This paper describes the IBM System Communications Division network and, using that network as an example, discusses some of the practical problems associated with providing computational and communication services to remote and local user communities.

Evolution of a laboratory communication network

by R. S. Moore

In 1975 a study group examined the data processing and communication requirements of the IBM laboratories in Europe and recommended that an organization be established to implement and support a network which would provide computing facilities to all sites on a coordinated basis. The European Laboratories Computational Network (ELCN) was established to fulfill that role. At the time the group reported, the implementation of a sister network had already begun in the United States.

Inevitably there have been changes in the organization of the laboratories since 1975, and these changes have required changes in the computational networks. This paper discusses the two IBM System Communications Division (SCD) networks and the coupling between them. The major sites served by the networks are shown in Figure 1.

The SCD network interconnects with other IBM internal networks for access to systems and data of other divisions. Primary interconnections are to the IBM European Information Systems (SWITCH) Network in Europe and to the Corporate Consolidated Data Network (CCDN) in the United States. In addition, all sites are nodes of the IBM internal job network,[1] which connects more than 200 systems at development and manufacturing locations.

Figure 1 Major IBM locations served by the SCD network

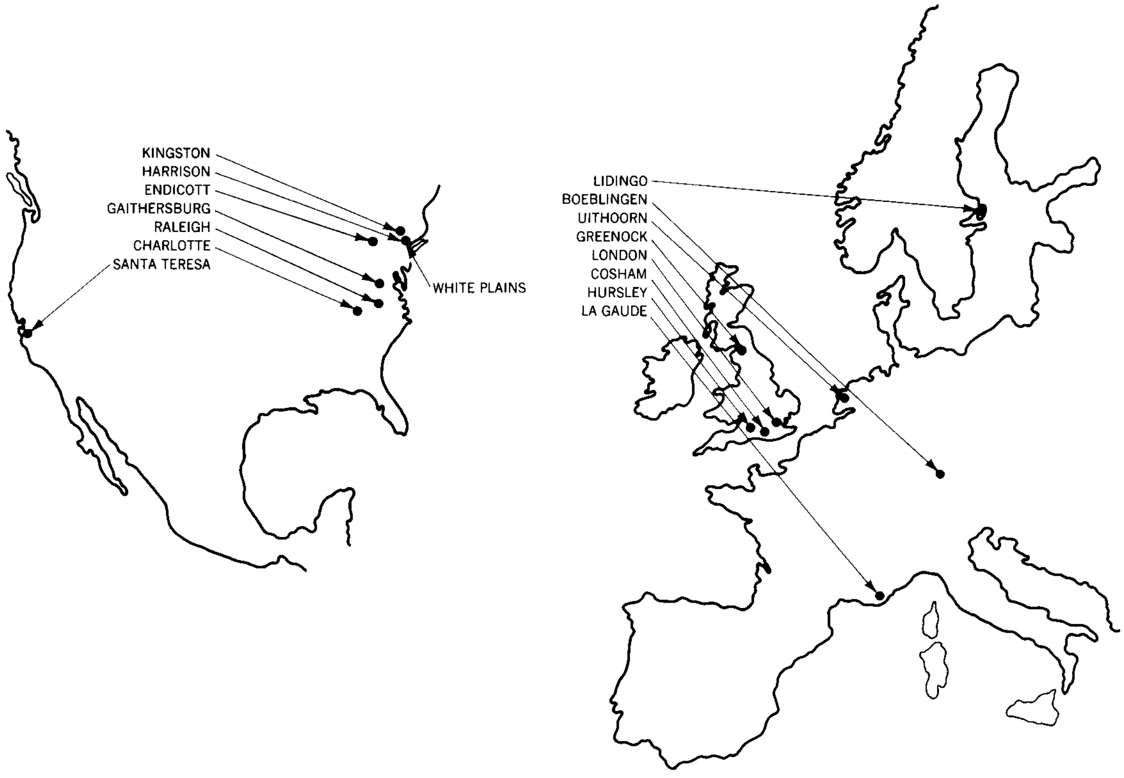

The paper discusses some of the reasons for establishing the network, outlines the network's basic organization and functions, and discusses in detail features of the interactive terminal network and techniques used for the measurement and management of service levels. Finally, it predicts the next steps in the evolution process, which will involve the addition of new function and migration to Systems Network Architecture (SNA), IBM's standard communication architecture.

Basis for the network

centralized development support

As technology advances, it becomes necessary to provide development systems of ever increasing power and complexity to assist at all stages in the logical and physical design of large-scale integrated circuits. In general, these systems require several large processors, such as the IBM System/370 Model 168, and skilled support teams. This type of system is an obvious candidate for centralization. Economic and skill requirements make it impractical to provide systems at every SCD site.

Engineering support systems for SCD users are now consolidated at two sites, Hursley in the UK and Kingston, New York. The centralization of engineering support has resulted in direct cost savings and standardization of techniques, and it has simplified the transfer of technology and projects between development groups.

