NETW 69
3398

Enabling High Availability over Multiple Optical Networks

Dimitri Staessens, Didier Colle, Ilse Lievens, Mario Pickavet, and Piet Demeester,
Ghent University — IBBT

Walter Colitti, Ann Nowé, and Kris Steenhaut, Vrije Universiteit Brussels

Ricardo Romeral, University Carlos III Madrid

ABSTRACT

Carriers are gradually adopting a network model that consists of MPLS-capable routers and OXCs interconnected by high bandwidth WDM links for transporting IP and Ethernet traffic. A control plane can be used to deliver dynamic circuit provisioning in the transport layer, as well as bandwidth provisioning and traffic engineering in higher layers. Globalization drives the quest for end-to-end QoS guarantees over different carrier networks. Recovery mechanisms are crucial to reach the high availability requirements of critical services. In this article, we share our vision on how to enable high availability services, spanning multiple networks using failure protection techniques while complying with administrative constraints.

INTRODUCTION

Communications services play a vital role in modern private, corporate, and institutional life. The transport networks on which companies depend for their critical business services are based on technologies such as wavelength division multiplexing (WDM), which drastically increased the bandwidth capacity at a low cost. This cost-effectiveness has driven the competition between operators to the point where there is little revenue to be made from classical phone/fax/data services. Looking to increase revenue, these operators are exploiting the capacity to introduce high bandwidth services such as broadband Internet access and digital television. This trend fuels the quest for converged network architectures; able to run all voice, Internet, and video services, commonly called triple-play. The scalability and robustness of the Internet protocol (IP) are the main reasons for its success; therefore, IP is the network protocol of choice for future networks, incorporating multi-protocol label switching (MPLS), which introduces powerful traffic engineering tools. The addition of a distributed control plane enables lightpath set up and tear down through inter- and

intra-layer signaling [1]. The generalized multi-protocol label switching (GMPLS) protocol suite [2] extends the MPLS label switching concept. Whereas MPLS labels are integers, GMPLS labels additionally can 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.

An important quality-of-service (QoS) parameter for a connection is its availability, which is the percentage of time the connection is effectively available with regard to the request. A common availability goal specified in a service level agreement (SLA) is *five nines*, referring to 99,999 percent availability, corresponding to a total annual downtime limit of a little over five minutes. One way to increase the availability is to introduce protection or restoration. Protection implies that a connection is protected against failures by setting up multiple connections that are reserved in advance. Restoration means that the network tries to revive the connection using available resources after a failure has occurred. Protection typically has higher recovery speed but lower availability than restoration and requires more resources [3]. In current optical networks the 1+1 protection scheme is the one most commonly adopted; reserving and enabling a back-up path for every lightpath in the network — often over a duplicate network infrastructure — one network called the *even* network, the other the *odd* network.

The odd and even networks often are built with equipment from different vendors. The operator may view these networks as being different *domains*. In this work, we consider the two networks to be one domain because there are no real restrictions on sharing information between them. However, for the operator, there is an issue of how to make these networks communicate with each other (e.g., E-NNI implementations).

Inter-domain routing and traffic engineering are gaining the attention of researchers. All standardization bodies have presented requirements for inter-domain routing. Important requirements are the administrative restriction

that no topology information may cross domain boundaries and that the inter-domain routing protocol must be independent of the intra-domain routing protocols. Different solutions are presented, based on GMPLS and the path computation element (PCE) architecture [4] and the Optical Border Gateway Protocol (OBGP). The PCE is an entity (node or process) that computes intra- and inter-domain paths on request. There can be multiple PCEs per domain, and path computation can run distributed over them. Most multi-domain network research was performed on the computation of single traffic engineered paths, for example, [5]. Protection is a recent topic, such as optimization of shared protection in multi-domain scenarios [6]. We based our approach on two observations. First, the optimal path cannot be found without sharing all the information, which violates the administrative restrictions, so every solution found will be sub-optimal. Second, every domain is operated by another company, trying to make a profit. It is opportune to optimize every domain, so that every carrier has its maximum profit from the current solution. The global solution, however, may be suboptimal. In a business sense, a carrier has no incentive to give up its profit so that its competitor can benefit. For this domain, an optimal solution can be computed.

