

 Open access • Journal Article • DOI:10.1109/MCOM.2005.1470813

Benefits of GMPLS for multilayer recovery — [Source link](#)





Bart Puype, J.-P. Vasseur, Adelbert Groebbens, S. De Maesschalck ...+4 more authors

Published on: 01 Jul 2005 - IEEE Communications Magazine (IEEE)

Topics: Optical Transport Network, Multiprotocol Label Switching and Network service

Related papers:

- [Recovery in multilayer optical networks](#)
- [Network Recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS](#)
- [GMPLS Technology and Its Application in WDM Optical Network](#)
- [Research on Application of GMPLS Network in Core Optical Network](#)
- [Study on design of optical communication network with GMPLS technology](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/benefits-of-gmpls-for-multilayer-recovery-2wdo3gjjo4>

Benefits of GMPLS for Multilayer Recovery

Bart Puype, Ghent University

Jean-Philippe Vasseur, Cisco Systems

*Adelbert Groebbens, Sophie De Maesschalck, Didier Colle, Ilse Lievens, Mario Pickavet, and
Piet Demeester, Ghent University*

ABSTRACT

IP-based backbone networks are gradually moving towards a network model consisting of high-speed routers that are flexibly interconnected by lightpaths set up by an optical transport network consisting of WDM links and optical cross-connects. Recovery mechanisms at both network layers will be crucial to reach the high availability requirements of critical services. In such a model, the GMPLS protocol suite can provide a distributed control plane that can be used to deliver rapid and dynamic circuit provisioning of end-to-end optical lightpaths. This article explains that it can be very beneficial to exploit this functionality to enhance the cost effectiveness of multilayer recovery significantly. Several practical case studies illustrate this concept and highlight the opportunities and challenges to be faced.

INTRODUCTION

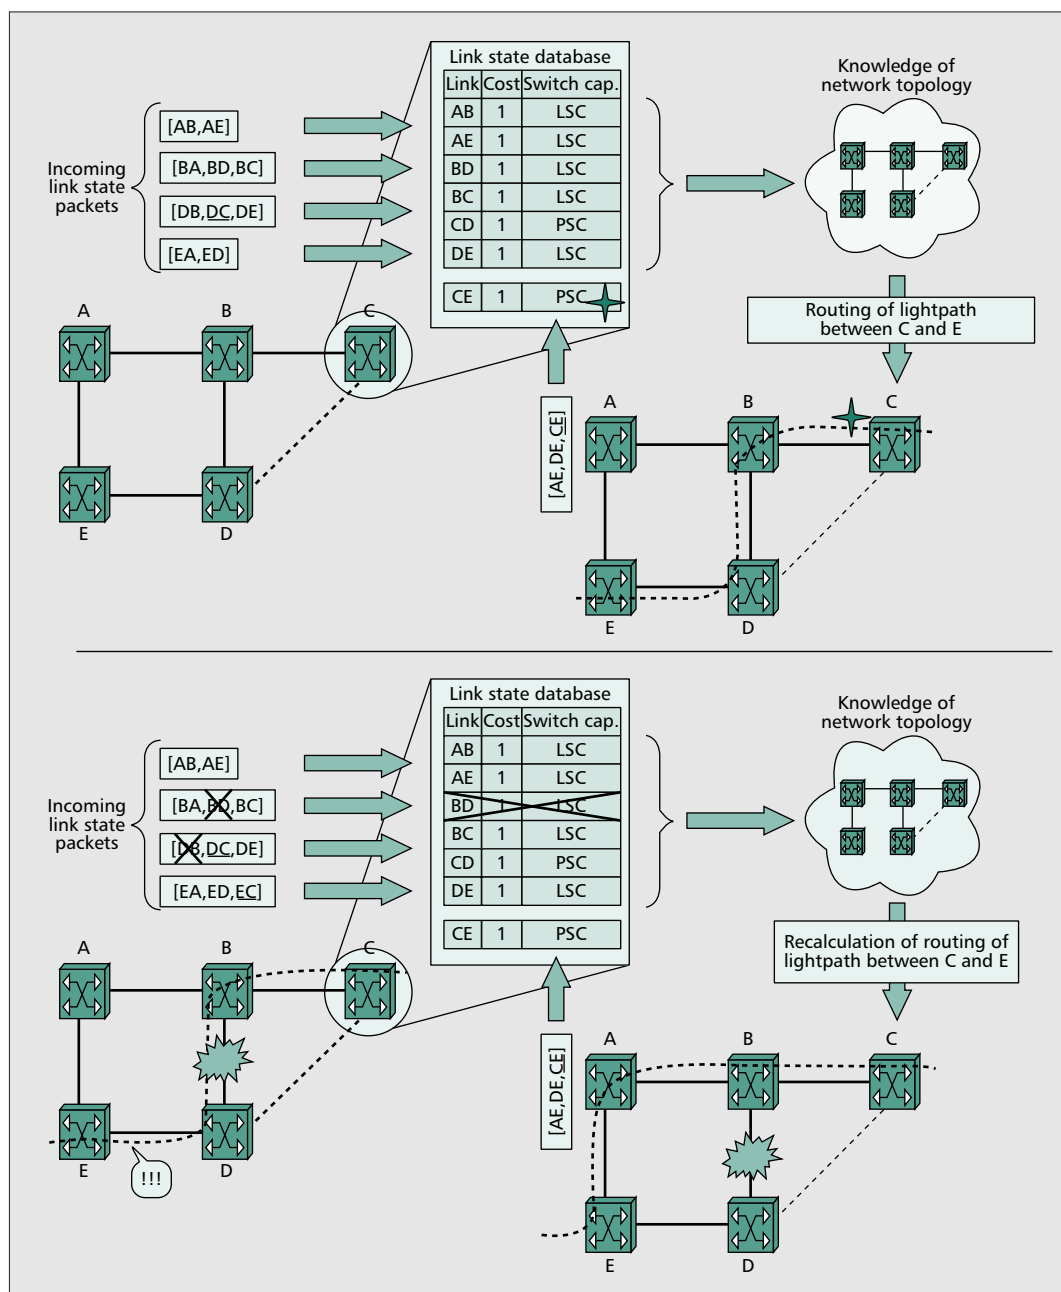
Backbone networks carrying IP traffic have traditionally been supported by optical transport networks (OTNs) that provided fixed transmission links between IP routers. Injected traffic is continuously shifting and changing in volume over time due to diurnal traffic fluctuation and overall traffic growth. As the traffic pattern becomes more and more dynamic, the ability to rearrange transmission capacity allows us to cope better with this trend. This requires that the OTN support setting up and tearing down optical connections (lightpaths) on demand in an automated way. Technological advances such as flexible optical cross-connects (OXCs), switching optical wavelength-division multiplexing (WDM) channels, have made this possible. The IP client network utilizes signaling to request lightpath setup or teardown, relying on the transmission network to perform the action. As such, an increased amount of intelligence is necessary in the OTN, which has led to the concept of intelligent optical networks (IONs).

