Resilience in Multilayer Networks

Piet Demeester and Michael Gryseels, University of Gent-IMEC
Achim Autenrieth, Technical University of Munich
Carlo Brianza, Italtel
Laura Castagna and Giulio Signorelli, SIRTI
Roberto Clemente and Mauro Ravera, CSELT
Andrzej Jajszczyk and Dariusz Janukowicz, ITTI
Kristof Van Doorselaere, Belgacom
Yohnosuke Harada, NTT

ABSTRACT The integration of different technologies such as ATM, SDH, and WDM in multilayer transport networks raises many questions regarding the coordination of the individual network layers. Especially in the area of network survivability, much can be gained by a better alignment of the healing actions taken by different network layers in case of outages. Survivability issues encountered in a multilayer environment include, among others: how to identify the original failure cause, how to appoint for each failure a layer responsible for its healing, how to let different layers interwork, and how to provide spare resources in an efficient way.

The resilience of telecommunication networks gets intense attention of all major players in the telecom world. The adoption of wavelength-division multiplexing (WDM), synchronous digital hierarchy (SDH), and asynchronous transfer mode (ATM) technologies in transport networks has resulted in an ever-increasing concentration of more traffic on fewer network elements. For instance, the failure of a 40-wavelength WDM system carrying 2.5 Gb/s SDH signals can affect up to 1,200,000 telephone calls. Consequently, outages cause more burdens than ever before, particularly to customers whose communications are of vital importance.

Network survivability, which is the ability of a network to recover traffic affected by failures, has gained a critical role in the design of telecommunications networks. A great deal of research has been done on the design of survivable network architectures and recovery schemes for the different transport technologies (see [1, 2] for surveys). Since the survivability techniques implemented today mainly address individual technologies, the study of solutions considering an integrated system is paramount for further development in this field. Some objectives for an integrated approach to multilayer survivability include:

- Avoiding contention between the different single-layer recovery schemes
- Promoting cooperation and sharing of spare capacity
- Increasing the overall availability that can be obtained for a certain investment budget
- Decreasing investment costs required to ensure a certain survivability target

This article presents results from the European research project PANEL,[1] which studied the coordination of recovery schemes present in different layers of a transport network. Note that this article does not include the circuit-switched layers; see [3].

A FRAMEWORK FOR MULTILAYER SURVIVABILITY STRATEGIES

Multilayer survivability starts from the viewpoint that a multitechnology network consists of a stack of single-layer networks. A client/server functional relationship holds between adjacent layers of the multilayer stack. Each lower network layer provides transport functionality to higher client layers. In each network layer with the necessary flexibility for rerouting, a single-layer recovery scheme may be deployed. The presence of recovery mechanisms in multiple layers of ATM-over-SDH-over-WDM networks has several drivers:

- Recovery schemes residing in lower layers (e.g., SDH) often enable more effective recovery from burdensome failures like cable cuts, while failures of higher-layer equipment (e.g., an ATM switch) requires additional resilience in the higher layer.
- The need for differentiation of service reliability (e.g., different grades for distinct customers or products) may result in the installation of recovery schemes closer to the network layer where traffic is actually injected in the network.
- The natural evolution of telecommunications networks may result in adding new survivable layers to the existing ones (e.g., optical layer survivability).

Past research [1, 2, 4, 5] recognized the importance of providing synergy between the healing actions of the different network layers. PANEL worked out a framework for multilayer survivability (Fig. 1), which compiles directions that can be taken by telecommunications network providers. The components of the framework can be classified in options pertaining to the single-layer recovery systems and options pertaining to multilayer interactions. The recovery options of individual technologies have been studied extensively in the scope of single-layer survivability [1, 2]. In the next sections we illustrate some of the multilayer options of our framework in the context of ATM-over-SDH transport networks. Nevertheless, many of the described survivability concepts are generic and applicable to other types of multilayer networks as well.

MULTILAYER RECOVERY APPROACHES

To define a multilayer recovery approach, the following question is raised: "For each failure, which network layers are responsible for its recovery?" We will illustrate some

[1] *Protection Across Network Layers, http://www.intec.rug.ac.be/Research/Projects/PANEL*

 0163-6804/99/$10.00 © 1999 IEEE

approaches with the example in Fig. 2 representing a two-layer network, where the SDH layer supports native SDH paths (e.g., A-B-C) and a client ATM layer. For instance, ATM path a-b-c (set up between ATM nodes a and c) utilizes ATM link capacity supported by SDH paths A-B and B-C.

