Lawrence Berkeley National Laboratory

Title

COMPUTER NETWORKS IN THE PHYSICAL SCIENCES

Author

Austin, D.M.

Publication Date

1976-11-01

Presented at the EDUCOM Conference
on Computer Networking in the University:
Success and Potential, Boston, MA,
November 11, 1976

LBL-5594

COMPUTER NETWORKS IN THE PHYSICAL SCIENCES

Donald M. Austin

November 10, 1976

Prepared for the U. S. Energy Research and
Development Administration under Contract W-7405-ENG-48

0 0 0 0 4 6 0 6 3 7 5

Computer Networks in the Physical Sciences

Donald M. Austin

Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720

This work was done with the support from the U. S. Energy Research and Development Administration.

I. INTRODUCTION

The ERDA (Energy Research and Development Administration) research community consists of eight multipurpose laboratories, more than a dozen single purpose laboratories and energy research centers, and over 100 universities distributed throughout the United States (see Fig. 1). In addition, a large number of private contractors undertake development and demonstration projects for ERDA. The nature of ERDA's mission requires the development of large scale modeling systems, analysis programs, data bases, and hardware facilities. These supply powerful computational techniques to a wide spectrum of research activities - from the investigation of the fundamental laws of nature to energy policy analysis. Computer science research in ERDA has created a variety of unique facilities, ranging from innovative computer hardware configurations, through sophisticated algorithm developments, to high level human-machine interface techniques. This research has brought computational and data management techniques to the ERDA community which form the basis for most of the theoretical and experimental analysis necessary for scientific research.

Since the ERDA community is widely distributed geographically, it has been difficult to provide the complete spectrum of computational and data management facilities equally to all researchers. The practical consequence has been, in some cases, duplication of effort; or, in other cases, lack of computational facilities. The expense involved in transporting specialized computer software among diverse operating systems is appreciable. It is widely acknowledged in the computer science community that that ratio of software development costs to hardware costs is approaching four to one. On the other hand, modern technology has led to the development of specialized hardware which is much too expensive to distribute on an equal basis to all ERDA sites. To achieve a balanced distribution of computer resources throughout the ERDA community, one solution is provided by a distributed computer network.

There are several kinds of networks being developed today. Among these are the star network, where one host computer controls all network activity in a centralized configuration with all other nodes connected directly to that host. The most basic kind of star networks consists of a large computer center supporting remote job entry (RJE) stations and interactive terminal ports. Most modern computer centers provide this basic level of network service. The RJE's typically have card readers, line printers, and perhaps a tape drive. Interactive terminal ports can be hardwired directly to the front end of the host computer or provide dial-up modems for remote access over telephone lines.

Another interesting configuration is the homogeneous network, consisting of two or more similar mainframes hooked together running a single operating system. This arrangement provides multiples of the usual facilities found in a single system, usually in a manner transparent to the user. This configuration can also be the center of a star network, as in most commercial time sharing services.

The most interesting and versatile configuration is the heterogeneous distributed network, where different computing services are made available to a distributed community of users, as in the ARPA network. This configuration is somewhat more difficult to use, since the nodes may have different operating systems, accounting algorithms, and so forth, but it does provide access to unique resources, including hardware, software, data bases, and personnel.

The ERDA Network Experimentation Project [1] was established to explore the use of distributed computer networks as a potential long term solution to the problems of reliable access to unique computational facilities throughout ERDA. The primary test vehicle for this research is the ARPA network (Fig. 2). The ERDA sites currently connected to the ARPA network are depicted in Fig. 3. Investigators from each site use the ARPA network, and sometimes TELENET (a commercial "value added" packet-switched network), to access remote time sharing systems and to transfer data files between sites. This activity has been increasing during the past year, and has produced many examples of uses of the network for conducting scientific research. [2]

II. CAPABILITIES OF A DISTRIBUTED NETWORK

The important capabilities inherent in a distributed computer network are the sharing of unique hardware, software, data bases, and personnel. The possibilities of remote scientific collaboration using networks are increasingly exciting as the development of device independent graphics systems, teleconferencing systems, and network mail facilities continues. The richer scientific environment offered by networks should provide a new style of scientific research in the near future. Some examples of these resources are given below.