In the United States, SCD has centralized support of other facilities. The Kingston site provides APL, text, and management information systems. Administrative system support for American SCD locations is centered in Raleigh, North Carolina. This centralization has resulted in economies of scale related to better utilization of hardware and skilled manpower.

common data-base access

The SCD network was developed to provide users at remote sites with access to a variety of application systems. Once installed, the network could provide many additional services with very little effort. Within each development location are planning groups that need information, such as market forecasts, from other sites. Prior to the installation of the network, access to such data was difficult, and delays of several days were common. Today, however, planners, programmers, and engineers have access to data on almost any processor in the division at whatever level it is required (for example, they can obtain documentation, search a data base, or update a data base). As users gain experience, they are able to modify their working methods to get maximal advantage from the facilities.

improved communication

A file transfer system, described below, makes it possible for network users to send messages with delivery times of a few minutes, or even to "talk" with one another using terminal-to-terminal facilities. Thus users can exchange documents for review, make comments or additions, and return the changed documents. There is a growing awareness of this capability among development groups, and this secondary usage may become the most important benefit of the network as its role in electronic document distribution develops.

improved mobility

Programming and engineering personnel often must attend meetings at other SCD sites. With the network, a programmer or engineer has access to his home system from the site he is visiting, so in many cases he can reduce the amount of material carried on a trip. Moreover, the network provides an excellent means of giving demonstrations at remote sites, and it allows traveling personnel to assist in diagnosing any problems that arise in their home systems.

improved security

The existence of the network as a separately managed entity ensures that consistent standards of data security can be applied to the transmission facility. Cryptography is used to protect sensitive data as it passes through the network.

ease of extension

Network facilities can give interactive terminal users access to any application and provide for the return of printed output. The facilities greatly simplify the problems associated with providing service to new locations. Installation of a single terminal cluster and a remote-job-entry work station is all that is required to provide basic service.

Network elements

The American and European networks are linked by two leased 9600-baud circuits, one between Hursley and Kingston, the other between Hursley and IBM's Santa Teresa Laboratory in San Jose, California. The networks have a single, coordinated operational strategy, and technical plans are closely aligned, although deviations do occur because of differences in the communications environments of the United States and Europe. This alignment ensures that development effort and operational experience can be shared easily.

File transfer system

It is often necessary to transfer large volumes of data (typically output listings) between network locations. This service is provided by the Network Job Entry (NJE) facility[1,2] of the MVS operating system[3] and the VNET facility[4] of the Remote Spooling Communications Subsystem[5] for host systems that use the VM/370 operating system.[6] An extensive discussion of the job transfer network is provided by Simpson and Phillips.[1]

All host systems are nodes on the network and can communicate with each other either directly or via intermediate store and forward. SCD nodes can also communicate with nodes owned and operated by other divisions. The design and configuration of this file transfer network is constrained by turnaround requirements and traffic volumes. (File transfer is not generally sensitive to short failures.) Links between systems use binary synchronous communication (BSC)[7] protocols and normally operate at 9600 bits per second (bps), although they range from 2400 to 56 000 bps. The only nonstandard facility is the optional use of Huffman data compaction[8] on high-cost routes where the additional CPU time required is justified.

Interactive terminal network

Figure 2 shows schematically the interactive network and its primary interconnections. A typical (primary) laboratory has 250 to 500 terminals installed at any one time, 50 to 100 of them using

Figure 2 Schematic diagram of the basic European-American network

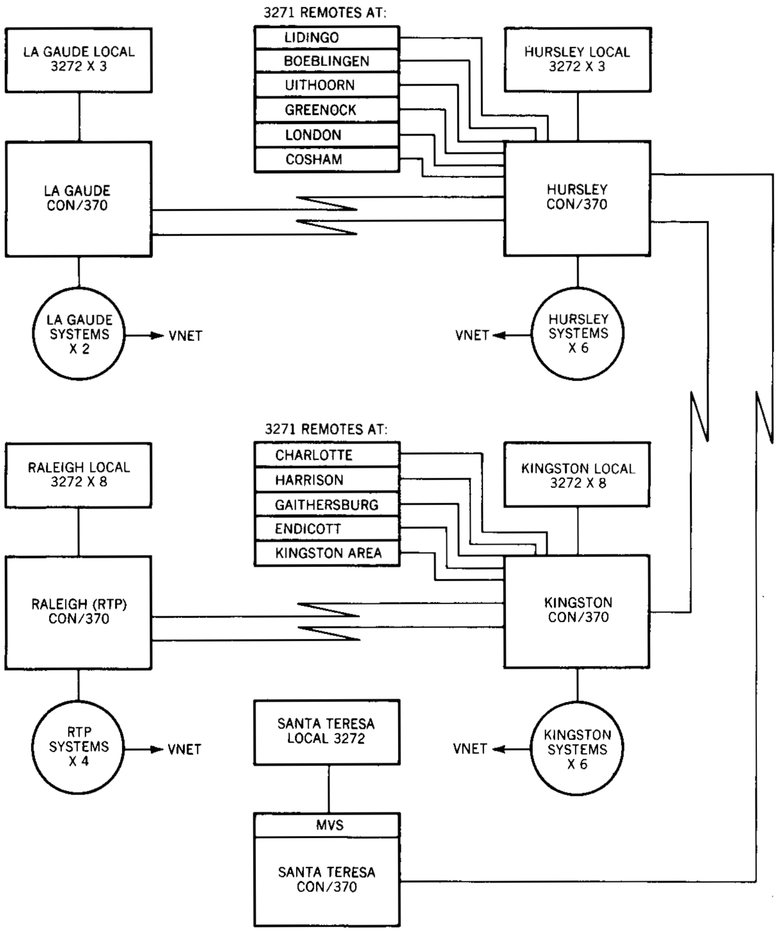

network facilities for either local application access or remote application connection. Internodal capacity normally is provided by multiple 9600-bps links, and BSC protocols are generally used. A number of basic network components provide the necessary facilities and meet service-level requirements. In general, the examples used in this section relate to the European environment, but most of the points illustrated are common to the United States and Europe.

