

A BRIEFING ON THE DISTRIBUTED ADAPTIVE
MESSAGE-BLOCK NETWORK

Paul Baran

April 1965

A BRIEFING ON THE DISTRIBUTED ADAPTIVE
MESSAGE-BLOCK NETWORK

Paul Baran *

The RAND Corporation, Santa Monica, California

LIMITATIONS OF CURRENT COMMUNICATIONS

We are witnessing evolution of our national military doctrine away from massive response--where only the "go" word need be transmitted--toward a more flexible response. This change creates new problems for the command and control system designer.

Figure 1 sketches the increased command and control in communications implied by flexible response doctrines. As we move to the right of the scale, increasing flexibility of response, the amount of military communications that must survive attack markedly increases. The degree of flexibility of response possible may be limited primarily by the availability of survivable communications. If this be true, our choice of military doctrine may be dictated more by the availability of communications than by desire.

We feel that communications available to the military may fall short in several respects (see Fig. 2). These include:

* Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

This paper was presented to The RAND Corporation's Air Force Advisory Group and Board of Trustees meeting, November 1964.

**DEMAND FOR
MILITARY
COMMUNICATIONS**




Fig. 1 — Command communications requirements posed by a flexible response doctrine

- SURVIVABILITY
- SECRECY
- ERROR RATE/QUALITY
- DELAY TIMES
- VOLUME
- FLEXIBILITY (USED-TO - -USED)
- COST

Fig. 2—How is command and control for flexible response limited by communications?

Survivability. Existing communications networks are, for the most part, highly vulnerable to overt and covert attacks directed against communications.

Secrecy. Most of our military communication is transmitted with negligible protection to the loss of critical information by an eavesdropper. Only a small portion of military communications is cryptographically protected--and the military is often highly constrained in subjects discussed over the telephone for fear of eavesdropping.

Error Rate/Quality. Those who must connect computers with other computers or with remote entry devices are not completely satisfied with the high error rate of the present-day communication networks. Without great care and much ancillary equipment, computers are very intolerant of errors caused by the communications links.

Delay Times. Partially because of the high costs of communication, military hard copy (paper) traffic is primarily a center-to-center operation with long delay times between end users.

Flexibility (User to User). We lack the ability to achieve complete flexibility of connection between users, particularly where secrecy is mandatory. We would prefer to be able to speak to anyone quickly without making prior arrangements, because we often do not know in advance to whom we wish to speak under a completely flexible response doctrine--almost by definition.

Volume. We note a rapidly increasing volume of military communications in the future. Our present communication plant is highly limited, relative to future demands, in the volume of traffic that it is able to transmit. We feel that the military would be more effective if communications were not always treated as a scarce resource.

Cost. Because of our methods of assigning cost, it is difficult to determine exactly what we are paying for military communications. One recent estimate held that we were spending about a billion dollars per year, and the cost increases about 15 per cent annually. Future communications costs will not be cheap.

BEATING THE VULNERABILITY PROBLEM

Figure 3A shows the centralized network--one method of building communication networks used in the past. Here, a group of stations (nodes) wishing to inter-communicate with one another are all tied to a central switching node. This central switching node establishes paths of communication from any node to any other node. Such networks require only extremely simple equipment but suffer from the disadvantage of being highly vulnerable. Destruction of a single central switching node destroys communications for the entire network.

With the decentralized network, in Fig. 3B, we reduce the severity of this problem by connecting each of the original group of nodes to a few nearby stations which act as sub-switching centers. These sub-switching centers are connected. Destruction of a single node no longer destroys all communication. However, the network is still less than wholly survivable since the destruction of a few of the sub-switching centers destroys network communications.

In this briefing, we consider in detail the use of the distributed network--a network in which each station is connected only to its nearest neighbors (Fig. 3C). We show that such a network is potentially more survivable than that in the previous example. We then consider the practical problem of how to build switching gadgets that allow any station to talk to any other station, by passing through a large number of nodes in tandem.

