

Network Security and Reliability: Emergencies In Decentralized Networks

A. Michael Noll

1. INTRODUCTION

The distinction between public and private networks has become quite blurred, as many of the papers presented in this volume, compiled by the Columbia Institute for Tele-Information, have shown. Public and private networks are interconnected and use each other's facilities to create a transparent "network of networks." As opposed to the past when the long-distance network was the sole responsibility, and was under the control of, AT&T, today's "network-of-networks" is a decentralized network in terms of control and overall responsibility. The Bell operating companies, the former independent telephone companies, AT&T, MCI, Sprint, and a host of other network suppliers and operators, all have this responsibility. The entire telecommunication network is in a constant state of change as it is reconfigured to meet the needs of its users. The use of private branch exchanges has been on the decline as business customers return to Centrex service provided by local telephone companies. The truly dedicated private networks of large businesses are being replaced by virtual networks provided by the many common carriers.

All telecommunication users -- private and public, business and residential -- take telecommunication for granted. That is, until something goes wrong and the network goes down! Then suddenly we all again learn that telecommunication is essential to today's economy. Network failures indeed can have catastrophic results, and such failures must be prevented and minimized through adequate attention to the security of the network and the individual networks that comprise today's "network of networks." With competition, the reliability of tomorrow's network is at greater risk, according to a report of the National Research Council.¹

This paper discusses the issue of network security and reliability and how these topics relate to emergencies in decentralized networks. The technology of a modern telecommunication network is described. Emphasis is given to the role of common channel signaling and the extra vulnerability that it creates along with the means for protection and service restoration. Technological, procedural, and cooperative solutions to ensure the reliability of the network are discussed.

2. NETWORK FAILURES

Three major failures occurred in AT&T's long-distance network in 1991. These three failures resulted in lengthy outages, hours in duration, that disrupted telecommunication

service in the affected major metropolitan areas and thus have focused attention on the reliability of today's telecommunication networks. All three of the failures in AT&T's network had a strong human component and were caused by management failures in procedures and policies. However, it is not only humans who cause network failures but technology as well; here, though, ultimately, human error is to blame.

Preceding years saw service failures caused by software errors and glitches in the signaling system that controls the overall operation of the network. One of these signaling system failures affected AT&T early in 1990. Other signaling system failures affected two local Bell companies on both the East and West coasts (interestingly, both local failures occurred on the same day). Long-distance carriers other than AT&T have also experienced various network failures, although the much smaller size of the total traffic carried by these carriers has resulted in much less impact and publicity.

These network failures serve to remind us that today's service and information economy is highly dependent on telecommunication, and any network failure of even a few hours strongly disrupts the conduct of business and even our personal lives. Two of AT&T's network failures in 1991 affected telecommunication service for two major airports thereby affecting airline service. Furthermore, today's technology has such a large scale that even a "small" failure has huge consequences. Over 30,000 telephone circuits are routinely carried over a single strand of optical fiber, one-tenth the diameter of a human hair. A single switching machine in a long-distance network can handle over 100,000 one-way trunks.

Network failures can occur for a variety of reasons. In most cases, the cause of a failure is accidental. However, network failure caused by deliberate acts or sabotage should not be ignored when planning for all contingencies. Deliberate acts of harm could come from either an external or an internal source, and a disgruntled employee could easily do much harm. The hardware could malfunction causing a network failure, although most key network hardware is duplicated and this redundancy greatly reduces hardware failure as a source of network failure. For example, the central processors and other key components in most switching machines are duplicated. Switching machines are operated by computer programs, or software, that is subject to a wide variety of errors, or bugs, when the software was written, ultimately by humans. Thus, even though hardware might be duplicated, usually the same software controls the hardware and thus a single software error could collapse the duplicated hardware. Hardware malfunctions can also affect software and lead ultimately to network failure.

Clearly, there are a wide variety of sources of network failure. The following tutorial is the basic architecture and workings of a modern telecommunication network in intended to help the non-technical reader to understand more fully these sources of network failure along with the means needed to protect against failure.

3. NETWORK TECHNOLOGY: A TUTORIAL

The network carries the signals generated by customer traffic from one place to another over a wide variety of transmission media, including optical fiber, microwave radio, and copper cable in the form of twisted pairs or coaxial. Various paths are switched together until a complete circuit between source and destination is formed to carry the signals generated by the customer traffic. The customer traffic usually consists mostly of speech signals generated by telephone instruments, but facsimile and data signals are also carried over the network.

Transmission and switching are the two major technologies involved in the switched telecommunication network. The various real and virtual paths that must be created to carry telecommunication traffic over the switched network must be set up and maintained during the duration of the call, and then dismantled at the completion of the call. Signaling is the aspect of telecommunication that deals with controlling the network to create, maintain, and dismantle these paths.

