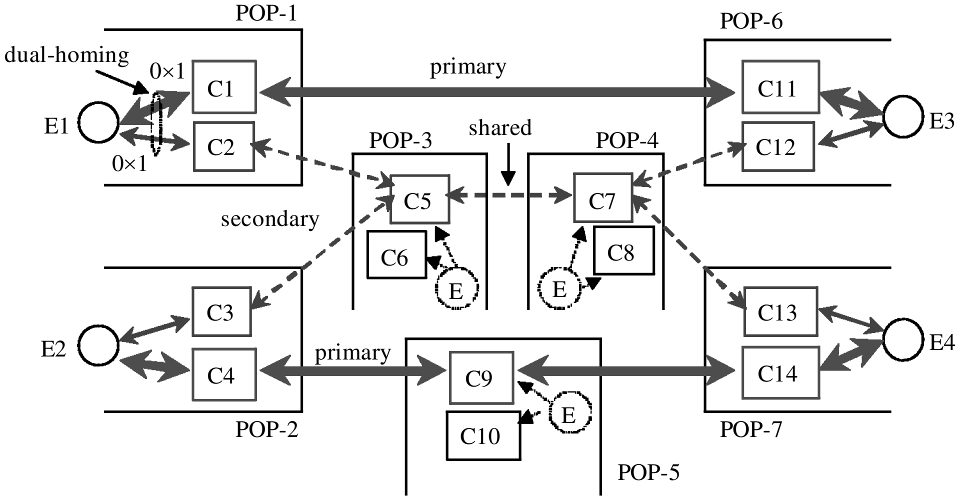

Figure 3. Service-layer sharing in conventional overlay architecture.

hardware hot-swap protocols and support in-service software upgrade. IP backbone architectures based on URRs will be discussed in § 5.

3. Overlay networking alternatives

In *overlay networking architectures*, the service (IP and MPLS) and transport (SONET/SDH, WDM) layers are designed separately, and their protection and restoration mechanisms are deployed and run independently from each other. Usually IP virtual topology and MPLS path layouts are designed first (even without any knowledge of the underlying fibre infrastructure), and the result translates into capacity demands for the design of the transport network. The service layer provides OSPF restoration and MPLS-based protection/restoration against router failures, whereas the transport layer provides SONET/SDH or WDM protection and restoration against transport node and link failures. Overlay networks are commonly deployed today when the ISP or network provider owns both the routers and the transport infrastructure, but their design, engineering, operation and maintenance, and evolution are assigned to separate organizations. These architectures are also common in the cases where the ISP (or network provider) owns only the router infrastructure and leases connections and capacity (and protection/restoration facilities) from other network or bandwidth providers. The advantage of overlay architectures is that they can provide a common transport infrastructure that accommodates a multi-service platform, for example when the ISP or network provider offers not only IP services, but also some other native data services (such as Frame Relay and ATM) and maybe voice.

Since restoration capacity is deployed at both layers, capacity inefficiencies may be obtained. Also, since the two layers are managed independently, multi-layer interactions between the various survivability mechanisms can occur (e.g. if MPLS-based protection/ restoration time-scales are close to SONET/SDH and WDM protection/restoration time-scales), unless the appropriate guard mechanisms are put in place.

(*a*) *Conventional overlay architecture*

To minimize the total capacity required by the service layer, *shared* restoration approaches are usually deployed at the IP/MPLS layers, i.e. the restoration capacity is shared among multiple node-pair connections. A possible implementation is presented in figure 3. This will be referred to as *conventional overlay architecture* throughout the paper. As shown, each edge router in the POP has dual-homing connectivity (via unprotected interfaces, also referred to as simplex or 0×1) to two core routers, which is used to establish primary and secondary paths that are *disjoint* at the MPLS layer, that is, the two LSPs can share the same fibre but belong to different wavelengths terminating in distinct routers. For example, the primary path between edge routers E1 and E3, in POP-1 and POP-6, is via core routers C1 and C11, whereas the (MPLS-disjoint) secondary path is via core routers C2 and C12 (the path also traverses core routers C5 and C7 located in POP-3 and POP-4, respectively). Similarly, the primary path between E2 and E4 is via C4, C9 and C14, and the secondary connection is via C3, C5, C7 and C13. As shown, the link between C5 and C7 is common to, and thus can be shared by, the restoration paths for edge router pairs E1–E3 and E2–E4. In fact, the network is designed to protect against *single* core router or edge router to core router interface failure. However, multiple failures can also be protected if the resources in their restoration paths are not shared.

The service-layer design is usually performed in two steps (via computer simulation): first, the POPs that are directly connected (i.e. one-hop away) at the IP/MPLS layer are identified. This is usually driven by their geographical location and the expected difficulty in providing a direct connection among them, and by capacity requirement. Second, this connectivity information is used to define primary and secondary MPLS paths for the POP pairs with traffic requirement. Disjoint primary and secondary shortest path (SP) connections that maximize sharing of restoration capacity are selected to minimize the total capacity requirement (Baroni *et al.* 2000*a*).

As presented in § 6, the capacity sharing achievable with this architecture results in significant reduction of the total number of network wavelengths with respect to the modified overlay approach described below. However, since the shared restoration capacity is accessed via the core routers, the interface cost between the core routers and the transport network elements (e.g. optical cross-connects (OXCs)) will play a key role on the cost of this architecture.

(*b*) *Modified overlay architecture*

An alternative approach is shown in figure 4, where restoration capacity sharing is achieved on a POP pair basis. Similarly to the previous case, each edge router in the POP has dual-homing connectivity (via 0×1 interfaces) to two core routers, to ensure primary and secondary paths that are *disjoint* at the MPLS layer. For example, the primary path between edge routers E1 and E4 is via core routers C1 and C5, whereas the secondary path is via core routers C4 and C7. As shown, for a given POP pair, the same core router is used (in each POP) for all the secondary paths (e.g. C4 in POP-i). This core router will be referred to as the *restoration* core router in POP-i for POP pair (i, j). In this way, all the remaining (*service*) core routers in the POP can be used to distribute the primary paths evenly, thus minimizing the impact of a core router failure on the restoration capacity (i.e. number of secondary paths)

