Recovery in Multilayer Optical Networks

Mario Pickavet, *Member, IEEE*, Piet Demeester, *Senior Member, IEEE*, Didier Colle, *Member, IEEE*, Dimitri Staessens, Bart Puype, Leen Depré, and Ilse Lievens

Tutorial

***Abstract*—The integration of different network technologies into a multilayer network, as in Internet-based networks carried by optical transport networks (OTNs), creates new opportunities but also challenges with respect to network survivability. In different network layers, recovery mechanisms that are active can be exploited jointly to reach a more efficient or faster recovery from failures. This interworking is also indispensable in order to overcome the variety of failure scenarios that can occur in the multilayer-network environment. A well-considered coordination between the different layers and their recovery mechanisms is crucial in order to attain high performance recovery. This paper provides an overview of multilayer recovery issues and solutions in an Internet protocol (IP)-over-optical-network environment, which is illustrated by quantitative case studies.**

***Index Terms*—Multilayer networks, protection, recovery, traffic engineering (TE).**

I. Introduction

BACKBONE networks carrying Internet protocol (IP) traffic, possibly enhanced with multiprotocol-label-switching (MPLS) functionality, are supported by optical transport networks (OTNs) that provide transmission links between IP routers. By applying wavelength division multiplexing (WDM), OTNs are capable of carrying many independent channels, which are carried on different wavelengths, over a single optical fiber. This allows the network to transport huge amounts of data that are needed for many current and future communication services, which play a very important role in many of our daily social and economical activities. For instance, strategic corporate functions show an increasing dependence on communication services.

Communication networks are subject to a wide variety of unintentional failures, which are caused by natural disasters, wear out and overload, software bugs, human errors, etc., as well as intentional interruptions due to maintenance [1]. As core communication networks also play a vital military role, key telecommunication nodes were favored targets during the Gulf War and could become a likely target for terrorist activity. For business customers, disruption of communication can suspend critical operations, which may cause a significant loss of revenue to be reclaimed from the telecommunications provider. In fact, availability agreements now form an important component of service level agreements (SLAs) between providers and customers.

In the cutthroat world of modern telecommunications, network operators need a reliable and maintainable network in order to hold a leading edge over the competition. Fast and scalable network recovery techniques are of paramount importance in order to provide the increasingly stringent levels of reliability that network operators demand for their future networks.

A multilayer transport network typically consists of a stack of single-layer networks. There is usually a client–server relationship between the adjacent layers of this stack. Each of these network layers may have its own (single-layer) recovery schemes. As will be shown in the following sections, it is important to be able to combine recovery schemes in several layers in order to cope with the variety of possible failures in an efficient way and to benefit from the advantages of the schemes in each layer. It is worth mentioning that implementing a multilayer recovery strategy does not mean that all the recovery mechanisms will be used at every layer.

As Internet traffic is continuously shifting and changing in volume over time (for instance, due to diurnal traffic fluctuation and overall traffic growth), there is an ongoing research towards creating optical networks with the flexibility to reconfigure transmission according to traffic demands. This requires the possibility to set up and tear down OTN-layer connections that implement logical links in the higher network layer in real time, which has led to the concept of intelligent optical networks (IONs). In addition to allowing the network to adapt to changing traffic demands, this flexibility in setting up lightpaths on demand turns restoration into a viable recovery option.

In the following sections, three generic approaches for providing recovery in multilayer networks (more specifically in IP-over-OTN networks) will be discussed: single-layer recovery schemes in multilayer networks (Section II) with the important issue of which layer in the network to provide the recovery scheme; static multilayer recovery schemes (Section IV-A) where recovery schemes at several network layers can be provided with the important issue of how to make them interwork (Section III); and the dynamic multilayer recovery strategies (Section IV-B) that exploit logical-topology

Manuscript received July 1, 2005; revised October 3, 2005. This research was supported in part by the European Commission through the IST-projects NOBEL and e-Photon/One and in part by the Flemish Government through the projects IWT-GBOU ONNA, IWT-ITEA TBONES, and FWO-project G.0315.04. The work of B. Puype was supported by the IWT through a Ph.D. Grant.

The authors are with the Department of Information Technology (INTEC), Ghent University–IBBT–IMEC, 9000 Gent, Belgium (e-mail: mario.pickavet@intec.ugent.be; piet.demeester@intec.ugent.be; didier.colle@intec.ugent.be; dimitri.staessens@intec.ugent.be; bart.puype@intec.ugent.be; leen.depré@intec.ugent.be; ilse.lievens@intec.ugent.be).

Digital Object Identifier 10.1109/JLT.2005.861118

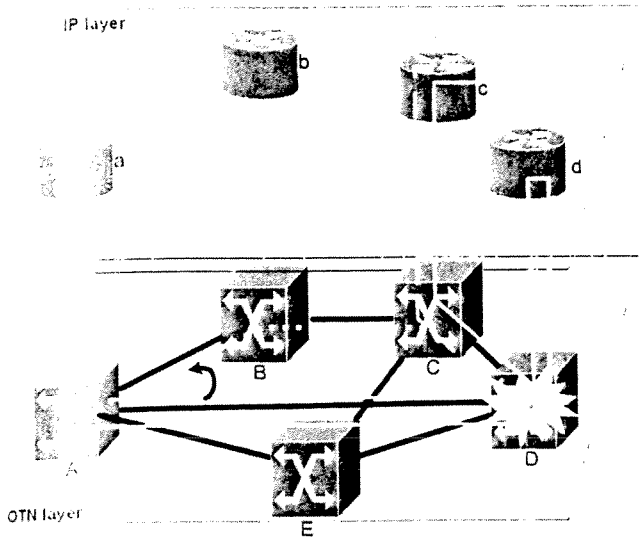

Fig. 1. Survivability at the bottom layer.

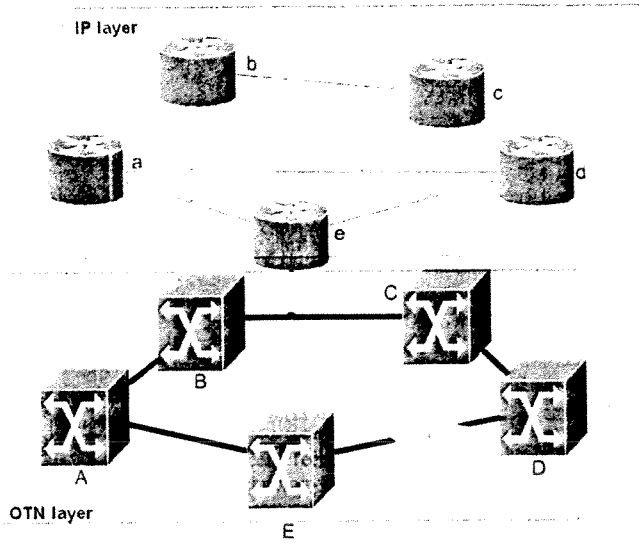

Fig. 2. Survivability at the top layer—secondary failures.

adaptations for survivability purposes. Some quantitative studies and comparisons between the different methods (Sections V and VI) will reveal the advantages and disadvantages of each approach.

II. Single-Layer Recovery in Multilayer Networks

This section discusses how recovery functionality can be introduced in multilayer networks by applying single-layer recovery schemes. The concepts and discussions are focused on a two-layer network but are mostly generic and therefore applicable to any multilayer network.

A. Survivability at the Bottom Layer

In this recovery approach, recovery of a failure is always done at the bottom layer of the multilayer network. In an IP-over-OTN network, for example, this implies that the 1 + 1 optical protection scheme, or any other recovery scheme that is deployed at the OTN layer, attempts to restore the affected traffic in case of a failure.

By recovering a failure at the bottom layer, this strategy has the benefit that only a simple root failure has to be treated and that the number of required recovery actions is minimal (the recovery actions are performed on the coarsest granularity). In addition, failures do not need to propagate through multiple layers before triggering any recovery action.

However, this recovery strategy cannot handle problems that occur due to failures in a higher network layer. Moreover, if a node failure occurs in the OTN layer [being an optical-cross-connect (OXC) failure], the OTN-layer recovery mechanism will only be able to restore the affected traffic that transits the failed bottom-layer node (being the OXC). The colocated higher-layer IP router will become isolated due to the failure of the OXC underneath; thus, all traffic treated by this IP router cannot be restored in the lower (optical) layer.

