THE PATH TO TODAY

Thursday, August 17, 1989
9:15 - 10:30

ARPANET: History, Technology, and Pioneers

Robert Taylor, Chair
Director, Systems Research Center, Digital Equipment Corp.
The ARPANET: Getting Started,

Paul Baran, Chairman of the Board, InterFax, Inc.
Early History "On Distributed Commmunications"

Lawrence Roberts, Chairman of the Board & CEO, NetExpress, Inc.
Conception and Birth of ARPANET

Frank Heart, Vice President, Bolt, Beranek and Newman
Early ARPANET Issues

J.C.R. Licklighter

EARLY HISTORY
"ON DISTRIBUTED COMMUNICATIONS"

UCLA SYMPOSIUM ON
ADVANCED COMMUNICATIONS
TECHNOLOGIES

PAUL BARAN

AUGUST 17, 1989

3001

BACKGROUND

1. IN THE LATE 50'S AND 60'S THE ICBM ERA BEGAN.

2. US/USSR TENSION STRESSED BY THE PARANOIC ADVANTAGE OF A FIRST STRIKE.

3. STABILITY REQUIRED CAPABILITY SURVIVING FIRST STRIKE.

4. BUT, A SURVIVABLE RETALIATORY CAPABILITY REQUIRES SURVIVABLE COMMUNICATIONS-- WHICH WE DIDN'T KNOW HOW TO BUILD...

THREE TYPE OF NETWORKS

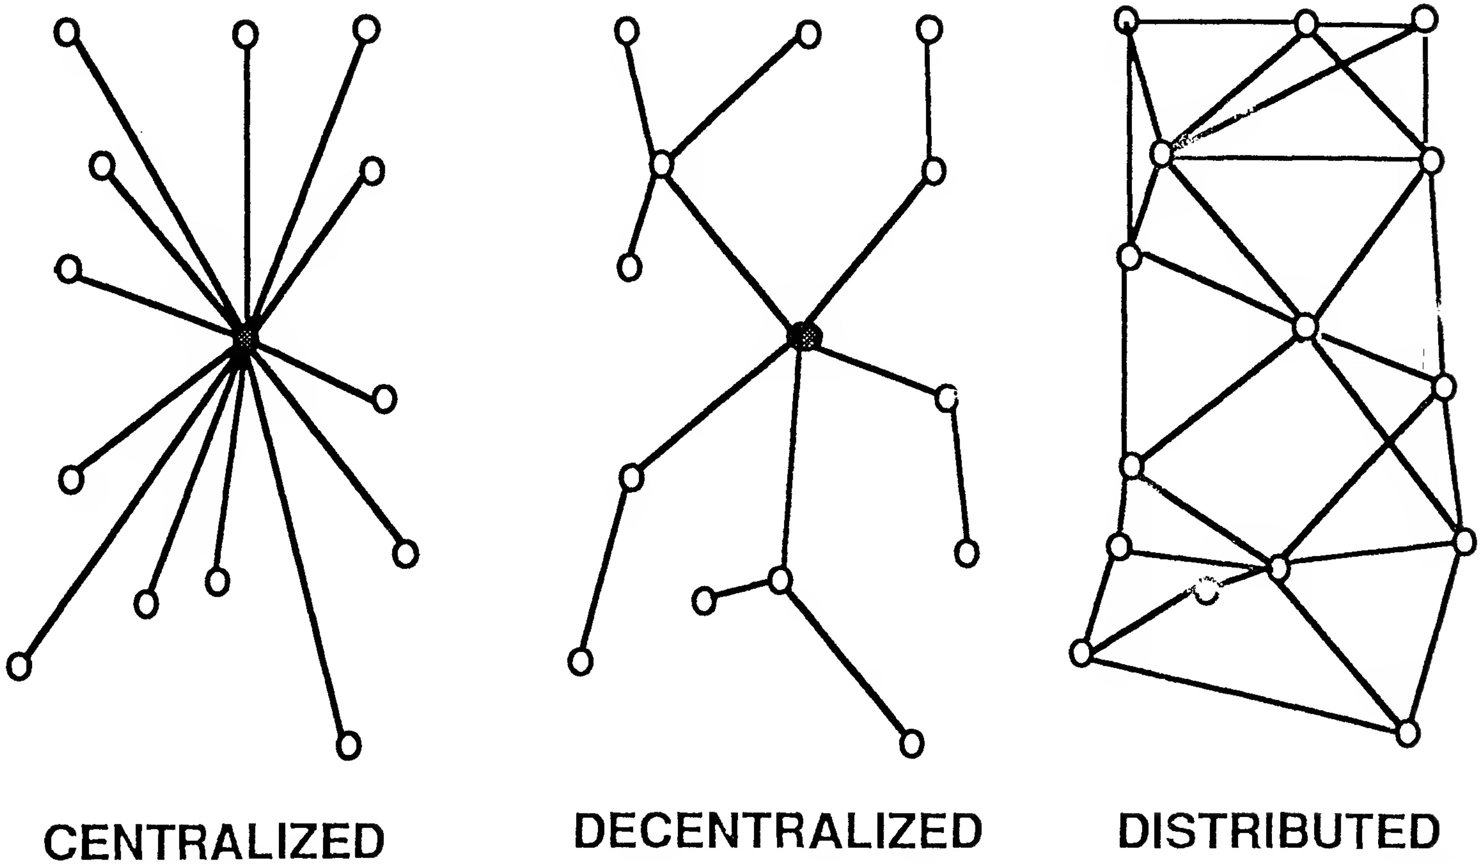

3011

THE KICKER: WE DIDN'T KNOW HOW TO BUILD DISTRIBUTED NETWORKS

1. NEEDED "PERFECT SWITCHING".
2. SIGNALS HAD TO TRAVERSE MANY NODES.
3. NO CENTRAL NODE ALLOWED.

VERSION 1: AM BROADCAST STATION NETWORK.

1. OBJECTIVE: "MINIMUM ESSENTIAL COMMUNICATIONS."