In the past, signaling was done over the same switched network that carried the customer traffic, and only when the complete end-to-end circuit was completed would the distant telephone be rung. Alternate routing to avoid congestion along the way was difficult. Costly voice circuits were connected even though actual talking had not yet begun. This type of signaling was cumbersome, and the network could not be easily and quickly reconfigured. Call completion times were lengthy, and costly facilities that normally would carry actual telecommunication signals were being used to carry signaling information. Furthermore, fraudulent users could generate false signaling information to avoid being billed. Today's telecommunication network separates the transmission and switching of customer traffic from the network control functions of signaling, as depicted in Fig. 1. This separation results in increased network security, reliability, and efficiency while also giving a higher level of service to network users along with new "intelligent" services. Signaling today is accomplished over a separate data network that carries only the signals needed for signaling, namely, the data signals that control the operation of the network. Today's telecommunication network thus can be envisioned as two separate networks, as shown in Fig. 2: one that carries solely the data signals needed to assign the real and virtual paths that carry customer-generated telecommunication traffic and another that provides the actual paths over which the customer-generated traffic is carried. This approach to signaling is called Common Channel Signaling (CCS). CCS allows for more efficient operation of the network and also benefits the consumer through faster call set-up times and through novel network services, such as 900 and 800 numbers. At the local level, CCS allows local telephone companies to transmit information identifying the telephone number of the person calling, although the display of this information to the called party has become quite controversial because of the concerns of many consumers over the privacy of their telephone numbers. The specific standard for CCS used in the United States is called Signaling System 7 (SS7). The signals sent over the CCS network are packet switched. With packet switching, a short data message carries information about its source and destination. The destination information is examined by switches along the way, and the packet of data is stored and forwarded gradually to the final destination. This is unlike the circuit-switched network that carries customer traffic and in which individual circuits are maintained for the entire duration of a call. Packet switching is particularly appropriate for the short, bursty messages that comprise most data signals.

A block diagram of the basic structure of a modern long-distance telecommunication network is shown in Fig. 3.² Switching machines at Service Switching Points - SSPs (formerly called Action Control Points and now Action Points - ACPs by AT&T) assign and switch customer traffic over transmission paths between the SSPs. If a particular transmission path is fully occupied, alternative paths can be assigned dynamically in a non-hierarchical manner. The traffic carried over the network is usually digital, and a large number of such digital signals are combined together through time division multiplexing to share transmission facilities. A basic telephone channel in digital form requires 56 kbps or 64 kbps. Twenty-four such digital signals are multiplexed together to create a single DS1 (or T1) signal at 1.54

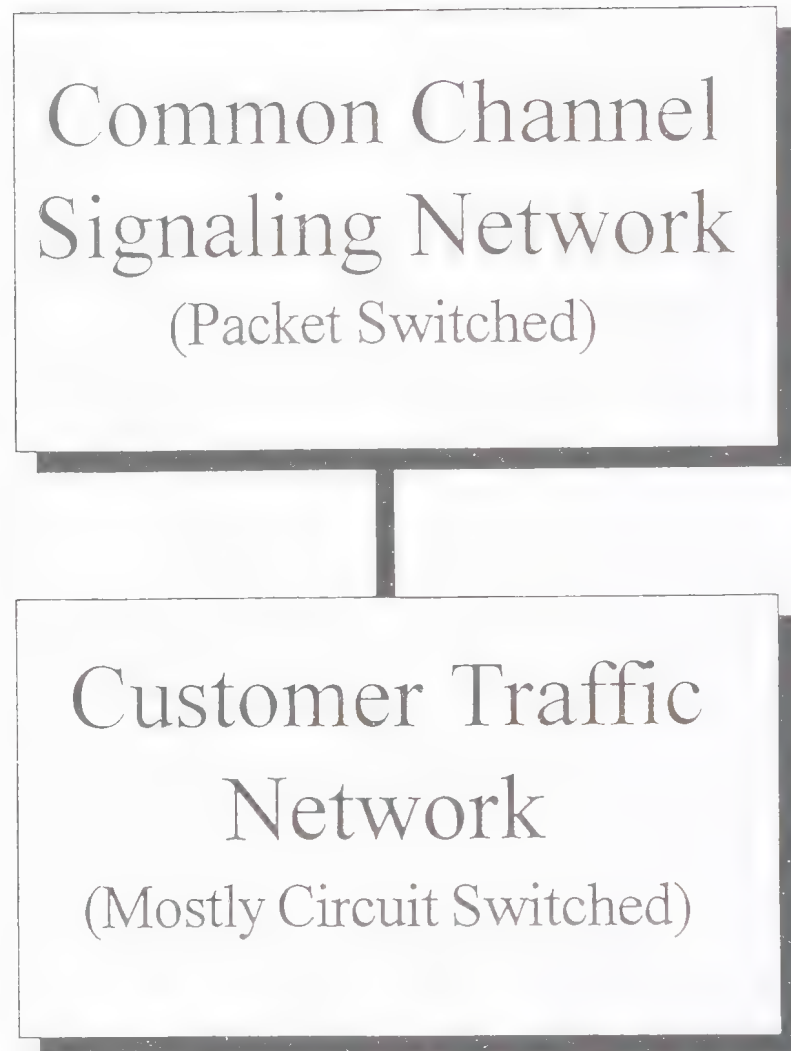


Figure 1

A modern telecommunication network can be conceptualized as two separate networks working together. One is the network that carries customer voice, data, and image traffic. The circuits carrying customer traffic are switched over this network. The other is a network that carries solely the signals needed to control the operation of the customer-traffic network. This control network is called the Common Channel Signaling (CCS) network and it carries short data messages that are stored and forwarded to their destinations in the network -- a technique known as packet switching.