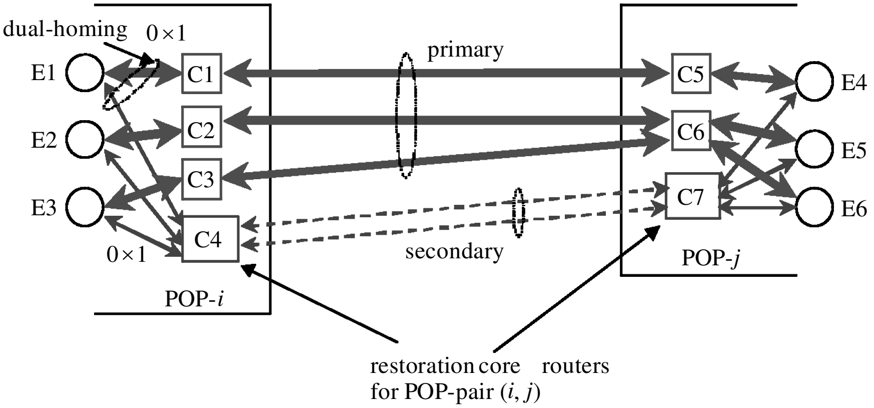

Figure 4. Service-layer sharing in modified overlay architecture.

needed for the POP pair. Throughout the paper, this approach will be referred to as *modified overlay architecture.*

Consider the example of figure 4, where POP-i and POP-j consist of $N_{\mathrm{CR}i} = 4$ and $N_{\mathrm{CR}j} = 3$ core routers, respectively.[†] Assume that the OC-48c demand between them translates to three OC-192c connections. As shown, the three OC-192c connections out of POP-i can be distributed among the $N_{\mathrm{CR}j} - 1 = 3$ service core routers. However, only two service core routers are in POP-j and, thus, one of them has to provide two OC-192c connections. Therefore, two secondary OC-192c connections are required between the POP pair (i, j) in order to protect against any core router failure in the source and destination POPs.

Therefore, the number of secondary OC-192c connections required for MPLS layer restoration for a given POP pair (i, j), S.OC-192$_{ij}$, can be written as (Baroni *et al.* 2000*a*)

$$\mathrm{S.OC\text{-}192}_{ij} = \max\left(\left\lceil \frac{\mathrm{P.OC\text{-}192}_{ij}}{N_{\mathrm{CR}i} - 1} \right\rceil, \left\lceil \frac{\mathrm{P.OC\text{-}192}_{ij}}{N_{\mathrm{CR}j} - 1} \right\rceil\right), \tag{3.1}$$

where P.OC-192$_{ij}$ is the number of primary OC-192c connections between POP pair (i, j), and $N_{\mathrm{CR}i}$ and $N_{\mathrm{CR}j}$ is the total number of core routers at POP-i and POP-j, respectively, and $\lceil x \rceil$ is the lowest integer larger than or equal to x.

This is based on the assumption, previously discussed, that the primary connections can be allocated among $N_{\mathrm{CR}} - 1$ core routers. In other words, although each core router can simultaneously act as both service and protection router, for any given POP pair, only one core router is always used (at each POP) for the secondary connections, as in figure 4. This condition can always be satisfied as long as the maximum demand among all the POP pairs is smaller (of the order of 50% or less) than the core router switching capacity, as for the network analysed in § 6. As shown in equation (3.1), the larger the number of core routers, the larger is the sharing of restoration capacity among them, which translates in a smaller number of secondary connections.

However, the number of core routers in any given POP depends not only on the total number of OC-48c add–drop traffic at the node, but also on the total number

---

[†] The nodes i and j may require a different number of routers because of a different amount of add/drop traffic.

of primary and secondary OC-192c connections out of the node, as described by the following equation:

$$N_{\mathrm{CR}i} = \left\{ \frac{2.(\mathrm{AD.OC}\text{-}48_i/4) + \sum_j (\mathrm{P.OC}\text{-}192_{ij} + \mathrm{S.OC}\text{-}912_{ij})}{\mathrm{RC}} \right\}, \tag{3.2}$$

where $\mathrm{AD.OC}\text{-}48_i$ is the total number of OC-48c add–drop connections at node i, and RC is the switching capacity of the core routers expressed in OC-192c (assumed equal for all of them).

As indicated by equations (3.1) and (3.2), the number of core routers and secondary connections are interrelated. Therefore, an iterative algorithm (via analytical calculation) based on the equations above was implemented to derive their values. The algorithm is described in details in Baroni *et al.* (2000*a*).

As shown in equation (3.2), the core router switching capacity RC is of primary importance in determining the number of core routers in the POPs.† In fact, the larger the switching capacity, the fewer the core routers deployed at each POP, and thus the more secondary connections needed. The influence of the core router capacity on the effectiveness of the modified overlay architecture was investigated by carrying out several network designs with different values of RC. It will be shown in § 6 that, for a given traffic demand, the value proposition of this architecture decreases as the router switching capacity increases.

The overlay designs previously described produce the number of OC-192c POP pair connections (primary and secondary) that are then used as capacity demand for the design of the transport layer. The latter is performed independently and several architecture alternatives are available (e.g. SONET and optical rings, and optical mesh), as described in § 5. In this analysis, an optical-mesh architecture was considered for the transport layer, for both overlay approaches, and the results are presented in § 6.

In both the overlay architectures previously described, the restoration capacity deployed at the service layer can be accessed by the edge routers (via the core routers) to carry pre-emptable traffic. The amount of pre-emptable traffic decreases as the capacity sharing increases. Also, the restoration capacity deployed at the transport layer (to protect against link failures) can be used for pre-emptable traffic. However, its accessibility depends on the interfaces between core routers and transport network elements (such as OXCs).

Since service and transport-layer failures are restored via separate mechanisms, the restoration performance of overlay architectures is twofold. According to the transport-layer architecture implemented, link failures will be restored via SONET/SDH or optical-ring protection or optical-mesh restoration. Restoration times of the order of 50 ms or a few hundred milliseconds are expected, respectively. Conversely, restoration times up to several hundred milliseconds or even a few seconds are expected for service-layer (core router and edge router to core router interface) failures, according to the IP/MPLS mechanism deployed and the number of interrupted LSPs.

---

† In the conventional overlay architecture previously described, the size and, thus, number of core routers in each POP has no effect on the service-layer design, as long as at least two core routers are available (see figure 3). However, this condition is always satisfied since it is the key requirement to ensure IP/MPLS path diversity.