2. NETWORK OF EXISTING AM BROADCAST STATIONS TRANSMITTING TELETYPE SIGNALS ON THE GROUND WAVE. (SKYWAVES DISAPPEAR UNDER ATMOSPHERIC NUCLEAR EXPLOSIONS.)

3. USED A SIMPLE FLOODING ALGORITHM.

4. SERIES OF BRIEFINGS THROUGH GOVERNMENT FOUND THAT "MINIMUM ESSENTIAL COMMUNICATIONS" WAS AN UNDEFINABLE CONCEPT.

5. EVERYBODY ALWAYS SEEMED TO WANT MORE CAPACITY

6. BACK TO THE DRAWING BOARD.

THE 1962 SYSTEM

1. OBJECTIVE: ALLOW EACH DEFENSE USER- ALL THE SURVIVABLE COMMUNICATIONS DESIRED.

2. USE THE HIGHEST STATE OF THE ART DIGITAL LINKS -- 1.54 MB/SEC. (T-1).

3. USE A MIXTURE OF SATELLITE, TELEPHONE, TV TERRESTRIAL COAX, ETC.

4. A COMMON SHARED NETWORK TO SERVE ALL USERS AND ALL MODALITIES. TWX, DATA, DIGITAL VOICE, ETC.

5. ALLOW A LARGE NUMBER OF USERS PER NODE.

TOWARD A SOLUTION: THE OBJECTIVE

1. TRAFFIC HAD TO ROUTE IT SELF THROUGH THE NETWORK.
2. HAD TO WITHSTAND NODE TRAFFIC OVERLOADS AND FAILURES.
3. HAD TO BE SELF ADAPTIVE.
4. HAD TO BE ERROR-FREE ON AN END-TO-END BASIS, NOTWITHSTANDING ANTICIPATED LINK ERRORS.

THE APPROACH

1. ALL SIGNALS HAD TO BE DIGITAL.

2. ALL TRAFFIC BROKEN INTO A STANDARD FORMAT "MESSAGE BLOCKS" (OF 1024 BITS.)

3. EACH "MESSAGE BLOCK" CONTAINED HOUSEKEEPING DATA TO PERMIT CARRYING INFORMATIONNECESSARY FOR ROUTING:

   START AND STOP FLAGS
   "TO" ADDRESS
   "FROM" ADDRESS
   "CRC ERROR DETECTION
   TIMES RELAYED
   SEQUENCE ORDER.