This is illustrated in the example of Fig. 1. Throughout this paper, we will label the top-level nodes lower case and the bottom-layer nodes in upper case. The considered network carries two traffic flows between client-layer nodes a and c. One traffic flow (a–d–c, indicated with a full line) transits the client-layer node d (using two logical links a–d and d–c), while the other traffic flow (a–c, indicated with a dashed line) uses a direct logical link from a to c and only transits the server-layer node D. Now let us assume that a failure occurs in the bottom layer, for example, the failure of node D. The figure illustrates that the server layer cannot recover the first traffic flow a–d–c. This is due to the fact that the client-layer node d is isolated due to the failure of D, which is terminating both logical links a–d and d–c. This failure can only be resolved at the higher layer. However, the second traffic flow a–c is routed over a direct logical link between nodes a and c. This logical link transits only the failing node D in the bottom layer, which means that this traffic flow can be restored by the bottom-layer recovery scheme (dotted line in the figure).

B. Survivability at the Top Layer

Another strategy for providing survivability in a multilayered network is to provide the survivability at the top layer of the network. In our example of an IP-over-OTN network, this could be the IP restoration technique or MPLS-based restoration (see [1] for a detailed overview of IP and MPLS recovery techniques). The main advantage of this strategy is that it can also cope with higher-layer failures. However, a major drawback of this strategy is that it typically requires a lot of recovery actions, due to the finer granularity of the flow entities at the top layer.

As a consequence of a single root failure in the lower layer, a complex scenario of secondary failures is typically induced in the higher network layer. This is illustrated in Fig. 2, where the failure of an optical link in the bottom layer corresponds with the simultaneous failure of three logical IP links in the top layer. Hence, these three logical IP links are part of a shared risk link group (SRLG) [2]. This implies that the recovery scheme in the top layer will have to recover from three simultaneous link failures, a quite-complex failure scenario. This is in clear contrast with a recovery scheme at the bottom layer that

would only have to cope with the simpler scenario of a single link failure.

Another disadvantage of a recovery at the top layer only is that the traffic injected directly in the lower layer (e.g., wavelength channels directly leased by a customer) cannot be recovered by the optical-network operator, even if the failure happens in the optical layer itself.

C. Variants

A slightly different variant on the strategy that applies survivability at the bottom layer is the survivability at the lowest-detecting-layer strategy. The lowest detecting layer is the lowest layer in the layered network hierarchy that is able to detect the failure. This implies that multiple layers in the network will now deploy a recovery scheme, but the (single) layer that detects the root failure is still the only layer that takes any recovery actions. With this kind of strategy, the problem of the bottom-layer recovery scheme not detecting a higher-layer failure is avoided, because the higher layer will detect the failure and will recover the affected traffic. However, this strategy requires that you can determine which layer is the lowest detecting layer (to avoid the condition where higher layers also react upon a lower-layer failure). Moreover, it still suffers from the fact that it cannot restore any traffic transiting a higher layer equipment isolated by a node failure in the layer below. With this strategy, the client layer in the example (Fig. 1) will deploy a recovery scheme, but the considered traffic flow a–d–c is still lost since this client-layer recovery scheme is not triggered by the occurrence of the node failure in the server layer. Therefore, although this strategy considers the deployment of recovery schemes in multiple layers, it is still considered as a single-layer survivability strategy in a multilayer network since for each failure scenario, the responsibility to recover all traffic is situated in only one layer (being the lowest one detecting the failure).

A slightly different variant of the strategy that provides survivability at the top layer is the survivability at the highest-possible-layer strategy. Since not all traffic have to be injected (by the customer) at the top layer, with this strategy, a traffic flow is recovered in the layer in which it is injected; in other words, the highest possible layer for this traffic flow. This means that this highest possible layer is to be determined on a per-traffic-flow basis. This survivability at the highest-possible-layer strategy is also considered as a single-layer survivability strategy for providing survivability in a multilayer network, even though it considers recovery schemes in multiple layers. Indeed, the recovery schemes in multiple layers will never recover the same traffic flow. Actually, this strategy deploys the survivability at the top layer strategy for each traffic flow individually.

III. Interworking Between Layers

In the previous section, some strategies are discussed that apply a single-layer recovery mechanism in order to provide survivability in the multilayer network. The advantages of these approaches can be combined by running recovery mechanisms in different layers of the network as a reaction to the occurrence of a single network failure. Generally, the choice of which layer(s) to recover the affected traffic due to a failure will depend on the circumstances, for example, on the failure scenario that occurred.

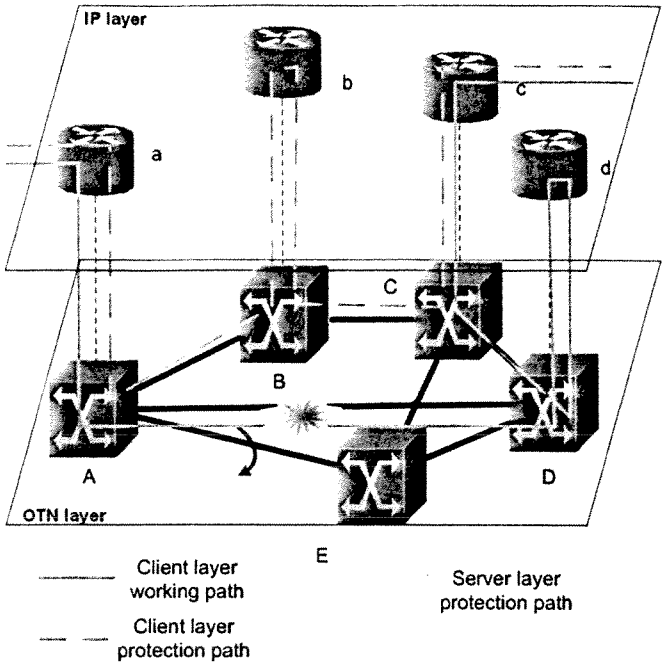

Fig. 3. Uncoordinated multilayer survivability strategy.

However, this interworking between layers requires some rules or coordination actions in order to ensure an efficient recovery process. These rules strictly define how the layers and the recovery mechanisms within those layers react to different failure scenarios and form a so-called escalation strategy. Several escalation strategies are discussed in this section: uncoordinated, sequential, and integrated escalation.

A. Uncoordinated Approach

The easiest way of providing an escalation strategy is to simply deploy recovery schemes in the multiple layers without any coordination at all. This will result in parallel recovery actions at distinct layers. Consider again the two-layered network (Fig. 3) with, for instance, the failure of physical link A–D in the server layer. This failure of the physical link will also affect the corresponding logical link a–d in the client layer, and hence affects the considered traffic flow a–d–c. Since the recovery actions in both layers are not coordinated, both the recovery strategy in the client layer and the recovery strategy in the server layer will attempt recovery of the affected traffic. This implies that, in the client layer, the traffic flow from a to c is rerouted by the recovery mechanism of the client layer, which results in a replacement of the failed path a–d–c by a new path (for instance, a–b–c). At the same time, the server layer recovers the logical link a–d of the client-layer topology by rerouting all traffic on the failing link A–D through node E. It is clear that, in this example, recovery actions in a single layer would have been sufficient to restore the affected traffic.

The main advantage of the uncoordinated approach is that this solution is simple and straightforward from an implementation and operational point of view. However, Fig. 3 shows the drawbacks of this strategy. Both recovery mechanisms occupy spare resources during the failure, although one recovery scheme occupying spare resources would have been sufficient. Spare resources are usually used to accommodate low-priority traffic during failure-free conditions, but this so-called "extra traffic" must be preempted when the spare resources are needed to recover from a failure. Hence, a repercussion of the uncoordinated approach is that more extra traffic than necessary is potentially disrupted. The situation can even be worse. For example, consider that the server layer reroutes the logical link a–d over the path A–B–C–D instead of A–E–D, then both recovery mechanisms need spare capacity on links A–B and B–C. If these higher-layer spare resources are supported as extra traffic in the lower layer, then there is a risk that these client-layer spare resources are preempted by the recovery action in the server layer, resulting in "destructive interference." In other words, none of the two recovery actions was able to restore the traffic, since the client layer reroutes the considered flow over the path a–b–c, which was disrupted by the server-layer recovery. The research done in [3] illustrates that these risks may exist in real networks—the authors prove that a switchover in the optical domain may trigger traditional client-layer protection. Moreover, such a multilayer recovery strategy can have significant repercussions on the overall network stability.