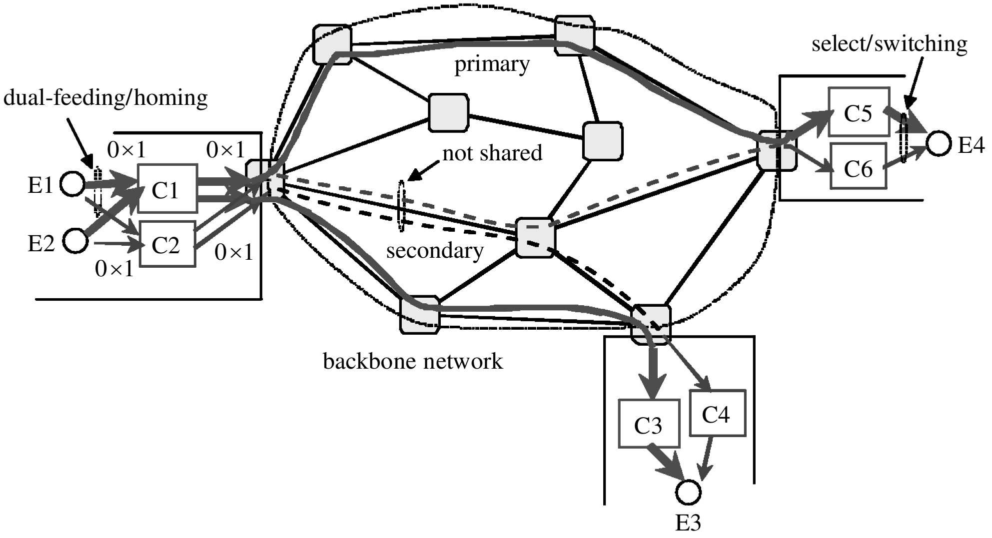

Figure 5. Service-layer networking architecture: 1 + 1 and 1:1 dedicated MPLS protection.

Overlay architectures provide two levels of bandwidth reconfigurability. At the transport layer, the level of wavelength reconfigurability is maximized as OXC-based mesh architectures are considered. At the IP/MPLS layer, packet level reconfigurability is available between the POPs where the OC-192c connections are terminated at the core routers.

4. Service-layer networking architectures

In *service-layer networking architectures*, networking functions such as routing and switching for service provisioning and restoration are performed at the service layers (IP and MPLS) by the edge and core routers (Makam *et al.* 1999). The only network elements deployed within the transport layer are WDM optical line systems (OLSs), to provide point-to-point OC-192c connectivity between the core routers. No additional networking functionality is assigned to the transport layer, and protection/restoration at the transport layers is *not* implemented. In addition to OSPF restoration, several MPLS-based protection and restoration solutions can be implemented: 1 + 1 and 1:1 dedicated protection, and shared restoration. As discussed below, knowledge of the fibre layout is key for the implementation of these approaches.

Service-layer networking architectures are now being considered by ISPs and network operators as an (often as more cost-effective) alternative to overlay networking architectures. Capacity efficiency is improved and the possibility of multi-layer survivability interactions is eliminated. However, in this case, IP/MPLS restoration mechanisms are used *also* in the case of link failures, which are critical given the extremely large number of interrupted LSPs. Therefore, restoration performance and scalability of service-layer networking architectures are crucial issues still to be determined.

(*a*) 1 + 1 *dedicated MPLS protection*

This restoration approach is depicted in figure 5. Each OC-48c signal out of the edge routers (carrying the LSPs between ingress and egress LSRs) is dual-fed into the duplicated core routers (via 0 × 1 interfaces), aggregated in OC-192c channels and diverse routed, via node and link disjoint physical paths, to the core routers at the destination end. Therefore, knowledge of the fibre infrastructure is required to ensure *physical* path diversity. At the receiving end, the two copies of the original OC-48c signal are delivered to the destination edge router, where a simple MPLS-based protection mechanism is implemented to provide restoration to failures in the primary LSP and core routers (Makam *et al.* 1999). The protection time is expected to be proportional to the difference in delay offered by primary and secondary (protection) LSPs. Since typical variance of the propagation time across a national backbone network is of the order of tens of milliseconds, the restoration time for the 1 + 1 dedicated MPLS protection architecture is expected to be of the same order.

The 1 + 1 dedicated MPLS protection is achieved by diverse LSPs and no sharing of restoration capacity is attained. Therefore, in this architecture, intermediate nodes perform *static* wavelength routing, interconnecting wavelength channels from incoming to outgoing OLS systems. Since switching functionality is not required at intermediate nodes, the simplest approach is to have through channels (optically) bypassing the core routers via patch-panel connections. This leads to a large saving in router ports and switching capacity, resulting in the cheapest way to implement this architecture. This is the case considered in the network design carried out in this work. However, this architecture provides only limited packet- or MPLS-level reconfigurability via the routers at the source and destination POPs, but no wavelength reconfigurability is available, a feature becoming increasingly important for network operators. Moreover, since the traffic is simultaneously fed to both primary and secondary paths, the 1 + 1 dedicated MPLS architecture does not support any pre-emptable traffic.

(*b*) 1:1 *dedicated MPLS protection*

As for the 1 + 1 dedicated MPLS protection architecture, in the 1:1 case each OC-48c connection is assigned physically disjoint primary and secondary paths. However, in this case, in fault-free state, the secondary LSP carries pre-emptable traffic so that the destination edge router receives both primary and pre-emptable traffic, increasing the network transport efficiency. In the case of a failure in the primary LSP or core router, the destination (egress) edge router switches to the secondary LSP after having requested the source (ingress) edge router to perform the same operation. This results in the pre-emptable traffic being dropped (Makam *et al.* 1999).

This approach requires a MPLS signalling mechanism between the ingress and egress edge routers, and thus the restoration time is expected to be longer than the 1 + 1 case previously described (and may be of the order of up to a few seconds). This is a penalty to pay for increased bandwidth efficiency. However, as the traffic demand increases, this solution may become extremely complex or even infeasible from a management viewpoint, due to the increasing amount of edge-router signalling that could flood the network in the case of a fibre cut.

Similar to the 1 + 1 case, the 1:1 dedicated protection approach allows for multiple simultaneous link or node failures to be restored, as long as, for each edge

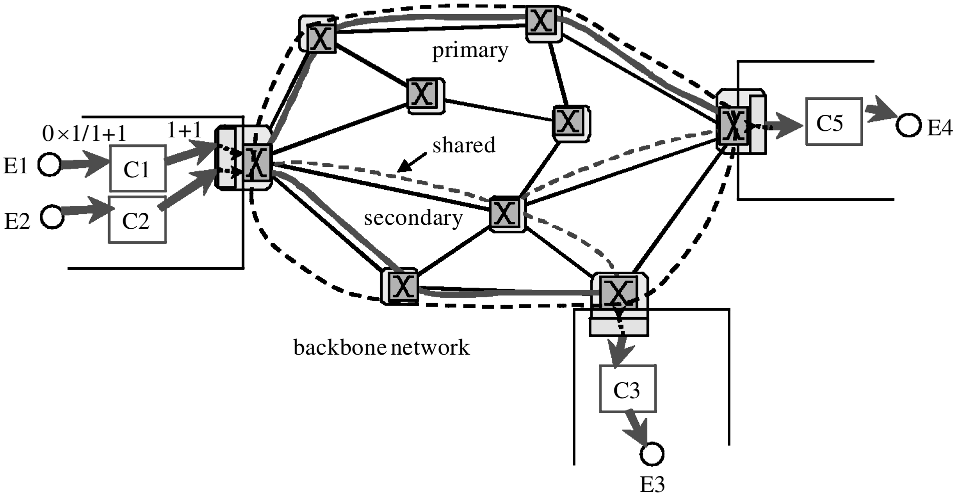

Figure 6. Transport-layer networking architecture with URRs: shared restoration.

router to edge router LSP, only one link or node between the primary and secondary paths is involved in the failures. Moreover, the 1:1 architecture provides the same reconfigurability characteristic of the 1 + 1 case.