telecommunications facilities

A 97 to 99 percent rate of availability can be expected on a typical international line (operating at 9600 bps), with about one failure per week. Figure 3 shows the distribution of failure duration on nine international lines over a period of 100 days. There were 218 failures, the average failure lasting 105 minutes. There were 55 failures in excess of the average. To reach an acceptable level of availability, the network is supported by controllers and technicians with a full range of test and diagnostic equipment. Despite

Figure 3 Distribution of failure duration on nine international lines

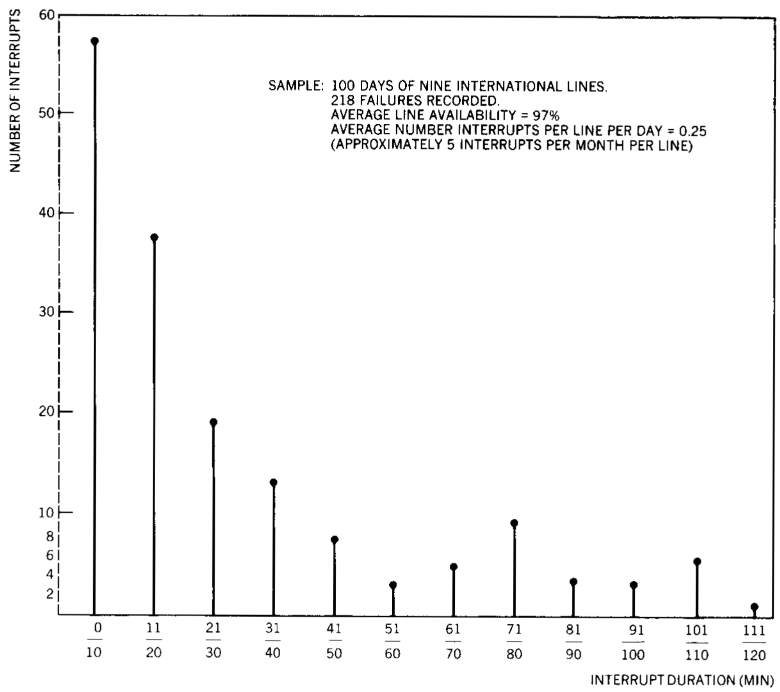

this technical support, it would be difficult to meet network service-level criteria by using a single line, so two diversely routed lines are made available wherever possible. This duplication effectively eliminates the line as a factor in analyzing the availability of the network's most critical services.

hardware

At major network nodes, a stand-alone processor provides terminal support and switching facilities. Current configurations make use of System/370 Model 155's in Hursley, Kingston, and La Gaude, and a Model 145 in Raleigh. To improve availability, the configurations have built-in redundancy (for example, in disk drives, operator consoles, and communication controllers), and the use of nonessential peripheral devices such as card readers, punches, tape drives, and printers is kept to a minimum.

software

The interactive terminal network is configured using a number of software components. Figure 4 shows how each component relates to the physical hardware. The software concentrator, CON/370, is a networking package originally developed to meet the national and international requirements of the IBM European Information Systems network. CON/370 interfaces in a well defined manner to a terminal controller, and it handles the routing of message segments through one or more network nodes to any available application interface. It provides extensive real-time status, trace, diagnostic, and statistical aids.