The addition of a distributed control plane allows lightpath setup and teardown through inter- and intralayer signaling. The generalized multiprotocol label switching (GMPLS) protocol suite specified by the Internet Engineering Task Force (IETF) [1] extends the MPLS label switching concept. Whereas MPLS labels are integers, GMPLS labels can additionally represent time slots in a time-division multiplexing (TDM) frame, wavelength or waveband on a fiber, fiber in a cable, or any further switching granularity.

Link state routing protocols such as Intermediate System to Intermediate System (IS-IS) and Open Shortest Path First (OSPF) have been extended to carry multilayer related information. Figure 1 illustrates the link state message exchange for a GMPLS network. Nodes learn about their neighbors through neighbor discovery as well as various link characteristics, flooding this information throughout the network periodically. The upper part of the figure shows node C receiving link state packets from all other nodes, allowing it to construct its link state database. As links can be of different types, the link state encompasses a field specifying multiplexing or switching capability of the advertised link (third column in the link state database). Optical links are marked lambda switch capable (LSC), while IP-level links are packet switch capable (PSC). Link C-D is such an IP link. The figure furthermore shows how the link state database can be used by a routing protocol to establish an additional lightpath, C-E, over the available LSC optical links (i.e., a connection in the optical layer), and advertise it to the network. It appears as a new IP layer PSC link in link state databases.

In addition to allowing the network to cope with changing traffic demands, the flexibility in setting up lightpaths on demand turns restoration into a viable recovery option. The lower part of the figure shows how one can provide GMPLS-based fault recovery. Over the same topology, a lightpath is now set up between nodes C and E. A failure occurs on

In the absence of a unified GMPLS driven control plane, meaning the network layers are controlled separately, the recovery schemes run independently and unaware of each other. This may lead to a situation where traffic is protected more than once.



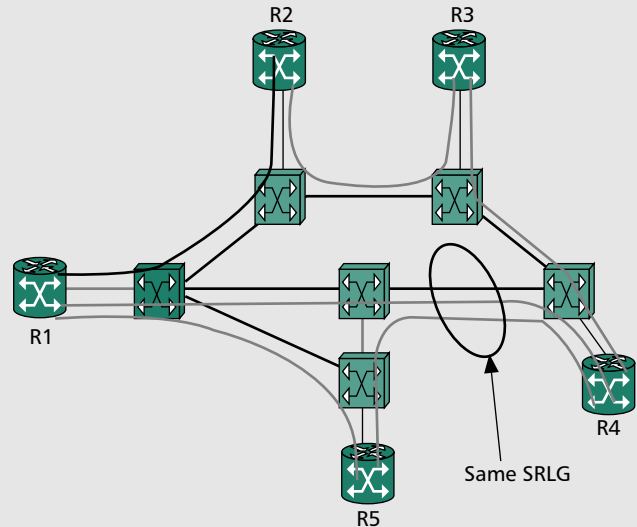
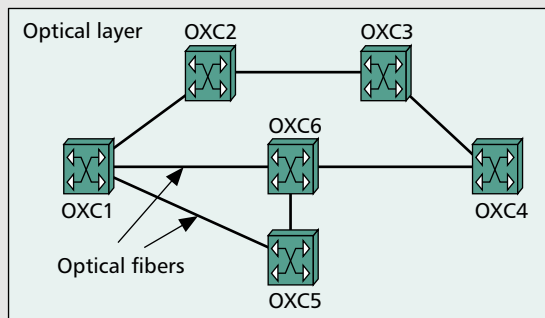
■ Figure 1. Link state routing and restoration in a GMPLS network.

the optical link BD. As lightpath C-E is routed over this failing link, it is also affected. In addition to nodes B and D, the lightpath endpoints (C and E) are notified of the failure using various protocols such as Link Management Protocol (LMP), BFD, or the hello-protocol of the Interior Gateway Protocol (IGP). They trigger the recovery process. The link failure(s) is reflected in the flooded link state packets, and all the nodes' link state databases are updated accordingly: LSC link BD is removed. This is used to recalculate the C-E lightpath routing, effectively restoring the CE PSC link in the IP network. Of course, restoration of an optical lightpath may take some time, which may trigger two routing convergence cycles (depending on optical restoration time and IGP tuning).

MULTILAYER RECOVERY CONCEPTS

In the classical case (IP over a static optical network), each network layer requires its own strategy. This is because although, for example, protection using disjoint backup paths in the optical layer offers large recovery granularity, it cannot protect against all failures. For example, a failing router in the IP layer cannot be recovered through optical layer actions. These types of failures require packet-level rerouting, or if rerouting is considered too time consuming, automatic switchover to backup label switched paths (LSPs) or routing (via tunneling) paths.

Resilience against failures requires a certain amount of backup (for protection) or spare (for restoration) capacity. In the absence of a unified GMPLS driven control plane, meaning the net-



■ **Figure 2.** *The Shared Risk Link Group concept.*

work layers are controlled separately, the recovery schemes run independently and are unaware of each other. This may lead to a situation where traffic is protected more than once: for example, once in the IP layer (redundant LSP or IP packet paths with sufficient capacity) and once optically (backup lightpaths).

Additionally, providing recovery for a number of network layers poses some complex failure issues. Figure 2 illustrates the Shared Risk Link Group (SRLG) concept, and refers to simultaneous secondary failures of multiple network elements caused by a single failure. Consider an optical fiber topology interconnecting six OXCs, used to provide IP link connectivity (lightpaths) to six IP layer routers, R1 through R6, in this multilayer network. IP links R1-R4 (OXC1-OXC6-OXC4) and R5-R4 (OXC5-OXC6-OXC4) share a common resource: fiber OXC6-OXC4. Failure of this fiber affects both these IP links, which thus share an SRLG.

