

Joint Quantification of Resilience and Quality of Service

János Tapolcai¹, Piotr Cholda², Tibor Cinkler¹, Krzysztof Wajda²,
Andrzej Jajszczyk², Dominique Verchere³

¹Budapest Univ. of Technology and Economics, Hungary;

²AGH Univ. of Science and Technology, Kraków, Poland, ³Alcatel R&I, France

E-mail: {tapolcai,cinkler}@tmit.bme.hu, {cholda,wajda,jajszczyk}@kt.agh.edu.pl,
dominique.verchere@alcatel.com

Abstract – In this paper a new concept called **Quality of Resilience (QoR)** is presented, which is based on the separation of **Quality of Service (QoS)** parameters into short-term quality factors, called **availability parameters**, and long-term quality parameters called **QoR parameters**. After dividing the time into intervals with length Δt , for each such an interval the service is considered available, if the user is satisfied. Simultaneously, the long term characteristic of the service is derived from the service downtime distribution. The choice of proper length of the Δt time periods is the key issue of the applicability of the model. With the downtime histograms the asymptotic of the service can be illustrated both at transport and service layer. Since the resilience mechanism of the network is reflected in the transport layer downtime histograms, this new characterization of the QoS helps to understand impact of different recovery schemes on the next generation services.

Index Terms – protection, quality of service, restoration, survivability attributes, quality of resilience.

I. INTRODUCTION

The problem of service differentiation based on recovery has been present for a few years, e.g., [1]-[8]. Although there exists a plethora of proposed methods, none of the frameworks has gained a common approval. The set of parameters used as a basis of the differentiation was characterized in [9], while an overview of these proposals can be found in [10]. In this paper, we propose an extension of the concept called *Quality of Resilience (QoR)* where QoS definitions are revised to meet short- and long-term characteristics while incorporating survivability and service availability issues.

Traditionally, the QoS classes are defined with the usage of the following parameters: delay, jitter, bit error ratio (BER), packet loss probability, throughput, etc. These parameters are called *QoS parameters*, and they represent the requirements of the users regarding the service. In the traditional Service Level Agreements (SLAs) both the long and the short term characteristic of the service is defined. An example is an SLA with the packet loss probability of 5% that is measured over a long period of time, while no minute with a ratio of packets lost exceeding 20% is acceptable.

In our QoR framework we propose to separate the short and long term quality factors. The resilience can be meas-

ured only with long term quality factors so that after the separation we can unambiguously quantify the effects of resilience.

We divide the whole duration of the connection or session into intervals Δt long. The quality metrics that can be evaluated within this time period Δt are the short-term quality metrics and they are called *availability parameters*. The parameters related to long-term characteristics, called *QoR parameters*, define criteria for the service duration of the connection. Fig. 1 shows the extended SLA where we have the short-term availability parameters and the long term QoR parameters.

The short term quality metrics measure the satisfaction of the user over a very short period of time (Δt). The dependence between the traditional QoS parameters and the user satisfaction can be very complex¹. In this paper we treat the user satisfaction as a binary function over a very short period of time Δt . In other words, the user is either fully satisfied with a service for a short period of time Δt or not at all². At such a fine granularity of time, we do not deal with partly satisfied users and restrict their judgment of the service into only two options. For example, in voice services merely those seconds matter in which the customers can clearly hear each other. Obviously, the length of the time period Δt should be very short, but long enough to decide if the service is satisfactory. For each time period Δt the satis-

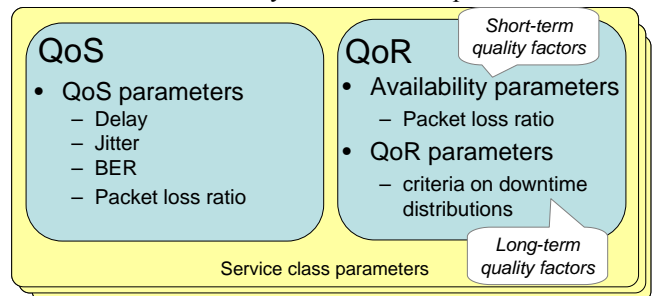


Figure 1. The service class parameters in SLA.