Figure 4 Relationships of network components to physical hardware

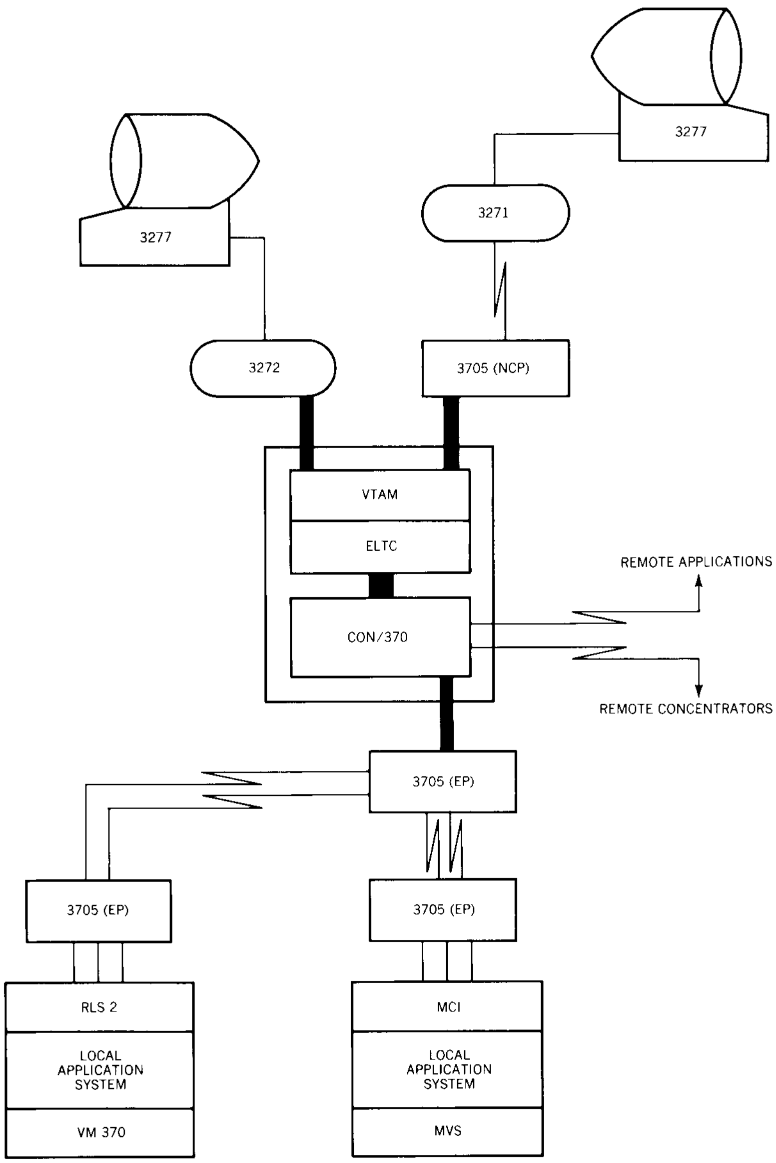

Transmission between CON/370 nodes is via multiple BSC channels which operate so that the load is spread dynamically and so that failure of a single line does not disrupt any terminal session. SCD-written extensions to the basic CON/370 package provide for full Huffman data compaction and encryption on a selective basis. CON/370 provides limited alternate routing facilities when the predefined primary connection fails. The design aim was to provide a high degree of reliability, optimal circuit utilization being a secondary consideration. Dynamic routing based on the instantaneous loading of the network is not provided.

terminal controller

The network's terminal controller, called ELTC (European Laboratories Terminal Controller), operates as a VTAM[9,10] application program. Any terminal user who requires access to a network application establishes a normal session with ELTC, which provides the support function necessary to control the terminal and pass the data to the host system that contains the application. For more than a year, ELTC has been in full production use in Hursley, La Gaude, and Kingston laboratories that support local IBM 3277 display stations attached via 3272 cluster control units. In addition, production use of ELTC has begun with terminals connected via an IBM 3705 communication controller and using the Network Control Program (NCP),[11,12] with a BSC connection to an IBM 3271 remote cluster control unit. Some features of ELTC are:

- *Session establishment.* The user is presented with an application selection panel as soon as the terminal is powered on. Entry of a valid application name establishes a connection to a currently active network application. Once a session is set up, the user need not be aware of ELTC, but he can take advantage of ELTC to provide functions not provided by some application systems. On receipt of a request for connection, ELTC communicates with CON/370 to determine whether the application is available. If not, the terminal user can employ other features to determine why and obtain such information as the expected repair time. If the application is available, ELTC acts as a relay program and passes data between the terminal and the application, using CON/370 to provide routing and transmission facilities.
- *Session monitoring.* ELTC handles every segment of terminal input and output. It is therefore an ideal facility for the accumulation of statistics. Statistical information relating to network-added delays, host system interaction times, and input and output data rates are accumulated and are available as real-time performance indicators or as long-term trend indicators. Examples of data gathered by ELTC are used in the discussion of service levels, below.
- *Escape facility.* Early experience with ELTC demonstrated that occasionally a user would wish to break a session temporarily in order to access a special network facility. To meet this requirement, ELTC provides an escape facility. The user simply depresses a programmed function key, his application session is suspended, and ELTC presents a series of options. Currently implemented escape options enable a user to print a copy of the information on his screen, control the allocation of a printer, and access another application or change terminals.
- *Segment concatenation.* Many application programs have evolved from a keyboard/printer version and do not take advantage of the 3277 display station's full screen management facilities. In some cases, screens are filled with many small,

independently written segments, causing excessive flicker, slowing the response, and wasting time by causing execution of multiple VTAM paths. When ELTC is sending output to a screen, it combines all message segments that have arrived for that screen and constructs a single output message.

MVS concentrator interface

ELCN was designed to support several different application systems, such as TSO,[13] APL,[14] ATMS,[15] and CICS,[16] operating with any of several access methods, including VTAM,[10] BTAM,[17] and TCAM.[18] When the network was designed, it was apparent that user pressure would force substantial change to accommodate new releases, and that customized interfaces with special hooks in standard code could not be maintained.

The only factor common to all applications was that they supported locally attached 3277 display stations. Therefore it was decided to implement a general-purpose interface to the MVS operating system, based on the concept of a virtual, locally attached 3277. The principle of this MVS concentrator interface, or MCI, is illustrated in Figure 5. As far as the application system and access method are concerned, all the network devices appear as locally attached 3277's. No change need be made to application code. A single change to MVS causes the interception of I/O operations to and from those addresses defined as network terminals. The change has proved reliable and easy to maintain, and the inefficiency associated with this type of interface has been more than offset by the simplification it provides. (Extension of this concept to support start/stop devices became necessary and is being implemented.)

Figure 5 Schematic diagram of the MVS concentrator interface (MCI)

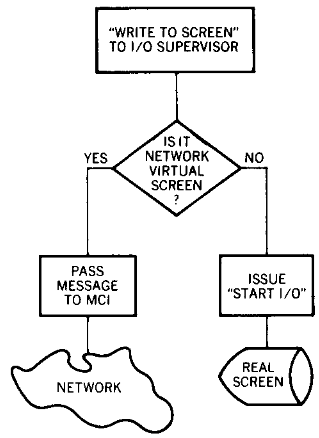

The design aim was to provide an application independent interface, and that aim was achieved in the steady state. However, real networks do exhibit failures, and it is necessary to introduce application dependent code to deal with failure and reconnection. MCI includes application monitor programs to deal with failures. Advanced features within MCI protect the user against network failure by recovering the session from the point of disconnection. MCI interfaces to one or more CON/370 nodes via any number of BSC lines.