We detail how single link and single node failures can be resolved in a multi-domain scenario in an efficient way, using a per-domain path computation method. Recently, the Common Control and Measurement Plane Internet Engineering Task Force (CCAMP IETF) group released RFC5152 [9], detailing per-domain path computation methods and signaling of both protected and unprotected inter-domain MPLS label-switched paths (LSPs) in a single-layer environment. Our approach contributes to this document by proposing a multilayer integrated approach for protected LSPs in GMPLS networks and computing back-up LSPs and working LSPs simultaneously in each domain to avoid crankback procedures in trap topologies.

MULTI-DOMAIN RESILIENCE CONCEPTS

The considered scenario is a multi-domain network composed from two layer domains with an IP/MPLS layer and an optical cross connect (OXC)-switched WDM layer. Each domain is running GMPLS and uses at least one PCE for path computations. The PCE continuously maintains a database of all topology and routing information in the domain. If there are multiple PCEs in a domain, the information may be distributed among them. All PCEs combined must be able to construct the total traffic engineering (TE) database for the domain.

Traffic is exchanged between domains using gateway nodes. These nodes also consist of an IP/MPLS router (label switched router [LSR]), for example, running BGP-4 and an OXC. The traffic exchanged between these domains is assumed to be exchanged between these routers. Although in theory it is possible to exchange traffic at the pure optical level, this is not considered in this article. In practice, it would present a great

challenge to directly interconnect OXCs from different vendors due to implementation differences in the equipment. Also, most operators would be reluctant to directly allow analog input into their network, because they would have difficulty monitoring and rerouting the signal.

For an end-to-end connection to be resilient over multiple domains, we must take into account different failure scenarios. For our solution, we set the following requirements regarding node and link failures. The border gateway nodes for the connection are handled separately and are called gateway failures.

First, we want to be resilient against single link or single node failures inside a domain. A node failure can be the OXC or the LSR (or both) at a specific site inside the domain. In single-domain studies, usually only single link/node failures are considered; however, when multiple networks are traversed, we think the connection should be resilient toward multiple failures, as long as there is only one node or link failure inside any domain. The reason for this requirement is that link failures are fairly common. These internal link or node failures should be handled without any interference of the connection head-end. This means that each working path segment in every domain should have its own protection.

Another requirement is to be able to recover from single border gateway node failures end-to-end. Border gateway failures are far less common than link failures, but they are extremely disruptive when they occur. This last requirement immediately imposes, among others requirements, that there must be (at least) two different ingress points into every domain.

Although it is possible to set up the back-up paths through different domains than the primary paths, there are some drawbacks. First, it is impossible to be absolutely sure that no physical resources are shared between the two networks. A typical example is when multiple network operators use the same bridge over a river for their optical fiber cables. A flood that removes the bridge from service disrupts all of the networks. Second, there is an administrative issue in having to discuss contracts with multiple partners. Therefore, we will route both working and back-up paths through the same network domains.

In summary, the connection should be protected against one failure per domain or a single gateway failure end-to-end. A concurrent failure of a border node and a link inside a domain is not guaranteed to be survivable in this approach.

In Fig. 1 we illustrate the general requirements stated previously using a simple example. We exclude node architecture and multilayer details at this point. Suppose we want to set up a reliable broadband connection (e.g., a virtual private network [VPN]) between a local area network (LAN) at Ghent University (UGent) and the LAN at the University of Carlos III in Madrid (UC3M). This involves a connection through the Belgian research network BELNET, the pan-European research network GÉANT2, and the Spanish research network RedIRIS. In every domain, we will set up an optical working path (primary, P_x). Each of these must be protected with a disjoint back-up path (primary back up, PB_x) to meet the requirement that the head-end

In practice, it would present a great challenge to directly interconnect OXCs from different vendors due to implementation differences in the equipment. Also, most operators would be reluctant to directly allow analog input into their network, because they would have difficulty monitoring and rerouting the signal.

Two-node-connected means that it is possible to find two node-disjoint paths between any two nodes in a graph. The two-node-connectedness of a domain is a very basic requirement, because any kind of node failure survivability within this network would require it to be two-node-connected.

■ Figure 1. Implementation of or requirements for a survivable connection.