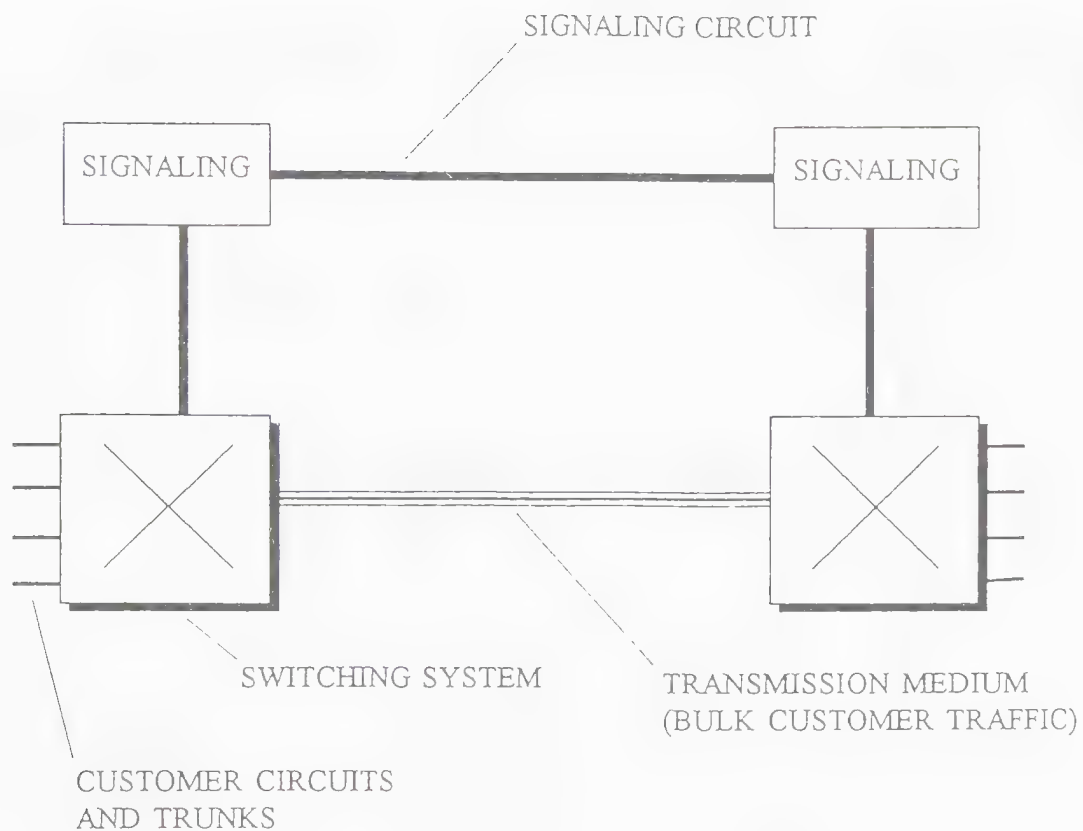


Figure 2

The switching systems responsible for switching the circuits carrying customer traffic are controlled by information sent over a separate Common Channel Signaling (CCS) network. The data sent over the signaling circuits is considerably less than the customer-generated traffic. The customer traffic in a long-distance network is carried in bulk over a variety of transmission media, including optical fiber, microwave radio, and copper wire.