In many locations a particular terminal normally is attached to a particular host system and requires only occasional access to another system. In this situation, there are obvious advantages to attaching the terminals to this host instead of to ELTC on the stand-alone concentrator system. Therefore ELTC is run as a VTAM application on an MVS host, and the user gains access to network applications via this "terminals out" interface between ELTC and MCI.

VM/370 Remote Laboratory Support

A BSC or channel-to-channel interface to one or more CON/370 nodes is provided by Release 2 of VM/370 Remote Laboratory Support (RLS2), an internally developed program. It provides similar function to the general-purpose interface to MVS. A network terminal user can log on to a remote VM system and use his terminal as a console. RLS2 also supports the VM DIAL command for access to a virtual machine. It supports the "terminals out" facility, so that a terminal attached to a VM system can access a remote application via the network.

integrated bulk data transmission

An integrated bulk data transmission system provides for the transmission of very large disk or tape files. It has two interfaces to the network: a full-duplex Execute Channel Program (EXCP) interface and a CON/370 interface which allows low-priority bulk data to be interleaved with high-priority terminal traffic. In this way, use is made of the spare line capacity available during off-peak hours, with no significant degradation in response time.

Service-level measurements

In 1977 SCD set up a task force to determine the most important service-level parameters. As a result, attempts were made to implement the service levels recommended. This section covers in detail the network-related parameters of most significance and gives examples of results achieved.

Basic service levels are defined for an arbitrary period accepted as each location's prime shift. The period normally defined is from 0830 to 1730 hours at the terminal site. Measurements of service-related parameters are made on a 24-hour basis, but commitments are made on the basis of the prime shift.

The availability of a particular service to an end user requires the simultaneous availability of a number of network elements. A series of elements is considered to form a "service aggregate." For a user in La Gaude to reach TSO in Hursley, for example, the network elements are defined as ELTC (La Gaude), CON/370 (La Gaude), lines (Hursley to La Gaude), CON/370 (Hursley), MCI (Hursley), and TSO (Hursley). The concept is shown in Figure 6.

Availability reporting currently is heavily dependent on CON/370, and unfortunately, failing CON/370 nodes may lose contact with sections of the network, so that statistics at any one node may be incomplete. A solution that requires a daily merging of statistics from several nodes has been considered, but rejected. A replacement method under development will require each network component to keep a record of its own availability with a resolution time of one minute, and to assume the responsibility of transmitting this data to one or more statistics-gathering points. The data can then be used to deduce the availability of each network element and any specified network aggregate.

Figure 6 Network elements between La Gaude and Hursley

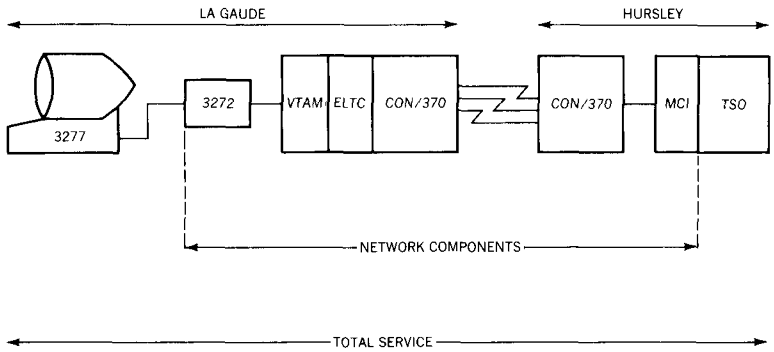

number of significant interrupts

In a rapidly changing environment it is apparent that changes to hardware and software can cause failures. Discussions with users indicate that there is justification for discounting very short network failures, provided that the user's session can be recovered without loss of data. Thus there is an incentive for network management to devise techniques for rapid recovery. The problem is in defining the term *very short*. At this point we are using a time of eight minutes or less; most users will wait about that length of time before leaving. Provided the user's copies of ELTC and CON/370 remain operational during a failure, it is possible to provide application status information. Service-level objectives have been set in terms of the number of interrupts longer than eight minutes, and progress has been made in providing automatic recovery mechanisms and smooth operating procedures to beat this deadline. Examples are:

- CON/370 will trap all abends, dump to disk, and restart within four minutes without operator intervention.
- The MVS concentrator interface (MCI) application monitor programs will reload upon failure.

There is still room for improvement in handling VTAM and OS/VS1 failures. Almost any problem with these systems results in a failure that takes more than eight minutes to detect, resolve, and correct. Experience to date shows that failures are infrequent, but that they typically last about 20 minutes.

ten-week rolling averages

The European section of the network currently aims at availability in the range of 96 to 98 percent, representing one or two hours of down time per week. Obviously a single long failure makes it impossible to meet the objective for that week. It is also clear that if a network is properly designed, long failures should not occur frequently. To provide a basis for trend analysis and a single figure of merit for network users, availability is reported on the basis of the average of the most recent ten weeks, as shown in

Figure 7 Network availability: (A) percentage of availability, Hursley TSO to La Gaude; (B) number of interrupts, Hursley TSO to La Gaude

