Journal of Network and Systems Management, Vol. 5, No. 2, 1997

Guest Editorial

Towards Fault Recovery and Management in Communication Networks*

D. Medhi[1] and D. Tipper[2]

Society's growing use of communication networks in both the public and private sector has led to an increased focus on reliability and survivability. Over the last decade, widely publicized network outages illustrated that a disruption of communications services can be very expensive to businesses and critical services (such as 911 service). The growing vulnerability of the public switched telephone network in the United States was initially discussed in 1989 by the National Research Council which noted "As we become more dependent on network, the consequences of network failure become greater and the need to reduce network vulnerabilities increases commensurately." With the popularity of the Internet, a variety of network failures has surfaced in the recent years. It is clear that fault management will grow in importance for the foreseeable future.

A variety of network failures are possible with typical events resulting in a failure being accidental cable cuts, hardware malfunctions, software errors, natural disasters (e.g., fire, hurricane, earthquake), human error (e.g., incorrect maintenance), and malicious attack (both hardware and software). Broadly, we can classify the failures as physical-type failures and software-type failures. First we discuss physical failures in this context. Consider, a fiber cable cut in a self-healing ring in a metropolitan area telecommunication network; the failure can be restored in milliseconds by changing the direction of flow (with appropriate availability of capacity). However, using such a ring approach on a national or an international scale is likely to be cost prohibitive. This may lead to an approach which allows a combination of ring and digital cross-connect systems for transmission layer restoration, although it should be kept in mind that the spare capacity provided for such restoration is unused during normal network

*Supported in part by National Science Foundation grant NCR-9506652.

[1] Department of Computer Networking, University of Missouri-Kansas City, Kansas City, Missouri 64110, USA. Email: dmedhi@cstp.umkc.edu

[2]Department of Information Science and Telecommunications, University of Pittsburgh, Pittsburgh, Pennsylvania 15260 USA. Email: tipper@tele.pitt.edu

1064-7570/97/0600-0101$12.50/0

operation and, thus, not generating revenue while sitting idle. In many networks, it is not cost effective to provide for transmission level restoration (e.g. network is based on leased capacity from various vendors, all of whom may not have facility restoration), or transmission level restoration is not applicable in certain failure scenarios (e.g. line card failure), thus fault recovery mechanisms must be considered at higher layers also. For example, in an ATM based network, one could deploy virtual path level restoration and virtual channel level restoration. All of the above restoration scenarios require additional capacity in each layer (over and beyond required for normal network operating conditions); the actual combination of additional capacity from each layer may be based on a quality-of-service objective under a network failure condition as well as whether the capacity is revenue producing in the event of a traffic overload (no failure). However, it may not always be possible to have adequate capacity for all failure scenarios—call admission control and traffic network management controls are also needed that are activated under a failure for network stability and operations. Given the possible range of failure types and network technologies, it is quite clear that a single solution approach will not be appropriate or cost-effective for all scenarios. Overall, this suggests that a multi-layer fault recovery procedure is needed with appropriate coordination among various layers for fault management.

It may be noted that the above multi-layer approach is primarily for physical type failure. In the above, we have discussed restoration in the context of capacity and rerouting around a failure. Software failures, or malicious attacks on software are possible causes of faults as well. Through this type of failure, the network is led to believe that there is either no capacity available (e.g. signalling software failure), or unnecessary or fictitious connections are generated which may clog the network. This type of failure thus requires a way to bring the affected network to a known state as well as security procedures to block protocolattacks.

Finally, there are two important issues that arise related to a failure: 1) coordination of alarms that may be generated (by different network elements) due to a failure, and isolation and localization of the fault, and 2) time required to restore from a failure (by taking appropriate actions) and the transient network behavior.

In this special issue of JNSM the articles published provide an overview of the state of the art on fault management presenting work addressing many of the issues noted above. This issue begins with a *Thresholds* article by L. Bernstein and C. M. Yuhas where the authors give an anecdotal view of practices used by some network providers to address a network fault.

The paper, "A Generic Model for Fault Isolation in Integrated Management Systems" by S. Kätker and K. Geihs, presents a methodology for fault isolation by integrating network, system, and service management layers.

In the paper "Sequential and Parallel Approaches to Incorporate Reliability in the Synthesis of Computer Networks," S. Chamberland and B. Sansó present a way to incorporate failure states in the design of computer networks by considering suitable trade-offs between capacity assignment and performance in the event of failures.

Since fast restoration from a failure is desirable, M. MacGregor, W. Grover and K. Ryhorchuk, in their paper "Optimal Spare Capacity Preconfiguration for Faster Restoration of Mesh Networks," present a way to preconfigure networks due to a transmission network failure so that when a failure occurs some of the preconfigured restoration paths can be used immediately, which then can be coupled with a real-time process if additional restoration paths are needed and also to lower the workload on a real-time process and the restoration recovery period.

Typically, in a high-speed network with automatic protection switching from the transmission network, a small detection and restoration time is desirable. D. Logothetis and K. Trivedi in "The Effect of Detection and Restoration Time on Error Recovery in Communication Networks," address this issue to minimize the loss of messages and show that detection and restoration times have a significant impact on network performance.

The paper "Empirical Evidence of Reliability Growth in Large-Scale Networks" by A. Snow and M. Weiss provides an analysis of the major outages in a real network to characterize the network failure behavior, and provides the reader with a perspective on occurrence of network failures. This information is useful in network modeling of future networks since it captures the failure patterndistribution.

Finally, R. Doverspike, in his *Report* article, updates the T1 committee work on network survivability and reports on the recent development to provide survivability in multi-layered networks.

We like to thank the authors for their work and the reviewers for the timely reviews. A special thanks to Manu Malek, the editor-in-chief, for allowing us thisopportunity.

D. Medhi is an Associate Professor in the Department of Computer Networking, a unit of the Computer Science Telecommunications Program at the University of Missouri-Kansas City (UMKC). He received his B.Sc. in Mathematics from Cotton College, Gauhati University, India, M.S. in Mathematics from the University of Delhi, India, and M.S. and Ph.D. in Computer Sciences from the University of Wisconsin-Madison in 1981, 1983, 1985 and 1987 respectively. Prior to joining UMKC in 1989, he was a member of the technical staff at AT&T Bell Laboratories, Holmdel, New Jersey, from 1987 to 1989. His research interests are in survivable network design, dynamic routing, broadband network design, large-scale optimization algorithms, wireless PCS, and network management. His publications have appeared in journals such as IEEE Trans. on Communications, IEEE/ACM Trans. on Networking, and Journal of Network and Systems Management, and

conferences such as IEEE Conference on Computer Communications (INFOCOM), IEEE Military Comm. Conference (MILCOM) and IEEE International Conference on Communications (ICC). He has served on technical program committees of several IEEE conferences.

D. Tipper is an Associate Professor of Telecommunications in the Department of Information Science and Telecommunications at the University of Pittsburgh. Prior to joining Pitt in the Fall of 1994, he was an Associate Professor in the Electrical and Computer Engineering Department at Clemson University. Between 1980 and 1982 he was employed as a System Engineer on NASA's space shuttle mission simulator by Singer-Link in Houston, TX. Professor Tipper is a graduate of the University of Arizona (Ph.D.E.E. 1988, M.S.S.E. 1984) and Virginia Tech (B.S.E.E. 1980). Dr. Tipper's current research interests include methods for improving network survivability, the development of efficient algorithms for nonstationary/transient queueing analysis, the design and analysis of network controls (e.g. routing, admission control, scheduling, etc.) and measurement based ATM traffic models. He is a member of INFORMS, Sigma Xi and a Senior member of IEEE.