

Innovative Technologies for Preventing Network Outages

Lawrence Bernstein

AT&T and companies engaged in similar communications enterprises are in business for two principal reasons: to satisfy customers by providing the products and services they desire, including uninterrupted end-to-end telecommunication services; and, in so doing, to earn the best possible return on investment, thereby rewarding shareholders and employees alike. As the fields of computing and telecommunications converge in the emerging new global networks, scientific advances often result from the aggressive pursuit of these goals.

Introduction

Operations Systems (OS) are required to support the maintenance, administration and management of global telecommunications networks. OS have evolved through three distinct developmental stages, or three waves, and we're now riding the crest of a fourth wave (Figure 1).

A service crisis precipitated the first wave, which occurred during the 1970s and early 1980s. It started within the New York Telephone system. Soon, however, it became apparent that the provisioning complexities there were a harbinger of problems that would become universal. One result of efforts to address the complexity dilemma was the centralization of work groups that had previously been dispersed—including surveillance systems, such as AT&T's Switching Control Center System. This consolidation of expertise allowed for more efficient use of the available technical resources.

During the 1980s, the need for cost reductions led a second wave of change in OS. Automation made it possible for systems to provide overall process control, replacing routine functions previously performed by humans. A good example of this second wave of OS is the Loop Maintenance Operations System, which performs the trouble-ticket management, automatic testing and analysis, dispatch and problem close-out, with control of the overall system process. Automation was achieved so successfully that investment

bankers were astonished when the seven Regional Bell Operating Companies—or "Baby Bells"—did not seek financing on the open market for several years after they were created. Savings from automation techniques provided sufficient equity for the companies to self-finance and expand.

An increasing desire to stimulate business growth by means of new service provisions and revenue opportunities fueled the third wave of change in OS management. At the start of the 1990s, the focus was on linking various operations systems to provide the Local Exchange Carriers (LECs) with a seemingly "seamless" overall functionality. As new services were introduced, flow-through maintenance and integrated engineering became fundamental to analyses and decision making. At the same time, development moved clearly in the direction of integrating voice, data, and video networks.

Despite the more efficient vital controls that are now built into OS, the current management model contains one major problem: OS remain separated and apart from the actual network. These telemetry networks (sets of OS), which are removed a layer from the actual network, are heavily relied upon to monitor, test, inform and track every network parameter. But these telemetry networks are themselves fragile, and they're removed from the actual network. Error messages generated by telemetry systems provide a good example of how this vulnerability can mani-

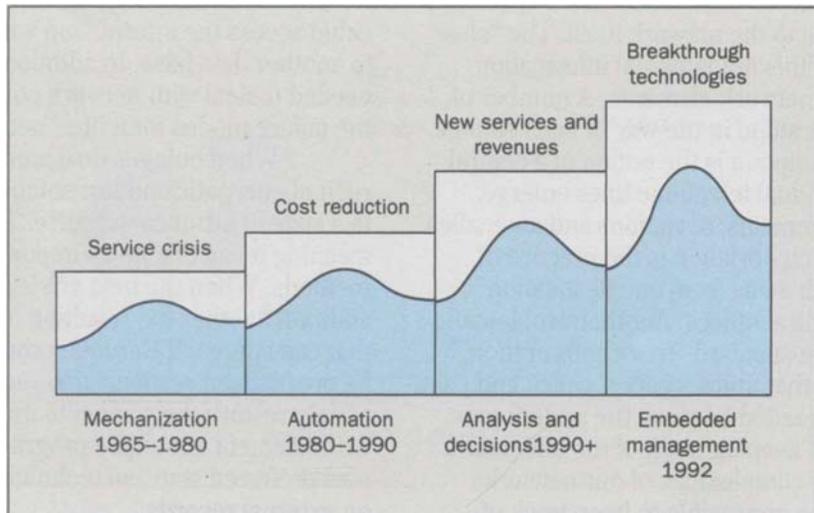


Figure 1. System diversity, achieved by combining software containing error-correcting codes with geographical dispersion, will become widely practiced in the near future.

fest itself. As the fault is often in the monitoring system, rather than in the actual network, there's a tendency to ignore error messages from telemetry systems. This tendency is a prescription for outages. The intervals between each wave of innovation have become progressively shorter, and the fourth wave is already underway. It's being driven by breakthrough technologies that permit management systems to be embedded in the network itself. Implementation of these innovations will require two radical shifts in thinking—one shift from a scientific viewpoint; the other from a business perspective.