¹ For instance, for voice communication in packet networks the *E model* defined by ITU-T Rec. G.107 is used.

² This is a similar approach as accepted in ITU-T Rec. E.800 [23] where similarly two states (available/not available) are defined for a service.

The value of the first bar is the probability of having an available service. Due to scalability problems in the downtime histogram, it is not drawn on the rest of the charts in the study. The value of A can be intuitively estimated on the basis of the chart since the bars sum to 1. In other words, if the area for $x > 0$ is higher, the probability of having available service is lower.

Fig. 3a shows the histogram of the downtime perceived at the transport layer. It concerns mostly circuit switched networks, where the unavailability is based on the detection of *Loss of Signal*, *Loss of Light* or *Loss of Frame* in SONET/SDH-like protocols [12]. It is mainly caused by failure and is recovered within a short period of time by a resilience mechanism and also called *recovery time histogram*. In this example, shared protection mechanism is performed for a single failure in European Reference Network (see also Fig. 4b). 40 ms was the shortest time to recover from a failure. However, after fatal network failures, the service was not available for several hours while the network was repaired.

Fig. 3b shows the histogram of the downtime perceived by the user. The unavailability at the service layer concerns mostly packet traffic. For example, the basis for the IP service availability function is a threshold of the *IP Packet Loss Ratio* (IPLR) performance measure. In our case, when IPLR exceeds this threshold in a time interval, the service is considered as unavailable. The unavailability perceived by a user is often caused by network congestion or any packet layer service degradation like delay, jitter, etc. (see the orange part of Fig. 3b). The unavailability inherited from transport layer is drawn with dark blue on Fig. 3b. The impacts of failures in IP network have been extensively studied in [11] where the authors show that the length of the unavailable period can be much longer than the length of the service outage at transport layer. In addition, TCP segments in transport layer which are lost during the outage are re-sent after the service is restored, which might lead to network congestion and degraded QoS. This phenomenon is also illustrated in the figures by repeating the dark blue part from Fig. 3a to Fig. 3b after shift it to a bit longer downtimes. Some of the failures that cannot be restored at the transport layer might be restored at the service layer (by IP restoration) after a few seconds. It is drawn light blue on Fig. 3b and reduces the probability of extremely long unavailability periods where the user should wait until the staff repairs the failed link compared to Fig. 3a. It is necessary to note that periods of time depicted in Fig. 3 are not identical with Δt periods used by the definition of the instantaneous service availability.

In most cases, the user cannot clearly distinguish the reason of the unavailability, e.g., in a VPN scenario, failures might lead to a complete network breakdown and the user perceives a 100% packet loss over a considerable amount of time. Note that with loose availability parameters, only a service broken due to a failure would be unavailable. Nev-

ertheless, QoS degradation has statistically larger impact on service degradation than network failures and the user cannot even perceive the reason of the unavailability of the service. For instance, the end user recognizes longer response times in the application layer, due to a certain level of TCP throughput degradation in the transport layer, which is caused by the increase of delay/jitter and packet loss in the IP layer, but not a failure in the physical layer. The eventual reason of these impairments might be either network congestion (due to traffic variability) or a failure. Nonetheless, in the first case, the network operator could deal with the QoS degradation, while in the second case, a resilience mechanism should be used.

III. RECOVERY TIME HISTOGRAM

In the remainder of the paper, we focus only on the recovery time histogram, which is the downtime histogram at the transport layer. For QoS parameters quantitative values are preferred. Thus, the number of downtime intervals (or bins of the histogram) is reduced through merging some of the intervals. While reducing the number of downtime intervals three principles should be kept. First, it should meet the requirements of the user. Secondly, it should meet all technical limitations of the service provider and thirdly there should be as few as possible intervals defined.

The significance of resilience as an important aspect of network operation has generated the necessity of analyzing its effect in service class definition. Therefore, the goal is to give a good approximation of the downtime distribution that can be simply derived, introducing a new performance metric to compare all types of restoration and protection solutions which can be offered to a client.