SURVIVABILITY OF THE DISTRIBUTED NETWORK CONFIGURATION

A "Monte Carlo" computer simulation is used to determine how the distributed network fares under attack. A




Fig. 3—Beating the vulnerability problem

group of nodes, as in Fig. 4A, is subjected to "attack" in which a certain percentage of the nodes is destroyed on a "random" basis. The number of nodes still able to communicate with one another is computed for two variables: the probability of destruction of a single node, and the degree of connectivity of the network. All measurements shown are for an 18x18 array of 324 nodes. Figure 4B shows the definition of the degree of redundancy or the connectivity used.

If, for example, each node connects to only two other stations, it is defined as a network of redundancy level "one." If twice as many links are used as this minimum possible connectivity, the network is said to have redundancy level "two." Examples are shown for redundancy levels of 1, 1-1/2, 2, 3, 4, 6, and 8. Throughout the exercise, it is assumed that a "perfect switching" ability exists. This allows all nodes to maintain communication with others regardless of the number of tandem nodes traversed.

Figure 4C demonstrates the high payoff for redundancy. The vertical axis represents the fraction of nodes that have withstood the physical attack and are in electrical communications with one another. The horizontal axis is the specified single node probability for destruction of a raid directed against the network. Figure 4C shows that low-redundancy networks fall apart under light attack. However, if the redundancy level is increased to values of only three or four, extremely tough networks can be built. By "tough" we mean that if a station survives the direct attack, it has a very good chance of being in communication

