Journal of Network and Systems Management, Vol. 13, No. 1, March 2005 (© 2005)
DOI: 10.1007/s10922-005-1844-7

Guest Editorial

Design and Operation of Survivable Networks

Wayne Grover[1] and David Tipper[2]

What do backhoes, sharks, fires, software, floods, hurricanes and train derailments have in common? They all can make communication networks fail. Today, the cut of a thumb-sized fiber-optic cable can disrupt millions of web applications, phone calls, banking services, flight bookings, and so on. At the same time new networked devices and applications such as; sensor networks, grid computing, mobile telephony, wireless enabled PDAs, peer-to-peer file-sharing, videoconferencing, interactive gaming, and virtual reality applications contribute to the creation of a more dynamic and unpredictable environment of time-and-space varying demand for transport connectivity and capacity. Traditionally, the ability of a system to continue to provide services in the presence of failures internal to itself is the art of high availability system design. In decades past, the emphasis was on designing high availability equipment and routing services over such elements in a non-redundant way. But in the last decade, driven by fiber-optics as the preferred—but vulnerable—physical medium, we have seen the advent of survivable networks as a specialized new area of network technology and network planning. A survivable network has abilities to continue providing services in the face of either internally arising, or externally inflicted failures. "Survivability" itself is only a qualitative term referring to the overall ability to carry on providing service in the face of such failures. Survivability can thus be derived through many measures ranging from the armoring of cables at the physical layer all the way up to a company splitting its traffic flows over multiple separate carriers in the service layer.

Survivability, to attain reliability, is therefore an essential property that must be designed into all networks, especially the backbone transport network. The service-layer or "user" networks that we perceive individually, such as the Internet, the phone networks, banking networks, travel networks, and so on, often ride

[1] TRLabs and the University of Alberta, Edmonton, Alberta, Canada, T6G 2V4. E-mail: grover@trlabs.ca; Home page: www.ece.ualberta.ca/~grover/

[2] Department of Information Science and Telecommunications, University of Pittsburgh, Pittsburgh, Pennsylvania. E-mail: tipper@tele.pitt.edu

1064-7570/05/0300-0007/0 © 2005 Springer Science+Business Media, Inc.

on a common transport network infrastructure. This trend toward "converged" networks where all applications and services are converted into IP packet flows and routed over DWDM optical transport has increased recently. The inherent ability of transport networks to recover from failures is thus crucial to commerce and society.

In general, all layers of a network need certain self-healing capabilities to address faults arising at their own layer. When this is done successfully, higher layers will never be aware of the failures that actually occurred at the lower layers. Obviously therefore, the transport layer, which is just above the passive physical layer is fundamentally important to invest with self-healing capabilities. The most common type of physical-layer failure is cable damage arising from natural or man-made causes; trench digging, construction work, craftsperson errors, ship anchors, sabotage, tress falls, earthquakes, rodents, fires, floods, etc. The sheer mileage of fiber-optic cable now deployed in ducts, direct-buried underground, or on overhead pole-lines, is so large that cable-cuts dominate all other sources of externally imposed network failures. Cable-cutting events occur virtually every few days in extensive networks with 50,000 or more route-miles of fiber. The surprisingly high rate of cable cuts, despite many measures for the physical protection of cables, is evidenced in the industry by the black humor of referring to backhoe equipment as "Universal Cable Locators."

As evidence of just how surprisingly frequent failures are, Snow [1] has reported that since 1992 there have been about 16 outages per month in the United States alone that each affected over 30,000 users. And increasingly, interesting (even bizarre) reports of cable cuts and their impact can be found daily on the Internet. A sampling on just one day in 2004 yields stories of ship anchors, train derailments, and the more typical cable dig-up events as well [2–6].