In an IP-over-ION scenario, however, recovery techniques may interact across layers so that double protection can be avoided. IP layer SRLG awareness can be built into this multilayer recovery scheme integration. Moreover, the ability to set up new lightpaths on demand makes the concept of optical restoration a more interesting option. In the next section we discuss some general concepts and methodologies concerning multilayer recovery.

STATIC MULTILAYER RECOVERY

A network failure propagating upwards (from server to client layers) through the network stack and causing secondary effects, can be dealt with in any of those layers. The choice of in which layer to recover the affected traffic depends on the exact circumstances. The rules of coordination to ensure efficient recovery interworking between the distinct layers during different failure scenarios are specified in a so-called escalation strategy.

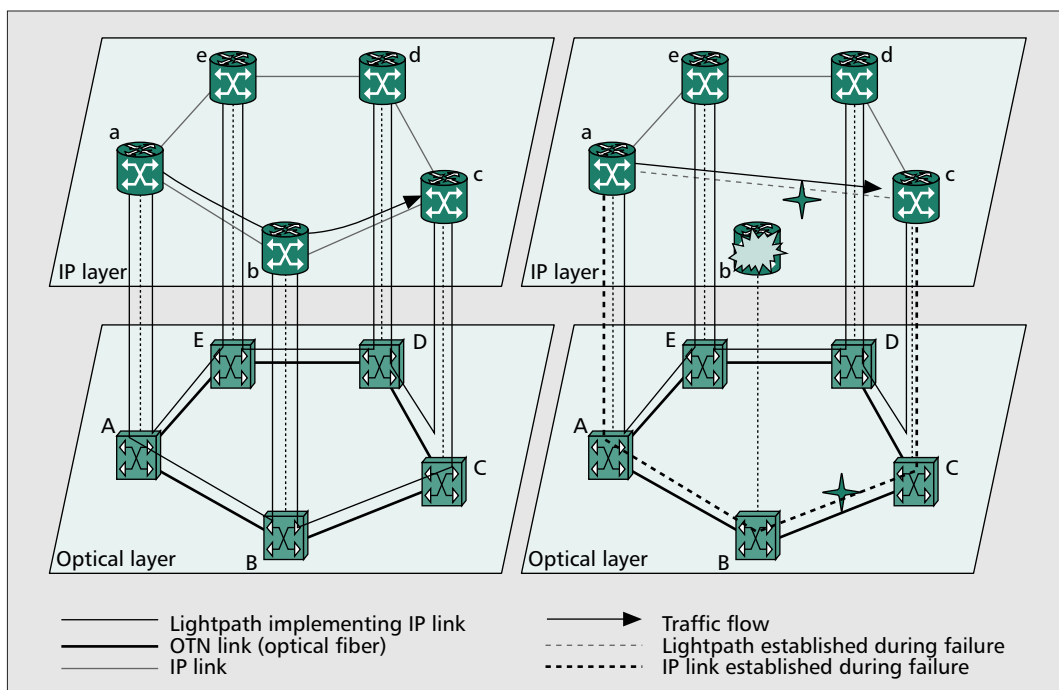
The easiest escalation approach is to simply deploy uncoordinated per-layer recovery

schemes, resulting in parallel recovery actions at distinct layers. For example, a severed fiber carrying lightpaths (IP links) results in both the optical recovery of the fiber (e.g., protection) and the recovery of the IP layer traffic (e.g., IP/MPLS restoration or protection). It delivers a simple and straightforward solution from an implementation and operational point of view, although it can cause spare resource occupation in multiple layers during a failure, lowering overall capacity efficiency. Additionally, one may see increased restoration time and some network oscillations due to race conditions.

More efficient is a timer-based sequential escalation strategy. Here recovery responsibility is escalated to the next layer once it is clear that the layer cannot recover the fault. However, concurrent attempts at recovery in different layers should be avoided by imposing chronological order on recovery mechanisms. There are two alternatives in terms of sequential escalation strategies:

- The so-called bottom-up escalation strategy whereby recovery commences in the bottom or lowest detecting layer, escalating upward for all affected traffic unrestorable by lower layers. The recovery acts on appropriate traffic granularities: coarse granularities (e.g., fibers) are handled first, recovering large volumes of affected traffic quickly, the remaining small part of the traffic at higher layers at finer granularities later on. Complex secondary failures (cf. SRLGs) are handled mostly at lower layers, limiting higher-layer recovery actions. Escalation can be coordinated with a recovery token passed to higher layers, or a holdoff timer set progressively higher as one moves upward in the layer stack [2].
- Conversely, with top-down escalation, recovery actions are instead initiated at the top or highest layer, going downward through the layer stack. Lower-layer recovery is triggered only when higher layers cannot restore all traffic. An advantage is that the

Although most flexible from a recovery point of view, combining several technologies in a single strategy is often practically unrealistic, both in terms of algorithmic complexity and implementation issues.



■ Figure 3. Example of dynamic multilayer recovery.

higher layer has a better overview of different service types, allowing it to better differentiate high-priority traffic and attempt to restore it first. Moreover, some failures may be recovered at a finer granularity, thus requiring a lower total volume of traffic to be rerouted (coarser rerouting may recover some unused capacity). Some efficiency problems may arise from, for example, a higher layer succeeding at recovering half of a lightpath's traffic flows, leaving the other half for the lower-layer recovery mechanisms, meaning that the spare protection capacity in the optical layer is only useful for half of the carried traffic.

A final more radical technique lies in the integrated approach, where coordination is established by combining the above into one integrated multilayer recovery scheme, implying the availability of a full overview of all the network layers. Although most flexible from a recovery point of view, combining several technologies in a single strategy is often practically unrealistic, in terms of both algorithmic complexity and implementation issues.

In [2] a multilayer network case study is presented where fiber and optical equipment failures are restored optically, and MPLS TE fast reroute handles IP/MPLS layer router failures and remaining link failures. Optical restoration uses shared backup paths to optimize required capacity compared to protection mechanisms, at the cost of increased recovery times and less deterministic behavior. Commercial implementations based on GMPLS/ION exist today. Furthermore, MPLS TE fast reroute [3] ensures that all LSPs are protected locally against link and node failures, meaning that for each protected facility a backup tunnel is presignaled. Such backup tunnels can be provisioned to provide bandwidth guarantees during failures and

can share bandwidth should they protect independent resources. The two schemes are coordinated bottom-up, using a holdoff timer.