In [4], the authors show a real-life example of network convergence problems that follow the impetuous use of the uncoordinated approach in an IP-over-OTN network, where the OTN-layer features the 1 + 1 path protection. They observed IP network convergence times after the occurrence of a link failure in the OTN layer. Although protection in the optical layer recovers a link within 20 ms, the recovery of the IP traffic that was transiting the link takes over 60 s in some cases. These slow recovery times are results of the IP-layer-topology discovery algorithms trying to rediscover the new IP-network topology, while the OTN layer is recovering by switching over to the backup fiber. More specifically, the authors show that intermediate system-intermediate system protocol (IS-IS) adjacency recovery may take up to 13 s, and IS-IS route recovery may take up to 18 s, and depending on the border gateway protocol (BGP) scanning timing, BGP routes recovery may take up 80 s if relevant interior gateway protocol (IGP) topology information is lost. Note that this problem can be solved easily with a sequential approach using a hold-off timer (see next section).

In summary, although simple and straightforward, letting the recovery mechanisms in each layer run without a coordinating escalation strategy has consequences on efficiency, capacity requirements, and even ability to restore the traffic.

B. Sequential Approach

In comparison with the uncoordinated approach, the sequential approach is a more efficient escalation strategy. Here, the responsibility for recovery is handed over to the next layer when it is clear that the current network layer is not able to perform the recovery task. For this escalation strategy, two questions must be answered: in which layer to start the recovery process and when to escalate to the next layer. Two approaches exist, the bottom-up escalation strategy and the top-down escalation approach, each having different variants.

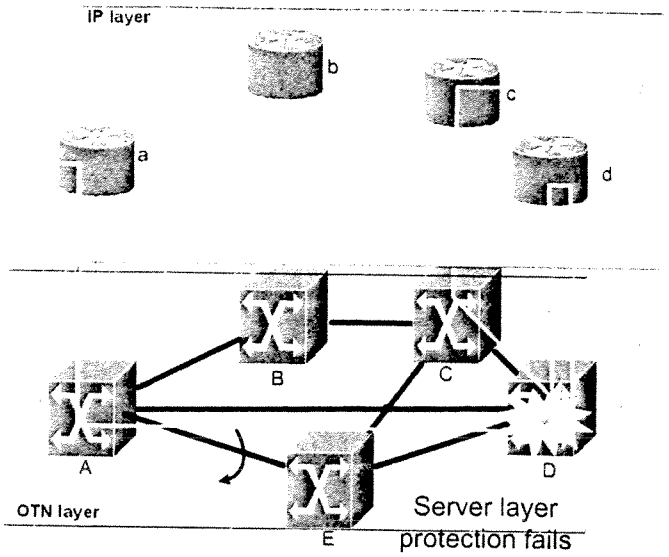

Fig. 4. Phase 1—recovery action in the server layer.

1) Bottom-Up Escalation: With this strategy, the recovery starts in the lowest detecting layer and escalates upwards. The higher-layer recovery scheme will only try to recover affected traffic that could not be recovered by the lower layer. The advantage of this approach is that recovery actions are taken at the appropriate granularity: First, the coarse granularities are handled, recovering as much traffic as soon as possible; and recovery actions on a finer granularity (i.e., in a higher layer) only have to recover a small fraction of the affected traffic. This also implies that complex secondary failures are handled only when needed. For instance, in the client-server example of Fig. 1, there is the failure of OXC D as the root failure. This corresponds with the simultaneous failure of three IP links (a–d, a–c, and d–c) in the client layer. Since the server-layer recovery mechanism can recover the flow a–c (by finding an alternative path for the optical-layer flow A–D–C), the client-layer recovery mechanism will only have to handle the recovery of the traffic over links a–d and d–c, being less complex than the simultaneous failure of three links.

This is illustrated in Figs. 4 and 5. The server layer starts with the recovery process (Fig. 4), attempting to restore the logical link a–d. The server layer fails in this recovery, since this logical link terminates on the failing node D. As such, the client-layer recovery scheme (Fig. 5) is triggered (the implementation of this trigger mechanism is discussed at the end of this section) to restore the corresponding affected traffic flow a–c (originally following the route a–d–c) by rerouting it over node b instead of node d.

An issue that must be handled in the bottom-up escalation strategy is how a higher network layer knows whether it is the lowest layer that detects the failure (so it can start with the recovery) or not (and has to wait for a lower layer instead). Typically, the fault signals that are exchanged to indicate a failure will carry sufficient information, so the layer where the

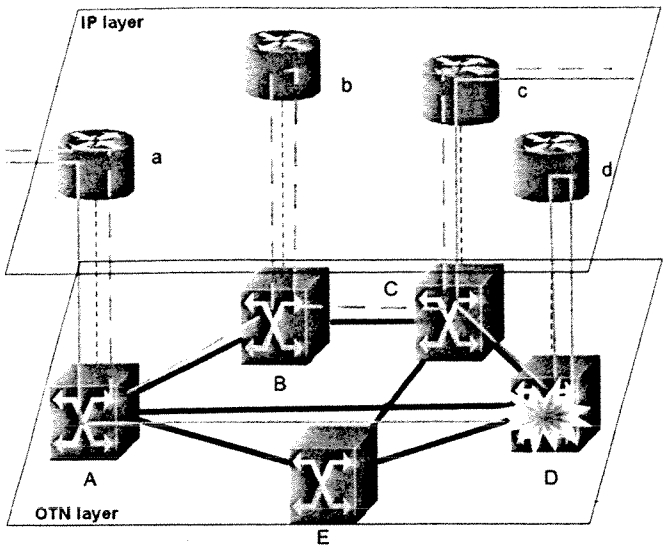

Fig. 5. Phase 2—recovery action in the client layer.

failure occurred can be derived. However, suppose that this is not the case. Assume that we have a four-layer network, where a failure occurs in the bottom layer. Assume that the failure is detected in all four layers at the same time and that the layer where the failure has occurred cannot be derived from those signals. This means that each of the higher layers can assume that it is the lowest detecting layer and start with the recovery. This can be overcome by appropriately using the mechanism of hold-off timers (see below), which are set progressively higher as we move upwards in the stack of layers. In this way, the recovery mechanisms in the higher layers will give their server layers an opportunity to perform the recovery.

2) Top-Down Escalation: With top-down escalation, it is the other way around. Recovery actions are now initiated in the highest possible layer, and the escalation goes downwards in the layered network. Only if the higher layer cannot restore all traffic, actions in the lower network layer are triggered. An advantage of this approach is that a higher layer can more easily differentiate traffic with respect to service types; therefore, it can try to restore high-priority traffic first. However, a drawback of this approach is that a lower layer has no easy way to detect on its own whether a higher layer was able to restore the traffic (an explicit signal is needed for this purpose). Thus, the implementation here is somewhat more complex and currently not applied. There is also a problem of efficiency, since it is very well possible that, for example, 50% of the traffic carried by a wavelength channel in an optical network is already restored by a higher-network-layer recovery mechanism; hence, protecting this wavelength in the optical layer as well is only useful for the other 50% of the carried traffic.

C. Implementation of an Escalation Strategy

The actual implementation of these escalation strategies is another issue. Two possible solutions are described here (for ease of explanation, the bottom-up escalation strategy is assumed in the following discussion).

The first implementation solution is based on a hold-off timer of T_{HO} seconds. Upon detection of a failure, the server layer starts the recovery immediately, while the recovery mechanism in the client layer has a built-in hold-off timer that must expire before initiating its recovery process. In this way, no client recovery action will be taken if the failure is resolved by the server-layer recovery mechanism before the hold-off timer expires. The main drawback of a hold-off timer is that recovery actions in a higher layer are always delayed regardless of the failure scenario. The challenge of determining the optimal value for T_{HO} is driven by a tradeoff between recovery time versus network stability and recovery performance.

The second escalation implementation overcomes this delay by using a recovery-token signal between layers. This means that the server layer sends the recovery token (by means of an explicit signal) to the client layer upon knowing that it cannot recover (all or part of) the traffic. Upon reception of this token, the client-layer recovery mechanism is initiated. This allows limiting the traffic disruption time, in case the server layer is unable to do the recovery. Compared to the hold-off-timer interworking, one disadvantage is that a recovery-token signal needs to be included in the standardization of the interface between network layers.

Note that, at the time of writing, only the timer-based approach was available in commercial networking products.

IV. Multilayer Survivability Strategies

A. Static Recovery Techniques

Multilayer survivability involves more than just coordinating the recovery actions in multiple layers. There is also the issue of spare resources, and how they have to be provided and used in an efficient way in the different layers of the network. One way or another, the logical (spare) capacity assigned to the recovery mechanisms that are deployed at higher network layers must be transported by the lower layer. There are several ways to realize this.