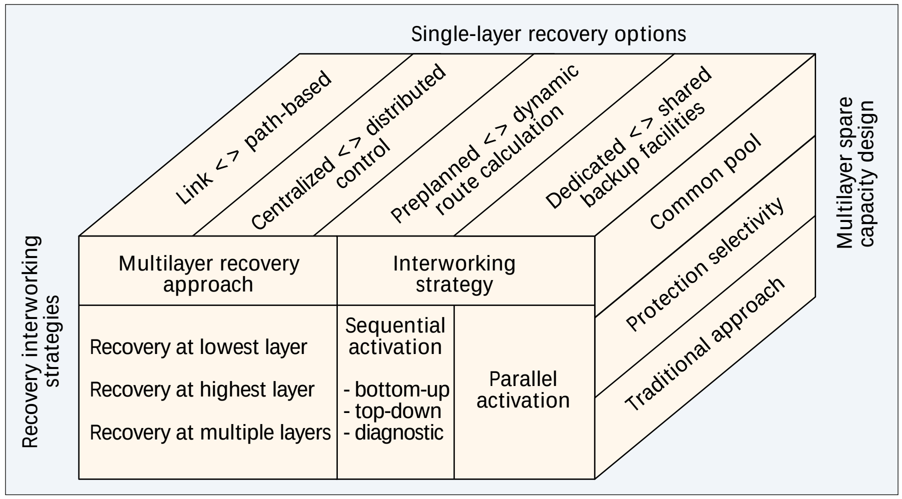

■ **Figure 1.** *The PANEL framework for multilayer survivability.*

Recovery at the Lowest Layer — A first recovery approach, denoted as *recovery at the lowest layer*, is to recover the affected services in the lowest possible layer. As such, the survivability is provided as close as possible to the *origin layer of the failure*. For example, cable cut A-B (Fig. 2.a) affecting ATM path a-c is resolved at the SDH layer. When ATM node b fails, the ATM recovery scheme needs to be activated to restore affected traffic. Because of the coarser switching granularity of lower layers, this recovery approach is simpler in the number of affected paths to reroute (compared to recovery at the highest layer, see below).

Due to the coexistence of multiple schemes, some interworking functionality needs to be provided to correctly assign and coordinate the recovery responsibilities of each recovery scheme. The break of cable A-B will result in a loss of the signal at SDH node B, but at the same time ATM node c will not receive valid ATM cells anymore. After 4 ms, the ATM node goes into an alarm state due to a loss of cell delineation. In the meantime, it may also have received alarms from the SDH layer, indicating that the fault cause is in the SDH layer. Indeed, in the event of an SDH layer failure, subsequent alarm messages will be passed through each upper layer after they are generated [6].

However, it may happen that the alarm signal from the SDH layer does not arrive fast enough (e.g., intermediate SDH network elements delay the alarm signal propagation, and the ATM layer defect may be detected first). As a result, the ATM layer would falsely conclude that the failure happened in its own layer and may incorrectly trigger the ATM layer recovery mechanism. However, the network resources are dimensioned to resolve SDH layer failures "at the lowest layer." This may create a situation for ATM layer resource competition, leading to network congestion and other unwanted behavior of the network. Recovery interworking is required to overcome this problem.

With recovery at the lowest layer, every survivable layer reserves some resources for the rerouting of affected paths. A single-layer recovery mechanism cannot acquire resources in other layers in real time without breaching the independent operation of different layers. This means that spare capacity in one layer needs to be supported at all times by lower layers. In our example, ATM layer capacity — provided by SDH paths — has to be reserved to reroute virtual paths (VPs) around a failing ATM node. In addition, some SDH layer capacity has to be reserved for SDH layer rerouting. The result is that each survivable layer decreases the actual utilization of the capacity, and the lowest layer may thus be poorly used with this recovery approach. This problem is further discussed later.