a. Hardware Resources

The 60 or so nodes on the ARPA network offer a wide variety of hardware capabilities which have never been collected in a single site accessible to the general scientific community (Fig. 4). These include, for example, the ILLIAC IV parallel processor machine at NASA-Ames, the CDC 7600's at LBL and BNL, and IBM 370/195's at UCLA, ANL, and RHEL for large scale computation. There are terabit mass storage devices available at CCA (the DataComputer), at LBL (the IBM Photodigital Chipstore), and at NASA-Ames (the Unicon device). Computer output to microfilm (COM) devices can be accessed at several sites (with appropriate arrangements for mailing the output back to the user), providing high quality graphics output for network users.

b. Software Resources

Current legend states that the cost of software development is often four times the cost of hardware. I am not exactly sure what this means, but one does usually find several hundred people writing programs for any large machine. One might expect that the standards for FORTRAN, PL/I, ALGOL, COBOL, etc. would obviate the claim for unique software, just as Esperanto was supposed to make Mr. Berlitz obsolete. For some time to come, I foresee that network access of special software has as good a future as Berlitz. After all, the international standard ASCII character codes have been adopted by every major computer manufacturer - except, of course, IBM, CDC, UNIVAC,....

It is a fact of life that really large scale software systems tend to take advantage of machine architectures and will thus be relatively difficult to transport. If they can be used over networks, there is no need for transporting. Figure 5 represents the concept of network access of special software.

c. <u>Data Bases</u>

The strongest case for distributed processing can be made for access to unique data bases. Consider, for example, the very large data bases containing the decennial census (on the order of 20 billion characters). This represents a vast amount of data of interest to a broad spectrum of users. To transport such a large data base to thousands of computer systems is a horrendous and wasteful task. While few people ever require the entire data base, they often need to sample freely from the data as the need arises. Clearly, if one could store the data efficiently on one system, and provide easy and efficient access to all users, a real savings in time and effort would accrue - perhaps even enough to finance the development of human oriented query languages, high level network protocols, archival mass storage devices, and other such developments required to make distributed networks function smoothly. Although a somewhat smaller set of users exists for scientific data bases, such as the elementary particle data base, the table of isotopes, neutron cross sections, and the physcial properties data base, the principle is the same.

d. <u>Research Collaboration</u>

Perhaps the most exciting consequence of computer networks is the possibility of enhancing research collaborations by scientists from several institutions. The field of high energy physics is particularly affected by this requirement, because there are only a few large particle accelerators available (such as Fermi Lab, SLAC/LBL, CERN, DESY, Rutherford, etc.). These machines must be used by all experimental physicists to do their work, and the expense and lengthy times involved in most experiments demand that the work be collaborative. Similar statements obtain for magnetic fusion energy devices, such as TOKAMAKs, TORMACs, Baseballs, and mirror machines. In order to do physics experiments, groups from several institutions join their efforts in setting up the experiment, taking data, and analyzing the data. The advent of computer networks linking the host computer systems of the home sites of these researchers has already had an impact on the quality of research collaboration. When several sites are involved in collaborative data analysis from a single experiment, it is mandatory that good communication be established between the groups, and that the analysis techniques are compatible and accurate. Heretofore, this function has been satisfied only through travel, telephone, and the U. S. mail. By sharing data, software, and comments over networks, this process has increased the quality and timeliness of scientific collaborative research.

In other fields, the comparison of models will become practical over networks, allowing consistency checks and model expansion to become feasible. Regional energy models can be hooked together to form a national model; local models can be "aggregated" to provide input to regional models; and such matters as the appropriate units of quantities can be decided upon collaboratively. These aspects, including the sharing of model data bases, make computer networks an exciting medium for doing scientific research in the very near future.