The most straightforward option is called duplicated protection and is depicted in Fig. 6 for a point-to-point example. (Note that we made an abstraction from the physical disjointness of working and backup paths in this conceptual example. The extension towards larger networks and the introduction of physical disjointness is straightforward.)

Each working IP link is transported via a lightpath in the OTN layer. To cope with OTN-layer failures, the lightpath is protected by a backup lightpath. To cope with IP-layer failures, the IP link is protected by a spare IP link (to be transported via the OTN layer as well). Moreover, if the spare capacity that is provisioned in the logical IP network is simply protected again in the underlying optical layer (backup lightpath for lightpath of spare IP link), we are coping with duplicated protection. Despite the reduced complexity, this is a rather expensive solution. Hence, investing in duplicated protection is very debatable and probably only meaningful in a few exceptional network scenarios.

The first possibility to save investment in physical capacity is by carrying the spare capacity in the logical higher-layer

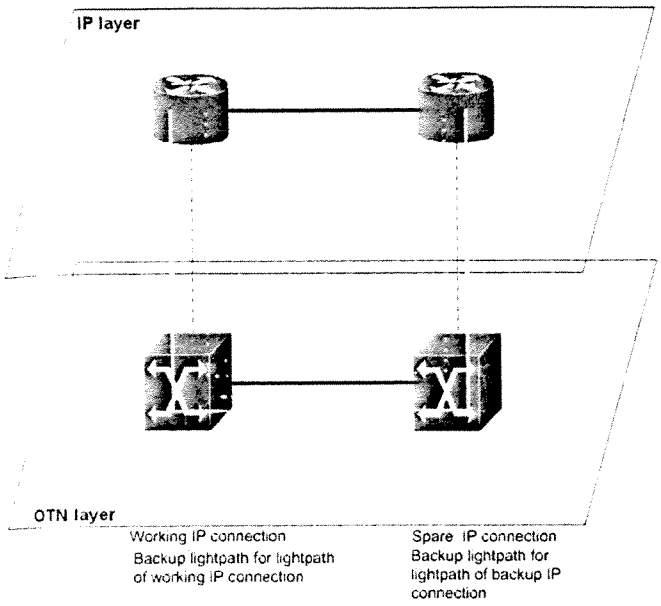

Fig. 6. Duplicated protection.

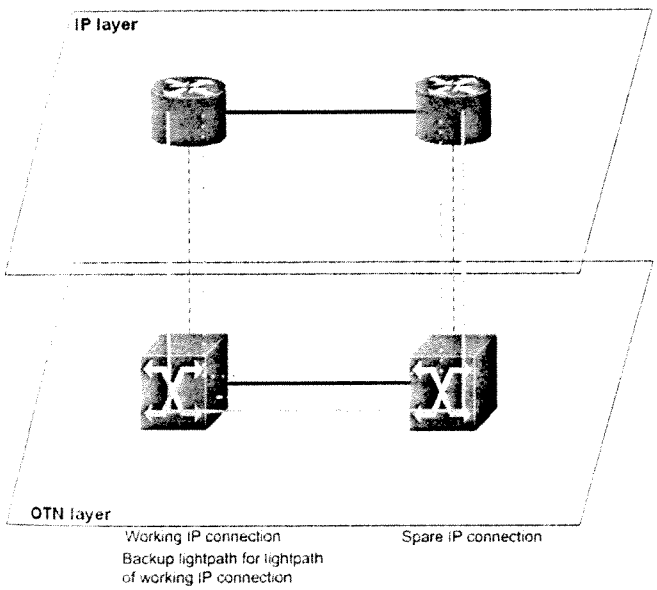

Fig. 7. Logical spare unprotected.

network allocated to the higher layer-network recovery techniques as unprotected traffic in the underlying network layer(s) (see Fig. 7 for the IP-over-OTN example).

This strategy, which is called logical spare unprotected, still allows protecting against any single failure: A cut of the bottom fiber (carrying the lightpath of the working IP link) would trigger the optical-network recovery, while a failure in one of the outer router line cards would trigger the IP-layer network recovery. A prerequisite for such a scenario is that the optical network supports both protected and unprotected lightpaths. It is crucial to guarantee that a single network failure is not able to affect simultaneously a working IP link and the unprotected lightpath that is carrying the IP spare capacity protecting that same link. Otherwise, the spare IP capacity would also become unavailable for recovery of the failure, and the recovery process would fail.

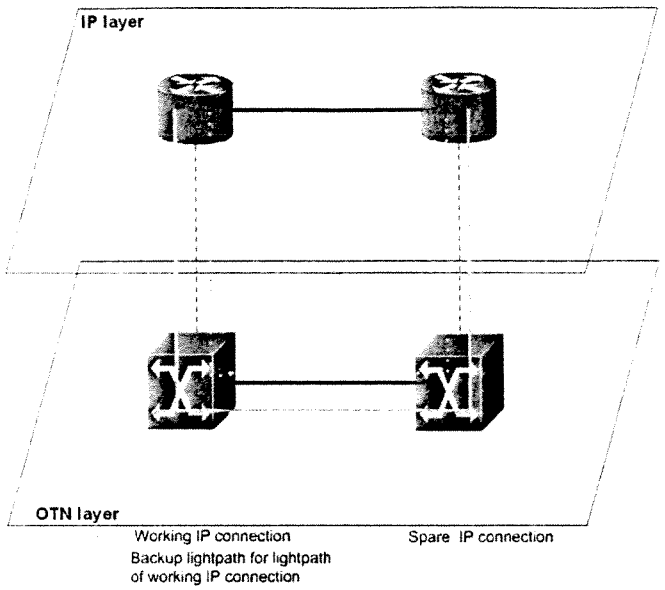

Fig. 8. Common pool strategy.

One step beyond simply carrying the spare capacity of the logical higher network layers as unprotected traffic in the underlying layer is to allow preempting this unprotected traffic by the network recovery technique of the lower network layer. This is the common pool strategy [5], and an example for an IP-over-OTN network is given in Fig. 8. OTN spare resources are provisioned for the optical protection of the lightpath implementing the working logical IP link. The lightpath implementing the spare logical IP link is then routed along the same (optical) spare resources. In case of a failure in the fiber carrying the working logical IP link, the optical protection will be triggered, which preempts the lightpath implementing the spare logical IP link. In that case, there is no problem in preempting this lightpath since it is not needed in the failure scenario. However, the preemption of lightpaths carrying logical spare capacity requires additional complexity. In summary, the common pool strategy provides a pool of physical spare capacity that can be used by the recovery technique in either the IP or the optical layer (but not simultaneously).

B. Dynamic Recovery Techniques

In the previous section, static multilayer recovery strategies have been discussed. They are called static because the logical-network topology (in an IP-over-OTN network, this is the IP-layer topology) remains unchanged (static) at the time of a failure. As such, the logical network must be provided with a recovery technique and the required spare resources for survivability reasons.

Dynamic multilayer survivability strategies differ from such static strategies in the sense that they actually use logical-topology modification for recovery purposes. This requires the possibility to set up and tear down lower layer-network connections that implement logical links in the higher network layer in real time. Therefore, optical networks will be enhanced with a control plane, which gives the client networks the possibility to initiate the setup and teardown of lightpaths in

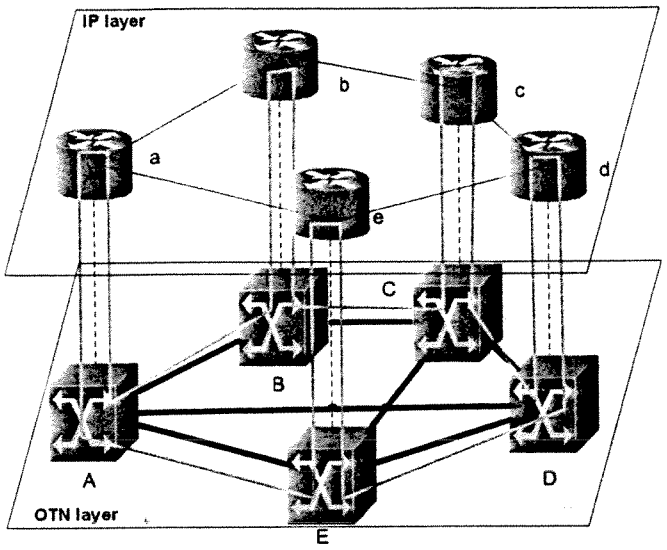

Fig. 9. Scenario before failure.

