

SOME PERSPECTIVES ON NETWORKS—PAST, PRESENT AND FUTURE

PAUL BARAN
Cabledata Associates
Palo Alto, California

(INVITED PAPER)

This paper reviews the initial motivation for the development of distributed networks using packet switching -- the need for military survivability. All the other highly desirable characteristics for distributed packet network configuration, such as economy, fell out as by-products. Early opposition to packet switching by virtue of its ill fit to a world of analog transmission and electromechanical switching is discussed. The subject is brought to date with a rationale and suggested approach for building future networks and systems in a more robust and distributed form. The goal, in part, is to reduce society's present excessive vulnerability to disruption by natural and man-made disasters ranging from juvenile mischief to terrorist attacks to wars. As an example of the underlying thesis, the paper considers the desirability, feasibility and some possible consequences of building international electronic mail systems on a distributed basis for low cost person-to-person communications.

1. INTRODUCTION

1.1 The RAND papers on distributed communications

The work that is the cause of this presentation today was a dozen or so memoranda I helped prepare while at the RAND Corporation 15 years ago [1]. This work laid out the notion of packet switching. It explored the payoffs for the most effective use of redundancy to achieve reliability in distributed switching; the use of standardized packet format size and the tradeoffs in its size; the header information required; adaptive routing; high speed switching; error detection and repeat transmission; and creation of virtual circuits, etc.

The motivation for this work was not to improve the state of the art of computer communications networking. Rather, it was a specific military need, one that did not lend itself to solution by then available technology.

In the late 1950's and early 1960's the two major world powers had developed intercontinental ballistic missile systems capable of carrying multi-megaton thermonuclear warheads. The technology of offense far eclipsed that of defense and a most dangerous and potentially unstable situation resulted. Only the country that unleashed its weapons first could be assured that it would survive.

Each nation relied upon receiving warning that their opponent had indeed fired their missiles before responding. With highly vulnerable command and control communications the dangers of a dangerous age are multiplied. Failure of these control systems, whether by accident, intent or mere stupidity, could tempt the isolated parts of the network to panic into irrevocable actions. What was needed was a defense system that could withstand the onslaught of a Pearl Harbor type attack and retain the capability for returning the favor in kind shortly thereafter. As unpleasant as this all is, the mutual building of tough, survivable defensive capabilities -- "second strike capability" -- neutralized much of the gain for opening hostilities. It is this single change of defense postures that cooled tensions among the great powers and made detente possible.

It was this need for survivability of communications that required plodding through new ground not previously explored. What was being sought were some insights into how to go about building

tough, robust networks that would not come unstrung when stressed. All the side benefits beyond survivability and extreme reliability for distributed networks built with packet switching, such as very low cost, evolved solely as by-products of the data processing required to be present in the basic network design. Once you distribute processing or add packet-by-packet intelligence to switching nodes, a wide range of additional capabilities emerges. Thus, it is possible to build a system whose performance is more reliable than its parts. There is a minor cost in some redundancy, but with care, the amount of this needed redundancy can be shown to be modest. Extreme survivability was found to be achievable with redundancy factors of about three. This may sound wasteful, but as the reliability required for each element was reduced, the overall cost was potentially lower than the minimal redundancy case where every element must be extremely reliable. Thus, the aim to achieve more robust networks provided some insight into building less expensive networks for comparable system performance.

The study of distributed communications was aimed at acquiring an understanding of tradeoffs between robustness, reliability, channel capacity and costs. A network built to withstand physical attacks that destroy half the communications links but permit the residual system to still operate in a coherent, effective manner, achieves approximately the same performance as a system in a benign world built of communications links that can be inoperative about half the time. Such highly intermittent, noisy links are potentially inexpensive in comparison to links that must operate 99.99% of the time. The tradeoff curves between cost and system reliability suggest that the most reliable systems might be built of relatively unreliable and hence low cost elements, if it is system reliability at the lowest overall system cost that is at issue. However, by force of habit this line of systems thinking is hard to swallow and component reliability is still being sought at all costs.

1.2 Some untouched issues

In the last 15 years some of the design issues received much attention. Others have been ignored. Those issues dealing with topics that could be mathematically modeled received most attention, while the more complex and more important topics have gone relatively untouched. In packet switching, where packets travel by different paths, encounter different delay queues and are occasionally mangled, it is important to be assured that low overall error

rates, say 10^{-12} are possible. A little theory plus some straightforward simulation showed early in the game that the flow and delay problems, while real, were of second order importance in a distributed network. Elegant suboptimization here could not possibly produce as much savings as the same effort spent elsewhere.