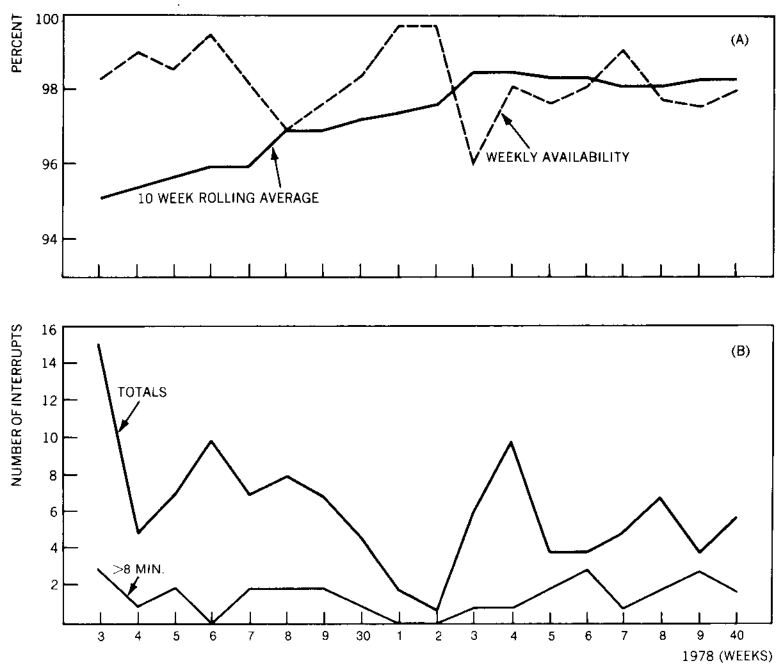

Figure 7. We have found that the trend line of this ten-week rolling average is a reasonable guide to service as perceived by the user. In addition, it provides a useful focus on excessive rates of change or relaxations of operational procedures.

Interaction time

There is little agreement on the ideal response time at a terminal. The definition of response time often is based on the interval between depression of the *enter* key and the display of the first output segment—the "primary response time." This measurement often is unrelated to user satisfaction, as the system is programmed to give an immediate (and trivial) response, such as movement of the cursor (in TSO), or a change of a status message (in VM), to signify that processing has begun. The user actually cares about the time from input to delivery of the final segment of the system's response—the "transaction time." For some applications the response is a single segment, in which case primary response and transaction time are equivalent. In other cases, multiple segment outputs are common. (For example, TSO for engineering users generates an average of two outputs per input, and APL provides an average of three outputs per input.) In some cases the segments are generated in rapid succession, and the "join together" code in ELTC recombines them before actual delivery.

Examination of actual interaction times indicates that the distribution of transaction times has a very long tail (see Figure 8). A

Figure 8 Distribution of transaction times

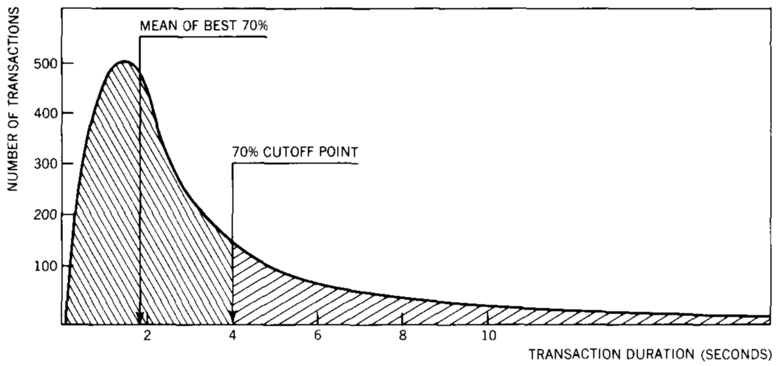

Figure 9 Average of best 70 percent of transactions, Hursley TSO to La Gaude users, in 15-minute periods during typical day

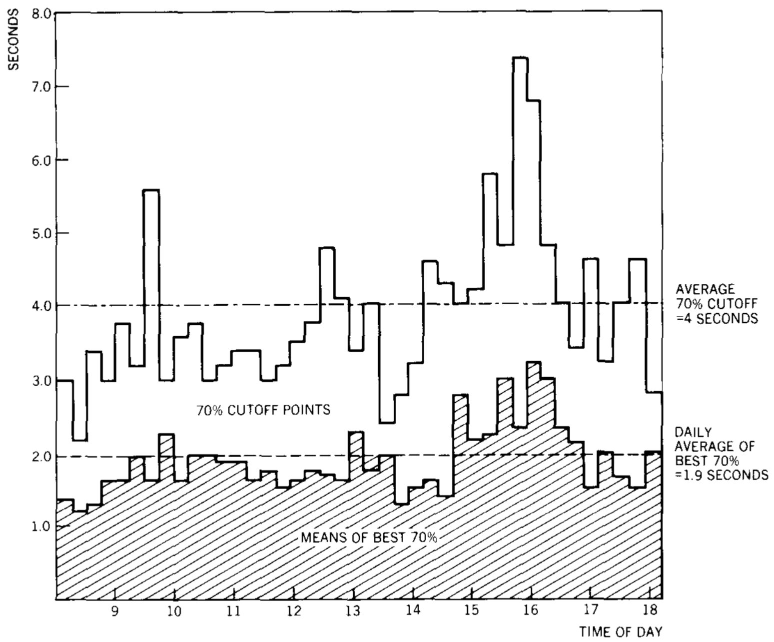

measure of average transaction time, therefore, is not a good performance indicator, as it is heavily weighted by the presence of (perfectly acceptable) long-process-time transactions in the tail region. After some experience of the real operational environment, we decided to take the average of the best 70 percent of transactions in any 15-minute period as a performance indicator, and to record the value of the transaction time at the 70 percent cutoff point. This data is available on line to terminal users and is summarized for trend analysis on a weekly basis. It serves as a sensitive measure of performance problems. Typical results are shown in Figure 9.