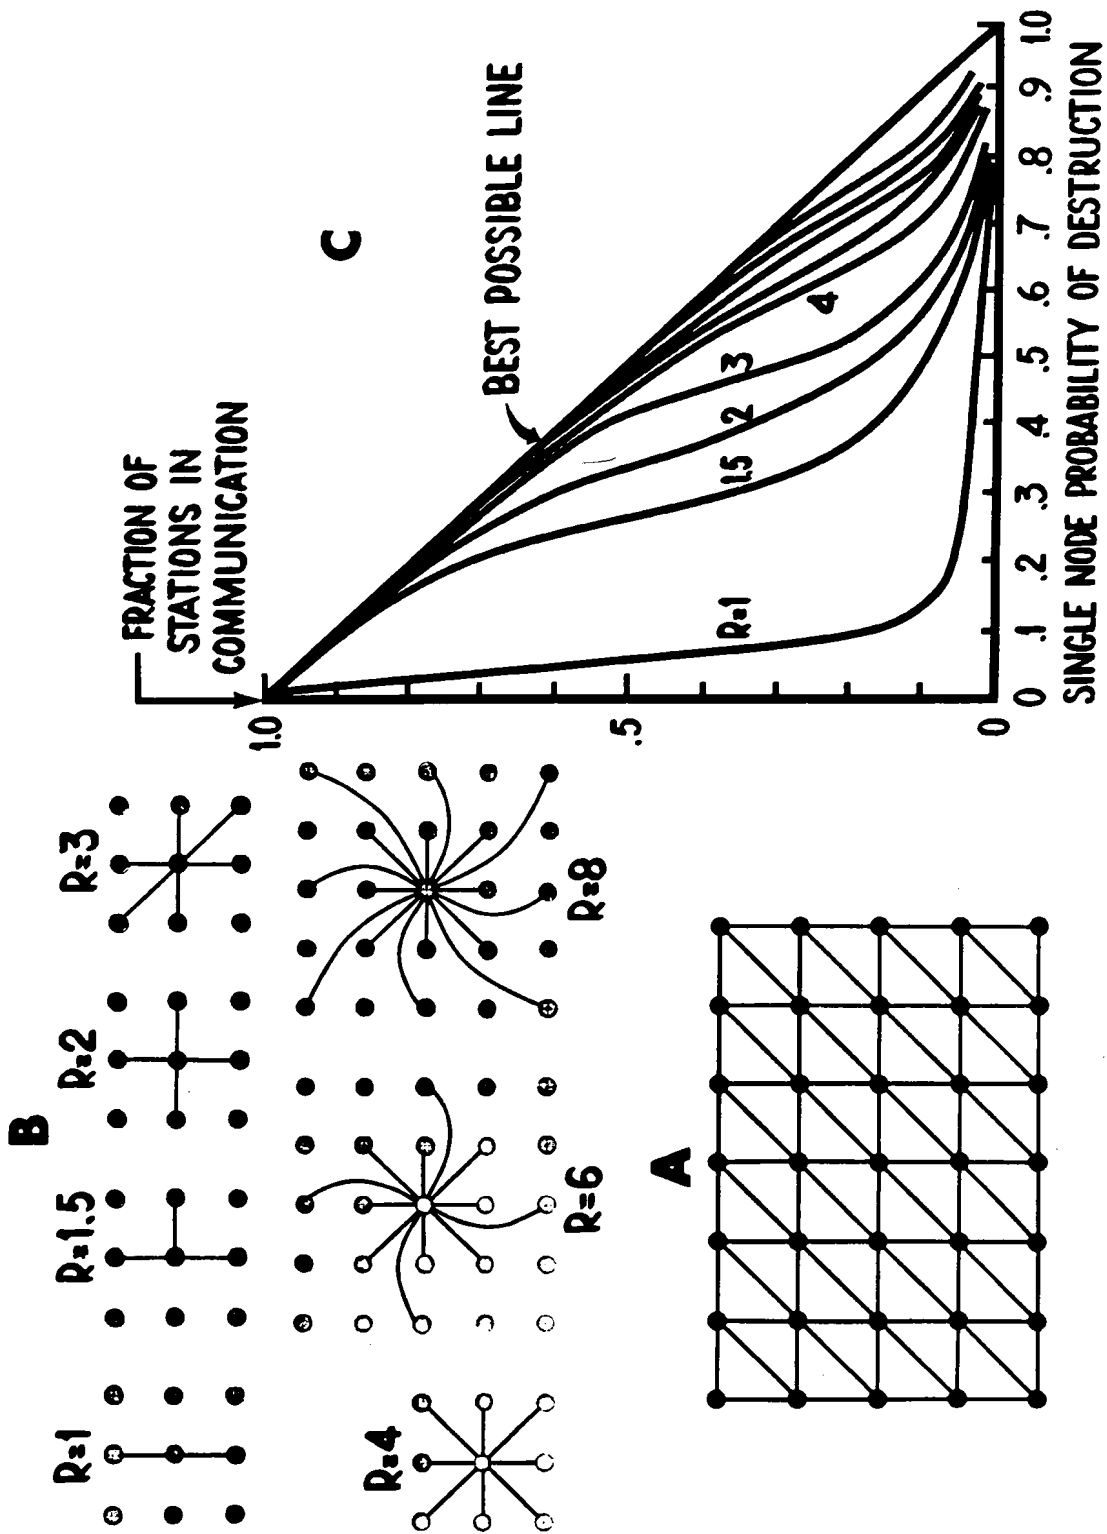


Fig. 4 — Survivability of a distributed network

with all the other surviving stations of the network. Two key points are to be noted from Fig. 4C. First, only a moderate amount of redundancy is required--on the order of three or four. More redundancy buys little. Second, such networks are able to maintain good post-attack communications even where perhaps half of the stations are destroyed.

Thus, it appears that the distributed network configuration is a good one to consider when we must build networks able to survive heavy attack.

PROBLEMS OF BUILDING A DISTRIBUTED NETWORK

To take advantage of the high survivability theoretically possible in distributed networks, we have to cross many tandem switching centers. This is difficult in present-day, analog-type networks because quality deteriorates in the tandem connections. We compound this problem if we wish to use a wide spectrum of emergency communication links while still trying to provide a high-quality, high-reliability system.

Our belief that we should be able to use unreliable links stems from the viewpoint that it should not make much difference whether "unreliability" comes from enemy attack, or from local electronic causes. In a future emergency, we might even wish to include television broadcasting links as well as conventional microwave links in our communications network. A low-altitude satellite is an example of an inherently unreliable link. When the satellite is overhead, we have a link; when the satellite is over the horizon, the link is out. From the system's point of view, this is just another unreliable link. In

the distributed network, the effective overall increase of system reliability depends upon the extensive use of unreliable elements.