For a link failure, both layers detect the problem, but the optical layer first initiates optical restoration (assuming sufficient spare capacity is available for a single fiber failure), while the IP/MPLS layer starts a holdoff timer. For successful restoration no recovery is attempted by the IP/MPLS layer, and the optical layer clears the fault before expiration of the holdoff timer.

For an optical node failure, optical restoration succeeds in recovering all affected traffic carried in transit lightpaths (again provided enough spare capacity is available). Routers attached to the node may become isolated, in which case the failure cannot be recovered solely in the optical layer. After the holdoff timer elapses, IP/MPLS fast reroute protection reroutes the set of affected packet LSPs traversing the isolated router (traffic destined to the isolated router is dropped). Recovery time here includes the holdoff time plus any additional time required in rerouting the LSPs.

For double link failures and IP/MPLS link and node failures, the recovery process needs to be performed at the IP/MPLS layer, and recovery time is equal to the previous case.

DYNAMIC MULTILAYER RECOVERY

For static strategies, the IP logical network topology is left unchanged (static) at the time of a failure, and no specific actions are taken to modify it. Dynamic multilayer recovery strategies rely on GMPLS/ION flexibility in (un)provisioning lightpaths to modify the logical topology for recovery purposes, avoiding advance establishment of logical network IP/MPLS spare resources (also solving the problem of whether and how to protect these in the optical layer).

The dynamic reconfiguration example of Fig.

3 shows the IP-over-optical layer stack before (left) and after (right) a failure occurs. Focusing on the IP/MPLS layer, let us assume that flow a-c is forwarded via intermediate router b in the failure-free situation. Before failure, the IP topology contains IP links ab and bc, provided by lightpaths A-B and B-C, respectively. In case of failure of router b, routers a and c detect the failure and regard these two logical links now as useless. They are torn down, freeing up associated optical layer capacity, which can in turn be used to set up a new direct IP link between routers a and c, implemented by a lightpath bypassing the failing router b. Note that the IP traffic from a to c (and potentially traffic from/to other nodes) will be rerouted along the alternative IP path(s) until the new lightpath is provisioned.

A challenge with dynamic multilayer recovery strategies involves the actual logical topologies to be realized during failures: not only has a failure-free logical topology to be dimensioned (as always), but logical topologies are also needed for each of the failure scenarios. These may be designed finding the most optimal topology and IP routing for each of the failures; this method is called global reconfiguration, and the IP topology is recomputed from scratch for each scenario (after removal of the failing network elements). An alternative is local reconfiguration, in which topology design instead begins from the failure-free one, where affected IP links and routers are first removed, and affected traffic is then rerouted over the remaining topology (but possibly adding IP links to do so). Although less optimal, the idea behind local configuration is that it lowers the amount of required reconfigurations and rerouted traffic during failures.

MTE-BASED RECOVERY

Multilayer traffic engineering (MTE) is a type of TE combining existing TE techniques (using TE LSPs, setting IP link weights) with the flexibility of a GMPLS-driven multilayer network. Not only can MTE reroute IP flows when traffic demands vary over time, it also allows logical topology reconfiguration when those variations exceed the range acceptable with simple rerouting. MTE's main objective is then to accommodate traffic demands in the most optimal way (in terms of quality of service [QoS], total throughput, etc.).

It does this by solving discrepancies between the logical topology configuration and the offered traffic pattern, usually caused by traffic fluctuations (leading to congestion). The MTE strategy can naturally also be used for problems that arise from network failures, whose impact on the logical topology will trigger reconfiguration, leading to the setup (and possibly tear-down) of IP links and rerouting of traffic. Such approaches are currently explored at the IETF, for example, by means of path computation element (PCE) techniques whereby path computation is involved across multiple layers to perform optimal rerouting choices [4].

Although not a dedicated recovery strategy, MTE-based recovery is a form of dynamic multilayer recovery, its resilience properties a by-product of the objective to cope with network

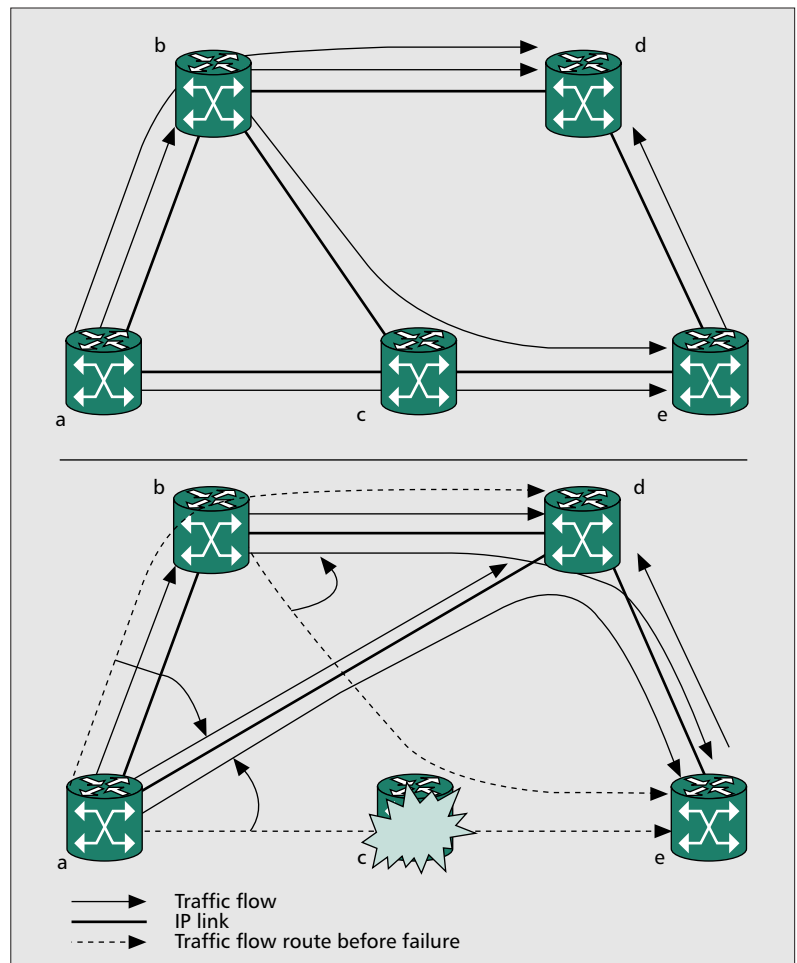
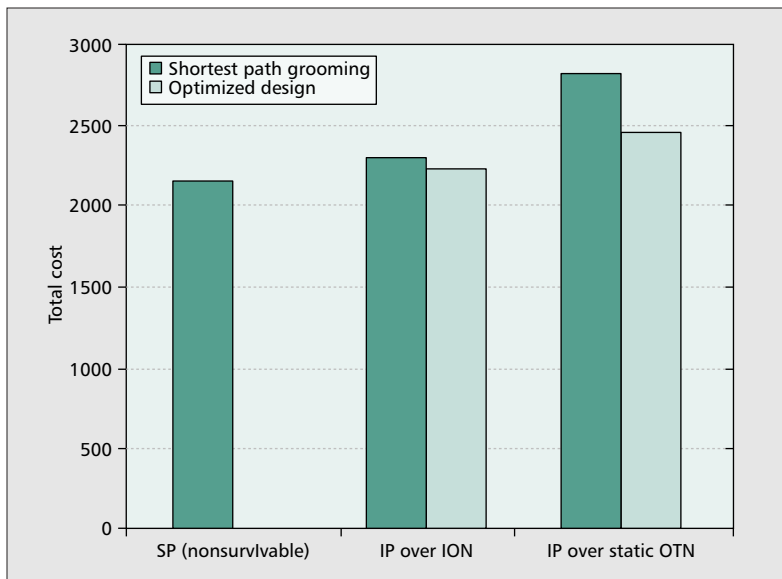


