NTIA Report 87-214
NCS Technical Information Bulletin 87-9

Fiber Optic Networks and Their Service Survival

Martin Nesenbergs

DEPARTMENT OF COMMERCE * UNITED STATES OF AMERICA *

U.S. DEPARTMENT OF COMMERCE
Malcolm Baldrige, Secretary
Alfred C. Sikes, Assistant Secretary
for Communications and Information

NATIONAL COMMUNICATIONS SYSTEM

NATIONAL COMMUNICATIONS SYSTEM

March 1987

PREFACE

The National Communications System, Office of the Manager, in Washington, DC, has directed the Institute for Telecommunication Sciences to perform a list of engineering tasks for an NSEP-enhancing development program required of commercial fiber optic systems using Federal Government rights of way. This report covers in part, Task 1, Telecommunications Engineering. Its main topics are survivable network configurations and fast service restoration under stress.

Administrative and technical monitoring of this study contract were performed by Messrs. M. L. Cain, A. H. Rausch, and E. Greene of NCS. Technical and management supervision of the program at ITS were provided by Dr. W. A. Kissick and Mr. J. A. Hull.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF FIGURES (con'd)

LIST OF TABLES

LIST OF ACRONYMS

ADCCP	Advanced Data Communication Control Procedures
ANSI	American National Standards Institute
AO&M	Administration, Operations, and Maintenance
ARPA	Advanced Research Projects Agency
ARPANET	ARPA Network
AUTOVON	Automatic Voice Network
BER	Bit Error Ratio
BH	Busy Hour
CCITT	International Consultive Committee for Telegraph and Telephone
CCS	Common Channel Signaling
CCSS	Common Channel Signaling System
CM	Connectivity Matrix
CONUS	Contiguous (or Continental) United States
CRC	Cyclic Redundancy Check
DCTN	Defense Commercial Telecommunications Network
DDN	Defense Digital Network
DSN	Defense Switched Network
DSX	Digital Signal (Format X=0,1,2,3,...)
FCS	Frame Check Sequence
FD	Full Duplex
FIPS	Federal Information Processing Standard
FOCS	Fiber Optic Communication System
FR	Facility Restriction
HDLC	High-Level Data Link Control
IF	Intermediate Frequency or Information Field

LIST OF ACRONYMS (con'd)

ISDN	Integrated Services Digital Network
ISO	International Standards Organization
kb/s	Kilo (thousand) Bits per Second
LAP	Link Access Protocol
LC	Link Capacity
Mb/s	Mega (million) Bits per Second
MLPP	Multilevel Priority and Preemption
NCS	National Communication System
NM	Network Management
NSDD	National Security Decision Directive
NSEP	National Security Emergency Preparedness
OH	Overhead
OSI	Open System Interconnection
RT	Routing Table
SDLC	Synchronous Data Link Control
SNA	Systems Network Architecture (IBM)
SOW	Statement of Work
SRI	Stanford Research Institute
TC	Traffic Carried
TO	Traffic Offered
VLSI	Very Large Scale Integration

FIBER OPTIC NETWORKS AND THEIR SERVICE SURVIVAL

Martin Nesenbergs*

The objective of this study is to look at fiber optic networks in a predominately functional domain and to assess their potential survivability advantages from that point of view. As a consequence, service survivability is emphasized far more than physical survivability, although physical existence of facilities is a definite prerequisite for all telecommunications services.

The need for a quantitative (or formal, or unique, or numerical) definition of the term "survivability" is addressed. The report proposes a partial solution to this problem. It introduces a network-related quantity, defined with the moments of the connectivity cross section histograms, that appears to possess many of the properties wanted for measuring and comparing survivabilities of different topologies. For lack of a better name, that quantity may be called the effective topological survivability index.

The fiber advantage of large data throughput, typically in tens of Mb/s, must be exploited when connectivity or other network status is in doubt. This is part of the network reconstitution or restoration issue. Outlines of procedures, protocols, and formats are given to achieve comprehensive network-wide restoral for small but still realistic networks. The information fields of extensive reconstitution data arrays are possible and advisable. If transmitted, received, and stored rapidly, and not processed in a lengthy manner, these data arrays are shown to offer unprecedented restoral opportunities. Through locally or regionally focused restoration processes the methods appear practicable even for very large networks.

The conclusion is that full-scale automation is essential. It should be distributed to all nodes of the network and its implementation should be with the very highest speed parallel processors. Any node that survives should be capable of both initiating and participating in the network restoration sessions. Thus, centralized hierarchical controls are to be avoided.

1. INTRODUCTION

This work is part of a larger study on National Security Emergency Preparedness (NSEP) for Fiber Optic Communication Systems (FOCS). As directed by the National Communications System (NCS) Technology and Standards Office, the broad objective of the overall study is:

> To establish "benchmark" specifications or guidelines to facilitate quantitative evaluation of fiber optic "Carrier's Carrier" proposed installations in accordance with NSEP requirements.

*The author is with the Institute for Telecommunication Sciences, National Telecommunications and Information Administration, U.S. Department of Commerce, Boulder, CO 80303-3328.

Two items of the above statement are taken as basic. They are: quantitative evaluation and NSEP requirements. Under the stated NSEP requirements, a key NCS trait to be evaluated is the survivability of essential, assured, services.

The background authority for this is found in several National Command directives. One can cite the Presidential Executive Order 12472 (April 1984), entitled "Assignment of National Security and Emergency Preparedness Telecommunications Functions." Among key responsibilities for the NCS, the Order directs the development of a national telecommunications infrastructure which is:

A. Survivable.
B. Responsive to NSEP needs of the President and the Federal Government.
C. Capable of satisfying priority telecommunications.
D. Consistent with other national policies.

The National Security Telecommunications Policy per NSDD-97 includes the statement: ". . . national telecommunications infrastructure must possess the functional characteristics of connectivity, redundancy, interoperability, restorability, and hardness necessary to provide a range of telecommunication services . . .".

Finally, the Technical Statement of Work (SOW) for this particular NSEP/FOCS project contains two relevant passages. In paragraph 2.0, under Purpose, the SOW states: ". . . the design parameters addressed (will) minimize interruptions of service . . . by proper attention to features which facilitate quick restoral of operation . . .". In paragraph 4.0, under Technical Requirements, Task 1, Telecommunications Engineering, the SOW calls for: "These studies will also provide the basis for assessing the need for additional network configuration capabilities to allow emergency restoral to be effected. (That is implied to include) restoration capabilities (e.g., alternative routing, etc.), subsystem modification (e.g., reconfiguration, etc.) . . . ."

There is clearly an abundance of topics to be addressed here. To proceed, we have limited the scope by selecting issues that appear more important, more basic, and seemingly more manageable within the initial given tasks. The so selected primary (or primitive) entities are:

I. Service.
II. Survivability.
III. Restorability.
IV. Functional network assets.

In as much as these functional qualifiers can assume too broad a scope, the intent here is to focus on rerouting, reconfiguration, exchanges of connectivity information, interoperability, and related selected algorithms.

Unfortunately, and as is well known, the objective process of quantification (i.e., assignment of measurable metrics) to many intuitively so acceptable concepts is difficult, if not impossible. As an example, consider the term "survivability." To resolve which of several system alternatives is more survivable and under what conditions, requires more than a subjective feeling of alleged experts. Yet a quantitative, practical or abstract, approximate or precise, workable engineering definition of survivability has so far eluded network designers and analysts. The general definition may well be an extremely difficult task. Past efforts have demonstrated that the term has many facets, that it depends on numerous variables (i.e., parameters), and that it varies widely with end-user missions, applications, technology, circumstances, location, as well as with time.

This report cannot claim a definitive solution for this dilemma to anyone's satisfaction. However, partial attempts are made in Section 2 to characterize the more apparent properties of the term "survivability," at least as said properties would pertain to the NSEP/FOCS needs. Section 3 extends the approach further by focusing on the continuity (or reconstitution, or rapid restoral) of end-to-end data, voice, or other services. Finally, Section 4 considers specific network examples to illustrate the use of the offered tools, as well as to remind one of end-user service realities and network design issues.

As a final comment in the Introduction, one must stress the advantageous transmission characteristics of the fiber optic medium. The physical properties (e.g., light absorption, scattering, wavelength effects, darkening, and recovery) that ensue from different deployment and implementation techniques are described in companion Reports prepared under this project. The properties are sufficiently unique, so that their correct exploitation is both a challenge and an opportunity for engineering survivable facilities.

Compared to other point-to-point transmission media, the fiber optic networks have enhanced potential for assured and survivable wideband communications service. For optimum benefits, however, the following conditions should be noted and issues resolved:

(1) While physical link connectivity is necessary, by itself it does not guarantee quick reconfiguration, rerouting, or uninterrupted service under stress. Survival and availability of nodal intelligence (e.g., control software) is also needed.

(2) Timely knowledge of link and node status must be made available for network controls, especially under unforeseen facility damage and/or traffic surge conditions.

(3) Automated management of the latest status information must be used for fast establishment and updates of connectivity, routing, traffic handling, and related operational tables throughout the network. In light of the huge bandwidth resource of the fiber, anything less than full automation would be a waste of precious restoration time.

(4) Essential status, such as connectivity update, information must be distributed everywhere, or as widely as possible. This will enable all connected nodes, including gateways to other networks, to be part of networkwide service restoral.

(5) The initial triggers for connectivity updates should ordinarily be generated either locally or remotely by automated processes. However, human overrides are to be permitted to accommodate testing and actions per command authority.

(6) Fiber optic transmission links enable orders of magnitude larger data rates than those offered by other media. This advantage makes it seemingly unnecessary to skimp or economize on the number of bits in the formats transmitted for control or protocol exchanges between nodes. It also suggests that, whether complex or simple, the node processes should be fast, so as to exploit efficiently the huge capacity of the fiber.

(7) Most advanced associated signaling schemes, viz., CCITT System #7 or other ISDN compatible CCS, could augment the crisis procedures delineated above. In fact, such functions as the exchange of the latest connectivity data now appear realistic on an unprecedented scale, with abundant message volumes and speed. Illustrations of this claim will follow later in the report.

2. SURVIVABILITY

2.1 Threat to Connectivity and Service

It is conventional to view the world of survivability in a reference framework of threats, facilities, pro-and-con tactics, and eventual system damages. For telecommunication networks the broad situation can be characterized as in Figure 1. Note that one starts on the left of this Figure with a given, or to be assumed, threat or stress scenario (Box 1). The stress causes two discernible effects on the candidate system: modified or increased offered traffic surges (Box 2), and potential damage to network facilities (Box 3). The damaged facilities can be links, nodes, and their support hardware, as well as software that is either centralized or distributed among the nodes.

This report is not concerned with Boxes 1 to 3. It starts with the premise that the NCS is faced with specified essential traffic requirements. The essential offered crisis traffic is likely to have message characteristics different from the ordinary. The crisis can result in general networkwide overload. Or the overload can occur in certain regions, during certain time intervals. In such cases one speaks of focused, geographic or temporal, traffic overloads. Their traffic surges or peaks can be quite high, easily exceeding excesses of 100%, yet they must be telecommunicated over the given, perhaps damaged, network facilities.

In Figure 1, our part of NSEP/FOCS survival work assumes that the traffic requirements are defined elsewhere. The study reported here pertains largely to Boxes 4 and 5. In particular, if the processes of Box 4 show little or no functional outages despite component damage, then one can speak of the network as at least functionally surviving. Or, when one damaged network suffers outages, it may still be operationally useful if supplied with adequate gateway facilities to other networks. Such gateways on a network serve two interoperation purposes: (a) they detour or take messages away from the

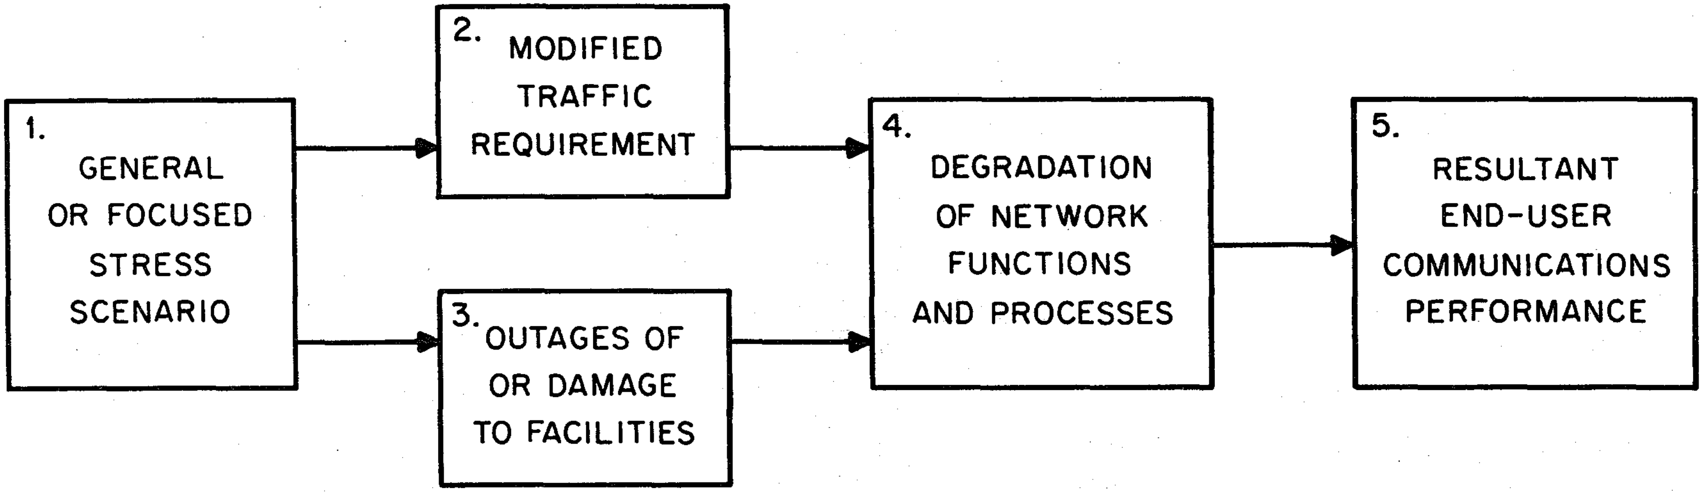
1. GENERAL OR FOCUSED STRESS SCENARIO
2. MODIFIED TRAFFIC REQUIREMENT
3. OUTAGES OF OR DAMAGE TO FACILITIES
4. DEGRADATION OF NETWORK FUNCTIONS AND PROCESSES
5. RESULTANT END-USER COMMUNICATIONS PERFORMANCE

Figure 1. The stress and survivability framework for end-user services.

network, and (b) they accept exterior messages from other networks to help them with their transport problems. But, survivability of systems and/or their internal functions is not the foremost concern of most end-users.

In critical conditions the term "survivability" can mean survivability of service to the end-users, e.g., the national command authorities. Because different end-users have different missions, their communications requirements may be unique and distinct. For instance, they may differ in their tolerance of message delays, bit error ratios (BER), or other measures of service quality. In general, therefore, survivability may be a mission-sensitive function of delay, BER, blocking probability, delivery of essential throughput, and other more or less standard performance parameters. But perhaps of largest significance to fiber optic systems is the basic role of network connectivity or its converse, disjointedness. Let us explain.

2.2 Fibers, Automation, and Service Continuity

Given the huge data rates of optical fibers, such as around 40 Mbps in most 1986 installations, a single operational fiber suffices to carry the data streams of many high-speed user terminals. If such a link is ON, the network either has or does not have the capability to benefit from the link's existence. Because the utility depends also on interoperability, routing, translation of address fields, user ID acceptance, and many other protocol-related things to be functional under stress. One concludes that the control functions of the network should possess the intelligence to permit the utilization of any and all connected live FOCS links. Switching or detouring to non-optical transmission trunks in other backup networks is to be provided through cooperative gateways. This all should be planned and implemented to the maximum degree possible, while anticipating either worst case, randomly distributed, or engineering target outages of network facilities.

Looking from the opposite point of view, if physical connectivity is not there and cannot be restored, neither is communications. The performance of both systems and user services, including their various objective and subjective performance parameters, become irrelevant when there is no connectivity. However, in the larger topology of national networks many component outages can occur before a complete, point-to-point or region-to-region loss of physical connectivity occurs. Some previously mentioned interconnecting paths may survive, but may also be hard to find by the human operators. If that were so, the connectivity might as well be physically nonexistent. The obvious answer to this dilemma is automated, state-of-the-art, distributed, dynamic, and very fast connectivity restoral.

Rapid restoral functions and issues are the main topics of what follows. They are addressed within the general framework of Figure 2. The comprehensive NSEP design and implementation benchmarking for the FOCS survivability enhancement is the main objective, as shown at the top. Below it, a very critical element is facility hardness engineering for FOCS. That topic is addressed and reported elsewhere. Similarly beyond the scope of this report are manual procedures, such as emergency- and chain-of-command (doctrine)-caused activities by human operators. The remaining quick and automated processes for crisis management and service restoral are considered here, but with a few notable exceptions. Whether automated or not, the functions of Administration, Operations, and Maintenance (AO&M), as well as Network Management (NM), are excluded here. Likewise for Common Channel Signaling (CCS). While it appears that Associated CCS might have advantages over Nonassociated CCS for FOCS, control signaling per se is not reviewed in this study.

Sections 3 and 4 will deal with service restoration aspects of survivability. Methods to ascertain the status of network connectivity and ways of using such status information to route and handle messages will be presented there. Before that, however, a few observations may be in order about the dependence of the often used term "network survivability" on mutual connectivity between network components (such as nodes).

2.3 Connectivity Cross Sections

For a service area (or node) A to communicate with another service area (or node) B, in principle, a number of independent (that is, separate and totally disjoint) message paths may be available in a network. Two end-to-end paths may be defined as independent when they have no links or nodes (excluding the end nodes) in common. See Figure 3. The bigger the overall connectivity cross section, i.e., the more separate paths there are between all significant interacting service areas, the more confident one feels about service survival --given fixed numbers of facility outages. For fixed locations A and B, conditional on connectivity, the connectivity cross section is a natural number, such as 1, 2, 3,.... In case of single homing by either A or B to the

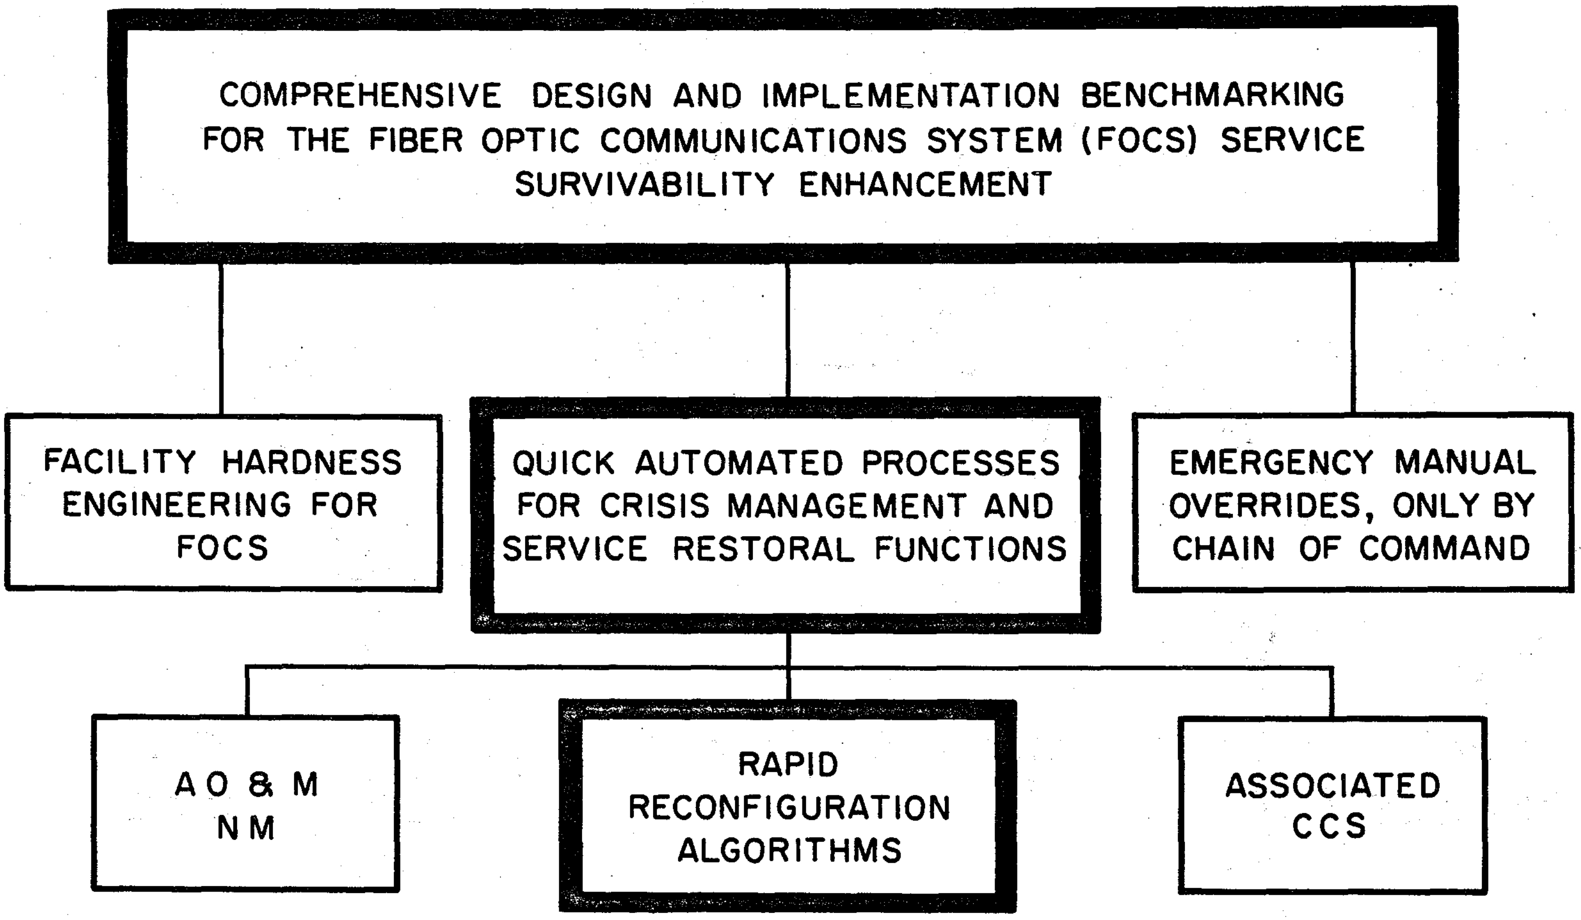
COMPREHENSIVE DESIGN AND IMPLEMENTATION BENCHMARKING FOR THE FIBER OPTIC COMMUNICATIONS SYSTEM (FOCS) SERVICE SURVIVABILITY ENHANCEMENT
FACILITY HARDNESS ENGINEERING FOR FOCS
QUICK AUTOMATED PROCESSES FOR CRISIS MANAGEMENT AND SERVICE RESTORAL FUNCTIONS
EMERGENCY MANUAL OVERRIDES, ONLY BY CHAIN OF COMMAND
AO & M NM
RAPID RECONFIGURATION ALGORITHMS
ASSOCIATED CCS

Figure 2. The role of control algorithms in rapid service restoral.

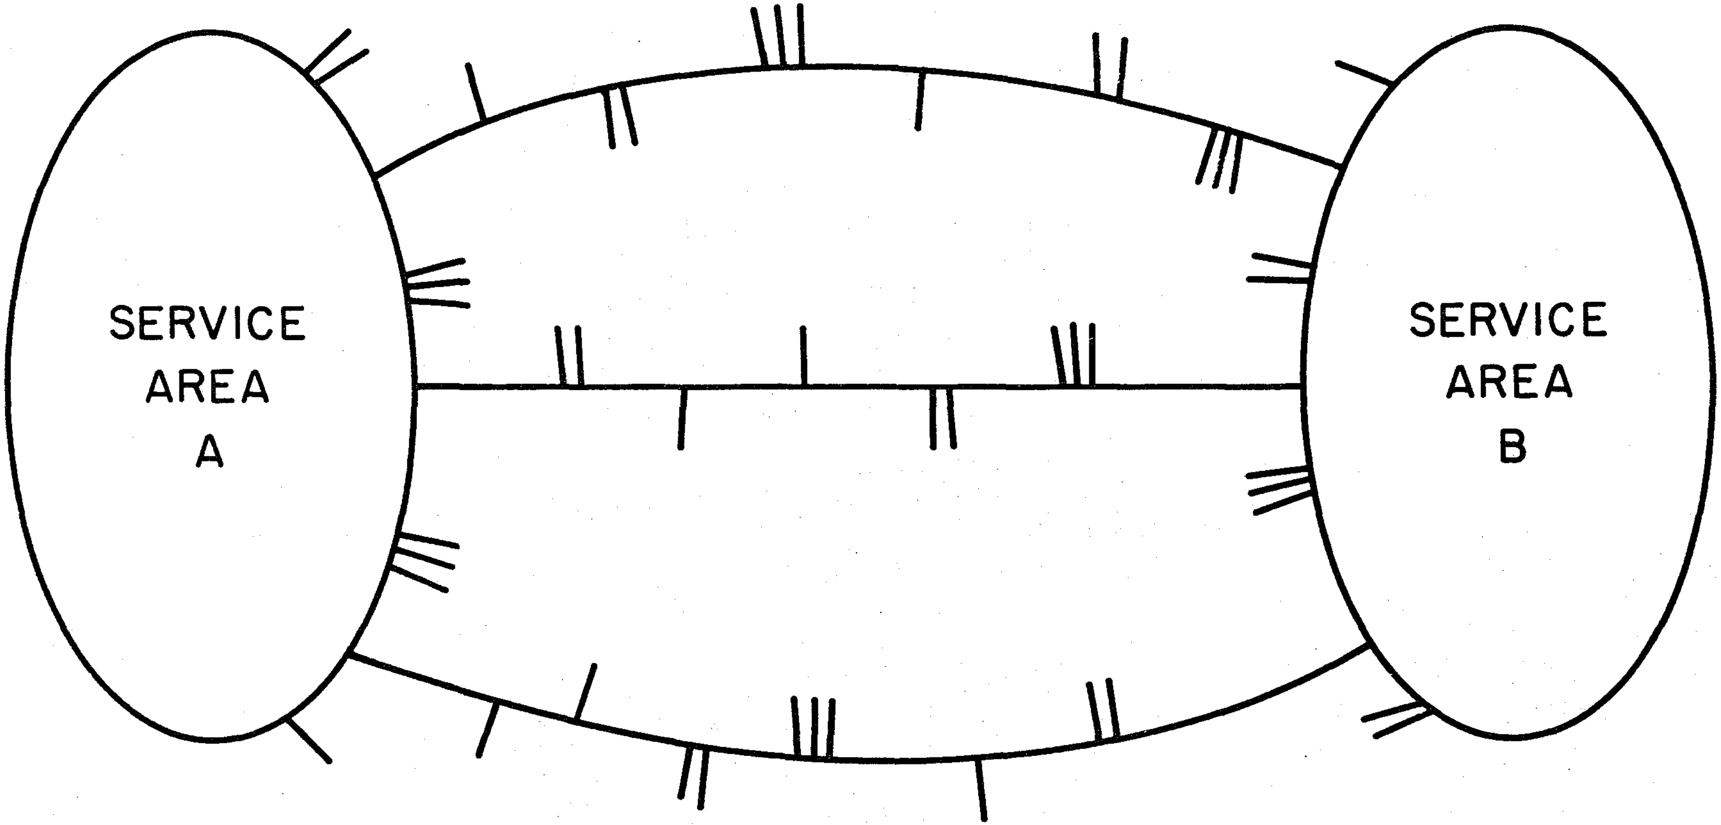

Figure 3. Connectivity cross section in terms of independent paths between A and B.

rest of the network, that number is 1. If the lesser of the two service area accesses to the net is double homing, then the cross section number is either equal to or less than 2, because the intervening network cannot increase that number. And so forth. For the ensemble of all possible (A,B)-pairs on a network, the cross section numbers form a distribution. Said distribution appears to contain essential information about the topological survivability of the physical network. Some properties of the distribution, such as extreme statistics or moments, can contribute towards a quantitative definition of network and/or service "survivability." We shall return to the distribution shortly. First, however, one must note that the determination of even a single (A,B)-dependent number should be done with care.

If an algorithm stubbornly insists that the shortest or most economical routes (e.g., least-cost routing) be part of the proceedings, then it may miss the true maximal connectivity. This hazard is illustrated in Figure 4. Between the two service areas, (A,B), the shortest distance in number of hops (sequential links) is 4, as drawn by the horizontal heavy line. If one preassigns this route to be part of the cross-section-seeking algorithm, then no other separate route is possible by definition. One would thus arrive at the fallacious conclusion that the connectivity cross-section index between A and B is 1. Inspection of Figure 4 reveals at a glance that there are instead 3 separate paths possible: these paths are all longer than 4, the paths go through all nodes not included in A or B, and they exhaust the triple homing of both A and B. Therefore more than 3 paths are impossible, and therefore the maximum connectivity cross section between A and B in this particular topology is 3. However, it should not be too difficult to construct an effective maxima-seeking algorithm that determines individual cross-section numbers for both existing and planned networks. Or, for synthesis of entirely new networks, constructive means could be devised that implicitly guarantee required connectivity cross-section numbers or better.

Return next to the distribution of cross sections that pertain to the connectivity of a network. In Figure 5, one finds a sample network of ℓ=41 links and n=20 nodes. It is tacitly assumed that, at least for time being, all links are full duplex (FD), bi-directional, so representative of typical fiber optic installations. This implies that between nodes or service areas the connectivity cross-section index is not dependent on the direction, A-to-B or B-to-A, taken in a path. A visual scan of the network suggests that some node

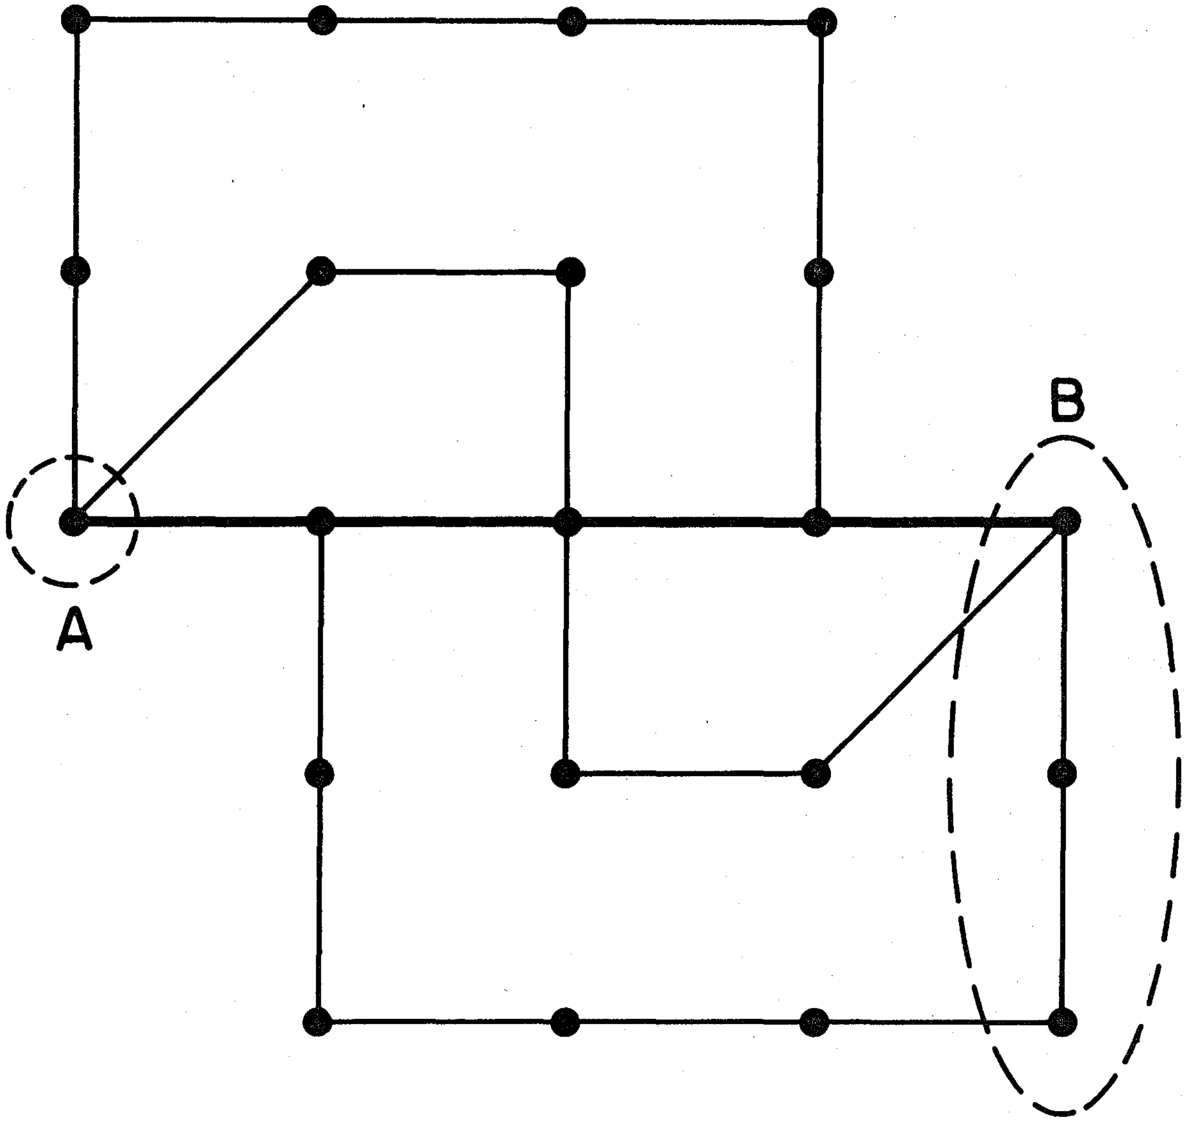

Figure 4. The cross section versus shortest route dilemma.