To meet the reliability problem caused by many unreliable links in tandem, we must do two things: 1) use digital modulation, and 2) employ a "store-and-forward" information transfer.

The only way that we know how to transmit signals via a large number of tandem repeaters without having the signal irrevocably corrupted is to use digital modulation. As long as the signal is even slightly stronger than the noise, we may re-form the signal before its next transmission hop. Recently developed digital computer techniques will also allow powerful processing of the digital stream for error detection, secrecy encoding, and automated switching. Thus, the network of the future that we will describe will use all digital transmission. It will, however, as will be shown, handle conventional analog input signals as well. We call it an "all-digital system" because all the processing and transmission in the network is conducted in an all-digital manner.

In the network to be described, it may be necessary to use links of different data rates being fed from a variety of different users operating with widely different data rates. Therefore, we will design our entire system around a standard package of bits which we shall call a "message-block." This new standard would allow users to enter the network at whatever data rate they wish, and yet make efficient use of various data-rate channels forming the network. This differs sharply from analog




Fig. 5—All digital network composed of mixture of links

transmission circuits where the bandwidth is identical all along the channel.

The links connecting the switching nodes operate at a very high data rate. Rapid switching decisions will be necessary at each switching node to direct each package (message-block) of bits to the next node. These message-blocks must be relayed quickly from station to station so that they are formed, transmitted, unpacked, and delivered to the recipient so as to create the illusion that a direct copper wire exists between himself and the sender.

HOT-POTATO ROUTING

A message-block is like a letter (Fig. 6). In our system it consists of 866 bits of information "rubber stamped" with the names of the addressee and the sender, together with some additional "housekeeping" bits, forming a package of 1024 digital bits. In our analogy, the addressee is the address on the letter; the sender corresponds to the return address, and the content of the envelope is the data being transmitted. As each letter has a cancellation date, we create an analogous quantity which we shall call the "handover-number." Every time a message-block goes from Switching Node to Switching Node along its way, this handover-number is incremented, providing a measure of path length of each message-block in the network.

We call the routing scheme we shall use a "hot-potato" routing scheme because each Switching Node acts as if the message-block (letter) it is handling is a hot potato which must be passed on quickly before it "burns" the Node.

WHAT A MESSAGE BLOCK IS... AN ANALOGY TO A LETTER





Fig. 6—The message—block concept

Visualize each Switching Node containing an electronic postman at a blackboard. Each postman uses the return addresses and handover-numbers of passing traffic to determine the apparent shortest path over each of its links to any station in the network. If the best path is busy, the second one is taken immediately without waiting. If the second-best path is busy or destroyed, the third-best path is taken, and so forth. Sometimes, it will be necessary even to return the message-block back over the link on which it arrived. Simulation shows that when each station in the network uses this simple policy, an extremely effective over-all switching action takes place that allows high volumes of traffic to be transmitted from any node to any other node over paths surprisingly close to being the shortest possible. We find that a message-block rarely (less than about one in a hundred million) takes a path longer than the peripheral number of links of the network. This means that, with proper use of the powerful automatic error-detection and repeat-transmission techniques that exist now, we can maintain a very low over-all error rate from user to user even though each message-block may travel by different paths over noisy links.

THE PAYOFF FOR USING AN ADAPTIVE ROUTING DOCTRINE


The simple routing doctrine described above permits the network to adapt to changes in traffic loading, and link and node destruction, automatically seeking the best routes in the network without human intervention. To appreciate the power of self-adaptation of such a network to its environment, consider the following example shown in Fig. 7. In Fig. 7A, ABLE, BAKER, CHARLIE, DOG, EASY,

TIME T₁



A

TIME T₂



B

Fig. 7—Adaptability to change of user location

FOX are the locations as shown. In Fig. 7B, ABLE moves to a different location. The network, which has learned where ABLE was, soon realizes that ABLE is transmitting from a new location and now delivers messages to ABLE at its new location.