Most of the funds available at the time for research were earmarked for support of education. Topics that could fit the rubric of education were relatively well supported. An undue number of Ph.D. candidates were strongly guided by their thesis advisors into producing displays of mathematical symbol manipulations seeking to squeeze a little more out of an assumed hypothetically perfect communications channel in a packet network. For a period of several years research funds were readily available at many major U.S. universities if the student's thesis appeared to have anything to do with information processing science and defense. By proper choice of project titles, military relevance could be implied while the resulting thesis could be read by a thesis advisor totally untrained in the real world of communications systems. The refining of minor issues that lent themselves to straightforward mathematical analysis became an academically acceptable field of exploration, while the more important system issues remained untouched.

I do not wish to deprecate any of this work. My concern is that there are styles in research that is funded just as there are in the length of women's dresses. When one research topic becomes too fashionable it tends to divert the support from other subject areas. Looking back at some of those topics briefly considered 15 years ago and then passed over, some topics appear to stand out as being more promising. For example,

- The concepts of dynamic bandwidth allocations as a function of content, user importance and allowable data aging.
- The real world issues in achieving super-reliability -- the system never failing -- by using higher degrees of connectivity and avoidance of statistical interdependency of failures.
- System design implications of universal multiple layer, low level cryptography.
- Overall systems economics tradeoffs present in use of very low cost, unreliable communications links.
- Design of future generation terminals (including the voice telephone) with packet formation and unpacking inherent in the basic terminal design.

I would be hesitant to recommend any of these topics to a Ph.D. candidate who wishes to acquire his degree painlessly. These topics are too conceptually ill-defined to lend themselves to neat mathematical formulation and they may be too "practical" or, euphemistically, "applied," to carry status in most schools of electrical engineering or computer science today. While greater pay dirt exists in the broader conceptually fuzzy issues, massaging of minor problems that lend themselves to mathematical analysis and whose answers are generally intuitively predictable to within 10 percent, has a greater outcome of assured success.

1.3 Classification

A question asked of me from time to time is "why wasn't the work on distributed communications classified?" I think there was an implicit feeling at RAND at the time to minimize the level of classification here. Only classification to

the minimum degree which local custom would comfortably permit was sought. An argument can be made that a more stable world situation exists if the adversary nation has the ability to survive attack and respond later. This reduces the proclivity for either party to take the opportunity of "winning" by going first. Hard as it is to imagine, it may be a non-zero sum game where it is to both parties' advantage to have a truly survivable command and control structure. The key objective of the military is preserving peace, not in fighting wars. A war is a failure to deter. Deterrence is not built by simply piling up more weapons, as the adversary's obvious counter is to build even more weapons in an uncontrolled arms race. But, a measure of stability does occur when each adversary has full respect for the damage the other side can inflict if attacked, thus making pre-emptive actions unrewarding. What at first may seem altruistic, in most instances is simply narrow, but long term best interest. It is doubtful if this early work would have any impact on command and control thinking at all if it were classified. Even making most of the early reports unclassified is not enough. Giving away ideas is not easy. Garbage and technical papers are two commodities of such low value in our culture that we must pay to have them hauled away. Fifteen years ago knowledge of digital processing by the communications community was limited. In particular, understanding of computer technology was absent in the governmental communications agencies that held operational responsibilities.

1.4 The experts' response

My most instructive experience with this emerging technology was the response by the communications community to it. I was in the Computer Science Department at RAND expounding building of non-hierarchical communications networks using computer technology. All signals were to be digital and processed on the fly in standardized blocks of 1024 bits. The notion was not taken seriously at all until the basic system concept was detailed. Then it was emotionally rejected as being thoroughly impractical and unworkable by some of the more respected people in the field of telecommunications at the time. This large and distinguished group included most of my friends at the telephone company and, at first, the Communications Department of RAND. On the other hand, I found many who thought that it all made sense. These people all had computer background -- I received particularly helpful support from the management of the Computer Science Department in the persons of the late John Williams, Willis Ware, Paul Armer and Keith Uncapher. The reason that it was necessary to write literally a two-inch thick pile of paper "On Distributed Communications" down to the transistor-by-transistor level was the response by communications experts not familiar with digital processing. They kicked, screamed, grumbled and worse. Their response tended to be emotional, often with anger, and rarely with humor. They were initially certain that the proponents did not understand how communications systems work. Part of their response can be appreciated by the realization that the telephone plant even at that time represented an investment in the tens of billions of dollars. When someone comes around and talks about building inexpensive communications networks using unreliable links and nodes, and of networks arranged willy nilly for extremely high survivability, it violates all their basic premises of network design. As this was in the days prior to the FCC allowing competitive communications networks, the cost of unamortized facilities in place was so much larger than that proposed sub-network to be built that the former could not be safely compromised at the expense of the latter.