For better illustration, in some chart we use the cumulative histogram. After normalization, the relative cumulative histogram is akin to the cumulative distribution function.

With the recovery time histograms the operator can get a very clear view of the characteristic of each recovery scheme that facilitates to find the proper solution for each QoS class. The evaluation of the recovery time histogram for each connection and recovery scheme is discussed below.

Generally, the routing algorithms are performed on a weighted graph that is created according to the given network. The transformed graph can be produced by modeling each network element in the original network as an arc in the graph. A Shared Risk Group (SRG) is defined as a group of network elements (i.e., links, nodes, physical devices, software/protocol identities, etc, or a mix of them) subjected to the common risk of a single failure. In this circumstance, each SRG of the original network can be represented by a set of arcs in the transformed graph. Similarly, a well-known example of modeling a node in the transformed graph is using the node splitting technique, where each node is split into two twin vertices in the transformed graph connected by an inner arc. In addition, each SRG is treated

as a possible failure event with an estimated probability of being unavailable. Moreover any combination of SRGs can be a “failure event” with a very small probability of being unavailable. The unavailability of multiple SRGs is the multiplication of the unavailability of the corresponding SRGs.

A histogram is evaluated for each connection, such that all failure events concerned are simulated and the histogram is updated. To process a failure event the following to options should be considered:

A failure event is recovered: Its chance is equal to the probability of the failure event multiplied with the probability of its successful recovery. In this case the recovery time is evaluated and the corresponding bin in the histogram is updated with the obtained resultant probability.

A failure event is not recovered: Its chance is equal to the probability of the failure event multiplied with the probability of its unsuccessful recovery. In this case the bin corresponding to the reparation time is updated with the obtained resultant probability.

The exploration of all failure events concerned is done in the following order: first, the SRGs that are involved in the working path are listed as failure events and processed. Obviously, the connection is available unless any of the listed failure events attacks the network. When the network is protected against single failures, all the dual failures are processed that attracts the working and the protection path at the same time. When the network is protected against dual failures, all failure events are processed that attacks the working, the first protection and the second protection route at the same time. etc..

IV. SIMULATION RESULTS

In this section, we show the recovery time histograms for various recovery schemes for single layer optical and MPLS networks. For simplicity, only one histogram is evaluated for each network and recovery scheme, which is the average of the histograms calculated for all connections. With the resultant histograms the operators are able to choose the proper recovery scheme for each QoS class.

A. Availability Model

The network component availability model of [15] was adopted with *Mean Time Between Failures (MTBF)* and *Mean Time To Repair (MTTR)* values shown in Fig. 4a. With these values, a resultant *MTBF* and *MTTR* of each SRG can be derived by applying the following calculation recursively. If an SRG has two network elements (with availability values of $MTBF_1$, $MTTR_1$ and $MTBF_2$, $MTTR_2$, respectively) and the failure of any of the components leads to the failure of the whole SRG, the resultant availability values are⁵:

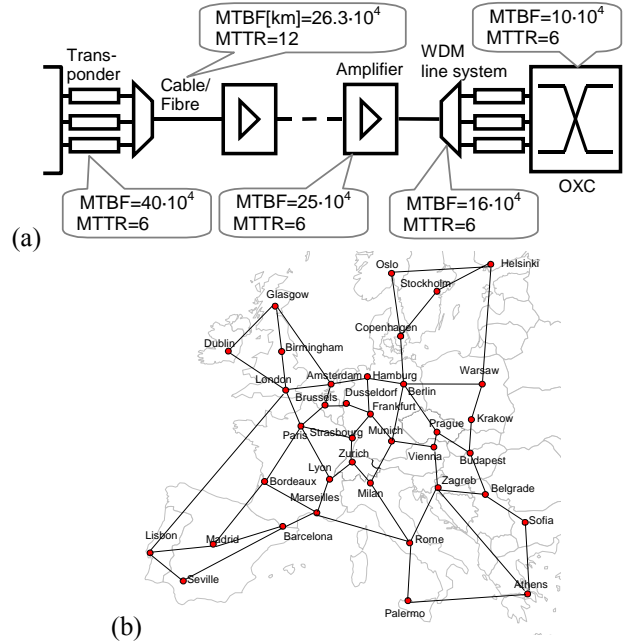


Figure 4. MTTR and MTBF [hours] values used in network component availability model (a) and the European Reference Network (E-Net) (b).