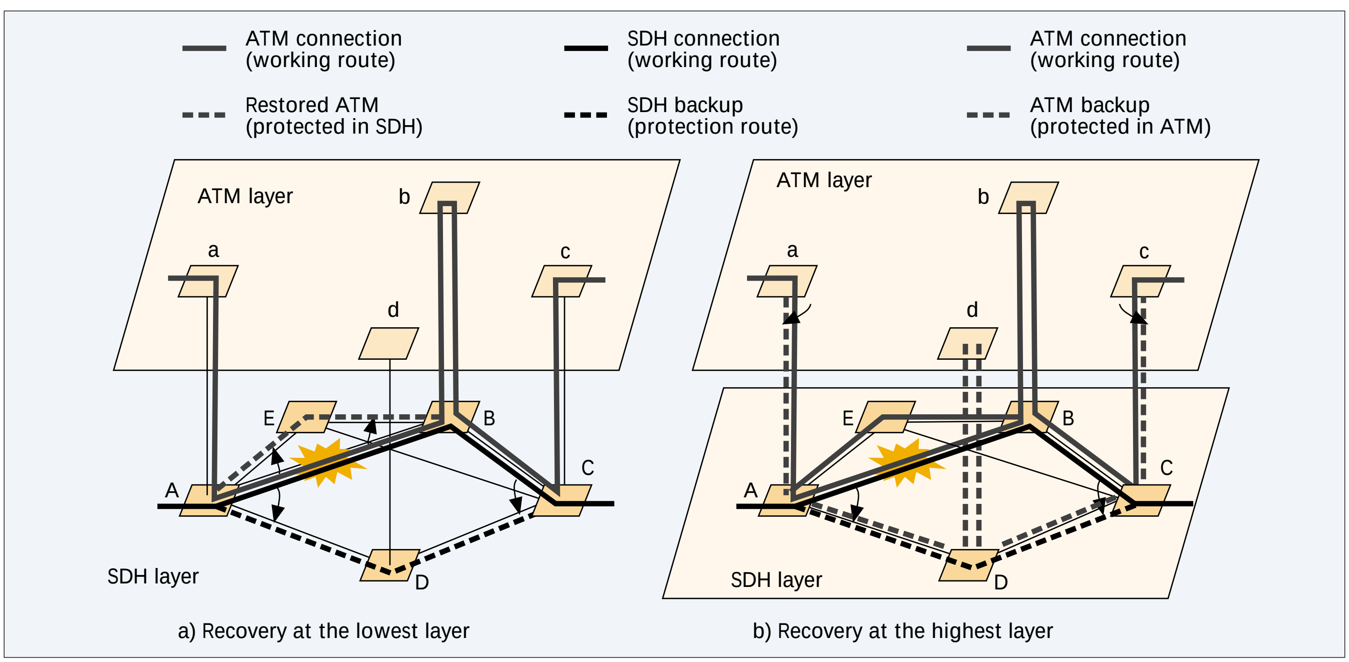

■ **Figure 2.** *Multilayer recovery approaches.*

Recovery at the Highest Layer — A second multilayer recovery approach, denoted *recovery at the highest layer*, consists of recovering disrupted traffic in a network layer closer to the origin of the traffic. A higher-layer recovery scheme can resolve failures happening in layers below. The example in Fig. 2b illustrates that an affected SDH path carrying SDH services is recovered by SDH network elements, while an affected ATM path (carried over different SDH paths) is recovered in the ATM layer. Providing resilience at the highest possible layer has the following advantages:

- A transport network often carries several service classes with different reliability requirements. It is generally easier to provide multiple reliability grades when the survivability schemes reside in higher layers.
- Since a single recovery scheme suffices to protect the service against failures occurring in every (lower) layer, the implementation complexity of interworking between different schemes can be avoided. In this case, interworking merely involves activating the responsible recovery scheme as fast as possible at the appropriate network layer.

On the other hand, the finer switching granularity of the higher layers complicates rerouting in the event of lower-layer failures, because many entities are then affected at the same time. For instance, a 2.5 Gb/s SDH cable break may affect many thousands of VPs in the ATM layer, each of which would have to be rerouted individually in this approach. This would probably slow down the recovery process. Second, when the server layer paths supporting the higher network layer are rather long, the recovery process may involve reconfiguration in network elements far away from the original failure cause. Most probably, special precautions will have to be taken to adapt the survivability scheme of the higher layer (to resolve the lower-layer failure scenarios). For example, with restoration at the ATM layer it is envisaged to assemble VPs sharing physical routes into VP groups as a technique to reduce the recovery efforts in case of physical failures.

Recovery Interworking/Escalation — In multilayer networks, several recovery schemes may be involved to resolve certain failures. An *interworking strategy* (also called *escalation strategy* [5]) consists of a set of rules describing when to start and stop, and how to coordinate, the activities of the different recovery schemes. Two options were identified concerning the activation: starting two or more recovery schemes in parallel or starting them sequentially.

The parallel strategy is fast and requires no communication or coordination between schemes. The sequential strategy may result in longer overall recovery times than parallel activation but is generally easier to control. A sequential strategy determines the order of activation of the schemes and coordinates the schemes, ensuring that they are activated at an appropriate moment. Two further questions need to be answered in case of *sequential interworking*: at which layer to start the recovery and when to "escalate to the next layer."

The question of where to start results in three approaches: *bottom-up*, *top-down*, and *in-between* (diagnostic). Bottom-up interworking starts at the network layer closest to the failure, ensuring very quick activation of the recovery mechanism. The upward escalation takes place upon expiration of a hold-off timer, reception of a recovery token [7, 8], or intervention of the management system. Top-down interworking always starts at the highest-layer network and escalates downward. For example, a fast protection scheme in the highest layer can be combined with a slower but more cost-effective restoration scheme in a lower layer. The latter scheme is used to restore the higher network layer to its original state, enabling it to cope with certain double failure scenarios. The decision to start at a specific layer depends on diagnostic interworking, among others, on received alarms and gathered survivability statistics.

A Quantitative Comparison of Recovery Approaches

In order to compare recovery approaches in a quantitative way, the PANEL project carried out recovery simulation and planning case studies on a variety of ATM-over-SDH network and traffic models. A comparison of multilayer recovery approaches is dependent also on the single-layer recovery options. Since our aim is to illustrate the multilayer aspect of the recovery approaches, the following case studies assume similar mechanisms to be deployed in the individual layers. We concentrate on path protection schemes (i.e., end-to-end recovery with dedicated backup paths) in both layers. Such schemes are often used for the resilience of premium-level traffic. The considered failure scenario included single cable cuts, single SDH node failures, and single ATM node failures.