4. COMPLETE INTELLIGENCE TO ROUTE EACH "MESSAGE BLOCK" TO BE INTEGRAL TO ALLOW ADAPTIVE SELF-ROUTING .

"IT AIN'T GONNA WORK BECAUSE..."

1. SERIES OF OVER 30 BRIEFINGS GIVEN THROUGHOUT GOVERNMENT, DEFENSE CONTRACTORS-- UNIVERSITIES, TELEPHONE COMPANIES, ETC.

2. 50% SAID IT WOULD WORK AND BECAME SUPPORTERS.

[illegible]. 50% SAID THAT IT CAN'T POSSIBLY WORK .

3. REASONS FOR NEGATIVE RESPONSES WERE
   A. DIDN'T UNDERSTAND DIGITAL PROCESSING, OR
   B. HAD A GENUINE CONCERN ABOUT ONE PORTION OF THE SYSTEM OR ANOTHER.

4. HAD TO ADDRESS ALL QUESTIONS RAISED BY MORE DETAILED ANALYSES. EACH QUESTION LED TO MORE QUESTIONS.

5. PRODUCED A THICK STACK OF REPORTS ON EACH ASPECT OF THE SYSTEM, INCLUDING SIMULATIONS.

CONCEPTS DESCRIBED IN THE REPORTS

IT IS FUN TO READ STUFF YOU WROTE 25 YEARS AGO.
YOU CAN FIND DESCRIPTIONS OF:

1. THE USE OF FIXED LENGTH PACKET COMBINING FULL ROUTING INFORMATION AND ERROR CORRECTION, AND WHY YOU WANT TO DO IT THIS WAY INSTEAD OF AN ALTERNATIVE APPROACH.

2. PACKET SWITCHING "MESSAGE BLOCKS" USING DYNAMIC ROUTING.

3. VIRTUAL NETWORK CONCEPT.

4. FLOW CONTROL

5. STATISTICAL MULTIPLEXING ECONOMY FOR MANY DIGITAL USERS BY SUPPRESSING SILENCE STATES.

6. ACHIEVING HIGH RELIABILITY USING UNRELIABLE LINKS.

7. SOME OTHER GOOD STUFF.

CONCEPTUAL BARRIERS

1. CONCEPT WAS HIGHLY CONTROVERSIAL.

2. HARDEST PEOPLE TO CONVINCE-- COMPETENT ANALOG TRANSMISSION ENGINEERS.

3. EASIEST TO CONVINCE-- DIGITAL TYPES.

4. ENCOUNTERED NIH TURF PROBLEMS:

5. RAND: "WHAT IN HELL IS SOMEONE IN THE COMPUTER SCIENCE DEPARTMENT SCREWING AROUND WITH COMMUNICATIONS. THAT'S OUR DEPARTMENT..."

6. AT&T: "IT CAN'T POSSIBLY WORK. AND, IF IT DID, DAMNED IF WE ARE GOING TO SET UP ANY COMPETITOR TO OURSELVES..."

CHRONOLOGY

1. 1960 LOW DATA RATE DISTRIBUTED NETWORK.

2. 1962 HIGH DATA RATE DISTRIBUTED NETWORK DEFINED.

3. 1964 PAPERS PUBLISHED OUTSIDE THE RAND/ UNIVERSITY/ DEFENCE AGENCY COMMUNITY.

4. 1965 PROJECT RAND FORMAL RECOMMENDATION TO THE AIR FORCE TO PROCEED.

5. 1965-6 AIR FORCE/MITER COMMITTEE REVIEWS AND MAKES RECOMMENDATION TO IMPLEMENT.

6. 1966 DOD GENERAL COUNSEL ASSIGNED ACTIVITY TO DCA. (AT THAT TIME DCA HAD NO DIGITAL CAPABILITY.)

7. PB REQUESTED DOD NOT TO PROCEED UNTIL A MORE SUITABLE AGENCY FOUND.

Conception and Birth of ARPANET

Lawrence Roberts
Chairman of the Board & CEO,
NetExpress, Inc.