$$MTBF_{SRG} = \frac{1}{\frac{1}{MTBF_1} + \frac{1}{MTBF_2}} \quad (2)$$

$$MTTR_{SRG} = \frac{MTTR_1 \cdot MTBF_2 + MTTR_2 \cdot MTBF_1}{MTBF_1 + MTBF_2} \quad (3)$$

The links of the network are described by two values: the ‘link interface’, which does not depend on the length of the fibre and the ‘link per km’ where the *MTBF* should be divided with the length of the fibre expressed in km. The *MTTR* does not depend on the length of the fibre in any cases. In GMPLS the link interface consists of a multiplexer and demultiplexer. The transponders are considered only at the terminals of the lightpath. The link per km value is related to a buried fibre and optical amplifiers placed every 100 km. The residual *MTBF* (and the *MTTR*) values for the SRGs assigned to the link are calculated according to Eq. (2) and Eq. (3) with the following values:

	MTBF			MTTR		
	node	link per km	link interface	node	link per km	link interface
GMPLS	10.0	238.0	8.0	6.0	11.4	6.0

B. Recovery Time Model

In order to give an average time approximation for the restoration of each SRG failure, the following recovery

⁵ The exponential distributions of failure times are assumed [24]. Eq. (3) is calculated as a weighted average value

where *MTTRs* are weighed by probabilities of different SRG failures.

time model of [16] was adopted. It is based on the MPLS recovery time sequence diagram of [3].

The average case total recovery time (t_r) expressed in ms for shared segment protection can be modeled as follows:

$$t_r^{[ms]} = (30 + f \cdot 0.01) + n \cdot 10 + \sum_{i=1}^{n-1} l_i^{[km]} \cdot 0.005 + |P_k| \cdot 10 + \sum_{i=1}^{|P_k|} l_i^{[km]} \cdot 0.01 \quad (4)$$

where f is the number of TDM LSPs interrupted by the link failure. Variable n denotes the number of OXC nodes, from (and including) the upstream node adjacent to the failure, to the branch node of the recovery segment. The length of the i -th link of the segment, after the branch node and before the merge node, is denoted by $l_i^{[km]}$. Note that the recovery time for shared path protection can also be represented using this analytical formula by setting the branch node (respectively the merge node) as the source (respectively destination) node. The term $|P_k|$ is the number of links on the k -th protecting segment of the failed k -th working segment (i.e., the number of links on the k -th protecting segment between the recovery branch node and the recovery merge node).

C. Numerical Example

As a comparison among single link, single link or node, dual link and dual link or node failures have been analyzed. Simulations are conducted on various reference networks and random topologies. Due to a lack of space the results are illustrated in details only for the pan-European GMPLS fibre-optical network defined by IST project LION and COST action 266 [17]. It has 28 nodes and 57 bi-directional links as shown on Fig. 4b. A traffic matrix in year 2006 is estimated according to [18]. A dynamic traffic pattern is generated according to the traffic matrix such that an Interrupted Poisson Process and Pareto inter-arrival times are integrated together with exponential holding time.

The goal of the simulation was to show illustrative recovery time histogram charts for each recovery scheme. Thus, 253 connection requests were lunched without blocking any demand. It was achieved by selecting only those connections requests where at least 3 SRG disjoint paths exist to make the dual failure scenario feasible. In addition the bandwidth of the connections was reduced to have a light loaded (12%) network, where the effects on routing caused by the lack of capacities are not significant. Inverse capacity proportion [19] is applied as a traffic engineering method to calculate link costs for each connection.