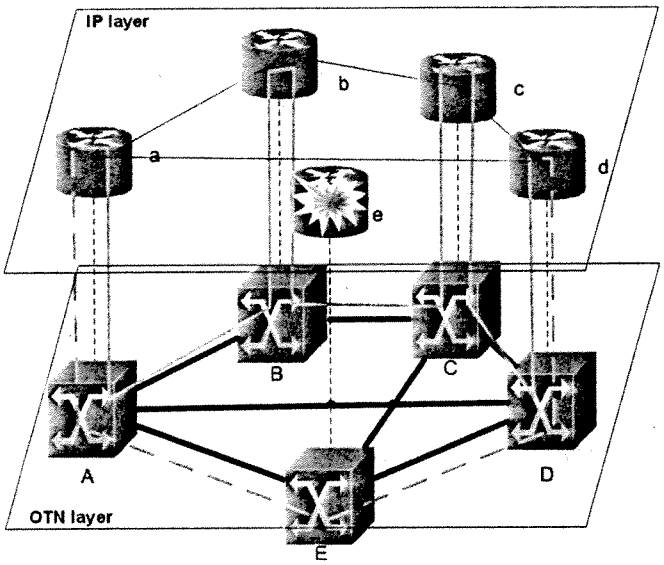

Fig. 10. Scenario after failure.

the optical layer. This is used to reconfigure the logical IP network in case of a network failure. This approach has the advantage that the logical-network spare resources should not be established in advance in the logical IP network; thus, the underlying optical network should not care about how to treat these client-layer spare resources. However, in the optical layer, a spare capacity still has to be provided to deal with lower layer failures such as cable cuts or OXC failures. Enough capacity is also needed in the optical layer to support the reconfiguration of the logical-IP-network topology and the traffic routed on that topology.

An illustration of a dynamic reconfiguration of the logical higher-layer topology in case of failures is given in Figs. 9 and 10 for an IP-over-OTN network. Initially, the traffic flow from router a to router d is forwarded via the intermediate router e. To this end, the logical IP network contains IP links a–e and e–d, which is implemented by lightpaths A–E and E–D in the OTN network. When router e fails, routers a and d will detect this failure and use the user-network interface (UNI) to request the optical layer for a teardown of the links a–e and e–d. The resulting free capacity in the optical layer can be used to set up a direct logical IP link from router a to router d. This is requested to the underlying optical network by requesting the setup of the lightpath between OXCs A and D. Thus, at the time of the failure, the logical-IP-network topology is reconfigured. As mentioned before, a special feature of the underlying optical network is needed for this: It must be able to provide a switched connection service to the client network quickly. Automatic switched optical networks (ASONs) [6], or more generally intelligent optical networks (IONs), have this particular feature.

V. Logical-Topology Design

A challenge with dynamic multilayer recovery strategies involves the actual logical topologies to be realized. Typically, a network scenario and a traffic demand will favor a certain logical topology (and a corresponding IP/MPLS routing) for the failure-free case. Network failures affecting part of this logical topology will require topology reconfigurations and rerouting to replace failing links and to circumvent the problem.

To illustrate the flexibility in logical-topology design, two clearly distinct methods can be proposed. The method of global reconfiguration considers each failure scenario separately. The IP topology is recomputed from scratch for each failure scenario (after removal of the failing network elements). The method of local reconfiguration follows a quite different approach. The topology design now starts from the failure-free case. For a particular failure scenario, the affected IP links and/or routes are first removed, and only the affected traffic is rerouted over the remaining topology (adding IP links where additional resources are needed). The idea behind local reconfiguration is that it lowers the amount of required reconfigurations and rerouted traffic in case of failures, since the new logical topology is derived from the failure-free one.

In addition to these two approaches, the design of the logical topology requires a certain computational ability, and this can be performed online as well as offline. Online reconfiguration and rerouting is more suitable to adapt to a changing network scenario and especially to changing traffic demands. Cross-layer traffic-engineering (TE) techniques can be used to redistribute lower-layer capacity to better cope with traffic or to optimize bandwidth throughput regardless of traffic pattern. However, because of its online nature, computation time must remain limited for the logical-topology update mechanism to retain its desired flexibility. The capability to deal with network failures should arise naturally from the process' goal to solve changing situations in the network automatically.

In cases with a more static traffic demand, one may prefer offline logical-topology design instead. In this case, the traffic demand is used to calculate a failure-free logical topology and a set of failure topologies corresponding with possible network failures. Note that, in this case, one needs to decide in advance which failures should be recoverable. More time is available because the design is done offline, so more extensive

topology-design algorithms can be utilized. This method is typically more optimal in terms of network cost, throughput, impact of failure on operation, etc.

In this section, examples of both online and offline methods will be discussed. The next section will present some results from case studies performed on these example recovery strategies.

A. Online Configuration

Multilayer TE (MTE) is a type of TE combining existing TE techniques [like link weights in IP or TE label switched paths (LSPs) in MPLS] with the lightpath-setup flexibility of the underlying optical layer. MTE cannot only reroute IP flows when traffic demands vary over time, it also allows on-the-fly logical-topology reconfiguration when those variations exceed the acceptable range for simple rerouting. The main objective of an MTE strategy is to accommodate traffic demands in the most optimal way, in terms of QoS, total throughput, network cost, etc.

This is realized by solving discrepancies between the logical-topology configuration and the offered traffic pattern. These discrepancies are usually experienced as congestion or QoS degradation. In addition to logical-topology configuration, routing of the IP traffic also needs to be performed. In fact, both two aspects of MTE will interact nontrivially since the IP routing is typically more complex than OSPF-based shortest paths (for example, taking advantage of explicit MPLS LSP routing in order to do IP-layer TE). As such, MTE shows some similarities with dynamic grooming. The MTE strategy can 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 teardown) of IP links and rerouting of traffic. The Internet Engineering Task Force (IETF) currently explores such approaches, for example, by means of path-computation-element (PCE) techniques whereby a path computation is involved across multiple layers to perform optimal rerouting choices [7].

Although not a dedicated recovery strategy, the MTE-based recovery is a form of dynamic multilayer recovery; its resilience properties are a by-product of the objective to cope with network problems (congestion or otherwise). However, the configuration of the logical topology is not predimensioned but decided upon online, at the time of the failure (or congestion).

Fig. 11 shows a simple example of MTE actions taken upon an IP router failure. Only the IP layer is shown (upper part: failure-free; lower part: router c fails) with 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 de), but also over newly setup capacity (link ad). The routes of flows are consequently decided by the MTE strategy, which tries to optimize IP-layer efficiency (IP-link filling) and optical-layer efficiency (amount of required lightpaths) and, in fact, finds a suitable compromise between both. In this case, the affected flow a–e was rerouted over two links, one of them a replacement IP link. Additionally, the new link ad attracts the

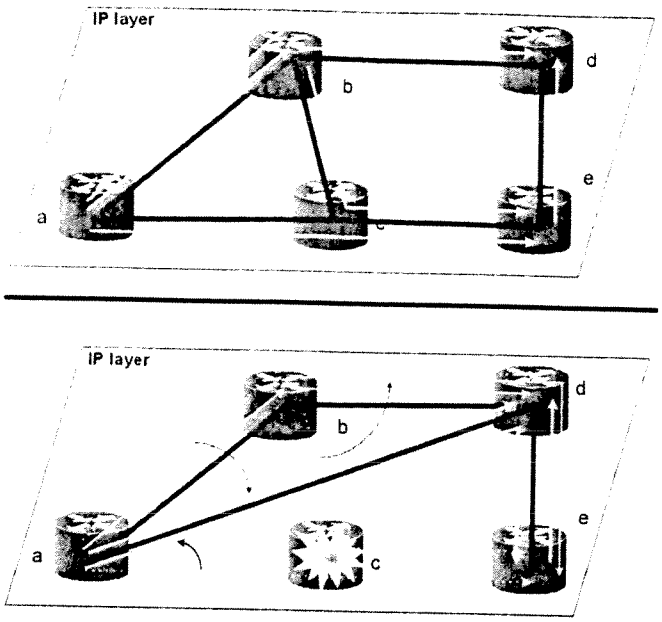

Fig. 11. Example of MTE-based recovery.

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 that we may see some secondary effects beyond the recovery of the affected traffic flows, similar to the effects of global versus local reconfigurations that were discussed earlier.

Note that such a dynamic reconfiguration mechanism may require dampening algorithms to create a certain amount of inertia to avoid triggering a network reconfiguration in case of a transitional failure. Consequently, an interesting viable 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 when the failure duration exceeds some period of time.