Recovery Performance Comparison — Figure 3a compares the performance of *recovery at the lowest layer* with *hold-off time* interworking and the *recovery at the highest layer*

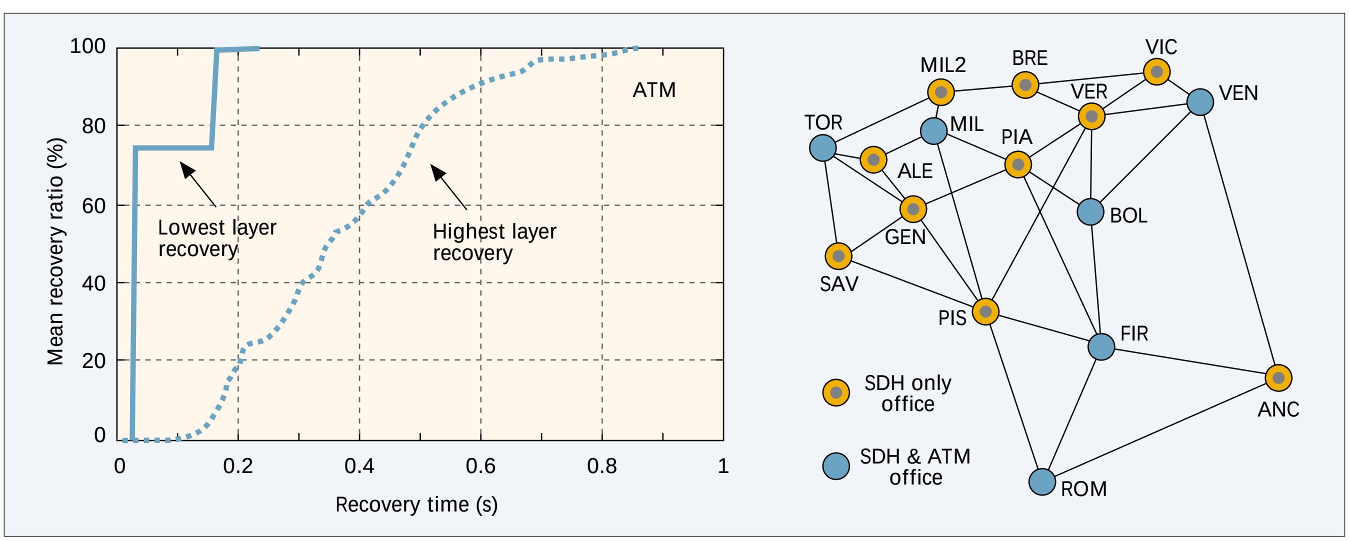

■ **Figure 3.** *a) Performance comparison; b) a 16-node example network.*

approach for a sample network with 16 SDH nodes, six ATM nodes, and 36 links (Fig. 3b). The graphs show the average recovery ratio of ATM layer connections for all complete site failures (measured at the ATM layer). Such performance values were obtained by simulating the recovery procedure for every site failure one after the other using a distributed, event-oriented multilayer simulation environment developed within PANEL [8]. The performance data of the simulation runs are collected and the mean recovery ratio is calculated from all results.

The most noticeable characteristic of the graphs is the two-step curve in *recovery at the lowest layer*. The majority of affected ATM connections are recovered by the SDH protection, which takes up to 50 ms and reaches an average recovery ratio of about 78 percent for the given network. ATM paths transiting the failed ATM and SDH site must be recovered at the ATM layer, since the SDH path supporting them is terminated at the failed node. Such multihop ATM paths are recovered by the ATM protection scheme. This scheme is triggered after expiration of the holdoff timer (100 ms). It needs 30 ms to setup the (preconfigured) translation table. Additionally, a few milliseconds signaling time are needed for each failed path due to the processing time of the recovery protocol. With recovery at the lowest layer, all ATM services were recovered after approximately 250 ms.

Because fast recovery of a single SDH connection recovers a very large number of ATM paths jointly, the performance of the recovery at the lowest layer interworking strategy is generally better than that of recovery at the highest layer. Single cable cuts, which are the most probable failures in a network, can be resolved completely in the SDH layer. The ATM layer only protects multihop ATM paths against transit node failures.

In *recovery at the highest layer*, all affected ATM paths are recovered in the ATM layer. No holdoff time is needed now, but the recovery is slower since a much higher number of ATM paths have to be recovered individually. The main influence on the ATM recovery ratio is not the number of nodes, but the number of affected ATM paths. In the example above with low ATM demand, all affected services were recovered after 900 ms. During simulations with more demands, this value increased to 3.5 s.

***Cost Comparison*[2] —** The recovery cost in the different cases was evaluated by assessing the amount of SDH and ATM equipment and fiber systems required by each recovery approach. The amount of resources was determined through a network dimensioning process, using network planning tools available within the Consortium [6]. The network cost originates from ATM equipment (backbone ATM DXC VP and its STM-1 interface cards), SDH equipment (DXC 4/4, STM-1 interfaces, 2.5 Gb/s line systems), and carrier cost, which includes the cost of fiber and regenerators. Actual equipment costs obtained from a survey of different ATM and SDH equipment manufactures were used [6]. All costs in this article are expressed in a normalized unit, corresponding to the cost of one electrical STM-1 port on an SDH line system.

In Table 1 (and Table 2), the investment costs related to the ATM demands are listed for recovery at the highest layer (and lowest layer, respectively) for a set of network and traffic scenarios (1)–(4). In all cases, recovery at the highest layer requires a larger investment budget (up to 20 percent more) than recovery at the lowest layer. There are several explanations for this outcome. First of all, with recovery at the lowest layer, only multihop VPs require additional ATM layer protection (and subsequent ATM layer spare capacity). This results in a lower ATM spare capacity requirement. Second, the approaches differ in the constraints that must be satisfied in finding working and backup VP routes. With recovery at the highest layer, the main and backup routes of a VP must be disjoint in both physical links and nodes. With recovery at the lowest layer, the only requirement is to have no transit ATM node in common. Hence, the recovery at the lowest layer approach leads to looser VP routing constraints and more room for optimization.