Although I later came to have an extremely high

regard for the intelligence of my friends in AT&T, I must confess to being slightly dismayed in those early days, when after a long and detailed explanation of how communications networks could be built around the concept of all digital networks including links and high speed packet switching, I would be given a tutorial starting with how a carbon button telephone instrument worked.

At the Bell Telephone Laboratories I encountered a somewhat better reception. Of those competent in the analog technology only, about 90% were difficult to persuade. On the other hand, those proficient in digital processing technology had no difficulty in understanding what was being said. One such person was Edward E. David, then of Bell Labs. One evening, as a favor to Willis Ware, Ed graciously undertook to serve as translator to a senior person at AT&T. After a few of my sentences, Ed turned to his colleague to explain what I had said in telephones. And, in turn, when his colleague responded in telephones, his remarks were translated back to me. After this experience, I went off to learn to speak telephone jargon, including Western Electric parts numbers.

1.5 A conceptual block?

Computer trained people rarely had difficulty with the concept or details; those who lacked such background generally encountered difficulties in understanding. Why should intelligent and competent experts in communications systems have difficulty understanding what was so obvious to those with similar background, but who also had some computer experience? In part, it was the difficulty faced in mentally incorporating a new technology with distinctly different characteristics into a god-given larger framework of existing systems architecture. The present form of network required all additional elements to work without disruption to all that has gone before, and all before to be usable to the maximum extent with the improvement.

When one begins to talk to analog communications people about using unreliable trunks with switches that open before the next tandem trunk is connected, the conceptual lack of fit with the existing system creates a mental block. Truth and reality in the world of technology rely heavily on prior background training. The most effective response is then demonstration. How then to show that the notion is workable?

2. SYSTEM EVOLUTION

2.1 The ARPANET

The first workable packet network built was the ARPANET. This was an experiment that was initiated, managed, pushed and pulled and made successful by Lawrence Roberts, then with the Department of Defense, Advanced Research Projects Agency (ARPA). Other work was beginning to take place in this field independently elsewhere, particularly by Davies and others at the British NPL [2]. The ARPANET initially sought to demonstrate that the notion of switching packets on a high speed store and forward basis would, in fact, work. ARPA prepared a request for proposal. Many companies bid. Bolt, Beranek and Newman won the contract, and the BBN group became the first implementers. It was Robert's strong personal interest, effective leadership and control of a large research budget at ARPA that made packet switching feasible. For a few year period many universities in the U.S. found that they could obtain research dollars most readily by proposing work that related to the ARPANET. This mere scale of effort assured that some good work would result, and much did.

The initial ARPANET experiment was a network to

interconnect computers at relatively low data rates so that ordinary telephone lines could be used to form the links. To reduce transit delays, 50 Kb bandwidths were used. What emerged became institutionalized, and is now synonymous with packet switching. Although the ARPANET is highly economic for data transmission, it only scratches the range of basic capabilities of this direction of communications system architecture. A few decades may have to pass before we begin to see the eventual evolution toward robust, large, universal services, high speed packet based networks operating at the 50 megabit per second rates and higher.

2.2 Future network architecture

Today's low data rate packet networks, while highly effective, should not be regarded as more than transitional developments. Opportunity for decimal magnitudes of improvement still await capture. But, it may take moving away from present approaches to consider future demands in fresh terms. The early work cited was not intended to describe an "ultimate" system. It was only intended to serve as an existence proof to suggest that there are other ways of approaching network design than the then solely considered hierarchical configuration.