Fig. 5 shows the recovery time histograms for no protection and dedicated protection scenarios. Without protection we had 0.00334 probability of losing the service. Dedicated 1+1 protection gives a very fast restoration in less than 20 ms and reduces the probability of losing

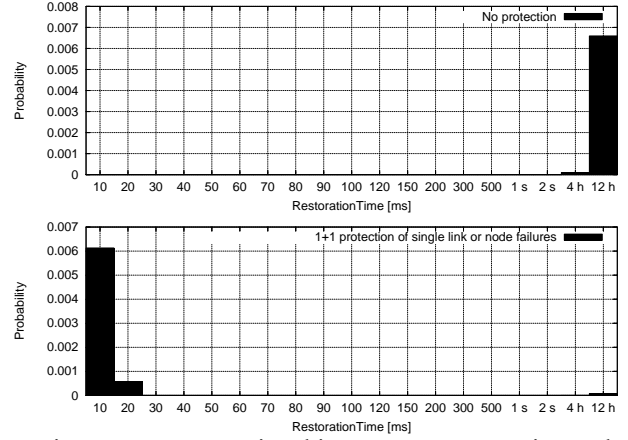


Figure 5. Recovery time histogram: no protection and dedicated protection.

the service more than 80 times (to 4.067×10^{-5}). Dedicated 1+1 has a very attractive recovery time histogram chart; even with its weakness it has good abilities to support services with high QoR parameters.

Fig. 6 shows the recovery time histogram for two shared path protection (SPP) approaches: in the first case, the working path is routed on the shortest path, and in the second step an SRG-disjoint protection route is selected to permit sharing of the protection capacity on links, which is for protecting SRG disjoint working routes. It requires two shortest path searches and is referred as “shared path protection with Dijkstras” in the figures (or “SPP Dijkstra”). This scheme is currently favoured in IETF deliberations for MPLS-layer protection and MPLS-controlled optical path protection [20]. The second SPP method is a single step approach, which calculates the minimum cost working and protection paths jointly. It uses ILP to ensure the optimality of the solution (and is referred as “SPP ilp”). The ILP can save on average appr. 15% of network resources and can decrease the blocking by ~5% [21], however it leads to lower availability (see Fig. 10). The lower availability comes from the fact that the ILP may route the working

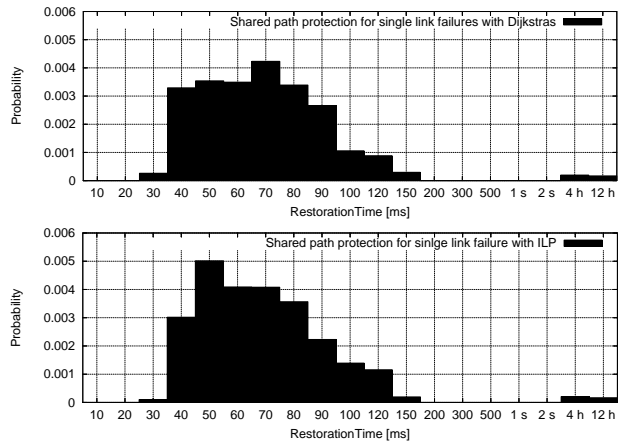


Figure 6. Recovery time histogram: shared path protection.




Figure 7 Cumulative recovery time histogram for SPP and SSP methods that protects link and node failures




Figure 8 Recovery time histogram: shared segment protection with restoration time constraints

path on a slightly longer route if it can save spare capacity for the protection route, and longer working route has a higher chance to be influenced by failures. When nodes are also protected against failures, this difference in availability value is 0.000216, which shows the real jeopardy using the sophisticated and highly capacity efficient methods.