(at UGent) will not be required to worry about failures inside, for example, the GÉANT2 network. If a failure along P_2 occurs, the path PB_2 is used. However, if a primary gateway between BELNET and GÉANT2 should fail, the GÉANT2 network cannot recover this by itself; therefore, we require another back-up path over another gateway. These back-up paths (back up, B_x) are used only in case of a gateway failure.

THE COMMON POOL MULTI-DOMAIN STRUCTURE

We now take a closer look at the proposed path set up. The back-up path B_x is used only when a gateway fails and therefore, can share resources with both P_x and PB_x . This yields better capacity consumption of the proposed scheme. The only restriction on B_x is that it cannot use the gateways that are connected by P_x and PB_x , because these gateways would present a single point of failure.

We consider only two LSPs, each having a dedicated ingress and egress point into every domain. The primary ingress and egress gateways support the working (primary) LSP; the secondary ingress and egress gateways support the back-up (secondary) LSP. If there are multiple back-up LSPs, we use a priority naming convention. We would use "tertiary," "quaternary," and so on, gateways, which are used respectively for the tertiary, quaternary, and so on, LSPs. In this article gateway priority will simply mean whether the gateway is a working or back-up gateway.

If the domain is two-node-connected, the proposed path structure always can be computed if the four gateways are given, but the priority of the egress gateways is not fixed. We have detailed mathematical proof of this claim; however this claim is beyond the scope of this article. Two-node-connected means that it is possible to find two node-disjoint paths between any two nodes in a graph. The two-node-connectedness of a domain is a very basic requirement, because any kind of node failure survivability within this network would require it to be two-node-connected.

Any algorithm used to compute the structure must have the network graph, the four gateways, and the priorities of the ingress gateways. The algorithm then must determine the egress gateway priorities and compute the paths. The fact that the structure can be computed in any two-node-connected network also means that we will never run into problems (e.g., trap topologies)

setting up this structure over multiple domains in a sequential way.

A simple method for computing the paths is to choose all priorities fixed and use the Suurballe algorithm [7] for the disjoint paths P_x and PB_x , then remove the primary gateway nodes from the graph, and compute the shortest path between the back-up gateway nodes for B_x . If this fails (this implies that the removal of the gateways disconnects the network graph), switch the priorities of the egress gateways, and a repetition of the algorithm must give a valid solution. This method has no sharing optimization, so we also designed a capacity usage minimization solution using integer linear programming (ILP) techniques.

Since the number of gateways in a domain is generally limited, all required path computations can be performed offline and can be given, for example, as input to a PCE in the network. Usually, such connections that require high availability are long-living connections with a fairly stable traffic load, subject only to some diurnal changes. In a dynamic operational environment, heuristics can be used, taking into account the available resources in the network at the time of the connection request if shorter-lived connections are required.

In Fig. 2, we show the capacity-sharing possibility using a simple example topology. When we look at the top solution, which is computed using the simple approach, it uses three hops for the working path P , four hops for the back-up PB , and three hops for the gateway back-up B . No sharing is done here, so the total solution uses ten hops. Now the bottom solution is optimized and uses three hops for the working path P , six hops for the back-up PB , and three hops for the gateway back-up B . Due to the sharing of three links, the total use of this solutions is nine hops. This relatively small gain is due to the restricted example topology; using a large topology can yield more significant improvements due to the increase in possible paths between the gateways.

The back-up path B_x is not protected. The reason for this is that effective protection of B after a failure of a primary gateway cannot be guaranteed in any two-node-connected network. This can be easily verified in a ring network, but denser topologies also can present problems. To guarantee effective protection of B after a gateway failure, a three-node-connected network will be required.

CONNECTION TIMELINE

We now detail how the survivable connection is set up regarding the full multilayer view of the

network, how different failure scenarios are explicitly handled, and how they affect the overall network state. Note that in the upper (network) layer, the LSP segments depicted in each domain are single-hop connections between gateways. In the lower (optical) layer, paths P_x , PB_x , and B_x are actual multihop paths implementing single-hop LSPs. Two failure scenarios are considered, being OXC failures inside a domain and OXC or LSR failure of a gateway. We now discuss our PCE/Resource Reservation Protocol with Traffic Engineering Extensions (RSVP-TE)-based approach here [8]. Note that this will require a new RSVP-TE object to specify gateways and their priorities. Therefore, we propose a gateway specification routing object (GSRO), and its most important contents are a domain identifier and two lists of gateways (an ingress list and an egress list) in priority order. Full protocol details and processing of this object are beyond the scope of this article.