Figure 10 Measurement of network-added delay

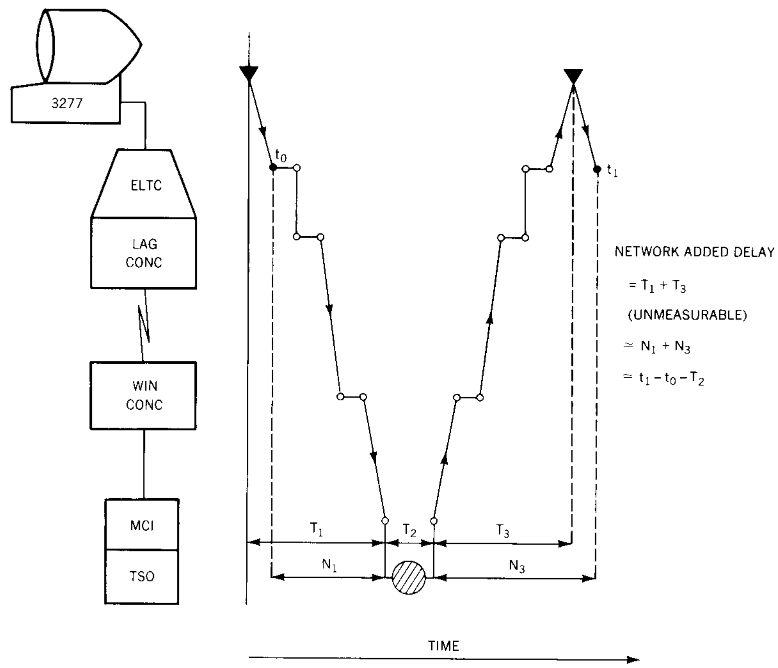

network-added delay

When the transaction time exceeds the targeted limit, the cause must be isolated to either the network component or the host system. As a measurement of network performance, attempts are made to measure the actual network delay. The mechanism is illustrated in Figure 10. Basically ELTC records the time when it detects an input from a terminal (t_0). This input is then transmitted across the network to the application interface. The application interface records the interval between passing the input to the application system and receiving the first segment of output (T_2). This host process time is appended to the message and transmitted back to ELTC. ELTC then displays the message on the screen and records the time at which VTAM signals successful display (t_1). ELTC is then able to compute the total round-trip time, subtract the host process time, and thereby calculate network-added delays.

This facility has only recently been activated; previously we had been using the primary response time as a quality indicator. Now we intend to use the same technique used for transaction times—that is, the 70 percent point for average and cutoff—on the basis that the longer delays are caused by exceptionally large screens and would distort the average. This 70 percent point can be moved easily, and we will be investigating the effect of alternative cutoff values. ELTC accumulates the response time statistics by application and separates local and remote terminals.

The future

The SCD network is continuously changing; new users and applications are being added at a rate which is dangerously near the maximum that can be accommodated without a decline in service levels. Of the major activities planned for 1979 and 1980, some of the more exciting are described briefly below.

workload sharing

The SCD network connects locations in Europe, the United States, and Japan. The nature of the development process is changing, so that more computing resources are required during the normal day shift. The consequence is pressure for many locations to install additional processors to meet the peak prime-shift demand. These systems would be lightly loaded during the second and third shifts. To avoid this situation, pilot projects are being set up to take advantage of the time differences between Hursley (UK), Kingston (East Coast United States), and Santa Teresa (West Coast United States). At 0900 hours in Hursley it is only 0100 hours in Santa Teresa, and the third-shift capacity of the Santa Teresa machines can be made available for prime-shift use by Hursley users. Similarly, second-shift capacity on the Hursley system can be made available for prime-shift use at Santa Teresa.

It is apparent that the projects must be chosen carefully; large CPU-time-consuming users with a small data-transfer requirement are most suitable. As distributed processing and distributed database techniques improve, and the cost of international bandwidth falls, the feasibility of this type of operation will also improve. The potential savings (in terms of hardware, space, and system support skills) are enormous. The problems related to operational procedures, reliability, and response times present a challenge which we believe can be met.

electronic document and facsimile distribution

The existing network can be used to transfer documents between skilled users, but initiation of such an operation requires familiarity with TSO or VM as well as careful planning with the recipient. Work now being completed is intended to provide an easy interface to the network for all users, especially secretarial staff using "mailbox" techniques. It is expected that this facility will be extended to support documents prepared on many IBM products, such as the recently announced 3730 distributed office communication system.

In addition to the improved exchange of text, work is in progress to attach digital facsimile scanners. Early implementation will simply capture the facsimile data as a normal file (a typical page will represent about 50 000 bytes) and use the file transfer system to send this data to one or more receivers. This service appears easy to implement and will allow the user to communicate diagrams and other uncoded data.

migration to SNA

The network provides communication between multiple systems using BSC line protocols, internally developed special-purpose code, and System/370 machines to provide communications facilities. This approach is expensive in terms of software maintenance and hardware. Recently announced standard software that makes use of Systems Network Architecture (SNA) concepts[19,20] offers a better alternative. We plan to move from today's implementation to VTAM/ACF[21] and SDLC line protocols.[22,23]

As the necessary hardware and software become available for internal use, a phased migration will begin. This migration will ensure that we are able to take maximal advantage of new terminal products, distributed systems, and host system applications without wasted internal development effort. Migration involves a significant amount of host application conversion, and it must be accomplished without major disruption of service. We estimate that two years will be required to complete the process.

network management

As the network grows, we must face the problems of its management. Time zones present peculiar difficulties because support available in a location is reduced in off-shift hours. Effort is being put into the development of software to simplify operational control and monitor the status of applications, equipment, and lines. A "network manager" system, to begin operation in 1979, will attempt to monitor the availability and utilization of the network in real time and allow the point of operational control to move as time changes. Efforts will also be made to follow the progress of jobs and files across the network, and to provide the user with a convenient way of reporting problems in a uniform way. Experience has taught us that new problems arise daily, and the ability to detect and isolate them is essential to providing acceptable service.