Fig. 7 and 8 show the recovery time histograms for shared segment protection (SSP). The first approach (referred to as “SSP Dijkstra”) is similar to the two step approach introduced for SPP, where in the first case the working path is routed on the shortest path, and in the second step an SRG-disjoint shared protection route is selected with a similar heuristic of [22]. The second approach can derive optimal or near optimal SSP solution with restoration time constraints. It uses an ILP formulation in a heuristic frame-

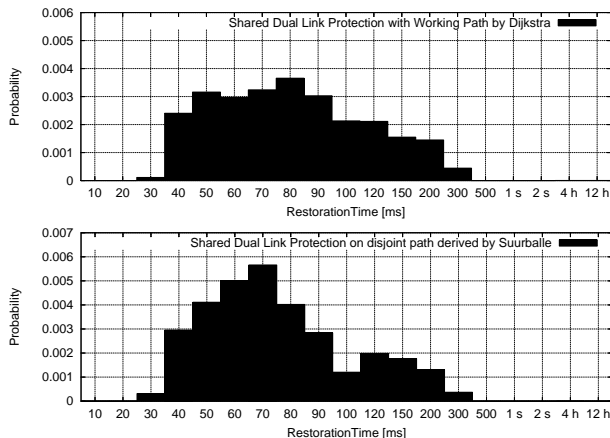


Figure 9 Recovery time histogram: shared dual failure protection

work [16] (and is referred to as “SSP rest const” below in the figures). The chart illustrates the usage of segments which results in the recovery time significant decrease. Since both “SPP Dijkstra” and “SSP Dijkstra” uses working path as a shortest path their service availability is similar. SSP with restoration constraints showed its excellent ability to guarantee short recovery for the price of finding longer working paths leading to lower service availability values (see also Fig. 10).

Fig. 9 shows the recovery time histograms for two shared protection scenarios protecting against dual failures (DF). The first two scenarios protect dual link failures only, while the third protects dual link or node failures. The first method (referred to as “DF Dijkstra”) is a generalisation of the “SPP Dijkstra”, where the working path is shortest path routed, and the first and second shared protection path is

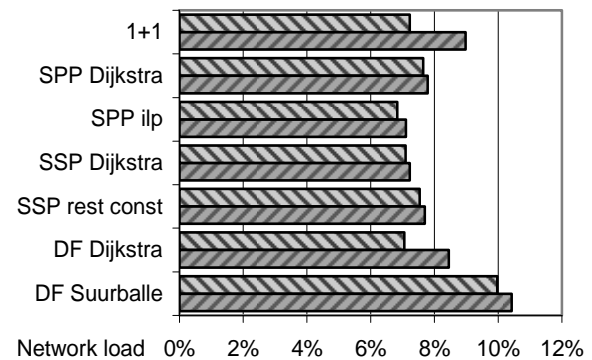
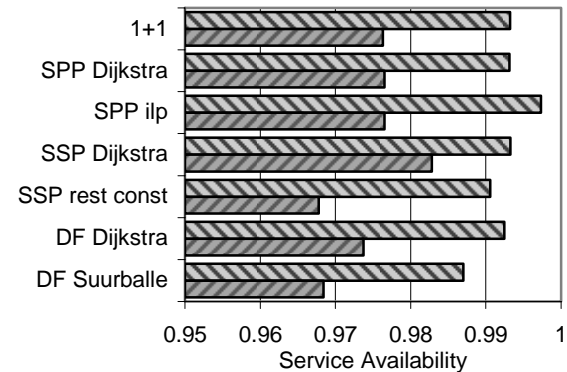
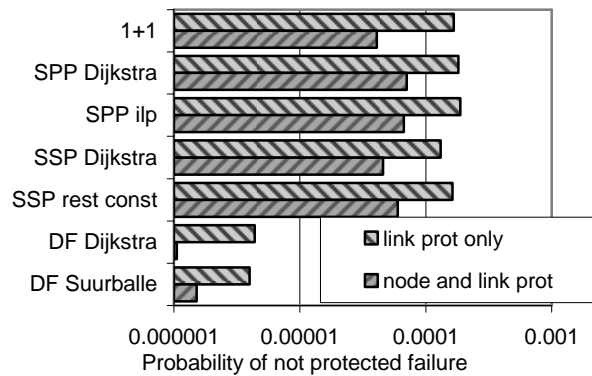


Figure 10. Overall comparisons of the selected recovery schemes.