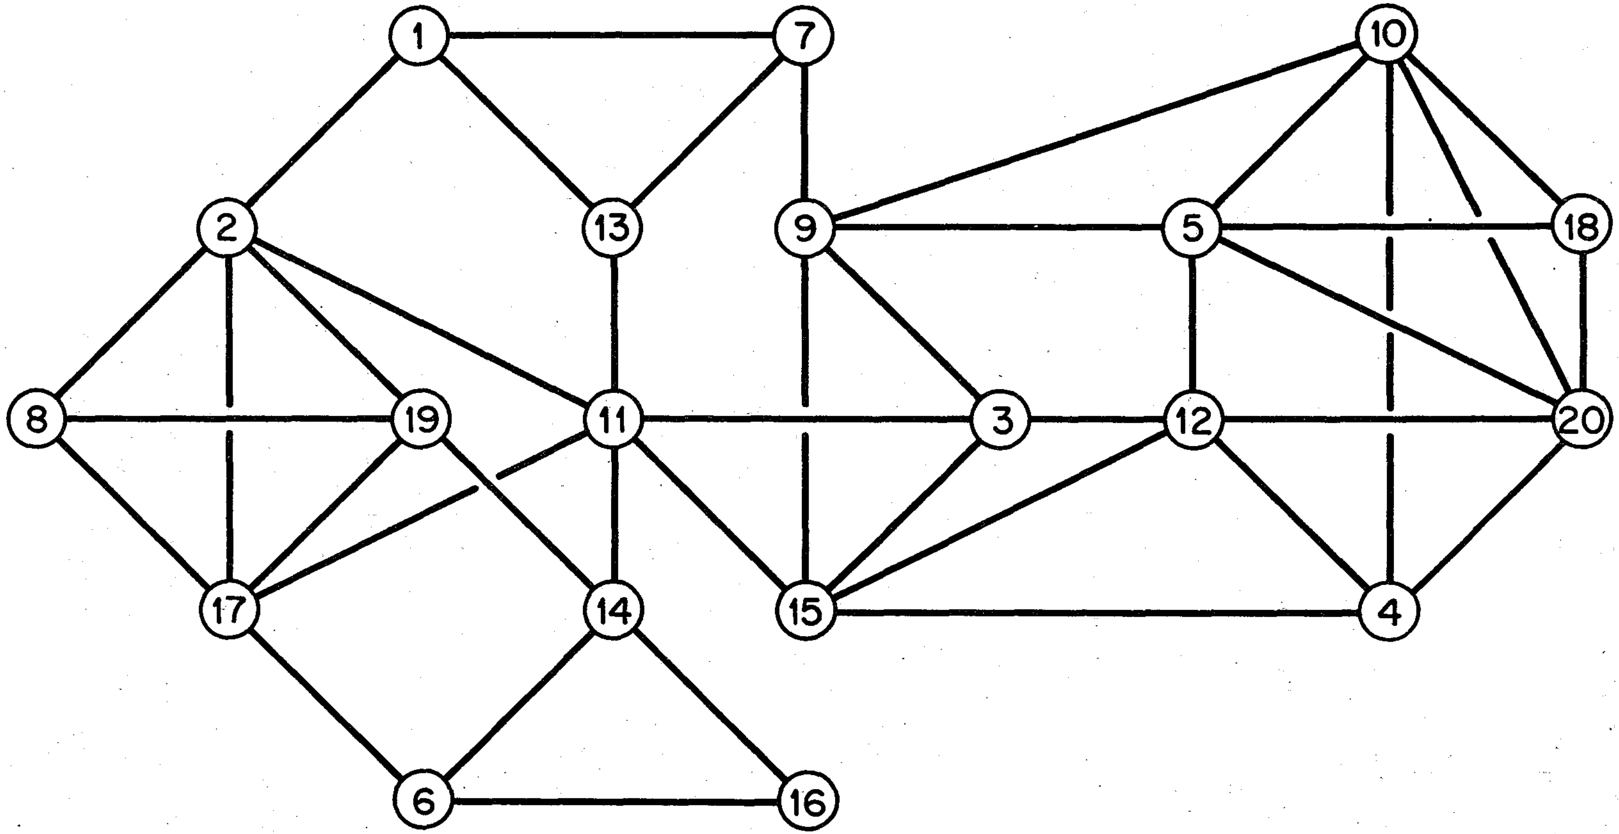

Figure 5. An illustrative network with 41 links and 20 nodes.

pairs, for instance nodes #6 and #16, or any nodes on the left paired with any node on the right, have a cross section of 2. Some other node pairs have connectivity cross sections as large as 5. An example of this are nodes #5 and #20. They can be linked by such independent paths as:

```
#5---#20,
#5---#10---#20,
#5---#12---#20,
#5---#18---#20,
#5---#9---#15---#4---#20.
```

It turns out that the range from 2 to 5 is the actual domain occupied by all 20(19)/2=190 cross sections of the network in Figure 5. The full density distribution has been computed and is plotted as a histogram in Figure 6, with x standing for the connectivity cross section. The minimum or guaranteed cross section of $\min(x)=2$ can be used as a numerical measure to compare relative survivabilities. But, so could the sample mean, which happens to be $m=2.47$, or the mean plus or minus some multiple of the standard deviation. The value of the sample standard deviation is roughly $s=0.76$.

For arbitrary networks with arbitrary broad or narrow histograms over x, a partial possible survival measure may be of the form $m-\theta s$, where $0<\theta<1$. For general application the θ coefficient could use either a physical or a statistical justification. Being unclear at present, that justification is left for future work. In the present example, by the way, an assignment of $\theta=1/4$ looks reasonable. Then by definition the survivability cross-section index for this network would be 2.28. The assignment of $m-s/4$ also fits the very concentrated distributions with $s<<m$, and in particular those homogeneous topologies that possess $s=0$.

The homogeneous networks that satisfy the $s=0$ property deserve a few comments. First, of course, the histogram for $s=0$ represents a degenerate distribution or a delta function with its entire mass concentrated at $x=m$. Second, the survivabilities of two networks with the same m are assumed indistinguishable, granted no other influential factors. Third, if a network alpha has a larger m value than network beta, then one says that network alpha is more survivable. Fourth, since the above strict ordering should apply for all $m=1,2,3,$ ... without limitation, one might as well assume this partial (topological) network "survivability" to be a linear function of m.

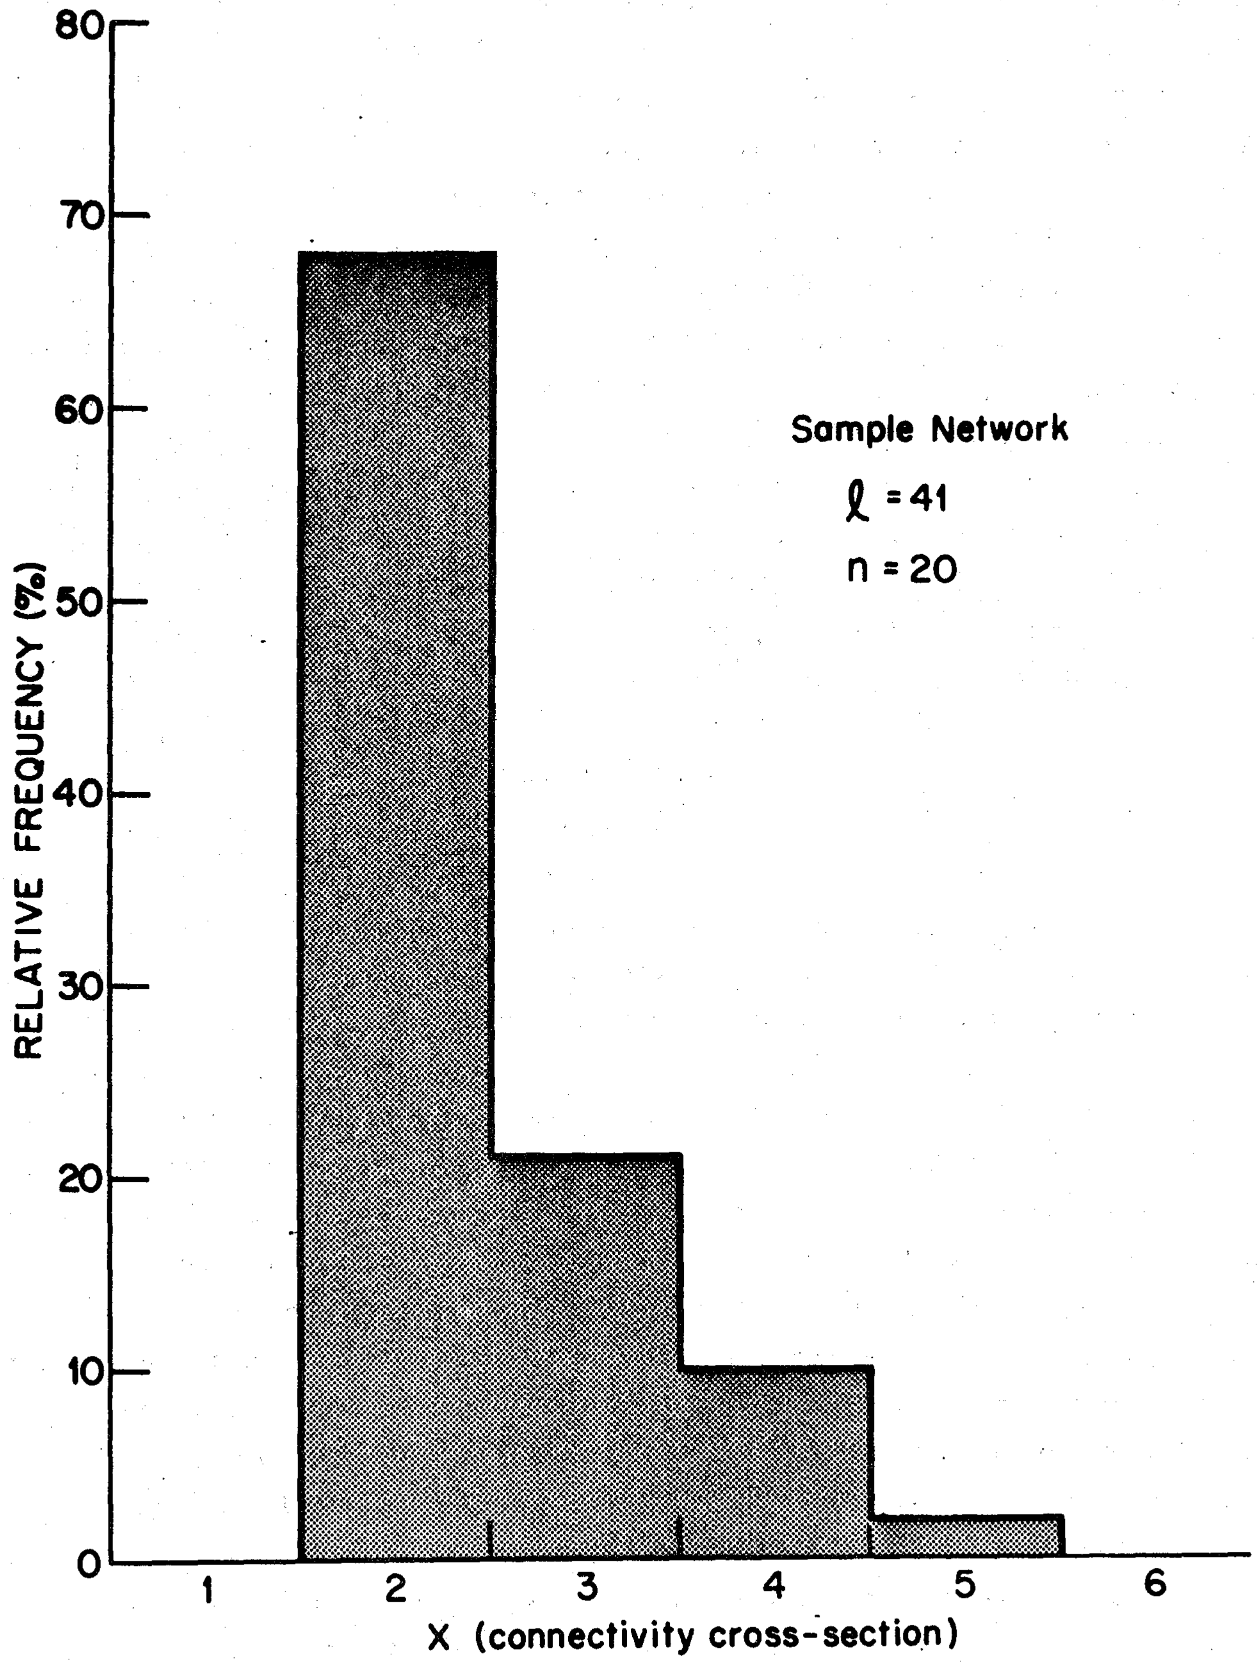

Figure 6. A histogram of connectivity cross sections.

Networks with s=0 can apparently be constructed for all integer values of m. We will demonstrate this by considering first infinite and later finite size networks.

Figures 7 and 8 show periodic s=0 structures of infinite size, i.e., with infinitely many links and nodes. The mean singular cross section, m, ranges from 1 to 4 in Figure 7. Of course, the uniformly periodic designs are not required to achieve s=0. All trees, infinite and finite, possess s=0 and m=1. Extensions to m values beyond 5 and 6 in Figure 8 can be done either by adding direct links between selective nonadjacent node pairs or by continuing with similar periodic patterns as introduced above.

Finite size topologies with vanishing standard deviation, s=0, are not much harder to conceive. Simple illustrative patterns for n=6 nodes are offered in Figure 9. Other less symmetric, less regular configurations appear possible. Thus, while there might be some exceptional {n,m} cases for which the s=0 condition cannot be met, overall there seem to exist numerous practical possibilities, especially for $0 \leq s << m$.

3. SERVICE RESTORATION

3.1 Connectivity Updates

The importance of connectivity to fiber optic networks was mentioned in previous sections. By itself connectivity may not be fully sufficient to guarantee any single or integrated service. However, it is an absolute necessity for communications between remote service areas. Here one addresses the dynamic issues that pertain to connectivity status changes and service restoration. Stresses can lead to damaged or disabled node and link facilities. Facilities can quite suddenly become disabled or at other times be restored to service. Furthermore, traffic serving facilities can become so congested as to be at least temporarily incapable of handling new traffic.

A complete specification of the term "restoration," such as in system reconfiguration or service restoration, generally involves too many application oriented parameters. Whatever form the definition may take, it is apt to be too cumbersome to satisfy every case encountered in practice. A simple scenario will suffice here. Assume that the following two criteria are of most importance: (a) connectivity through the network facilities, and (b) acceptable service quality to the end users. The restoration process is said to be really needed when either one or both of the criteria are not met.

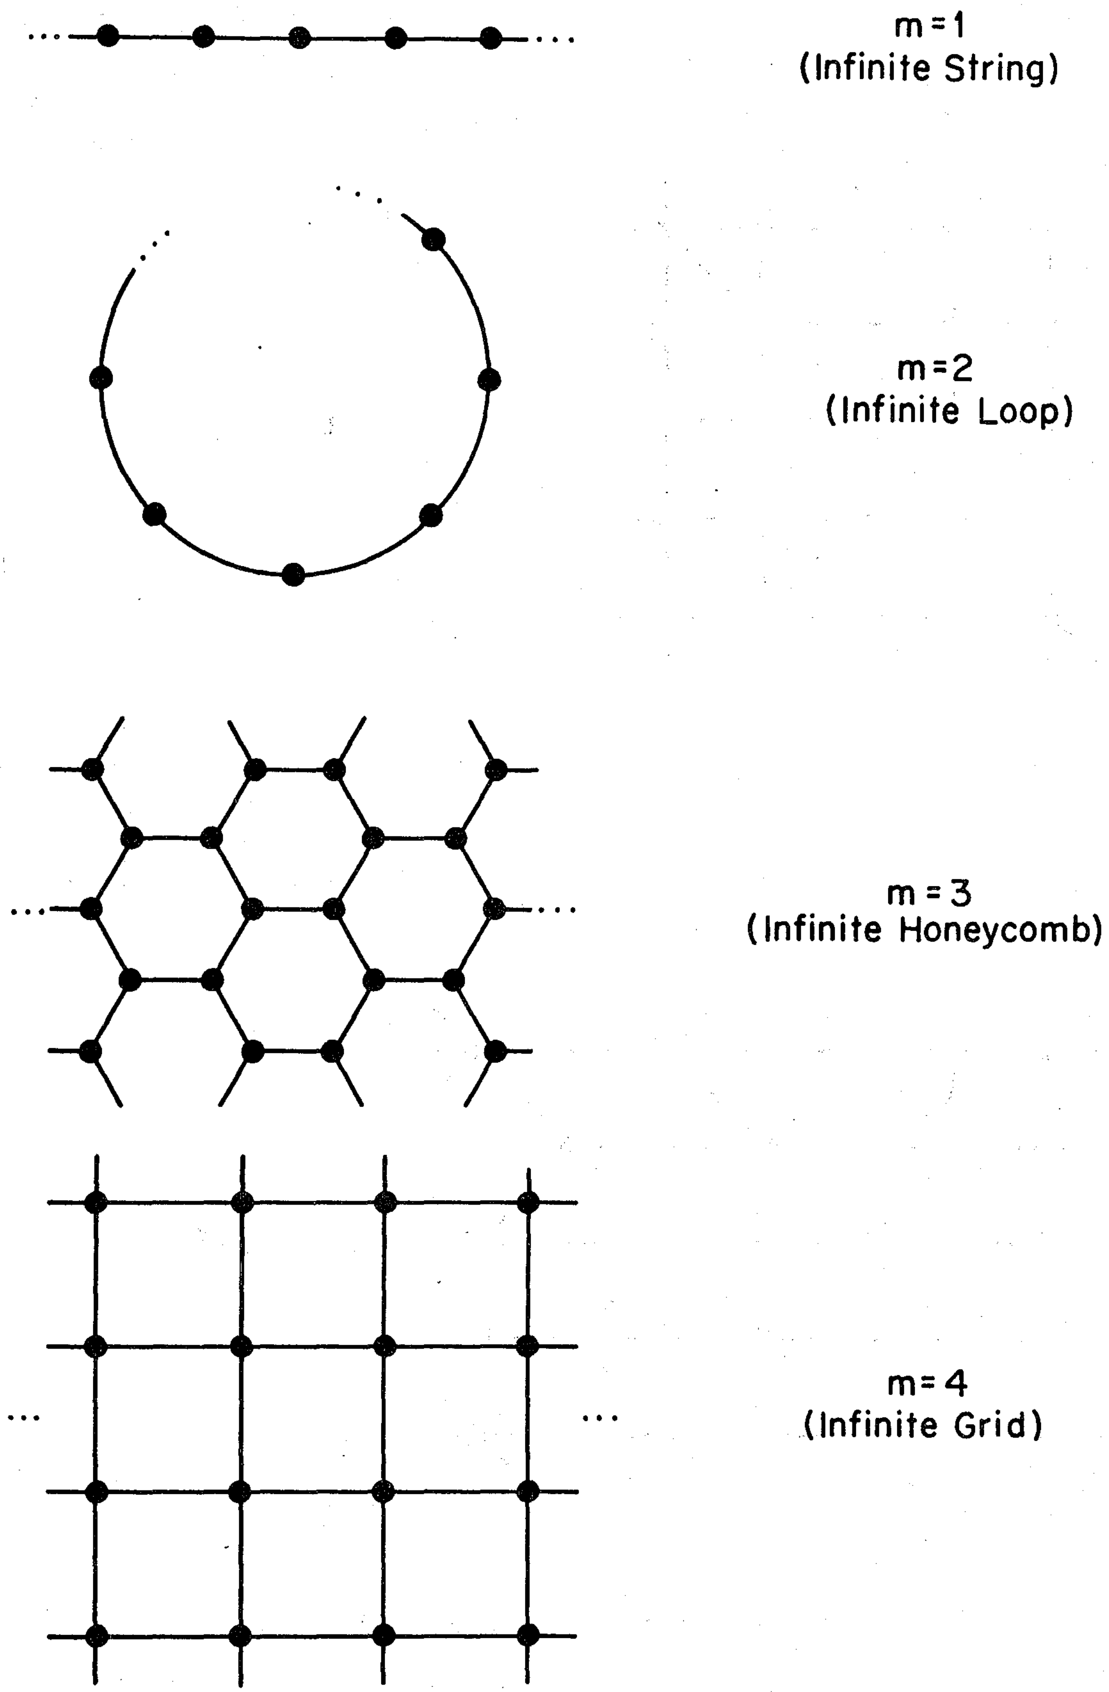

Figure 7. Infinite s=0 topologies for m=1,2,3, and 4.

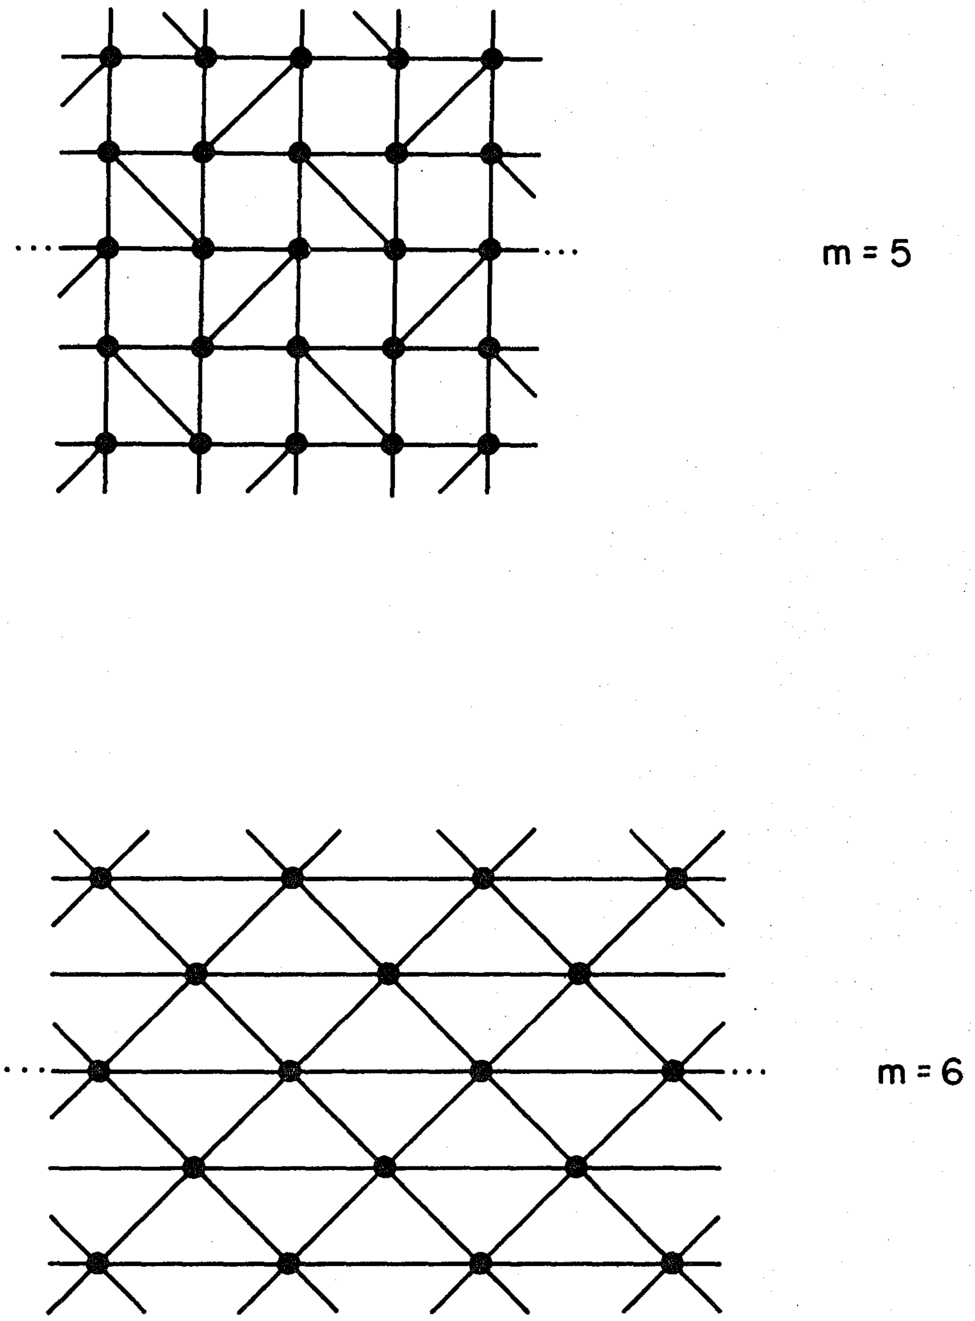

Figure 8. Infinite topologies s=0 for m=5 and 6.

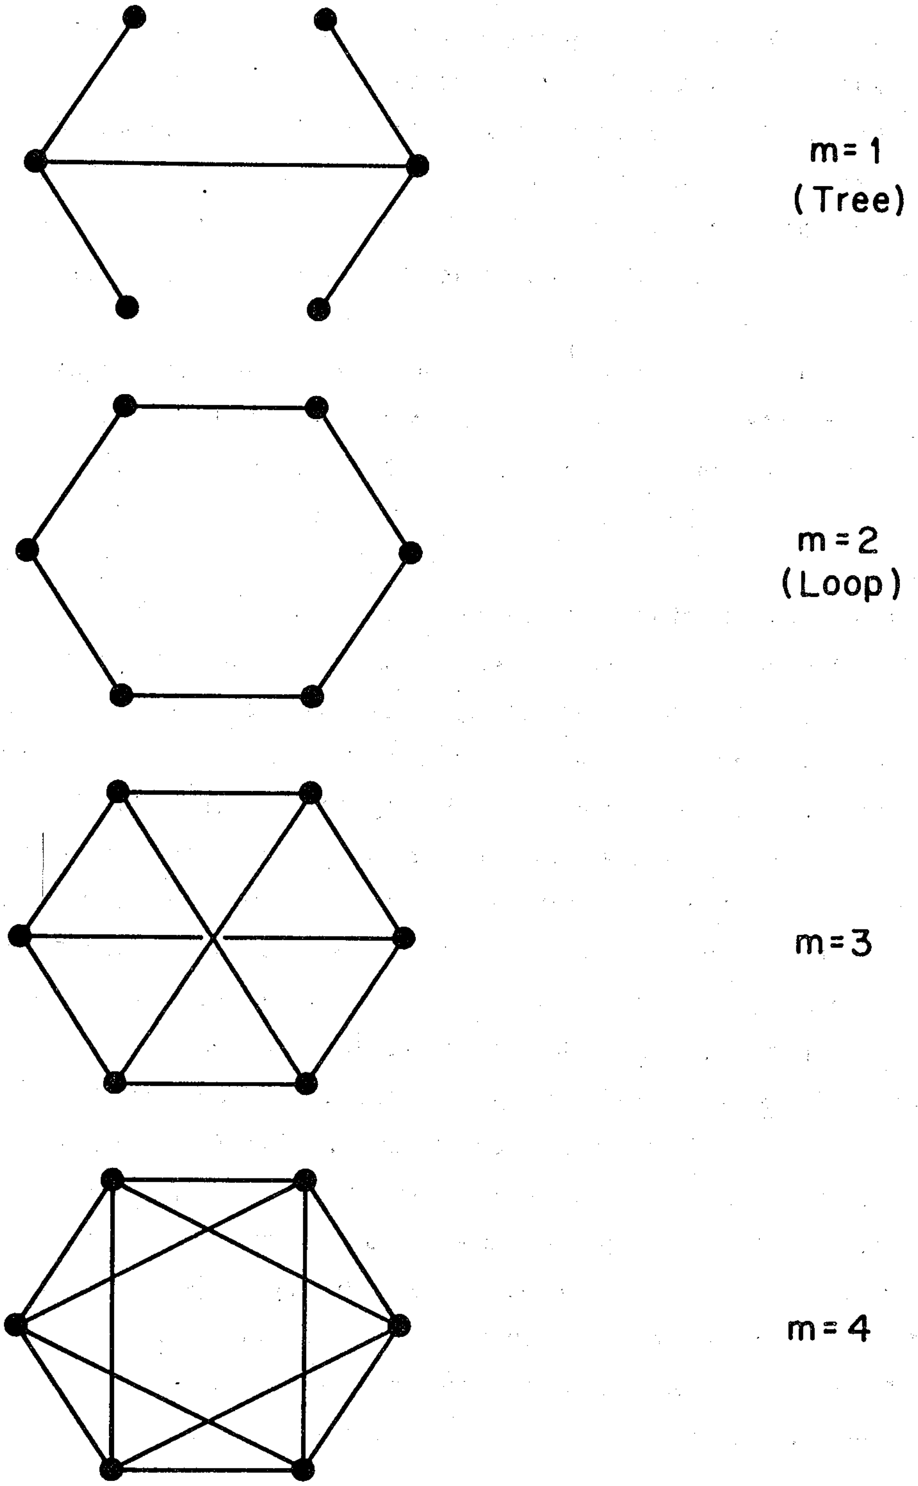

Figure 9. Finite s=0 topologies with six nodes and m=1,2,3, and 4.

Otherwise restoration activities can be initiated by various AO&M/NM functions, not excluding steps taken in response to faulty information.

One view of the restoration process is therefore a sequence of events in the (Q,C)-plane, where C stands for connectivity and Q denotes quality of service. This point of view is illustrated in Figure 10. As shown, the x-axis is the quality of service, or Q, expressed in some arbitrary or abstract units. The y-axis is the network connectivity, C. The connectivity can be quantified with the help of the previously introduced connectivity cross section histograms, see Section 2.3.

The (Q,C)-plane is divided into two parts. The acceptable region, found to the upper right, consists of all (Q,C) points where both Q and C are good enough. The unacceptable region consists of all other points and is normally found towards left and down in Figure 10. When things are working well, one is in the acceptable region. A point (Q_0,C_0) represents such a case at time T_0. Let a damaging stress event start immediately after time T_0, so that a short time later at T_1 the process finds itself under stress and already in an unacceptable state (Q_1,C_1). At this time, while the system is still under attack, some trans-attack restoration attempts may begin. Let T_2 be the time when the attack to the network ceases. The system in now shown to be at (Q_2,C_2), still in the unacceptable region. From now on the so-called post-attack restoration takes place. Eventually, at time T_3, the system and the service are returned to the acceptable region, say to some locus (Q_3,C_3). Although here the operations may or may not be as good as they were in the initial pre-attack situation, the restoration can stop at T_3.

The time interval T_3-T_0 is called the restoration time. As a function of network size, damages, manual tasks, automation, facility speed, and sundry protocol factors, the restoration time can vary widely. Perhaps from milliseconds to weeks. Methods to estimate actual numbers will be addressed later. First, however, we simplify the topic by selecting one particular network facet for deeper scrutiny. That most important facet is network connectivity. Whenever connectivity changes are suspected, with or without justification, potential data base updates and associated protocol activities may be called for.

To restore efficient routing and delivery of message traffic everywhere, the latest connectivity alterations must first be established. Next, the pertinent connectivity updates must be broadcasted as widely and as promptly as

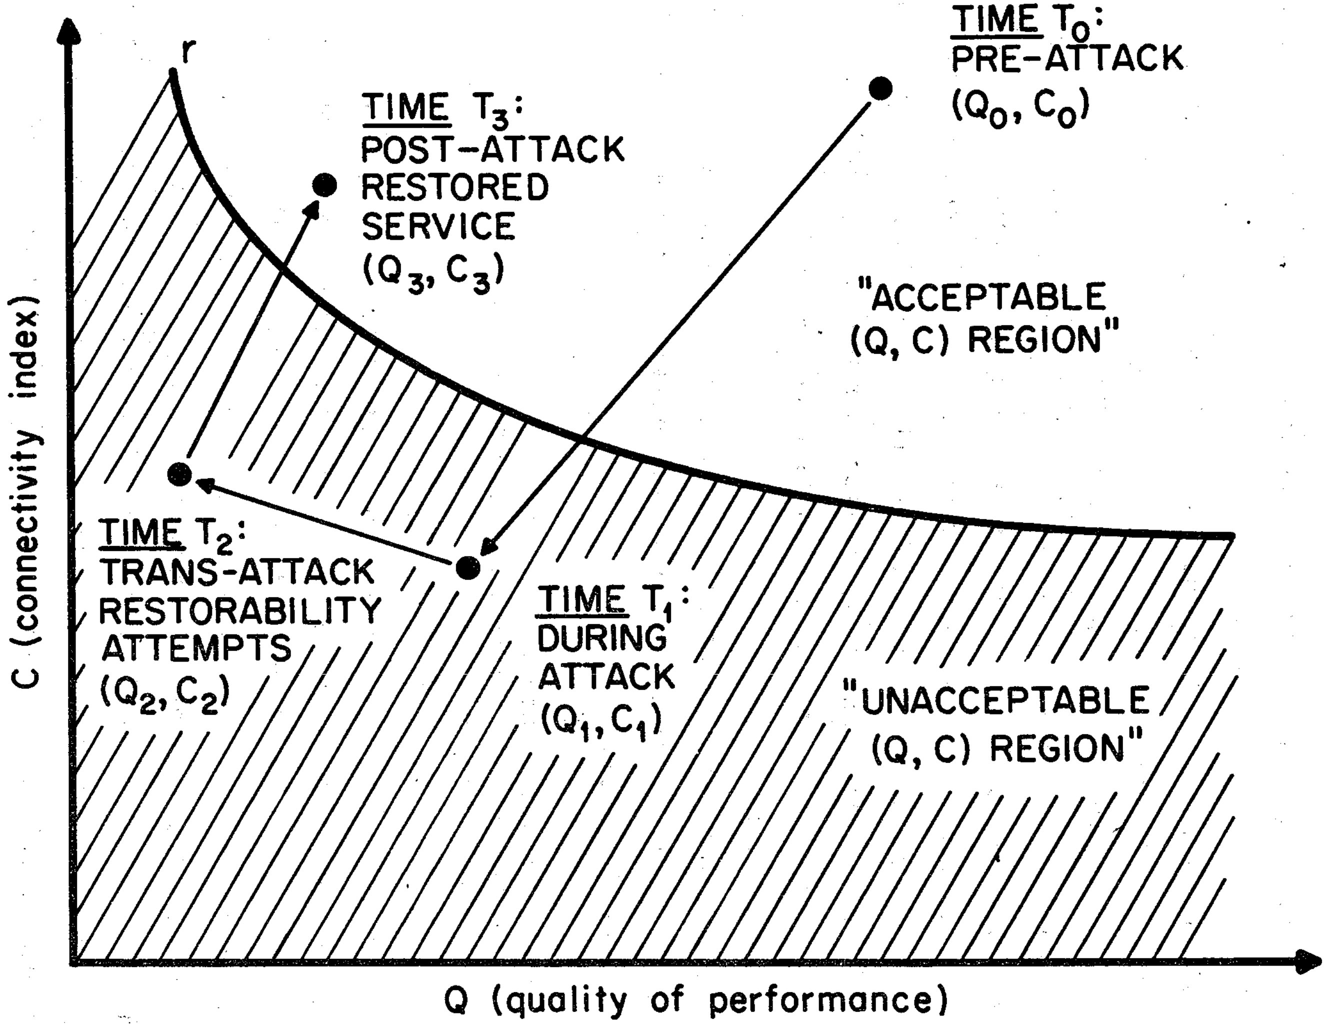

Figure 10. One view of the restoration process.

possible. Both collection and distribution of update information takes place during the restoration time window, T_3-T_0, and quite naturally should involve all assigned operational network control centers in small enough networks. In larger networks, time and other resource management arguments may lead to tailormade limitations of the search-and-broadcast domains. These "update domains" can be local geographic regions, national hierarchical centers, or functionally determined sets of nodes.

Because of the large fiber bandwidth and VLSI speeds, it appears possible to execute quite sophisticated, connectivity related, search-and-broadcast algorithms in a negligibly short time interval. For present purposes let that interval be less than one second. If regularly scheduled, perhaps no more than approximately once per hour, or if triggered by stress events, the time so wasted seems acceptable for most NCS scenarios. An illustration of these search-and-broadcast events is given in Figure 11. On FOCS networks the events are intended to convey link and node outages, connectivity updates, routing table revisions, gateway selections for bypass networks, and other restoration functions to be determined.

The individual connectivity update interval has a fine structure. Thus in Figure 11 the interval is divided into three phases. These are preliminary functional phases defined as:

(1) The start phase.
(2) The actual search-and-broadcast phase.
(3) The stop phase.

All three phases must have limited durations to meet the postulated less than one second objective for their sum. Since networks can be of varied sizes and shapes, the most crucial time limitation is likely to affect the search-and-broadcast phase. There are many alternatives to implement the three phases. In Figure 12 one finds a diagram of activities at a particular node. The node can be either the "initiating node," that is, the one that broadcasts the very first update message, or it can be any other node that happens to receive a subsequent delayed copy (e.g., on the first, second, third, ..., hop) of the update request.