Figure 4. Example of MTE-based recovery.

problems (congestion or others). However, the configuration of the logical topology is not predimensioned but chosen online at the time of failure.

Figure 4 shows a simple example of MTE actions taken upon IP router failure. Only the IP layer is shown (upper part, failure-free; lower part, router c fails): the logical topology and some IP/MPLS traffic flows. The failure of router c affects both flows b-e and a-e (as they were forwarded via c), triggering traffic reroute and lightpath setup. In the case of flow a-e, rerouting happens partly over existing IP links (link d-e), but also over newly setup capacity (link a-d). The routes of flows are of course decided by the MTE strategy, trying to optimize IP layer efficiency (IP link filling) and optical layer efficiency (amount of required lightpaths), and in fact finding a suitable compromise between them. In this case, the affected flow a-e was rerouted over two links, one a replacement IP link. Additionally, the new link a-d attracts flow a-d, since it is a more optimal direct path for that flow (making the assumption of equal link costs in this example). This means we may see some secondary effects beyond the recovery of the affected traffic flows, similar to the effects of global vs. local reconfiguration for dynamic multilayer recovery.

Note that such dynamic reconfiguration may



■ **Figure 5.** Comparison of static OTN and dynamic (ION) multilayer recovery.

require dampening algorithms to avoid triggering network reconfiguration in case of transitional failure. Consequently, a viable interesting option consists of relying on higher-layer recovery with minimal spare resources (which would lead to degradation of some traffic during failure) combined with the MTE strategy to dynamically reconfigure the network should the failure duration exceed some period of time.

SIMULATION RESULTS

In this section we illustrate some of the concepts from the previous section by means of two simulation studies. In the first we compare the performance of static and dynamic multilayer recovery strategies in terms of total network cost. For the second case we take an existing MTE strategy and evaluate how it performs for network failures.

COMPARISON OF STATIC AND DYNAMIC MULTILAYER RECOVERY STRATEGIES

A tool was developed that designs IP over ION logical topologies for certain scenarios: the failure-free case and single IP router failures. This is an optimization problem that incorporates many aspects. Traffic grooming (multiplexing finer-granularity flows into larger-granularity bandwidths) is one of them. Logical topology design is closely related; a mathematical formulation is given in [5]. We consider design with the restriction of shortest-path routing; a seminal paper on such optimization is [6]. Also, spare and working capacity are jointly optimized in the design; see [7] for information on this.

The resulting total multilayer network capacity requirements were then calculated. For example, the required capacity for each optical link is then the maximum link capacity of all failure-free and IP failure cases. So we combine all scenarios into a single worst case scenario to derive a logical topology (and associated resources) for coping with single IP router failures.

For a static multilayer recovery strategy where on-demand setup of lightpaths is unavailable (static OTN), this dimensioning requires setup of all backup capacity (corresponding to the worst case scenario) in advance. For the dynamic case, however, we can rely on GMPLS signaling and reduce the number of lightpaths established at any one time, allowing reuse of IP-OXC interface cards, wavelength, OXC ports, and line systems between the failure scenarios.

This leads to significant cost reduction, quantified by a simulation study [8] based on several scenarios, one concerning an OTN consisting of 12 nodes and 17 bidirectional links. It is a down-sized version of a pan-European reference network topology [9]. The traffic demand applied is modeled by a corresponding symmetric matrix with a volume of 630.6 Gb/s [9], modeling pan-European traffic for the year 2002 based on population numbers, and the work and home activities of users.

As a base scenario, we establish a nonsurvivable groomed logical topology design, using shortest path IP layer routing and dimensioning only for working traffic, allowing us to determine the additional cost associated with providing resilience in the first place. Next to this base case, two optimized designs were calculated based on IP (re)routing for both a static OTN and a GMPLS signaled ION, coping with single IP router failures. IP flow routing is handled by assigning “all one” weights to logical IP links. Rerouting will lead to IP link bandwidth up/downgrade (and thus number of interface cards used and lightpaths set up), and thus logical topology reconfiguration.

Figure 5 gives the results of the performance for these three schemes, showing the cost of the various resources (interfaces, ports, line systems). The algorithm that optimizes the logical topology design starts from a full demand mesh (with direct IP links for all non-zero demands) and deletes all IP links along edges whose removal improves a certain objective value (related to the total cost). For each network scenario, we show the optimized design cost as well as the grooming base case cost calculation, where the starting topology was instead the non-survivable shortest path grooming design with additional capacity to reach a network solution with enough bandwidth for survivability. Comparing this with our optimization algorithm, we observe that a network operator would be charged an additional optimized cost of only 3 percent for IP rerouting over ION, but 15 percent for IP rerouting over static OTN (compared to the nonsurvivable “grooming” base case). Additionally, we see that IP rerouting over static OTN performs much worse if it is adopted on a nonoptimal logical design, but even after multilayer cost optimization it is still at a disadvantage with respect to IP rerouting over ION.