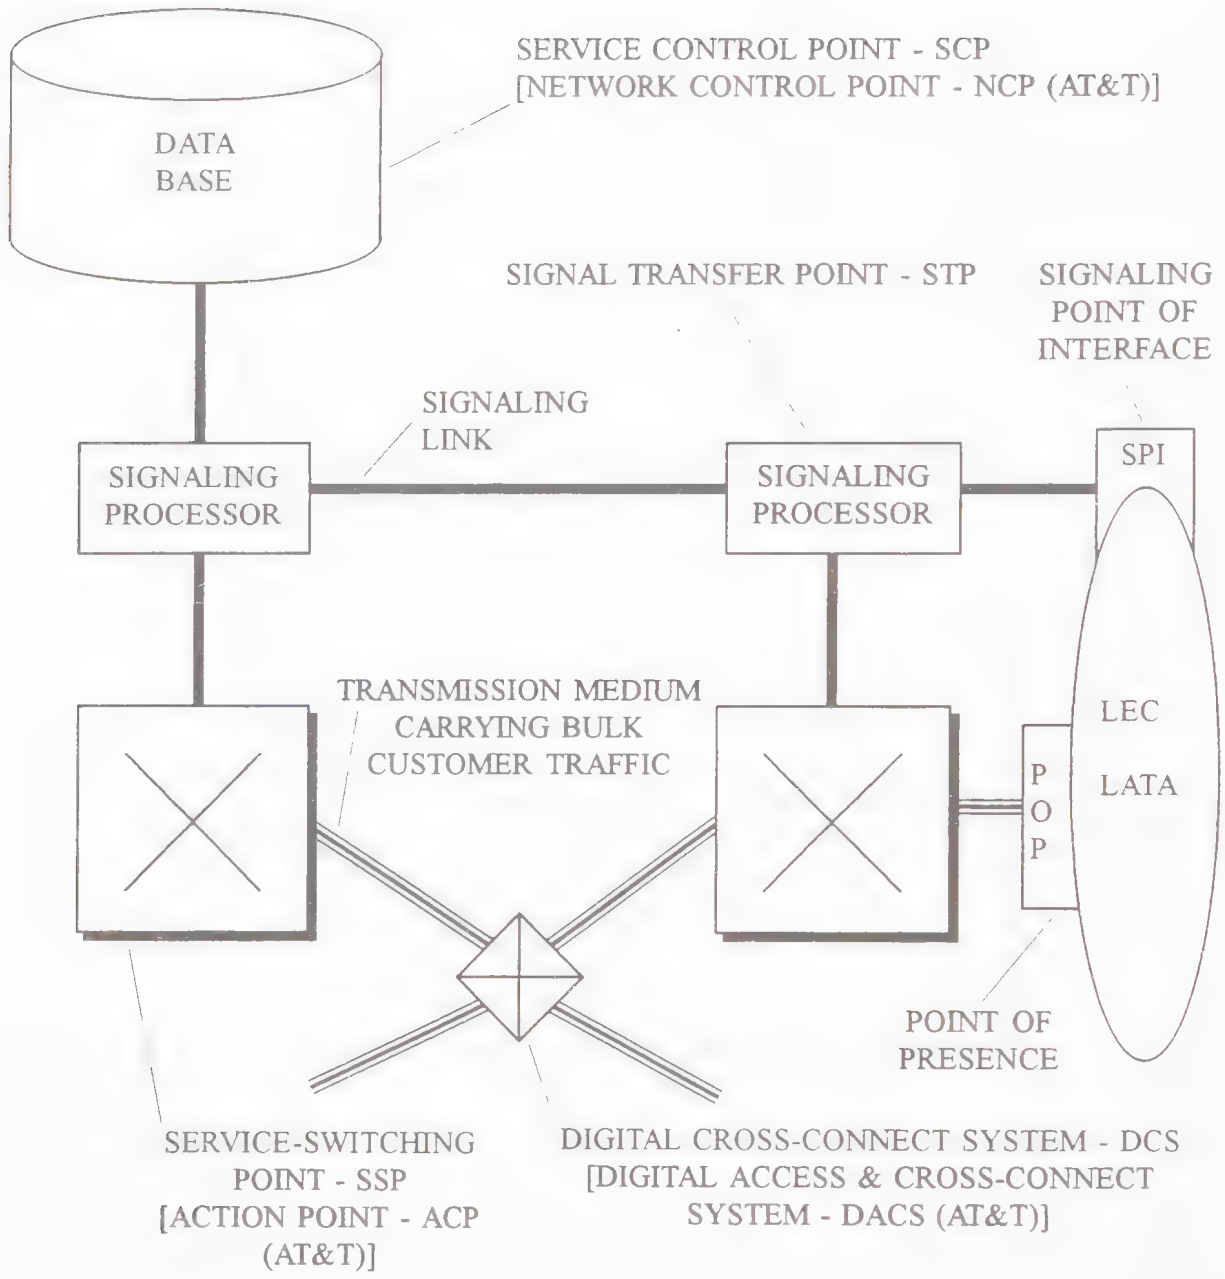


Figure 3

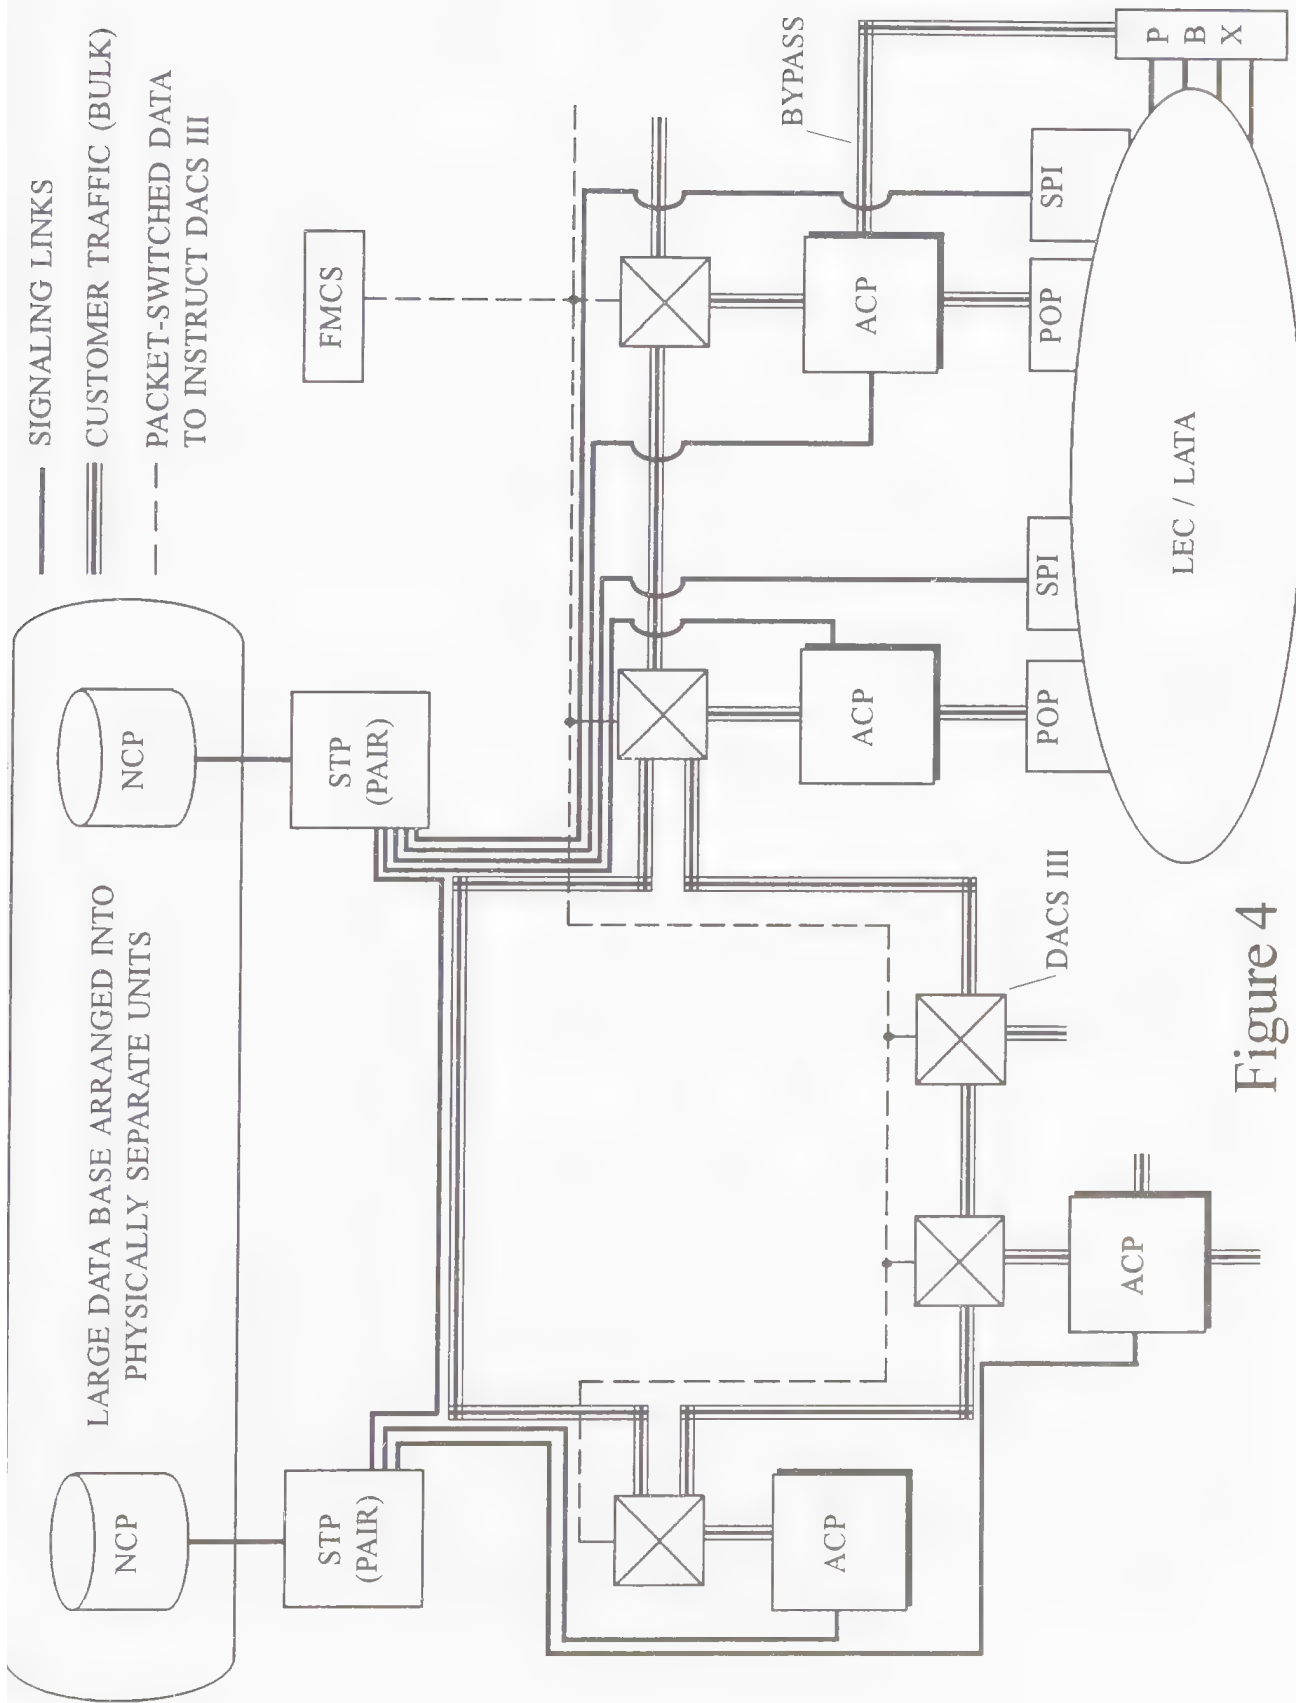


Figure 4

Mbps. Higher-level digital signals are at the so-called DS3 (or T3) level, operating at a rate of approximately 45 Mbps and containing 672 voice channels. Only at the local level where customers are connected to their local serving office is the traffic mostly analog carried over twisted pairs of copper wire.

A separate network carries the signals needed to control the assignment of transmission paths so that customer signals can be carried from the calling party to the called party. This signaling network interfaces with the switching machines at the SSPs. The signaling signals themselves are switched through the signaling network at nodes called Signal Transfer Points (STPs). Customer and routing information is stored in the signaling network at nodes called Service Control Points - SCPs (called Network Control Points (NCPs) by AT&T). Signaling links connect STPs to each other and to the SSP switching machines.

Most long-distance networks carry vast amounts of digital traffic consisting of tens of thousands of multiplexed digital channels over optical fiber. It sometimes is necessary to be able to reconfigure these higher-level digital signals. In the old days, such reconfiguring of telephone circuits would be done at patch-cord panels. The electronic version of such patch-panels for multiplexed digital signals are called Digital Cross-connect Systems (DCSs) (called Digital Access and Cross-connect Systems - DACS by AT&T). They are located in the transmission media so that higher-level digital signals can be re-routed from one transmission path to another. Such re-routing is performed for such purposes as an emergency, routine maintenance, or major network re-routing. The customer traffic carried over long-distance networks connects to the local networks operated by the Local Exchange Carriers (LECs) at so-called Points of Presence (POPs). The LECs are responsible for providing local service within Local Access and Transport Areas (LATAs). A Signal Transfer Point (STP) in the long-distance network connects to the local network at a Signaling Point of Interface (SPI).

3.1. An Example: AT&T's Network

The overall architecture of AT&T's long-distance network is depicted in Fig. 4. It is described as an example of a modern telecommunication network with advanced technological and service features and considerable use of technology to ensure reliability and fast restoration of service in case of a network failure. Fig. 4¹ shows that some large business customers could have their own direct access to the long-distance networks operated by the Interexchange Carriers (IXCs) and thus bypass the LECs for long-distance service. Most local networks are converting to CCS; AT&T's long-distance network was the first to use CCS.