The activity of every node is under the control of its own internal timer, T_i. In Figure 12, the inherent separation of the start, search-and-broadcast, and stop phases, can refer basically to the start, count, and stop states of

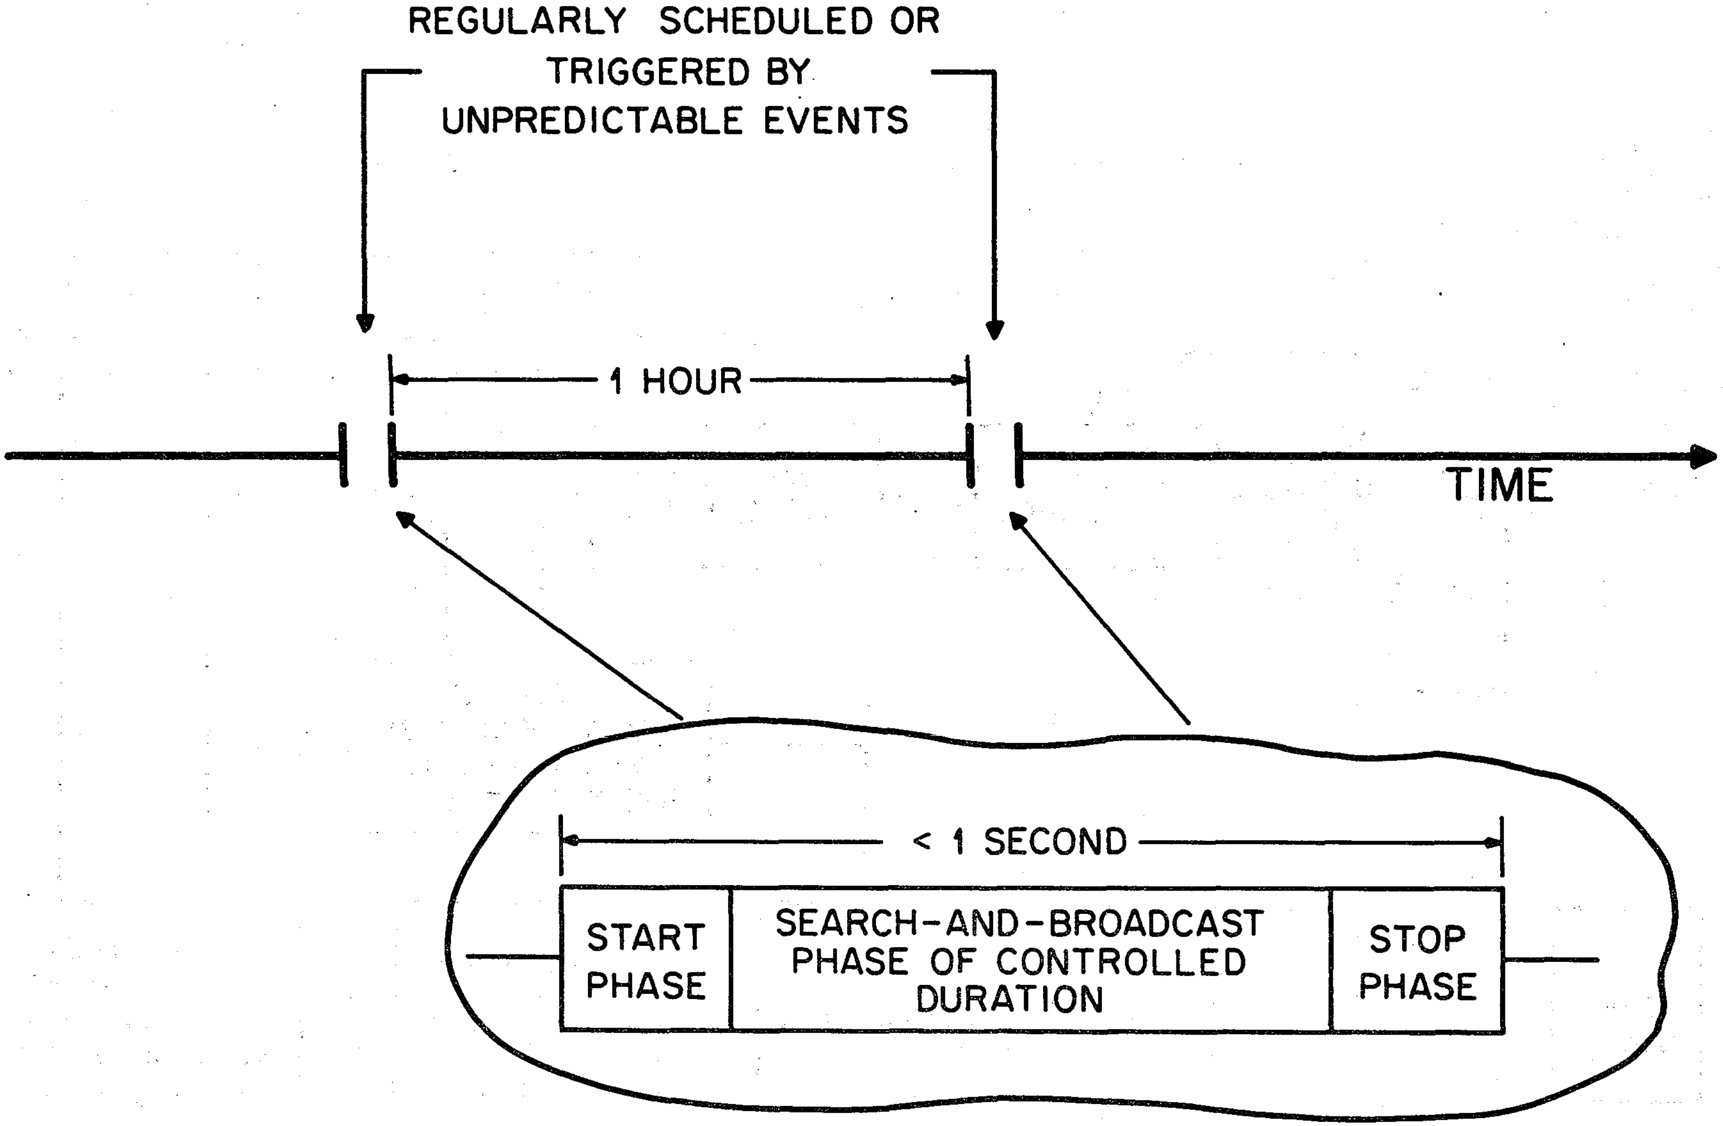
REGULARLY SCHEDULED OR TRIGGERED BY UNPREDICTABLE EVENTS
1 HOUR
TIME
< 1 SECOND
START PHASE
SEARCH-AND-BROADCAST PHASE OF CONTROLLED DURATION
STOP PHASE

Figure 11. Search-and-broadcast intervals for network updates.

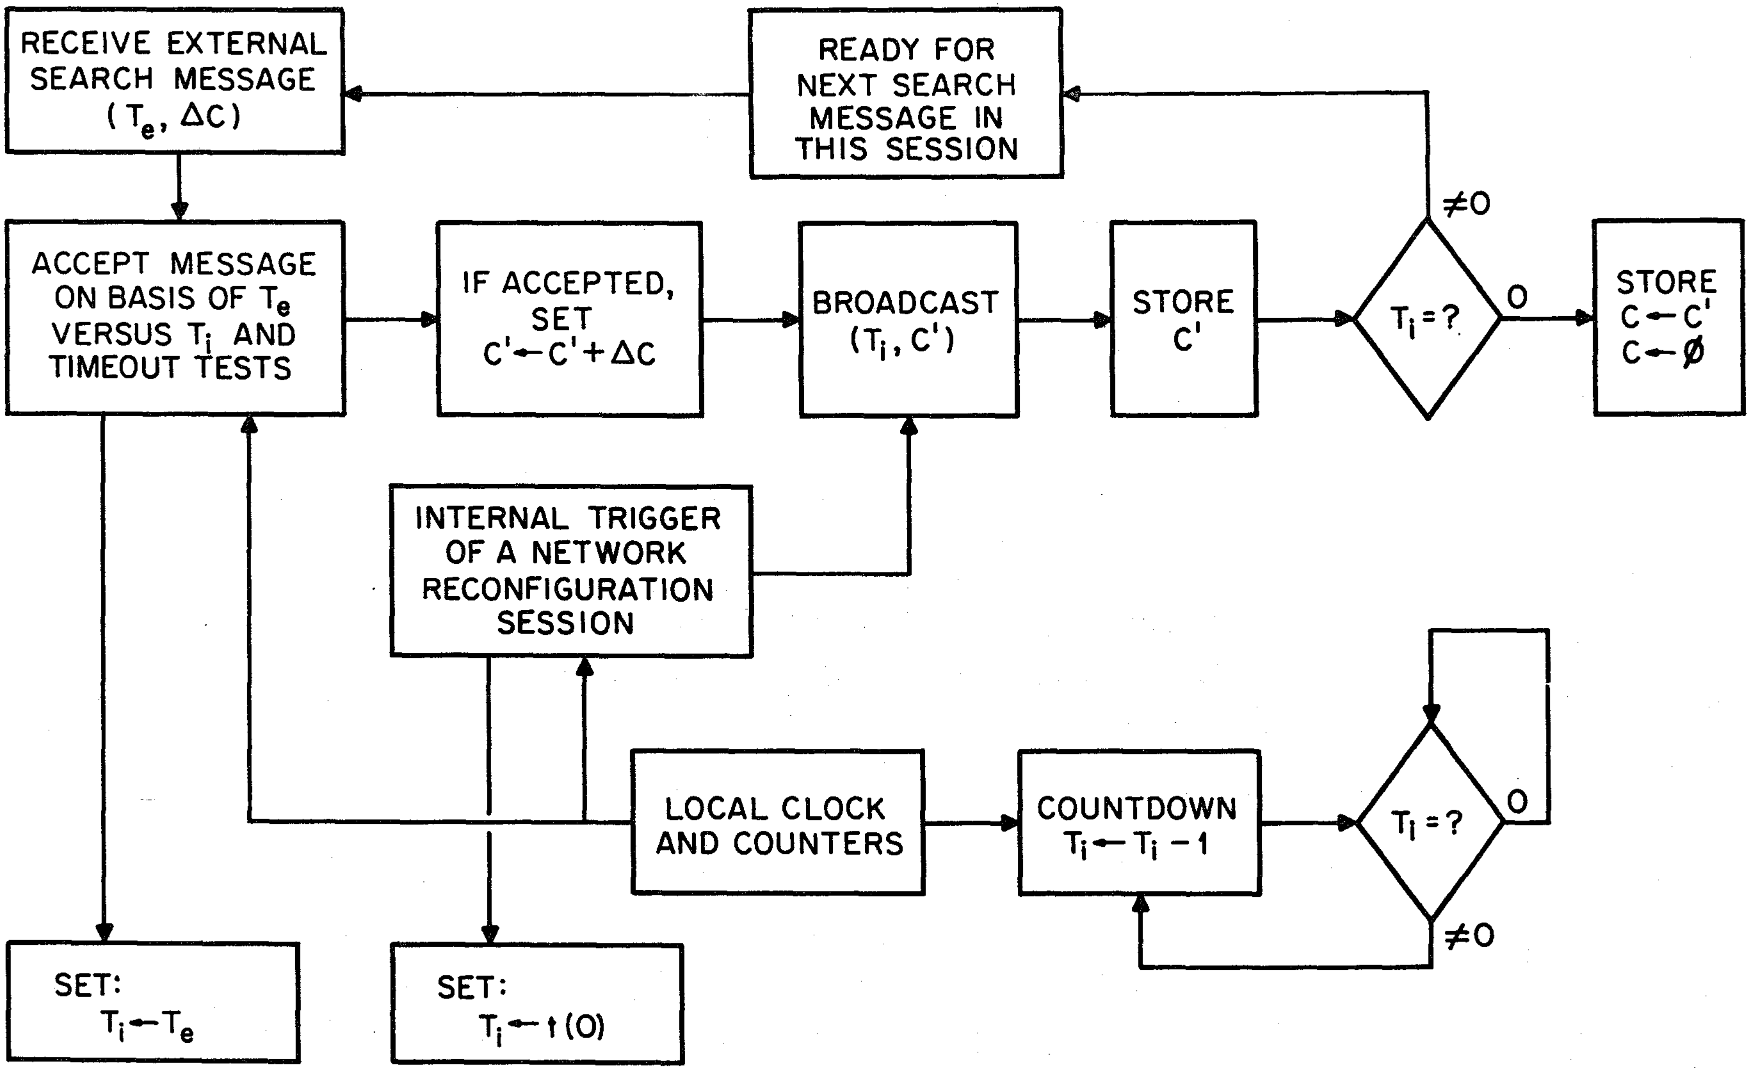
RECEIVE EXTERNAL SEARCH MESSAGE (T_e, ΔC)
READY FOR NEXT SEARCH MESSAGE IN THIS SESSION
≠0
ACCEPT MESSAGE ON BASIS OF T_e VERSUS T_i AND TIMEOUT TESTS
IF ACCEPTED, SET C'←C'+ΔC
BROADCAST (T_i, C')
STORE C'
T_i = ?
0
STORE C←C' C←Ø
INTERNAL TRIGGER OF A NETWORK RECONFIGURATION SESSION
LOCAL CLOCK AND COUNTERS
COUNTDOWN T_i←T_i − 1
T_i = ?
0
≠0
SET: T_i←T_e
SET: T_i←t(0)

Figure 12. Node activities during the three update phases.

that internal timer. When a node receives a timer count from any other node, it recognizes that as one of many possible external timers, T_e.

The initiating node begins by setting its own T_i to a specific initial value of $t(0)>1$. Thereafter it counts down from $t(0)$ to 1 while sending and receiving latest network status updates. When its timer hits $T_i=0$, the node settles on the last update as the final one. It ceases the search-and-broadcast business and stops.

The reason for choosing a countdown instead of a countup is not all that significant. Either count can be used, if one has strong preferences. However, countdown does enable the initial node to determine independently and rather arbitrarily the total duration of the session, without worries about threshold settings here and there in the network. The noninitiating nodes are only obliged to count down and to stop at 0, no matter what the circumstances. Intuitively, this is more important for the larger networks with very many nodes, links, and an unknown mix of facility outages.

Summary of the initiating node's activities in Figure 12

After recognizing from a local, internal, or any other valid source a trigger signal that a network reconfiguration session is in order, the node sets T_i to $t(0)$ as mentioned. It broadcasts, that is, it sends out over all outgoing links a two-part message (T_i, C'). Here C' is the initial node's first knowledge about the network's connectivity or other pertinent information. It may therefore consist of very little, perhaps a mere statement that the node itself exists. The node also stores C' in its data base. If there are valid reasons, the store C' function could take place at the same time or even precede the broadcast (T_i, C') function. It then verifies whether T_i is still larger than zero and is now ready for an echo from anywhere in the network.

Several things can happen next. First, the worst case can be a total silence---no broadcast returns from anywhere. The internal timer, that depends on local clocks and counters, would run out, i.e., reach $T_i=0$, and the node would come to the conclusion that it is isolated from whatever remains of the network. Second, an echo could materialize, but wrong in some respect and therefore not acceptable. Examples of this may be completely garbled texts, bad formats, or strange timer counts, now referred to as external T_e, that are out of range for the session triggered by the original $t(0)$. Such erroneous messages would be ignored by the initiating node, while its own timer would continue the countdown. And finally, the third possibility: a good-looking message can arrive. That message would also have the two-part format, now denoted as (T_e, ΔC). A deciding factor either to accept or to reject the new message could be the comparison of T_e with T_i. If the timers agree, the message is accepted, and in the network status table C' is replaced by whatever new is learned from ΔC. One denotes this with the symbol $C' <= C' + \Delta C$. Now the initiating node has finally learned something new. It is in a position to rebroadcast that new knowledge to the rest of the network. At the same time it must store the new C' and verify whether the internal timer has or has not expired. If not, and $T_i > 0$, the node activity continues by remaining within the loop shown in the upper part of Figure 12. If $T_i = 0$, an exit from the loop takes place. As a final act, the permanent network connectivity status C is replaced by what has been learned in the update session. One denotes that by $C <= C'$. And, of course, in preparation for the next such update session, it is usually a good practice to clear the C' table. That is indicated by the right-to-left arrow, $C' <=$ in the diagram.

Summary of the noninitiating node's activities in Figure 12

The activity chart of Figure 12 pertains to all nodes, including those that do not initiate a reconstitution session. Before a broadcast is received by such a secondary node, the pertinent node restoration process resides in a passive state with an empty update table $C' =$. The passive state is the "Ready for Next Search Message" state on top center of Figure 12.

When such a node receives an apparent externally generated search message, it first validates its format, (T_e, ΔC). If it is not valid, the message is ignored and the node remains in the same passive state. If the message is valid, the node joins the search-and-broadcast session in progress. It sets its internal timer equal to the received count, namely $T_i <= T_e$. It starts its own countdown towards zero, and it stores the just received network status information per $C' <= C' + \Delta C$. It joins the flood of broadcasts by sending over all links, including the one over which the last message was received, its newly constructed version of (T_i,C'). The secondary node is now in the loop. It does all the things called for in Figure 12 and in the proper sequence. It checks whether the internal timer is larger than zero and thus continues the session. Or if $T_i=0$ is true, it exits the loop and executes the $C<=C'$, $C'<=\phi$, tasks of the stop phase.

Generalizations beyond connectivity

Restoration process depends strongly on the establishment of the existing connectivity status. But as noted elsewhere, other information may also be needed. In fact, one, two, three, ... , arrays of different data may be involved in specific network restoration schemes. For brevity, we call this the vector or super-vector of arrays, V. The search-and-broadcast message format then can be visualized as the status of two quantities:

{ Counter, Vector V }.

Connectivity C can be generally thought as part of V. However, since knowledge of C evolves from and contributes to the basic search-and-broadcast process itself, it appears expedient to emphasize certain parts of C as an addendum to V. If so, the format is also equivalent to:

{ Counter, Vector V, Connectivity C }.

The latter view is taken in the expanded message format demonstration of Table 1. It is assumed that an arbitrary node, N, initiates the process by setting the timer to t(0)=i and by broadcasting the message inside the { ... } parentheses. That message is: i, the original count; $V(0|N)$, information vector V at initial step 0 for node N (This initial V may be nothing more than "I am N and this is a 10 Mb/s line."); and [N(0)>-], which identifies this as the broadcast from node N at step 0. Inside the square brackets, the symbol > stands for messages being transmitted from the node on the left to the node on the right. The - denotes the unknown remains of the network.

At the next step, the timer is at i-1 and the vector V is variously augmented at those nodes, N, where the broadcast has been heard and is being forwarded. Symbolically, V becomes $V(1|N)$. The next symbol, [N(0)>N(1)], shows that there is a link from the left to the right node. Symbol [N(1)>-] shows that at step 1 the broadcast continues. And so forth, until the timer reaches zero and the broadcast ceases. At every node that was involved in the procedure, $V(i|N)$ is the final state of the super-vector. In the set of

Table 1. Expanded Search-and-Broadcast Message Format at an Arbitrary Node

Counter		Message	
t(0)=i	"Start"	{i:	V(0|N), [N(0)>-}
i-1		{i-1:	V(1|N), [N(0)>N(1)], [N(1)>-]}
i-2		{i-2:	V(2|N), [N(0)>N(1)], [N(1)>N(2)], N(2)>-]}
•			•
•			•
•			•
•			•
1		{1:	V(i−1|N), {N(0)>N(1)], [N(1)>N(2)], ..., N(i-2)>N(i-1)], [N(i-1)>-]}
0	"End"	{0:	V(i|N), [N(0)>N(1)], [N(1)>N(2)], ..., [N(i-1)>N(i)]}

[N(k)>N(k+1)]'s, the same node can occur many times. The same pair of nodes can also show up in different brackets and on reversed sides of the arrow, >. When that happens, a bi-directional (i.e., FD) direct channel exists between the pair. In the text to follow, such links will be identified with the = symbol between the nodes.

An example

An illustration of the process starts with a damaged state for the earlier network of Figure 5. Assume that stress events have caused outages of nodes #17 and #19. What remains is a topology of 33 links and 18 nodes. That to-be-restored damaged network is shown in Figure 13. Let it be node #8 that first senses local connectivity problems and commences the restoration process. This fact is indicated by N(0)=8. To keep the search-and-broadcast example short and simple, assume that the initial counter is set by local controls to t(0)=3. The subsequent string of events is listed in Table 2 and proceeds as follows:

* With the count 3, the initial node N(0)=8 broadcasts the first message, including its own ID in [8>-].

* At count 2, only node 2 has received, stored, modified, and is in the process of broadcasting the updated message. For all it knows, the network may be nothing more than a directed link from node 8 to node 2, plus what else was found in vector V.

* At count 1, three nodes are active. They are nodes 1, 8, and 11. Node 1 knows that 8 can send to 2, and that 2 can send to 1. It proceeds to broadcasts that fact. Node 8 learns that it can send to and receive from 2. It equates that, [8=2], with an FD link. But it continues broadcasting everything it knows because the count is larger than zero. Node 11 finds out that 8 can go to 2, and that 2 can go to 11. It also continues broadcasting.

* At count 0, the final state is reached. Only a set of nine nodes (i.e., 1,2,3,7,8,11,13,14, and 15) have either received or transmitted messages. All those activities stop now. This feature is recognized in the message format by the absence of [...>-] at the

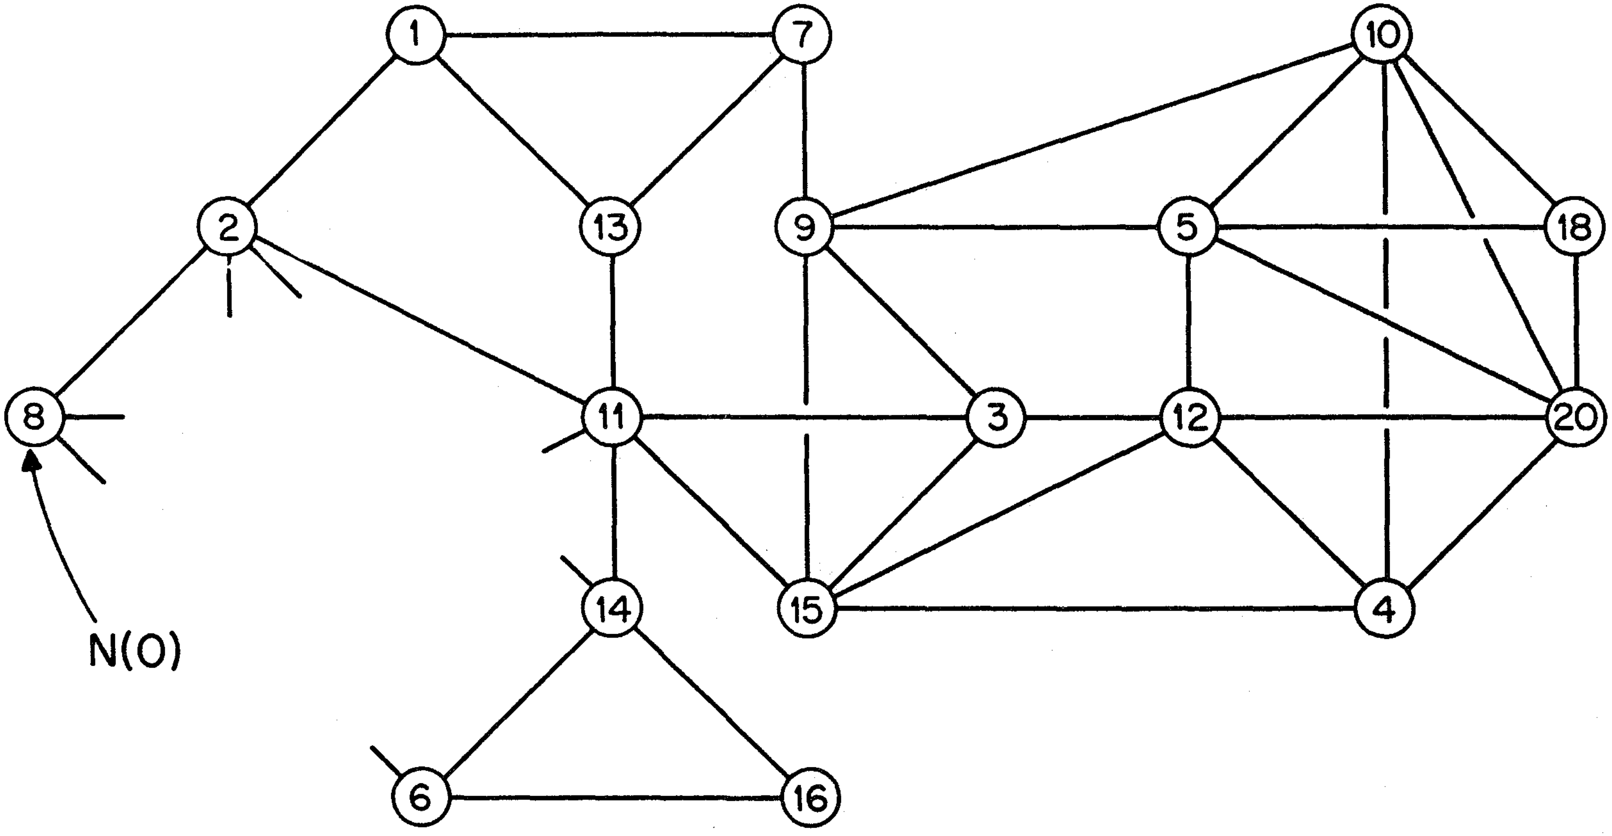

Figure 13. The same network as in Figure 5, but with nodes 17 and 19 disabled.

Table 2. Search-and-Broadcast Example for N(0)=8 t(0)=3, and Nodes 17, 19 Disabled

<table>
<tr><th>Counter</th><th>At Node</th><th>Message</th></tr>
<tr><td>3</td><td>8</td><td>{3: V(0|8), [8>-]}</td></tr>
<tr><td>2</td><td>2</td><td>{2: V(1|2), [8>2], [2>-]}</td></tr>
<tr><td>1</td><td>1</td><td>{1: V(2|1), [8>2], [2>1], [1>-]}</td></tr>
<tr><td></td><td>8</td><td>{1: V(2|8), [8=2], [8>-]}</td></tr>
<tr><td></td><td>11</td><td>{1: V(2|11), [8>2], [2>11], [11>-]}</td></tr>
<tr><td>0</td><td>1</td><td>{0: V(3|1), [8>2], {2>1]}</td></tr>
<tr><td></td><td>2</td><td>*{0: V(3|2), [2=1], [2=8], [2=11]}</td></tr>
<tr><td></td><td>3</td><td>{0: V(3|3), [8>2], [2>11], [11>3]}</td></tr>
<tr><td></td><td>7</td><td>{0: V(3|7), [8>2], [2>1], [1>7]}</td></tr>
<tr><td></td><td>8</td><td>*{0: V(3|8), [8=2]}</td></tr>
<tr><td></td><td>11</td><td>{0: V(3|11), [8>2], [2>11]}</td></tr>
<tr><td></td><td>13</td><td>{0: V(3|13), [8>2], [2>1], [1>13], [2>11], [11>13]}</td></tr>
<tr><td></td><td>14</td><td>{0: V(3|14), [8>2], [2>11], [11>14]}</td></tr>
<tr><td></td><td>15</td><td>{0: V(3|15), [8>2], [2>11], [11>15]}</td></tr>
</table>

* Locally completed

very end. In this example, most nodes have learned very little about the remaining network connectivity. Thus, node 1 still has not established whether the links between 8 and 2, or 2 and 1, are FD or not. Other nodes, like 4, 5, 6, ... , are not even aware that a search-and-broadcast session had been conducted. This is because the original counter was set so low. Only two nodes, namely nodes 2 and 8, have completed their search to the point of knowing their full local connectivity.

One can group together all those nodes that actually get involved in a session. That group is roughly centered around the initiating, also called the primary, node. The size of this group depends on the setting t(0) and on the speed of the timer. (The speed can be assumed to be the same for all nodes.) Given smaller size networks, one can require that the coverage group be the entire network. Efficient designs of such a strategy may call for timing that is necessary and sufficient:

(a) To find the desired connectivity information of the entire topology.

(b) To distribute the whole information parcel to every corner of the network.

Assuming that (a) and (b) are satisfied, every node in the small network knows the connectivity properties of every other node in the network. Since this feature cannot be expected to extend to many realistically large topologies, there appears to be a natural distinction between the broadcasting coverage properties of small and large networks. We treat the two cases separately.

Small network properties

Different length fiber links and different speed devices at nodes cause different message delays in the network. The well-known solution to this timing problem is by synchronization techniques that coordinate clocks and counters at dispersed locations. To simplify the timing issue we assume from now on a discrete and idealized world of time epochs: where the signal propagation over every link takes exactly one unit of time; where that unit of time is the same as used in the internal timers of all nodes; and where processing delays are negligible at all intelligence handling nodes.

If so, then the number of links traveled in a signal path is indistinguishable from the time consumed, and vice versa. In particular, the initial timer setting T_1=t(0) is equivalent to a light signal hopping or traversing t(0) links sequentially, in any order, and in any direction. To proceed one needs notation and several definitions. These are given next:

ℓ = Number of links in the network. Normally these will be bi-directional full duplex (FD), although there might be reasons at some time (e.g., when faced with other transmission media) to violate this rule.

n = Number of nodes in the network.

d(i,j) = Distance between nodes i and j. For all the paths possible between i and j, there is at least one with the smallest number of hops. That number is the distance between the two points. While theoretically there could be nodes separated by infinite distances, the practical networks considered here consist of node pairs with finite distances.

d = Diameter of the network. From all the node pairs (i,j) in the network, there is at least one pair with the largest distance d(i,j) between the two points. Understandably, that largest distance is called the diameter of the network. For the realistic networks considered here all diameters are finite integers.

dmax(X) = Largest possible diameter of any network that is required to satisfy constraint X. If the constraint refers to the number of nodes and links, it is indicated as (X)=(n,ℓ). If only one or the other is known, the argument is either (n,.) or (.,ℓ).

[r] = Integer part of the real number r. As an example, [3.14]=3.

There are a number of potentially useful network properties that involve the above entities. The following properties are presented without proof:

PROPERTY I. The time needed for a complete search-and-broadcast session in any network is either 2d or 2d+1.

PROPERTY II. For all constrained networks with $n \geq 2$ nodes and $n-1 \leq \ell \leq n(n-1)/2$ links, realizable upper bounds can be placed on the largest possible network diameter. Certain known bounds depend on the constraints as follows.

If n is fixed, then $\mathrm{dmax}(n,.) = n-1$.

If ℓ is fixed, then $\mathrm{dmax}(.,\ell) = \ell$.

If both n and ℓ are fixed,

then $\mathrm{dmax}(n,\ell) = n-[r]$,

where $r = \{1+\sqrt{9+8(\ell-n)}\}/2$.

PROPERTY III. Under the assumptions made and excluding time spent in bypass networks or gateways, it suffices to set the initial timer to

$$t(0) = 2n - [\sqrt{9+8(\ell-n)}],$$

where either known values or upper bounds are used for n and/or ℓ.

Property I establishes that the initial timer setting depends almost entirely on the network diameter, when the objective of the session is to achieve complete connectivity status exchange over the whole network. A network that is under stress, unfortunately, can have indeterminately many outages of links and nodes and thus an unknown diameter.

Property II attempts to place maximum values on the diameter, given specified node or link counts. Assume that an undamaged network is as illustrated in Figure 5. It has n=20 nodes and ℓ=41 links. Before any stress or damage occurs, such a known topology has a known diameter of d=6. Starting from any node as the search-and-broadcast trigger, therefore, the initial timer setting of t(0)=13 will suffice. However, what if the connectivity is unknown, while the node and link counts are the same 20 and 41, respectively? Under such premises, Properties II and III assert that a considerably higher dmax(20,41)=13 and t(0)=27 may now be necessary for some topologies.

Next assume that an undetermined number of links have been disabled, but that due to some reason, such as hardening, all the n=20 nodes have survived. Property II now admits that an even higher dmax(20,.)=19 may occur, whence a t(0)=39 may be advisable.

Finally, let there be reliable information that no more than 8 of the 20 nodes could have survived, with link status totally unknown. One could stay with the original link count $\ell=41$, but that would be unrealistic. No more than 8*7/2=28 links are possible for n=8 nodes. If one lets n=8 and $\ell=28$, one obtains dmax(8,28)=1 and t(0)=3, as should always be expected for a fully connected network. Hence it may be safer to apply only the node constraint and deduce dmax(8,.)=7, t(0)=15.

Properties of larger networks

Due to network size or related reasons a complete network investigation is not attempted here. Instead, the search-and-broadcast activity is to cover a certain limited part of the network. That part can consist of geographically adjacent or distant sites, the latter being imbedded in a set of other nodes for which the restoration authorities have little or no concern. For ease of graphical representation we shall visualize the so involved nodes as being adjacent to each other. Again, there is to be a primary node that initiates the session. Other nodes have a secondary role in that they merely react to the received search broadcasts. The activities of both types of nodes can proceed according to the scheme shown in Figure 12.

The sizes and shapes of the restoration search regions result in geometric constructions akin to that of circles and ellipses in finite point topologies. Figure 14 depicts in a very stylized and symmetric way some of the issues associated with coverage regions in a large grid network. The choice of the regular square connectivity between nodes is not particularly important. If the initial timer is set to t(0)=6 at the primary node A, the search can propagate only six hops from A and not beyond. That boundary is identified as the outer shell and it encircles point A as shown. Nothing propagates past it.