These adaptive routing doctrines were found to work efficiently when simulated. For example, in a network of 49 stations in which no node knew where the other nodes were at a time $T = 0$, it was found that within one second of scaled real-world time, all stations were efficiently routing traffic to all other stations. This network was also examined for performance while stations were destroyed and traffic in process. In summary, it is felt that the very simple doctrine described is able to provide highly sophisticated operation of a network of stations.

THE DISTRIBUTED ADAPTIVE MESSAGE-BLOCK NETWORK:

A SPECIFIC SYSTEM

To this point, we have described some of the underlying concepts of a new communication network, using the hot-potato routing doctrine to switch blocks of data from user to user. Let us consider a specific application of a communication network to provide common user services for a large number of military users in the future. In particular, we have examined a network of 400 Switching Nodes and 200 concentrating stations called Multiplexing Stations. The number of Switching Nodes was chosen to provide a highly survivable network, while the number of Multiplexing Stations provides service to at least 100,000 separate users operating a wide variety of input devices,

including start/stop teletype machines and good quality high-information-delta-modulation (HIDM) voice telephone. Each Multiplexing Station in Fig. 8 is connected to three switching centers.

Figure 9 represents one way in which the user visualizes his communications. Each user or subscriber has a terminal device, such as a telephone, connected to a Multiplexing Station. The Multiplexing Station is similar to a PBX or telephone central office. Signals from the subscriber are packaged into message-blocks and transmitted at high speed through the network of Switching Nodes terminating at the called subscriber. The entire network of Switching Nodes can be considered a black box comprising a distributed transmission plant that allows delivery of message-blocks from itself to the remote Multiplexing Station with minor regard for individual node failures caused by enemy action or reliability problems. How the message-blocks travel is of no importance to the end users.

Although the entire system being described is basically a store-and-forward-transmission system, it can in turn be viewed by the users as a black box containing a fixed time delay of about one-half second. Aside from this half-second delay, the illusion is maintained that a "copper-wire real-time" circuit exists between the two users.

How are analog voice signals sent over this all-digital system? The signal from the telephone microphone is converted into a digital signal for transmission. In our network, we shall use HIDM for voice signals. This is a simple digital modulation system that allows efficient digital voice signal reconversion representing a balance between




Fig. 8 — Interconnection between multiplexing stations and switching nodes




Fig. 9—User's view of the communication network

good (not excellent) quality, low data-rate, and economy of circuitry; 19,200 bits per second appears to provide highly intelligible speech. Push-buttons on the telephone change the connections of a synchronous counter circuit within the telephone to generate a repeating binary pattern as each push-button is sequentially depressed. Our digital telephone contains one additional circuit that is of interest. When no voice signal is present (between words or when the other party is speaking), the output signal is suppressed. No message-block need be sent unless a voice sample has been heard in the last one-twentieth of a second. Thus, a good quality telephone conversation represents a binary load to the network of only about 5000 bits per second per digital telephone. (Only one person speaks at a time, and about 50 per cent of the time of voice transmission is silence.) Unlike vocoder conversation systems, there is no loss of naturalness of the voice, yet the data rate taken to transmit the voice is comparable to a very high-quality vocoder.

Secrecy equipment is most economically provided as an integral part of switching apparatus. Both switching and cryptographic processing require digital equipment that can be time-shared.

Figure 10A defines end-to-end encryption in which clear text is locally encrypted and transmitted to an end secure area containing a crypto decoder. The use of an identical key in the crypto coder and decoder permit deciphering the transmitted stream. Where many tandem stations are required, link-by-link encryption can be used. Each message originator holds a key only to his local switching center. Each such




Fig. 10—Types of encryption systems

switching center has encryption equipment for its outgoing line to its adjacent Switching Node. This arrangement does not require the message originator and the end addressee to have identical keys. However, as the information is available in the clear at each switching center, one must have great confidence in the integrity of the personnel at each and every center. Fig. 10B shows a more secure arrangement. It uses two layers of cryptography: one for the text and another for the headings needed to direct traffic. It requires less concern about security at intermediate stations. Providing separate keys between message originators and end addressees still stands in the way of the full flexibility we would like. Therefore, we will design the system somewhat akin to the double encryption of Fig. 10C.