From the scientific point of view, the change is analogous to delving into the physiology of networks, so to speak, rather than being concerned only with their larger anatomy. It's a shift in emphasis from form to function. It means thinking in terms of routing, for example, rather than switching; of transport, instead of transmission; of detailed network management, not just overall operations systems.

From a business perspective, the coming innovations spell the difference between just keeping up and jumping out in front. Based on the current installed investment in network operations (which is twice that of customer sales and service), the savings from OS in a typical LEC amounts to about \$4 per line per year. If the best current OS practices were fully implemented, the potential savings in 1994 and beyond is \$23 per line per year. Use of the breakthrough technology could push the savings as high as \$90 per line. These savings might

be applied to improving customer service, and to network customization, helping telephone companies meet the competition in the local exchange networks.

Four Obstacles

Four broadly defined conditions have until now prevented the embedding of OS management into a network. The first of these is the "messy, real-life factor," or how to overcome the myriad of assumptions that have evolved, over many years, into the ways that networks are currently managed. Next, there's the "funnel factor." It's a result of the vastly increasing volumes of information that are now funneled onto fiber and through switches. Then there are the "software and digital factors." As the life blood of our increasingly complex systems, software must be as "pure"—or fault-free—as possible. And digital systems, by their inherent nature, present a whole set of additional challenges. Despite the obvious challenges, however, promising new developments have been found for dealing with each of these obstacles.

Messy, real-life factors. All of the details encompassed by all of the physical and logical networks that have accumulated, from the time that Alexander Graham Bell first strung cable in Boston until today, serve to create collective obstacles—the messy, real-life factor.

Every OS existing today, regardless of its sophistication, is a descendent variation of placing 3- by 5-inch notecards in a shoe box. These variations exist, fundamentally, because critical data on the network

arrangement is external to the network itself. The “shoe box” could disappear if this fundamental information were embedded in the network elements. A number of basic assumptions have stood in the way of such change.

One basic assumption is the notion of a central office from which individual telephone lines emerge. While the assumption remains, deviations and anomalies from it abound. One such deviation is the practice of “dual homing,” in which a line from one destination is routed through two central offices. Another problematic assumption is use of the standard “tree configuration,” with a tracking system that utilizes paper, pencil and pain reliever! The latter is needed because the system provides no reliable way of keeping track of the end points.

The increasing complexities of our networks mean that soon it will be impossible to keep track of *every* connection; to anticipate *every* future need; to find *every* creative field solution that subverts standard practices. The present network arrangement, which is fixed and relies on external records kept by fallible humans, guarantees error-prone rearrangement and data-base inaccuracies.

Added to the difficulties, which are inherent in the sheer size and layers of rearrangement that have been added over the years, are the new network elements—such as data bursts—that must be transmitted. Technicians must learn new methods, from finding a misplaced protocol card (which can be considerably more difficult than tracing a dial tone), to working with fiber, which means learning about photons and sophisticated test devices. Analog to digital; copper to fiber; voice to data and video: Each of these changes have added to the complexity of the networks, themselves, and of the services they offer.

Network OS has simply not kept pace. When the systems are not able to diagnose and/or repair themselves, humans must be trained to perform the required maintenance. The increasing complexities of such maintenance work also pose a significant danger to the systems.

The network must, itself, be capable of indicating what and where its problem is, calling attention to any special characteristics of the situation. Outside plant conditions should be administered, right down to the administrative details, based on data that is provided by the network itself. Simple, yet vital, messages—like “beware of dog” and “terminal closet on third floor”—can be stored in the network element’s local memory. Technicians

could access the information without having to move to another data base. In addition, expert systems are needed to deal with network complexities, i.e. what are the failure modes for a fiber network?

When outages do occur, the tendency is for a rush of energetic and fast solutions, leaving the system in a state of artful contrivance. This state is abnormal, meaning it can’t be measured or tested by the usual methods. When the next crisis occurs, a similar response adds yet another fix, resulting in a system architecture that can approach Baroque proportions. There needs to be pre-planned responses to such outages; fixes that ultimately restore the system to its original state. Each network element should be programmed to announce its own preferred state, so technicians do not have to rely on external records.

Tom Daugherty’s^{1,2} intelligent access network controller provides an asynchronous, time-division multiplexing capability for the outside plant. It can be accessed from terminals housed in neighborhood pedestals, for example, or from connector boxes located on the customers’ premises. A practical application of Fraser’s INCON technology³, the controller establishes, in real time, the route that a call will take. It also contains adequate memory to store the needed administrative data, and can be embedded into existing fiber networks.