CONNECTION SET UP

The first step is that the head-end sends an RSVP-TE PATH message to the tail-end, saying it wants to set up a survivable connection through the intermediate domains (Domain 1 and Domain 2). The scenario is shown in Fig. 3. To do this, the head-end, for example, performs a lookup in its IP table or open shortest path first (OSPF)-TE database to determine its preferred gateway toward the destination. The source sends the PATH message (with as destination the IP address of node D) toward its preferred gateway e_{1s} . This gateway determines the next hop (for instance, using its BGP table or using the explicit route object [ERO] in the PATH message), being the primary ingress gateway for Domain 1, i_{11} , and forwards the message. The primary gateway i_{11} then consults the PCE for Domain 1. This PCE determines the secondary ingress gateway, i_{21} , and computes the optical paths P_1 , PB_1 , and B_1 . This also determines the primary egress gateway e_{11} and back-up egress e_{21} . The path information for P_1 , PB_1 , and B_1 is stored in a temporary database, and the routing information is returned to i_{11} . The information returned by the PCE to i_{11} contains the four gateway nodes (i_{11} , e_{11} , i_{21} , e_{21}) and the path information for P_1 and PB_1 . A GSRO with the gateway information and an ERO with a path key for P_1 are attached to the PATH message. Path keys are an alias for an explicit path and provide a means for keeping paths confidential. The PATH message is sent to e_{11} . Then e_{11} determines

Figure 2. Resource sharing optimization.

the next domain (Domain 2) using its BGP table or by consulting the PCE and forwards the PATH message to i_{12} . Note that any internal topology information that may be stored in a record route object (RRO) should be filtered out by the egress gateways. The gateway i_{12} consults the PCE of Domain 2 and also sends the information contained in the GSRO to the PCE. The PCE first determines the gateway i_{22} with this information. The rest of the path computation (e_{12} , e_{22} , P_2 , PB_2 , and B_2) is the same as in the first domain, and a second GSRO and ERO are attached. The PATH message is finally forwarded by e_{12} to the destination node via i_{1d} .

Upon reception of the PATH message containing the two GSROs and two EROs, the destination replies with a reservation (RESV) message. This RESV message is sent to e_{12} via i_{1d} and forwarded to i_{12} , and then the paths P_2 and PB_2 are reserved in the optical layer according to the procedure detailed in [10]. If the set up is successful, the RESV message is forwarded to Domain 1, and the paths P_1 and PB_1 are reserved in the same way. Finally, the source receives the RESV message. Then, the working LSP is reserved and can be used.

Next, the back-up LSP is signaled in the same way using the same identifier in the session object as used for the primary LSP. The source sends a PATH message containing the two GSROs to i_{21} via e_{2s} . This gateway i_{21} then consults its PCE. The PCE notices that the request is coming from the back-up gateway for this con-

Figure 3. Multilayer view of the setup.

■ Figure 4. Failure resolution.

■ Figure 5. Simulation network topology [11].

nection (by inspection of the RSVP-TE session object) and replies with the route information for B_1 . The PATH message with an ERO is sent to Domain 2 (after possible filtering of the RRO by the egress node); an ERO with path key for B_2 is added, and the message is forwarded toward the destination. The destination replies with a final RESV message (the GSROs can be discarded), and the back-up path is set up using standard ERO messages. When the head-end receives the RESV message, the set up is complete.

The reason to perform the path reservation in two phases is that RSVP-TE does not allow path set up initiated by any node other than the path head-end (meaning that node A cannot set up a connection between nodes B and C).