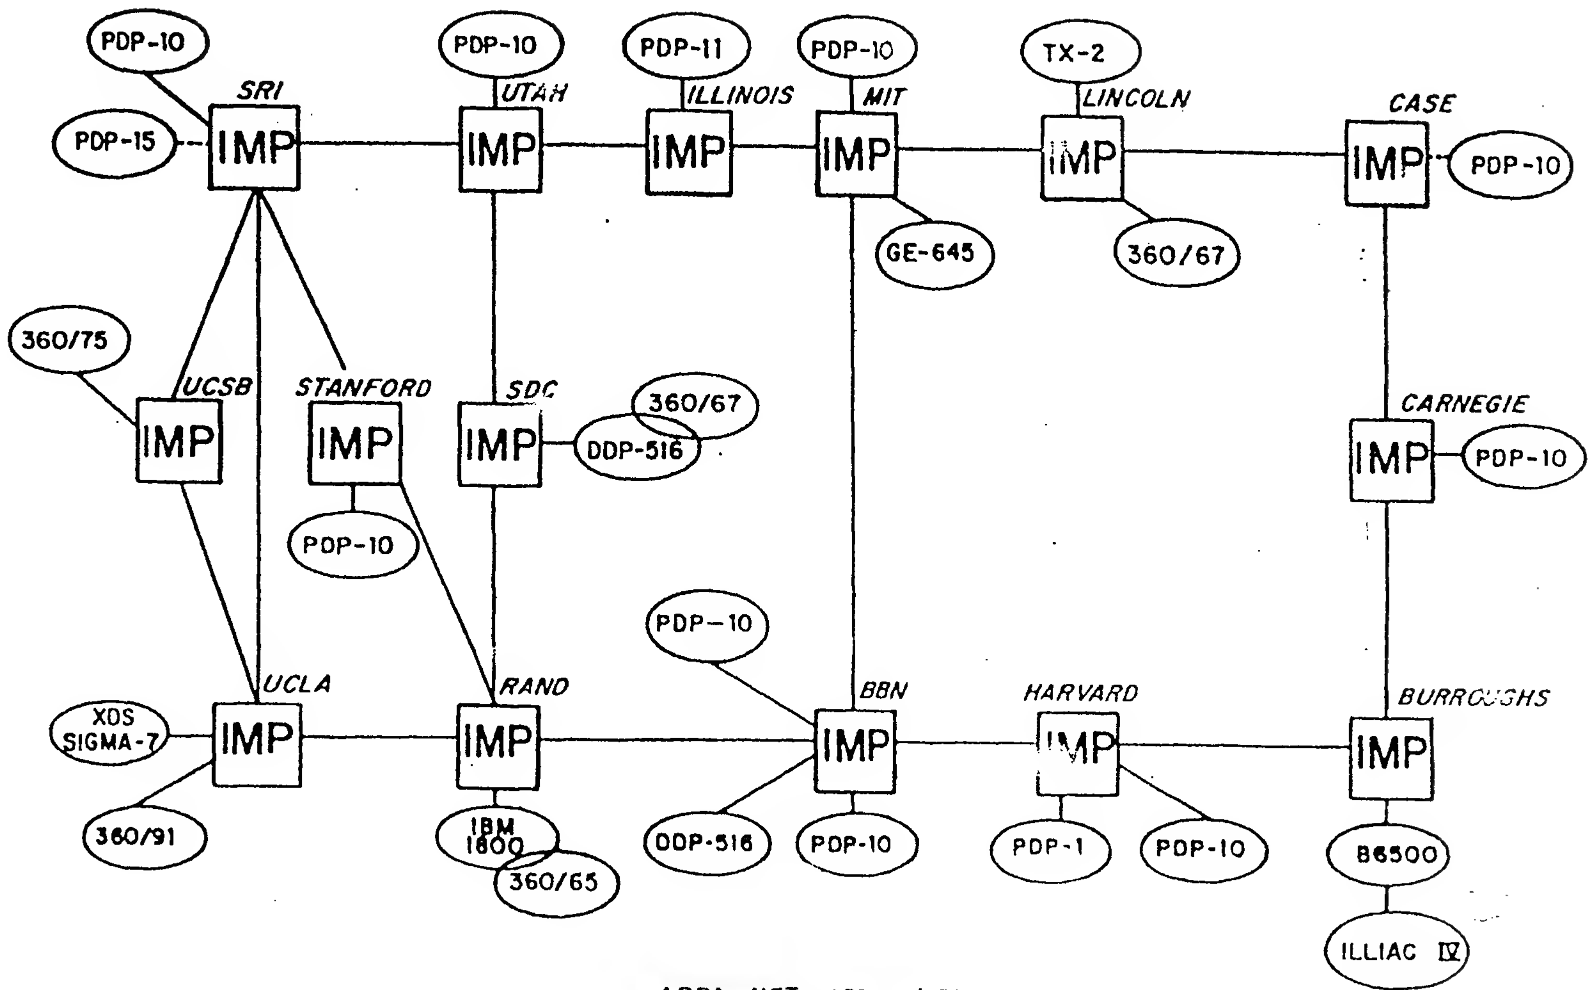

ARPA NET, APRIL 1971

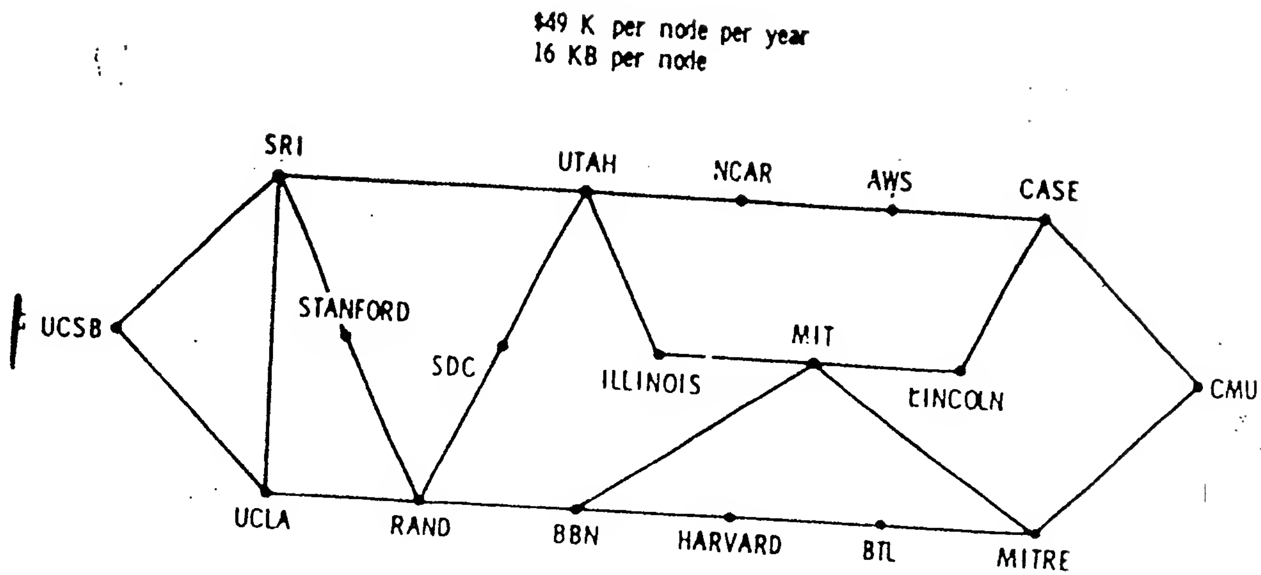

Figure 1—ARPA network initial topology

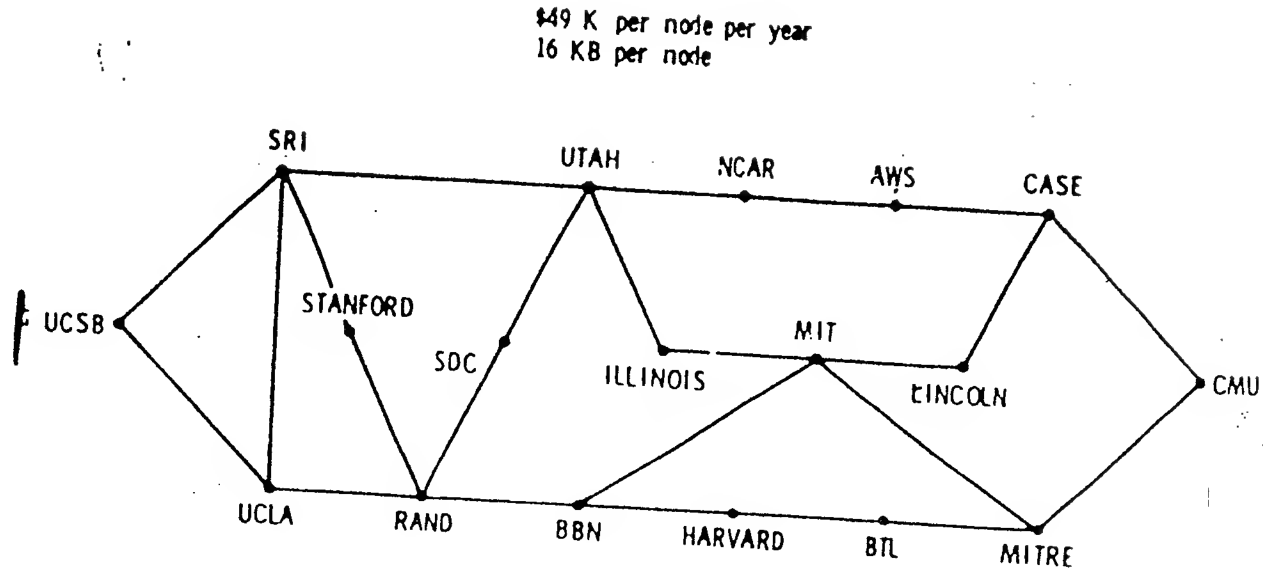

Figure 1—ARPA network initial topology

Early ARPANET Issues:

Frank Heart
Vice President,
Bolt, Beranek and Newman

An Emphasis on Reliability—I

- Ruggedization
- Similar IMPs
- Protected memory
- Power fail interrupt—auto restart
- Watchdog timer—auto reload
- No switches/buttons

subdivided them into 128 byte packets, and added a header specifying destination and source addresses; then, based on a dynamically updated routing table the minicomputer sent the packet over whichever free line was currently the fastest route toward the destination. Upon receiving a packet, the next minicomputer would acknowledge it and repeat the routing process independently. Thus, one important characteristic of the ARPANET was its completely distributed dynamic routing algorithm on a packet-by-packet basis, based on a continuous evaluation within the network of the least delay paths, considering both line availability and queue lengths.