From a network design perspective, the 1 + 1 and 1:1 cases are the same architecture, and, thus, will be considered as one only in § 6.

5. Transport-layer networking architectures

In *transport-layer networking architectures*, the transport layer not only provides point-to-point connectivity between core routers, but also adds networking functions such as switching and routing for service provisioning and protection/restoration. To this aim, digital (i.e. SONET/SDH) or optical (i.e. WDM) add–drop multiplexers (ADMs) and cross-connects (XCs) are deployed together with the OLSs.

These architectures are based on single-homing POP architectures that rely on ultra reliable routers, URRs (discussed in § 2). In this case, each connection out of an edge router is homed to a single URR (see figure 6). A protected 1 + 1 interface may be selected in order to ensure interface failure protection. (Similarly, protected interfaces will most likely be used between the core routers and the transport network element.) Single-homing architecture will likely lead to simplified network operation and management compared with the case with dual-homing POP architectures. (It is important to note that the network design analysis for transport-layer architectures presented in this work also holds in the case of a novel dual-homing POP architecture (based on duplicated core router and select function) proposed in Baroni *et al.* (2000*a*).)

As shown in figure 6, the backbone network provides transport-layer restoration against a single link or node failure. In this analysis, only shared transport restoration approaches are considered for the backbone (i.e. inter-POP) portion of the network, namely optical-mesh restoration and SONET/SDH and optical-ring protection, as described below.

(a) Optical shared mesh architecture

In the optical shared mesh architecture, each edge router to edge router connection is assigned primary and secondary paths that are *physically* link- and node-disjoint among themselves. Multiple link- and node-disjoint primary paths can share the restoration capacity among their secondary paths, thus reducing the total network restoration capacity (Van Caenegem *et al.* 1998; Doshi *et al.* 1999; Baroni *et al.* 1999). This is made possible by deploying wavelength switching elements (i.e. OXCs) at the network nodes. Most likely, the OXCs will be deployed not only at the service nodes, but also at the transport hubs, to provide network reconfigurability. In this way, the restoration capacity can be accessed via all the network nodes, not only the service nodes, resulting in increased sharing and, therefore, reduced restoration capacity.

The primary paths can still optically bypass the OXCs at the intermediate nodes, whereas the secondary paths need to be terminated as switching is needed for sharing the restoration capacity. However, the primary paths will also be terminated if network reconfigurability is necessary. This is the case considered in this analysis (see figure 6). The core routers are always bypassed at the intermediate nodes, by both primary and secondary connections.

Restoration time of the order of several tens of milliseconds (or up to a few hundred milliseconds) is expected, according to network size and traffic demand (Agrawal *et al.* 2000).

(b) SONET/SDH ring architecture

In this approach, the backbone portion of the network consists of 4-fibre SONET bi-directional line switched rings (4F SONET BLSR) at OC-192c rate. Therefore, in this architecture, the transport network elements are SONET/SDH add–drop multiplexers. The primary paths are expressed through the intermediate nodes by means of an optical patch-panel. Although limiting the network reconfigurability, this approach tends to minimize the number of ADMs needed.

The traffic hand-off between rings can be implemented in several ways, according to the level of reconfigurability required. To provide reconfigurability, in this analysis the two ADMs are assumed to be interconnected via a digital cross-connect. Similarly, the traffic entering and leaving the backbone network is passed through the XC. This choice reflects the market availability of a single SONET/SDH network element with both ADM and XC functionalities.

SONET/SDH ring protection guarantees very fast restoration time (of the order of 50 ms). Moreover, multiple failures can still be restored if they belong to different rings. Another important feature of 4F rings, as considered here, is that the restoration capacity can be utilized to perform span switching, to enable local maintenance on a fibre basis without the need for network-wide connection rerouting.

SONET/SDH mesh restoration is not considered here given its limited deployment by network operators.

(c) Optical channel shared protection ring architecture

As opposed to multiplex section approaches (SONET or optical MS/SPRing for example), the optical channel shared protection ring (OCh/SPRing) is designed to protect against failures on a per-optical channel (OCh) basis (Manchester *et al.*

1999). Key to this implementation is the monitoring of OCh-level failure indication signalling, for example via OCh overhead. This also leads to another advantage with respect to SONET MS/SPRing, that is, a given OLS can be shared by multiple OCh/SPRings belonging to different physical rings. Moreover, it is important to note that the different OLSs belonging to a given ring can be equipped differently, according to the individual capacity requirement; that is, OCh/SPRing architectures allow service providers to invest as they grow.

Two and four fibre OCh/SPRings can be deployed according to the availability of 2F and 4F OLSs, respectively. In this analysis, 4F OCh/SPRing with OC-192c wavelengths are considered. The transport network elements are optical add–drop multiplexers and ring switches (ORSs). Similar to the SONET/SDH case, the primary paths are expressed through the intermediate nodes by means of optical patch-panel. Again, different approaches can be considered for the traffic hand-off between rings. In the case considered here, the two ORSs are interconnected via an optical patch-panel providing $1+1$ dedicated protection between the two. Alternatively, an optical cross-connect (OXC) can be placed between the ORSs or even for traffic add–drop to increase wavelength reconfigurability.

Optical rings are expected to provide a similar restoration time to SONET/SDH rings (of the order of 50 ms). Also in this case, multiple failures can still be restored if they belong to different rings.

In transport-layer networking architectures, restoration capacity can be utilized for the transport of pre-emptable traffic. In the case of POP architecture based on URRs, the protection channel of the $1 + 1$ interface from edge router (via core router) to the transport network element (e.g. OXC) is used to insert the pre-emptable traffic.