AT&T operates a large, switched network providing long-distance service to residential and business customers. This network is controlled by software so that a wide variety of customized services can be offered. Such a software controlled network is quite flexible and can be configured to provide customized services that mimic the dedicated private-line services of the past. For those customers who demand dedicated facilities, AT&T offers "special services." AT&T's network supplies standard switched service, 800 service, 900 service, and a variety of customized business services. AT&T conceptualizes its network as having three layers: (1) bulk transmission at the DS3 level, including DACS III operating at the DS3 level to reconfigure the network as required; (2) a switching fabric consisting of No.4 ESS™ machines; and (3) the common channel switching network.

Considerable redundancy exists in the AT&T network as a means to achieve reliable service. Individual switching machines at the ACPs, CCS switches at the STPs, and databases

at the NCPs all include duplicated multi-processors and other duplicated hardware to achieve redundancy in their internal operation. Some ACP switching machines operate in pairs so that one could handle the traffic of the other in case of an emergency. NCPs are paired in primary and back-up configurations. CCS switches at the STPs are also paired with each operating at no more than 50% capacity so that any one of the pair can handle all the traffic controlled by the other. Transmission paths are replicated to achieve redundancy, and three physically-separate routes carry East/West traffic across the country.

There currently are 22 STPs (operating in 11 pairs) and about 300 NCPs in AT&T's signaling network. Customer traffic is carried over a network consisting of over 100 No.4 ESS™ digital switching machines at the ACPs. Calls are switched in a dynamic non-hierarchical manner through no more than one intermediate ACP. The mainstay transmission medium of AT&T's network is optical fiber, and digital microwave radio is used to achieve redundancy and also to reach low-traffic locations. Nearly 140 million calls are handled each day over AT&T's worldwide network. AT&T's CCS network uses from one to sixteen 56 kbps data circuits to carry the CCS signals between STPs. Compared to the tens of thousands of circuits that carry customer traffic between ACPs, the CCS network is quite "thin." However, even so, AT&T states that its CCS network is the largest private packet switched network.

AT&T's digital cross-connect switch is called a Digital Access and Cross-connect System (DACS). There are in the order of 200 DACS IIIs in the AT&T network operating at the DS3 level. In case of an emergency, the DACSs can be commanded to switch and re-route traffic around faults in the network. Software, called RAPID, detects any failures, assesses the spare capacity needed to restore service, issues the appropriate commands to the DACSs, and tests the digital circuits both before and after the reconfiguration. The commands can be sent either over terrestrial fiber paths or, as a back-up, over a communication satellite and received at the DACSs by VSATs (Very Small Aperture Terminals). AT&T calls this automated system for the fast restoration of service FASTAR™; presently, the first DS3 channel is restored in under one minute and additional DS3 channels are then restored every few seconds. So that there is more than one way to reach any location along the network, transmission routes in the AT&T network are configured in the form of large loops. DACSs are located at the points where one loop touches another and also in links from one loop to another.

4. LOCAL-INTEREXCHANGE DIFFERENCES

By their very name, long distance networks cover great distances and are very intensive in terms of costly transmission routes and facilities, particularly since these routes are usually duplicated to ensure redundancy. Common channel signaling was originally invented by AT&T as a means to ensure the more effective allocation of transmission channels. Long distance networks are relatively thin in terms of switching. Local networks are quite different. Local networks do not cover great distances but do have to switch many calls. Local networks thus require much more switching than long distance networks. Redundancy usually is less available at the local level; for example, there is only one local loop and one local office serving a customer. The problems of effective use of costly transmission facilities are much less of a problem at the local level and, not surprisingly, CCS has been much slower at being introduced at the local level. One motivation for CCS at the local level is the new

"intelligent" services that it facilitates along with compatibility with interconnection to long distance networks.

5. THE SIGNALING SYSTEM: A NEW VULNERABILITY

We have seen how in the past the circuit for a long-distance call was connected progressively according to information sent along the lines. Today's telecommunication networks are controlled through a far better system: common channel signaling (CCS). We have seen how CCS sends data signals over a dedicated data network, and these data signals are received and interpreted by CCS computers and network switching machines that then allocate trunks and transmission facilities to carry the actual telecommunication traffic of the customers.

AT&T's and other carrier's networks use CCS, as do many local telephone companies, and soon the entire public telecommunication system of the United States will utilize CCS. Private corporate networks will want direct access to CCS to give them more direct access and control over various network features and functions. All this will certainly be to the benefit of the consumer and business users. However, there will also be a risk that some software glitch could transmit an erroneous signal or traffic indication that would collapse the entire network thereby bringing telecommunication to a total halt in the country. This situation is somewhat similar to what happened in the electric power industry. In the distant past each power station was independent and connected only to its local customers. Over time, power stations were connected to an electric power grid, or network. This interconnection created a situation in which some glitch could collapse the whole grid, as demonstrated by the Northeast power failures of the past. Computerized control systems with appropriate human intervention now minimize the risk of such catastrophic failures in the electric power industry.

6. NETWORK SECURITY AND RELIABILITY

Transmission media need to be protected. The cables that carry thousands of twisted pairs of copper wire in the local loop are pressurized both to keep out moisture and also to enable any leaks or breaks to be detected. Unlike copper wire, optical fiber is very difficult to tap and thus inherently offers security. Transmission routes are often duplicated to offer security in terms of alternate routing in time of emergency.