MTE STRATEGY AS A RECOVERY MECHANISM

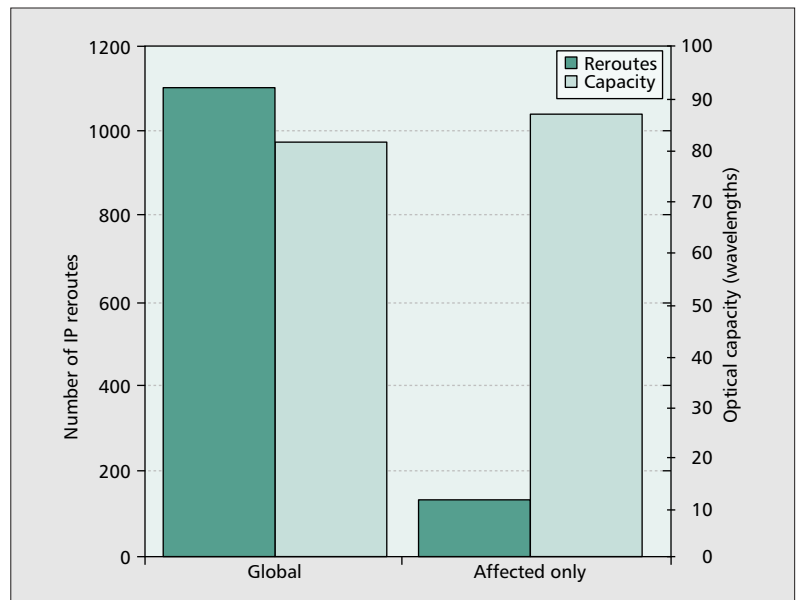
Multilayer TE and routing is achieved by shifting transit traffic from the IP/MPLS to the optical layer; [10] presents a GMPLS-based hardware implementation. Various MTE-type strategies exist. For example, the integrated approach [11] routes incoming traffic requests using a multi-

step algorithm, and identifies some main problems of multilayer routing: wavelength continuity constraints and bandwidth granularity discrepancies. The mechanism in [12] separates the problem into logical topology design (offline), dynamic routing (online), and bandwidth adjustment (traffic-driven) modules. In [13] routing is performed over a single graph that represents the IP/MPLS and optical networks, its main contribution being the cost model to integrate both layers in a single model.

MTE relies on multilayer routing to adapt to changing traffic demands (which are a prerequisite; static demands will be accommodated more optimally by offline provisioning). It requires fast signaling protocols, and one needs to take into account limitations on algorithm complexity, and the possibility of undesired effects such as network instability and QoS degradation during flow reroutes (loss, jitter). The MTE proposed strategy [14] will rely on IP-based costs only, since an overlay model is considered (in contrast to the integrated model in [11, 13]). Optical routing is delegated to an optical RWA algorithm (first-fit in this case). IP/MPLS traffic is routed into a full mesh of LSPs in the logical layer, their path calculated over a virtual full mesh (corresponding with the more sparsely meshed logical topology), which serves to represent the high flexibility in lightpath setup. The cost function will optimize IP layer performance and depends on IP link load. It serves to attract traffic flows to IP links such that they are all moderately loaded, avoiding overloaded links as well as a large number of lightly loaded ones (which reduce bandwidth efficiency). The introduction of higher cost for links with a load below a certain threshold removes traffic flows from those links, thus lowering the total number of links in use. The logical topology configuration is a side-effect of the virtual mesh routing: the remaining virtual mesh IP links that carry traffic are then the only ones to be actually set up, thereby optimizing the logical topology for the traffic pattern at hand. This means all aspects of MTE are handled through a single path calculation — no separate mechanisms as in [12] are necessary.

No backup path calculation is performed in the IP layer (although optical layer recovery will likely be present). This MTE strategy is normally triggered proactively [15] on network performance degradation (to this end it monitors link loads), but some small alterations allow it to be triggered on network faults as well. The integrated view in [13] allows one to calculate backup paths over a multilayer network, but in this case our overlay view separates IP/MPLS (i.e., MTE) and optical recovery. We have simulated the MTE strategy actions for single router failures on a 28-node 41-fiber meshed pan-European reference topology [9].

As mentioned previously, the online optimizing character of the MTE strategy may lead to some secondary effects such as rerouting of unaffected traffic flows and even setup/teardown of lightpaths unrelated to a certain router failure. Live rerouting of a traffic flow requires signaling in order to update the GMPLS forwarding tables for its LSP, possibly causing some QoS



■ Figure 6. Effect of MTE optimization scope on IP and optical performance.

degradation. While acceptable for globally reoptimizing dynamic multilayer recovery schemes, it has real impact on network performance in the case of MTE-based recovery schemes. This is because here the reconfiguration is performed online, and the extent of the rerouting and lightpath setup/teardown will affect total convergence time (time between failure occurrence and stabilization of the traffic flow routes).

To illustrate this, we performed a simulation on both the MTE strategy discussed above and a slightly modified version of it where, on detection of a failure, only the affected traffic flows are rerouted and allowed to receive replacement lightpaths. The results are shown in Fig. 6, averaged for all single-router failures, where “affected only” stands for the modified version of the algorithm, similar to a local dynamic recovery scheme. We show both IP layer performance in terms of number of IP reroutes during recovery, indicating convergence speed, and optical layer performance through total optical resource usage (amount of required wavelengths after convergence). We can see a very large improvement in the convergence characteristics of the MTE strategy by using appropriate signaling to identify failure-triggered congestion as such, in exchange for a slight increase of total optical resource usage. This is because the affected only scheme will fixate the route of unaffected flows, so some lightpaths may end up being used less efficiently.

CONCLUSION

The introduction of GMPLS in IP over optical networks offers new possibilities in providing resilience against network failures. We have discussed various issues concerning multilayer survivability, and shown the differences between static (with advance provisioning) and dynamic recovery where the logical topology is reconfigured at the time of failure. A special case of dynamic recovery is multilayer traffic-engineer-

We can see a very large improvement in the convergence characteristics of the MTE strategy by using appropriate signaling to identify failure-triggered congestion as such, in exchange for a slight increase of total optical resource usage.

ing-based recovery, where the recovery features arise naturally from the MTE's goal to keep the network running correctly. A first case study illustrated the advantages a dynamic strategy has over a static scheme for IP layer restoration against router failures. MTE-based recovery was explored in a second case study, where we improved an MTE strategy to reduce the performance impact due to traffic rerouting during failures.