6. Network design and economic analysis

A typical nationwide US backbone with projected IP traffic in *ca.* 3 years was considered for the network design and economic analysis. The network consists of about 50 nodes, considering both service POPs and transport hubs, with a POP-to-POP demand at the OC-48c level and aggregated demand of *ca.* 12 Tb s^{-1} (on average each node pair had two OC-192c connections) (Baroni *et al.* 2000*b*). Two-fibre (2F) mesh designs were carried out for the service IP/MPLS and transport OXC cases (and, thus, also for the overlay architectures), whereas four-fibre (4F) designs were performed for the SONET/SDH and optical ring transport cases. OLS systems with up to $80 \times 10\ \text{Gb s}^{-1}$ 4F-wavelengths (or equivalently $160 \times 10\ \text{Gb s}^{-1}$ 2F-λs) were considered, with optical amplification every 50 miles and digital regeneration every 250 miles. Two different sizes were assumed available for the integrated SONET/SDH ADM/XC: 32 and 128 4F $10\ \text{Gb s}^{-1}$ ports, respectively. Similarly, OXC with 256 and 1024 bidirectional (i.e. 2F) ports were considered.† The switching capacity of the core routers was assumed to be $\text{RC} = 480\ \text{Gb s}^{-1}$, corresponding, for example, to 96 OC-48c ports facing the edge routers and 24 OC-192c ports facing the backbone network.‡

---

† The size of the OXC does not depend on the bit-rate of the input/output signals, as its optical switching fabric is transparent to the bit-rate.

‡ As previously discussed, the modified overlay is the only architecture where the core router capacity directly impacts the backbone design, as it influences the number of secondary connections (see equation (3.2)). Therefore, two different designs were carried out for the modified overlay architecture, considering $\text{RC} = 480\ \text{Gb s}^{-1}$ and $\text{RC} = 1\ \text{Tb s}^{-1}$, respectively.

A near-optimal network design tool was utilized for the network design (Doshi *et al.* 1995). For both overlay architectures presented in § 3, an OXC-based optical mesh (with shared restoration) was designed for the transport layer. In the case of service-layer networking architectures, no distinction exists between 1 + 1 and 1:1 MPLS dedicated protection from a backbone design (i.e. resource allocation) viewpoint. Also, the same ring design was used for both the SONET/SDH and optical ring transport architectures.

Some of the key design results are presented in table 1. Total costs are normalized and expressed relative to the conventional overlay network architecture. The major cost components modelled in this analysis included equipment interfaces (I/Fs), IP and transport-layer switch fabrics, optical transponder units (OTUs, i.e. opto-electronic digital 3R regenerators), OLSs and line repeaters (WDM amplifiers). However, the cost of fibre in the ground was not included. Neither was the edge routers cost, which is common to all the architectures. The projected cost of each component was derived by applying multiplying factors to estimate the technology cost reduction. (To study the dependency of the results on the cost assumptions, a sensitivity analysis was carried out for the main components, as described below.)

Consider the three overlay architectures first. As shown, the modified overlay with RC = 480 Gb s^{-1} and RC = 1 Tb s^{-1} requires *ca.* 2% and *ca.* 18% more OC-192c wavelengths than the conventional overlay, respectively. A similar trend applies if wavelength-miles are considered. Therefore, service-layer restoration capacity sharing at the *network level* (as in the conventional overlay) is more efficient than that at a *POP pair level* (as in the modified overlay) for RC $\geqslant$ 480 Gb s^{-1}. (A design of the modified overlay architecture with a smaller value of RC (e.g. RC = 240 Gb s^{-1}) would probably result in better restoration capacity sharing.) Thus, fewer OLS end terminals, line repeaters and add–drop, through and regeneration OTUs† are needed, as shown in table 1. However, in the conventional overlay architecture, the (service-layer) restoration capacity sharing is achieved by terminating the OC-192c wavelengths into the core routers at the intermediate nodes, as discussed in § 3. Therefore, a large number of expensive OC-192c interfaces and router ports are required to pass the wavelengths from the OXCs (transport layer) to the core routers. Similarly, a large number of core router fabrics are needed (326 in the table). For this reason, the conventional overlay was found to be more expensive than the modified overlay architectures. However, as shown, for a fixed traffic demand the value proposition of the modified overlay decreases as the router switching capacity increases: the cost savings with respect to the conventional overlay decrease from 17% to 6% as RC increases from 480 Gb s^{-1} to 1 Tb s^{-1}.‡

Consider now the service and transport-layer networking architecture results of table 1. As shown, the service-layer architecture and the (SONET/SDH and optical) ring cases result in *ca.* 25% and *ca.* 64% more OC-192c wavelengths and OC-192c

---

† Add–drop OTUs convert 1.3 μm SONET short reach optical signals to 1.5 μm WDM-compatible wavelengths and vice versa. OTUs used to regenerate 1.5 μm WDM-compatible wavelengths are referred to as through OTUs or regeneration OTUs according to whether they are located at an intermediate node or at a regeneration site (every 250 miles in the network links), respectively.

‡ The modified overlay architecture was also designed considering the extreme case with *only* two core routers in each POP (the value of RC to satisfy this condition was *ca.* 4 Tb s^{-1}). In this case, no sharing of restoration capacity is achievable at the service layer. (From equation (3.1), it is clear that the number of secondary wavelengths is equal to the number of primary wavelengths for $N_{\text{CR}} = 2$.) This modified overlay was found to be 7% more expensive that the conventional overlay.

Table 1. *Network design results*

(The headings in the first column are (1) total 2F OC-192c wavelengths; (2) total 2F OC-192c wavelengths × miles; (3) average 4F OC-192s per link; (4) 4F OLS end terminals; (5) 4F repeaters; (6) 4F add–drop OTUs; (7) 4F regenerator + through OTUs; (8) integrated SONET/SDH ADM+XC 32 4F OC-192c ports; (9) integrated SONET/SDH ADM+XC 128 4F OC-192c ports; (10) 4F OC-192c WDM compatible TX/RX pair; (11) OXC–256 2F ports; (12) OXC–1024 2F ports; (13) ORS; (14) OC-192c 0×1 I/F; (15) OC-192c 1+1 I/F; (16) core routers (480 Gb s^{-1}); (17) core routers (1 Tb s^{-1}); (18) core router OC-48c ports; (19) core router OC-192c ports; (20) normalized total costs.)

case	conventional overlay	modified overlay (480 Gb s^{-1} router)	modified overlay (1 Tb s^{-1} router)	service layer: 1 + 1 or 1:1 MPLS	transport layer: optical mesh (OXC)	transport layer: SONET/SDH rings	transport layer: optical rings (OCh/SPRings)
(1)	16 648	17 062	19 588	14 853	11 285	18 518[†]	18 518[†]
(2)	6.67M	7.21M	8.31M	6.39M	4.81M	7.12[†]	7.12[†]
(3)	113	115	132	100	76	127	127
(4)	626	638	714	572	494	620	620
(5)	1 067	1 102	1 243	994	846	887	887
(6)	16 648	17 062	19 588	2 360	11 285		4 202
(7)	9 822	10 868	12 544	15 920	7 282	17 209	17 209
(8)						8	
(9)						61	
(10)						4 202	
(11)	4	1	1		10		
(12)	63	62	67		40		
(13)							2 101
(14)	10 716	3 648	4 146				
(15)					2 360	2 360	2 360
(16)	326	171		175	175	175	175
(17)			98				
(18)	14 596	14 596	14 596	14 596	14 596	14 596	14 596
(19)	10 716	3 648	4 146	4 720	4 720	4 720	4 720
(20)	1.00	0.83	0.94	0.63	0.65	0.81	0.73

[†] These network designs were done on a 4F basis, but the units are expressed on a 2F basis for consistency in this table.