The technical and operational success of the ARPANET quickly demonstrated to a generally skeptical world that packet switching could be organized to provide an efficient and highly responsive interactive data communications facility. Fears that packets would loop forever and that very large buffer pools would be required were quickly allayed. Since the ARPANET was a public project connecting many major universities and research institutions, the implementation and performance details were widely published. (10, 11, 12, 13, 14). The work of Leonard Kleinrock and his associates at UCLA on the theory and measurement of the ARPANET has been of particular importance in providing a firm theoretical and practical understanding about the performance of packet networks.

The ARPANET was first demonstrated publicly at the first International Conference on Computer Communications (ICCC) in Washington, D.C. in October 1972. Robert Kahn of BBN organized the demonstration installing a complete ARPANET node at the conference hotel, with about 40 active terminals permitting access to dozens of computers all over the U. S. This public demonstration was, for many (if not most) of the ICCC attendees, proof that packet switching really worked. At this time, it was difficult for many experienced professionals to accept the fact that a collection of computers, wideband circuits, and minicomputer switching nodes (equipment totaling well over 100 pieces) could all function together reliably. The ARPANET demonstration lasted for three days and clearly displayed its reliable operation in public. The network provided highly reliable service to thousands of attendees during the entire duration of the conference.

Industry Reception

From the first time I distributed a description of packet switching outside the computer research community (the 1967 paper) until about 1975, the communication industry's reaction was generally negative since this was such a radically different approach. In some of the initial technical speeches I gave, communications professionals reacted with considerable anger and hostility, usually saying I did not know what I was talking about since I did not know all their jargon vocabulary. The most common technical flaw suggested (before the ARPANET was built) was that the buffers would quickly and catastrophically run out. After the ARPANET was operating successfully, their pitch changed to be that packet switching would never be economic without the government subsidy. Paul Baran reported the same reaction to his papers when he presented them; this reaction was the major reason his proposals never moved the military. Donald Davies reported a somewhat less angry response from the British Post Office, more one of mild interest but no serious consideration.