ACKNOWLEDGMENTS

Bart Puype and Didier Colle thank the IWT for its financial support for their Ph.D. and postdoctoral grants, respectively. The work was partly funded by the European Commission through projects IST-NOBEL and IST-ePhoton/ONE; and by the Flemish Government through projects IWT-GBOU ONNA, IWT-ITEA TBONES, and FWO G.0315.04.

REFERENCES

- [1] E. Mannie et al., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," draft-ietf-ccamp-gmpls-architecture (2002).
- [2] J.-P. Vasseur et al., *Network Recovery, Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*, Morgan Kaufmann Series in Networking, Elsevier, 2004.
- [3] P. Pan et al., "Fast Reroute Techniques in RSVP-TE," draft-ietf-mppls-rsvp-lsp-fasteroute, work in progress.
- [4] Path Computation Element Workgroup, <http://www.ietf.org/html.charters/pce-charter.html>
- [5] B. Ramamurthy et al., "Design of Virtual Private Networks (VPNs) over Optical Wavelength Division Multiplexed (WDM) Networks," *Opt. Networks*, vol. 3, no. 1, Jan/Feb 2002.
- [6] B. Fortz et al., "Optimizing OSPF/IS — IS Weights in a Changing World," *IEEE JSAC*, vol. 20, no. 4, May 2002 pp. 756–67.
- [7] R. R. Iraschko et al., "Optimal Capacity Placement for Path Restoration in STM or ATM Mesh Survivable Networks," *IEEE/ACM Trans. Networking*, vol. 6, no. 3, June 1998, pp. 325–36.
- [8] A. Groebbens et al., "Logical Topology Design For IP Rerouting: ASONs versus Static OTNs," submitted, *Photonic Network Commun.*, 2005.
- [9] IBCN, reference topologies for the IST LION project and COST266 action: http://www.ibcn.intec.ugent.be/css_design/research/projects/IST_FP5/NRS/index.html
- [10] K.-I. Sato et al., "GMPLS-Based Photonic Multilayer Router (Hikari Router) Architecture: An Overview of Traffic Engineering and Signaling Technology," *IEEE Commun. Mag.*, vol. 40, no. 3, March 2002, pp. 96–101.
- [11] M. Kodialam et al., "Integrated Dynamic IP and Wavelength Routing in IP over WDM Networks," *Proc. IEEE INFOCOM 2001*, pp. 358–66.
- [12] P. Iovanna et al., "A Traffic Engineering System for Multi-layer Networks Based on the GMPLS Paradigm," *IEEE Network*, vol. 17, no. 2, Mar./Apr. 2003, pp. 28–37.
- [13] T. Cinkler et al., "Fairness Issues of Routing with Grooming and Shared Protection," *Proc. ONDM 2004*, pp. 665–84.
- [14] B. Puype et al., "Optical Cost Metrics in Multi-Layer Traffic Engineering for IP-over-Optical Networks," *Proc. ICTON 2004*, vol. 1, pp. 75–80.
- [15] B. Puype et al., "Multi-layer Traffic Engineering in Data-centric Optical Networks, Illustration of Concepts and Benefits," *Proc. ONDM 2003*, pp. 221–26.

BIOGRAPHIES

DIDIER COLLE received a M. Sc. degree in electrotechnical engineering (option, communications) from Ghent University in 1997. Since then, he has been working at the same university as a researcher in the Department of Information Technology (INTEC). He is part of the research group INTEC Broadband Communication Networks (IBCN) headed by Prof. Piet Demeester. His research led to a Ph.D. degree in February 2002. He was granted a postdoctoral scholarship from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen) for

the period 2003–2004. His research deals with design and planning of communication networks. This work focuses on optical transport networks, to support the next-generation Internet. Until now, he has actively been involved in several IST projects (LION, OPTIMIST, DAVID, STOLAS, NOBEL and LASAGNE), COST actions 266 and 291, and the ITEA/IWT TBONES project. His work has been published in more than 100 scientific publications in international conferences and journals.

SOPHIE DE MAESSCHALCK received an M.Sc. degree in electro-technical engineering (option communication technology) in 1998 from Ghent University, Belgium. In 1998 she started working as a research assistant at INTEC, Ghent University. In 2003 she obtained her Ph.D. Her research interests included the design, planning, and evaluation of multilayer IP-over-optical transport networks and broadband access networks. She was involved in the ACTS project SONATA, in the IST projects LION, BREAD, and MUSE, and in COST action 266. Currently she is working at Proximus, one of the Belgian mobile operators.

PIET DEMEESTER [SM] finished his Ph.D. thesis at INTEC, Ghent University, in 1988. At the same department he became group leader of the activities on metal organic vapor phase epitaxial growth for optoelectronic components. In 1992 he started a new research group on broadband communication networks. The research in this field has already resulted in more than 300 publications. In this research domain he was and is a member of several program committees for international conferences, such as ICCCN, the International Conference on Telecommunication Systems, OFC, ICC, and ECOC. He was Chairman of DRCN '98. He was chairman of the Technical Programme Committee ECOC '01. He has been Guest Editor of three special issues of *IEEE Communications Magazine*. He is also a member of the Editorial Board of *Optical Networks Magazine* and *Photonic Network Communications*. He was a member of several national and international Ph.D. thesis commissions. He is a member of ACM and KVIV. His current research interests include multi-layer networks, QoS in IP networks, mobile networks, access networks, grid computing, distributed software, network and service management, and applications (supported by FWO-Vlaanderen, the BOF of Ghent University, the IWT, and the European Commission). He is currently a full-time professor at Ghent University, where he is teaching courses in communication networks. He has also been teaching different international courses.

ADELBERT GROEBBENS received in 2000 a Master's degree in computer science (option communication technology) at Ghent University, Belgium. In October 2000 he was granted an FWO aspirant scholarship from the Flemish government. Since then, he has worked as a researcher in the Broadband Communications Networks Group (IBCN) at INTEC, Ghent University. In March 2005, he received a Ph.D. degree in computer science from Ghent University. His interests include operations research and software engineering techniques that can be applied for modeling and solving problems encountered in the design and planning of communication networks. His Ph.D. research was focused on IP-over-optical transport networks to support the next-generation Internet. He has published over 20 journal or conference papers in this field. He participated in the European IST-project Layers Interworking in Optical Networks (LION), the European ITEA project Transparent Backbone Optical Network Simulator (TBONES), and some Flemish interuniversity projects.