But, inside the outer shell the exchange of status messages is not at all uniformly good. Nodes at or on the outer boundary manage to hear only the updates that have been generated on the most direct single route or on several routes from A. Such nodes may learn little about most other nodes inside the outer shell. On the other hand, when a secondary node is near enough to the

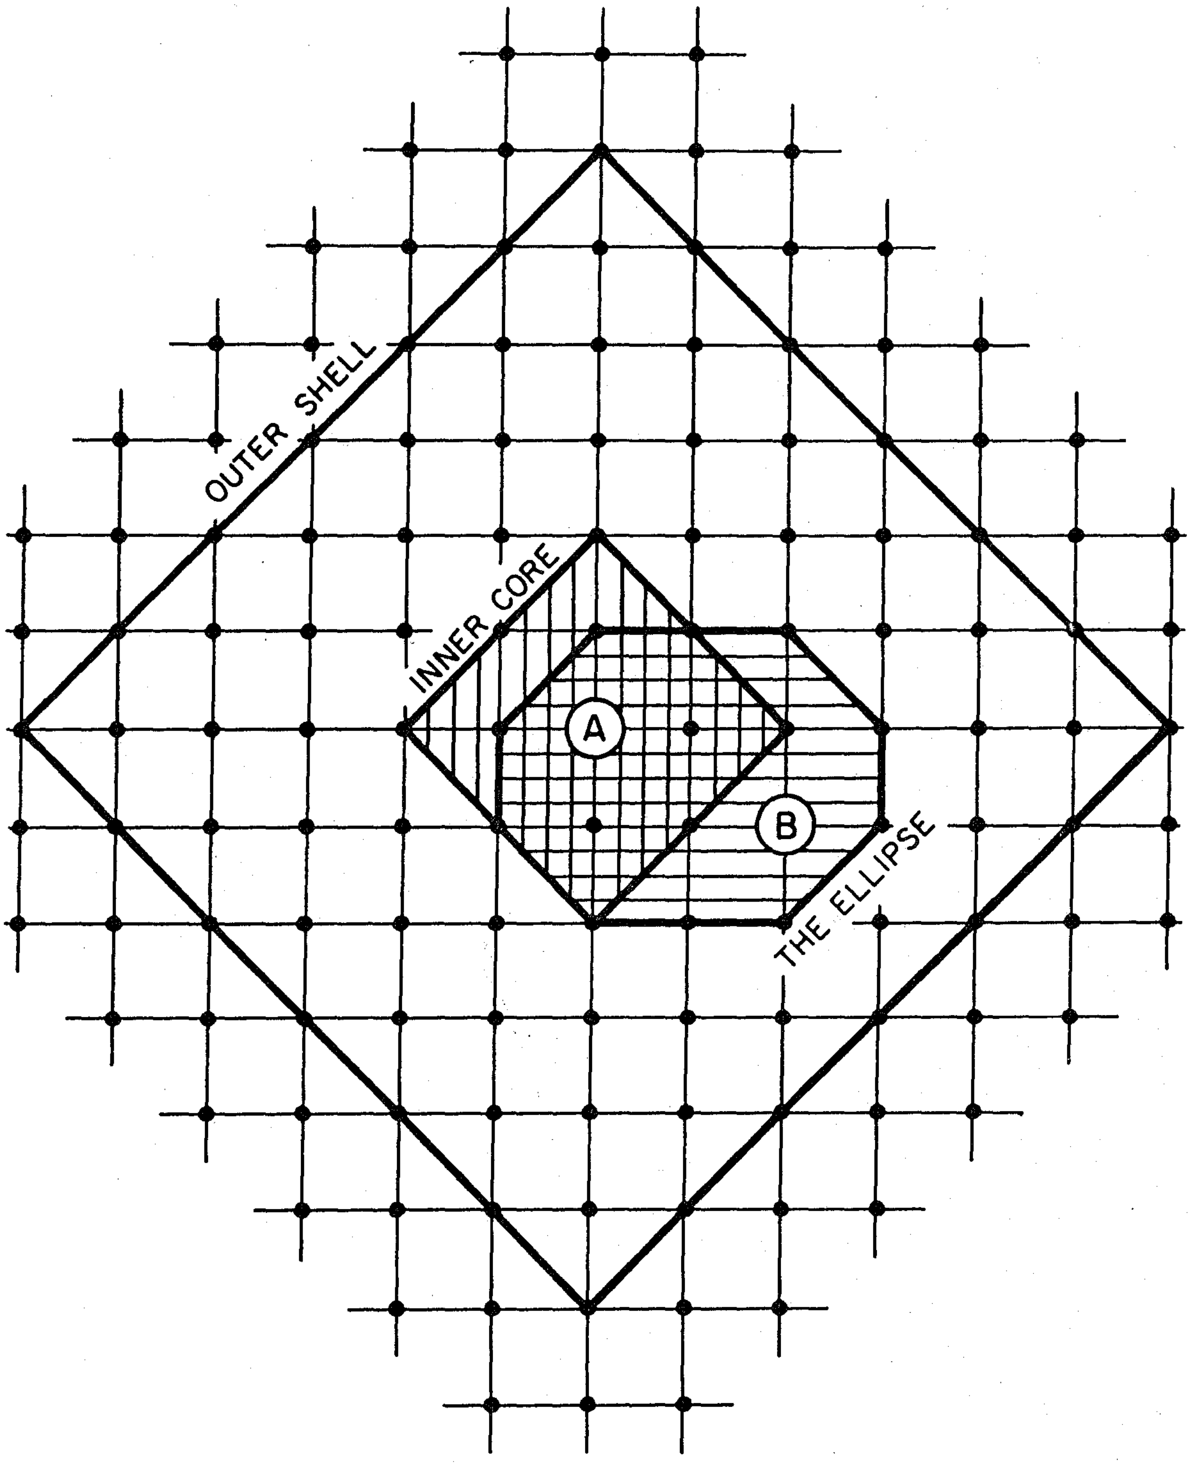

Figure 14. The coverage ellipse for t(0)=6 in an idealized topology.

primary node A, it may learn everything there is to know about a limited neighborhood containing itself and A.

Two lesser neighborhoods surrounding node A appear to be of significance. In Figure 14 they are called the "inner core" and the "coverage ellipse."

The inner core has the property that every point in it learns the connectivity status of every other node in the core, and vice versa. For instance, the top vertex of the core learns that the furthest bottom vertex is connected to the rest of the core after the search message from A has hopped twice down and four times up, which is allowed by t(0)=6. For the initial timer setting of t(0)=x, the inner core consists of all nodes that are at the distance equal to or less than [x/3]. In that sense the core resembles a circle. The corresponding core radius in Figure 14 is 2.

The coverage ellipse has related properties. In addition to the primary node A it includes a given secondary node B. One defines the ellipse for points (A,B) as the set of all nodes whose sum of distances to A and to B is less than or equal to a given constant. The two points A and B are thus quite analogous to the focal points of ordinary planar ellipses. In the case of Figure 14, that constant is t(0)=6. When selecting points (A,B) and the timer setting t(0), one must satisfy $d(A,B) \leq t(0)$. The usefulness of the ellipse comes into play when the primary node A wishes to establish routing to an important secondary node B. All routes within the ellipse that terminate on A and B can serve to connect the two.

Depending an restoration details, one may wish to treat either the inner core, or the coverage ellipse, or both, as a manageable small subnetwork. In that case, the initial timer settings for the small subnetwork are to be determined by the desired subnetwork diameter. The known, unknown, and assumed topological properties of the subnetwork, such as the numbers of nodes or links, affect the session duration as per Properties I, II, and III listed in the section devoted to the small network properties. They are useful to ensure that in a large network an intended small local search does not mushroom into an unbounded global affair.

3.2 Scope of Automation

For fast service restoral under stress, automation must be designed into the NSEP/FOCS nodes. To be useful and effective, the latest and fastest technology (e.g., VLSI and parallel processing) should be distributed throughout the network. Centralization of automation and processor power into a few switches or super-computer centers should be avoided because of clear susceptibility threats. As emphasized earlier, there is no fiber bandwidth-caused pressure to minimize or to worry about the optimum number of bits exchanged between nodes. The fiber has plenty of data throughput capacity, as long as it is not destroyed or seriously impaired.

For the above reason, number crunching of the same or similar restoration data should not be duplicated at different nodes. Instead, results deduced at one node should be broadcasted to all interested nodes and quickly stored away. Old (or not assuredly relevant) data should be simply erased. One can perceive this broadcast process as something far richer than a mere test and acknowledgement of link connections between nearest neighbors. To emphasize that point, we will next outline a menu of network status descriptors. The menu will be quite abundant in that it contains far more than what is currently used in existing network status monitoring and reporting. It may also be more than enough for most fiber optic networks.

The general network status specification, as to be done by vector V earlier, is divided into six information fields. They are:

(1) Network Link Capacities.

(2) Connectivity Matrix.

(3) Traffic Carried.

(4) Traffic Offered.

(5) Facility Restriction.

(6) Routing Tables.

If one wants to expand further, fields such as node capacities, gateway status, plus others, could be added to the network status array. What follows is a brief discussion of these information fields. The discussion covers the scope, sizing, nomenclature, and automation oriented methods to handle their data in present and future fiber optic networks.

(1) Network Link Capacities

The Link Capacity is abbreviated as LC. For a network of n nodes, one can conceivably be faced with asymmetric one-way (i.e., simplex or half-duplex) links between any of the directed n(n-1) node pairs. In such a situation, an n X n array or matrix more than suffices to represent the LC status of a network. Most networks, especially those with fiber optic links, are fortunately full duplex (FD) and with a rather sparse connectivity. Then the number of links ℓ may be significantly less than n(n-1)/2, and the full-blown matrix consists mostly of zeroes. We prefer to represent LC as an ordered list or as a table of real numbers. Each number depicts the useful data rate on a particular link.

A particular node in a larger network can have different involvements with LC tables. Thus, the node may process and store LC's for every near and far corner of the network. This is called the Global option or (G). Or the node may be entirely satisfied with the knowledge of LC values for a limited local region. This is called the Local option (L). Finally, there may be a total absence of interest in LC data. This vacuous option corresponds to None Specified and whenever pertinent is indicated by the (-) symbol.

(2) Connectivity Matrix

The Connectivity Matrix is abbreviated as CM, or in equations simply by C. It is typically a symmetric binary matrix of n X n dimensions. If there is a direct link between nodes i and j, then the (i,j)-th element in the matrix is 1. If there is no direct connection, the element is 0.

Just as in the case of the LC table, the CM can have coverage that extends over the entire network. Then it is clearly Global (G) in character. Or its range may be limited to a Local (L) domain. The most common example of this arises when the node in question, perhaps little more than a concentrator, knows only the directions of its terminating trunks. If the node is totally ignorant about its own connectivity in the network and about its nearest neighbors, the (-) symbol is used.

(3) Traffic Carried

Traffic Carried is abbreviated as TC. We envision TC to be a table, where individual entries describe the traffic on individual links. The TC table therefore has length ℓ. The entries in the table can be further split into two information fields:

* The total amount of traffic carried on the links.

- The distribution of the traffic types carried (e.g., voice, data, packet switched, circuit switched, etc.).

Several real numbers are needed for TC representation. One number may be enough to show the link congestion from the actual carried load in terms of any relative or absolute units. The customary traffic engineering approach is to design the TC loads for the Busy Hour (BH). In the NSEP/FOCS stress scenario, equivalent "stressed BH's" may have to be considered. Moreover, some three or four real numbers appear adequate to depict the mix of traffic types.

The distinction between Global (G), Local (L), or None Specified (*), TC information exchanges may have to be made here just as it was done for LC and CM earlier.

(4) Traffic Offered

The abbreviation for Traffic Offered is TO. It measures the arriving traffic at individual nodes and is expressed in appropriate units (e.g., Mb/s) per node. Therefore TO is a table of n items, one item for each node. As noted before for TC, several real numbers per item serve to model TO. They are needed to describe:

- The total offered load per node.
- The distribution of different TO loads.

A new distinction to be made here is the conceptual separation of "exterior" TO that comes to a node from without and "internal" TO that arrives from anywhere within the network. The effect of TO is to test the ability of a node to handle the traffic and, almost simultaneously, to determine how effectively the terminating transmission links manage to carry the load. Failure of either function results in some manifestation of congestion, which can involve message buffering, rerouting, delays, and partial blockage with associated service degradation or even message loss.

In a stable network under benign conditions, all or nearly all offered nodal traffic should be promptly carried by the links. Then there is a negligible message loss and TO is mapped onto TC. Unfortunately, this case may have little relevance to anticipated stress scenarios.

Finally, several numbers are needed for representation of traffic amount and types at each node. And the previously introduced specifiers, (G), (L), and (-), distinguish between global, local, and "none" modes of TO automation, respectively.

(5) Facility Restriction

Facility Restriction is abbreviated as FR. It pertains to different message classifications, priorities, and security protection levels. Sensitive service requirements may call for access screening and preemption, such as MLPP, or for message and address encryption. There may be reluctance to use certain exposed transmission paths, especially when hardened or safer facility options are available. To satisfy FR needs of a live traffic in a changing network, information about the current FR status must be exchanged.

We view FR as a qualifier or a restriction on the network connectivity matrix, CM. It is then a matrix with slightly less than n X n dimensions. To distinguish between several dozen of anticipated restriction classes, individual terms in the FR matrix should be one or two 8-bit bytes. The (G), (L), and (-) automation qualifiers apply.

(6) Routing Tables

Routing Tables are abbreviated as RT. The totality of such tables can provide detailed paths from every transmitting node to every receiving node in the n-node network. Hence the RT consists of an n X n matrix, where the individual elements are one, two, three, or more preferred route listings between a pair of nodes. A listing can be a sequence of nodes or, equivalently, a sequence of labeled links. As such, the elements in the matrix are groups of tables or smaller sub-matrices. They in turn consist of lists of integers, where the length of a list may depend on the diameter of the network and on the routing algorithm employed.

Much has been researched and published on routing strategies in data, as well as voice, networks. As noted, modifications may be appropriate for fiber optic media. First, there is the fundamental issue whether RT's should be locally computed or broadcast through the network.

We prefer the substantive broadcast approach because of its time savings and fiber utilization properties. Second, the spectrum of RT coverage, such as (G), (L), or (-) must be resolved. Third, the automation algorithms themselves must be specified, at least in broad terms. Many algorithms are possible, but rather generally we prefer to distinguish between Fixed (F), Adaptive (A), and what can be called Minimal (M) routing specifications.

Algorithms in class (F) can produce one or more routes, of varied detail, but all fixed for a reasonable period of time. If there are enough different route choices in the (F) list, and if the message handling machines are permitted to try many choices, the scheme (F) may be reasonably survivable. Option (A) relies on the distributed intelligence among the n nodes to guide the message adaptively from one end to another. It can be supported by a variety of survivability arguments in the damaged network scenario. The (M) algorithm is anything minimally acceptable, as the name implies. Given addresses, such as from node A to node Z, and if the message finds itself at some node X, the (M) rule may merely tell where not to go to avoid going around in circles.

Table 3 lists the six network status descriptors and gives a very approximate estimate of their array sizes at a typical node. Sharper sizing estimates require either more specific network designs or arbitrary assumptions about their properties. But even in its present rough form, the table merits a few explanations. To begin, the network is assumed to have ℓ links, n nodes, and a diameter d. As previously stated, there are several loose properties that relate d to both ℓ and n. However, too much depends on what one knows about the network topology. And even then, the dependence is typically manifested as rather imprecise approximations or bounds.

The first two columns in Table 3 repeat the sequential numbering and acronyms for the six components. The third column identifies the array type. It is either a one-dimensional table (i.e., a list or a vector) or a two-dimensional matrix. The elements in the arrays can be real numbers, binary numbers, or non-negative integers. This is indicated in the fourth column. The number of digits for any of the number representations remains to be engineered by network designers. However, one expects that in practice 2- or 3-decimal digit approximations should suffice to convey the magnitudes of LC, TC, and TO. The binary CM and FR cases differ. The direct connectivity question between a pair of nodes can be decided by a single bit, viz., are they linked or not. The facility restriction is more involved. One may have to determine the number of essential services, user classes, and restrictive facility types. Since all of that appears undetermined at this stage, asterisks are entered everywhere in row #5.

Table 3. Sizing of the Status Information Arrays

Component		Array Type	Numbers Used	Approximate Array Size		
#	Acronym			(G)	(L)	(-)
1	LC	Table	Real	ℓ	$2\ell/n$	0
2	CM	Matrix/Table	Binary	$n(n-1)/2$	$n-1$	0
3	TC	Table	Real	ℓ	$2\ell/n$	0
4	TO	Table	Real	n	$(2\ell/n)+1$	0
5	FR	Matrix	Binary	*	*	*
6	RT	Matrix	Integer	$2dn(n-1)$	$4(n-1)$	$n-1$

*Undetermined

Columns 4, 5, and 6 list approximate array sizes for the three strategies: Global (G), Local (L), and None (-). The sizes refer to what is stored at one node in the quiescent state between restoration epochs. Once restoration begins, the six new arrays can start practically from zero and build up to the same or a modified quiescent state.

Take the (G) column first. Since there are ℓ links, ℓ numbers suffice to list all their capacities. Because of the n nodes and the fact that the full n X n matrix contains a redundant symmetry and n main diagonal terms, $n(n-1)/2$ bits actually suffice for the connectivity matrix. Entries ℓ for TC, and n for TO, are based on the premise that only the total traffic volumes, carried or offered, need be monitored. If the previously discussed distinction between different traffic categories, plus their statistics, are to be included, then these sizes should be magnified some 2 to 10 times. Finally, row #6 in column (G) shows a quantity $2dn(n-1)$. This assumes that on the average 4 different routes are to be listed from each of the n origination nodes to any of n-1 terminations, and that the node reversal in the FD environment utilizes the same set of 4 alternate routes. The average length of each route is assumed to be d/2.

The fifth column represents the Local coverage (L). In as much as we do not wish to get involved in what is or is not a local area for arbitrary networks, we offer a rather elementary definition here. For a fixed node A, let all outgoing/incoming links and all directly linked nodes constitute the "local" domain. The array sizes again apply to the volume of numbers needed at an individual node. Consider the LC row. Each link has two ends that terminate at some of the n nodes. It is true that the counts of homing links can and do vary widely among nodes. However, whenever they are added up over the whole network, there must be on the average exactly $2\ell/n$ link ends per node. Their data rates are given here as real numbers. Similar arguments apply to the derivations for the CM, which uses only one row of the connectivity matrix, and to TC, which is similar to LC under (L). In the offered traffic or TO row, the situation differs only slightly. One number depicts the total offered traffic at any node. However, on the average, the node has $2\ell/n$ nearest neighbors. Thus $(2\ell/n)+1$ real numbers represent all TO states within a unit radius from any network node. The latter traffic substream can come from end-user facilities, gateways, or temporary patches to other nets. At the very bottom of the (L) column, again 4 alternate routes are assumed to each of the n-1 potential addressed nodes.

The very last column in Table 3 alleges no specification (-) of restoration numbers almost everywhere. This feature is noted in rows 1 to 4, where empty arrays are shown.

The situation is unclear in the case of facility restriction, that is row #5, where some minimum information must be kept and possibly updated to assure that sensitive information stays on appropriate restricted channels. Similar observations pertain to the minimal routing instructions that may be necessary to avoid looping of paths in the arbitrary, stressed or damaged, network. In row #6 and column (-), n-1 integers are judged to suffice.

As the concluding comment on Table 3, it must be recognized that the 6 component arrays discussed are by no means required to be all of the same category, such as all (G), or all (L), or all (-). Perhaps, some yet to be determined mix of (G), (L), and (-) is optimum for NSEP/FOCS. If one were to sample six times from three categories, there could be a total of 3^6=729 possible different options to be considered. Admittedly, among the universe of 729 choices, only some 50 to 100--or perhaps 10%--would make sense for a fiber optic network. Note: In such a construction one samples LC, CM, TC, TO, and FR from {(G),(L),(-)}, while picking RT from the previously introduced set {(F),(A),(M)}.

Four illustrative choices for this six-dimensional "super vector" of arrays are shown in Table 4. The first example is called "maximal." It supplies global (i.e., over the entire network) information for LC, CM, TC, TO, and FR, as well as enhanced adaptive RT capability for the array vector. This maximal option is likely to be the most expensive. In fiber optics practice, a far more attractive choice is that called "realistic" in Table 4. It has only one global feature, namely the complete connectivity matrix. It keeps track of link capacities and facility restrictions locally only, and it ignores all traffic statistics. As far as routing tables go, it implements a reasonably simple algorithm that relies largely on the available global CM. For lack of a better word, the third illustration is called "example." It retains only two arrays: a local CM and a fixed RT. The connectivity matrix here has a limited range. To find routes beyond that range, the fixed routing table merely indicates in which direction to exit from the local area. The fourth and final option gets rid of all arrays, except for a minimal instruction set for

Table 4. Four Illustrative Choices for the Six-Dimentional Array Vector

$$
V \equiv \begin{pmatrix} LC \\ CM \\ TC \\ TO \\ FR \\ RT \end{pmatrix} : \underbrace{\begin{pmatrix} G \\ G \\ G \\ G \\ G \\ A \end{pmatrix}}_{\text{"MAXIMAL"}} , \underbrace{\begin{pmatrix} L \\ G \\ - \\ - \\ L \\ A \end{pmatrix}}_{\text{"REALISTIC"}} , \underbrace{\begin{pmatrix} - \\ L \\ - \\ - \\ - \\ F \end{pmatrix}}_{\text{"EXAMPLE"}} , \underbrace{\begin{pmatrix} - \\ - \\ - \\ - \\ - \\ M \end{pmatrix}}_{\text{"MINIMAL"}}
$$

routing. The latter is to prevent loops and unacceptably (e.g., infinitely) long message routes in the network.

3.3 Distance Matrix

An entity more appropriate for the final stages of the automated restoration process is the distance matrix, D. It could have been appended to the previous field of the six arrays, which were the topic of section 3.2 and particularly of Table 3. However, its unique, somewhat novel, and potentially useful properties, make a separate section appropriate for D.

For a network of n nodes, one defines D as the n X n matrix whose (i,j)-th element is the distance d(i,j) from node i to node j. For FD links, the matrix D is symmetric about its main diagonal. The main diagonal itself contains all zeroes. In more general, cost- and mileage-dependent networks, the distances between nodes can be arbitrary real numbers. Here, as noted earlier, individual links are said to be of unity length. On any given route, then, the distance traveled is the same as the count of links. For the previous network example of Figure 5, the distance matrix D is shown in Figure 15. For a direct comparison with the related connectivity matrix, the corresponding C matrix for the same network is displayed in Figure 16.

As the restoration process begins and knowledge about C spreads in the network, a similar development pertains to D. However, there is a major difference. Except for some favorable or trivial topologies, the growth of D appears intuitively to start at a much slower pace than does the spread of C data. The problem can be visualized by considering arbitrary nodes N, S, E, and W, in a network with many more than four nodes, but with an unknown topology. If N initiates the search-and-broadcast session, how long does S have to wait before it gets some indication of the approximate distance between E and W? When is S assured that it has obtained the true and final value for the E-to-W distance? If N, S, E, and W are allowed to range over the entire network, what is the required worst-case delay for the final network D data to be dispersed everywhere? It is no surprise that the total required time for D dissemination is the same as for C. That conclusion follows from mapping properties, whereby C uniquely determines D, and vice versa.

The following formal properties apply to C and D:

(1) Matrix D contains n zeroes and they all are on the main diagonal.

j

i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	0	1	3	4	3	3	1	2	2	3	2	4	1	3	3	4	2	4	2	4
2	1	0	2	3	4	2	2	1	3	4	1	3	2	2	2	3	1	5	1	4
3	3	2	0	2	2	3	2	3	1	2	1	1	2	2	1	3	2	3	3	2
4	4	3	2	0	2	4	3	4	2	1	2	1	3	3	1	4	3	2	4	1
5	3	4	2	2	0	5	2	5	1	1	3	1	3	4	2	5	4	1	5	1
6	3	2	3	4	5	0	4	2	4	5	2	4	3	1	3	1	1	6	2	5
7	1	2	2	3	2	4	0	3	1	2	2	3	1	3	2	4	3	3	3	3
8	2	1	3	4	5	2	3	0	4	5	2	4	3	2	3	3	1	6	1	5
9	2	3	1	2	1	4	1	4	0	1	2	2	2	3	1	4	3	2	4	2
10	3	4	2	1	1	5	2	5	1	0	3	2	3	4	2	5	4	1	5	1
11	2	1	1	2	3	2	2	2	2	3	0	2	1	1	1	2	1	4	2	3
12	4	3	1	1	1	4	3	4	2	2	2	0	3	3	1	4	3	2	4	1
13	1	2	2	3	3	3	1	3	2	3	1	3	0	2	2	3	3	4	3	4
14	3	2	2	3	4	1	3	2	3	4	1	3	2	0	2	1	2	5	1	4
15	3	2	1	1	2	3	2	3	1	2	1	1	2	2	0	3	2	3	3	2
16	4	3	3	4	5	1	4	3	4	5	2	4	3	1	3	0	2	6	2	5
17	2	1	2	3	4	1	3	1	3	4	1	3	3	2	2	2	0	5	1	4
18	4	5	3	2	1	6	3	6	2	1	4	2	4	5	3	6	5	0	6	1
19	2	1	3	4	5	2	3	1	4	5	2	4	3	1	3	2	1	6	0	5
20	4	4	2	1	1	5	3	5	2	1	3	1	4	4	2	5	4	1	5	0

Figure 15. Distance matrix D for the network of Figure 5.

j

i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	1	1	0	0	0	0	1	0	0	0	0	0	1	0	0	0	0	0	0	0
2	1	1	0	0	0	0	0	1	0	0	1	0	0	0	0	0	1	0	1	0
3	0	0	1	0	0	0	0	0	1	0	1	1	0	0	1	0	0	0	0	0
4	0	0	0	1	0	0	0	0	0	1	0	1	0	0	1	0	0	0	0	1
5	0	0	0	0	1	0	0	0	1	1	0	1	0	0	0	0	0	1	0	1
6	0	0	0	0	0	1	0	0	0	0	0	0	0	1	0	1	1	0	0	0
7	1	0	0	0	0	0	1	0	1	0	0	0	1	0	0	0	0	0	0	0
8	0	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	1	0	1	0
9	0	0	1	0	1	0	1	0	1	1	0	0	0	0	1	0	0	0	0	0
10	0	0	0	1	1	0	0	0	1	1	0	0	0	0	0	0	0	1	0	1
11	0	1	1	0	0	0	0	0	0	0	1	0	1	1	1	0	1	0	0	0
12	0	0	1	1	1	0	0	0	0	0	0	1	0	0	1	0	0	0	0	1
13	1	0	0	0	0	0	1	0	0	0	1	0	1	0	0	0	0	0	0	0
14	0	0	0	0	0	1	0	0	0	0	1	0	0	1	0	1	0	0	1	0
15	0	0	1	1	0	0	0	0	1	0	1	1	0	0	1	0	0	0	0	0
16	0	0	0	0	0	1	0	0	0	0	0	0	0	1	0	1	0	0	0	0
17	0	1	0	0	0	1	0	1	0	0	1	0	0	0	0	0	1	0	1	0
18	0	0	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	1	0	1
19	0	1	0	0	0	0	0	1	0	0	0	0	0	1	0	0	1	0	1	0
20	0	0	0	1	1	0	0	0	0	1	0	1	0	0	0	0	0	1	0	1

Figure 16. Connectivity matrix C for the network of Figure 5.

(2) Matrix C has n+2ℓ ones, n of them being by definition on the main diagonal.

(3) For a connected network, the diameter d is the largest element in D. The network illustrated in Figures 5, 15, and 16 has the diameter d=6. This is best seen from matrix D.

(4) The network diameter can also be deduced from matrix C. For k=1,2,3, ... , one has d=k, if and only if, k is the smallest integer such that the logical exponent, C^k, is the matrix of all ones. If no such k can be found, the network is either disjointed or it cannot be a finite network. (Note: In logical matrix exponentiation, as in logical multiplication over the binary field, the products and sums of terms are replaced by logical AND and OR functions, respectively.)

(5) Matrix C follows easily from D. If the (i,j)-th element in D, namely d(i,j), is either 0 or 1, the corresponding (i,j) term in C is 1. If d(i,j) is neither 0 nor 1, the term in C is 0.

(6) Given C, matrix D can be deduced by several methods. Any of the methods for finding shortest paths between nodes can be used. An example: Starting from C, the well-known "rooted tree" algorithm at node i generates the entire i-th row of matrix D. Or one can use a formal mix of both logical and arithmetic operations to express,

$$D = dC^d - \sum C^k,$$

where:

- Sum $\sum$ goes from k=0 to k=d-1.

- C^0 = I (the identity matrix).

- As in item (4), exponents of C are logical constructs.

- Scalar products (e.g., dC^d), sums, and differences of matrices are ordinary arithmetic operations over integers.

The computation of D is therefore straightforward, at least in principle, assuming that C is known. A direct implementation of the matrix equation of (6) is outlined in Figure 17. In this diagram, as explained above, the required mix of logical and arithmetic operations is to take place. Thus, d is a positive integer; C, I, and Π are binary matrices; while D and $\sum$ are integer matrices. The ease of the above calculations depends primarily on the size of the network, i.e., on the magnitude of n. Improved algorithmic methods might also help, but that is a topic more appropriate for the next section.

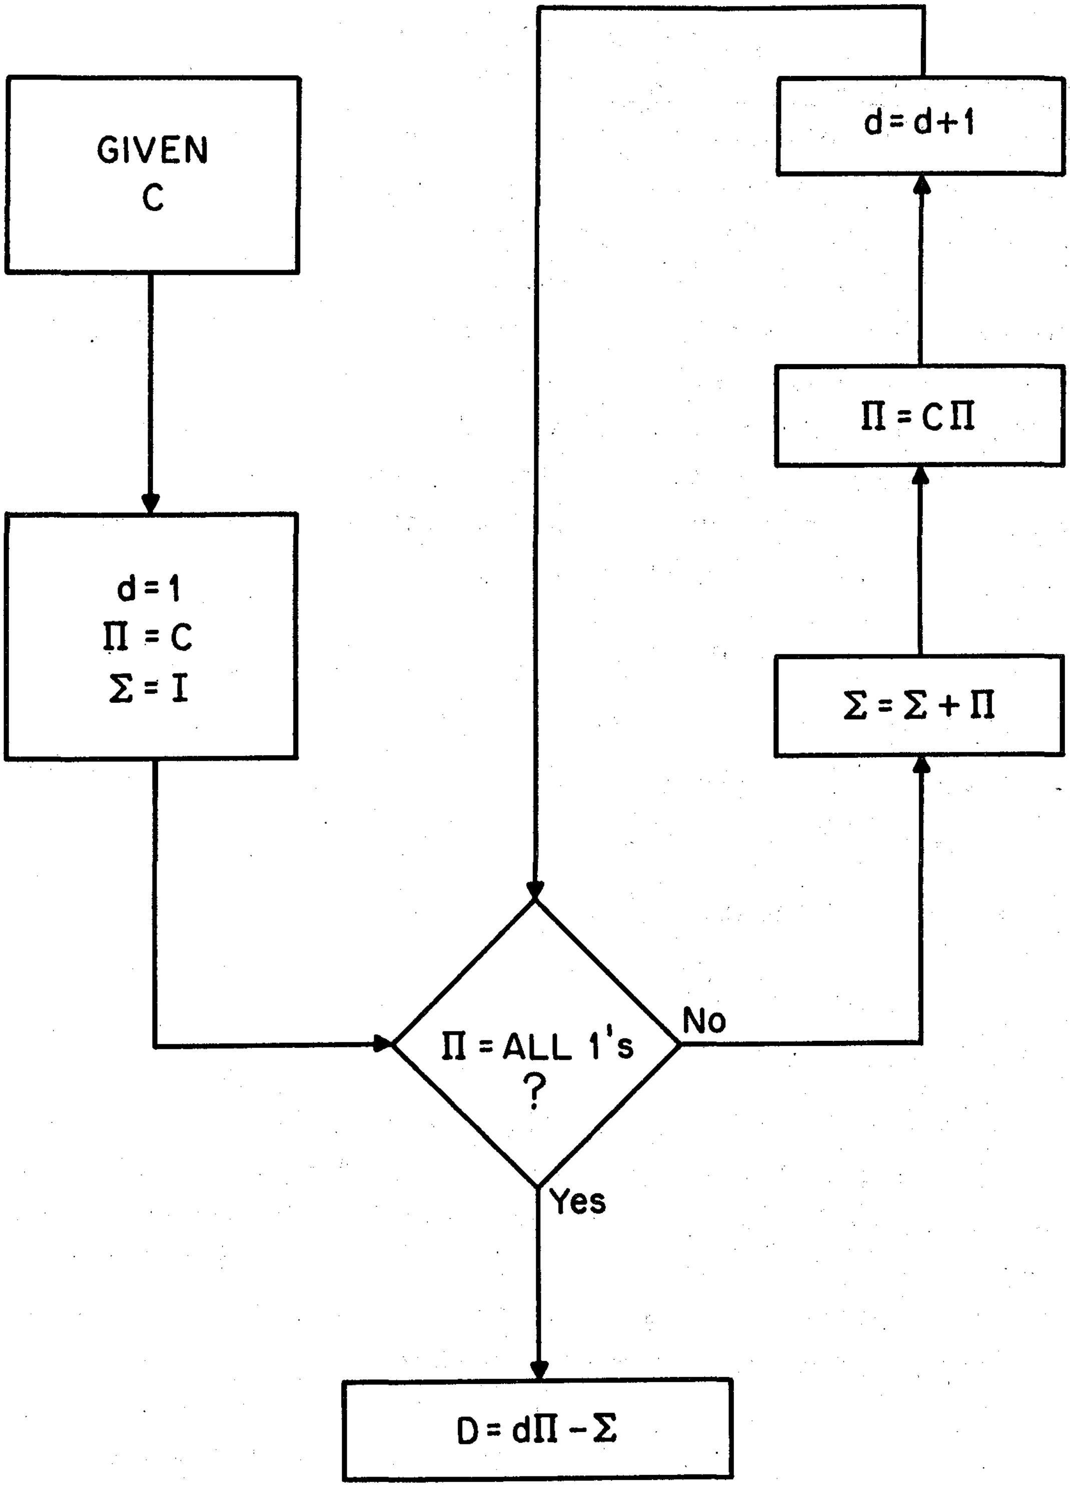

Figure 17. Computation of D, given C.

3.4 Control and Algorithms

Standard reference models and link controls

Controls, protocols, and associated standards play an important role in all domestic and international data and other network operations. Unfortunately, even for fiber optic systems, so broad a topic is simply beyond the scope of this study. In a nutshell then, the FOCS controls should be compatible with and transparent to other essential, present and planned, NCS data communications protocols.

Organizations, such as ANSI, CCITT, FIPS, ISO, plus others, have been working on the so-called Open System Interconnection or OSI network architecture. In its familiar form, this OSI Reference Model consists of seven layers that are numbered from bottom up:

Layer 7	Application
Layer 6	Presentation
Layer 5	Session
Layer 4	Transport
Layer 3	Network
Layer 2	Link
Layer 1	Physical

Despite the efforts of many organizations, progress along the lines of standard terms, designs, and procedures, has been slow for the upper layers. Far more has been accomplished on the lower three layers. This appears fortuitous for the present effort, as clearly it is the physical fiber (Layer 1), the operations of the fiber link (Layer 2), and the operations of the fiber linked network (Layer 3), that are of immediate concern to this program.

The physical characteristics of the fibers, such as stress-resistant engineering, hardness, manufacture, installation, and maintenance practices, are reported in companion reports of this project. That covers Layer 1.

The general end-to-end connectivity is in the domain of Layer 4. However, from the carrier's-carrier point of view associated with FOCS, the network control, whether it is for circuit-switched, packet-switched, frame, or datagram service, belongs in Layer 3. The recently popular packet interface, X.25, appears to be an eventually important future standard. Functionally comparable and relevant to ISDN planners worldwide is the CCITT's Common Channel Signaling System (CCSS) No. 7. Like X.25, CCSS No. 7 involves all three lowest OSI layers. Both systems can be projected unto two operational planes: the user message plane (U-plane) and the network control plane (C-plane). According to this, the planned fiber optic network would handle its end-user data on the U-plane and its restoration, plus other more or less crucial housekeeping tasks, on the C-plane. In as much as a "user" may be another network, the general U- and C-planes may occasionally interact. A significant part of C-plane functions deals with control of data links, thus with Layer 2.

As an example of interest to search-and-broadcast activities over FOCS links, consider a part of protocol history associated with Layer 2. In the early 1970's IBM announced its Synchronous Data Link Control (SDLC), which is still used in IBM's SNA. Then ANSI modified it and called it the Advanced Data Communication Control Procedure (ADCCP). Next, ISO also modified SDLC and named it the High-level Data Link Control (HDLC). Around this time, the U.S. Government gave the ADCCP the status of a Federal Standard. Internationally, CCITT reviewed, modified, and adopted HDLC for its Link Access Procedure (LAP) as part of the X.25 network interface standard. However, during all that time changes were made in the standards, especially in HDLC. To be more compatible, CCITT recommended an updated version and called it LAP B. Other changes also have occurred and continue to occur at present time.

Today, the official version of ADCCP is found in the Federal Standard 1003A --- "Telecommunications: Bit Sequencing of the American National Standard Code for Information Interchange in Serial-by-Bit Data Transmission." Its equivalent versions are also given in the FIPS Publication 71, ANSI Standard X3.66-1979, ISO Standards IS4335 and IS3309, and CCITT X.25 Level LAP B.