	(1)	(2)	(3)	(4)
ATM equipment cost	2357	6332	5817	17,553
SDH equipment cost	2075	5690	5768	16,488
Carrier cost	1385	3456	4713	13,031
ATM service cost	5817	15,478	16,298	47,072

■ **Table 1.** *Investment costs for recovery at the highest layer.*

	(1)	(2)	(3)	(4)
ATM equipment cost	1411	3294	3044	8000
SDH equipment cost	2193	5749	6373	17,792
Carrier cost	1304	3570	4672	13,301
ATM service cost	4908	12,613	14,090	39,093

■ **Table 2.** *Investment costs for recovery at the lowest layer.*

Furthermore, the SDH network layer supports both working and spare ATM capacity. With recovery at the highest layer, no SDH spare resources have to be provided for ATM demands, but the extra deployment of the ATM spare capacity results in higher occupation of SDH facilities by the ATM layer. In fact, the total SDH capacity required for the ATM demands is about the same for both recovery approaches. With recovery at the lowest layer, part of the spare capacity for ATM service resilience is only put in the SDH network layer, and only has an impact on the needed equipment (and installation costs) of one network layer instead of two. Concluding, it seems cheaper to resolve physical and SDH failures using recovery at the SDH layer instead of dedicating expensive ATM capacity for this task.

COMMON POOL SURVIVABILITY

Since transmission and switching resources are very expensive, optimization of spare resources is very important for operators (and indirectly for customers as well), because it results in containment of the costs and consequently also of the service tariffs. As explained before, multilayer survivability implies providing multiple spare capacity pools, each dedicated to a particular network layer. Since capacity of each survivable layer is carried by its server layers, this results in a reservation of resources in all layers below. With traditional capacity planning, this results in poor utilization of the lowest layer. Moreover, the server layer potentially also protects the spare capacity of the client layer, which does not always result in an equal increase of service availability. Such redundant protection can be avoided by supporting working and spare client layer capacity through different server paths, and treating them differently in the server layer. For example, when protection selectivity is available in the server layer, paths carrying client layer spare capacity can be left unprotected.

Nevertheless, the server layer then still dedicates some

[2] *More details about this planning case study can be found in [6].*

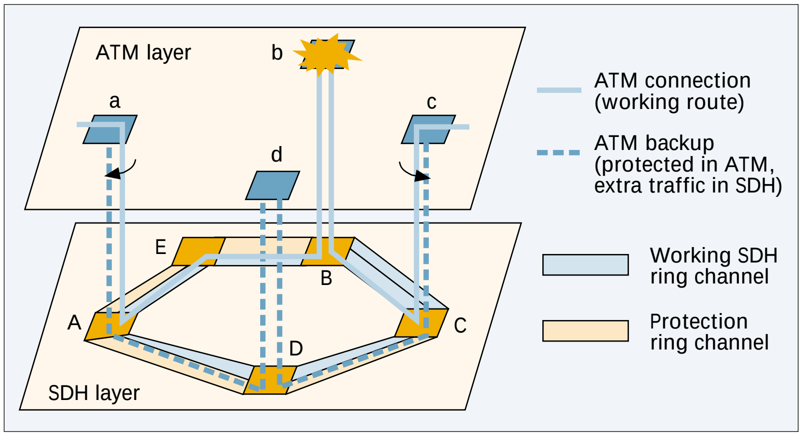

■ **Figure 4.** *Common pool survivability — an ATM failure scenario.*

In summary, the spare capacity required for ATM layer resilience is treated as *unprotected preemptible traffic* in the SDH network layer. As such, the SDH spare resources are reused for extra resilience at the ATM layer. A traditional approach deploys two spare capacity pools, each dedicated to a specific layer (i.e., ATM or SDH). The difference in reliability is that with a common pool it is no longer possible to provide resilience against common failure of an ATM node and a cable cut. Since such failure scenarios happen rarely, the effect on service availability is very small. In contrast, the savings in SDH equipment achieved with a common pool approach can be quite substantial [10].

resources to carrying the client layer spare capacity. The utilization of the server layer resources can be further improved by sharing spare capacity across network layers as described in [10]. The basic idea of *common pool survivability* is to treat the spare capacity of the client layer as extra traffic (i.e., carried on unprotected preemptible paths) in the server layer. The spare capacity at the server layer is planned to protect the server layer paths carrying the actual traffic. With common pool survivability, the server layer spare capacity is reused by a higher-layer recovery scheme. Little or no additional server layer resources are thus required to support the client spare capacity, which is now carried in the reserve capacity provisioned for server layer survivability. The utilization of network resources is better than with any other approach.

The concept of common pool survivability is illustrated with a simple network example shown in Figs. 4 and 5, considering a network with a meshed ATM layer on top of SDH bidirectional rings. Although the figure considers, for the sake of readability, a single SDH ring, it has also been verified for a network consisting of interconnected rings. In fact, it should be stressed that the common pool survivability concept is generic and can be applied to other types of multilayer networks as well.