Today's telecommunication networks are switched and controlled by computers. These computers are a prime source of network vulnerability to failure. Although in most cases, the vulnerability is accidental, deliberate attempts to sabotage the computers in order to collapse the network can not be ignored. The security of these computers thus is essential. A fair amount of research effort has been developed over the past decades to the general topic of computer security, and the results of this research are applicable to the computer aspect of network security.

At the most obvious level of security, telecommunication gear and equipment facilities must be physically secured. Appropriate physical access safeguards must be installed. Without physical safeguards any person intent on harming the network need only "pull the plug" on a switching machine. At another level, the maintenance of appropriate audit trails is essential, both to determine when some deliberate attack is underway, and also to help in performing

a "post mortem" to determine the causes of a specific network failure. The software that controls the switching machines and the signaling system is a major source of network vulnerability. Computer programs are designed and written by fallible humans and hence are subject to error. A simple software error might have catastrophic effects on the network and might even propagate through the network via the signaling system. Telecommunication programs are usually very large, and their production requires the considerable time and effort of many programmers. This human effort must be appropriately managed and all errors eliminated. However, errors will occur, either because of unanticipated situations or because of simple human fallibility. Automatic techniques to verify that a computer program does only what is intended and is free of error are being researched but are not yet foolproof. Meanwhile, the only real solution is to expect errors but minimize their effect through appropriate safeguards and testing.

7. SAFEGUARDS

The very technology that has increased the scope and impact of network failures also plays a major part in protecting telecommunication network. Computers monitor the operation of networks and automatically notify human network supervisors in the case of some abnormality. Presented with appropriate information, the human supervisors then instruct the control system to take appropriate action. Human intervention continues to have an essential role since otherwise the automated monitoring system might take inappropriate action and escalate the final harm to the network. Technology can only be trusted so far, in my opinion.

7.1. Gateways: Protection Against Outsiders

The existing common carriers at all levels -- local exchange and interexchange -- fully realize the essential importance of protecting their individual networks against failures from both accidental and deliberate sources. In essence, the existing carriers have created a security fence around their operations. As they interconnect their signaling systems, the security fence broadens to include others equally secure. The real threat is when others currently not part of the existing league want direct higher-level entry and access to the network and signaling system. Pressures from large business customers for such access are already starting. Large businesses and cellular operators want access to the customer databases stored at the NCPs. One possible solution to the potential threat posed by giving access to the signaling system to outsiders is the use of so-called gateways. All outsiders would be required to go through a gateway to gain access to the signaling system. A gateway would be a pair of STPs functioning as a boundary between the common channel signaling system and the outsiders. The gateway would be an overseer that would examine and monitor signaling messages sent by the outsiders to be certain that no inadvertent or deliberate harm to the network would occur. Clearly, research and development is needed to develop such gateways and to assure that they function as intended in protecting the network from harm.

7.2. Network Testing

The "network of networks" resulting from many interconnected long-distance, local, and private networks is a complex affair. Many of these networks are already controlled by Common Channel Signaling (CCS), and those that are not, are quickly migrating toward CCS. The "network of networks" thus will be interconnected not only in terms of the

transmission and switching of customer signals but also in terms of the signals that control the overall operation of the networks. This interconnection at the signaling level creates the need to be absolutely certain that any hardware and software additions or changes will not in any way harm any network to which it is attached.

Both intra-network and inter-network integrity must be assured. The Exchange Carriers Standards Association sponsors a Network Operations Forum with broad representation from the telecommunication industry. Guidelines are developed for the testing, maintenance, and installation of access networks and the interface between Local Exchange Carriers (LECs) and Interexchange Carriers (IXCs.) Bellcore developed an Internetwork Interoperability Test Plan for the Association's members to test and evaluate different network failure scenarios. In the first phase, test facilities at Ameritech, AT&T, Bellcore, Northern Telecom, NYNEX, and Sprint will be interconnected to test the response of existing systems to various failure scenarios.

Many interexchange carriers (such as AT&T, MCI, and Sprint), local exchange carriers, and vendors already have extensive testing programs before installing any new hardware or software. Cooperation between these bodies would make very good sense as a way to achieve an industry-wide testing program before any new hardware or software is installed in the network. Other than Bellcore, existing standards organizations or even the FCC might act as the facilitators to create the necessary forum for cooperation between competitors. Clearly, such cooperation will be challenging given the highly proprietary nature of any new product or system. Will a vendor be willing to release some new piece of hardware or software for testing by its competitors? Should some industry-wide neutral body be able to perform the needed tests and also safeguard the proprietary nature of a new product or system?

7.3. Procedural Solutions

In February 1992, a consortium of about fifteen major telecommunication carriers signed an agreement of mutual aid to restore service in the case of "critical disruption to their telecommunications networks supporting the New York City Metropolitan Region" (Agreement dated February 18, 1992). The agreement is the first of its kind and stipulates the procedure to be followed in time of a network emergency affecting high-capacity transmission facilities. The details of the agreement were determined by the Mayor's Task Force on Telecommunications Network Reliability, chaired by the Commissioner of the New York City Department of Telecommunications and Energy.

In the event of the loss of critical telecommunication facilities affecting New York City, the affected carrier notifies the Commissioner. If the failed facilities cannot be restored within two hours, the Commissioner is involved in declaring an "Emergency." The other carriers have then agreed to work with the affected carrier to make facilities available, for "reasonable and customary out-of-pocket expenses," to restore service. If more than one carrier is affected by an emergency and sufficient capacity is not available to restore the failed facilities, the Commissioner has the authority to allocate available facilities across the failed common carriers. The surviving carriers do not have any obligation to provide facilities if they do not have any excess capacity available. The Commissioner has no real authority to force cooperation, but the Agreement is a strong statement of intent to cooperate under the realization that a catastrophic network failure ultimately affects everyone and the common good of New York City is best served by cooperation at a time of emergency. The New York City Telecommunications Department also has published a catalog of various telecommunication services that may be useful to businesses in an emergency, such as call

forwarding, or that may be used to create redundancy to avoid catastrophic failures or limit their effects.