Consider an application of ADCCP, or one of its modified versions, to the link protocols of a fiber optic network. Assume that restoration functions in crisis conditions are to preempt the usual message and control activities. Then, at least in principle, the link level accesses at a single automated network node may be as depicted in Figure 18. In the diagram, the node is served by three links. It suffices to view them as physical links, although in some rare instances they could be virtual links. Each link is shown to have a separate message buffer. This need not be so in actual implementation, because

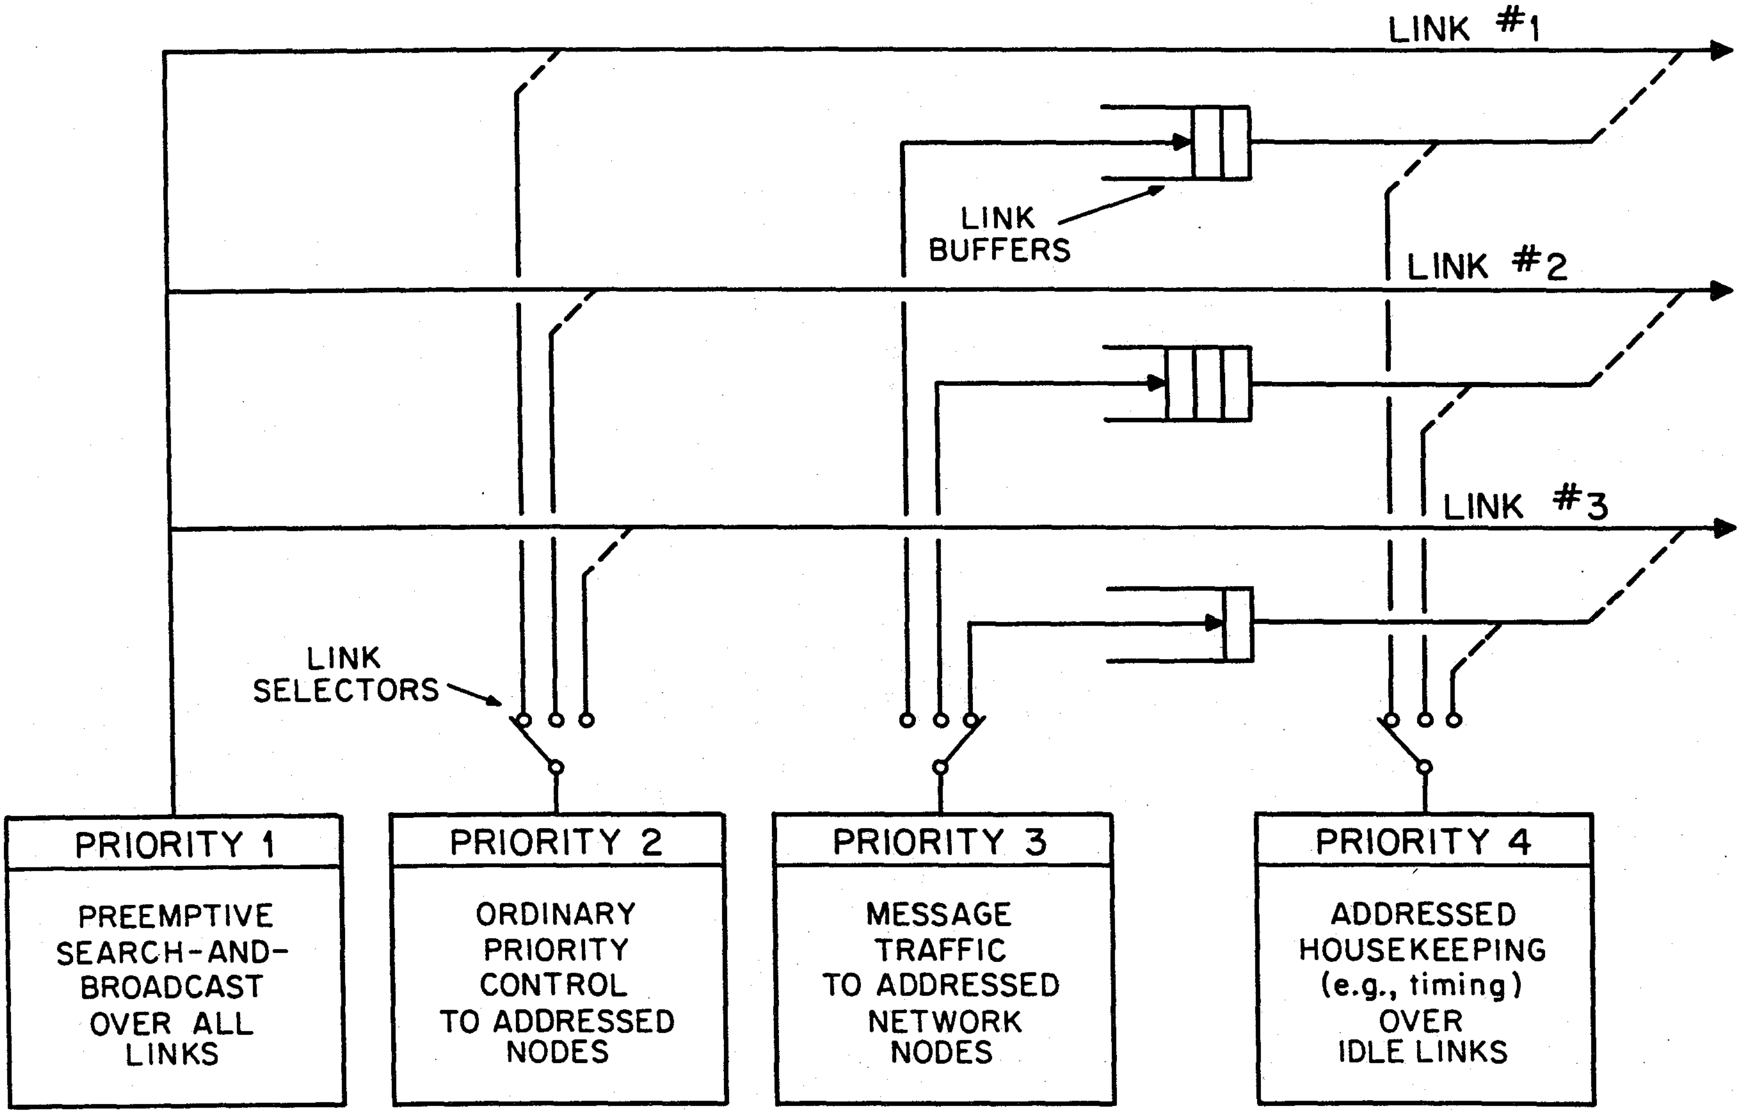
LINK #1
LINK BUFFERS
LINK #2
LINK #3
LINK SELECTORS
PRIORITY 1
PREEMPTIVE SEARCH-AND-BROADCAST OVER ALL LINKS
PRIORITY 2
ORDINARY PRIORITY CONTROL TO ADDRESSED NODES
PRIORITY 3
MESSAGE TRAFFIC TO ADDRESSED NETWORK NODES
PRIORITY 4
ADDRESSED HOUSEKEEPING (e.g., timing) OVER IDLE LINKS

Figure 18. Preemption priorities at an automated restoration node.

end-user messages can come from a variety of MLPP preemptive and priority queueing service classes. However, if the screening of user classifications is done in advance of link access, all but the very highest priority preempting messages can be stored in the node buffers. Then in principle, buffering for link access is as shown in Figure 18.

Search-and-broadcast messages that carry out the network reconstitution are given the highest priority, namely #1. They preempt all other messages. This is indicated by the horizontal solid lines that point to the right. A slanted broken line indicates a lower priority. It can be preempted by traffic on the solid line. Thus, priority #1 preempts #2, priority #2 preempts #3, and so forth.

If the restoration broadcasts under stress have priority #1 and the ordinary message traffic has priority #3, that leaves two levels (#2 and #4) for other network functions, such as perhaps CCS, AO&M/NM, and so forth. These are control functions and they are presumed to belong to one of two priority groups. One group contains all control information, such as link signaling, that is a prerequisite for any message (user data or CCS) transport. This is priority #2. The second control group consists of less urgent deliverables that can wait until the links become idle. Record keeping, facility monitoring, even timing, can be in this priority #4 category. This lowest priority is preempted by every other priority class.

Another noted feature in Figure 18 is the absence of link selector switches for search-and-broadcasting. Broadcasts, by definition, are sent everywhere. All other message types, however, are addressed to a particular destination. They may require specific routes, hence specific exit links from a node.

At this time it seems premature to select one particular message format for FOCS. User and traffic requirements will eventually identify the best format from several candidates. However, a likely outline of the data format can be suggested. Figure 19 shows a general frame format of the ADCCP type. The frame is meant to be of variable length, largely due to the variation in the size of the information field. Thus, it differs from the various fixed packet length systems that have been designed to facilitate interworking.

The frame starts and ends with a unique flag of fixed length. Between the flags, there are four fields: the address field, the control field, the information field, and the Frame Check Sequence (FCS). Usually the address and

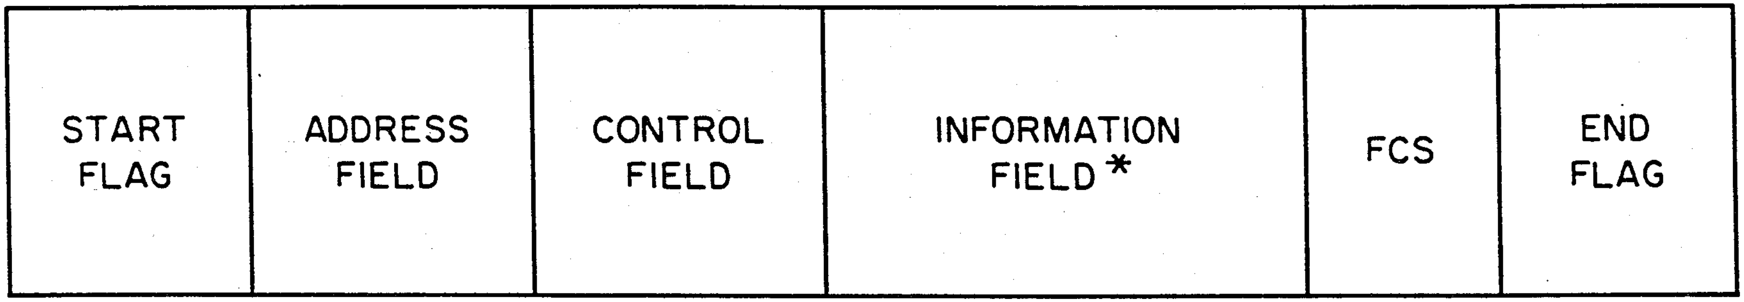

*Variable length

Figure 19. Frame format for ADCCP type of link level control.

control fields are also of fixed length, while the information fields may vary. That feature may be retained in the FOCS, with the understanding that the occasional long arrays needed for restoration and related controls are delegated to the variable information fields of one or more so identified, dedicated, frames. The required frame identification can be done as part of an expanded control field, as we shall shortly see.

The address field is basically filled in by the end users. As expanded in Figure 20, the address field divides naturally into two parts: the sender's address shows where the message comes "from," and the would be destination address is where the message is going "to."

Figure 20 suggests that the control field could be split into three parts. The so-indicated parts are associated with frame (message or user) priority, route specification, plus both "to be continued" and length specifiers for the immediately following variable information field. Included in the priority subfield could be the ID's for restoration and control priorities of Figure 18, as well as the MLPP grades of the more essential users. Thus, the control field would immediately alert all nodes en-route that this particular frame is not an end-user message, or an inter-node signaling package, but instead is a highest priority restoration broadcast. Additional parts of the control field can deal with the above mentioned facility restriction (FR), message or user identity screening, crypto support elements for secure data substreams, and other needed matters as they arise.

The information field has both its contents classified, its length prescribed, and its continuation state assigned by the control field. The contents, as already noted in the discussion of Figure 18, can be one of at least four: the highest priority restoration search, the urgent controls that preempt user messages, the actual user messages (with their own MLPP), and finally the lowest priority housekeeping that in turn is preempted by everyone else.

Finally, the end flag is preceded by a Frame Check Sequence (FCS), also variously called the error detection, redundant parity check, or the cyclic redundancy check (CRC) sequence. The job of the FCS is to verify that all parity bits agree (e.g., they are 0 for even parity). If one or more parity bits disagree, then some bit errors have been detected within the frame and in almost all applications the message cannot be accepted. In the standard ADCCP, HDLC, and similar link protocols, the receiving node uses the FD capability of

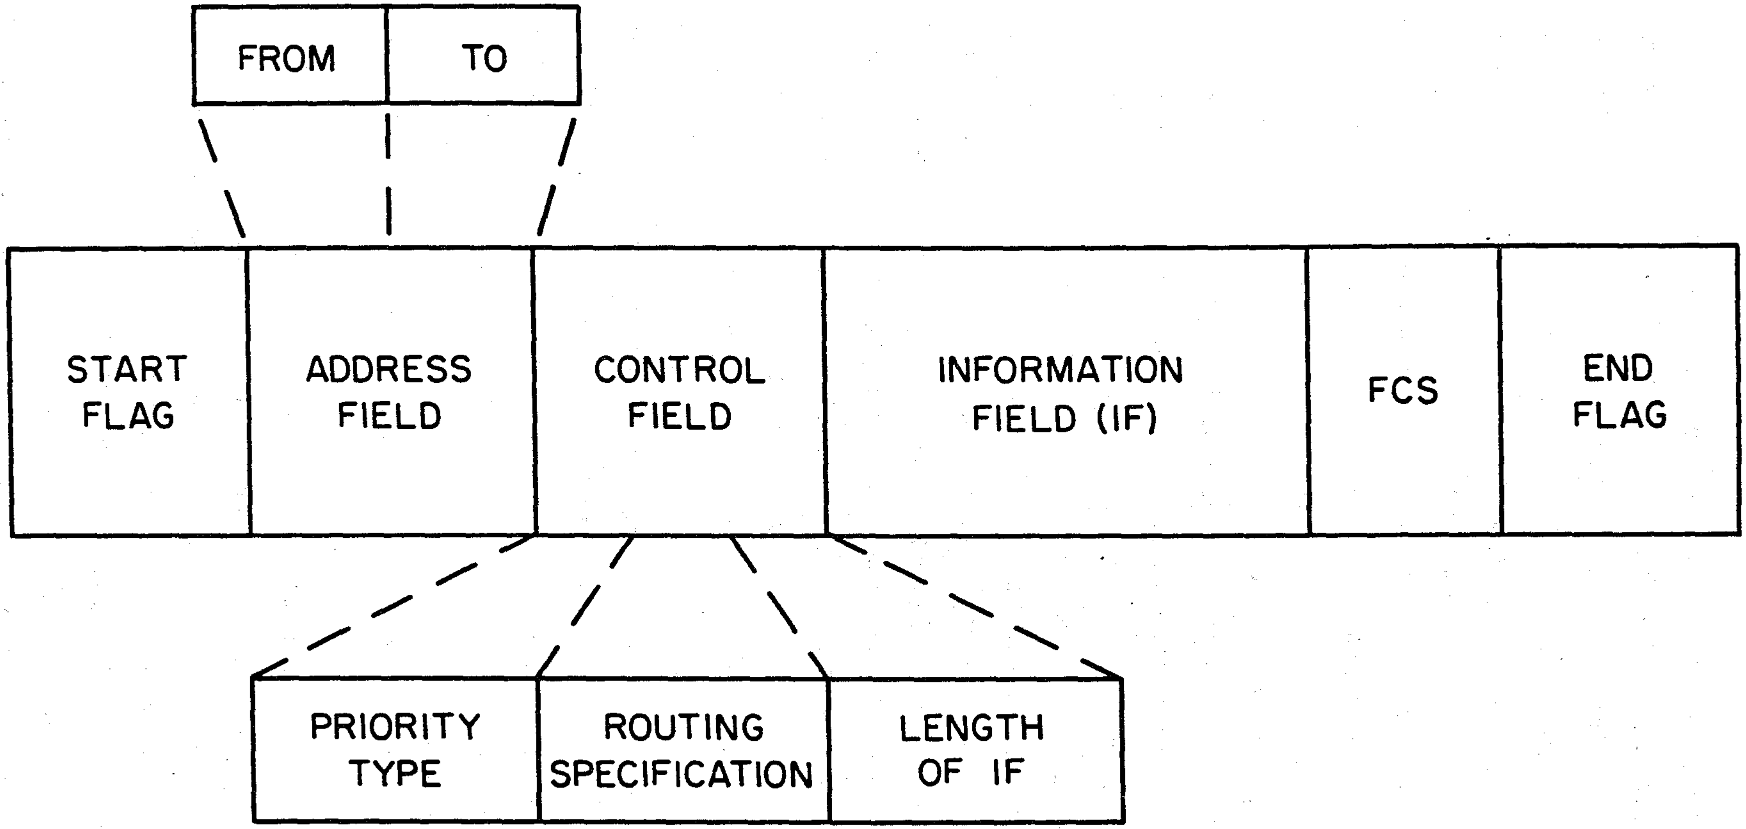

Figure 20. Potential modification of the ADCCP frame format.

the link to notify the sending node of the error detection event. The protocol then asks for a repetition of the previously transmitted frame.

Connectivity and disjoint paths

Connectivity cross section was the topic of section 2.3. A statistical treatment of all cross section numbers in a network led to a histogram that in part has properties desired for a definition of network survivability. Future work should address the development of efficient algorithms to generate the histograms or their moments rapidly.

Consider again a network topology of n nodes and ℓ links. For every pair of nodes, i and j, let their respective orders, i.e., the number of directly linked neighbors, be o(i) and o(j). As before, denote by d(i,j) the distance between the two nodes. Then the connectivity cross section x between i and j is upperbounded as $x \leq M$, where

$$M = \min \left\{ o(i),\ o(j),\ \left[(n-2)/(d(i,j)-1) \right] \right\}.$$

Note: The square brackets in this expression stand for the integer part of their contents.

To deduce the actual value of x, efficient algorithms appear to be missing. One can conceive a brute force approach, where systematically and in an increasing order all combinations of 1,2,3,..., other nodes are deleted, followed by a trial that discerns whether i is, or is not, still connected to j. Clearly, if k is the smallest number that causes i and j to become disconnected, then x must equal k. Moreover, when all cases less than M are exhausted and connectivity has prevailed throughout, there is no need to look at k=M. One concludes that the number of connectivity trials, T, is bounded by

$$T \leq \binom{n-2}{1} + \binom{n-2}{2} + \ldots + \binom{n-2}{M-1},$$

for every pair of nodes.

To assemble the entire network histogram, to perform the statistical moment analysis, and to derive the connectivity or the topological survivability index, one may have to process as many as Tn(n-1)/2 test cases. The number of basic machine steps per test case can be estimated as not less than a hundred for practical size networks. Thus the brute force approach presents a big job. Ways and algorithms should be found to reduce the complexity of this task.

Disjoint path searches in the past have been separated into methods that seek out node-disjoint or link-disjoint paths. Algorithms by Even and Kleitman pertain to this aspect of node and arc connectivity, as the case may be. The methods are complex, but still not directly applicable to our needs. Perhaps more promising are the random search or Monte Carlo algorithms discussed by many.

Historically, several algorithms have evolved to distinguish specific node or link arrangements (also called points and vertices or arcs and lines, respectively) according to a particular given rule. These could be distances from a given node (e.g., the root node in case of the rooted-tree network construction), links subjected to excessive carried traffic, or any facility status. Such algorithms are called labeling algorithms by some. One of the oldest is the Dijkstra algorithm for finding the shortest path between two specified nodes in a network.

The sum of traffic capacities over disjoint paths leads to the concept of maximum flow between any pair of nodes. The often quoted "Max-Flow Min-Cut Theorem" asserts that the maximum flow can be exactly equal to, but not more than, the minimum cut that separates the two nodes. The best known and allegedly the simplest algorithm for finding the maximum flow, as well as the minimum cut, is due to Tanenbaum, Malhotra, and co-workers.

Routing algorithms

Every network that performs addressable point-to-point message delivery, does so by some explicit or implicit means of routing. Many routing schemes exist. Yet it is perhaps not fair to say that there are as many routing schemes as there are networks. While a number of similar networks may share the same algorithm, there may also be other networks that under light or heavy congestion conditions have options to switch between several route selection algorithms. Distinctions are made between fixed (static or nonadaptive) and adaptive routing. In the older, or smaller, or hierarchical networks, the intelligence or routing control is found to be centralized in unique network control centers and systems. In other networks, the controls may be distributed to various degrees.

It is not the purpose of this section to review at length the existing routing algorithms. Nor is it possible at this time to select routing schemes for NSEP/FOCS networks. That engineering job remains to be done in the future. However, to illustrate what some of the options might look like, the next two subsections present shortest path methods that are based on either the availability of the connectivity matrix C or the distance matrix D.

Shortest path methods for matrix C

Given the connectivity matrix C, each link has a length or weight of unity. To find the shortest path between nodes under such a premise, one can do no better than find the path with the least number of links. Although they are quite capable of handling links of any (e.g., nonbinary) length, the previously mentioned labeling algorithms can be utilized here as well. Of course, the methods can now be simplified by merely applying labels of 1, 2, 3, ... , that correspond to the counts of links in a path.

Figure 21 presents the so-called rooted tree algorithm that goes beyond just finding the shortest path between two nodes. Starting from a given root node, this method develops a tree network with desired shortest paths to all other nodes. For n nodes the tree network has n-1 links. Since it has no loops, it allows only one path from one node to another, however, the paths from any node to the root are guaranteed to be of minimum length.

The algorithm outlined in Figure 21 performs an iterative search over the distance index i. For each i it does two things:

* First it finds the "ring" R(i) of all nodes that have the minimum distance of i to the root.

* Second, it prepares a table that lists R(i) on one side and a single parent from the previously prepared inner ring R(i-1) for every node in R(i), on the other side.

The second listing of node pairs is crucial because it shows the linkage in the rooted tree network. Chaining together of consecutive [R(i-1),R(i)] links produces the complete rooted tree. In using such routing tables, either forward or reverse list searches can turn out to be more expedient. But, that topic is more appropriate under data base management issues, and thus is beyond the scope of this project.

The rooted tree algorithm has been applied to the twenty node network first introduced in Figure 5 and with the connectivity matrix C shown in

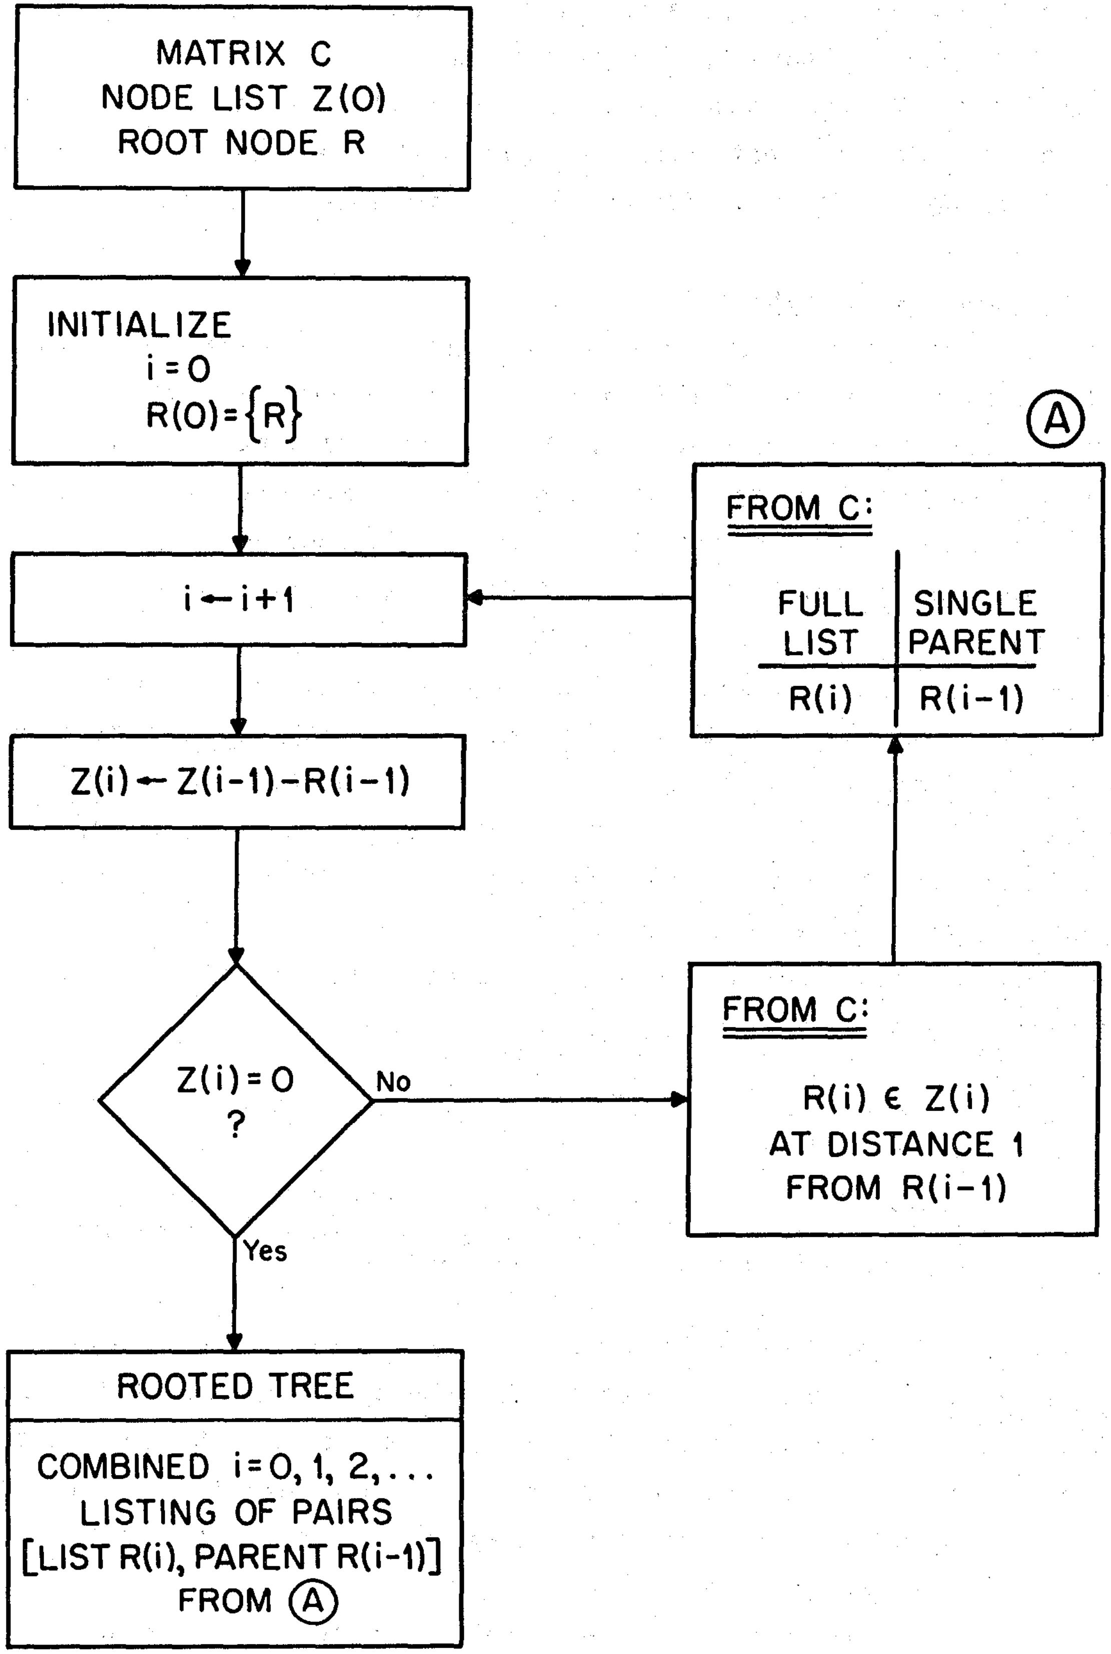

Figure 21. Simple outline of the rooted tree algorithm.

Figure 16. With the root node assumed to be node #1, the resultant rooted tree is given in Figure 22. Note, that this application to a single root has generated only n-1 or 19 shortest paths out of a grand total of n(n-1)/2 or 190 possible paths. Thus, while properly producing a three-hop path from #3 to #1, it fails utterly to reveal the direct link from #3 to #11. A total of n-1 or 19 roots must generally be processed by this algorithm in order to construct a single copy of the shortest routes for all pairs. If alternate or other backup routes are required, the rooted tree method can be expanded accordingly.

Shortest paths from matrix D

The distance matrix D can be used to find routes between any two nodes in the network. In particular, the route in question could be the shortest of all routes. Denote the source node as s and the terminating node as t. Then a readout of the (s,t)-th element of matrix D gives the value of the shortest distance between the two points, namely d(s,t). But, more appears possible with matrix D.

To illustrate the procedures to come, assume the twenty-node network of Figure 5. Its D matrix is as shown in Figure 15. Let s=7 and t=16. Then from D, d(s,t)=d(7,16)=4. Thus the shortest possible path takes 4 hops. Let us find that 4-hop path by manipulating the stored matrix D.

Define a twenty-dimensional binary vector $V=\{v_i\}$ as follows. Let

$$\begin{aligned} v_i &= 1 && \text{for } i = s \text{ and } t, \\ &= 0 && \text{otherwise.} \end{aligned}$$

In this vector V, 1's are inserted in locations s=7 and t=16. All other locations contain zeros:

$$V = (00000\ 01000\ 00000\ 10000).$$

Next consider the product V*D. It is also a vector of 20 terms, where each term, say the k-th (k=1,2, ... ,20), is a sum of two distances. Those are the distances from k to s, and from k to t. One writes therefore,

$$V*D = \{d(s,k)+d(k,t)\}.$$

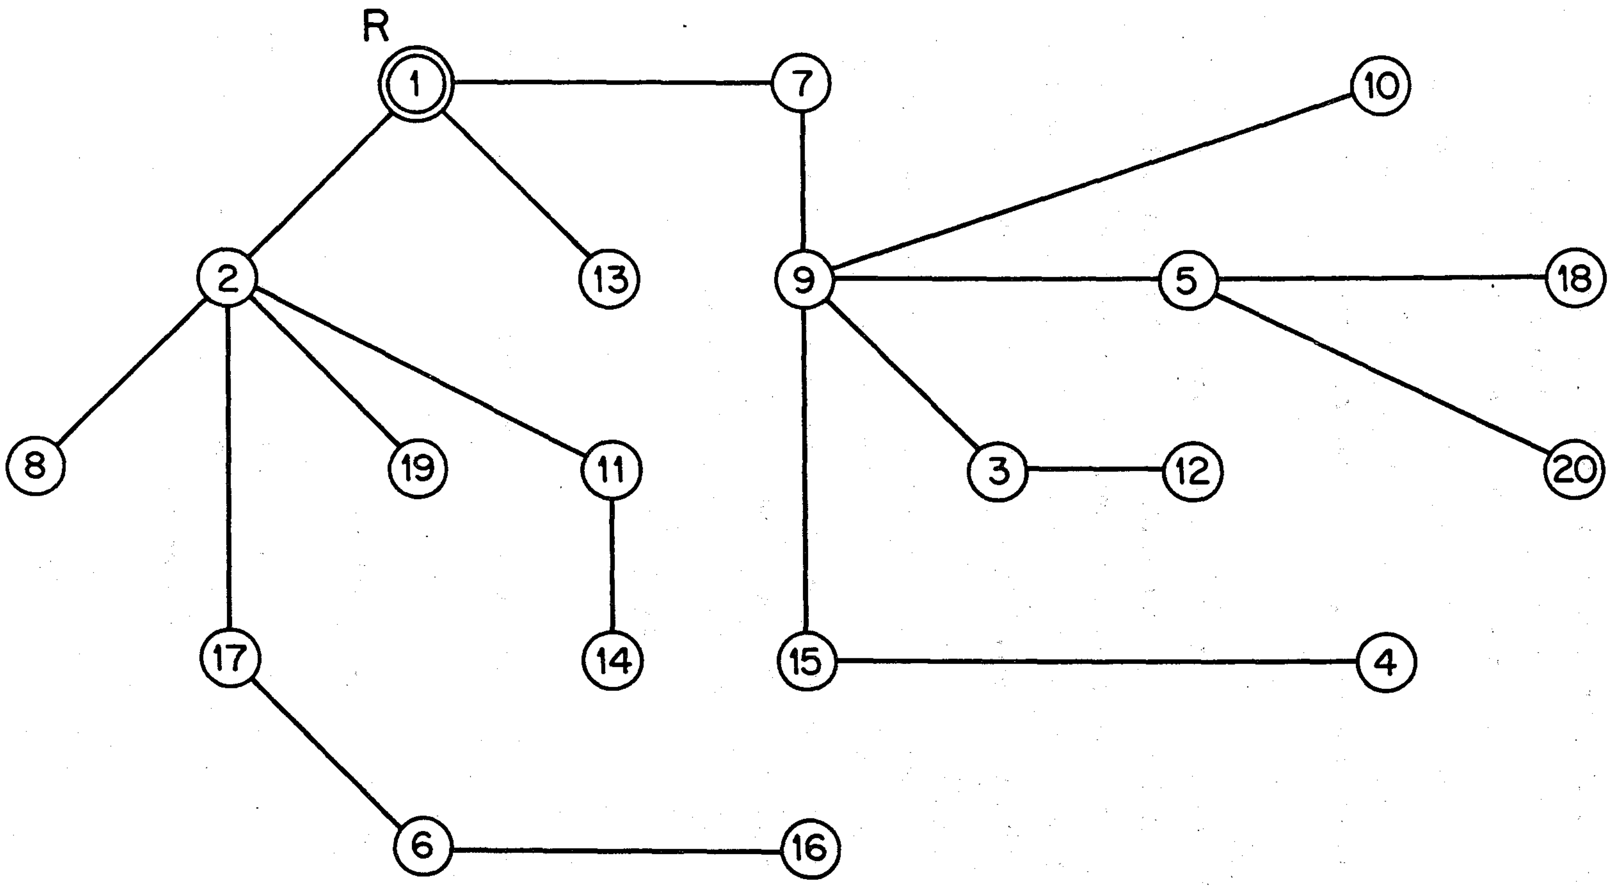

Figure 22. A tree rooted at R = 1.

In the assumed network example,

$$V*D = (55577\ 54657\ 47445\ 45958).$$

It is clear that only those nodes, k, for which

$$d(s,k) + d(k,t) = d(s,t)$$

is true, can be on the shortest path between s and t. There have to be at least d(s,t)+1 such candidate nodes among the elements of V*D for the construction to be valid. Inspection of the V*D example verifies that it is indeed the case. There are five elements that equal 4. They are in the location set

$$X = (7,11,13,14,16)$$

of the above V*D vector.

All the members x of X are candidate nodes for the desired route. By design, they include the two end nodes s=7 and t=16. A simple look at either d(s,x) or d(x,t) in matrix D reveals how the X set is to be permuted (i.e., resequenced) to trace out the actual route through the network. Thus in our example:

x	d(s,x)	d(x,t)
7	0	4
11	2	2
13	1	3
14	3	1
16	4	0

List X therefore contains all the candidates x for the tandem nodes, but in an apparently wrong order. However, that is easily remedied by looking up the correct sequence either from the d(s,x) or the d(x,t) column, for the forward or reverse direction, respectively. The route from 7 to 16 is therefore a mere permutation of the set X, namely (7,13,11,14,16), that corresponds to d(s,x) column being nothing more than (0,1,2,3,4) and d(x,t) being (4,3,2,1,0). An inspection of the path from node #7 to #16 in Figure 5 confirms the procedure.

The above example is fortunate in the sense that it has a unique shortest path from s to t. In general, there can be several alternate paths between a pair of nodes. The procedure then has to be modified to select a single route from all the shortest path possibilities. The next example illustrates the problem and its solution.

For the second example consider the same network of Figure 5 and its D matrix from Figure 15. Assume s=8 and t=20. Then the minimum distance d(8,20)=5, the binary vector V is

V = (00000 00100 00000 00001),

and the tandem distance vector V*D is

V*D = (65556 76566 55765 85765).

Instead of the required d(8,20)+1=6 nodes, one finds nine nodes that satisfy the d(8,x)+d(x,20)=5 constraint. More than one shortest path is possible in the tandem node locating set

X = (2,3,4,8,11,12,15,17,20).