The design and management of networks that are inherently reliable and robust against untoward events such as those listed above was the central theme of the Fourth International Workshop on Design of Reliable Communication Networks DRCN 2003, held in the spectacular venue of the Banff Centre in Banff National Park, Alberta, Canada. This event continued the successful series of prior DRCN workshops in Brugge, Belgium (1998), Munich, Germany (2000), and Budapest, Hungary (2001). TRLabs (at the University of Alberta) (www.trlabs.ca) served as the organizing and under-writing institution to host DRCN for the first time outside Europe. The geographic distribution of DRCN 2003 attendees was remarkable: the total of 148 registrants hailed from 26 countries including Australia, Belgium, Canada, France, Germany, Hungary, Italy, Japan, Singapore, Spain, Sweden, Switzerland, the United Kingdom, and the United States. The technical program consisted of 60 papers from the open call (out of 136 submissions), six invited presentations and four tutorials.

DRCN 2003 benefited from a pre-announced plan to publish a Special Issue of JNSM based on the theme of "Design and Management of Highly Reliable

Networks and Services." Accordingly, six papers were selected for further review and extension following the conference for inclusion in this Special Issue of JNSM.

The first paper "Different Algorithms for Normal and Protection Paths" by R. Gupta, E. Chi, and J. Walrand, considers the issue of whether having different routing algorithms for selecting working and backup paths is beneficial. The work is set in a context where survivability is obtained through the principle of routing a working primary path and then establishing a disjoint secondary or backup path. The straightforward approach is to use the same type of shortest path algorithm once for each sub problem, simply adding a restriction on the second instance not to re-use any links (or nodes) of the first route found. Intuitively, however, one should be able to do better than this, especially if capacity sharing is considered. A set of different routing algorithms for normal and protection paths are, therefore, studied via simulation for various scenarios. The effect of reconfiguration of backup paths to reduce network congestion is also studied. It is shown that using two different routing algorithms for normal and backup paths is advantageous under certain conditions.

The second paper, entitled "Demand-wise Shared Protection for Meshed Optical Networks," by A. M. C. A. Koster, A. Zymolka, M. Jager, and R. Hulsermann, considers an interesting compromise between dedicated and fully shared protection capacity, with the specific limitations of switching and transmission in current optical networks in mind. In this approach, spare capacity is shared amongst the individual lightpaths belonging to the overall bundle of demand between an end-node pair, but not among different demands. Hence, the total demand flow between each node pair has a spare capacity dedicated to it, but within the flows, protection capacity is shared. This approach aims to combine the advantages of both dedicated and shared protection for meshed optical networks. As part of this strategy, the authors explore and advocate the deliberate diversification of normal traffic demand paths to increase the sharing of protection path capacity. Computational results are given illustrating the cost benefits of the approach.

The third paper, by E. Kubilinskas, M. Pioro, and P. Nilsson, is entitled "Design Models for Robust Multi-Layer Next Generation Internet Core Networks Carrying Elastic Traffic." This work considers two-layer survivability design problems in which WDM and MPLS layers are jointly planned for a combined survivability strategy. The aspect of "elastic traffic" is an interesting and useful additional dimension to this problem which recognizes a realistic tendency in the Internet—that the aggregation of applications using the bandwidth over a certain route tend to be able to exploit however much bandwidth is provided for them.

The next paper, "Design of Reliable IP/GMPLS Networks an Integrated Approach," by F. Mobiot, B. Sanso, and A. Girard, was one of the "Best Paper Award" recipients at the DRCN 2003 Workshop. This paper treats the problem of design of survivable IP networks over a transport network such as GMPLS. In particular, the authors consider how to design the IP layer logical network for a normal QoS and a degraded QoS in the event of failures in the transport network. The QoS is determined using an effective bandwidth, network calculus approach. An optimization-based formulation of the logical network synthesis problem is given and a Lagrangian relaxation solution technique proposed. Numerical results are given to illustrate the effectiveness of the integrated design approach.

The fifth paper, “Comparing Restoration Concepts Using Optimal Network Configurations with Integrated Hardware and Routing Decisions,” by S. Orlowski and R. Wessaly, is about making a careful cost comparison between link and path restoration. It is well known that path restoration is inherently somewhat more capacity efficient. And yet, especially for an optical network, span restoration may be considerably simpler and more reliable to engineer from a transmission standpoint because of its relative locality of action. But span restoration (or protection) seems to be an often-overlooked option, because it is so widely appreciated that path restoration will require less spare capacity. Thus, it makes sense to ask what the cost difference may actually be, despite the capacity differences, when real equipment details are brought into consideration. The authors arrive at an important message for us: when all things are considered, span restoration is essentially as cost-effective as path restoration. This means that in practice as long as some type of shared-mesh scheme is chosen (as opposed to rings or 1+1 APS), then one can expect to reap essentially all possible cost benefits.