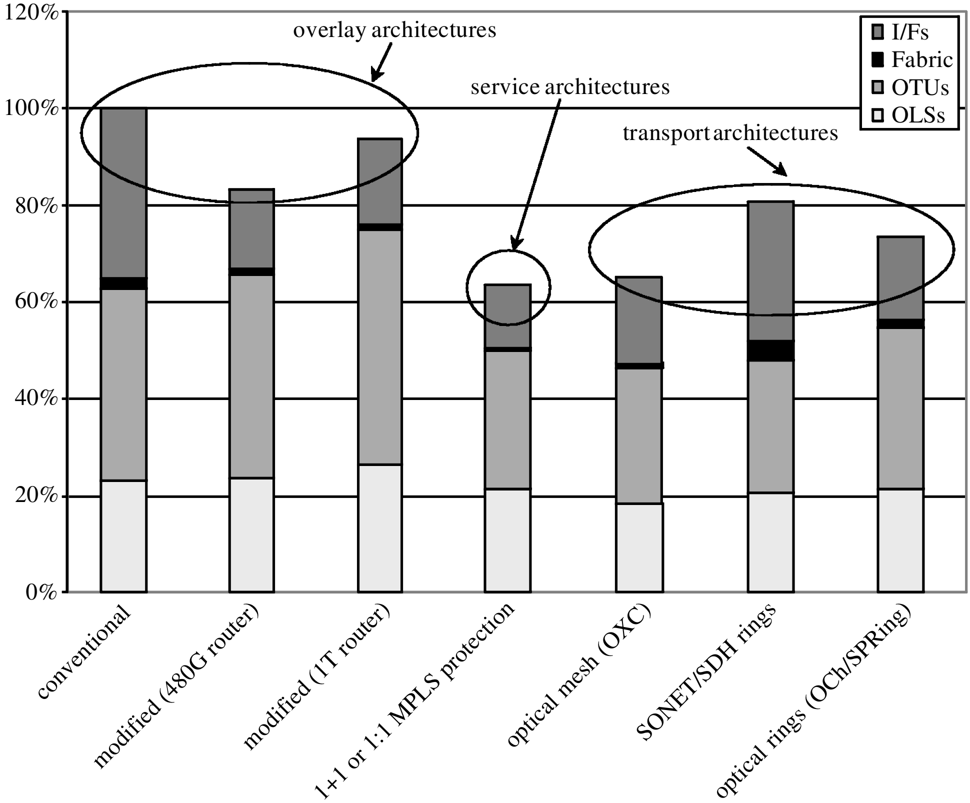

Figure 7. Relative cost and components breakdown
(relative to conventional overlay architecture).

wavelength-miles than the optical-mesh case, respectively. This demonstrates the large wavelength sharing achievable by OXC-based mesh restoration, as reflected by the smaller number of OLS terminals and repeaters. Also, it is important to note the large difference (24%) in the *average* number of 4F wavelengths in a network link between the service-layer and optical-mesh architectures. This is the number of channels to be protected in the case of a link failure (by local selection or end-to-end switching in the $1+1$ and 1:1 MPLS dedicated protection case, respectively, and by end-to-end wavelength re-routing in the optical-mesh case). Since, in the optical-mesh architecture, each wavelength is terminated in the OXC at each node, the number of 4F add–drop OTUs is equal to the total number of OC-192c wavelengths in the network links. Conversely, in the service-layer architecture, the OC-192c wavelengths are transferred (at the intermediate nodes) from one OLS to the next via patch panel. Therefore, a large number of through OTUs are needed at the intermediate nodes. The number of 4F add–drop OTUs is given only by the wavelengths added and dropped at the end service POPs, which is based on the traffic demand.

All the service and transport-layer architectures require the same number of OC-48c and OC-192c ports in the core routers. Also, transport architectures based on POP configuration with duplicated core routers and select function require the same number of core routers as the service-layer architecture. This is the case shown in table 1. If POP architectures with URRs were implemented, the number of core routers would halve. However, the cost of an URR (with duplicated switching fabric and software controller) can be considered as approximately twice the cost of a simple router. Therefore the result of the cost analysis for the transport-layer architectures holds also for POPs based on URRs.†

Table 1 shows that the service-layer and the optical-mesh architectures are the lowest in cost, and *ca.* 35% less than the conventional overlay. The two architectures can be considered the same in cost, as the 2% difference is within the tolerance range of the design results. However, the analysis would greatly favour the transport architecture if the fibre costs (e.g. right of way and/or installation costs) were considered, since the service-layer architecture requires *ca.* 30% more OC-192c wavelength-miles than the optical mesh. In addition, the greater OC-192c wavelengths in the service-layer case results in a larger amount of spare equipment (lasers, amplifiers, ...) and larger maintenance cost. Finally, it is important to note that the optical mesh provides *full* (OXC-based) wavelength-level reconfigurability, an extremely desirable feature for large backbone carriers. Therefore, this architecture is expected to offer lower operation costs and shorter intervals for wavelength provisioning. The maximum amount of pre-emptable traffic that can be inserted in the OXC mesh architecture is *ca.* 60% of the total service traffic.

The SONET/SDH and optical ring architectures were found to cost *ca.* 19% and *ca.* 27% less than the conventional overlay, respectively. As discussed in § 5, they provide partial reconfigurability. However, it is important to note that the 4F ring architectures provide span switching capability for maintenance purposes, which is not guaranteed for the 2F mesh designs.