2.3 Why tougher systems?

What sort of changes would be helpful in next generation networks? Of course higher data rates and lower cost are desirable, but of all parameters, I would opt for greater emphasis on robustness. In a world where natural and man-made damage is not unknown, networks tend to be vulnerable to damage by even small groups seeking to disrupt a nation's infrastructure. In a world we regard as being at peace, bombings are commonplace in a dozen countries every day. Some are set to attract the attention of the news media for political aims. A few individuals bent on changing their own governments, or a small nation seeking to change international order may well find telecommunications and other utility structures to often be their most cost effective targets. We tend to build minimum redundancy systems designed only for a benign world. This makes mischief and civil disruption most cost effective against the highly capital intensive shared utilities: common carrier transportation, electrical power and telecommunications. Information transmission structures tend to be highly vulnerable to sophisticated attack. Our large systems are generally well designed to cope with natural disasters, but tend to underplay the probability of and response to man-caused system failures.

2.4 Approach

We have learned a little about the payoff for distributed processing. We understand the use of redundancy in error correcting digital processing and in the building of reliable organisms. Perhaps it is time to consider transferring some of these concepts to the architectural design of our information systems and even our social order system. In the larger view, could we have a more stable world if it were organized along the same lines as distributed information and control systems? Is the concept of the sovereign level of government each centralized in a single vulnerable node, such as a Washington or a Moscow, a wise choice in the very long term future? Change cannot take place rapidly here. Try as we might to build a stable and peaceful world, national governments with inviolable internal sovereignty may be expected to be around for at least the next hundred years. During this time we must live with nationalistic sovereign nations with the right to do whatever they please within their borders, even in a world where the weapons are extremely destructive and

portable. Most governments are run by reasonable people primarily concerned with betterment of their own citizens' lives. Nuclear proliferation is occurring and we must face the long range reality that in the family of nations there are occasions, however rare, in which the power is in the hand of an erratic or even crazy leader.

3. COMPUTER DATA BECOMING PEOPLE DATA

3.1 Hidden use of data networks

Let us move beyond the negative issues of the need and directions for building more robust systems for the moment and consider a promising development -- the development of sophisticated electronic message services. This has come about by the broadened use of data networks in general, and the ARPANET in particular.

In the next several years the early packet networks of the world will begin to interconnect. The resulting overall network will be more effective and useful for international communications than ever before possible. It could well become more important than voice telephone connection. Most people believe the networks are conveying digital data between computers. More likely the major portion of the data traffic is language text. It is messages exchanged between humans with relatively unsophisticated intermediate processing. Increasingly computer terminals, and the connected computers, are primarily facilitating the exchange of human-to-human messages.'

Those knowledgeable about these matters are disposed to look the other way and pretend that it is only "data" being sent to avoid a sticky regulatory battle. By accident of history, the tariff structure used by the telephone and telegraph administrations throughout the world for language text transmission is based upon telegram transmission costs. The tremendous decline in transmission cost in terms of bit-kilometers has not been directly reflected to the network users. The savings by new transmission technology are, throughout the world, generally used to cross subsidize other services felt to be more socially desirable. Rate disparities are commonplace. Bit-kilometer tariffs between countries are disproportionately expensive in comparison to flows within a nation's borders. The telephone user pays less for his bits than the data user. A three minute international telephone call using a 56 Kb digital PCM channel can transmit 3 x 60 x 56 000 or over 10 million bits or about 7000 pages of text. This can be an awful lot of record traffic for a few dollars, but the old tariffs in existence block such applications. Thus, those who transmit data at these bargain rates do not want to talk about their loophole. In the process, however, we are concealing a major evolutionary development -- a breakthrough in low international communications costs for alpha-numeric text.

3.2 Pro and anti-people technology

Some observers of mass communications consider communications technology in terms of being pro or anti unfettered person-to-person communications. For example, reusable magnetic tape video recorders are called "pro-people" because any individual can publish at a person-affordable cost (under \$1000) to control their own communications medium. Other forms of video recording that require an expensive shared reproduction facility costing in the millions of dollars, such as that proposed by CBS, the Phillips/MCA disks, or the RCA disks, are called "anti-people." The high cost of the master producing machines takes power from individuals. This power is concentrated into the hands of those who can afford to own an expensive duplicating

machine -- large organizations or governments. Amateur CB radio is "pro-people" communications. National network Tv broadcasting is regarded as "anti-people" as it permits control by institutions.

The proponents of "pro-people," or small scale technology (versus conventional "anti-people," or large scale technology) hold an implicit belief. That is that a society organized on a distributed, or a more person-to-person basis, can be more stable, less constraining and more adaptable to change than more impersonal, hierarchical, monolithic organizations.