III. EXAMPLES OF APPLICATIONS

The following examples were taken from experiments already conducted or in the planning stage under the ERDA Network Experimentation Project. The difficulties involved have been deliberately underrepresented for the purpose of this illustration. These difficulties fall mainly in the areas of low effective transmission rates and the lack of adequate system documentation available. Nevertheless, the examples serve to illustrate the potential of computer networks for scientific research.

a. Partial Wave Analysis

Collaborators at Lawrence Berkeley Laboratory (LBL) and Rutherford High Energy Laboratory (RHEL) are just beginning to analyze data on $\bar{p}p$ - 2 pi experiments at 16 different energies. The analysis proceeds in a series of six more or less distinct steps involving two network nodes (LBL and RHEL) and another two potential nodes (CERN and BNL) (Fig. 6). The experiment proceeds as follows:

i) Data is sent to RHEL from LBL (over the network), CERN (over an RJE link) and from BNL (currently, by air mail).

ii) The first phase of the analysis is carried out on RHEL's IBM 360/195, producing a large set of coefficients for the partial wave amplitudes.

iii) The file of coefficients is transferred to LBL over the ARPA network. At LBL, a cluster analysis program divides the solution set into several groups. This is a lengthy process requiring human judgement, so during this phase the collaborators use network mail facilities to communicate on a daily basis.

iv) The results of the cluster analysis are subjected to a continuity fit program on LBL's CDC 7600 for further selection of solutions.

v) These results are then transferred back to RHEL, where they are used as weak constraints for a pole extraction fitting program (to select resonance states).

vi) The final solutions are input to a graphics display program at RHEL. The displays are sent to LBL over the TELENET connection to a storage scope display terminal.

Clearly, step i) is a potential network process, since BNL is almost ready to connect the central computer facility to the network, and the amount of data contributed by each site is small enough to transmit over the network.

In several of the steps, the potential exists for using available analysis software wherever it exists on the network. For example, a cluster analysis program developed by a chemist at LLL could be used for that portion of the analysis.

b. Track Sensitive Target Experiment

The Track Sensitive Target (TST) [3] is a perspex box of hydrogen immersed in a neon-hydrogen bubble chamber at RHEL. The TST provides isolated protons for a pion beam to interact with, and the outgoing gamma rays are detected by the neon-hydrogen chamber. This experiment was a collaboration between LBL, RHEL and Turino.

The film from this experiment was divided three ways, with LBL, RHEL, and Turino each analyzing a third of the events (Fig. 7). It was necessary to determine any biases produced by the three separate analysis systems, so close collaboration results was required throughout the analysis process. The long delays inherent in the postal system (a minimum of ten days) tend to produce undesirable short cuts and oversights. Using the network mail facility and the capability of transmitting graphics displays over the network, LBL and RHEL were able to maintain daily communication. In contrast, Turino had to send two physicists to RHEL for several months to complete their work. This exercise, which was one of the first research experiments carried out over the LBL link to the network, demonstrates one of the most powerful applications of the network. It does not speak, of course, to the loss to the researchers of two months in the lovely Oxford countryside.

c. Mathematical Software Portability

The Numerical Software Group at Argonne has been developing reliable, transportable mathematical software for several years [4] (Fig. 8). The EISPACK package for solution of eigenvalue problems is perhaps the best known of their efforts. During the development of mathematical software packages, considerable effort goes into making the software portable and accurate to a certain precision on several machines (IBM, CDC, UNIVAC). By gaining access to other computers via the network, this task can be greatly facilitated. The first effort to test this concept was a particular routine in the minimization package, MINPACK. The routine, Davidon's Optimally Conditioned Optimization Routine, was compiled and run on LBL's CDC 6600, and the output transferred back to Argonne for appraisal. It has been estimated by the director of that project that the network would have saved up to eight months in the development of EISPACK, and an even greater potential savings exists for the development of MINPACK.

d. Gateways Between Networks

The Magnetic Fusion Energy network [5] (MFEnet, formerly called the CTRNet, Fig. 9) has been established by ERDA to provide a central facility for the support of research in this area. This star network provides access to major participating laboratories and universities throughout the country. However, there remain many sites with somewhat smaller MFE programs, particularly universities, which do not have reliable access to the MFEnet.

A network gateway is a host having access to two or more networks, and which supplies the necessary protocol translation software to allow messages to flow from one network into another (Fig. 10). An example of such a gateway is the RATS [6] system developed by LLL. RATS is an ARPA network host, and through its dial-out facilities, has access to the MFEnet centered at LLL (as well as to commercial networks, such as TELENET and TYMNET).