I learned a major lesson from that experience; people hate to change the basic postulates upon which considerable knowledge has been built. In the case of packet switching, the first postulate to change was the statistical nature of the traffic - data versus voice. The second was that computing was expensive. Some people find it is impossible to consider such a major jolt to their memory organization - they avoid it with putdowns if possible, if not with anger. Other people are more willing to reconsider, but for everyone it requires considerable effort. Those of us proposing packet switching all came from the computing field and did not need to change lots of prior concepts and knowledge. Many of those in the communications field still have not accepted packet switching. (The fight is heating up again as voice packet switching starts to be considered.)

ARPANET Growth

As soon as the first four nodes were brought up and tested in December 1969 the network grew very rapidly. One year later, in December 1970, the network had grown to 10 nodes and 19 host computers. By April 1971, there were 15 nodes with 23 host computers. The topology was as shown in Fig. 1. By this time, it was clear that connecting terminals directly to the network through a PAD-type device was important. Such a device was designed and built in 1970/1971, and the first Terminal Interface Procesor (TIP) was added to the network in Aug 1971. This permitted users with no computer to select a computer from all those around the country. In many cases having the user attach his terminal to a TIP and access even his own host(s) through the network was found to be more reliable. This was the start of a trend which today is almost the rule, workstations should attach to a network, not a computer. By January 1973, the network had grown to 35 nodes of which 15 were TIP's and was connected to 38 host computers. Fig. 2 shows a map of the network at this point. Network traffic had grown rapidly in 1972, from 100,000 packets/day to 1,000,000 packets/day, exceeding the original estimates I made in 1967. Also in 1973, the first satellite link was added to the network with a TIP in Hawaii; later in the year a pair of TIP's were added in Europe. In September 1973, the network was up to 40 nodes, 45 host computers with internode traffic of 2,900,000 packets/day, had clearly reached a stage of operational stability, heavy usage, and was by any measure a major success. It was at this point that I left ARPA to spread the technology to the commerical world.

The ARPANET has continued to grow since 1973, with 111 host computers in 1977 and more than 400 hosts in all the interconnected networks by 1983. The network has become a utility for both ARPA and the Department of Defense as a whole. Research has continued into internetting and packet radio but little change has occurred in the basic network.

Early Technical Issues

- Emphasis on autonomy from hosts and other use
- One host per IMP ——> multiple hosts per IMP
- Local host ——> distant host ——> very distant host
- Adaptive routing—a favorite problem
- Propagating program releases through network
- Intra-site communication—a surprise
- A tiny program—6000 words
- Synchronizer bugs

An Emphasis on Reliability—II

- Retransmission with powerful check sum
- IMPs good at checking phone lines
- Remote program-controlled crosspatching
- Remote debugging
- Extensive introspection and reporting
- 24-hour monitoring
- TIP auto dial-in

IMP Development Group Photo

there were 11 reports written for the Air Force in Rand Memorandum group, of which a couple were class... ...tunately the others were very sparcely published in the scientific press. thus their impact on the actual development of packet switching was mainly supportive, not sparking its development - that happened independently at Rand, NPL and ARPA.

The First Network Experiment

Convinced that computer networking was important, the first task was to set up a test environment to determine where the problems were. Thus, in 1966, I set up two computer networks between Lincoln Laboratory's TX-2 computer and System Development Corporation's Q-32 computer using a 1200 bps dial channel (high speed those days). Each computer was operating in time-sharing mode and permitted any program to dial the other computer, log-in, and run programs much as it would execute a subroutine call. The experiment showed that there was no problem getting the computers to talk to each other and use resources on the other computer - time-sharing operating systems made that easy. The real problem uncovered was that dial communications based on the telephone network were too slow and unreliable to be operationally useful. This work, jointly authored with Tom Marill, was published in the AFIPS FJCC proceedings, Nov. 1966 (2). The lesson learned - **a new data communications network is needed in order to successfully network computers.**