3.3 Building an international people-to-people network inexpensively

Would a "pro-people" international communications network be socially useful? Is it feasible? Could use of in-place telephone circuits plus a little small scale technology make it possible to build low cost people-to-people written text systems crossing national boundaries? Let us consider this wild thought a bit, including the applications, the technology, and its likely source of evolution.

Conventional communications networks represent large scale investments and are generally national monopolies. One impact of packet switching is to reduce the cost of entry into the bit transmission and remote processing business. It lends itself to permitting a small amount of raw transmission bandwidth to be purchased and subdivided among many non-co-located users. We must consider five major components of our hypothetical person-to-person packet switched network: 1) the user's terminal; 2) a local distribution system; 3) an IMP, or packet switching node; 4) long distance transmission lines or satellites; and 5) a host computer, somewhere.

In fathoming the future it is necessary to reconsider our old view as to the economics of complex electronics. The LSI revolution is far along and its offspring the microprocessor, is here together with powerful support chips, plus very low cost memory. The hardware to build a workable switching node, or IMP, for a packet system even today is now affordable by the computer hobbyist. Very low cost terminals will soon be here. These can be bootlegged by acoustic coupling to the existing telephone network. Since the long distance telephone line cost is tolerable if we buy a voice circuit and divide it up among the potential users all the hardware ingredients will soon exist for a people-affordable network. We shall consider its application, but first let us ask "who might build such a system?"

3.4 Who may be the builders?

The computer hobby is a phenomenon less than two years old. Computer clubs abound in the U.S. At this time 250 stores in the U.S. specialize in selling microcomputers and low cost peripherals. The computer hobbyist group appears to be comprised of those who work in the computer field by day with their avocation following their vocational interest at night. These are highly competent and imaginative individuals, and are not totally impoverished. Each issue of the present hobby literature is probably richer in economically useful ideas than can be found in journals of the professional computer literature. Hobbyist software tends to be better written than commercial programs. A powerful Basic language jammed onto a single ROM chip or two is a common achievement. Unlike the commercial sector, hobbyists actively trade their software. It is the hobbyists who developed the most useful low cost magnetic media data exchange standard -- "the Kansas City Standard" -- for cassettes. By this I mean that one can build more cost effective

equipments using the hobbyist standard than commercial data processing standards. The computer hobbyists are probably having more profound effects on technology development than even the radio amateur operators did in the early days of radio -- when the spectrum above 200 meters was regarded as useless, only to be developed by the amateurs. We may well be entering a new era where small groups of individuals using their own funds can be the builders of the digital communications and processing systems of tomorrow. Of course, large companies, well funded government laboratories and other institutions will always be around to take claim for the commercialized versions of any innovations as has always been the case. We may have to watch this phenomenon carefully to see how it occurs. What we do know is that the basic elements needed are now down to the pocketbook level of the hobbyists. The major stumbling block that might hold back this person-to-person channel of development is the rigorous grip that the communications industry holds everywhere in the world by virtue of government monopoly regulation. This nationalization or heavy administrative control came from the past historic control of communications by the military agencies of government starting in the day of the optical semaphore telegraph. Perhaps it is time to reconsider the use and control of communications channels and extend personal freedom by allowing the public access rights to its own communications resource.

3.5 What will it be used for?

How is the new network going to be used? If the American scene is correct, we might see games and messages. The messages could well replace the postal mail. Of course, you may argue, we already have a good international mail system. Yes, a letter can generally be exchanged between any two people in much of the developed world in anywhere from a few days to perhaps two weeks. With a good electronic message system you can interact several times a day with virtually instantaneous transmission. A different form of communications results when messages are exchanged in hours in lieu of days at a time. With proper technology, the cost of electronic messages can be cheaper than postal rates.

3.6 Impact

Any time you change things by a decimal order of magnitude, you change quality as well as quantity. The automobile = 100 kilometers/hour; the airplane = 1000 kilometers/hour. Neglecting secondary characteristic differences, such as the airplane's ability to travel over water, the simple factor of 10 can provide revolutionary change in the structure of society. With a drastic increase in the speed of international person-to-person mail, a revolutionary change can take place; not simply in the speed and cost of transmitting written text, but also in the way letters are written and how they are used. With a good text editing and retrieval system there is less need for typists to retype letters before transmission, thus serving to both shorten the length of the messages and to accelerate turnaround time.