Network investigators at NYU and ANL [4] were able to gain access to the main MFE CDC 7600 through the ARPA network by using the RATS gateway. Although this was done on a test basis only, it proves the feasibility of internetwork communication as a mechanism for research sharing on an even broader scale. Clearly, any system which supports more than one network has the potential of linking these networks together for an even broader set of resources.

e. Teleconferencing

The PLANET-II teleconferencing system [7] developed by the Institute for the Future is being used for daily communication by nearly a dozen collaborative groups within the ERDA community (Fig. 11). PLANET provides a bookkeeping service for messages, a survey or vote-taking service, a review service, and several other handy features which make communication among large groups easy. Since a printed transcript is always available, it far outweighs the telephone conference for many purposes. A sample of the groups using the teleconference facilities includes:

i) The ERDA Interlaboratory Working Group for Data Exchange, involving researchers from eight laboratories and ERDA Headquarters;

ii) Three panels on the ERDA Network Experiment Project, including people from six laboratories and universities discussing the implementation, application, and objectives of network use;

iii) An international group developing a transportable data management system (the Berkeley Data Management System);

iv) An international group developing robust mathematical software.

It is significant that these groups, which formerly met at most on a quarterly basis, are now keeping in daily contact over the network. The facilities available to conferees include the ability to enter print files or data files into the conference, collaborate on group reports, broadcast results to all participants immediately, communicate asynchronously (at any time, independent of time zone differences), and to save transcripts of the conference proceedings for further distribution.

IV. THE FUTURE

There are several serious problems to be faced in existing networks. Primarily, the transmission rates are too slow for many applications, the heterogeneous operating systems are unfathomable to the casual user, and the availability of resources is not well documented. The course of computer science research in networking must take all these impediments into account if networking is to ever reach its full potential to the scientific community.

The future of communications technology seems particularly bright today. The use of satellites, helical wave guides, and the like make megabit/second transmission a five year goal. Other possibilities include the piggybacking of data, voice, video, and facsimilie signals on high bandwidth media. It seems certain that the large scale use of high bandwidth communication media will bring the price to within reasonable limits very soon.

The human-machine interface is perhaps the most troublesome problem faced by network researchers. After all these years, the computer is still a variety of dumb beast which simply refuses to bow to human languages. The availability of intelligent terminals, with local storage and powerful microprocessors, can allieviate this problem somewhat, but what is still needed is the universal job control language (another Esperanto).

Machine-machine interfaces show some signs of being solved. The international CCITT X25 [8] protocol is an encouraging sign of cooperation on this problem. Probably the next step is to interface machine directly with the communication medium, when that medium becomes capable of handling speeds comparable to other peripherals.

There is much work to be done to make distributed computing transparent to the user, but I am confident that the next few years will provide great strides in this direction.

REFERENCES

1. Austin, D. M. and Hall, D. E., "Computer Science Research in Resource Sharing Over Networks," in Selected Examples of ERDA Research in the Mathematical and Computer Sciences, Ed. M. E. Rose, ERDA-76-118, 1976.

2. Hall, Dennis E. and Austin, Donald M., A Summary of Findings of the ERDA Network Investigators Panel, LBL-5322, June, 1976.

3. Austin, Donald M. and Hall, Dennis E., Status Report on ERDA Network Experiments, UCID 3817, Feb., 1976.

4. Amiot, L. W. and Bare, R. A., Status of Argonne ARPANET Implementation and Experiments, ANL Internal Report, Sept., 1976.

5. Killeen, John, "The CTR Computer Network," in Proceedings of the Berkeley Workshop on Distributed Data Management and Computer Networks, LBL-5315, May, 1976.

6. Donnelley, James E., "Controlling Transactions Between Distributed Computer Resources," in Proceedings of the 5th Texas Conference on Computing Systems, University of Texas, Austin, Oct., 1976.

7. Vallee, J., Johansen, R., Lipinski, H. and Wilson, T., "Pragmatics and Dynamics of Computer Conferencing: A Summary of Findings from the FORUM Project," in Proceedings of the Third International Conference on Computer Communication, Ed. P. K. Verma, Toronto, Aug., 1976.