The steps of establishing a telephone call connection between two network users appear in Fig. 11. As in the conventional telephone system, the calling subscriber lifts his telephone off the hook indicating that service is desired, and immediately hears the dial tone (in our case, the dial tone is silence). The calling subscriber depresses push-buttons on his telephone. This transmits the called address via the calling Multiplexing Station to the called Multiplexing Station. If the called subscriber is connected, his telephone rings. Either a busy signal returns, or an agreement to accept the call. In our network, we add one feature. The calling Multiplexing Station transmits to the called Multiplexing Station something we call a "crypto start number." In each Multiplexing Station, we hold a key reserved for conversations between it and each of about




Fig. 11—Operations performed in setting up a call between network users

1000 Multiplexing Stations. The crypto start number corresponds to the number of calls made between the calling and called Multiplexing Station that particular day. If both Multiplexing Stations agree that this is in fact the 1003rd call, then the digital key base assigned for this combination of two stations is modified by a second key, and a new key base created to be used only for the call being established. Since the number of modifying operations performed to the crypto start number is large (and only a small part of the possible sequence generated is ever used) information from one telephone call cannot be used to help break a later call. The significance of this statement is that it is safe for the network to handle mixed classified and unclassified traffic. Information from one telephone call cannot be used to determine the key sequence in order to decipher the next telephone call. After this exchange, the two crypto generators are in synchronism, and the remainder of the call continues under cryptographic control. While we have shown only one level of cryptography, another level on a link-by-link basis is included in the Multiplexing Station and at each Switching Node to every other Switching Node.

Although we have discussed two levels of cryptography, with keys reserved on a per-call basis, our primary protection comes from still a third source, i.e., a form of cryptography commonly described in the open literature as "autokey." In the autokey procedure, the preceding text forms part of the key. Thus, it is necessary to correctly receive all the text of a message before being able to decode the next arriving symbol. Autokey, while being a

very powerful secrecy system, is not practical in a network containing sources of error. A single error destroys subsequent communication. But in the network we are describing, a very low error rate is expected--perhaps one error per year per user. Thus, we can afford to use the autokey principle for our end-to-end cryptographic protection. Let us recall that message-blocks will travel by different paths through the network. A fraction of a second after a call is established, the probability that all message-blocks will arrive by the same paths is very low. The significance of this statement is that only the end Multiplexing Station will ever see all the message-blocks transmitted after a fraction of a second. Persons tapping lines between stations can never be privy to all sequential message-blocks; they have a high probability of missing at least one. Thus, even if an enemy agent had all the keys to the network and listened from a central point, he would be unsuccessful because only the end Multiplexing Station receives all the message-blocks, and autokey requires that all message-blocks be correctly received.

COST

One skates on thin ice while describing the cost of a large system, never built, using new techniques. But this does not absolve the system designer of the obligation to discuss costs. This is ours (Fig. 12):

The described network might be built to serve 100,000 users at an annual cost of about \$60 million per year, including amortization. These are not conservative costs, as they assume that we will take advantage of the low-cost,

Purchasing major network

ITEM	QUANTITY	APPROXIMATE UNIT PRICE	TOTAL
SWITCHING NODES	400	\$ 150,000	\$ 60,000,000
MULTIPLEXING STATIONS	200	300,000	60,000,000
END TERMINAL DEVICES	100,000	200	20,000,000
LINKS (AIRPLANE MI)	120,000	400	48,000,000
			<u>\$ 188,000,000</u>

Cost summary

ITEM	TOTAL
RESEARCH AND DEVELOPMENT	\$ 23,700,000
ENGINEERING AND INSTALLATION	22,760,000
PURCHASE OF MAJOR NETWORK INVESTMENT	188,000,000
TOTAL INITIAL NETWORK INVESTMENT	\$ 235,000,000
BASIC ANNUAL NETWORK COST (10 YEAR BASIS)	\$ 60,000,000

Fig. 12—Approximate cost estimate

unreliable equipment that can be built and used in such a network. But, even though we will use unreliable equipment, we expect to achieve better reliability than we are accustomed to today. Systems such as this can be built only if one appreciates that unit reliability and systems reliability are two separate things, and that in the properly designed system, system reliability can be greater than unit reliability, unlike some systems that have been built. Building a lot of cheap, in lieu of a lesser amount of expensive, equipment is an art the military has not really practiced for a long time.