By its nature, data is a universal language. Certainly the Arabic numeral set is universal. Even the character sets of the world's written languages tend to map onto a much smaller set than the number of languages themselves. A single Roman language (with a handful of local character exceptions) combines the separate languages of the entire Western world. Many more people will be able to communicate across the written text language barriers than can the oral language barrier. (Many more read, albeit slowly and painfully, the language of others than can hear it and understand it in real time.)

3.7 Subversive objective

By undercutting the existing tariff structures for record traffic by bypassing government controls, we may be able to do more to create effective and international cooperation at the person-to-person level than all the grandiose institutions that have been tried in the past, including the United Nations, international broadcasting, grandstand games such as the Soyuz-Apollo linkup, the Peace Corps and large scale diplomatic agencies taking turns entertaining one another at the expense of the taxpayers of the hosting nations.

4. CONCLUSIONS

I have discussed two seemingly divergent topics in the latter section of this paper: building tougher network structures and encouraging the development of person-to-person communications around the world bypassing rigid national control structures. They are really related issues as both are directed towards a more stable world order.

Very long term world stability can occur only if the vast economic and social disparities between nations can be reduced. This will take improved international communications, not only for governments and businesses, but also for people. Still, it might take a hundred years or more. In the interim, civil comfort may require understanding of how to build more robust information systems for societies' infrastructure.

Information is different than all other resources. If I sold you a physical commodity, I would no longer have it. Sharing information benefits all. After I give information to you, I still have it to use for myself. There are greater chances for devising a non-zero sum game for information activities in international intercourse than ever possible in physical resource trading. Improved information flows may well be more important than physical resources. In the highly developed countries of the world today more than half of the GNP is developed via the transmission and processing of information. Thus, the scope for economic impact can be great. God has sprinkled mineral resources very unevenly, but brains surprisingly uniformly throughout the world, considering the different needs to survive in different cultures. Altruism has its bounds. While it is nice to have a long range view of a world with reduced economic tensions (never forget that both Adam Smith and Karl Marx were only economists), there are heavy prices to be paid for any such transition.

To give an example close to home; with the future almost zero cost computer communications capability, it will be possible to program a computer from anywhere. Can you imagine a world where most of the software will be written by the lowest cost labor available, irrespective of what country it may be found? Clearly, it is a mixed blessing, with most of us who read this having very much to lose in the short term while the rewards may be long in coming and can be felt only by others. Maybe this is what statesmanship is all about.

REFERENCES

- [1] Paul Baran, On distributed communications: I. Introduction to distributed communications networks, The RAND Corporation, Santa Monica, California, Memorandum RM-3420-PR, August 1964.
- Sharla P. Boehm and Paul Baran, On distributed communications: II. Digital simulation of hot-potato routing in a broadband distributed communications network, RAND Corp. Memorandum RM-3103-PR, August 1964.

J.W. Smith, On distributed communications: III. Determination of path-lengths in a distributed network, RAND Corp. Memorandum RM-3578-PR, August 1964.

Paul Baran, On distributed communications: IV. Priority, precedence, and overload, RAND Corp. Memorandum RM-3638-PR, August 1964.

Paul Baran, On distributed communications: V. History, alternative approaches, and comparisons, RAND Corp. Memorandum RM-3097-PR, August 1964.

Paul Baran, On distributed communications: VI. Mini-cost microwave, RAND Corp. Memorandum RM-3762-PR, August 1964.

Paul Baran, On distributed communications: VII. Tentative engineering specifications and preliminary design for a high-data-rate distributed network switching node, RAND Corp. Memorandum RM-3763-PR, August 1964.

Paul Baran, On distributed communications: VIII. The multiplexing station, RAND Corp. Memorandum RM-3764-PR, August 1964.

Paul Baran, On distributed communications: IX. Security, secrecy, and tamper-free considerations, RAND Corp. Memorandum RM-3765-PR, August 1964.

Paul Baran, On distributed communications: X. Cost estimate, RAND Corp. Memorandum RM-3766-PR, August 1964.

Paul Baran, On distributed communications: XI. Summary overview, RAND Corp. Memorandum RM-3767-PR, August 1964.

- [2] D.W. Davies, K.A. Bartlett, R.A. Scantlebury, and P.T. Wilkinson, A digital communications network for computers giving rapid response at remote terminals, ACM Symposium, Gatlinburg, October 1967.