

Chapter 1. Introduction

In recent years the world has witnessed exploding demands on communication networks. In addition to more conventional communication devices such as the telephone and television, the development of the Internet provided for interconnectivity of personal computers. This interconnectivity enabled a transition of the personal computer from computing device into a new kind of personal communication device, ultimately creating an enormous world-wide network of these PCs. The booming trend of the Internet has overwhelmed the existing communication infrastructure and has placed demands non-existent just a decade ago.

To meet some of these demands, there has been a global effort to merge the segregated networks into one integrated single network capable of merging long standing foes -- data, voice and video -- into a Broadband Integrated Services Digital Network (B-ISDN). Prior to the integration, data was generally transported by separate and relatively slow packet-switched networks over existing telephone lines; voice was transported via circuit-switched networks; and video was broadcast over separate cable networks. The integration of these services into one network that delivers all of the disparate services will increase the overall efficiency of networks, as well as raise the potential for higher revenues, thus creating a greater ability to meet the growing customer demand for these services.

However, the same integration will make these networks operate at higher risk, since the amount as well as the importance of information flowing through will increase. This condition, in turn, will increase the importance of the networks to offer reliable and continuous performance. Furthermore, just offering a capability to transfer data is becoming an insufficient condition for

good performance of communication networks; the birth of numerous new services such as e-commerce and video-conferencing have further expanded requirements where the quality of the service offered is becoming just as important as the service itself.

In other words, it is not only important that the data can get from point A to point B, but also how fast and with what quality it can get there. Therefore, anything that impairs the network's capability to offer a service, as well as the quality of that service, also plays an extremely important part in the overall trend of analyzing and retrofitting today's networks and planning, and designing future networks to make them more immune to any kind of failure.

The escalating demands also have resulted in much work on categorizing, classifying, and analyzing network failures and, although somewhat ambiguous, classification of failures has received a lot of attention from the research community and telecommunications standards-setting bodies. The classification has been focused primarily on complete or hard failures of networks or their parts, looking at whether a service is offered; it has basically not addressed the possibilities of the effects of soft failures on network performance which could potentially affect the quality of an existing service.

Generally, soft failures include a very broad range of conditions, which do not completely terminate a link, connection or a service. Rather, in a subtle way, they impair and reduce the overall performance of a network. Due to overwhelming end user demand for quality, maximum bandwidth and speed of connectivity, it is of paramount importance that existing as well as future networks perform at optimum levels. Even more important than extracting the maximum bandwidth out of the infrastructure is the influence soft failures can have on BER, which in turn may wreak havoc on quality of service (QoS). Under

these strenuous conditions, the possibility of “slight” performance decreases potentially caused by soft failures is becoming an important factor.

Well-designed, installed, and operated fiber optic networks should not, in principle, experience soft failures. In real networks, however, there is the potential for soft failures such as unexpected or unaccounted attenuation from microbends and cracks, and improper fiber and/or connector use resulting from improper installation. It is expected that the horizontal cabling associated with fiber to the desktop will be much more exposed and thus vulnerable to a variety of soft failures. Furthermore, the multilayered nature of computer network protocols only adds new possibilities for creating different kinds of soft failures, as well as for hiding and masking the existing ones.

This research investigates the suspected “gray area” of network performance that lies somewhere between optimal performance and complete failure, and thought to be affected by such soft failures as mentioned above. To determine the existence of such an area, soft failures are emulated in a laboratory network setup by inserting variable attenuation in the fiber connecting two ATM switches. By applying gradual attenuation to mimic soft failures, and then using a powerful software tool to generate and monitor network traffic, it is possible to detect deterioration or degradation of network performance.

1.1. Research Objective

As mentioned before, a great deal of previous research has been done in classifying failures. Based on the work of ANSI Technical Subcommittee T1A1, three features -- unservability (U), duration (D), and weight (W) – can be used to describe network failures as catastrophic, major or minor [1]. Catastrophic

network failures are often triggered by such disastrous occurrences as earthquakes, floods, tornadoes, power grid failures, or war. Major cable backbone failures can occur as a result of fire, terrorism, or tandem switch failure. Minor single link or path failures often are due to fiber failure, equipment failure, and shelf or unit failure.