SUMMARY

Let us again turn to Fig. 2 and compare what we have proposed against our future communication problems.

- 1) Survivability. We propose the use of an all digital distributed network using adaptive routing able to withstand heavy destruction and operate effectively after attack.
- 2) Secrecy. We propose the use of cryptographic provisions built into the switching apparatus forming an integral part of the system. We propose the use of encryption schemes even more secure than the keys themselves.
- 3) Error/Rate/Quality. We propose the use of integral automatic error-detection repeat transmission means to allow the user-to-user error rates of less than 10^8 bits. This is several orders of magnitude better than now found in practice.
- 4) Delay Times. With the exception of the half-second delay in transmission time (approximately equivalent to what is noted on a high-altitude synchronous satellite transmission circuit) we permit automatic user-to-user transmission without switching delays.

- 5) Volume. We propose a network with orders-of-magnitude more data transmission capability than our present military communications network.
- 6) Flexibility. We propose instantaneous user-to-user communications among a large number of potential users without requiring pre-arrangement of circuit assignment or keys.
- 7) Cost. While this is an expensive system, the cost appears to be roughly only about six per cent of what we may now be paying for communications in the military. The price may be a bargain for what it buys. It might even be an economical system for certain commercial data-transmission applications.

Advantages and disadvantages of the distributed network concept are listed in Figs. 13 and 14, respectively.

- NETWORK USES ADAPTIVE LEARNING
- LESS VULNERABLE
- MEETS FUTURE MILITARY REQUIREMENTS
- HANDLES BROAD MIX OF INPUT DEVICES
- LARGE NUMBER OF VARIOUS TYPES OF USERS
- NO CUMULATIVE VOICE DISTORTION - LOW ERROR RATE
- CAN USE MIXTURE OF LOW COST, EVEN NOISY
UNRELIABLE LINKS
- SILENCE PERIODS SUPPRESSED
- ALL USERS CRYPTOGRAPHICALLY PROTECTED
- HIGHLY IMMUNE TO SOPHISTICATED SABOTAGE
- HOLDING KEYS DOES NOT ALLOW EAVESDROPPING
- COST COMPARABLE TO SOFT ANALOG NETWORKS

Fig. 13— Advantages of distributed network

- DIFFICULT TO EXPLAIN CONCEPT
- MUST UNDERSTAND COMPUTERS TO EVALUATE FEASIBILITY
- SYSTEM NEVER BEFORE BUILT
- IS EXPENSIVE TO SIMULATE
- LARGE NETWORK REQUIRED
- ONE - HALF SECOND DELAY IN VOICE TRANSMISSION
- PERFORMANCE VULNERABLE TO POOR DESIGN
- LOW COST WILL HINGE ON CAREFUL DESIGN
- EXISTING TERMINAL DEVICES SOMEWHAT EXPENSIVE
- ANALOG/DIGITAL CONVERSION REQUIRED
- NEED FOR SURVIVABLE, SECURE, ERROR - FREE,
COMMUNICATION NOT ALWAYS APPRECIATED
- CRYPTOGRAPHY NOT YET OFFICIALLY REVIEWED
- LARGE SYSTEMS COST A LOT OF MONEY

Fig. 14— Disadvantages of distributed network

APPENDIX

Further information may be found in the following publication series, entitled On Distributed Communications:

- I. Introduction to Distributed Communications Networks, Paul Baran, RM-3420-PR.