A team led by Daugherty^{1,2} has proven the feasibility of providing the electronics, using time-division multiplexing, needed to mix voice telephone calls with video and data bursts in the network outside plant. The controller maintains information on the end user, as well as the peculiar characteristics of the transmission, so that each telephony experience can be tailored to the specific needs of the application, without sacrificing either response time or audio quality.

Embedded knowledge is used by the controller to establish virtual circuits, which become real when a call is placed. The particular path that the call takes is determined at call set-up, based on the available physical paths, network utilization at that instant, and any special transmission characteristics that are needed. The trick is in coupling the call set-up with the administrative data, and then in signaling the route through local Asynchronous Transfer Mode (ATM) designs.

The funnel factor. This phenomenon is the result of higher bandwidths and increased capacity inherent in fiber (as opposed to copper) networks. Combined with

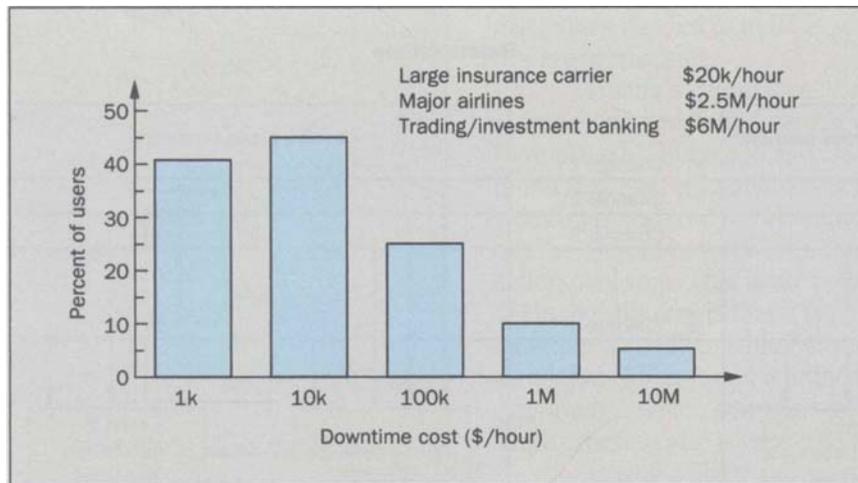


Figure 2. Network reliability is vitally important, as is shown in studies of the effects of major outages on large business customers.

faster switches, fiber permits far greater volumes of information to be carried over a given transport line, or through a particular switch. When any one element fails, therefore, the effect is disproportionately bad. And attempts to route the huge amounts of traffic around a failure point can induce congestion in other, otherwise unaffected, areas of the network.

“Focused overloads” are another interesting phenomenon related to the funnel factor. The tremendous capacity of the network encourages the telephone companies to sell new services. Problems occur when, for example, tens of thousands of customers in a given region try to order tickets to an upcoming concert at 9 a.m. on the first day that they go on sale. In another “focused overload,” AT&T has managed traffic during natural disasters by limiting calls coming into the disaster area, thereby allowing calls to fan out from the point of the crisis.

Other techniques designed to overcome the funnel factor include ring diversity and protection switching. Ring diversity permits calls between central offices to travel in either a clockwise or counter clockwise direction around the ring, to avoid a break. In protection switching, two fibers are used to back up each other. Calls are switched back and forth between the two lines, to make sure each is error-free. This practice does not increase network reliability, however, as each switch induces a burst of noise of sufficient duration to destroy sessions being created by the host computer.

Some operations or businesses, like airlines, banks, or brokerage houses, require special diversity

management to ensure against the loss of access to the network (Figure 2). In such cases, we can create two or more logical circuits. The only problem is that both fibers carrying the different logical circuits may be housed in the same sheath, vulnerable to the same back hoe. In one instance, for example, three expensive alternative cables were found grouped together inside a sheathing located on a railroad bridge. All three cables were equally vulnerable to a single, potentially disastrous, train wreck. In another potentially dangerous practice, different logical circuits have been multiplexed onto the same fiber.

In analog systems, the need to spread reserve capacity physically around the network is obvious. In physically compact, highly efficient fiber networks, however, reserve capacity is not so easily handled. However, Gitlin’s diversity-coding algorithm^{4,5} offers an elegant approach. It allows the use of error-connecting codes, like those designed into computers (called Bose-Chaudhuri codes) to obtain route diversity without the need for one-to-one redundancy (Figure 3). Use of this approach can bring the experience of managing analog systems back into play.