A working ATM VP, a-c, routed via ATM node b, is transported via the working SDH paths A-E-B and B-C, as shown in Fig. 4. To protect this working VP against transit ATM node breakdowns, a backup ATM VP a-c routed via ATM node d is preestablished. This backup VP is transported via the SDH paths A-D and D-C. However, these SDH paths are not routed along working SDH ring channels but as extra traffic along the protection channels of the ring. If the ATM DXC b fails (Fig. 4), the working ATM VP a-c is rerouted via its preassigned backup VP. This involves a simple reservation and reconfiguration process at the ATM network layer only. After recovery completion, ATM VP a-c is transported over SDH protection channels without requiring a reconfiguration of any of the SDH add-drop multiplexers (ADMs).

On the other hand, when a cable cut (Fig. 5) occurs, the SDH ADMs access the protection channels to reroute the disrupted working SDH paths. Hence, the SDH paths supporting the ATM spare capacity may be preempted in order to free resources for the ring protection protocol. This implies the temporal unavailability of some part of the ATM spare capacity until the failure is repaired. There is no need for activation of the ATM recovery, since the SDH protection restores all affected VPs.

Multipart Survivability in SDH-over-WDM Networks

Currently, WDM technology is mainly used on a point-to-point basis to increase the available transmission capacity on existing fibers. The throughput growth possible with WDM technology increases the impact of failures, since many more connections fail simultaneously. In the future, when optical ADMs and cross-connects become available, there will be opportunities for full optical paths between offices far from each other. A cable cut then results in a complex failure scenario in the SDH layer. If recovery is left to SDH, a single cable cut may require many switching actions at different locations.

Providing survivability at the WDM layers becomes inherently attractive as network throughput increases [8, 11]. Optical recovery schemes recover multiple affected SDH links at once and entail less coordination than SDH recovery. In the last few years, the research on WDM recovery schemes has therefore intensified, and optical protection switching has become a topic of interest in the standardization bodies [9]. Because of the functional similarity of SDH and WDM networks, the survivable architectures in the WDM layer are likely to operate in a similar way to the standardized SDH architectures [9].

In reality, optical networking (and optical recovery schemes) will not be deployed at once in all parts of the network. Such WDM islands may then be interconnected by pure SDH transmission links (or links with point-to-point WDM). The end-to-end survivability of an SDH connection thus cannot be ensured by WDM protection only. In addition, it is

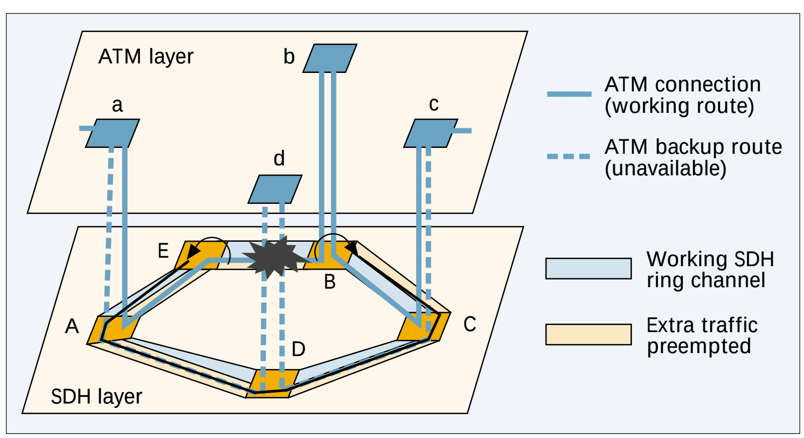

■ **Figure 5.** *Common pool survivability — an SDH failure scenario.*

very difficult to detect certain failures (e.g., those causing slow signal degradation) at the WDM layer. Additional SDH layer survivability will thus remain necessary in the network.

Based on the above rationale, several protection strategies can be proposed. The first option is to have a *protected SDH layer over an unprotected WDM layer*. For instance, in the current networks where WDM is used for point-to-point systems in well-defined areas, it is reasonable to leave the protection to SDH, since protection mechanisms and management systems of the SDH layer are very mature.

A second option is to have *an unprotected SDH over a protected WDM layer*. Whenever WDM is not yet deployed in all parts of the transmission network, this option uses WDM protection whenever possible and SDH protection in other parts of the network. As such, this strategy partitions the network into different survivable subnetworks, and failures are resolved within their subnetwork. Although such partitioning enables survival of failures happening in different subnetworks, it requires special precautions at the gateways, which can be expensive to implement. Also, SDH layer failures within the WDM islands are not covered.

These latter issues are not present with a *protected SDH layer supported by a protected WDM layer*. However, if a combination of the two protection schemes is planned without any precautions, the spare capacity of the SDH layer results in a higher amount of wavelengths. Moreover, these wavelengths (partly carrying spare SDH capacity) are protected in the WDM layer. This option may be too costly since WDM layer protection of SDH spare capacity is not always useful. Redundant protection can be avoided when working and spare SDH capacities are supported by separate wavelengths, using protection selectivity at the WDM layer. Nevertheless, supporting SDH spare capacity by unprotected WDM wavelengths still results in poorer utilization of the WDM layer than in the other options. A common pool configuration, where SDH spare capacity is supported on preemptible wavelengths in the WDM layer, obtains optimal utilization.