8. CCITT - COM CII Proposal for Revised Draft Recommendation X.25, France and UK Post Office, November, 1975.

ERDA FIELD ORGANIZATION

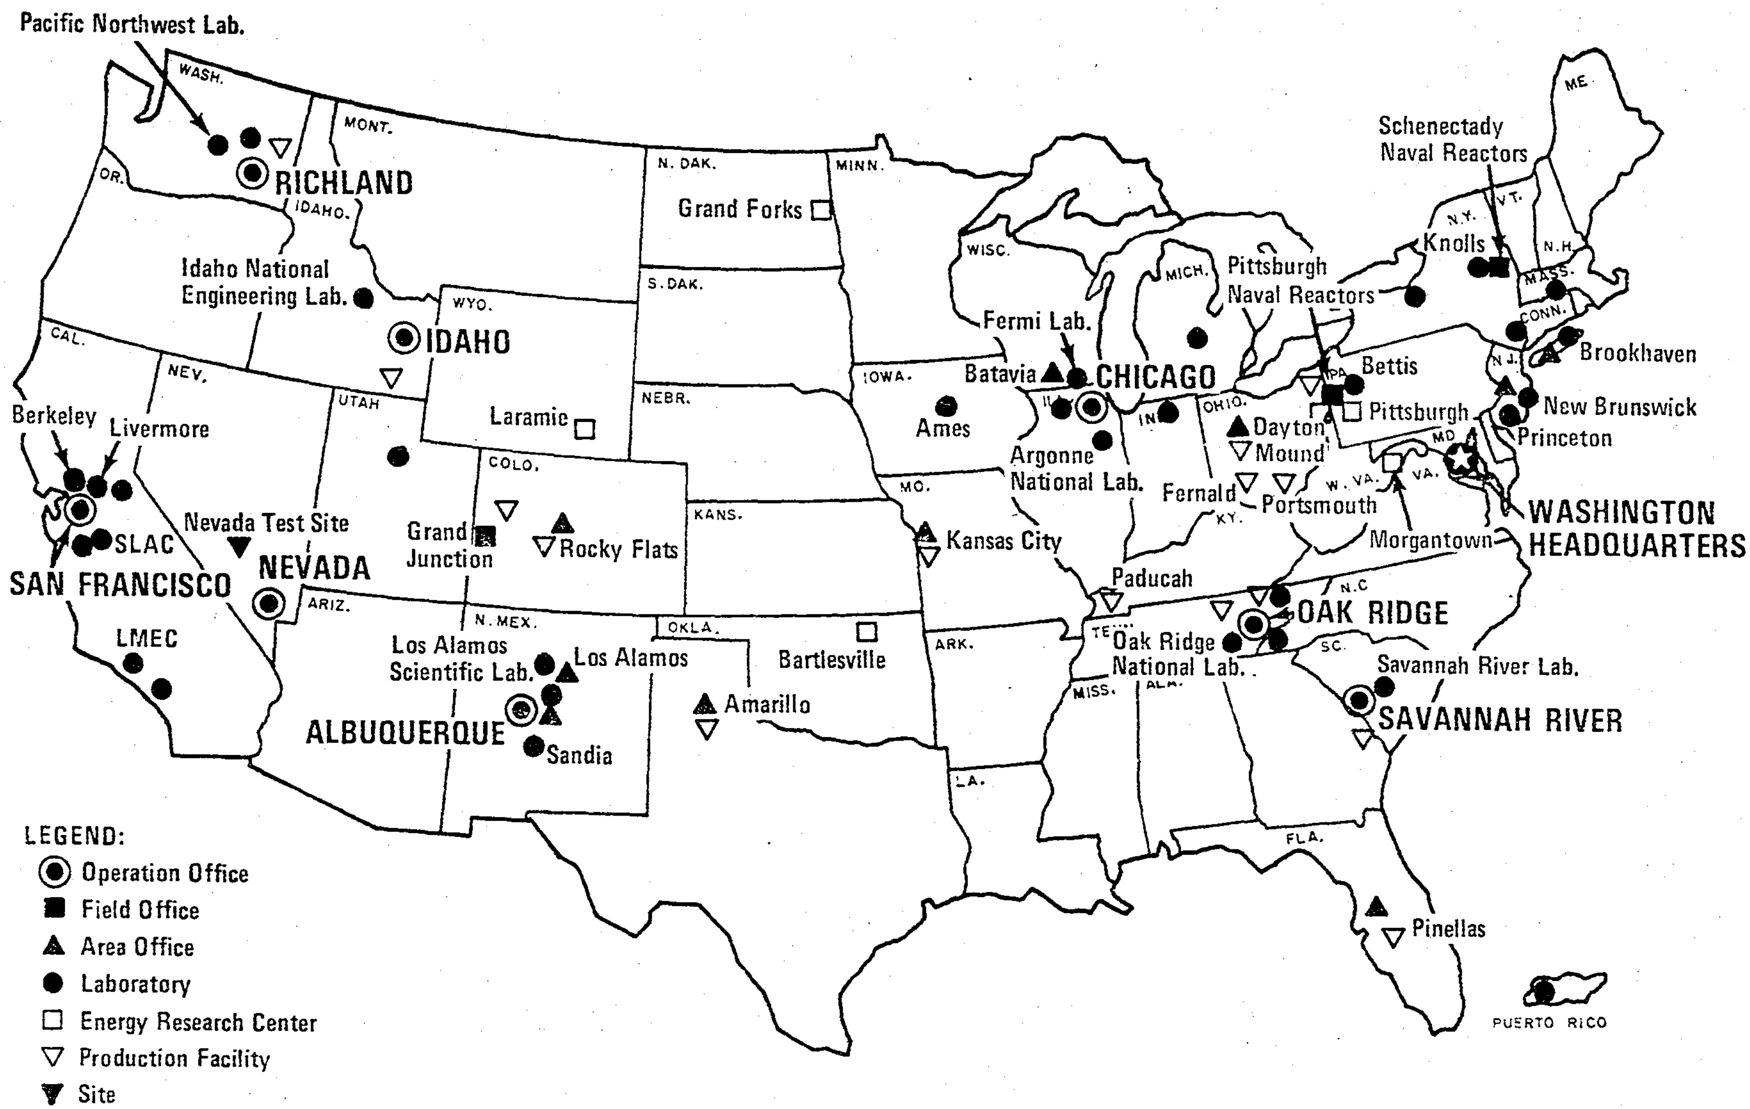

Fig. 1

February 1976