Although discussing all of the details is beyond the scope of this article, there are some issues worth mentioning. First, it must be possible that the head-end specifies all (abstract) nodes or gateways it wants to use in loose EROs and/or GSROs (in case back-up gateways are specified). Because these objects implicitly or explicitly specify the egress router priorities, the PCE should be able to choose their priority if the preferred connection is not possible (and notify the head-end). Second, for resource sharing, RSVP-TE uses an association object. The path B must be associated with both PB and P for resource sharing. Third, for a failure of the primary path, there are two possible restoration solutions: the path PB initiated by the gateway and the path B initiated by the head-end. Because there are two solutions, there is a race condition between these two back ups. In an ideal scenario, the intra-domain restoration performed by the gateway should be under 50 ms, and the head-end should not notice the interruption. However, any escalation strategy developed for multilayer race conditions, such as hold-off timers [3] can be adopted readily for this multidomain race condition. Fourth, the path structure is computed as a single entity; only the signaling is performed in two phases to comply as much as possible with current standards. There is little possibility that after the working LSP has been set up, the reservation of the path B would fail in some domain. This is the reason why the PCE stores the information for all three paths during the set up of the working LSP, and why the back-up LSP should be signaled with the same session identifier as the working LSP. When (in Domain x) the RESV message is successfully returned for P_x and PB_x , the resources for B_x must be marked as *used* by the PCE to avoid blocking of B_x during its PATH-RESV phase. Should the path B_x fail (for instance, due to a failure in the short timeframe

Figure 6. Simulation results.

between the RESV of P_x and PB_x and the PATH-RESV of B_x), the PCE for that domain should try to find an alternate path for B_x . If this fails, the entire session should be rolled back, and the connection blocked. Note that this is a highly unlikely scenario. Fifth, in failure-free operation, the head-end has a working LSP and a back-up LSP available for the connection. Note that, although both are disjoint in the IP/MPLS layer, it is possible that the back-up LSP optical segment (B_x) shares optical resources with the working LSP (P_x) or its protection lightpath (PB_x) in some domain x . This means that these LSPs cannot be considered equal. The sharing of resources for protection between multiple layers in a network is known as *common pool sharing* [3]. This is why we call this solution a common pool multi-domain multilayer protection solution.

FAILURE SCENARIOS

In Fig. 4, we illustrate two examples of considered failures. In the top example, we show what happens when a failure affects a primary path P in Domain 2. In this case, the traffic along P is redirected over its back-up path PB , and the working LSP survives. First, the reserved resources along PB are activated for PB , which means the back-up path B is preempted. All resources in this domain along P are released. Now the preemption of the back-up segment B renders the back-up LSP invalid.

In the bottom example, an LSR failure along the working path is considered. In our approach, this means a failure of the working LSP and all traffic is routed over the back-up LSP. This is a head-end operated recovery action. The head-end activates the back-up LSP, overriding the PB paths in every domain, and thus PB resources are automatically released. After that, the working path P resources are released in every domain, which must be done through control plane signaling. If there are enough resources, each domain could try to further protect the B path after this failure. An OXC failure along the working path has the same end result.

SIMULATION RESULTS

We evaluated the computation of the structure on a pan-European network taken from the e-Photon/One project [11]. There are 17 domains to be interconnected over a physical infrastructure with 67 nodes and 120 links (average degree 3.58). Minimum degree is two, maximum degree is six (Fig. 5).

The domains largely correspond with the countries, and each node is a gateway for a certain domain, so the number of gateways varies between two (Portugal) and six (Germany) gateways. Some countries are grouped together in order to have at least two gateways available.

Figure 5 shows the ILP path result for the connection between the Czech Republic (Praha, Ostrava) and Spain (Madrid, Gijon). The total number of hops consumed is 16, whereas the simple approach uses 18 hops, an 11 percent improvement.

Simulation results are depicted in Fig. 6. Figure 6a shows the average efficiency of the ILP solution versus the simple solution. The total number of hops for interconnecting all 17 domains is 1527 (100 percent) using the simple solution. The ILP optimized solution uses 1397 hops, a 8.9 percent gain. The ILP clearly accomplishes this by choosing longer PB paths to allow maximum sharing. Figure 6b gives a view for every domain-pair. There are 17 domains or 136 domain pairs. We note that sometimes the ILP and simple solution coincide (100 percent), but if a large part of the network must be traversed, the ILP can boast serious improvements. The connection of Holland to Spain has a 30 percent improvement (70 percent).

CONCLUSIONS

In this article we presented a method for enabling high availability over multiple optical networks. This method relies on a path structure to be set up in every intermediate domain that is able to cope with both internal failures and a

The structure is fairly easy to compute, and given enough time and resources, an optimal solution is possible in every domain. The very simple heuristic we used performs well on average, but an ILP solution can have advantages for longer connections.

single gateway failure end-to-end. It also keeps head-end recovery actions to a minimum. The structure is fairly easy to compute, and given enough time and resources, an optimal solution is possible in every domain. The very simple heuristic we used performs well on average, but an ILP solution can have advantages for longer connections. We also proposed a new RSVP-TE object, the GSRO, to enable signaling of multiple ingress/egress points into a domain.