PANEL GUIDELINES

Recovery at the highest layer is the recommended approach when:
- Multiple reliability grades need to be provided with fine granularity
- Recovery interworking can not be implemented; (e.g., because of equipment limitations or when client and server network layer are owned by different operators)
- The survivability schemes in the highest layer are more mature than in the lowest layer

Recovery at the lowest layer is the recommended approach when:
- The number of entities to recover has to be limited/reduced
- When the lowest layer supports multiple client layers and it is appropriate to provide survivability to all services in a homogeneous way
- The survivability schemes in the lowest layer are more mature than in the highest layer
- When it is difficult to ensure the physical diversity of working and backup paths in the higher layers (e.g., when client and server layer are provided by different operators, the server network operator may not want to diffuse information on his network topology)

Recovery at the lowest layer may require explicit coordination between recovery schemes; the hold-off time based interworking mechanism is the best compromise between recovery performances and implementation complexity.

With recovery at the lowest layer, multiple spare capacity pools are present in the network for resilience of higher layer traffic. When higher layer spare capacity requirements are relatively large, the actual utilization of the lower layers can become poor. In that case, using unprotected or preemptible server layer paths (i.e., common pool) to carry the client spare capacity is recommended to alleviate redundant protection and remain cost-effective.

In SDH-over-WDM networks, where WDM recovery schemes are only available in parts of the network (evolutionary scenario):
- The disadvantage of leaving the protection completely in the SDH layer is that it is difficult to exploit the full potentials of optical networking
- The disadvantage of shifting the protection completely to the WDM layer (without any SDH protection in WDM parts) is that the SDH paths are protected in a segmented way
- SDH protection over WDM protection offers the highest degree of reliability, but a cost-effective implementation may only be achieved if working and spare SDH capacity are treated differently in the WDM layer (protection selectivity or better common pool).

CONCLUSIONS

This article presents outcomes of the PANEL research project, which studied survivability in multilayer transport networks. The major challenges for a telco in a multilayer environment are to determine which layer(s) is responsible for each failure (recovery approach), to coordinate single-layer recovery actions (escalation strategy), and to plan multiple spare capacity pools in an appropriate way. In order to provide overall service resilience in a cost-effective way, service providers can determine a suitable strategy based on the PANEL guidelines (see box).

ACKNOWLEDGMENT

This work has been supported by the European Commission through the ACTS project AC205/PANEL and by the Flemish Government through IWT project ITA/950214/INTEC.

REFERENCES

[1] T.-H. Wu, "Emerging Technologies for Fiber Network Survivability," *IEEE Commun. Mag.*, vol. 33, no. 2, 1995, pp. 58–74.
[2] ANSI T1, "A Technical Report on Network Survivability Performance, rep. no. 24, Oct. 1993.
[3] K. R. Krishnan, R. D. Doverspike and C. D. Pack, "Improved Survivability with multi-layer dynamic routing," *IEEE Commun. Mag.*, July 1995, pp. 62–68.
[4] D. Johnson, "Survivability Strategies for Broadband Networks," *Proc. IEEE GLOBECOM '96*, London, U.K., 1996.
[5] L. Nederlof *et al.*, "End-to-end Survivable Broadband Networks," *IEEE Commun. Mag.*, vol. 33, no. 9, 1995, pp. 63–70.
[6] M. Gryseels *et al.*, "Optimal Design for Service Resilience in ATM on SDH Backbone Networks," *Proc. IEEE ATM '98 Wksp.*, Fairfax, VA, May 1998.
[7] J. Manchester, "Fault Detection and Propagation in Transport Networks," *Proc. 1st Int'l. Wksp. Design of Reliable Commun. Networks*, Brugge, Belgium, 1998.
[8] A. Autenrieth et al., "Simulation and Evaluation of Multilayer Broadband Networks," *Proc. 1st Int'l. Wksp. Design of Reliable Commun. Networks*, Brugge, Belgium, 1998.
[9] A. McGuire and P. Bonenfant, "Standards: The Blueprints for Optical Networking," *IEEE Commun. Mag*. vol. 36, 1998, pp. 68–78.
[10] M. Gryseels *et al.*, "Common Pool Survivability for ATM over SDH Ring Networks," *Proc. 8th Int'l. Symp. Network Planning*, Sorrento, Italy, Oct. 1998.
[11] O. Gerstel, "Opportunities for Optical Protection and Restoration," *Proc. OFC '98*, San Jose, CA, 1998.

BIOGRAPHIES

PIET DEMEESTER (piet.demeester@intec.rug.ac.be) received his PhD degree from the University of Gent at the Department of Information Technology (INTEC) in 1988. He worked for the Interuniversity Micro Electronics Center (IMEC) and became professor at the University of Gent in 1993 where he is currently teaching telecommunication networks and where he is responsible for the research on communication networks. He is senior member of IEEE and he was or is member of several technical program committees. He

was involved in about 15 European ESPRIT, RACE and ACTS projects. He was co-editor of the special issue of the IEEE Communications Magazine: "Optical Networks Research in Europe, 1997." He was chairman of the First International Workshop on the design of Reliable Communication Networks in 1998 (DRCN98) and he is member of the editorial board of the journals: "OPTICAL NETWORKS MAGAZINE" and "JOURNAL PHOTONIC NETWORK COMMUNICATIONS." He published over 300 publications in the field of optoelectronics and broadband networks. His current interests are related to broadband communication networks (IP, ATM, SDH, WDM, access) and include network planning, network and service management, telecom software, internetworking, etc.