Introduces the system concept and outlines the requirements for and design considerations of the distributed digital data communications network. Considers especially the use of redundancy as a means of withstanding heavy enemy attacks. A general understanding of the proposal may be obtained by reading this volume and Vol. XI.

- II. Digital Simulation of Hot-Potato Routing in a Broadband Distributed Communications Network, Sharla P. Boehm and Paul Baran, RM-3103-PR.

Describes a computer simulation of the message routing scheme proposed. The basic routing doctrine permitted a network to suffer a large number of breaks, then reconstitute itself by rapidly relearning to make best use of the surviving links.

- III. Determination of Path-Lengths in a Distributed Network, J. W. Smith, RM-3578-PR.

Continues model simulation reported in Vol. II. The program was rewritten in a more powerful computer language allowing examination of larger networks. Modification of the routing doctrine by intermittently reducing the input data rate of local traffic reduced to a low level the number of message blocks taking excessively long paths. The level was so low that a deterministic equation was required in lieu of Monte Carlo to examine the now-rare event of a long message block path. The results of both the simulation and the equation agreed in the area of overlapping validity.

- IV. Priority, Precedence, and Overload, Paul Baran, RM-3638-PR.

The creation of dynamic or flexible priority and precedence structures within a communication system handling a mixture of traffic with different data rate, urgency, and importance levels is discussed. The goal chosen is optimum utilization of the communications resource within a seriously degraded and overloaded network.

- V. History, Alternative Approaches, and Comparisons, Paul Baran, RM-3097-PR.

A background paper acknowledging the efforts of people in many fields working toward the development of large communications systems where system reliability and survivability are mandatory. A consideration of terminology is designed to acquaint the reader with the diverse, sometimes conflicting, definitions used. The evolution of the distributed network is traced, and a number of earlier hardware proposals are outlined.

- VI. Mini-Cost Microwave, Paul Baran, RM-3762-PR.

The technical feasibility of constructing an extremely low-cost, all-digital, X- or K_u -band microwave relay system, operating at a multi-megabit per second data rate, is examined. The use of newly developed varactor multipliers permits the design of a miniature, all-solid-state microwave repeater powered by a thermoelectric converter burning L-P fuel.

- VII. Tentative Engineering Specifications and Preliminary Design for a High-Data-Rate Distributed Network Switching Node, Paul Baran, RM-3763-PR.

High-speed, or "hot-potato," store-and-forward message block relaying forms the heart of the proposed information transmission system. The Switching Nodes are the units in which the complex processing takes place. The node is described in sufficient engineering detail to

estimate the components required. Timing calculations, together with a projected implementation scheme, provide a strong foundation for the belief that the construction and use of the node is practical.

VIII. The Multiplexing Station, Paul Baran, RM-3764-PR.

A description of the Multiplexing Stations which connect subscribers to the Switching Nodes. The presentation is in engineering detail, demonstrating how the network will simultaneously process traffic from up to 1024 separate users sending a mixture of start-stop teletypewriter, digital voice, and other synchronous signals at various rates.

IX. Security, Secrecy, and Tamper-Free Considerations, Paul Baran, RM-3765-PR.

Considers the security aspects of a system of the type proposed, in which secrecy is of paramount importance. Describes the safeguards to be built into the network, and evaluates the premise that the existence of "spies" within the supposedly secure system must be anticipated. Security provisions are based on the belief that protection is best obtained by raising the "price" of espied information to a level which becomes excessive. The treatment of the subject is itself unclassified.

X. Cost Estimate, Paul Baran, RM-3766-PR.

A detailed cost estimate for the entire proposed system, based on an arbitrary network configuration of 400 Switching Nodes, servicing 100,000 simultaneous users via 200 Multiplexing Stations. Assuming a usable life of ten years, all costs, including operating costs, are estimated at about \$60,000,000 per year.

XI. Summary Overview, Paul Baran, RM-3767-PR.

Summarizes the system proposal, highlighting the more important features. Considers the particular advantages of the distributed network, and comments on disadvantages. An outline is given of the manner in which future research aimed at an actual implementation of the network might be conducted. Together with the introductory volume, it provides a general description of the entire system concept.