Figure 7 shows the relative cost component breakdown by the four major categories: interfaces (I/Fs), fabric, OTUs and OLSs. All percentages in the figure are relative to the conventional overlay case. Interface costs include all the port costs, whereas the fabric cost include OXC, integrated SONET/SDH ADM+XC, ORS and IP router fabrics. The OTU costs include all add–drop, through and regeneration OTUs. Finally, the OLS costs include the end terminals and repeaters. As shown, the OTU costs account for about one-third to almost one-half of the total cost in each case. The remaining cost is almost equally shared between interfaces and OLSs. Also, it is important to note that the fabric cost is a very small factor in the network cost for all the architectures considered: at most it accounts for 4% in the SONET/SDH ring case.

Given the importance of the OTUs on the network cost, a sensitivity analysis was carried out to evaluate how the relative costs would change if the initial cost assumptions were not met. However, the OTUs are based on 10 $\mathrm{Gb\,s^{-1}}$ optics that is also present in the OXC and ORS I/Fs, and, to a lesser extent, in the integrated SONET/SDH ADM+XC interfaces. Therefore, a general sensitivity analysis to changes in the cost of 10 $\mathrm{Gb\,s^{-1}}$ optics was carried out. As shown in figure 8, the service-layer and optical-mesh and ring architectures decrease in cost relative to the conventional overlay case when a 50% decrease in 10 $\mathrm{Gb\,s^{-1}}$ optics costs is applied, whereas the SONET/SDH case increases. Also, in this case, the service-layer and optical-mesh network costs are the same. This reflects the greater percentage of

---

† As it will be discussed below, the core router fabric cost is a very small percentage of the total network cost (less than 1%, see figure 7). Therefore, even large variations in the URR cost have little impact on the total network cost.

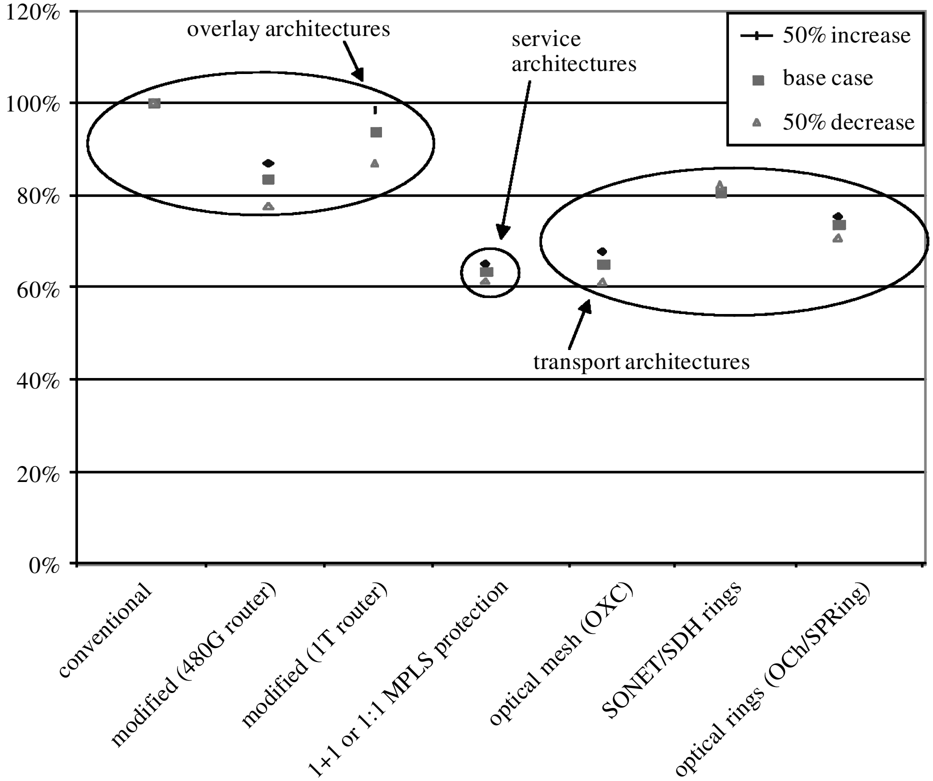

Figure 8. Sensitivity to changes in 10 Gb s^{-1} optics cost (relative to conventional overlay architecture).

10 Gb s^{-1} optics in the optical-mesh case. For 10 Gb s^{-1} cost decreases greater than 50% (not shown here), the optical-mesh architecture becomes the lowest in cost.

A sensitivity analysis was also performed for the core router costs. As shown in figure 9, the overall costs relative to the conventional overlay case are largely insensitive to IP router cost changes for the service and OXC cases and move up somewhat for the SONET/SDH and optical ring cases as IP costs decrease. This trend reflects the relatively larger percentage to total cost that IP routers contribute in the overlay case.

7. Conclusions

The IP traffic demand is experiencing a dramatic growth worldwide. However, given the global reach of the Internet, IP applications present key differences with respect to traditional voice traffic: the distribution is more biased towards long-reach connections and the traffic forecast is much less predictable. These requirements are driving ISPs and network operators to investigate architectural alternatives for scalable and flexible IP backbones.

In this work, several overlay, service and transport-layer network architectures were analysed and compared for multi-terabit IP networks. The comparison was based on

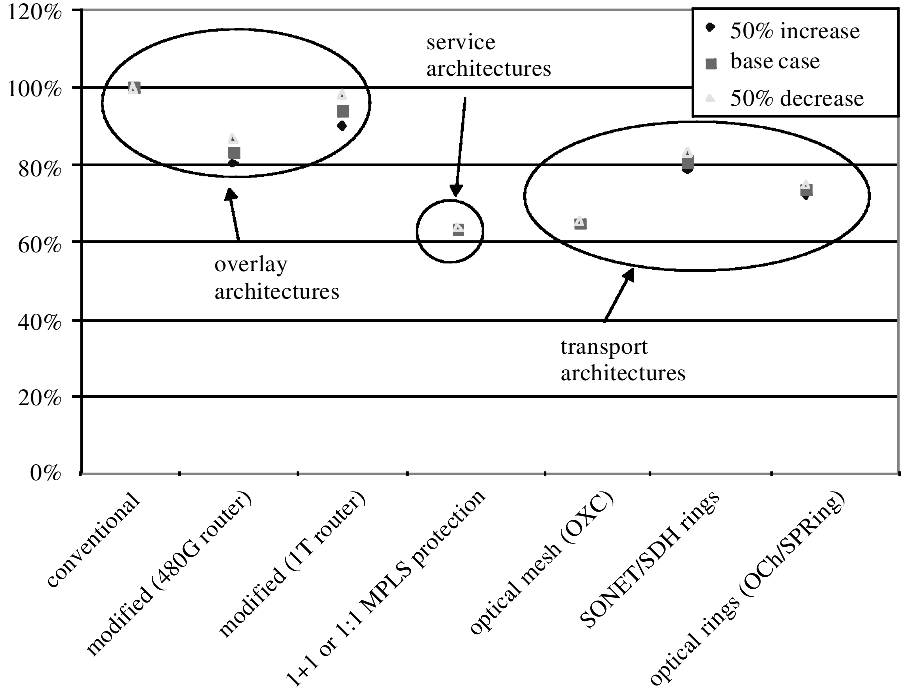

Figure 9. Sensitivity to changes in IP router cost (relative to conventional overlay architecture).

multiple factors, such as network capacity, cost, restoration strategy, reconfigurability and accommodation of pre-emptable traffic.