selected as the shortest link-disjoint path-pair, such that sharing of the protection capacity is permitted in links protecting SRG disjoint working routes. A shortest link-disjoint path-pair can be calculated with Suurballe's algorithm [25]. In the second method, Suurballe's algorithm is used to calculate 3 minimum cost disjoint paths (it is referred to as "DF Suurballe"), and the shortest one is selected as the working route (while the two protection routes are derived in a similar way). The drawback of the first method in comparison to the second one is that it may fail even if there are 3 disjoint paths between the source and destination nodes, although the second has a lower blocking. The probability of having a failure that cannot be restored is $\sim 10^{-6}$ when dual link failures are protected and $\sim 10^{-7}$ when all combinations of two network elements (that can be a link or a node) are protected. The resilience against dual failures requires 40-50% more network resources (see also Fig. 10) leading to higher network utilisation and longer working routes. As a consequence, we have lower service availability than in the single link protection case. The recovery time in case of dual link failures is almost two times longer compared to the case when only a single working path is set up. Another important result presented in Fig. 19 is the huge improvement in service availability by protecting node failures besides link failures.

V. CONCLUSION

In this paper a new concept is presented, which unambiguously quantify the effects of the recovery methods on the service quality. The main idea of the paper is to evaluate a recovery time histogram on simulated European transport network, providing a detailed comparison over various resilience mechanisms and solutions.

ACKNOWLEDGEMENTS

This work has been done within the EU FP6 IP IST-NOBEL (<http://www.ist-nobel.org>) framework. The authors thank György Lajtha, Stefan Bodamer, Didier Colle, Håkon Lønsethagen, Inge-Einar Svinnset, Árpád Szilávik, Attila Mihály, Csaba Antal and András Császár for their helpful comments on the paper

REFERENCES

[1] C. V. Saradhi, M. Gurusamy, L. Zhou, "Differentiated QoS for Survivable WDM Optical Networks," *IEEE Opt. Comm. Suppl. to IEEE Comm. Mag.*, Vol. 42, No. 5, May 2004, pp. S8-S14.
 [2] O. Gerstel, G. Sasaki, "Quality of Protection (QoP): A Quantitative Unifying Paradigm to Protection Service Grades," *Opt. Netw. Mag.*, Vol. 3, No. 3, May/June 2002, pp. 40-49.
 [3] A. Autenrieth, "Recovery Time Analysis of Differentiated Resilience in MPLS," *Proc. DRCN 2003*, pp. 333-340, Banff, Canada, Oct. 2003.
 [4] R. Clemente, L. Serra, G. D'Orazio, G. Cosmo, "A Framework for Class of Service Definition in GMPLS-