This state-of-the-art idea will reduce the sensitivity to single failures by allowing digital communications to be routed in different directions without requiring total redundancy. In an “n” circuit, for example, an “n+1” circuit will carry the error-correction codes. These codes will be computed for each byte from corresponding bytes in the customer’s transmission. If one circuit breaks, its customers’ data can be recreated by taking the difference

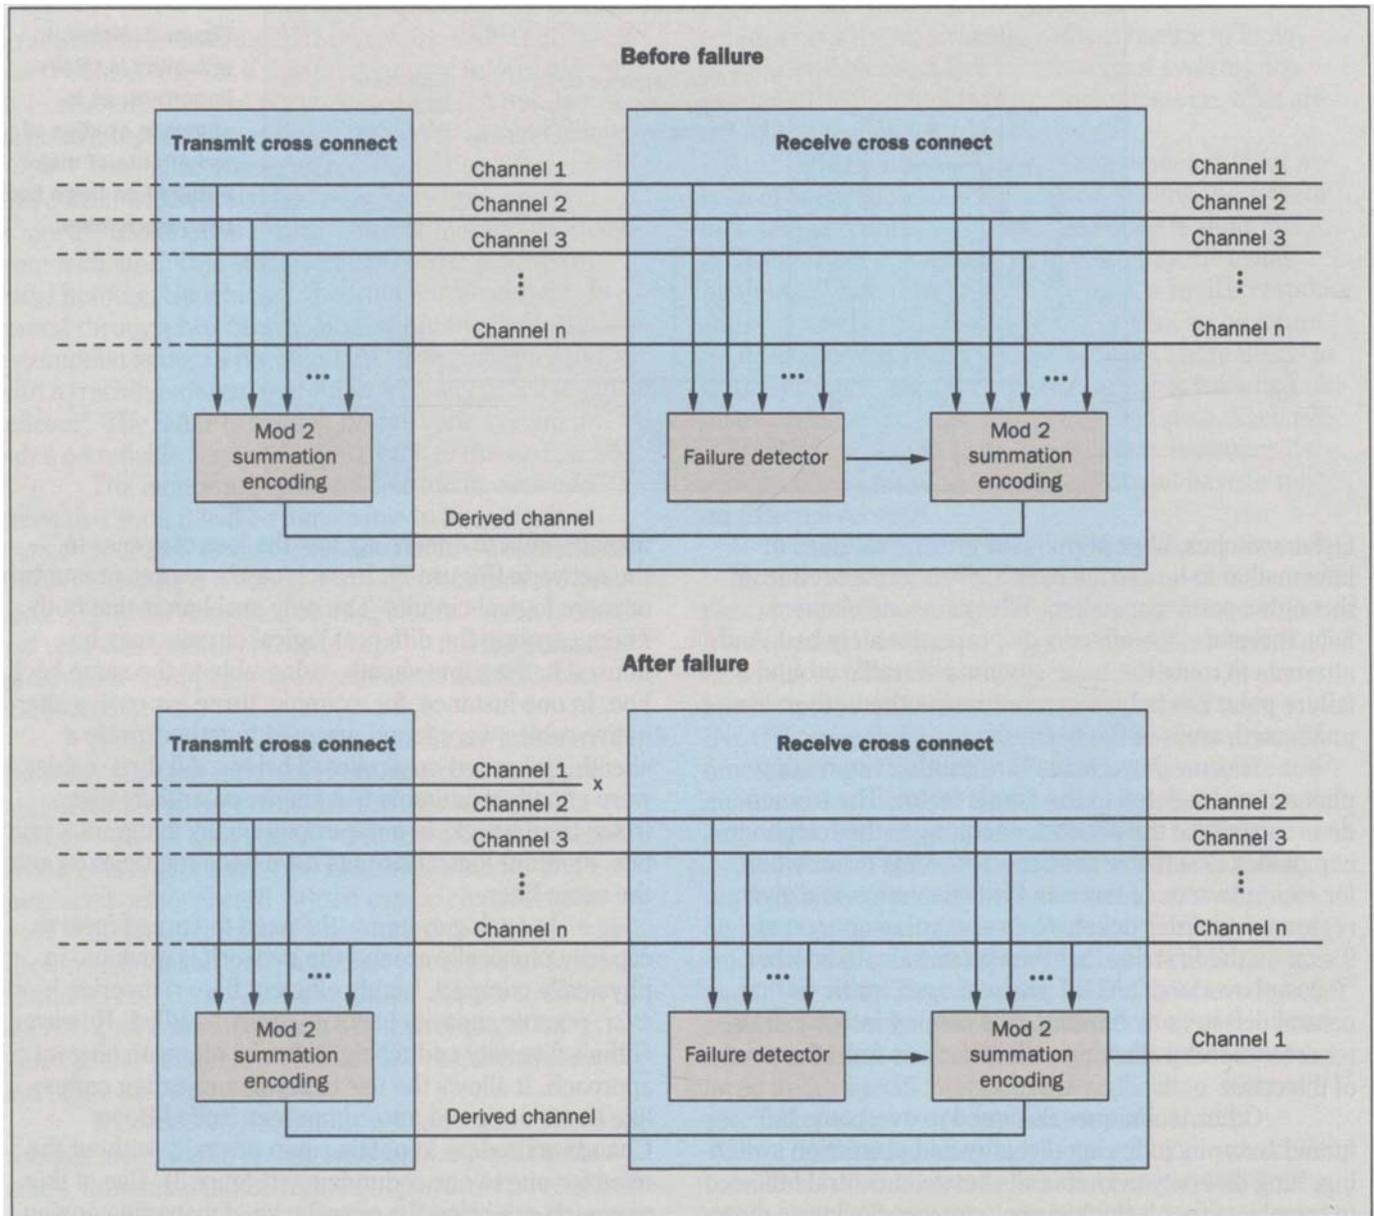


Figure 3. The 1960's New York service crisis ushered in the era of Operations Systems to help people cope with the accelerating complexity of providing better services. Increasingly, OS management is now being embedded into the network itself.

between the error-correction circuit and the remaining circuits that did go through.

Special, very fast circuits need to be designed for the general cases, and code-detection algorithms must be adapted for digital communications. Problems of data synchronization and diversity mapping still need to be worked, but this new theory holds the promise of breaking the funnel factor. If so, it would allow a return to "the

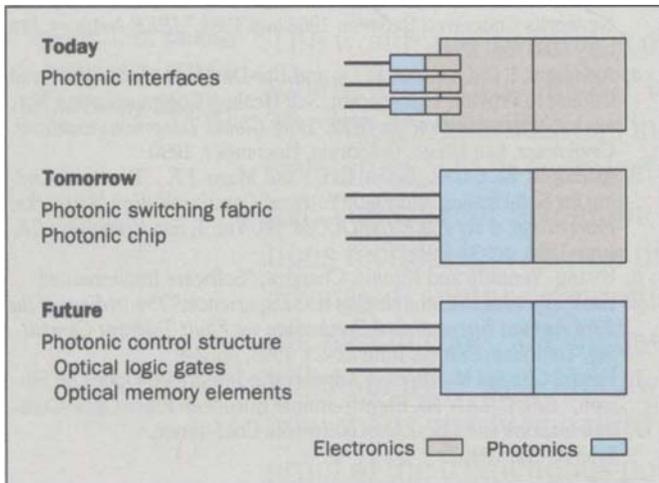


Figure 4. The role of photonics in telecommunications will increase in concert with expanding markets for video and multimedia services.

good old days” of robust networks, when a single failure did not incur the risk of catastrophe.

Networks designed with this technique will have much different management requirements. A re-emphasis on design and planning will lessen the need for self-healing networks. Although some self-healing will still be required, this design technique will provide a balance, with engineering trade-offs determining the best solution to any given situation.

The software factor. Software is our life blood and the source of profound advances. But no one can deny that it is error prone and likely to become more so with increasing complexity, as in client/server systems. Useful software is the abstraction of a problem and its solution, which is conditionally stable for the operational range that has been tested. This definition has spawned thousands of view graphs and millions of words, but programs are still prone to hangs and crashes. This conclusion is borne out by studies at Carnegie-Mellon University’s Software Engineering Institute (SEI), which paints a bleak picture of reliability. Using a 5-point scale that places reliability between 1, Ad Hoc Chaotic, and 5, Process Controlled, the SEI studies rank most U.S. software between 1 and 2. AT&T software ranked 1.8.

Reliability requires both the development of robust software, with all of the prototyping, building and testing techniques that are needed to achieve it, and the

disciplines needed to build systems whose characteristics are predictable.