With the aid of D one can sort out and tabulate the relative shortest distances for all tandem location candidates x in X:

x	d(s,x)	d(x,t)
2	1	4
3	3	2
4	4	1
8	0	5
11	2	3
12	4	1
15	3	2
17	1	4
20	5	0

When sorted according to increasing d(s,x), this yields:

d(s,x)	x
0	8
1	2,17
2	11
3	3,15
4	4,12
5	20

The situation now is clearer. Being at distances d(s,x)=0 and 5, the end nodes 8 and 20 represent no ambiguity. Moreover, there is no ambiguity at d(s,x)=2, as the path has only one alternative there, namely node x=11. Thus nodes #8, #11, and #20 offer no other alternatives at their respective d(s,x). Distance step d(s,x)=1 allows the choice between two routes to go from node #8 to #11. It is evident that, in the absence of other routing conditions, both choices must be equally good. One can then arbitrarily pick either #2 or #17. Let us select #2.

Distance steps d(s,x)=3 and 4 present the first instance where a little care is advised. One simply cannot pick #3 and #4 and expect that this will be the wanted shortest route. In fact, as Figure 5 shows, #3 is not directly linked to #4 and such a five-hop path is impossible.

One may, however, select any one of the four nodes among (3,4,12,15). After that selection it remains to ascertain the shortest path existence to the already fixed points (i.e., either #11, or #20, or both). Let us arbitrarily choose #3 for the tandem node at step d(s,x)=3. Since #3 is known to be directly linked to #11, nothing is needed here. But one must yet find whether #4 or #12 is the linkage needed to go from #3 to #20. Matrix D establishes d(3,4)=2 and d(3,12)=1. Therefore node #12 provides the final link for the above process. The five-hop end-to-end route is (8,2,11,3,12,20) and it agrees with the d(s,x) sequence of (0,1,2,3,4,5).

The explanation by examples, as given above, has not addressed several key questions. One question, that seems to be of rather general interest in the unity- or integer-distance networks, pertains to the random selection of the next node. Assume that at every distance d(s,x) between x=s and x=t there is more than one, perfectly valid, candidate tandem node for the shortest length route to be determined via matrix D. In a schematic fashion, such a situation is illustrated in Figure 23. Given that the first (or the next) node to be picked can be anywhere on the d(s,x) scale, are there any advantages for some selection strategies over others? The advantage could amount to quicker or less complicated procedures for routing with D. The strategies themselves could differ from each other in how they generate the sequence of next picks. For instance, one could use a half-distance rule, where the first random pick is at distance x_1, such that $d(s,x_1)=d(s,t)/2$; the second pick is at x_2, such that $d(s,x_2)=d(s,x_1)/2$, and so forth. One could also select a node where their multiplicity is the largest, which would be at the x where d(s,x) is maximum in Figure 23. Perhaps there exists, what one could call, an optimum binary splitting scheme for routing with D in a general class of network. Further study, both with a general scope and in detail, remains to be done.

Summary. To identify shortest paths with the aid of the distance matrix D, one can execute iteratively and until finished the following six tasks:

(1) Given source and termination nodes, s and t, construct the end-identifier vector V.

(2) Multiply vector V with matrix D to generate the tandem distance vector V*D.

(3) Extract from V*D the set X of all candidate nodes for the shortest path.

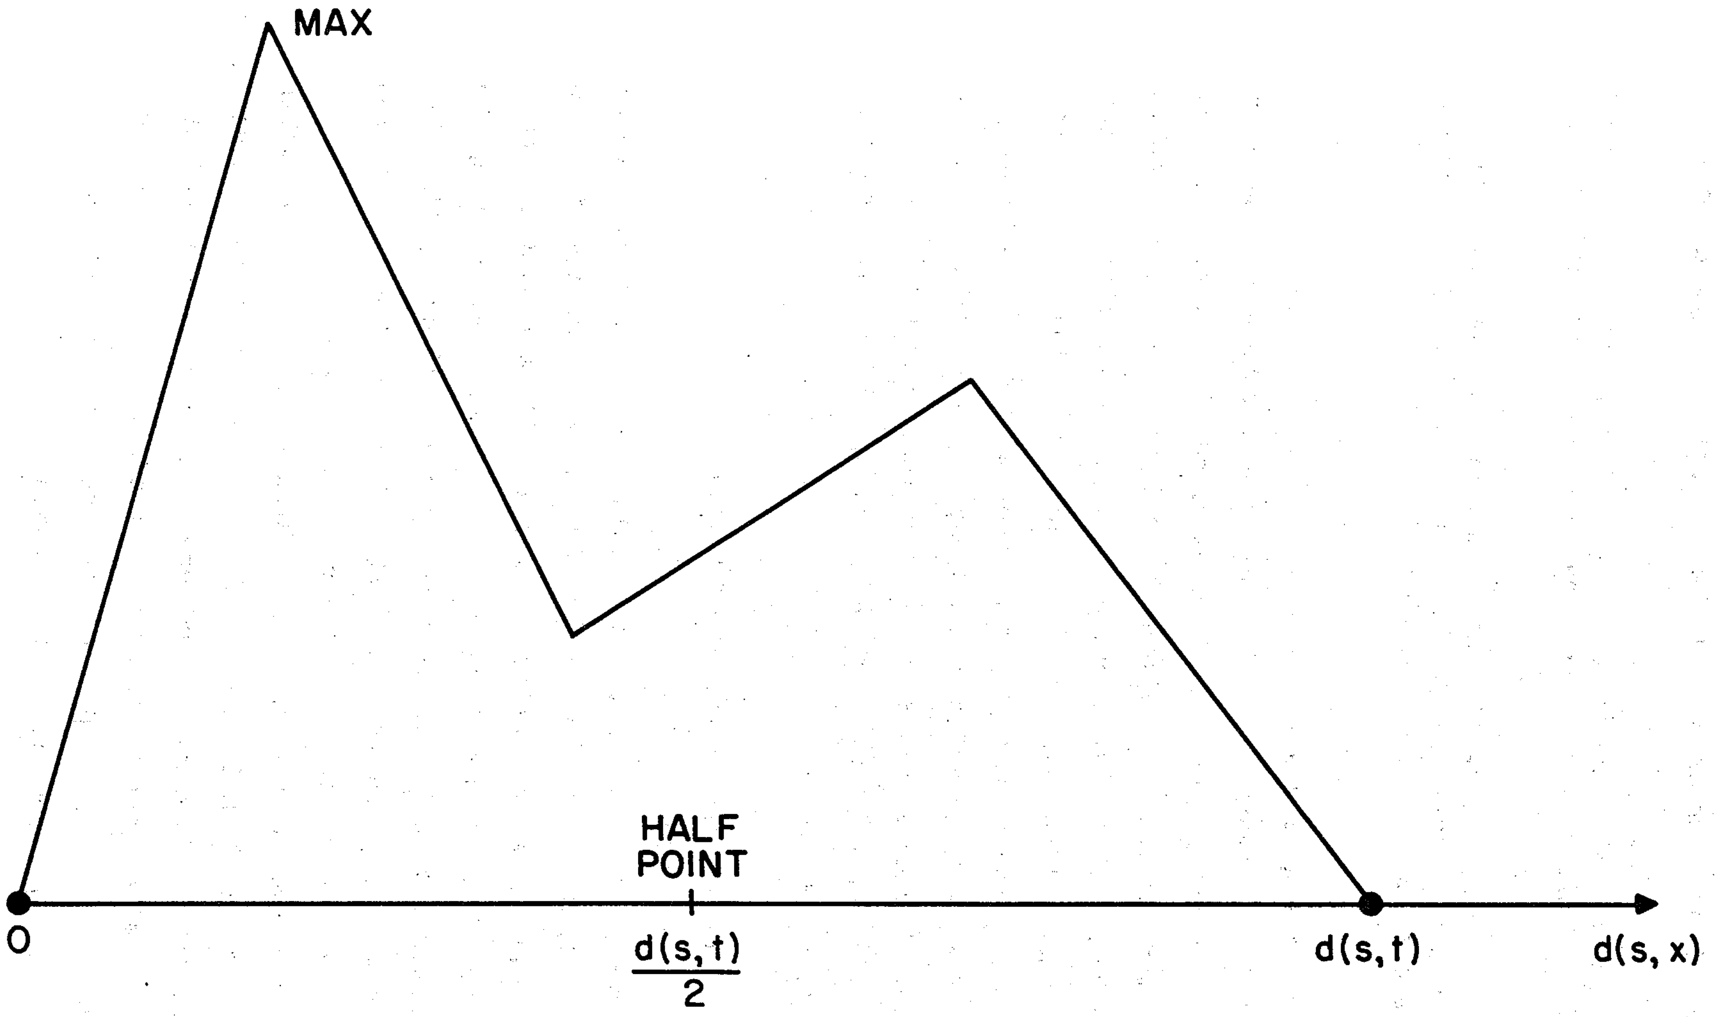

Figure 23. Schema for multiplicity of candidate tandem nodes.

(4) Rearrange list X in increasing order of d(s,x), where x denotes an element of X.

(5) Identify and set aside unique nodes in X. If there is a unique node at every d(s,x), define the shortest path as the rearranged list X. EXIT.

(6) Perform the random node selection algorithm on that non-empty set of nodes whose multiplicity is larger than one. Redefine one or more x nodes as pseudo s and t nodes. RETURN to task (1).

3.5 Propagation and Transmission Times

For the reconstitution process that includes the exchange of the six-dimensional array vector (see section 3.2), the required time budget could be of concern. In this brief section it is shown that for the expected geographic size fiber optic networks, based on the realistic local versus global choices of Table 4, the total restoration time is apt to be a fraction of a second. That total consists of sequenced link sessions, the detailed time budgets of which are likely to be dominated by propagation delays and not by the number of bits used in the arrays. The node delays are assumed to be negligible.

In any medium with a constant index of refraction, the propagation time for a light signal is by definition

(Distance) * (Refractive Index) / (Speed of Light).

The refractive index for typical glass in the visible region is assumed to be 1.50. The speed of light is taken as 300,000 km/s, corresponding to propagation in the vacuum. Then the optical delay is .001 s = 1.00 ms for a fiber that is 200 km = 125 miles long. The 125-mi number is quite representative of distances between many neighboring cities in the USA, especially on both coasts. In the Western States distances are greater and 500-mi separations are not uncommon. Then the propagation time must be around 4.00 ms. Finally, if one were to span the entire CONUS with a continuous fiber link, one might conjure a 3,000-mile circuit from New York to San Francisco. The resultant 24.00-ms delay represents an extreme upper bound. For realistic fiber optic deployments along major interstate highways, it seems far more prudent to assume that the optical propagation delays per link will be somewhere in the middle of the 1.00- to 10.00-ms range.

Next consider the transmission or modulation time needed to send the multidimensional array data for network reconstitution. Assume that the n-node, ℓ-link, d-diameter network has the sizing of the status arrays as summarized in Table 3. Let all integers and real numbers be roughly approximated with the nearest quantized 8-bit numbers. In terms of n and ℓ, let d be upperbounded as per PROPERTY II of section 3.1. Then, for every given network size in terms of n and ℓ, as well as any Global (G) or Local (L) coverage combination for the six vector V components, one can estimate the number of bits involved in a single search-and-broadcast transmission from one node to another.

The results of such a calculation for small networks are shown in Table 5. The rows represent the number of nodes, such as n=8, 12, 16, 20, and 24. The columns are the links ℓ=16, 20, 24, 28, 32, and 36. For each (n,ℓ)-pair, four numbers are tabulated. The uppermost, denoted as (L)*, stands for the simple Local option without routing tables (RT). It requires the least exchange of status bits. The second number, denoted as (L)**, also represents the Local option, but includes the RT. The third number, (G)*, gives Global attention to everything--including traffic--but to the exclusion of RT.

The fourth number, identified by (G)**, includes the RT in the Global option. It calls for the most data to be transmitted, and probably would never be implemented. The reason for that is the fact that RT's develop slowly. They are only finished when the t(0) timer has expired or when the restoral is done. Before that, acquisition of partial routes are of doubtful value. And, of course, when finally all nodes possess the same complete copy of the connectivity matrix, there remains the option to generate all needed RT's locally with a common algorithm.

All four quantities are based on the array-size formulas given previously in Table 3, plus one 8-bit word for the identification of the broadcasting node. To illustrate how the indivisual numbers are derived, consider the example n=8 and ℓ=16.

For (L)*:

$$48\ (\ell/n) + n + 15 = 96 + 8 + 15 = 119.$$

For (L)**:

$$48\ (\ell/n) + 33n - 17 = 96 + 264 - 17 = 343.$$

Table 5. Sizing of the Array Vector V in Bits

		ℓ=16	ℓ=20	ℓ=24	ℓ=28	ℓ=32	ℓ=36
n= 8	(L)*	119	143	167	191	***	***
	(L)**	343	367	391	415		
	(G)*	356	420	484	548		
	(G)**	3,940	3,108	2,276	1,444		
n=12	(L)*	91	107	123	139	155	171
	(L)**	443	459	475	491	507	523
	(G)*	426	490	554	618	682	746
	(G)**	19,434	17,386	15,338	13,280	11,242	11,306
n=16	(L)*	79	91	103	115	127	139
	(L)**	559	571	583	595	607	619
	(G)*	512	576	640	704	768	832
	(G)**	54,272	50,496	46,720	42,944	39,168	35,872
n=20	(L)*	***	83	93	103	112	122
	(L)**		691	701	711	720	730
	(G)*		678	742	806	870	934
	(G)**		110,118	104,102	98,086	92,070	86,054
n=24	(L)*	***	***	87	95	103	111
	(L)**			823	831	839	847
	(G)*			860	924	988	1,052
	(G)**			195,164	186,396	177,628	168,860

* Without routing tables (RT)
** With routing tables
***Topoligically impossible

For (G)*:

$$16\ell + n(n+15)/2 + 8 = 256 + 92 + 8 = 356.$$

For (G)**:

$$16\ell + n(n+15)/2 + 8 + 16\ dn\ (n-1) = 256 + 92 + 8 + 3{,}584 = 3{,}940,$$

because, as shown earlier in Section 3.1, the value $d=4$ applies here.

The delays caused by the various bit counts depend further on the amount of overhead (OH) involved, as well as on the data rates of the respective FOCS channels. To keep things simple, assume two OH values: 0% (idealistic) and 100% (more realistic). Furthermore, assume that one can have only two channel throughput rates: either a rather modest 10 Mb/s or a more state-of-the-art 25 Mb/s. Then the transmission times may be summarized as shown in Table 6.

Table 6 brings together all the estimates of Table 5 and therefore applies to node counts in the range $8 \leq n \leq 24$ and link counts in the range $16 \leq \ell \leq 36$. They are deemed to be practical ranges for implementation. From the coverage options quoted, both Local options appear easily relizable. Likewise, the Global option without RT, also identified as (G)* in Table 5 and believed to be within the ballpark of the earlier "realistic" choice from Table 4, appears possible. Under the most conservative circumstances shown, its transmission time is only .21 ms. That is almost negligible when compared with the expected one-to-ten milliseconds for optical propagation.

Early in the reconstitution search, i.e., when the timer is near t(0), the individual V fields are the shortest, as is clear from Table 2. They grow and eventually achieve their maximum size towards the end, i.e., when the timer approaches 0. For networks with diameter $d \leq 50$, it is sufficient to set the initial timer to t(0)=100. The total time consumed in a search-and-broadcast session is then less than or at most approaching one second. The one-second time budget was the conjectured target number associated with Figure 11.

4. REPRESENTATIVE NETWORK CASES

4.1 Relatively Small Networks

For the first small network example, consider the topology shown in Figure 24. The fiber optic transmission facilities here are aligned with the major interstate highways in the Western States, bounded by an approximate rectangle that has Albuquerque, New Mexico; Buffalo, Wyoming; Sioux Falls, South Dakota; and Oklahoma City, Oklahoma at its corners. The basic network

Table 6. Transmission Time in ms of the Array Vector V for Practical Size Networks*

	10 Mb/s Link		25 Mb/s Link	
	0% OH	100% OH	0% OH	100% OH
Local option Without RT	.02	.04	.01	.02
Local Option With RT	.08	.17	.03	.07
Global Option Without RT	.11	.21	.04	.08
Global Option With RT	19.52	38.03	7.81	15.61

*This applies to all $8 \leq n \leq 24$ and $16 \leq \ell \leq 36$.

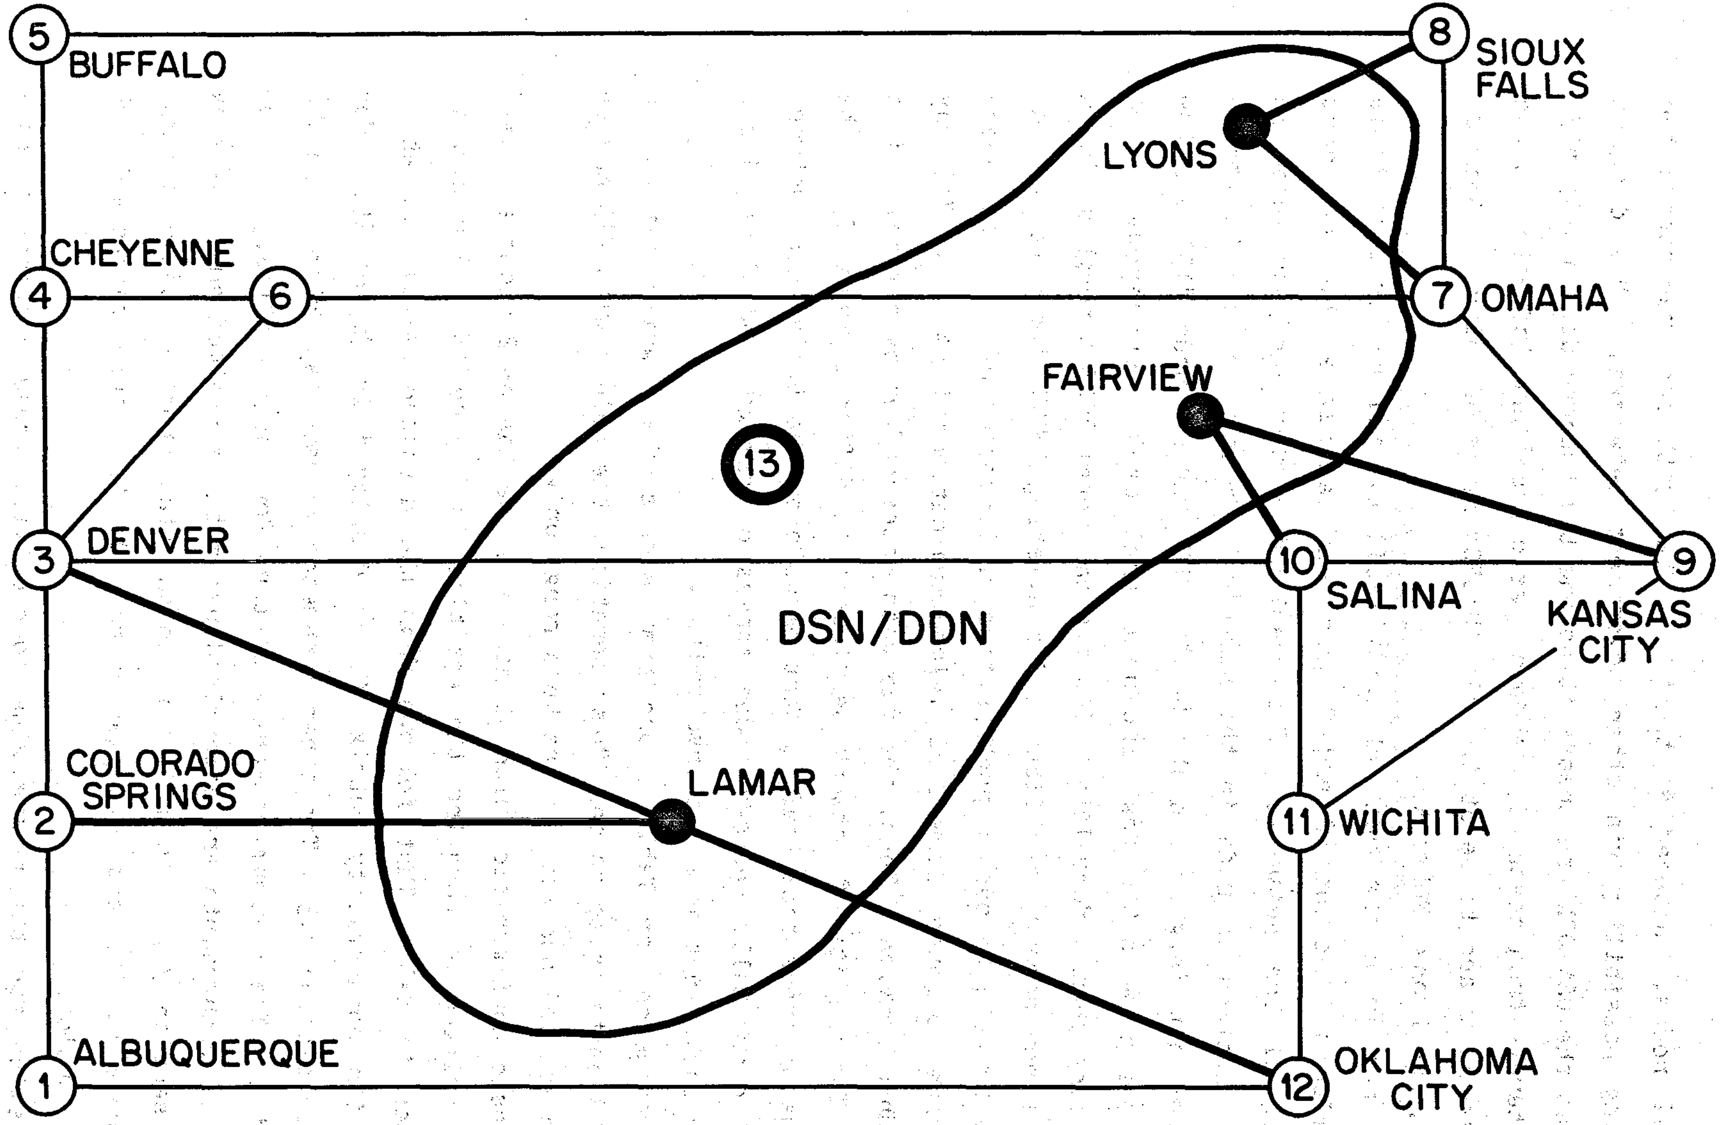

Figure 24. A twelve node network with DSN/DDN backup.

consists of n=12 nodes, numbered #1 to #12. The 13-th node represents the alternate routing possibility via other Government networks. The diagram refers to this as the DSN/DDN and treats it as one big "supernode," namely #13. Three access gateways from seven of the dozen basic nodes are indicated to the DSN/DDN. The gateways are to be at the old AUTOVON sites of Fairview, Kansas; Lamar, Colorado; and Lyons, Nebraska.

In accordance with the network survivability and restoration tools introduced earlier in Sections 2 and 3, one can calculate the pertinent characteristics of the network. The resultant numbers are presented in Table 7. A distinction is made in the numerical columns between the basic twelve-node fiber net (e.g., without the DSN/DDN backup, node #13) and the augmented thirteen node setup that includes that backup supernode. The distinction is more than a formality. The typical total FOCS throughput is in tens of megabits per second, per fiber. The existing DSN/DDN trunks, on the other hand, may be able to manage perhaps tens of thousands of bits per second each. The planned Defense Commercial Telecommunications Network (DCTN) may offer a 1.544 Mb/s service. Thus, the two data rates differ by several orders of magnitude. The utility of network services will differ accordingly. The normal peacetime or pre-attack traffic that flows over the twelve node network may have to be drastically reduced under stress or damage conditions (i.e., when the slower backup networks are involved).

In the conclusion for Table 7 one must stress that, while the column without node #13 pertains to normal Mb/s operation, the column with #13 may be restricted to essential services in the kb/s range only.

For the basic configuration without node #13 one has n=12 nodes and ℓ=16 links, as is verified by counting the entities in Figure 24. The actual network diameter is d=5. That follows from the distance matrix, D, which together with its connectivity matrix, C, is displayed in Figure 25. For definitions, see section 3.3. The other parameter values also follow from the previous material. The maximum network diameters for the specified, known or unknown, survival constraints are determined using the so-called small network properties of section 3.1. As will be demonstrated shortly, these maximum diameters, namely 9, 11, and 16, play a role in the timer management for network reconstitution.

The statistics for the networkwide connectivity cross sections involve means and standard deviations over all possible 12*11/2 = 66 point-to-point

Table 7. Characteristics of the Western State Major Highway Network

Characteristic	Symbol	Without Node #13	With Node #13
Number of Nodes	n	12	13
Number of Links	ℓ	16	23
Actual Network Diameter	d	5	4
Maximum Constrained Diameters:			
Given n and ℓ	$d_{max}\ (n,\ell)$	9	8
Given n	$d_{max}\ (n,.)$	11	12
Given ℓ	$d_{max}\ (.,\ell)$	16	23
Connectivity Cross Section for the Network:			
Mean	m	2.32	2.77
Standard Deviation	s	.47	.60
Effective Index	$m-s/4$	2.20	2.62

$$C = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix}$$

$$D = \begin{bmatrix} 0 & 1 & 2 & 3 & 4 & 3 & 4 & 5 & 3 & 3 & 2 & 1 \\ 1 & 0 & 1 & 2 & 3 & 2 & 3 & 4 & 3 & 2 & 3 & 2 \\ 2 & 1 & 0 & 1 & 2 & 1 & 2 & 3 & 2 & 1 & 2 & 3 \\ 3 & 2 & 1 & 0 & 1 & 1 & 2 & 2 & 3 & 2 & 3 & 4 \\ 4 & 3 & 2 & 1 & 0 & 2 & 2 & 1 & 3 & 3 & 4 & 5 \\ 3 & 2 & 1 & 1 & 2 & 0 & 1 & 2 & 2 & 2 & 3 & 4 \\ 4 & 3 & 2 & 2 & 2 & 1 & 0 & 1 & 1 & 2 & 2 & 3 \\ 5 & 4 & 3 & 2 & 1 & 2 & 1 & 0 & 2 & 3 & 3 & 4 \\ 3 & 3 & 2 & 3 & 3 & 2 & 1 & 2 & 0 & 1 & 1 & 2 \\ 3 & 2 & 1 & 2 & 3 & 2 & 2 & 3 & 1 & 0 & 1 & 2 \\ 2 & 3 & 2 & 3 & 4 & 3 & 2 & 3 & 1 & 1 & 0 & 1 \\ 1 & 2 & 3 & 4 & 5 & 4 & 3 & 4 & 2 & 2 & 1 & 0 \end{bmatrix}$$

Figure 25. Connectivity and distance matrices for the basic twelve node network of Figure 24.

cross sections. The terms and numbers are as introduced in section 2.3. The mean or average connectivity cross section is m=2.32. The corresponding standard deviation is s=.47. For a numerical assessment of topological survivability, the effective index or metric assumed here becomes m-s/4=2.20.

By itself the value of the index, e.g., 2.20, has certain limitations. For instance, it is difficult to interpret its network traffic handling, economics, or other implications. However, things become more tangible when two or more indices are compared. For the same investment in node and link plant, a larger cross-section index implies a potential for enhanced survivability in terms of network connectivity (i.e., to avoid the network being cut into disjointed parts by facility outages). As noted, the index says little or nothing about the relative traffic capabilities of the compared configurations. A glance at the two columns in Table 7 make that point clear. The capacity of the DSN/DDN and its gateway feeds determines the data rates through node #13.

The augmented network consists of the basic twelve-node FOCS plus the DSN/DDN. In the rightmost column of Table 7, it has n=13, ℓ=23, d=4, and other numbers as listed. The most notable may be the effective cross-section index of m-s/4=2.62. Instead of the earlier basic value of 2.20, this shows a quantitative improvement of 0.42. This is a significant number in that it implies many cases of connectivity survival, where connectivity is impossible in the earlier twelve-node setup. Analyses of threat scenarios are beyond the scope of this study. Therefore, realistic probability numbers for connectivity failures under stress are not possible.

A final comment on the effective index m-s/4: As described and discussed, the index treats all nodes and all links in a equal and uniform way. Up to this point, there is no built-in discrimination for users, sites, or nodes of different priority. There is also no built-in distinction between high and low data rates, device speeds, facility restrictions, and so on. If such features are to be part of topological design, the definition of the connectivity cross section could be modified by weighing individual message paths in a to-be-determined discriminatory manner.

Suppose next that the network in the Western States suffers stress and damage that result in basic node outages at Denver, Cheyenne, and Kansas City, plus the loss of the AUTOVON gateway at Lamar. The remaining operational parts of the network are illustrated in Figure 26.

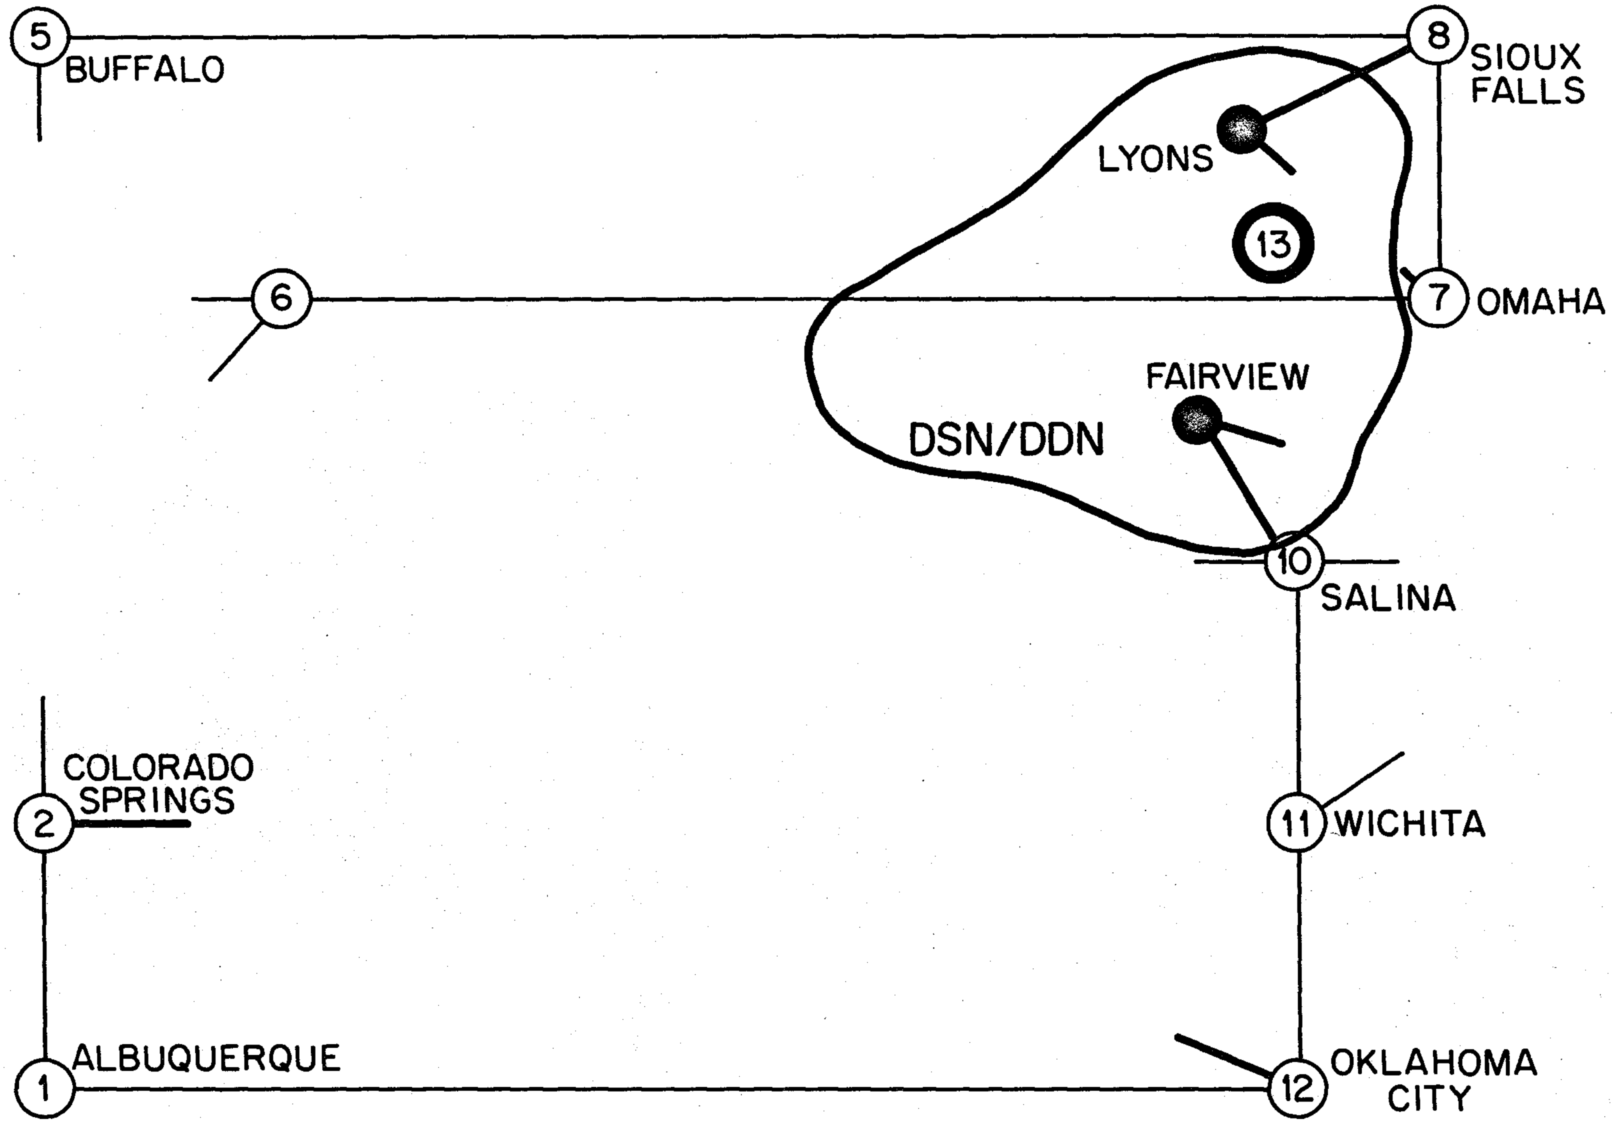

Figure 26. Surviving parts of the twelve node network.

Although nine of the original twelve FOCS nodes are operational, they are now faced with several problems. First, they are unaware of who else is or is not disabled. Second, initially they do not know the new connectivity. For example, somehow they must learn that the former basic network has been separated into two disjoint pieces. Third, the status of the backup DSN/DDN is questionable. Fourth, the routing tools, such as the connectivity and distance matrices must be re-established. And there may be other difficulties to be resolved, perhaps too numerous and varied to be anticipated here. Much of that is part of the restoration function to be covered subsequently.

First, however, note the dilemma of how to represent the characteristics of a disjointed topology. Things like diameters are not finite, unless one deals with the individual segments separately. Zero cross sections, which result from a network being cut into two or more disjointed subnetworks, are of doubtful statistical value. For that reason, it seems justified to skip the descriptions of the disjointed basic remains. Instead, Table 8 characterizes only the connected network that results from the incorporation of the backup node #13.

Including the backup node #13, the number of nodes is now $n=10$. The number of links, $\ell=9$, is the least permitted to ensure connectivity. Mainly because of the low number of working links, the actual diameter is relatively large (compare with Table 7). The three constrained diameters, however, are all lower and actually indistinguishable. However, by assumption, none of the diameters are known before the restoration process is completed. Since this is a tree topology, there is exactly one possible path for every pair of nodes. Thus, $m=1$, $s=0$, and $m-s/4=1$.