Last and not least, is the paper by P. Hegyi, M. Maliosz, A, Lad’anyi, and T. Cinkler entitled “Virtual Private/Overlay Network Design with Traffic Concentration and Shared Protection.” This work considers the design of virtual private networks (VPNs) with shared backup protection over a transport network that supports resource partitioning. Different design modes are investigated depending on whether normal and protection paths are optimized locally for each VPN individually, or globally for all VPNs jointly. Three routing-based heuristic algorithms are proposed for finding shared protection paths. Simulation results are given to illustrate the tradeoffs between the design modes and heuristic algorithms.

The Guest Editors are happy to have served in the DRCN 2003 organizing committee—Grover as General Chair and Tipper as Technical Program Chair—and we were delighted to have the official cooperation of JNSM and the guidance of Manu Malek in both promoting DRCN and in developing this Special Issue. We hope these papers offer the readers a view of the important research topics addressed in DRCN 2003. We want to also thank the authors for their contributions, patience, and timeliness in meeting the deadlines. Finally, we want the JNSM readership to know that DRCN carries on and we would like to encourage further linkage and participation between DRCN and JNSM. The planning for the next DRCN—DRCN 2005—is already underway. See www.drcn.org and please consider either attending and/or submitting your own ongoing research for possible inclusion at DRCN 2005.

REFERENCES

1. A. P. Snow and M. W. Thayer, Defeating telecommunication system fault-tolerant designs, *Proceedings of the Third Information Survivability Workshop (ISW 2000)*, Boston, Massachusetts, USA, October 24–26, 2000.
2. I. Tham, ZDNet UK, (September 20, 2001) [Online, Anchor-Draggers Cut Asia's Internet Pipe, available: http://news.zdnet.co.uk/internet/0,39020369,2095715
3. G. Wearden, ZDNet UK, (November 26, 2003) [Online, Cable Failure Hits UK Internet Traffic, available: http://news.zdnet.co.uk/communications/networks/0, 39020345,39118125,00.htm, accessed January 20, 2004.
4. Quebec Scientific Information Network, (January 29, 2003) [Online, Forestville–Rimouski Underwater Cable Repaired, available: http://www.risq.qc.ca/nouvelles/nouvelle_item.php?LANG=EN&ART=1231, 20 accessed January 20, 2004].
5. W. McAuliffe, ZDNet UK, (August 3, 2001) [Online, Train Crash Could be to Blame for Internet Derailment, http://news.zdnet.co.uk/business/0,39020645,2092503,00.htm, accessed January 20, 2004].
6. L. Bowman and M. Broersma, ZDNet News, (June 10, 1998) [Online], Severed MCI Cable Cripples the Net, http://zdnet.com.com/2100-11-510740.html, accessed January 20, 2004.

Wayne D. Grover is the author of "Mesh-based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking," published by Prentice-Hall; Professor in the Department of Electrical and Computer Engineering at the University of Alberta; and Chief Scientist for Network Systems at TRLabs. He has 25 years of experience in industry and academia and issued patents on 26 topics and over 100 other refereed publications. He is an IEEE Fellow "for contributions to survivable and self-organizing broadband networks" and Fellow of the Engineering Institute of Canada. His current interests are focused on p-cycles and the protected working capacity envelope (PWCE) concept, both topics which he originated. He teaches courses on transport network technology, survivable network design, introductory communications, and telecommunication systems at the University of Alberta.

David Tipper is an Associate Professor in the Department of Information Science and Telecommunications at the University of Pittsburgh. He also holds the position of Professor II of Informatics at Molde College in Norway. Prior to joining Pitt in 1994, he was an Associate Professor of Electrical and Computer Engineering at Clemson University in Clemson, SC. His current research interests are network design and traffic restoration procedures for survivable networks, network control techniques, and performance analysis.