MICHAEL GRYSEELS (michael.gryseels@intec.rug.ac.be) received an M.Sc. degree in electrical engineering in 1996 from the University of Gent, Belgium. Since 1996 he has been working as a researcher for the Fund for Scientific Research of Flanders (F.W.O.) at the Department of Information Technology (INTEC), University of Gent, where he is studying multilayer networks. Currently, he is a visiting researcher with Telstra Corporation, Australia. His main research interests include modeling, planning, and evaluation of multitechnology network architectures.

ACHIM AUTENRIETH (achim@lkn.e-technik.tu-muenchen.de) received his Dipl.-Ing. degree at the Munich University of Technology, Germany, in 1996. Since then he has been a member of the research staff at the Institute of Communication Networks at the Munich University of Technology, where he was responsible for the ACTS Project PANEL from 1996 to 1998. He is currently working on his Ph.D. in the area of multitechnology resilience, including IP.

CARLO BRIANZA (carlo.brianza@italtel.it) received his Dr. Eng. degree in electronics at Politecnico di Milano and a Master's in information technology at the Cefriel research institute in Milan in 1993. He joined Italtel S.p.A. Central R&D in 1994 and was responsible for Italtel participation in both the PANEL and MISA ACTS projects. At the end of 1998 he joined the Operations and Support Systems division of INFOSTRADA S.p.A. His working experience encompasses all TMN layers.

LAURA CASTAGNA (L.Castagna@sirti.it) received a doctorate degree in electronic engineering from Politecnico di Milano in 1997. She joined the Sirti R&D Dept. in 1996 as a stager to study new methodologies to plan an ATM multiservice network. Since 1997 she has worked in the R&D Department of Sirti;she is engaged in planning ATM, SDH, and WDM networks, in interworking between IP and WDM, and as a delegate to ATM-Forum and ITU-T SG 12/13/15.

GIULIO SIGNORELLI (G.Signorelli@sirti.it) graduated in electronic engineering, specialization telecommunications, from University of Pavia, Italy, in 1992. He worked in the electrooptics laboratories of this university. In 1995 he joined the R&D Division of Sirti. He was involved in the transport network area dealing with SDH and WDM network planning. He participated in Telecom Italia SDH planning projects and several turnkey network projects for other operators. He is a member of ITU and ETSI study groups involved in optical networking. He has recently jointed Infostrada, the new Italian telco operator, where he is involved in the SDH and WDM network startup phase.

ROBERTO CLEMENTE (roberto.clemente@cselt.it) received his degree in electronic engineering from Politecnico of Torino in 1994. During the next year he joined CSELT, the corporate telecommunication research center of Telecom Italia, and started working in the Network Planning and Management Division. His interests are in the dimensioning of transport SDH/WDM networks, addressing both the long distance and regional network shares. He was involved in several EURESCOM and ACTS projects and has been co-author of papers regarding transport network planning and single- and multilayer recovery schemes in several conferences for the last three years.

MAURO RAVERA (mauro.ravera@cselt.it) received a degree in electronic engineering from the Politecnico of Torino in 1988; during the same year he joined CSELT, the corporate telecommunications research center of the Telecom Italia group. Since 1992 he has worked in the Network Planning and Management Department, where he has been responsible for a project which contributed to the fundamental plans of the access and transport network of Telecom Italia.

ANDRZEJ JAJSZCZYK [F'99] (Andrzej.Jajszczyk@itti.com.pl) received his M.S. degree in electrical engineering in 1974, and his Ph.D. degree in communications in 1979. He is a professor at the University of Technology and Agriculture, Bydgoszcz, Poland. His research interests include high-speed networking as well as network and service management. He has published over 150 papers and six books, and holds 19 patents. He was an editor of *IEEE Transactions on Communications* and founding editor of *IEEE Global Communications Newsletter*. He is Editor-in-Chief of *IEEE Communications Magazine*. He has organized several symposia, and served on numerous program and organizing committees.

DARIUSZ JANUKOWICZ (janukow@unitel.com.pl) finished his Master's studies at the University of Technologies in 1997. He worked on management of communication networks at the Institute of Communication and Information Technologies. He was involved in the PANEL project in the management-related part. Since 1999 he is involved in the development of Internet access networks and switching management at the Telecommunication Studies and Design Bureau "UNITEL" Ltd., Poznan, Poland.

KRISTOF VAN DOORSELAERE (kristof.van.doorselaere@is.belgacom.be) graduated as an engineer from the University of Gent in 1996. In that year he joined the research group of Prof. Piet Demeester at the University of Gent and studied there the fault management and survivability in WDM/SDH/ATM transport networks. In 1998 he left for Belgacom and is working there on the long-term leased line architecture.

YOHNOSUKE HARADA (harada@netscrg.tnl.ntt.co.jp) graduated from the Kyoto University Faculty of Engineering, and his major is on automatic control. After finishing his Master's degree, he entered NTT Electrical Communication Laboratories. He has been working on teletraffic study, network planning algorithms, network planning tools, dynamic routing algorithms, development of network operation systems, security architecture, ATM standardization, and overseas ATM application trials with AT&T, DT, and KDD. He is currently participating in multimedia applications development in social science.