At the Federal level, the National Communications System (NCS) is responsible for ensuring the integrity and responsiveness of telecommunications from the perspective of national security and emergency preparedness.⁴ Military tactical transportable microwave radio gear, communication satellites, cellular phones, and high-frequency radio are some of the technologies that can be used in an emergency, such as an earthquake or hurricane to provide emergency communication until the public switched network and other conventional systems are restored. The Modification of Final Judgment gave Bellcore responsibilities related to National Security Emergency Preparedness (NSEP). In meeting these responsibilities, Bellcore serves as a central body to coordinate the efforts and activities of the BOCs related to NSEP. Service restoration, resource allocation, disaster response, the development of operational plans for NSEP, and joint government/industry planning are some of the specific activities coordinated by Bellcore.

8. CONCLUDING THOUGHTS

Private networks are today defined by software and rarely utilize dedicated facilities. Private and public customer traffic are carried together over one network. The network security requirements for different kinds of customers might very well be quite different. Should the network be designed to offer the highest level of security to all or is it possible to offer different levels to different customers on the same network? Might higher levels of security be offered to some customers at the expense of others? These and other interesting questions of public policy need not to be forgotten.

In the time before divestiture and telecommunication competition, AT&T had the sole responsibility for the operation and integrity of the Nation's telecommunication system. The administrative control of the network was centralized. However, the administrative control of today's "network of networks" is fragmented across many competing common carriers and is truly decentralized. One can only wonder whether some form of centralized administrative control and oversight over today's decentralized network is needed and, if so, who should perform that function. Some interexchange carriers have network operation centers where the entire operation of their networks can be instantaneously monitored. Displays of the status of various transmission routes and switching systems, of the traffic being carried, and of the signaling system are used to help human operators supervise the network and take appropriate action if needed in time of an emergency. Similar centers would help all carriers -- interexchange and local exchange -- monitor the status of their networks and assist restoration in time of emergencies. One even wonders whether some national center to monitor the entire network on a more global basis would be useful and whether the operation of such a center would be a meaningful role for the FCC -- particularly at a time when regulation is being reduced. The very technology that has made the network appear more vulnerable clearly can safeguard the network too. However, in the end, people and human error will most likely be the cause of most network emergencies and failures. The use of technology to protect against human shortcomings continues to be a challenge.

REFERENCES

- A Guide to Contingency Services*, New York City Mayor's Task Force on Telecommunications Network Reliability, January 1992.
- Bodson, Dennis and Eleanor Harris, "When The Lines Go Down," *IEEE Spectrum*, March 1992, pp. 40-44.
- Growing Vulnerability of the Public Switched Network*, National Research Council, Washington DC: National Academy Press, 1989.
- Holste, D.J., "Creating A Network That Heals Itself," *AT&T Technology*, Vol. 5, No. 2, 1990, pp. 42-47.
- Noll, A. Michael, *Introduction to Telephones & Telephone Systems (Second Edition)*, Norwood MA: Artech House, Inc., 1991.

ENDNOTES

¹ *Growing Vulnerability of the Public Switched Network*, National Research Council, Washington DC: National Academy Press, 1989.

² The switching nodes in a network are known as Service Switching Points (SSPs), or as Action Points (ACPs) using AT&T terminology. The signaling computers are at nodes known as Signal Transfer Points (STPs). A large data base containing routing, billing, and customer information can be accessed by the Common Channel Signaling system. This data base is known as a Service Control Point (SCP), or as a Network Control Point using AT&T terminology. The signaling information is sent over signaling links. The transmission media carrying customer traffic can be reconfigured using a Digital Cross-connect System (DCS), or a Digital Access and Cross-connect System (DACS) in AT&T terminology. The DCS is able to reconfigure digital circuits carrying 50 Mbps of customer traffic. A SSP connects to the local network operated by a Local Exchange Carrier (LEC) at a Point of Presence (POP); a STP connects to the local network at a Signaling Point of Interface (SPI).

³ The overall architecture of AT&T's long-distance network operates at three levels: (1) transmission media, that can be reconfigured by DACS IIIs carrying bulk customer traffic between ACPs; (2) No.4 ESS™ switching machines located at ACPs; and (3) a CCS network consisting of STPs and NCPs. In case of a failure in the network, the appropriate DACS IIIs can be reconfigured to carry traffic around the fault. The information necessary to reconfigure the DACS IIIs is sent over an AT&T packet-switched, terrestrial data network from a Facility Monitor and Control System (FMCS). A communication satellite system with Very Small Aperture Terminals (VSATs) at each DACS III serves as a back-up to this terrestrial data network. The signaling links in AT&T's CCS network consist of from 1 to 16 56 kbps, two-way, packet-switched data circuits. A business customer with a Private Branch Exchange (PBX) can bypass the LEC's network and connect directly to AT&T at an appropriate ACP. STPs operate in pairs so that any one of the pair can take over all the work of the other in case of an emergency. Transmission paths and processors in the No.4 switching machines are duplicated in AT&T's network to increase reliability. Local traffic enters AT&T's network at at least two POPs and two ACPs, again to offer alternative routes in case of an emergency.

⁴ Bodson, Dennis, and Eleanor Harris, "When the Lines go Down," *IEEE Spectrum*, March 1992, pp. 40-44.