The restoration process can be initiated by any one of the surviving nodes. For the purposes of discussion, assume that it is node #2, namely Colorado Springs, that senses problems with the Denver connection and is the first to initiate the called-for restoration session. Node #2 is likely to be entirely in the dark about the status of the damage. Therefore, it can be presumed to know nothing about any of the revised topology parameters in Table 8. However, node #2 can have in its memory the pre-crisis numbers of Table 7. Using said numbers, there are at least two ways to set the initial timer, $t(0)$:

Table 8. Characteristics of the Damaged Western State Network

Characteristic	Symbol	With Node #13
Number of Nodes	n	10
Number of Links	ℓ	9
Actual Network Diameter	d	8
Maximum Constrained Diameters:		
Given n and ℓ	$d_{max}\ (n,\ell)$	9
Given n	$d_{max}\ (n,.)$	9
Given ℓ	$d_{max}\ (.,\ell)$	9
Connectivity Cross Section for the Network:		
Mean	m	1.00
Standard Deviation	s	.00
Effective Index	$m-s/4$	1.00

* With or without #13, the largest of the two dmax(n,ℓ) values is 9. On the basis of the node and link counts of the undamaged network, one could set t(0)=19.

* Considering damages, one may take the optimistic view that perhaps all 13 augmented nodes could have survived, the only outages being on the network links. It would follow then that the constrained diameter dmax(n,.)=12 and t(0)=25.

Figure 27 illustrates that both timer settings are adequate for the damaged, as well as for the undamaged, scenario. They exceed the time required by all twelve FOCS nodes to reconstruct the complete connectivity, plus other matrices, over the entire network. The listing is based on #2 being the initiator. The original undamaged results are shown in solid black. The damaged network is drawn as empty or white rectangles. Note that in the undamaged scenario it is node #12 that takes the longest to collect all data. That maximum delay is 10 hop units. For the damaged scenario, the longest time of 16 hop units is needed for node #2, all under the assumption that node #2 is the initiator. No damaged network search times are shown for nodes #3, #4, and #9, because they are supposed to be defunct by assumption.

For the second small network example consider a backbone type network that extends over the entire CONUS, but has only n=12 nodes. As indicated in Figure 28, the nodes are located in the major US metropolitan areas. The inter-node transmission plant is assumed to have ℓ=18 links. The geographic layout appears quite typical of existing and planned long distance communications facilities in North America. One finds the heaviest concentration of service needs in the North-East, followed by a lesser market on the West coast. A relatively sparse two-link connectivity interconnects the two coasts.

The main topological characteristics of this hypothetical backbone network are listed in Table 9. When counted in actual link hops, the network diameter is a rather modest d=4. The maximum constrained diameters, depending on what is claimed to be known about n or ℓ, range from 8 to 18, all of which appears reasonable. However, there seems to be a weakness associated with the survivability, that is, with the connectivity cross-section aspect, of this

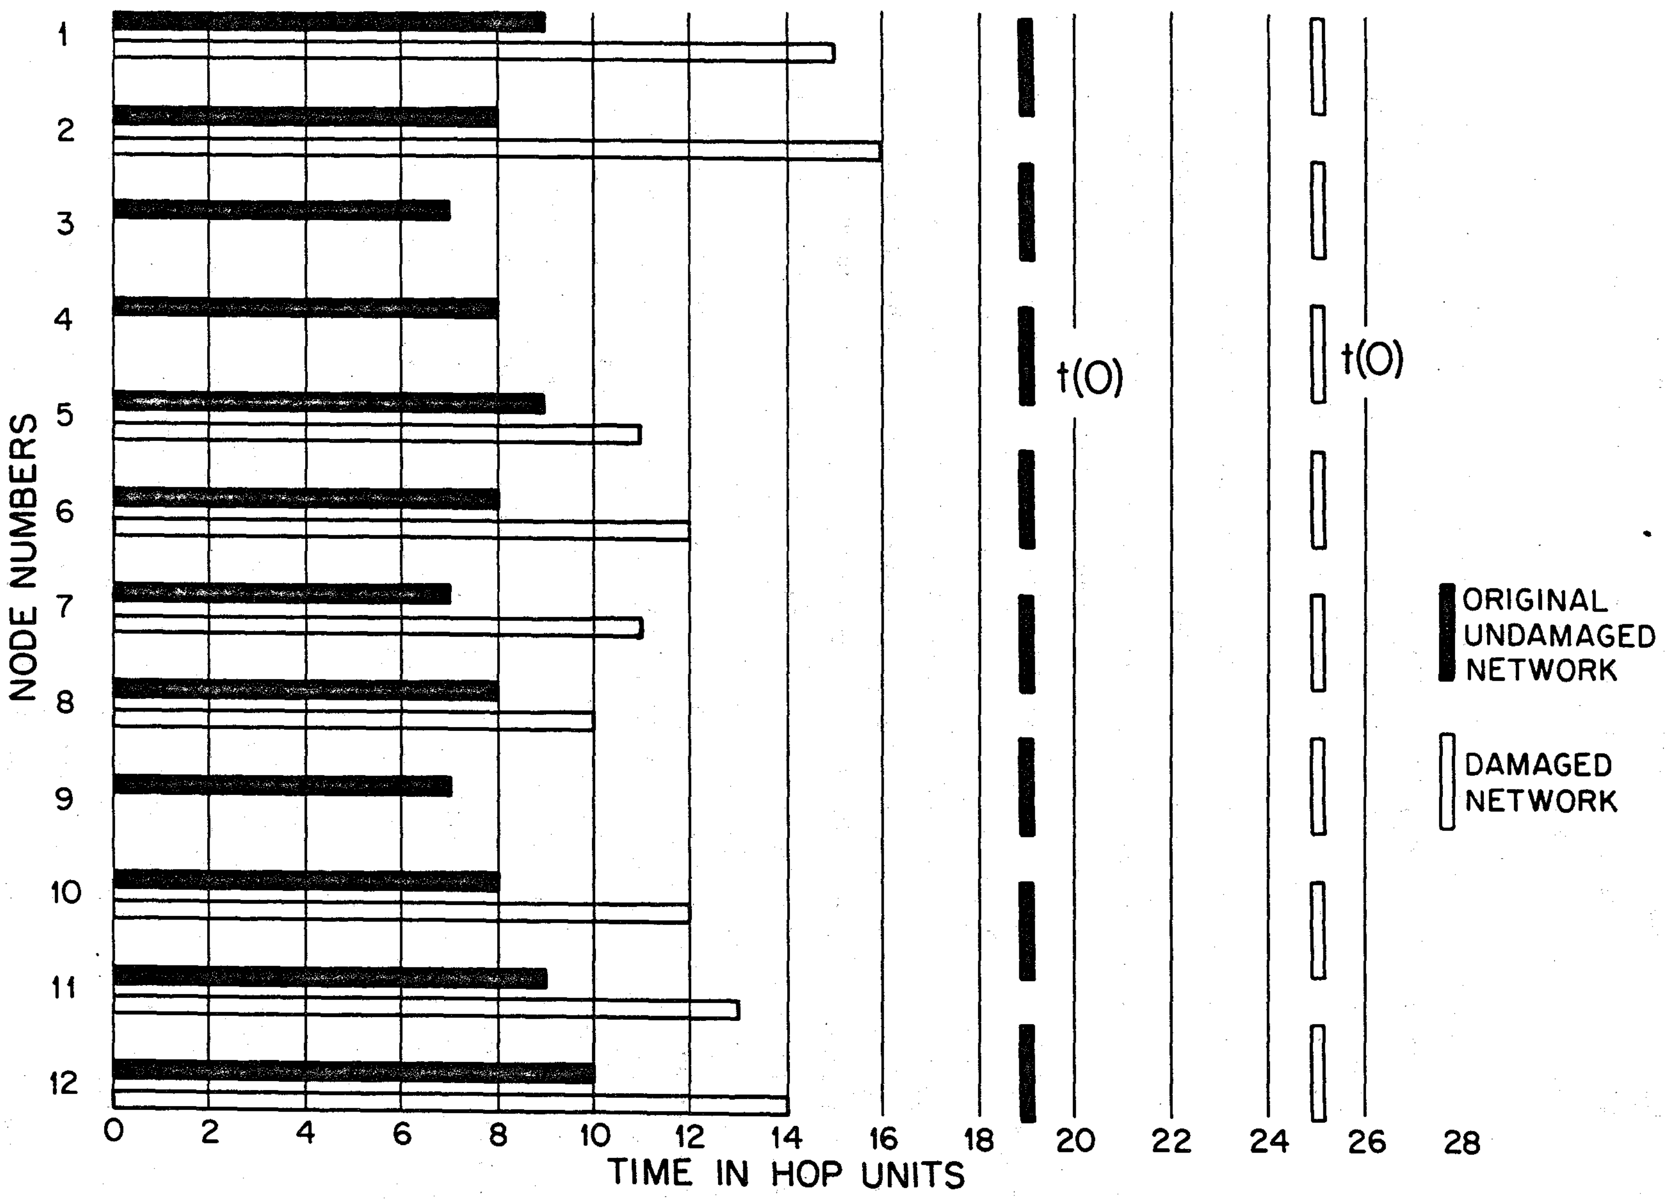

Figure 27. Times needed for complete status updates at the various network nodes.

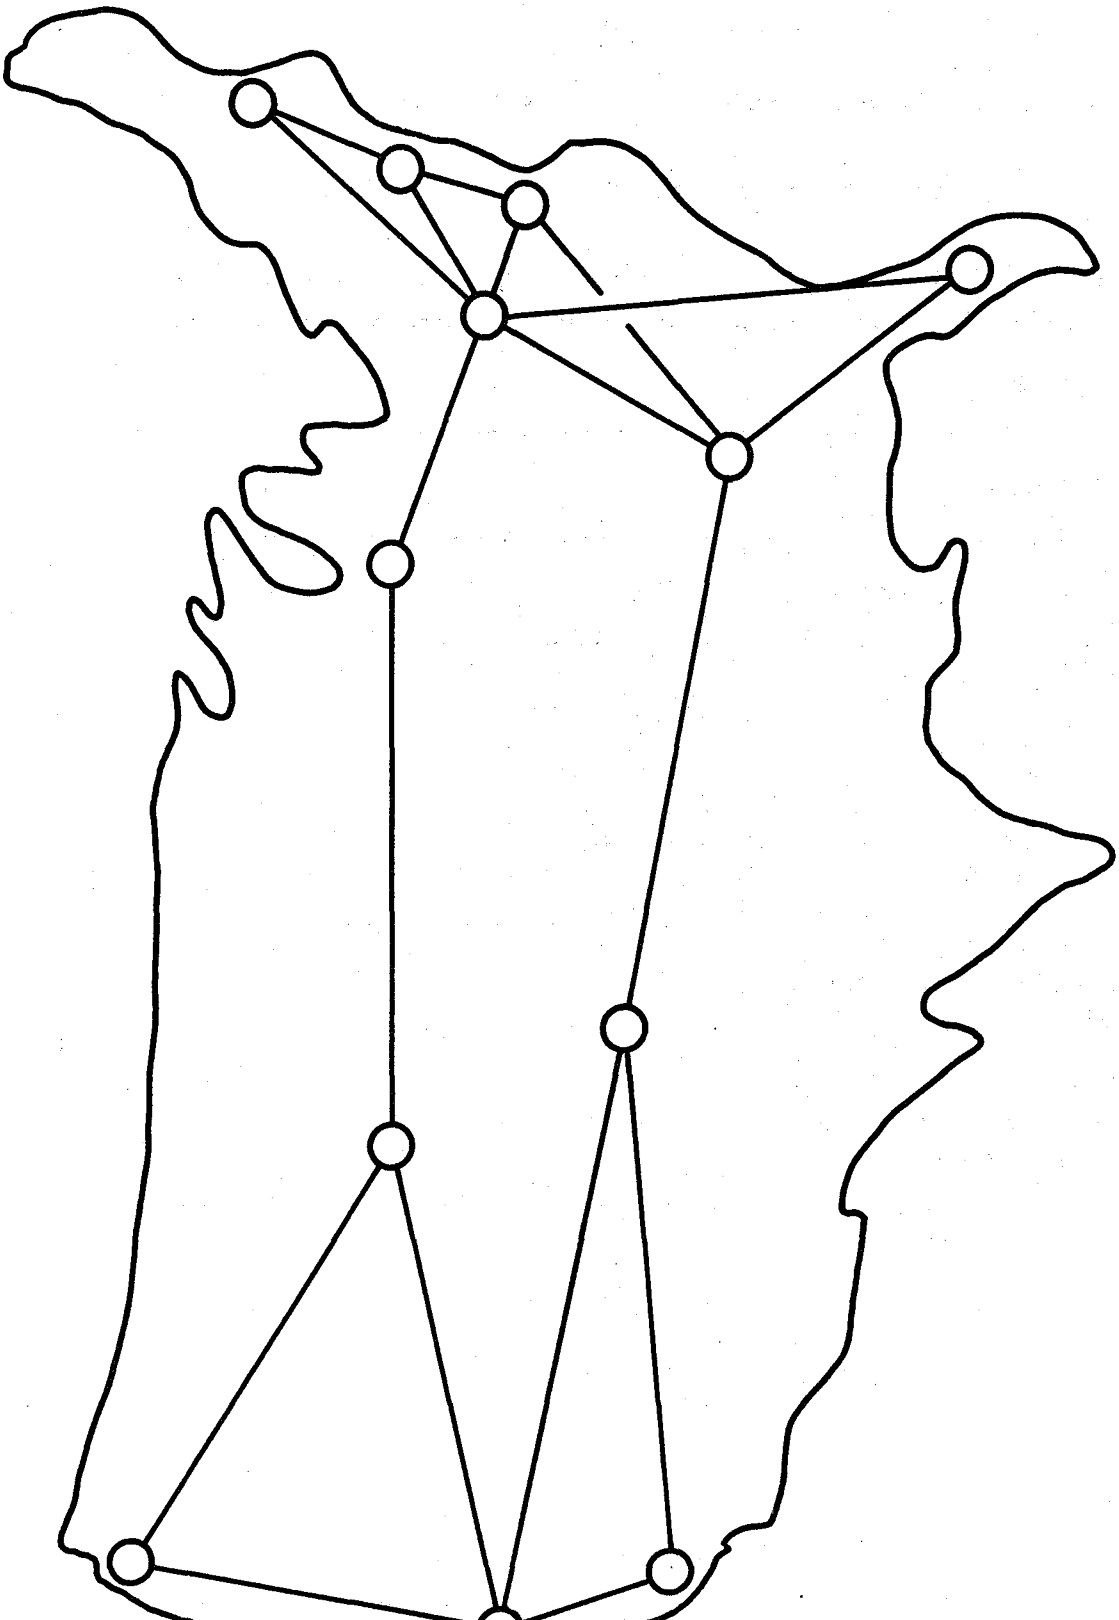

Figure 28. Hypothetical backbone network for CONUS.

Table 9. Characteristics of Two CONUS Backbone Networks

Characteristic	Symbol	For Figure 28	For Figure 29
Number of Nodes	n	12	
Number of Links	ℓ	18	
Actual Network Diameter	d	4	
Maximum Constrained Diameters:			
Given n and ℓ	$d_{max}\ (n,\ell)$	8	
Given n	$d_{max}\ (n,.)$	11	
Given ℓ	$d_{max}\ (.,\ell)$	18	
Connectivity Cross Section of the Network:			
Mean	m	2.09	3.00
Standard Deviation	s	.34	.00
Effective Index	$m-s/4$	2.01	3.00

particular network configuration. The effective cross-section index is only around 2, as can be anticipated for the thin East-to-West connectivity.

Ignoring the economic, traffic, population projection, and other service arguments, the connectivity cross section can be noticeably enhanced by a simple rearrangement of the links between the same identical nodes. Figure 29 represents one possible alteration. For the new topology, the number of nodes, the number of links, and therefore the four diameters, are the same as for the earlier network in Figure 28. Only the connectivity cross sections differ. That all is apparent in the lower right-hand column of Table 9. The mean and the effective cross section index for the enhanced version are both equal to 3. In order to disconnect one part of the backbone from another, instead of 2, now 3 or more combinations of node and/or link outages must be affected.

4.2 Large Networks

Specific large networks are difficult to illustrate for two reasons. First, the graphics become too unwieldy for ordinary human comprehension. And second, the derivation of network characteristics (e.g., the moments of the connectivity cross section) requires increasing detail and volumes of calculation. This section presents one network topology example that is based on the ARPANET (now called DARPANET) that existed approximately 10 years ago. This example is followed by a few observations about other networks so large that one can manage to focus only on small fractions of them. For many practical purposes the latter networks can be considered infinite.

Figure 30 outlines a network of 57 nodes and 71 links. The node numbers are the same as in the alphabetical listing of the early ARPA network sites. Thus for example:

Node #1	. .	Aberdeen
Node #2	. .	AFWL
•		
•		
•		
Node #57	. .	Xerox.

For simplicity, the site names, the node packet handling roles and functions (i.e., TIP, IMP, etc.), and the transmission circuit types (i.e., terrestrial, satellite, etc.) are deleted from the graph.

The characteristics of the above network are summarized in Table 10. The number of nodes is n=57 and the number of links is ℓ=71. That includes

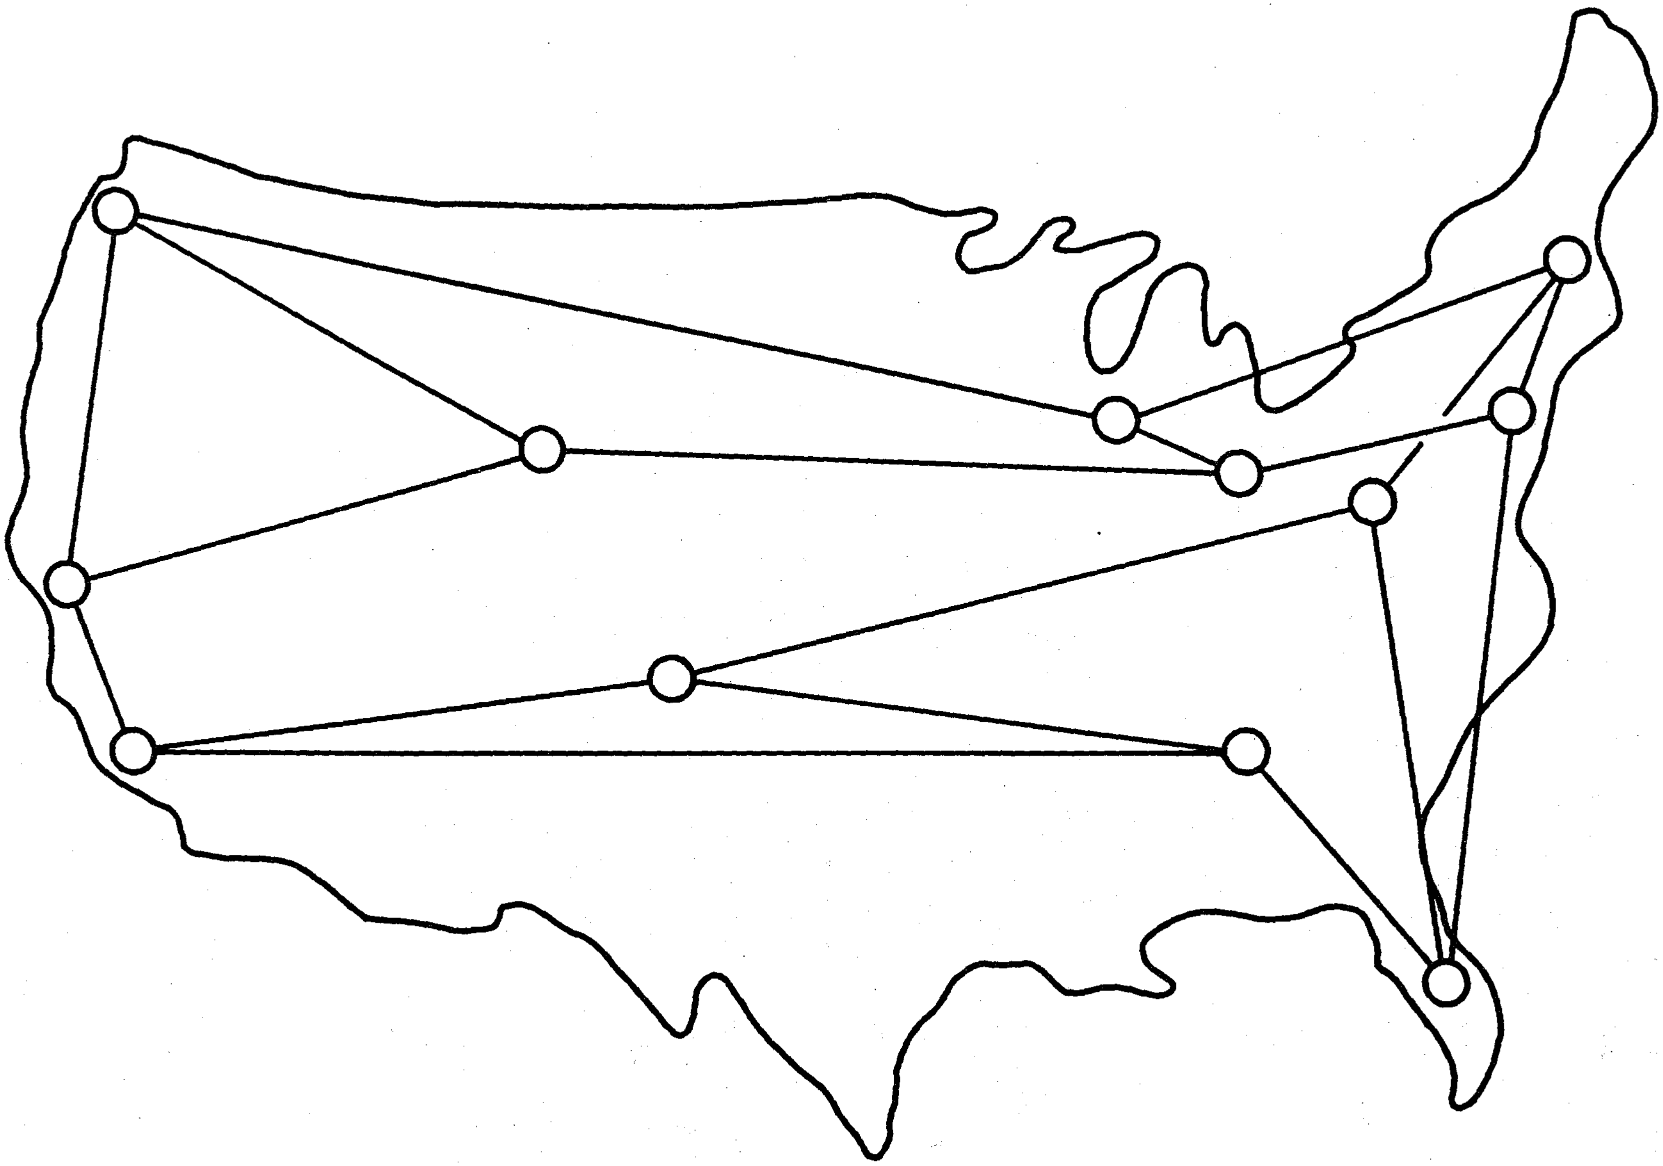

Figure 29. Backbone topology with enhanced survivability.

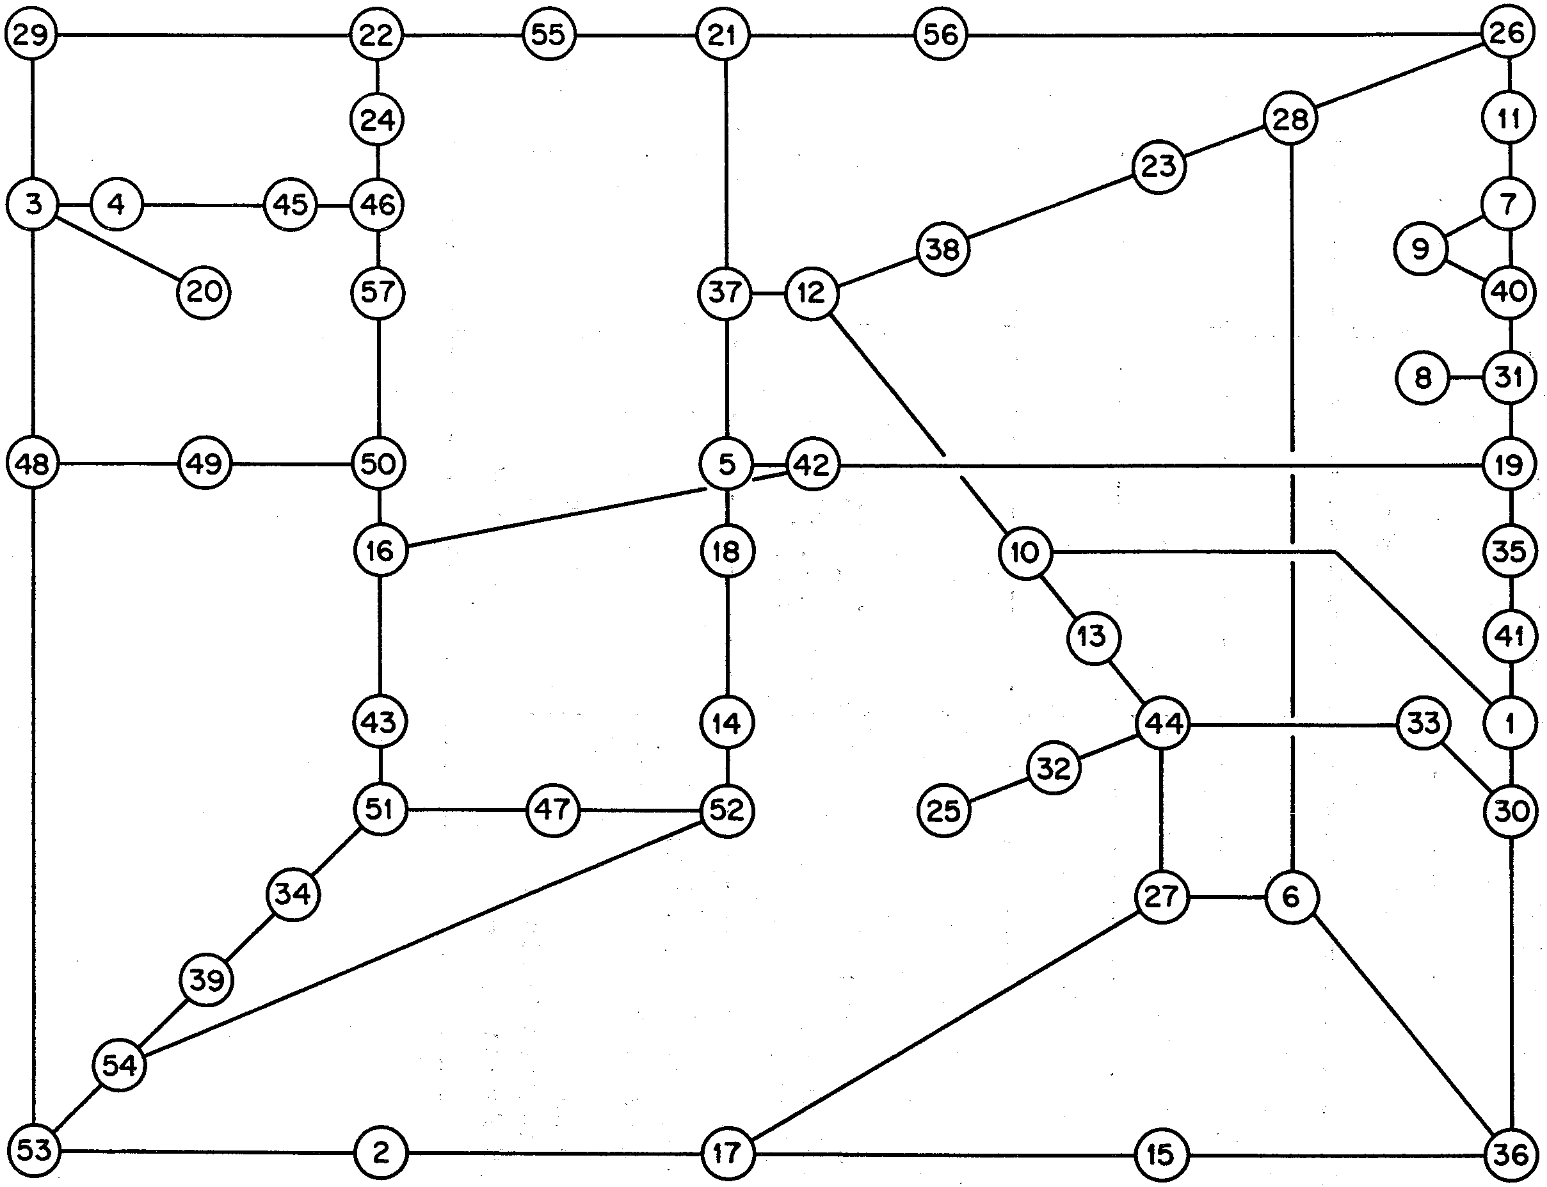

Figure 30. An ARPA network logical map from 1976.

Table 10. Characteristics of the ARPA Network, Vintage 1976

Characteristic	Symbol	Value
Number of Nodes	n	57
Number of Links	ℓ	71
Actual Network Diameter	d	11
Maximum Constrained Diameters:		
Given n and ℓ	d_{max} (n,ℓ)	51
Given n	d_{max} $(n,.)$	56
Given ℓ	d_{max} $(.,\ell)$	71
Connectivity Cross Section of the Network:		
Mean	m	2.12
Standard Deviation	s	.61
Effective Index	$m-s/4$	1.97

overseas nodes in Hawaii and in the United Kingdom, as well as a few deployed satellite circuits. The actual network diameter is $d=11$. This largest distance is realized by counting unity hops, for instance, between node #25 (London) and node #46 (SRI). The three maximum constrained diameters are much larger than 11, but all are in a rather small range between 51 and 71. This implies that in a stress or crisis-caused restoration session, the initial timer setting $t(0)$ should be around one hundred in order to cover the entire damaged or so-suspected network.

A tradeoff to be earnestly considered here pertains to picking a much lower value for the timer, say $t(0)=10$. This would mean that the restoration data exchange would be limited to a mere fraction of the topology, unless, of course, said topology was appropriately disfigured. In many situations, such a limitation may constitute a drawback. However, the negative aspect must be weighed against the potential positive benefits of the shorter scheduled or unscheduled reconstitution sessions. By the way, in networks such as ARPANET or others that utilize adaptive, locally implemented, routing techniques, the advantage of Global tables may not be all that significant.

The connectivity cross section statistics have also been investigated for the network of Figure 30. The results reveal that the mean cross section is $m=2.12$, the standard deviation is $s=0.61$, and the effective index is $m-s/4=1.97$. For assessment of overall topological survivability, the old ARPANET therefore has basically two separate paths between any pair of representative nodes. If one were curious about the available opportunities to increase the effective index to 3, the following can be observed. For a fixed number of nodes, say $n=57$, one must implement at least 86 links to get $m-s/4=3.00$. On the other hand, if the number of links were fixed at $\ell=71$, the number of nodes would have to be reduced to 47 or less. For arbitrary n and ℓ, a necessary (but not always sufficient) condition for $m-s/4=x$ is $2\ell \geq nx$. Thus, in larger networks like this, increases in networkwide cross section are likely to require numerous properly installed links. This may help to carry the offered traffic, but the network design is almost certain to be complex and expensive.

When the finite topology of interest is embedded in an infinite network, the initial timer setting must be such as to cover not only the finite subnetwork, but also a considerable region surrounding it. After all, the interior finite subnetwork might become locally disjointed, while numerous alternate paths could exist in the surrounding infinite network.

In section 3.1, under the heading of "Properties of larger networks," a few elementary topological constructs called the outer shell, the inner core, and the coverage ellipse were introduced. The definitions apply to any given network topologies around one or two starting nodes A and B.

To demonstrate the application of these concepts to the reconstitution process, consider Figure 31. The unbounded stylized network contains an infinite number of nodes located at every crosspoint on the grid. On that infinite grid let there be only five nodes, namely A, B, C, D, and E, that need to establish communications with each other. One refers to their eventual mutual interconnections, whatever their shape, as a local subnetwork.

Let node A initiate the restoration search by sending out the first message and by appending the timer count, say t(0)=9. The timer setting immediately determines the outer region, or shell, and the inner region, or core, around node A. Both play a role in the restoral process for the local subnetwork. Let us explain.

Nodes A and B are the foci of the (A,B) ellipse. The links interior to the ellipse are shortest path candidates to connect A to B. Likewise for node pairs (A,C), (A,D), and (A,E). Links interior to all four ellipses are therefore prime candidates for linking all five nodes, A to E, and thus finding the connectivity for the desired subnetwork. Such a target network can be defined as the intersection of the four coverage ellipses. The intersection is shown in Figure 31 by the accented lines. One can view the so-embedded connectivity as a starting subnetwork to restore connectivity. In a sense then, the choice of t(0)=9 is rather fortunate, as it happens to be just big enough to include all five nodes in the four ellipses. Note that this intersection of ellipses also coincides with the inner core defined earlier in Figure 14.

Larger search regions and larger t(0) values are needed to find alternate routes that are outside the inner core. Candidate regions can be the union of the coverage ellipses, the outer shell (see Figure 14), or any expanse determined with larger t(0) settings by the originating timer. One possible approach, consistent with the t(0)=2d+1 rule for small networks (i.e., Property I in section 3.1), is to define an equivalent subnetwork diameter within the larger network. A maximum (over all node pairs) of the minimum (within the infinite network topology), applied to the finite set of nodes, may be adequate

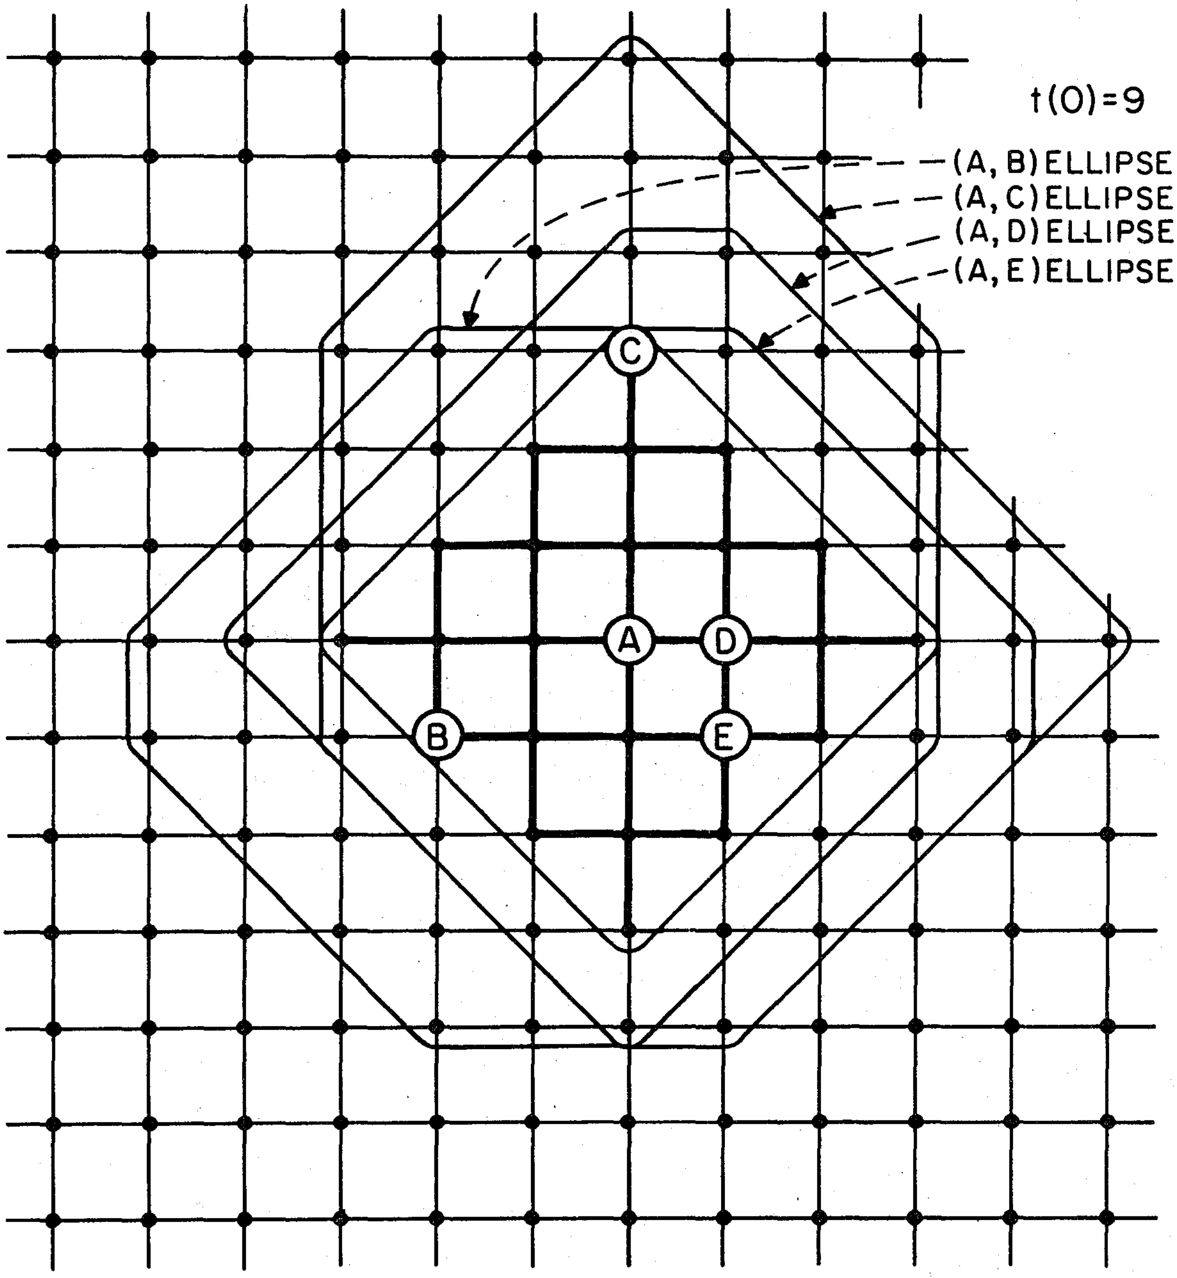

Figure 31. The smallest embedded subnetwork to serve nodes A, B, C, D, and E.

to determine the new "d." In the example of Figure 31, such a definition would result in

$$d = d(B,C) = 6,$$

and in a corresponding timer setting of $t(0)=13$. The so-generated search domain would certainly extend outside the coverage ellipses illustrated. It can be made even larger, perhaps as high as $t(0)\leq 100$, the only real upper bound being the system time budget imposed in section 3.5.