The network design and economic analysis of a typical nationwide IP backbone showed that OXC-based transport-layer architectures with full wavelength reconfigurability result in the same cost as service-layer network architectures with dedicated MPLS protection. The latter, however, do not provide any wavelength reconfigurability. In both cases, the cost reduction with respect to a traditional overlay design was *ca.* 35%.

It was also shown that architectures with SONET/SDH or optical rings are 12% and 24% more expensive than the optical-mesh architecture. However, the latter provide span switching capability for maintenance purposes and ensure 50 ms protection.

The sensitivity analysis results show that larger cost reduction in 10 Gb s^{-1} optics would make the optical-mesh architecture even more attractive. On the other hand, changes in IP router costs would not result in a significant advantage to any particular architecture.

In conclusion, transport-layer networking architectures are anticipated to be better suited for large IP backbone than service-layer architectures, particularly in regard to restoration performance. This is due to the fact that future backbone networks will carry a very large number of wavelengths per fibre (80, 160 or even more), and, as shown in this analysis, a single link failure due to fibre cut would result in the need for restoring these dozens of optical channels, as opposed to the need to restore up to thousands of LSPs in network architectures based on MPLS restoration.

The author acknowledges colleagues John O. Eaves, Manoj Kumar, M. Akber Qureshi, Antonio Rodriguez-Moral and David Sugerman, who greatly contributed to this work. He also thanks the reviewers for numerous comments that helped improve the paper.

References

Agrawal, N., *et al*. 2000 Demonstration of distributed mesh restoration and auto-provisioning in a WDM network using a large optical cross-connect. *Proc. OFC 2000, Baltimore, MD, USA, March 2000*, paper PD37.

Anderson, J., Manchester, J. S., Rodriguez-Moral, A. & Veeraraghavan, M. 1999 Architectures and protocols for IP optical networking. *Bell Labs Tech. J.* **4**, 105–124.

Baroni, S., Bayvel, P., Gibbens, R. J. & Korotky, S. K. 1999 Analysis and design of resilient multifiber wavelength-routed optical transport networks. *IEEE J. Lightwave Technol.* **17**, 743–758.

Baroni, S., Eaves, J. O., Kumar, M., Qureshi, M. A., Rodriguez-Moral, A. & Sugerman, D. 2000*a* Analysis and design of backbone architecture alternatives for IP optical networking. In *IEEE J. Selected Areas Commun.* Special Issue on 'Protocols and architectures for next generation optical WDM networks'. (In the press.)

Baroni, S., Qureshi, M. A., Rodriguez-Moral, A. & Sugerman, D. 2000*b* Backbone network architectures for IP optical networking. *Proc. OFC 2000, Baltimore, MD, USA, March 2000*, vol. 1, pp. 159–161.

Black, D. P. 1999 *Building switched networks*. Addison-Wesley.

Demester, P. M., *et al*. 1999 Resilience in multilayer networks. *IEEE Commun. Mag.* **37**, 70–77.

Doshi, B. T., Dravida, S. & Harshavardhana, P. 1995 Overview of INDT—a new tool for next-generation network design. *Proc. IEEE GLOBECOM '95*, pp. 1942–1946.

Doshi, B. T., Dravida, S., Harshavardhana, P. & Qureshi, M. A. 1998 A comparison of next-generation IP-centric transport architectures. *Bell Labs Tech. J.* **3**, 63–85.

Doshi, T., Dravida, S., Harshavardhana, P., Hauser, O. & Wang, Y. 1999 Optical network design and restoration. *Bell Labs Tech. J.* **4**, 58–84.

Doverspike, R. D., Philips, S. & Westbrook, J. R. 1999 Future transport network architectures. *IEEE Commun. Mag.* **37**, 96–101.

Gerstel, O. 1998 Opportunity for optical protection and restoration. *Proc. OFC '98*, pp. 269–270.

Johnson, D., Hayman, N. & Veitch, P. 1999 The evolution of a reliable transport network. *IEEE Commun. Mag.* **37**, 52–57.

Kaufman, D., Asif, S., Kasim, A. & Mahoney, M. 1999 The optical Internet—a new network architecture. *Proc. NFOEC '99*, vol. II, pp. 93–104.

Loyd Jones, M., Butler, R. K. & Szeto, W. C. 1999 Sprint long distance network survivability: today and tomorrow. *IEEE Commun. Mag.* **37**, 58–62.

Makam, S., Sharma, V., Owens, K. & Huang, C. 1999 Protection/restoration of MPLS networks. IETF Draft ⟨draft-Makam-mpls-protection-00.txt⟩. See http://www.ietf.org.

Manchester, J., Bonenfant, P. & Newton, C. 1999 The evolution of transport network survivability. *IEEE Commun. Mag.* **37**, 44–51.

OECD 1998 Internet infrastructure indicators. Report no. DSTI/ICCP/TISP(98)7/FINAL (28 October 1998).

Opalka, Z. & Soman, S. 1999 Building the new network infrastructure: will the new super routers have what it takes? Available at http://www.nexabit.com/architecture.pdf.

Rosen, E. C., Viswanathan, A. & Callon, R. 1999 Multiprotocol label switching architecture. IETF draft ⟨draft-ietf-mpls-arch-05.txt⟩. See http://www.ietf.org.

Sato, K.-I. & Okamoto, S. 1999 Photonic transport technologies to create robust backbone networks. *IEEE Commun. Mag.* **37**, 78–87.

Van Caenegem, B., Van Parys, W., De Turck, F. & Demeester, P. M. 1998 Dimensioning of survivable WDM networks. *IEEE J. Selected Areas Commun.* **16**, 1146–1157.

Wauters, N., Ocakoglu, G., Struyve, K. & Falcao Fonseca, P. 1999 Survivability in a new pan-European carriers' carrier network based on WDM and SDH technology. *IEEE Commun. Mag.* **37**, 63–69.

Wright, A. & Yarlagadda, R. 1999 Field trial of OC-48 ATM directly over DWDM. *Proc. NFOEC '99*, vol. I, pp. 133–137.