A straightforward approach to MTE is achieved by shifting transit traffic from the IP/MPLS to the optical layer; Sato [8] presents a generalized-multiprotocol-label-switching (GMPLS) [9], [10] based hardware implementation. Various MTE and multilayer routing strategies exist. For example, the integrated approach in [11] routes incoming traffic requests using a multistep algorithm. It identifies some main problems of multilayer routing: optical-layer wavelength continuity constraints (because of the high cost of full wavelength translating OXCs) and electrical/optical bandwidth granularity discrepancies. The mechanism in [12] separates the problem in logical-topology design (offline), dynamic routing (online), and bandwidth adjustment (traffic driven) modules. In [13], online routing is performed over a single graph, which represents the IP/MPLS and optical networks—its main contribution being the cost model to integrate both layers in a single model. In Section VI, we will present a strategy assuming an IP-over-optical overlay model and largely based on IP-layer load-dependent costs.

B. Offline Optimization

Fig. 12 sketches how the required OTN resources can be calculated for a static and a dynamic multilayer recovery scheme

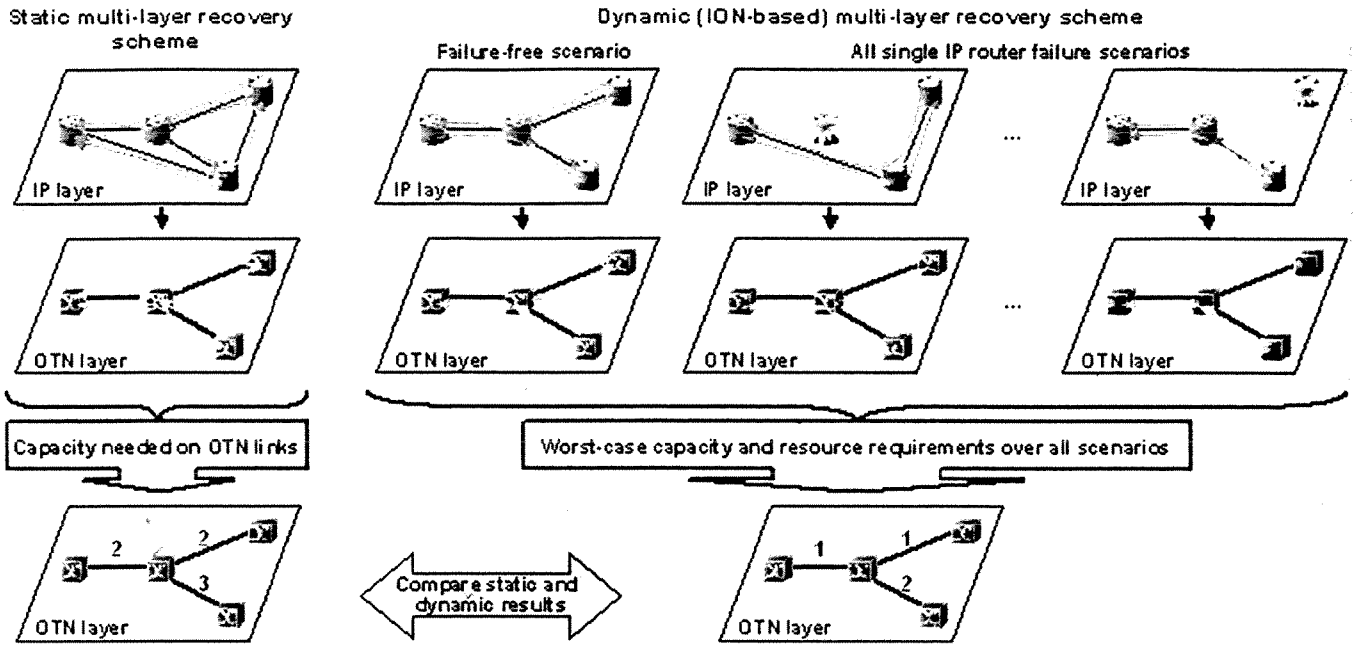

Fig. 12. (Left) Static multilayer resilience scheme versus (Right) dynamic multilayer resilience scheme using ION flexibility.

(in the IP layer, some working and spare LSPs are shown; the IP topology in the static case has to be biconnected to allow MPLS recovery of router failures). The bottom part of the figure shows the calculation results: the actual resource requirements on the OTN links for both strategies, showing a significant improvement in the case of dynamic multilayer recovery.

The dynamic approach has the advantage of being highly efficient in terms of required backup capacity. There are definitely some issues and challenges. Compared to the static multilayer approach, the recovery time is likely to be significantly higher. Indeed, with the static multilayer approach, the routers adjacent to the failed or isolated router can quickly detect the failure and the network can converge in a short period of time by means of a fast IP routing or an MPLS TE fast reroute. The dynamic multilayer recovery approach on the other hand requires the IP router(s) to signal the setup of new IP links via the UNI interface, the routing and signaling in the optical layer, and finally, the setting up of IGP router adjacencies over the newly established IP link(s). In particular, such an approach requires several precautions in order to prevent "false-positive" alarms that could lead to network instabilities. Indeed, upon a router failure, the network should be quickly reconfigured to limit the impact of traffic disruption and/or QoS degradation due to congestion, but at the same time, it would be undesirable to trigger a complex set of recovery mechanisms involving several layers for a temporary router failure. Thus, the trade-off between fast recovery time and network stability is hard to determine. Note also that a dynamic multilayer recovery mechanism still requires some extra equipment capacity in the IP layer.

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 [14]. We consider a design with the restriction of shortest path routing, a seminal paper on such optimization is presented in [15]. Also, spare and working capacities are jointly optimized in the design (see [16] for further details).

VI. Case Studies

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

A. MTE Strategy as a Recovery Mechanism

MTE relies on multilayer routing to adapt to changing traffic demands. It requires fast signaling protocols, and one needs to take into account the limitations on algorithm complexity and the possibility of undesired effects such as network instability, QoS degradation during flow reroutes (e.g., loss, jitter), etc.

The MTE strategy described in this section will rely on IP-based costs only, since an overlay model is considered (in contrast to the integrated models in [11] and [13]). Optical routing is delegated to an optical routing-and-wavelength-assignment (RWA) algorithm (first-fit in this case) [17]. IP/MPLS traffic routes are calculated assuming a full mesh of LSPs in the logical layer; their path calculated over an artificial full mesh, which serves to represent the high flexibility in lightpath setup. The cost function (Fig. 13) will drive IP-layer performance optimization. It serves to attract traffic flows to IP links such that they are all moderately loaded, discouraging overloaded links [load above "high load threshold (HLT)"] as well as a large number of lightly loaded ones [load below "low load threshold (LLT)," which reduces bandwidth efficiency].

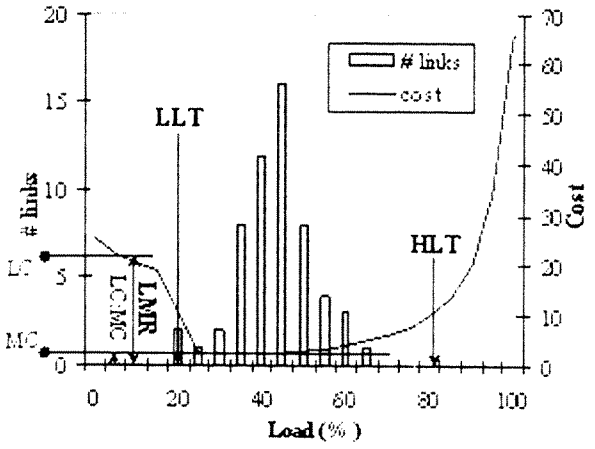

Fig. 13. IP-layer cost function and sample link load histogram.

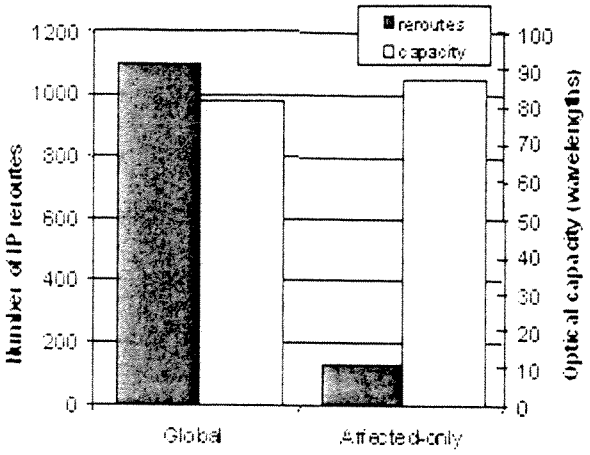

Fig. 14. Effect of MTE optimization scope on IP and optical performance.