Conclusions

This paper has attempted to introduce some of the factors that influence the design, implementation, and management of an international network that encompasses several time zones. The technical problems of transporting bit streams between points are generally soluble. Practical experience has resulted in a management concept that separates, as far as possible, the management of individual host systems from management of the telecommunications network. This concept makes it possible to treat the network as an independent subsystem with well defined parameters which can be tracked and controlled in terms of cost and performance. Technical problems are easy to overcome when compared with issues of user education and operational liaison. Operational experience with the SCD network has begun to have significant influence on its future. The role of the network as a provider of remote computing facilities is entering a consolidation phase. The future is in the expansion of its role as a powerful communications facility and its role as a vehicle for balancing resources across time zones.

ACKNOWLEDGMENTS

The design and implementation of the network have been a team effort. I cannot hope to list all those who have contributed. I would particularly like to thank Allan Weis for providing the initial driving force for ELCN, Anthony Proudman (Manager of ELCN) and Barbara Myers (SCD Director of Information Systems) for their continuing support of the network's developing telecommunications role, Tom Gilbert and his staff (Communication Facilities, Kingston) for their many contributions to the success of the network, and my own ELCN staff in Hursley and La Gaude for doing the real work.

We have borrowed software where we could: CON/370 and the integrated bulk data transmission system from IBM Information Services Ltd., of Cosham, UK (I particularly thank Bert Barker for his help), and RLS1 from the IBM Field Engineering Center in Chicago. Ralph Snyder and Lyn Long have been particularly helpful in providing the code for RLS2.

CITED REFERENCES

1. R. O. Simpson and G. H. Phillips, "Network job entry facility for JES2," *IBM Systems Journal* **17,** No. 3, 221-240 (1978).
2. *Network Job Entry Facility for JES2 Version 5 Program Product General Information*, IBM Systems Library, order number GC23-0010, available through IBM branch offices.
3. *OS/VS2 MVS Overview*, IBM Systems Library, order number GC28-0984, available through IBM branch offices.
4. E. C. Hendricks and T. C. Hartmann, "Evolution of a virtual machine subsystem," *IBM Systems Journal* **18,** No. 1, 111-142 (1979).
5. *VM/370 Remote Spooling Communications Subsystem (RSCS) User's Guide*, IBM Systems Library, order number GC20-1816, available through IBM branch offices.
6. *IBM Virtual Machine Facility/370 Introduction*, IBM Systems Library, order number GC20-1800, available through IBM branch offices.
7. *General Information—Binary Synchronous Communications*, IBM Systems Library, order number GA27-3004, available through IBM branch offices.
8. J. H. Mommens and J. Reviv, *Coding for Data Compaction*, Research Report RC 5150, IBM Thomas J. Watson Research Center, Yorktown Heights, New York 10598 (November 1974). (ITIRC No. 75A00475.)
9. H. R. Albrecht and K. D. Ryder, "The Virtual Telecommunications Access Method: A Systems Network Architecture perspective," *IBM Systems Journal* **15,** No. 1, 53-80 (1976).
10. *Introduction to VTAM (Virtual Telecommunications Access Method) (Level 2)*, IBM Systems Library, order number GC27-6987, available through IBM branch offices.
11. W. S. Hobgood, "The role of the Network Control Program in Systems Network Architecture," *IBM Systems Journal* **15,** No. 1, 39-52 (1976).
12. *Introduction to the 3704 and 3705 Communications Controllers*, IBM Systems Library, order number GA27-3051, available through IBM branch offices.

13. *OS/VS2 TSO Terminal User's Guide*, IBM Systems Library, order number GC28-0645, available through IBM branch offices.
14. *APL Language Reference*, IBM Systems Library, order number GC26-3847, available through IBM branch offices.
15. *System/360 ATMS-II General Information*, IBM Systems Library, order number GH19-6025, available through IBM branch offices.
16. *Customer Information Control System/Virtual Storage (CICS/VS) Version 1, Release 3 General Information*, IBM Systems Library, order number GC33-0066, available through IBM branch offices.
17. *IBM System/360 Operating System Telecommunications Access Method*, IBM Systems Library, order number GC30-2004, available through IBM branch offices.
18. *OS/VS TCAM Program Reference Summary*, IBM Systems Library, order number GY30-1024, available through IBM branch offices.
19. J. H. McFadyen, "Systems Network Architecture: An overview," *IBM Systems Journal* **15,** No. 1, 4–23 (1976).
20. *Systems Network Architecture General Information Manual*, IBM Systems Library, order number GA27-3102, available through IBM branch offices.
21. *Introduction to Advanced Communications Function*, IBM Systems Library, order number GC30-3033, available through IBM branch offices.
22. R. A. Donnan and J. R. Kersey, "Synchronous data link control: A perspective," *IBM Systems Journal* **13,** No. 2, 140–162 (1974).
23. *Synchronous Data Link Control General Information*, IBM Systems Library, order number GA27-3093, available through IBM branch offices.

Reprint Order No. G321-5098.