Huang’s fault-tolerant software library⁶ reuses an external “daemon” (the people who name these things have played “Dungeons and Dragons” far too much) to monitor the process and detect snags. The library provides a pre-ordered set of recovery mechanisms that permit the system to work even though the individual transaction may stop. This is not new, in itself. But the beauty of Huang’s library is that it is general purpose; it can run on either a single machine or multiple machines; it is a standalone library; and it implements “n-version programming,” which compares several versions of the same program and accepts the best match.

Huang’s library allows the software to be assembled as it would with any other utility, so a special software design is not needed for this purpose. Various recovery mechanisms can be chosen. Though not yet commercially available, this library is being used, with some success, within AT&T.

The digital factor. During a lunch conversation, former Bell Labs Executive Director Bob Lucky once replied “you can’t” to the question: “How can one predict outages for digital systems?” The difficulty is in finding a theoretical basis for predicting failure by extrapolating from directly observed performance. Digital systems either work or they don’t; they respond to checks with a maddening, “healthy, healthy, healthy, dead” syndrome.

One approach to finding some basis for prediction requires analysis of analog failure modes and extrapolation, since digital systems are assemblages of analog components. Lasers, for example, are analog devices. By measuring changes in the power needed to send information with a specified error rate, one can determine when a laser will fail. Increases in power output indicate trouble ahead. In fiber, we can not measure the current on the line, but must look for micro bends, instead, which might cause attenuation.

The physics of photonics allow on-line testing without interfering with the customers’ data. The test signal can be sent on a different wavelength, or color, down the same fiber as other data without conflict and without removing the circuit from service. This opens the way for truly pre-emptive maintenance.

Today, photonic interfaces exist for connecting fiber to copper. They attach to lasers to translate photons

to electrons (Figure 4). Applying Cohen's photonic testing chip⁷ for live circuits, we can use photonic switching to inject different wavelengths on a fiber line. This technique is based on a technology called the Silicon Optical Bench, and it will be in use in 1994.

Conclusion

Just as Marshall McLuhan changed the way we think about things with "the medium is the message," we network managers need to change our perspective to "the network is the data base." Doing this means moving away from the belief that an absence of trouble reports means that the system is totally healthy, and toward a more pre-emptive form of maintenance.

With fiber, capacity is no longer a limiting factor. The added capacity must be devoted to embedding network management, as well as new services. This will permit us to manage more sophisticated networks at lower cost. Networks that are self-diagnosing, self-healing and intelligent will free up staff time to focus on customer needs and introduce new services, which will lead to happy customers.

References

1. Daugherty, T. H., et al., Patent pending, Communications Access Network Routing.
2. Daugherty, T. H. et al, Patent pending, Method and Apparatus for Establishing Connections in a Communications Access Network.
3. Fraser, A.G., "Early Experiments with Asynchronous Time Division Networks: A Retrospective Account of Three Experimental ATDM Networks Conceived Between 1969 and 1984." *IEEE Network*, January 1993, pp. 12-26.
4. Ayanoglu, I, C.-L., Gitlin, R. D., and Bar-David, I., "Analog Diversity Coding to Provide Transparent, Self-Healing Communication Networks," *Proceedings of the IEEE 1990 Global Telecommunications Conference*, San Diego, California, December, 1990.
5. Ayanoglu, E., I, C.-L., Gitlin, R. D., and Mazo, J.E., "Diversity Coding for Self-Healing and Fault-Tolerant Communication Networks," *Proceedings of the IEEE INFOCOM '90*, Vol. 1, San Francisco, CA, June 1990, pp. 95-104.
6. Huang, Yennun and Kintala, Chandra, "Software Implemented Fault Tolerance: Technologies and Experience," *Proceedings of the 23rd Annual International Symposium on Fault Tolerant Computing*, Toulouse, France, June 22-23, 1993, pp. 2-9.
7. Henry, Charles H., "Recent Advances in Integrated Optics on Silicon," *EFOC/LAN 90*, Eighth Annual European Fibre Optic Communications and Local Area Networks Conference, June 27-29, 1990.

(Manuscript approved August 1993)

Lawrence Bernstein is Operations Systems Vice President, AT&T Network Systems, and Executive Director at Bell Laboratories in Warren Township, New Jersey. He is responsible for technology supporting systems development and network management. Mr. Bernstein joined the company in 1961. He has a B.S.E.E., from Rensselaer Polytechnic Institute, in Troy, New York, and an M.S.E.E. from New York University, in New York, New York.