5. CONCLUSIONS

5.1 Summary

This technical report consists of essentially four sections plus an extensive bibliography. They all pertain to specific service survivability aspects of fiber optic networks. Of particular importance is the continuity of end-user data services. Physical survivability is by all means necessary, but as such is not emphasized here. Instead, functional concepts and methods that use distributed automation are the focus for rapid connectivity restoration under stress. The latter is to guarantee minimum service outages for a fiber network that is under stress or has suffered some node and link outages.

Section 1 is an introduction to the background and objectives of this work. It lists seven conditions or issues that variously influence and in some cases determine the scope of the subsequent sections.

Section 2 deals with survivability. Threats to network connectivity and services affect all networks including fiber optic networks. The special fiber advantages of high data rates, especially when used in conjunction with latest high-speed integrated circuit facilities (e.g., VLSI-based processors, memories, switches, etc), have the potential to minimize the service downtime to users. Automation is essential here. To compare network survivabilities of different configurations, a quantitative measure for the term "network survivability" is defined. The chosen survivability metric depends only on the topology, i.e., on the arrangement of nodes and links.

The process starts with a definition of a unique number x that depicts the connectivity cross section between a particular pair of nodes. Since there are many node pairs in a network, one can assemble a histogram over all x's. Given said histogram of all connectivity cross sections, one can further calculate its mean m, standard deviation s, and, if need be, other higher moments of x. In Section 2 it is suggested that a quantity, m-s/4, be called the effective cross-section index and that it be viewed as a relative measure of survivability for the given topology. Ranking of indices provides an ordering of relative network survivabilities. If one is tasked to improve the survivability index of a network, this can be achieved by increasing m and/or by decreasing s.

Section 3 is devoted to service restoration. External and internal, automatic or manual, events can initiate certain network procedures that amount to confirmation of old or discovery of new modified network connectivity and performance status. These restoration intervals are meant to be short, perhaps on the order of one second. They are to recur infrequently. When scheduled, their repetition rate should not exceed once per hour. When unscheduled, the restoration sessions would be triggered by any valid stress or emergency conditions.

Node activities during the restoration events are based on the premise that the optical fiber has a huge throughput capability. Thus there is no immediate need to minimize or to economize on the number of signaling bits exchanged between surviving nodes. The restoration protocols need not follow the precedents set by sophisticated and clever protocols designed for other transmission media, where signaling bits are a scarce commodity. For a starter, any simple restoral protocol may be contemplated here. It does not seem to matter whether the protocol is or is not wasteful of signaling bits. Latest and fastest parallel implementations, however, should be used in the node processors.

To outline a possible scheme for nodal activities during a restoration window, the node functions are divided into two classes. The first class consists of primary actions that initiate the search-and-broadcast sessions and specify their durations. The second class consists of everything else done by all other nodes involved in the process. Special timers are used for session initiation, continuation, and termination. The same timers also appear useful for exclusion of unwanted interference or break-ins by other, nearly simultaneous, reconstitution sessions. For smaller networks, the timer management is related to the size of the network diameter. The diameter, however, can only be approximated or upperbounded because of constraints imposed by uncertainty or partial knowledge about the surviving elements in the new topology. For larger networks, where exhaustive searches are unrealistic, local coverage areas are as determined by the initial timer setting.

The initial message from the primary node can be quite short. Not knowing who else is operationally alive, or who is connected to whom, the first broadcast may contain only the sender's own identity. Thereafter, the message grows rapidly in size as more and more is learned about the surviving configuration. The ultimate scope of network data to be exchanged can include: Network Link Capacities, Connectivity Matrix, Traffic Carried, Traffic Offered, Facility Restriction, and perhaps Routing Tables. In toto, this can be a sequence or a vector of up to six large arrays. The array sizes are estimated for several choices of global or local, maximal or minimal, array coverage scenarios.

An array, called the distance matrix or D, is introduced. It need not be sent back and forth, as it appears well suited for local computation at the session conclusion. Matrix D can be used in a distributed fashion to automate adaptive routing procedures.

Section 3 also includes brief considerations of link-level protocols, standard frame formats, as well as routing and shortest path algorithms. It concludes with an assessment of time budgets for optical signal propagation within fibers versus the transmission (i.e., modem) time. The total session is apt to be dominated by the propagation delays, yet for the network configurations postulated here the 1-second session objective seems realizable.

Section 4 presents illustrative network examples and is divided into small and large network cases. The first small example is a twelve-node network along major interstate highways in the Western United States. Under normal and under damage conditions, its survivability and restoration characteristics are discussed. A slightly different approach is taken to the second example, which is a hypothetical CONUS backbone with a typically thin East-to-West connectivity. Disregarding economics and traffic, for the same number of nodes and links, the effective connectivity index can be substantially increased by a simple rearrangement of the connecting links. For a larger network, the now outdated 1976 ARPANET is considered. That network map shows 57 nodes and 71 links. Its complexity is probably near the limit of how far one should go via paper and pencil, to illustrate the aforementioned survivability and reconstitution concepts.

5.2 Remaining Issues

A number of issues and questions have been encountered in this brief study. Only a few have been adequately covered in the present report. Of the remaining issues, the following six seem to require further attention.

(1) Broader definition of "survivability."

A general, quantitative, joint function-and-service oriented definition is still needed. The definition given here focuses on topology and should not be construed as the final solution. Other definitions should be investigated. One envisions the desired ultimate definition as a useful tool to be employed in the assessment of system alternatives. To what extent such a goal can be realized, is not at all clear. Perhaps, basic algorithms, or heuristics, or approximations should be developed first to make substantial progress on this issue.

(2) Algorithms for cross sections.

Despite previous work in graph theory, there is no simple and quick algorithm for computation of individual, point-to-point, cross sections in an arbitrary network. There is also no direct way known to get the moments of the cross section histograms. In real life all links are not full duplex, nor are they of equal importance. The corresponding asymmetry and weight generalizations appear needed, but so does a more thorough justification for the m-s/4 metric to represent the one and only admissible connectivity index. Perhaps other functional properties of the histogram or something entirely different should be used.

(3) Search area limitation in large networks.

When faced with one very large network or with several networks interconnected via gateways, the question arises of how best to keep the restoration search within controlled bounds. To avoid human patch-throughs, automated standard schemes should be developed.

(4) Applications of the distance matrix.

On a preliminary level, the distance matrix shows the diameter of the network itself or of its possible subnetworks. Combined with the capability to identify required shortest path distances, when specified tandem nodes are assumed between the source and the destination, there seems to be an opportunity here for accelerated identification of alternate routes. Methods to split and select best tandem points (see Figure 23) need to be defined and analyzed.

(5) Design of node processes for a fiber optic network.

For planning, design, and implementation purposes, the node processor functions and features should be defined in increasing detail. That design can be based on existing, off-the-shelf, already available, standard software and hardware. Or it could pursue more innovative approaches, perhaps along the lines discussed here or found anywhere in the latest computer and/or fiber optic network literature.

(6) Standard and nonstandard protocol issues.

To plan for service continuity, interoperability between different networks must be considered. To that end, existing and planned interface standards must be included as effectively as possible. The pertinent work of such standard setting National and International organizations, as ANSI, CCITT, IEEE, ISO, and others, applies to this issue. The earliest use of the seven layer OSI Reference Model is recommended in order to avoid potential future incompatibilities. Initial efforts should be directed towards the lower levels, especially to the Link and Network levels. Planning for the selection of ISDN options should commence. As an issue, ISDN has not been emphasized above. However, future efforts may have to be directed at ISDN.

6. BIBLIOGRAPHY

Albanese, A., Editor (1985), Special Issue on Fiber Optics for Local Communications, IEEE J. Sel. Areas in Commun. SAC-3, pp. 813-962.

ANSI (1984), FDDI token ring media access control, Standard X3T9/84-100/X3T9:5/83 - 16 Rev. 7.2 (Oct. 26).

ANSI/IEEE (1985), Carrier sense multiple access and collision detection, Draft International Standard, ISBN 471 - 82749-5.

Appel, J.J. (1986), Network operations--A major opportunity in evolving digital networks, IEEE Commun. Mag. 24, January, pp. 39-42.

Atkins, J.D. (1980), Path control: The transport network of SNA, IEEE Trans. Commun. COM-28, pp. 527-538.

Ayoub, J.N., and I.T. Frisch (1970), Optimally invulnerable directed communication networks, IEEE Trans. Commun. COM-18, pp. 484-489.

Babcock, W.L. (1985), Fiber-optic integration into modern telecommunication network design, Telecommunications, December, pp. 64-72.

Baran, P. (1964), On distributed communications networks, IEEE Trans. Commun. Systems CS-12, pp. 1-9.

Baran, P. (1965), On survivability of networks, IEEE Trans. Commun. Technology COM-13, pp. 379-380.

Basch, E.E., and T.G. Brown (1985), Introduction to coherent optical fiber transmission, IEEE Commun. Mag. 23, May, pp. 23-30.

Bell, P.R., and K. Jabbour (1986), Review of point-to-point network routing algorithms, IEEE Commun. Mag. 24, January, pp. 34-38.

Bell, T.E. (1986), Technology'86--Communications, IEEE Spectrum 23, January, pp. 49-52.

Benes, V.E. (1965), Mathematical Theory of Connecting Networks and Telephone Traffic (Academic Press, New York, NY).

Benes, V.E. (1966), Programming and control problems arising from optimal routing in telephone networks, Bell System Tech. J. 45, pp. 1373-1438.

Biondi, E., L. Divieti, C. Roveda, and R. Schmid (1968), Optimal solution for a class of assignment and transportation problems, 2nd Int. Conf. Computing Methods in Optimization Problems, San Remo, Italy. Paper published in 1969 as part of Computing Methods in Optimization Problems-2, L.A. Zadeh, L.W. Neustadt, and A.V. Balakrishnan, Editors (Academic Press, New York, NY).

Boehm, B.W., and R.L. Mobley (1969), Adaptive routing techniques for distributed communications systems, IEEE Trans. Commun. Technology COM-17, pp. 340-349.

Boesch, F.T., and I.T. Frisch (1968), On the smallest disconnecting set in a graph, IEEE Trans. Circuit Theory CT-15, pp. 286-288.

Boesch, F.T., and R.E. Thomas (1970), On graphs of invulnerable communication nets, IEEE Trans. Circuit Theory CT-17, pp. 183-192.

Brosio, A., F. Gagliardi, L. Lambarelli, G. Panarotto, D. Roffinella, and M. Sposini (1985), A reconfigurable high-speed optical system for integrated local communications, IEEE J. Sel. Areas in Commun. SAC-3, pp. 825-834.

Burr, W.E. (1986), The FDDI optical data link, IEEE Commun. Mag. 24, May, pp. 18-23.

Cantor, D.G., and M. Gerla (1974), Optimal routing in a packet-switched computer network, IEEE Trans. Commun. COM-23, pp. 1062-1069.

Cavers, J.K. (1975), Cutset manipulation for communication network reliability estimation, IEEE Trans. Commun. COM-23, pp. 569-575.

Cederbaum, I., and I. Paz (1973), On optimal routing through communication nets, IEEE Trans. Commun. COM-21, pp. 936-941.

Chang, H. (1986), Prospects for optical fiber communications in China, IEEE Commun. Mag. 24, April, pp. 18-24.

Chou, W., and H. Frank (1970), Survivable communication networks and the terminal capacity matrix, IEEE Trans. Circuit Theory CT-17, pp. 192-197.

Cohen, L.G. (1986), Trends in U.S. broad-band fiber optic transmission systems, IEEE J. Sel. Areas in Commun. SAC-4, pp. 488-497.

Cypser, R.J. (1978), Communications Architecture for Distributed Systems (Addison-Wesley Publishing Co., Reading, MA).

Dash, J., G. Duerksen, S. Federico, L. Forman, A. Russo, and R. Warren (1985), FT3C optical fiber communication system: Preliminary draft of EMP test and assessment final report. Prepared for the Defense Nuclear Agency by AT&T Technologies, Federal Systems Division, Greensboro, NC.

Davies, D.W. (1972), The control of congestion in packet-switching networks, IEEE Trans. Commun. COM-20, pp. 546-550.

Davies, D.W., D.L.A. Barber, W.L. Price, and C.M. Solomonides (1980), Computer Networks and Their Protocols (John Wiley & Sons, New York, NY).

De Bortoli, M., A. Moncalvo, and G. Pellegrini (1985), Multiplexing in the optical distribution network: a technical and economic comparison, CSELT Technical Reports XIII, pp. 175-178.

Decina, M., and D. Vlack, Editors (1983), Special Issue on Packet Switched Voice and Data Communication, IEEE J. Sel. Areas in Commun. SAC-1, pp. 961-1139.

Dennis, J.B. (1964), Distributed solution to network programming problems, IEEE Trans. Commun. Systems CS-12, pp. 176-184.

Dijkstra, E.W. (1959), A note on two problems in connexion with graphs, Numer. Math. 1, pp. 269-271.

Dinn, N.F., A.G. Weygand, and D.M. Garvey (1986), Digital interconnection of dissimilar digital networks, IEEE Commun. Mag. 24, April, pp. 12-17.

Duc, N.Q., and E.K. Chew (1985), ISDN protocol architecture, IEEE Commun. Mag. 23, March, pp. 15-22.

Even, S. (1975), An algorithm for determining whether the connectivity of a graph is at least k, SIAM J. Comput. 4, pp. 393-396.

Finley, M.R. (1984), Optical fibers in local area networks, IEEE Commun. Mag. 22, August, pp. 22-35.

Fischer, M.J. (1983), A performance model for loss systems with crisis calls, IEEE Trans. Commun. COM-31, pp. 1212-1216.

Fishman, G.S. (1986), A Monte Carlo sampling plan for estimating network reliability, Operations Research 34 (4), pp. 581-594.

Ford, L.R., and D.R. Fulkerson (1962), Flows in Networks (Princeton University Press, Princeton, NJ).

Frank, H. (1967), Vulnerability of communication networks, IEEE Trans. Commun. Technology COM-15, pp. 778-789.

Frank, H. (1974), Survivability analysis of command and control communications networks--Parts I and II, IEEE Trans. Commun. COM-22, pp. 589-595 and 596-605.

Frank, H., and W. Chou (1970), Connectivity considerations in the design of survivable networks, IEEE Trans. Circuit Theory CT-17, pp. 486-490.

Frank, H., and I.T. Frisch (1970), Analysis and design of survivable networks, IEEE Trans. Commun. Technology COM-18, pp. 501-519.

Frank, H., and I.T. Frisch (1971), Communication, Transmission, and Transportation Networks (Addison-Wesley, Reading, MA).

Fratta, L. (1975), Basic analytical techniques and routing procedures. Paper published in NATO Advanced Study Institute Series, Computer Communication Networks, R.L. Grimsdale and F.F. Kuo, Editors, pp. 35-62 (Noordhoff Integnational Publishing, Leyden, Netherlands).

Fratta, L., M. Gerla, and L. Kleinrock (1974), The flow deviation method: An approach to store-and-forward communication network design, Networks 3, pp. 97-133.

Frisch, I.T. (1963), Optimum routes in communication systems with channel capacities and channel reliabilities, IEEE Trans. Commun. Systems CS-11, pp. 241-245.

Frisch, I.T. (1975), Technical problems in nationwide networking and interconnection, IEEE Trans. Commun. COM-23, pp. 78-88.

Gallager, R.G. (1976), Basic limits on protocol information in data communication networks, IEEE Trans. Inform. Theory IT-22, pp. 385-398.

Gallager, R.G. (1977), A minimum delay routing algorithm using distributed computation, IEEE Trans. Commun. COM-25, pp. 73-85.

Gallager, R.G. (1981), Applications of information theory to data communication networks. Paper published in NATO Advanced Study Institute Series, New Concepts in Multi-User Communication, J.K. Skwirzynski, Editor, pp. 63-82 (Sijthoff & Noordhoff International Publishers, Alphen aan den Rijn, Netherlands).

Geer, N., and P. Stromecky (1981), Gridnet simulation: Volume I, Model descriptions and operating instructions, GE Final Report No. 81NV008, GE Space Systems Division, Hunstville, AL (Prepared under contract for the National Bureau of Standards, Gaithersburg, MD 20234).

Gerla, M. (1984), Controlling routes, traffic rates, and buffer allocation in packet networks, IEEE Commun. Mag. 22, November, pp. 11-23.

Gill, A., and I.L. Traiger (1968), Computation of optimal paths in finite graphs, 2nd Int. Conf. Computing Methods in Optimization Problems, San Remo, Italy. Paper published in 1969 as part of Computing Methods in Opimization Problems-2, L.A. Zadeh, L.W. Neustadt, A.V. Balakrishnan, Editors (Academic Press, New York, NY).

Green, P.E., Jr., (Editor) (1980), Special Issue on Computer Network Architectures and Protocols, IEEE Trans. Commun. COM-28, pp. 409-677.

Greene, T.V., and L.C. Brown (1969), Route control in AUTOVON electronic switching centers, IEEE Trans. Commun. Technology COM-17, pp. 442-446.

Grieco, D.M. (1977), Comparison of network search schemes, IEEE Trans. Commun. COM-25, pp. 459-462.

Hakimi, S.L. (1969), An algorithm for construction of the least vulnerable communication network or the graph with the maximum connectivity, IEEE Trans. Circuit Theory CT-19, pp. 229-230.

Hara, E. (1983), A fiber-optic broadband LAN/OCS using a PBX, IEEE Commun. Mag. 21, October, pp. 22-27.

Harary, F., and E.M. Palmer (1973), Graphical Enumeration (Academic Press, New York, NY).

Henry, P.S. (1985), Introduction to lightwave transmission, IEEE Commun. Mag. 23, May, pp. 12-16.

Hitchner, L.E. (1968), A comparative study of the computational efficiency of shortest path algorithms, Operations Research Center Report ORC 68-25, University of California, Berkeley, CA.

Hofacker, R.Q., and I. Jacobs (1986), The wideband world of fiber, AT&T Bell Laboratories Record, March, pp. 14-21.

Hsieh, W.N., and I. Gitman (1984), Routing strategies in computer networks, IEEE Computer Mag. 17, June, pp. 46-56.

Huynh, D., H. Kobayashi, and F.F. Kuo (1977), Optimal design of mixed-media packet-switching networks: Routing and capacity assignment, IEEE Trans. Commun. COM-25, pp. 148-157.

IEEE (1982), Local Area Network Standards, CSMA/CD Access Method and Physical Layer Specification, IEEE Project 802-3-82/0.1 10 (New York, NY).

IEEE (1985a), Token passing bus access method and physical layer specification, Standard 802.4 (ISO/DIS 8802/4).

IEEE (1985b), Token ring access methods and physical layer specifications, Standard 802.5 (ISO/DIS 8802/5).

Ilyas, M., and H.T. Mouftah (1985), Performance evaluation of computer communications networks, IEEE Commun. Mag. 23, April, pp.18-29.

Kao, C.K. (1982), Optical Fiber Systems: Technology, Design, and Applications (McGraw-Hill Book Co., New York, NY).

Keck, D.B. (1985), Fundamentals of optical waveguide fibers, IEEE Commun. Mag. 23, May, pp. 17-22.

Keiser, B.E., and E. Strange (1985), Digital Telephony and Network Integration (Van Nostrand Reinhold Co., New York, NY).

Kleinrock, L. (1972), Communication Nets (Dover Publications, New York, NY).

Kleinrock, L. (1975), Queueing Systems, Volume I: Theory (John Wiley & Sons, New York, NY).

Kleinrock, L. (1976), Queueing Systems, Volume II: Computer Applications (John Wiley & Sons, New York, NY).

Kleinrock, L. (1978), Principles and lessons in packet communications, Proc. IEEE 66, pp. 1320-1329.

Kleinrock, L. (1985), Distributed systems, Comm. of the ACM 28, pp. 1200-1213.

Kleitman, D. (1969), Methods for investigating the connectivity of large graphs, IEEE Trans. Circuit Theory CT-16, pp. 232-233.

Koyama, M., and T. Itoh (1981), Present and future of large capacity optical fiber transmission, IEEE Commun. Mag. 19, May, pp. 4-10.

Lai, W.S. (1985), The integrity of data delivery in computer networks, IEEE Trans. Commun. COM-33, pp. 1222-1224.

Lawler, E.L. (1976), Combinatorial Optimization: Networks and Matroids (Holt, Rinehart and Winston, New York, NY).

Leon-Garcia, A., (Editor) (1986), Special Issue on Network Performance Evaluation, IEEE J. Sel. Areas in Commun. SAC-4, pp. 773-996.

Li, V.O., and J.A. Sylvester (1984), Performance analysis of networks with unreliable components, IEEE Trans. Commun. COM-32, pp. 1105-1110.

Linfield, R.F., and M. Nesenbergs (1985), Military access area characterization, NTIA Report 85-185, U.S. Department of Commerce, National Telecommunications and Information Administration, Institute for Telecommuncation Sciences, Boulder, CO 80303-3328.

Lokerson, D.T. (1983), A communications system which prioritizes the queues, service channels, and traffic, IEEE Trans. Commun. COM-31, pp. 113-118.

Lundgren, C.W., and P.S. Venkatesan (1986), Applications of video on fiber cable, IEEE Commun. Mag. 24, May, pp. 33-49.

Lynch, J.F. (1985), Trends in fiber optics, Signal Mag. 39, June, pp. 33-44.

Maione, T.L., and D.D. Sell (1977), Experimental fiber-optic transmission system for interoffice trunks, IEEE Trans. Commun. COM-25, pp. 517-523.

Malhotra, V.M., M.P. Kumar, and S.N. Maheshwari (1978), An O(1V1^{3}) algorithm for finding maximum flows in networks, Inf. Proc. Lett. 7, pp. 277-278.

Mansuripur, M., J.W. Goodman, E.G. Rawson, and R.E. Norton (1980), Fiber optics receiver error rate prediction using the Gram-Charlier series, IEEE Trans. Commun. COM-28, pp. 402-407.

McQuillan, J.M., G. Falk, and I. Richer (1978), A review of the development and performance of the ARPANET routing algorithm, IEEE Trans. Commun. COM-26, pp. 1802-1811.

McQuillan, J.M., I. Richer, and E.C. Rosen (1980), A new routing algorithm for the ARPANET, IEEE Trans. Commun. COM-28, pp. 711-719.

Merlin, P.M., and A. Segall (1979), A failsafe distributed routing protocol, IEEE Trans. Commun. COM-27, pp. 1280-1287.

Midwinter, J.E. (1982), Optical fiber communications systems development in the UK, IEEE Commun. Mag. 20, January, pp. 6-11.

Minoli, D., and E.H. Lipper (1980), Cost implications for survivability of terrestrial networks under malicious failure, IEEE Trans. Commun. COM-28, pp. 1668-1674.

Misra, K.B. (1970), An algorithm for the reliability evaluation of redundant networks, IEEE Trans. Reliability R-19, pp 146-151.

Moran, J.R. (1985), Future fiber optic systems, Signal Mag. 39, August, pp. 88-91.

Morriss, B. (1985), Planning for NS/EP telecommunications, Signal Mag. 39, August, pp. 67-71.

Nesenbergs, M. (1978), Bounds on the number of possible networks, IEEE Trans. Commun. COM-26, pp. 1315-1316.

Netes, V.A. (1981), Selection of generalized quality indicators of communications networks, Telecomm. and Radio Engineering (Translated from Russian), 35/36, pp. 10-13.

Netravali, A.N., and Z.L. Budrikis (1985), A broadband local area network, AT&T Tech. J. 64, pp. 2449-2465.

Niiro, Y., and H. Yamamoto (1986), The international long-haul optical-fiber submarine cable system in Japan, IEEE Commun. Mag. 24, May, pp. 24-32.

Ogawa, K. (1982), Considerations for single-mode fiber systems, Bell System Tech. J. 61, pp. 1919-1931.

Ohnsorge, H., Editor (1986), Special Issue on Broad-Band Communications Systems, IEEE J. Sel. Areas in Commun. SAC-4, pp. 425-648.

Peach, D.F. (1986), Trends toward a more stress-resistant fiber optic telecommunication installation, NTIA Technical Memorandum 86-116, U.S. Department of Commerce, National Telecommunications and Information Administration, Institute for Telecommunication Sciences, Boulder, CO.

Personick, S.D. (1978), Fiber optic communication--A technology coming of age, IEEE Commun. Mag. 16, March, pp. 12-20.

Personick, S.D., (Editor) (1978), Special Issue on Fiber Optics, IEEE Trans. Commun. COM-26, pp. 945-1126.

Personick, S.D., (Editor) (1983), Special Issue on Fiber Optic Systems, IEEE J. Sel. Areas in Commun. SAC-1, pp. 353-575.

Pierce, J.R. (1978), Optical channels: Practical limits with photon counting, IEEE Trans. on Commun. COM-26, pp. 1819-1821.

Pollack, M. (1965), A bibliography of communication network studies, IEEE Trans. Commun. Technology COM-13, pp. 552-554.

Prisco, J.J., and R.J. Hoss (1985), Fiber optic regional area networks, IEEE Commun. Mag. 23, November, pp. 26-39.

Prosser, R.T. (1962), Routing procedures in communication networks--Part I: Random procedures and Part II: Directory procedures, IRE Trans. Commun. Systems CS-10, pp. 322-329 and 329-335.

Rey, R.F., Editor (1984), Engineering and Operations in the Bell System (AT&T Bell Laboratories, Murray Hill, NJ).

Rice, D.H., and G.E. Keiser (1985), Application of fiber optics to tactical communication systems, IEEE Commun. Mag. 23, May, pp. 46-57.

Riordan, J. (1960), The enumeration of trees by height and diameter, IBM J. of Res. and Develop. 4, pp. 473-478.

Robertazzi, T.G., and P.E. Sarachik (1986), Self-organizing communication networks, IEEE Commun. Mag. 24, January, pp. 28-33.

Ross, F.E. (1986), FDDI - a tutorial, IEEE Commun. Mag. 24, May, pp. 10-17.

Rudin, H., and H. Mueller (1980), Dynamic routing and flow control, IEEE Trans. Commun. COM-28, pp. 1030-1039.

Rudin, H., Editor (1981), Special Issue on Congestion Control in Computer Networks, IEEE Trans. Commun. COM-29, pp. 373-535.

Rudin, H. (1985), An informal overview of formal protocol specification, IEEE Commun. Mag. 23, March, pp. 46-52.

Ryan, J.S., Editor (1985), Special Issue on Telecommunications Standards, IEEE Commun. Mag. 23, January, pp. 6-62.

Salz, J. (1985), Coherent lightwave communications, AT&T Tech. J. 64, pp. 2153-2209.

Salz, J. (1986), Modulation and detection for coherent lightwave communications, IEEE Commun. Mag. 24, June, pp. 38-49.

Scerbo, L.J., J.P. Varachi, H. Murata, and P.L. Pope, (Editors) (1986), Special Issue on Engineering and Field Experience with Fiber Optic Systems, IEEE J. Sel. Areas in Commun. SAC-4, pp. 649-772.

Schroeder, M.A., and K.T. Newport (1986), A methodology for quantifying network connectivity using a graph theory approach, MITRE Tech. Report MTR 9278, MITRE Corp., Rome, NY (Prepared under contract to Rome Air Development Center, Griffiss AFB, NY 13441-5700).

Schwartz, M., and T. Stern (1980), Routing techniques used in computer communication networks, IEEE Trans. Commun. COM-28, pp. 539-552.

Schwartz, M.A., and D.G. Messerschmidt, (Editors) (1981), Special Issue on Maintenance, Control, and Protection of Remote Electronics, IEEE Trans. Commun. COM-29, pp. 1413-1445.

Schwartz, M.I. (1984), Optical fiber transmission--From concepts to prominence in 20 years, IEEE Commun. Mag. 22, May, pp. 38-48.

Segall, A., P.M. Merlin, and R.G. Gallager (1978), A recoverable protocol for loop-free distributed routing, Conference Proc. of ICC'78, pp. 3.5.1-3.5.5 (Toronto, Canada).

Segall, A. (1983), Distributed network protocols, IEEE Trans. Inform. Theory IT-29, pp. 23-35.

Siperko, C.M. (1985), LaserNet--A fiber optic intrastate network (Planning and engineering considerations), IEEE Commun. Mag. 23, May, pp. 31-45.

Sproule, D.E., and F. Mellor (1981), Routing, flow, and congestion control in the Datapac network, IEEE Trans. Commun. COM-29, pp. 386-391.

Stanley, I.W. (1985), A tutorial review of techniques for coherent optical fiber transmission systems, IEEE Commun. Mag. 23, August, pp. 37-53.

Stern, T.E. (1976), A class of decentralized routing algorithms using relaxation, Proc. IEEE Nat. Conf. Telecommunications, Dallas, TX, pp. 42.1-1 to 42.1-5.

Suh, S.Y., S.W. Granlund, and S.S. Hegde (1986), Fiber-optic local area network topology, IEEE Commun. Mag. 24, August, pp. 26-32.

Sze, D.T. (1985), A metropolitan area network, IEEE J. Sel. Areas in Commun. SAC-3, pp. 815-824.

Takasaki, Y., M. Tanaka, N. Maeda, K. Yamashita, and K. Nagano (1976), Optical pulse formats for fiber optic digital communications, IEEE Trans. Commun. COM-24, pp. 404-413.

Tanenbaum, A.S. (1981), Computer Networks (Prentice-Hall, Inc., Englewood Cliffs, NJ).

Toida, S. (1973), Generation of all cutsets of a bipath network, IEEE Trans. Commun. COM-21, pp. 1414-1417.

Tymes, L.W. (1981), Routing and flow control in the TYMNET, IEEE Trans. Commun. COM-29, pp. 392-398.

Weber, J.H. (1962), Some traffic characteristics of communication networks with automatic alternate routing, Bell System Tech. J. 41, pp. 1201-1247.

Wilkov, R.S. (1972), Analysis and design of reliable computer networks, IEEE Trans. Commun. COM-20, pp. 660-678.

Williams, T.G. (1963), The design of survivable communications networks, IEEE Trans. Commun. Systems CS-11, pp. 230-241.

Yen, J.Y. (1968), Some algorithms for finding shortest routes through the general network, 2nd Int. Conf. Computing Methods in Optimization Problems, San Remo, Italy. Paper published in 1969 as part of Computing Methods in Optimization Problems-2, L.A. Zadeh, L.W. Neustadt, and A.V. Balakrishnan, Editors (Academic Press, New York, NY).

FORM **NTIA-29**
(4-80)

U.S. DEPARTMENT OF COMMERCE
NAT'L. TELECOMMUNICATIONS AND INFORMATION ADMINISTRATION

BIBLIOGRAPHIC DATA SHEET

1. PUBLICATION NO.
NTIA Report 87-214
NCS TIB 87-9

2. Gov't Accession No.

3. Recipient's Accession No.

4. TITLE AND SUBTITLE

FIBER OPTIC NETWORKS AND THEIR SERVICE SURVIVAL

5. Publication Date
March 1987

6. Performing Organization Code
ITS.N1

7. AUTHOR(S)
Martin Nesenbergs

8. PERFORMING ORGANIZATION NAME AND ADDRESS
National Telecommunications and Information Admin.
Institute for Telecommunication Sciences
325 Broadway
Boulder, CO 80303

9. Project/Task/Work Unit No.

10. Contract/Grant No.

11. Sponsoring Organization Name and Address
National Communications System
8 S. Courthouse Rd.
Arlington, VA 20305

12. Type of Report and Period Covered

13.

14. SUPPLEMENTARY NOTES

15. ABSTRACT *(A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here.)*

The objective of this study is to look at fiber optic networks in a predominately functional domain and to assess their potential survivability advantages from that point of view. As a consequency, service survivability is emphasized far more than physical survivabily, although physical existence of facilities is a definite prerequisite for all telecommunications services.

The need for a quantitative (or formal, or unique, or numerical) definition of the term "survivability" is addressed. The report proposes a partial solution of this problem. It introduces a network-related quantity, defined with the moments of the connectivity cross section histograms, that appears to posses many of the properties wanted for measuring and comparing survivabilities of different topologies. For lack of a better name, that quantity many be called the effective topological survivability index.

(con'd)

16. Key Words *(Alphabetical order, separated by semicolons)*

17. AVAILABILITY STATEMENT

☒ UNLIMITED.

□ FOR OFFICIAL DISTRIBUTION.

18. Security Class. *(This report)*
Unclassified

19. Security Class. *(This page)*
Unclassified

20. Number of pages
121

21. Price:

☆ U.S. Government Printing Office: 1980—678-495/529

ABSTRACT (con'd) Block 15

The fiber advantage of large data throughput, typically in tens of Mb/s, must be exploited when connectivity or other network status is in doubt. This is part of the network reconstitution or restoration issue. Outlines of procedures, protocols, and formats are given to achieve comprehensive network-wide restoral for small but still realistic networks. The information fields of extensive reconstitution data arrays are possible and advisable. If transmitted, received, and stored rapidly, and not processed in a lengthy manner, these data arrays are shown to offer unprecedented restoral opportunities. Through locally or regionally focused restoration processes the methods appear practicable even for very large networks.

The conclusion is that full-scale automation is essential. It should be distributed to all nodes of the network and its implementation should be with the very highest speed parallel processors. Any node that survives should be capable of both initiating and participating in the network restoration sessions. Thus, centralized hierarchical controls are to be avoided.