ARPANET Development

The chance to develop and build a major computer network experiment based on radically new communications technology came within a few months. I was asked to take over the responsibility of the ARPA Information Processing Techniques (IPT) office and manage and build its programs. ARPA was sponsoring computer research at leading universities and research labs in the U.S. These projects and their computers provided an ideal environment for an experimental network project; consequently, the ARPANET was planned during 1967 with the aid of these researchers to link these projects' computers together. One task was to develop an computer interface protocol acceptable to all 16 research groups. A second task was to design a new communications network technology to support 35 computers at 16 sites with 500,000 packets/day traffic. The initial plan for the ARPANET was published in October 1967 at the ACM Symposium on Operating System Principles in Gatlinburg Tennessee (3). The reasons given at that time for establishing a computer network were:

A. **Load Sharing:** Send program and data to remote computer to balance load.
B. **Message Service:** Electronic mail service (mailbox service).
C. **Data Sharing:** Remote access to data bases.
D. **Program Sharing:** Send data, program remote, e.g. Supercomputer.
E. **Remote Service:** Log-in to remote computer, use its programs and data.

The communications network design was that of the now conventional packet network: Interface Message Processors (IMP's) at each node inteconnected by leased telecommunication lines providing a store and forward service on very short messages. The main difference from later packet nets was that the IMP's were located at the computer sites and connected by a short parallel cable rather than a communications line interface.

Also presented at the Gatlinburg Symposium was Donald Davies's first open publication on the NPL packet network concepts presented by Roger Scantlebury, "A Digital Communication Network for Computers Giving Rapid Response at Remote Terminals." (4) It detailed the concept of a high level packet net with high capacity nodal switches and interface computers in front of mainframe computers. This was the first time that either Davies or I knew anything about the work of each other since our 1965 contact. The NPL paper clearly impacted the ARPANET in several ways. The name "packet" was adopted, much higher speed was selected (50 Kilobit/sec vs. 2.4 Kilobit/sec) for internode lines to reduce delay and generally the NPL analysis helped confirm the concept of packet switching.

Another confirmation of the basic concepts came from finally being able to read the Rand reports on distributed communications. The Rand work was very detailed, since it covered the whole network including microwave and one valuable analysis on routing. Their hot-potato routing algorithum was a useful starting point for the ARPANET routing design.

During 1968, a request for proposal was let for the ARPANET packet switching IMP equipment and the operation of the packet network. The RFP was awarded to Bolt Beranek and Newman Inc. of Cambridge, Massachussetts, in in January 1969. The RFP specified the general packetswitching concept, packet size, and interface protocol so that bidders could not totally change the system concept, to circuit or message switching for example. The two largest computer companies to receive the RFP no bid it because they didn't have mini-computers with which to make an economic bid. BBN bid the Honywell 516 mini-computer which was ideal for the task in 1969. Significant aspects of the network's internal operation, such as routing, flow control, software design, and network control were developed by a BBN team consisting of Frank Heart, Robert Kahn, Severo Ornstein, William Crowther, and David Walden. By December 1969 four nodes of the net had been installed and were operating effectively.

The first set of detailed papers covering the ARPANET were published in May 1970 at the AFIPS SJCC. (5-9) These papers reported the motivation and economics (5), the detailed design of the IMP (6), the network delay analysis and experience (7), the topological design programs and results (8), and the host-to-host protocol (9). These papers showed the world for the first time that packet **switching works**, that it is economic and that it is reliable and virtually **error free.** They also provided a complete description of how a working network was designed. As such these papers were the technical and motivational basis for many other network experiments around the world.

A Little Later

- Major efforts on host protocols
- Evolution of monitoring and control
- The TIP—and other network access devices
- Mixed line speeds and satellite links
- The Pluribus multiprocessor

More General Issues

- A climate permitting proposal variations
- A leader with considerable freedom at ARPA
- Control of both net and hosts
- Low levels of bureaucracy in government and at BBN
- Amazingly strong people at other contractors and early hosts
- AT&T coordination point
- Good fortune at Honeywell
- A success catastrophe
  - It worked, reliably and as specified
  - People quickly became dependent upon net, despite modest demonstration goal