The introduction of higher cost "LC," for links with a load below a certain threshold (compared to the cost "MC" for moderate loads), removes traffic flows from those links, thus lowering the total number of links in use. The ratio "LMR" between LC and MC is important because it determines how traffic can be diverted from a direct IP link in the artificial-mesh topology towards multihop routes; LMR corresponds with the typical maximum length in hops of such a multihop path.

The logical-topology configuration is a side effect of the artificial-mesh routing calculations: The remaining artificial-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 that all aspects of MTE are handled through a single path calculation—no separate mechanisms, some of them possibly offline as in [12], are necessary.

No backup path calculation is performed in the IP layer (although optical-layer recovery is likely to be present). This MTE strategy normally triggers proactively [18] 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 [19].

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 not related to the router failure. The live rerouting of a traffic flow requires signaling in order to update the GMPLS forwarding tables for its LSP, which possibly causes some QoS degradation. While acceptable for globally optimizing dynamic multilayer recovery schemes, it has real impact on network performance in the case of MTE-based recovery schemes. This is because the reconfiguration is performed online in this case, and the extent of 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 strategies 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. 14, 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 the number of IP reroutes during recovery indicating convergence speed, as well as 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 keep the route of unaffected flows, so some lightpaths may end up being used less efficiently.

B. Comparison of Static and Dynamic Multilayer Recovery Strategies

An offline tool was developed to design IP-over-ION logical topologies for certain scenarios: the failure-free scenario and single-IP-router-failure scenarios. The resulting total multilayer-network capacity requirements were then calculated. For example, the required capacity for each optical link is the maximum link capacity out of all scenarios (failure-free scenario and IP-router-failure scenarios). Thus, we combine all scenarios into a single "worst-case scenario" to derive a logical topology (and associated resources) that is able to cope with single IP router failures.

For a static multilayer recovery strategy where setup of lightpaths is unavailable (static OTN) upon demand, this dimensioning requires setup of all backup capacity (corresponding to the worst case scenario) in advance. However, for the dynamic case, we can rely on the GMPLS signaling and reduce the number of lightpaths established at any time (see Fig. 12), allowing reuse of IP–OXC interface cards, wavelengths, OXC ports, and line systems between the failure scenarios.

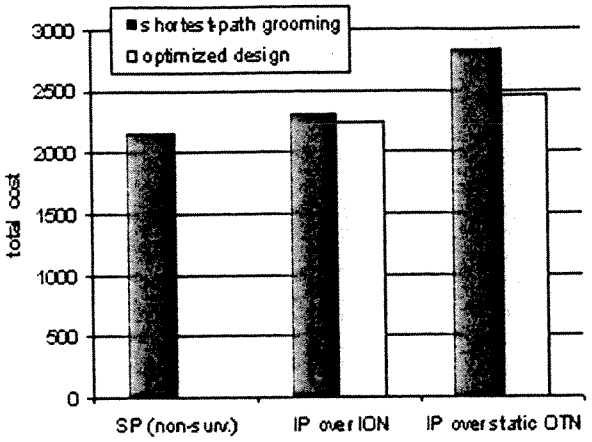

Fig. 15. Cost comparison of different strategies.

This leads to a significant cost reduction, quantified by a simulation study [20] based on several case studies. One of these studies concerns an OTN consisting of 12 nodes and 17 bidirectional links. It is a downsized version of a pan-European reference network topology [19]. The applied traffic demand is modeled by a symmetric matrix with a volume of 630.6 Gb/s, modeling pan-European traffic for the year 2002 based on population numbers and the work and home activities of users.

Fig. 15 gives the results of the performance for five different cases, including the cost of various resources (interfaces, ports, line-systems, etc.). As a benchmark (left bar on Fig. 15), a logical topology is designed based on the shortest path IP-layer routing and grooming when no survivability is taken into account. The same logical topology is applied to two survivable cases taking single IP router failures into account: IP over ION and IP over static OTN (second and fourth bar in Fig. 15). This leads to additional resources in order to reach a network solution with enough bandwidth for survivability.

Next to these base cases, two optimized logical-topology designs were also calculated for IP over ION and IP over static OTN (third and fifth bar in Fig. 15). The algorithm that optimizes the logical-topology design starts from a full demand mesh (with direct IP links for all nonzero demands) and deletes all IP links along edges whose removal improves a certain objective value (related to the total cost).

This comparison study reveals that IP rerouting over ION comes at a small extra cost when compared to the nonsurvivable case. The extra cost is considerably higher for IP rerouting over a static OTN. Additionally, we see that IP rerouting over a 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 clear disadvantage with respect to IP rerouting over ION.

VII. Conclusion

In multilayer IP-over-optical networks, various recovery mechanisms are at our disposal. This paper outlined the opportunities and challenges that are to be addressed. It was shown that recovery at multiple layers is necessary to reach a high availability. To coordinate these mechanisms at different layers, a sequential layer interworking approach with a hold-off timer or a recovery-token signal is crucial.

To provide multilayer recovery, a major distinction is based on the static or dynamic nature of the OTN. In case of a static logical topology, several alternatives exist with respect to spare-capacity provisioning: duplicated protection, logical spare unprotected, and common pool. The latter alternative turns out to be quite efficient in terms of bandwidth usage, albeit with additional complexity. On the other hand, the introduction of ASON or GMPLS functionality allows for dynamic logical-topology adaptations and opens new opportunities for recovery. When a failure occurs, new lightpaths can be set up or torn down to alleviate the failure repercussions. A case study on a pan-European network reveals that this additional flexibility can highly reduce the spare capacity needed for failure recovery. MTE techniques for coping with dynamic IP traffic patterns can be used in a natural way for a failure recovery as well. The MTE strategy was improved to reduce the performance impact due to traffic rerouting during failures and tested on a pan-European reference network.

References

[1] J.-P. Vasseur, M. Pickavet, and P. Demeester, *Network Recovery*. San Francisco, CA: Morgan Kaufmann, 2004.

[2] D. Papadimitriou *et al.* (2002, Jun.). *Shared Risk Link Groups Encoding and Processing*. [Online]. Available: www.ietf.org. Internet draft, work in progress.

[3] N. Wauters, G. Ocakoglu, K. Struyve, and P. F. Fonseca, "Survivability in a new pan-European carriers' carrier network based on WDM and SDH technology: Current implementation and future requirements," *IEEE Commun. Mag.*, vol. 6, no. 8, pp. 63–69, Aug. 1999.

[4] C. Guillemot, "VTHD French NGI initiative: IP and WDM interworking with WDM channel protection," presented at the IP over DWDM Conf., Paris, France, 2000.

[5] P. Demeester *et al.*, "Resilience in multi-layer networks," *IEEE Commun. Mag.*, vol. 37, no. 8, pp. 70–76, Aug. 1998.

[6] ITU-T Recommendation G.807/Y.1302. (2001, Jul.). *Requirements for automatic switched transport networks (ASTN)*, Geneva, Switzerland, ITU-T Standardization Org. [Online]. Available: www.itu.int

[7] A. Farrell, J.-P. Vasseur, and IETF Path Computation Element Workgroup. [Online]. Available: http://www.ietf.org/html.charters/pce-charter.html

[8] K.-I. Sato *et al.*, "GMPLS-based photonic multilayer router (Hikari Router) architecture: An overview of traffic engineering and signaling technology," feature topic on "Optical Switching", *IEEE Commun. Mag.*, vol. 40, no. 3, pp. 96–101, Mar. 2002.

[9] W. Alanqar *et al.* (2003, Dec.). *Requirements for Generalized MPLS (GMPLS) Routing for Automatically Switched Optical Network (ASON)*. [Online]. Available: www.ietf.org. Internet draft, work in progress.

[10] D. Papadimitriou *et al.* (2003, Nov.). *Requirements for Generalized MPLS (GMPLS) Signaling Usage and Extension for Automatically Switched Optical Network (ASON)*. [Online]. Available: www.ietf.org. Internet draft, work in progress.

[11] M. Kodialam *et al.*, "Integrated dynamic IP and wavelength routing in IP over WDM networks," in *Proc. IEEE Int. Conf. Computer Communications (INFOCOM)*, Anchorage, AK, 2001, pp. 358–366.

[12] P. Iovanna *et al.*, "A traffic engineering system for multilayer networks based on the GMPLS paradigm," *IEEE Network*, vol. 17, no. 2, pp. 28–37, Mar./Apr. 2003.