ACKNOWLEDGMENTS

This work was partly funded by the European Commission through projects, Information Technologies Society-Next generation Optical network for Broadband European Leadership (IST-NOBEL) phase 2, NoE-ePhoton/One+, Building the Future Optical Network in Europe (BONE), and the Belgian fund for scientific research, project FWO 3G057808.

REFERENCES

- [1] B. Puyet et al., "Benefits of GMPLS for Multilayer Recovery," *IEEE Commun. Mag.*, vol. 43, no. 7, July 2005, pp. 51–59.
- [2] A. Farrel, "GMPLS, Architecture and Implementations," *Morgan Kaufmann Series in Networking*, Elsevier, 2005.
- [3] J.-P. Vasseur et al., "Network Recovery, Protection and Restoration of Optical, SONET-SDH, IP, and MPLS," *Morgan Kaufmann Series in Networking*, Elsevier, 2004.
- [4] A. Farrel, "A Path Computation Element (PCE)-Based Architecture," RFC 4655, Aug. 2006.
- [5] F. Aslam et al., "Interdomain Path Computation, Challenges and Solutions for Label Switched Networks," *IEEE Commun. Mag.*, vol. 45, no. 10, Oct. 2007, pp. 94–101.
- [6] L. Truong et al., "Dynamic Routing for Shared Path Protection in Multidomain Optical Mesh Networks," *J. Opt. Net.*, vol. 5, 2006, pp. 58–74.
- [7] J. W. Suurballe and R. Tarjan, "A Quick Method for Finding Shortest Pairs of Disjoint Paths," *Networks*, vol. 14, 1984, pp. 325–36.
- [8] D. Awduche et al., "RSVP-TE: Extensions to RSVP for LSP Tunnels," RFC 3209, Dec. 2001.
- [9] J.-P. Vasseur et al., "A Per Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)," RFC 5152, Feb. 2008.
- [10] J. P. Lang et al., "RSVP-TE Extensions in Support of End-to-End GMPLS Recovery," RFC4872, May 2007.
- [11] <http://qosip.tmt.bme.hu/~mesko/e1net>

BIOGRAPHIES

DIMITRI STAESSENS () received his M.Sc. degree in computer science in 2004 from Ghent University. In 2005 he joined the optical networking research group of the Department of Information Technology, under a grant from the Interdisciplinary Institute for Broadband Technology (IBBT). His research focuses mainly on resilience in multidomain networks and resilience in transparent optical networks. He participated in multiple European FP6 and FP7 projects, such as IST-NOBEL and NOBEL phase 2, NoE e-Photon/One+, NoE Bone, Strep Dynamic Impairment Constraint Networking (DiCONET), and several national projects on optical networking.

DIDIER COLLE received an M.Sc. degree in electrotechnical engineering (option: communications) from Ghent University in 1997 and a Ph.D. degree in 2002. Since 1997, 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 Professor Piet Demeester. He was granted a postdoctoral scholarship from the Institute for the Promotion of innovation through Science and Technology in Flanders (IWT-Vlaanderen) for the year 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. He has actively been involved in several IST projects (Layers Interworking in Optical Networks [LION], OPTIMIST, DAVID, STOLAS, NOBEL, and LASAGNE), in the COST-action

266 and 291, and in the International Technology Education Association (ITEA)/IWT TBONES project. His work has been published in more than 100 scientific publications in international conferences and journals.

ILSE LIEVENS received an M.Sc. degree in electrotechnical engineering, focusing on telecommunications, in 1994 from Ghent University. She then joined the Department of Information Technology at Ghent University, where she obtained a Ph.D. degree in 2000 in the INTEC-IBCN research group. She is currently working as a post-doctoral assistant in the IBBT-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 inter-university projects (IWT GBOU ONNA). She is author and co-author of several publications in conference proceedings and journals (ECOC, IEEE JSAC, IEEE Commun. Magazine, PNC, etc.).