Most of the recovery solutions for these failures have been incorporated in the lower layers of networks such as SONET and ATM, and usually include individual layer restoration features such as automatic protection switch (APS), self-healing ring (SHR), self-healing network (SHN), or failure-resistant virtual path (FRVP).

The multilayered nature of computer networks (for example, WDM/SONET/ATM/IP) makes the problem of network restoration more complex; nevertheless, recovery solutions have also been offered such as layer escalation, subnetwork escalation, and scheme escalation.

Soft failures are not as easy to define and have not yet been studied as extensively as hard failures. Most of the higher layer protocols assume that the fiber is a perfect medium and that errors are either the fault of other network malfunctions or there is a complete failure of the physical layer. To further complicate soft failure situations, if a certain level of degradation does exist in the fiber, such issues as congestion, contention, number of acknowledgements and retransmissions, proprietary equipment and code incompatibility make it even more difficult to determine the actual cause for a decrease in performance.

In addition, there is no “standard” network architecture – networks are all configured differently. As a result, it would be too impractical, maybe virtually impossible, to examine soft failures for all possible existing networks. Therefore,

the goal of this thesis project is to provide insights into network performance under certain specific soft failure conditions with very controlled system parameters, rather than accomplish a conclusive study that is applicable to all networks. The hope is that after gaining insights into degradation versus performance, certain conclusions can be extrapolated to network performance in general.

To achieve this, the scope and size of the testbed network was limited, and real hardware and software were used for the setup, as opposed to employing a simulated network. It was determined that a “real” system would elicit better responses and would result in more useful, specific data.

The test set-up included two ATM switches (Olicom CrossFire 9100 and 9200), two PCs (450 MHz Pentium III processors with 128 MB of RAM), and two NICs (RapidFire 6162 ATM 155 PCI), all connected by OC-3 multimode optical fiber links. Furthermore, two variable attenuators were inserted between the two switches and the same set of tests was executed for different levels of attenuation.

The ATM switches were configured to run LANE, where the 9200 operated as the primary LANE administrator and the 9100 was the secondary. The emulated protocol was Ethernet (IEEE 802.3). This configuration offered a SONET/ATM/LANE/Emulated Ethernet/IP/TCP/Application protocol stack.

The application layer scripts and network performance tests were performed using Ganymede’s Chariot software [2]. Chariot is a software tool designed to test end-to-end network performance of complex and multiprotocol networks, and is distributed on the PCs as end users of the network.

For this project, four main tests were selected; three emulated Bader¹ benchmark or classic transactions, internet applications, and multimedia data; the fourth was a test designed for “stressing” the network by running 16 simultaneous connections between the two end points. Then, each test was repeated for increased attenuation levels until the ports on the ATM modules displayed a red light as an indication of insufficient optical power to the switch.

Performance is evaluated based on the following parameters which are observed and compared for different attenuation levels: throughput, transaction rate, and response time for TCP-based applications (Tests 1 and 2), and throughput, lost data and percent bytes lost for UDP-based applications (Tests 3 and 4).

1.2. Thesis Organization

Chapter 2 is a brief history of computer networks, and includes descriptions of the standardization initiatives ISDN and B-ISDN.

Chapter 3 provides an overview of computer networks, including topics such as OSI, SONET/SDH, and ATM. Bandwidth management of ATM networks is also described, as well as insights into different service categories offered by ATM, quality of service and flow control issues.

Chapter 4 gives a description of network failures, their classifications, network restoration techniques and "soft" failures.

¹ Chariot names this class of application scripts after Lance Bader who initially specified them for IBM's System Network Architecture (SNA). Now, they can generally be applied to all network architectures.

Chapter 5 offers an explanation of the application layer of the computer network. It describes the scripts used for generating network traffic, as well as the metrics used for network performance evaluation.

Chapter 6 describes the configuration of the experimental setup, as well as the actual equipment used. It further explains the parameters of each component in the setup, and the response variables on which the evaluation model is based, such as attenuation, throughput, transaction rate, response time and lost data.

Chapter 7 presents the performance test results. It offers the obtained data for individual response variables, as well as the correlation between the network performance and the attenuation. It also attempts to relate the given attenuation through the fiber with the corresponding bit error rate.

Chapter 8 gives the conclusions and recommendations for further research.