[13] T. Cinkler *et al.*, "Fairness issues of routing with grooming and shared protection," in *Proc. Optical Network Design and Modelling (ONDM)*, Ghent, Belgium, 2004, pp. 665–684.

[14] B. Ramamurthy *et al.*, "Design of virtual private networks (VPNs) over optical wavelength division multiplexed (WDM) networks," *Opt. Netw. Mag.*, vol. 3, no. 1, pp. 59–67, Jan./Feb. 2002.

[15] B. Fortz *et al.*, "Optimizing OSPF/IS—IS weights in a changing world," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 4, pp. 756–767, May 2002.

[16] R. R. Iraschko *et al.*, "Optimal capacity placement for path restoration in STM or ATM mesh survivable networks," *IEEE/ACM Trans. Netw.*, vol. 6, no. 3, pp. 325–336, Jun. 1998.
[17] B. Puype *et al.*, "Optical cost metrics in multilayer traffic engineering for IP-over-optical networks," in *Proc. Int. Conf. Transparent Optical Networks (ICTON)*, Wroclaw, Poland, 2004, vol. 1, pp. 75–80.
[18] ——, "Multilayer traffic engineering in data-centric optical networks, illustration of concepts and benefits," in *Proc. Optical Network Design and Modelling (ONDM)*, Budapest, Hungary, 2003, pp. 221–226.
[19] S. De Maesschalck *et al.*, "Pan-European optical transport networks: An availability based comparison," *Photonic Netw. Commun.*, vol. 5, no. 3, pp. 203–225, May 2003.
[20] A. Groebbens *et al.*, "Logical topology design for IP rerouting: ASONs versus static OTNs," *Photonic Netw. Commun.*, 2005, to be published.

Mario Pickavet (S'94–A'99–M'04) received the M.Sc. and Ph.D. degrees from Ghent University, Ghent, Belgium, in 1996 and 1999, respectively, all in electrical engineering with specialization in telecommunications.

Since 2000, he is a Professor at Ghent University, where he is teaching telecommunication networks and algorithm design. His current research interests are related to broadband communication networks [wavelength division multiplexing (WDM), IP, generalized multiprotocol label switching (GMPLS), optical packet switching (OPS), and optical burst switching (OBS)] and include design, long-term planning, and routing of core and access networks. His research interests also include operations research techniques that can be applied for routing and network design, and he is currently involved in the European Information Society Technologies (IST) projects "All-Optical Label Swapping Employing Optical Logic Gates in Network Nodes (LASAGNE)" and "Optical Networks: Towards Bandwidth Manageability and Cost Efficiency (e-Photon/ONe)." He has published over 100 international publications, both in journals (IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (JSAC), *IEEE Communications Magazine*, JOURNAL OF LIGHTWAVE TECHNOLOGY, *European Transactions on Telecommunications*, *Photonic Network Communication*, *Journal of Heuristics*) and in proceedings of conferences. He is the coauthor of the book *Network Recovery: Protection and Restoration of Optical, SONET–SDH, IP, and MPLS* (San Francisco, CA: Morgan Kaufmann, 2004).

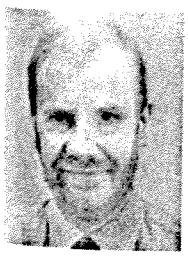

Piet Demeester (M'89–SM'98) received the Ph.D. degree from the Department of Information Technology (INTEC), Ghent University, Ghent, Belgium, in 1988.

In the same department, he became the Group Leader of the activities on metal organic vapour phase epitaxial growth for optoelectronic components. In 1992, he started a new research group on broadband communication networks. The research in this field resulted in more than 300 publications. He is currently a full-time Professor at Ghent University, where he is teaching courses in communication networks. He has also been teaching in different international courses. His current research interests include multilayer networks, quality of service (QoS) in IP networks, mobile networks, access networks, grid computing, distributed software, and network and service management and applications [supported by the Fund for Scientific Research-Flanders (FWO–Vlaanderen), the Bijzonder Onderzoeksfonds BOF) of Ghent University, the Institute for the Promotion of Innovation by Science (IWT), and the European Commission].

Dr. Demeester was and is a member of several program committees of international conferences, such as the International Conference on Computer Communications and Networks (ICCCN), *Optical Fiber Communication (OFC)* Conference, International Conference on Communications (ICC), and European Conference on Optical Communications (ECOC). He was the Chairman of the Design of Reliable Communication Networks (DRCN) in 1998. In 2001, he was the Chairman of the Technical Programme Committee of ECOC. He was the Guest Editor of three special issues of the *IEEE Communications Magazine*. He is also a member of the Editorial Board of the Journals *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 the Association for Computing Machinery (ACM) and KVIV.

Didier Colle (A'01–M'03) received the M.Sc. degree in electrotechnical engineering (communications) and the Ph.D. degree, in 1997 and 2002, respectively, all from the Ghent University, Ghent, Belgium. He was granted a postdoctoral scholarship from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen) in the period 2003–2004.

Since 1997, he has been working at the same university as a Researcher in the Department of Information Technology (INTEC). He is a member of the research group INTEC Broadband Communication Networks (IBCN) headed by Prof. P. Demeester. His research deals with the design and planning of communication networks. This work focuses on optical transport networks (OTNs), to support the next-generation Internet. He has actively been involved in several Information Society Technologies (IST) projects [Layers Interworking In Optical Networks (LION), Optical Technologies in Motion for the IST Programme (OPTIMIST), DAta and Voice Integration over Dense Wavelength Division Multiplexing (DAVID), Switching Technologies for Optically Labeled Signals (STOLAS), Next-Generation Optical Network for Broadband in Europe (NOBEL) and All-Optical Label Swapping Employing Optical Logic Gates in Network Nodes (LASAGNE)], in the COST-Action 266 and 291, and in the Information Technology for European Advancement (ITEA)/Institute for the Promotion of Innovation by Science (IWT) TBONES project. His work has been published in more than 100 scientific publications in international conferences and journals.

Dimitri Staessens received the M.Sc. degree in computer science from Ghent University, Ghent, Belgium, in 2004.

After a short period of researching numerical algorithms for solving Sturm-Liouville problems, he focused his research towards applications of graph theory for providing survivability in future generation optical networks. He has been active in several projects including TBONES and the Next-Generation Optical Network for Broadband in Europe (NOBEL).

Bart Puype received the M.Sc. degree in electrotechnical engineering from Ghent University, Ghent, Belgium, in 2002. In 2004, he received an Institute for the Promotion of Innovation by Science (IWT) Ph.D. scholarship.

Since 2002, he has been working as a Research Assistant at the Department of Information Technology of Ghent University. 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 Institute for the Promotion of Innovation by Science (IWT)/Information Technology for European Advancement (ITEA) TBONES and the European 6th framework program Information Society Technologies (IST) Next-Generation Optical Network for Broadband in Europe (NOBEL) projects.

Leen Depré received the M.Sc. degree in computer science from the Ghent University, Ghent, Belgium, in 2004. Her thesis "Design and Simulation of Replica Placement Algorithms Based on Congestion Detection" considered the distribution and duplication of the content in content distribution networks.

Since 2004, she has been working as a Research Assistant at the Department of Information Technology (INTEC) [INTEC Broadband Communication Networks (IBCN) research group] at Ghent University. Her research interests cover broadband communication networks. She is currently involved in the Information Society Technologies (IST) Next-Generation Optical Network for Broadband in Europe (NOBEL) project. She is involved in a study group that focuses on multidomain and multilayer survivability of IP and optical networks.

Ilse Lievens received the M.Sc. degree in electrical engineering (focusing in telecommunications) and the Ph.D. degree, in 1994 and 2000, respectively, all from the Ghent University, Ghent, Belgium. Her Ph.D. thesis "Use of Distributed Rerouting in Meshed ATM Networks" considered the design of rerouting algorithms for survivability in meshed ATM networks, focusing on distributed autonomously working techniques.

She then joined the Department of Information Technology at Ghent University in the Broadband Communication Networks Group (IBCN research group). 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., Information Society Technologies (IST) Layers Interworking In Optical Networks (LION), IST Next-Generation Optical Network for Broadband in Europe (NOBEL), Information Technology for European Advancement (ITEA) TBONES], and in national interuniversity projects [Institute for the Promotion of Innovation by Science (IWT) GBOU ONNA].

Dr. Lievens is author and coauthor of several publications in conference proceedings and journals [European Conference on Optical Communications (ECOC), IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (JSAC), *IEEE Communications Magazine*, *Photonic Network Communications* (PNC), etc.].