MARIO PICKAVET has been a professor at Ghent University since 2000, where he teaches courses on discrete mathematics, multimedia networks, and network modeling. His current research interests are related to broadband communication networks (WDM, IP, GMPLS, OPS, OBS) and include design, long-term planning, techno-economical analysis, and routing of core and access networks. He gives special attention to operations research techniques that can be applied for routing and network design. He has published over 150 international journal and conference publications. He is one of the authors of the book [Network Recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS].

PIET DEMEESTER [SM] is a full-time professor at Ghent University where he teaches courses in communication networks. He is the head of the Broadband Communication Networks group (www.ibcn.intec.ugent.be). The group is part of the recently established IBBT (www.ibbt.be). His research interests include: multilayer IP-optical networks, mobile networks, end-to-end quality of service, grid computing, network and service management, and distributed software and multimedia applications. He has published over 500 papers in these areas in international journals and conference proceedings. In this research domain, he was and is a member of several program committees of international conferences, such as: OFC, ECOC, ICC, Globecom, Infocom, and DRCN.

WALTER COLUCCI received his Master's degree in engineering in electronics and telecommunications from the University of Perugia, Italy. Currently, he is a Ph.D. student in the Electronics and Informatics (ETRO) department, Vrije Universiteit Brussels. He is also a staff member of the Computational Modeling Lab (CoMo). His major research activity is in the field of multilayer traffic engineering in ASON/GMPLS networks.

ANN NOWÉ received an M.S. degree from Ghent University, Belgium, in 1987, where she studied mathematics with a minor in computer science. She received a Ph.D. degree from Vrije Universiteit Brussels (VUB), Belgium, in collaboration with Queen Mary and Westfield College, University of London, United Kingdom, in 1994. The subject of her dissertation was located at the intersection of computer science (AI), control theory (fuzzy control), and mathematics (numerical analysis, stochastic approximation). Currently, she is a professor at VUB. Her major areas of interest are artificial intelligence (AI)-learning techniques, in particular, reinforcement learning and learning in multi-agent systems. She is also the director of the Computational Modeling Lab, CoMo, at VUB.

KRIS STEENHAUT received a Ph.D. degree from Vrije Universiteit Brussels in 1995. Currently, she is a lecturer in telecommunications at Erasmushogeschool Brussels and at Vrije Universiteit Brussels. Her research interests include the application of AI techniques for the control of telecommunication networks. She is a member of the ETRO department.

RICARDO ROMERAL received his M.Sc. in telecommunications engineering from the University Carlos III of Madrid (UC3M) in 2001 and his Ph.D. in inter-domain protection schemes in 2007. He is an assistant professor in teletraffic and switching subjects at the University Carlos III of Madrid. His research interests include traffic engineering and reliability in inter-domain IP/MPLS networks.

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) | [Purchase History](#) | [Cart](#) | [Sitemap](#) | [Help](#)

Welcome University of Gent

[Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)[SUPPORT](#)

Results for "((enabling high availability over multiple optical networks)<in>metadata)"

Your search matched 1 of 1818669 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

[e-mail](#) [printer friendly](#)

» Search Options

[View Session History](#)[New Search](#)

» Key

Indicates full text access

IEEE JNL	IEEE Journal or Magazine
IET JNL	IET Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IET CNF	IET Conference Proceeding
IEEE STD	IEEE Standard

Modify Search

[Search](#)☐ Check to search only within this results setDisplay Format: ☒ Citation ☐ Citation & Abstract[IEEE/ET](#)[Books](#)[Educational Courses](#)[Application Notes \[Beta\]](#)

IEEE/ET journals, transactions, letters, magazines, conference proceedings, and standards.

[view selected items](#)[Select All](#) [Deselect All](#)

- ☐ 1. Multidomain optical networks: issues and challenges - Enabling high availability over multiple optical networks
Staessens, D.; Colle, D.; Lievens, I.; Pickavet, M.; Demeester, P.; Colitti, W.; Nowe, A.; Steenhaut, K.; Romeral, R.;
[Communications Magazine, IEEE](#)
Volume 46, Issue 6, June 2008 Page(s):120 - 126
Digital Object Identifier 10.1109/MCOM.2008.4539475
- [AbstractPlus](#) | Full Text: [PDF](#)(295 KB) IEEE JNL
[Rights and Permissions](#)

[Help](#) [Contact Us](#) [Privacy & Security](#) [IEEE.org](#)

© Copyright 2008 IEEE - All Rights Reserved