ILSE LIEVENS received an M.Sc. degree in electrical engineering, focusing on telecommunications, in 1994 from Ghent University. She then joined the INTEC, Ghent University, where she obtained a Ph.D. degree in 2000 in the IBCN research group. Her Ph.D. thesis, *Use of Distributed Rerouting in Meshed ATM Networks*, looked at the design of rerouting algorithms for survivability in meshed ATM networks, focusing on distributed, autonomously working techniques. She is currently working as a post-doctoral assistant in the IBCN research group. Her research interests involve broadband communication networks, focusing on the design, reliability, and survivability of IP and optical networks. She is and has been involved in several European projects (e.g., IST LION, IST NOBEL, ITEA TBONES), and national interuniversity projects (IWT GBOU ONNA). She is author and co-author of several publications in conference proceedings and journals (ECOC, *IEEE JSAC*, *IEEE Communications Magazine*, PNC, etc.).

MARIO PICKAVET is a full-time professor at Ghent University, where he teaches telecommunication network courses. His current research interests are related to broadband communication networks (WDM, IP, (G-)MPLS, OPS, OBS). In this context he is currently involved in the European IST projects STOLAS, LASAGNE, and ePHOTON/ONe. He has published in several international publications, both journals and conference proceedings.

BART PUYPE (bart.puype@intec.Ugent.be) received his M.Sc. degree in electro-technical engineering in 2002 from Ghent University. Since then he has been working as a research assistant, preparing his Ph.D., at INTEC, Ghent University. Starting in 2004, he received an IWT Ph.D. scholarship. His main interests are in the field of communication networks, focusing specifically on multilayer traffic engineering and multilayer resilience in IP-over-optical networks. He is currently involved in the IWT/ITEA TBONES and European 6th framework program IST NOBEL projects.

JEAN-PHILIPPE VASSEUR is a system architect at Cisco Systems where he works on IP/MPLS architecture specifications, focusing on IP, TE, and network recovery. He holds an engineering degree from France and an M.S. from Stephens Institute of Technology, New Jersey. Before joining Cisco, he worked for several service providers in large multiprotocol environments. After two years within the EMEA technical consulting group focusing on IP/MPLS routing, VPN, and traffic engineering designs for service providers, he joined the CISCO Engineering team as a technical leader with a particular focus on IP, MPLS traffic engineering, and recovery mechanisms. He is an active member of the IETF, co-chair of the IETF PCE Working Group, and coauthor of several IETF specifications. He is a regular speaker at various international conferences, and is involved in various projects in the area of IP and MPLS. He has also filed several patents in the area of IP and MPLS and is coauthor of *Network Recovery* (Morgan Kaufmann, July 2004) and *Definitive MPLS Network Designs* (Cisco Press, March 2005).

IEEE COMMUNICATIONS MAGAZINE

CALL FOR PAPERS

ETHERNET TRANSPORT OVER PUBLIC WIDE AREA NETWORKS

BACKGROUND

Ethernet's dominance as a LAN technology creates the natural desire to interconnect Ethernet LANs through the public telecommunications network. Frame Relay and ATM have been the most common current technologies for WAN connections. Frame Relay, however, has scalability issues in terms of higher data rates and virtual LAN connectivity, and ATM has certain disadvantages in terms of complexity. Some high-speed WAN connections today use a PPP mapping for Layer 2, which requires termination of the customer Ethernet frames. For simplicity of processing and maximum transparency, however, the WAN interconnection should ideally preserve the customers' Ethernet frames through the WAN. Evidence of the interest in this area is seen in number of standards bodies and forums that have been working on it. Specifically, the ITU-T is working on a family of standards for Ethernet transport (Study Group 15) and the operations and performance of these networks (Study Group 13). The Study Group 15 work focuses on the required Layer networks, protocols, and carrier network equipment for Ethernet transport in addition to general recommendations to describe Ethernet networks and services into terms that can be used by public carriers. ATIS Committee OPTXS (formerly T1X1) is actively providing the North American input to the Study Group 15 work. IEEE is considering some of the access and metropolitan network aspects in the 802.17 Resilient Packet Ring (RPR), 802.3 Ethernet, and 802.1 bridging working groups. The Metro Ethernet Forum (MEF) is working on the definition of customer Ethernet services and the IETF is looking at Ethernet transport from the perspective of Layer 2 Virtual Private Networks. While work is continuing in each of these bodies, many of the foundational standards have been approved.

This feature topic is a follow-up to the one that appeared in the March 2004 issue. Since that time, great strides have been taken in terms of standards approval and carrier field experience with Ethernet WAN services. It is the goal of this feature topic to capture that progress, giving the Communications Magazine's readers an updated view of Ethernet WAN transport.

Scope of Contributions

This feature topic is intended to provide tutorial information to the Communications Magazine readers on Ethernet Transport and related issues associated with it. Papers are solicited in, although not limited to, the following areas:

- Experiences with Ethernet transport services and trials
- Network architecture alternatives and analysis for Ethernet transport networks
- Operations and performance issues associated with Ethernet transport
- Enabling technology for Ethernet transport
- Ethernet service types required by residential, small/medium enterprise, and large enterprise customers
- Ethernet services complementing or competing with existing leased line services
- Impact of Ethernet transport services on existing corporate networks
- Layer 1, Layer 2, and Layer 3 switching alternatives for virtual private line and LAN services

SCHEDULE FOR SUBMISSIONS

Submission Deadline: July 5
 Notification of Acceptance: September 12
 Final Manuscript Due: September 31
 Publication Date: November issue

Manuscripts must be submitted through the magazine's submissions Web site at:

<http://commag-ieee.manuscriptcentral.com/>

On the Manuscript Details page (under Submit First Draft of a New Manuscript) please click on the drop-down menu to select **November/Ethernet Transport over WANs**

FEATURE TOPIC EDITORS

Steve Gorshe
steve_gorshe@pmc-sierra.com

Glenn Parsons
gparsons@nortel.com

Mike Truskowski
truskows@cisco.com

Li Mo
mo.li@zteoptical.com

Italo Busi
Italo.Busi@alcatel.it