based Meshed ASTN," *Proc. DRCN 2003*, Banff, Canada, Oct. 2003.
 [5] M. Tacca, *et al.*, "Differentiated Reliability in Optical Networks: Theoretical and Practical Results," *IEEE/OSA Journal of Lightwave Technology*, Vol. 21, No. 11, November 2003, pp. 2576-2586.
 [6] S. Arakawa, *et al.*, "Design Method of Logical Topologies with Quality of Reliability in WDM Networks," *Phot. Netw. Comm.*, Vol. 5, No.2, Mar. 2003, pp. 107-121.
 [7] E. Calle, *Enhanced Fault Recovery Methods For Protected Traffic Services in GMPLS Networks*, PhD Dissertation, Universitat de Girona, Spain, Feb. 2004.
 [8] P. Cholda, A. Jajszczyk, K. Wajda, "A Unified Framework for the Assessment of Recovery Procedures," *Proc. IEEE HPSR 2005*, Hong Kong, 12-14 May 2005; <http://www.kt.agh.edu.pl/~cholda/Papers/HPSR2005.pdf>
 [9] P. Cholda, A. Jajszczyk, B. E. Helvik, A. Mykkeltveit, "Service Differentiation Based on Recovery Methods," *Proc. 2nd EuroNGI Workshop on Traffic Engineering, Protection and Restoration for NGI*, Rome, Italy, April 21-22, 2005; <http://www.kt.agh.edu.pl/~cholda/Papers/EuroNGI2005.pdf>
 [10] J. Tapolcai, P. Cholda, T. Cinkler, K. Wajda, A. Jajszczyk, A. Autenrieth, S. Bodamer, D. Colle, G. Ferraris, H. Lønsethagen, I.-E. Svinnset, and D. Verchere, "Quality of Resilience (QoR): NOBEL Approach to the Multi-Service Resilience Characterization," *Proc. 1st IEEE/CreateNet International Workshop on Guaranteed Optical Service Provisioning GOSP 2005 (co-located with 2nd International Conference on Broadband Networks BROADNETS 2005)*, Boston, MA, Oct. 7, 2005; <http://www.kt.agh.edu.pl/~cholda/Papers/GOSP2005.pdf>
 [11] G. Iannaccone, C. Chuah, R. Mortier, S. Bhattacharyya, C. Diot, "Analysis of link failures in an IP backbone", *Proc. of the 2nd ACM SIGCOMM Workshop on internet Measurement*, Marseille, France, Nov. 2002
 [12] J.-P. Vasseur, M. Pickavet, P. Demeester, *Network Recovery. Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*, Morgan Kaufmann Publishers, San Francisco 2004.
 [13] J. Marescaux, J. Leroy, M. Gagner, F. Rubino, D. Mutter, Didier, M. Vix, S. E. Butner, M. K. Smith, "Transatlantic Robot-Assisted Telesurgery," *Nature* 413, Sep. 2001, pp. 379-380.
 [14] Y. Bar-Cohen, C. Mavroidis, M. Bouzit, B. Dolgin, D. L. Harm, G. E. Kopchok, R. White, "Virtual reality robotic telesurgery simulations using MEMICA haptic system", *Proc. SPIE Smart Structures and Materials Conference*, Newport, CA, March 2001, Paper No. 4329-47
 [15] S. Verbrugge, D. Colle, P. Demeester, R. Huelsermann, M. Jaeger, "General Availability Model for Multi-layer Transport Networks," accepted to *DRCN 2005*.
 [16] J. Tapolcai, P.-H. Ho, D. Verchere, T. Cinkler, "A Novel Shared Segment Protection Method for Guaranteed Recovery Time," accepted to *BROADNETS 2005*.
 [17] LION and COST 266, "Reference networks, 2003,"

Part of the European Information Society Technologies (IST) Fifth Framework program.

[18] M. Vaughn, R. Wagner, "Metropolitan network traffic demand study," *Proc. 13th annual meeting Lasers and Electro-Optics Society (LEOS 2000 annual meeting)*, Rio Grande, Puerto Rico, November 2000, Vol. 1, pp. 102-103.

[19] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," *Proc. IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000, pp. 519–528.

[20] W. Grover, J. Doucette, M. Clouqueur, D. Leung, and D. Stamatelakis, "New options and insights for survivable transport networks," *IEEE Communications Magazine*, vol. 40, pp. 34–41, January 2002.

[21] P.-H. Ho, J. Tapolcai, H. T. Mouftah, and C. -H. Yeh, "Linear Formulation for Path Shared Protection", *Proc. IEEE International Conference on Communications (ICC)*, Paris, France, June 2004, in Optical Networking Symposium.

[22] D. Xu, Y. Xiong, and C. Qiao, "Protection with Multi-Segments (PROMISE) in Networks with Shared Risk Link Groups (SRLG)", *Proc. 40th Annual Allerton Conference on Communication, Control, and Computing*, 2002.

[23] ITU-T Recommendation E.800: *Terms and Definitions Related to Quality of Service and Network Performance Including Dependability*, August 1994.

[24] K. S. Trivedi, *Probability and Statistics with Reliability, Queuing, and Computer Science Applications*, John Wiley & Sons, Inc., New York 2001.

[25] R. Bhandari, *Survivable Networks. Algorithms for Diverse Routing*, Kluwer Academic Publishers, Dordrecht, The Netherlands 1999.