ARPANET GEOGRAPHIC MAP, AUGUST 1976

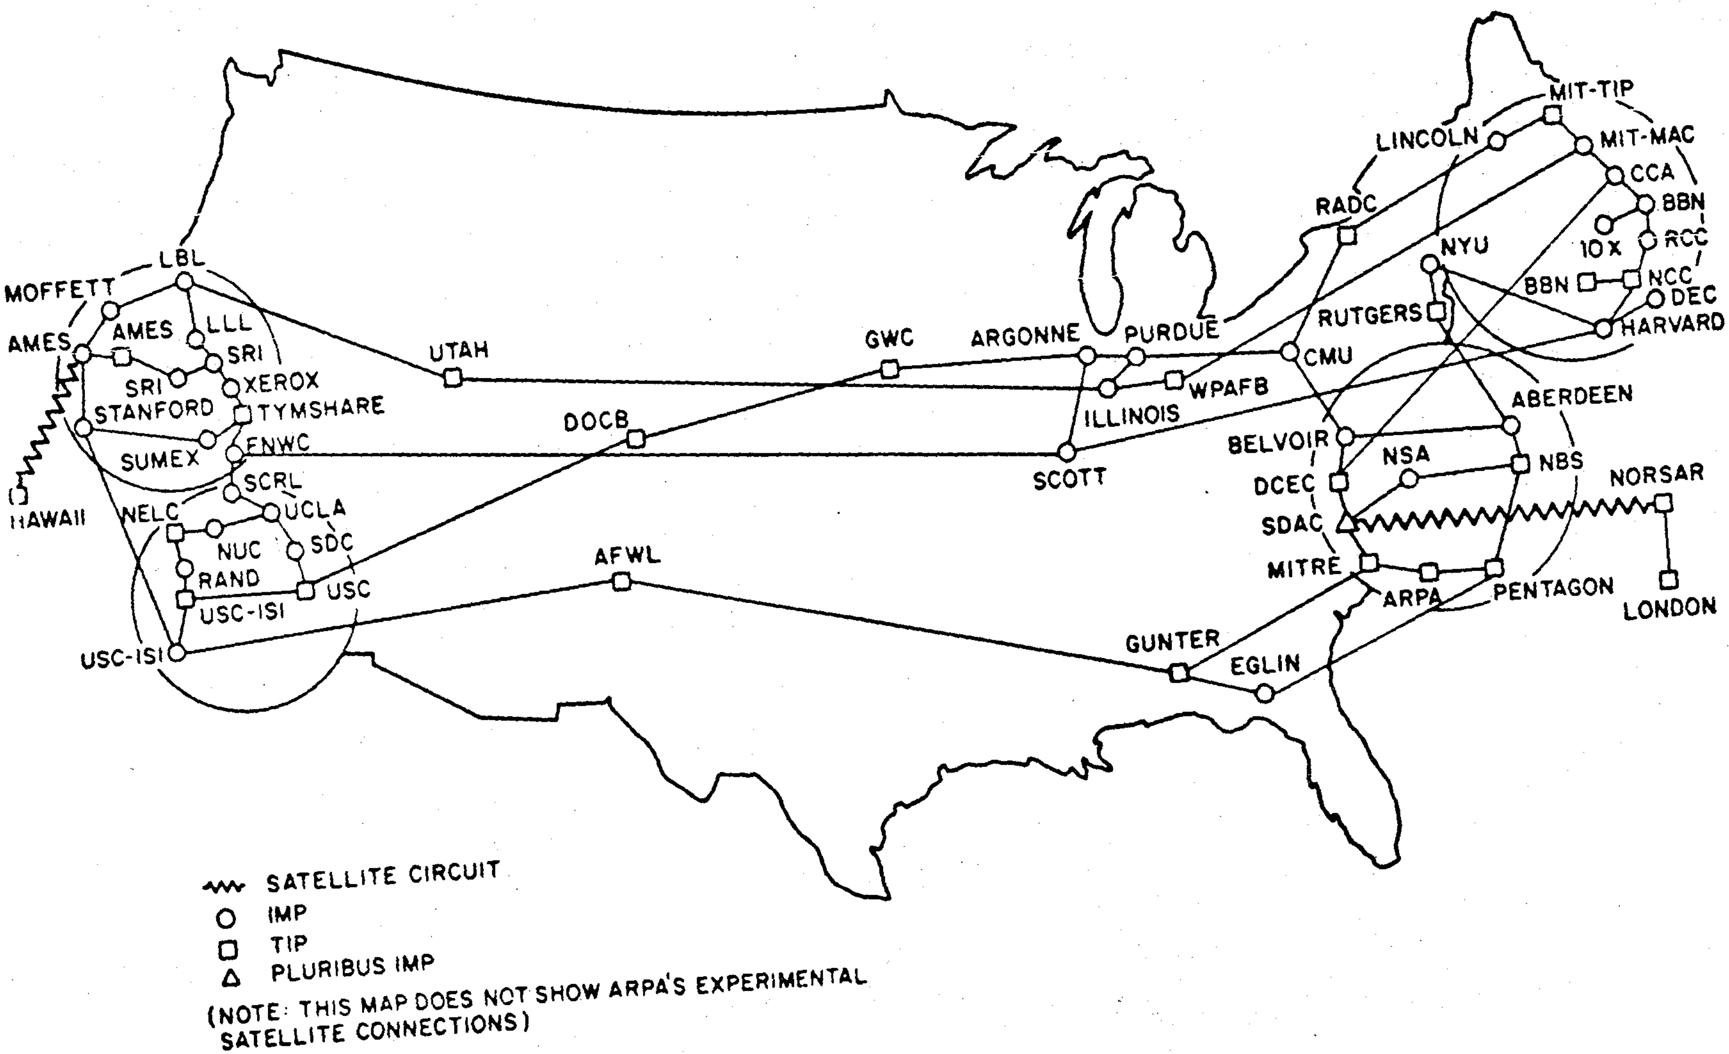

Fig. 2

ERDA SITES ON THE ARPA NETWORK

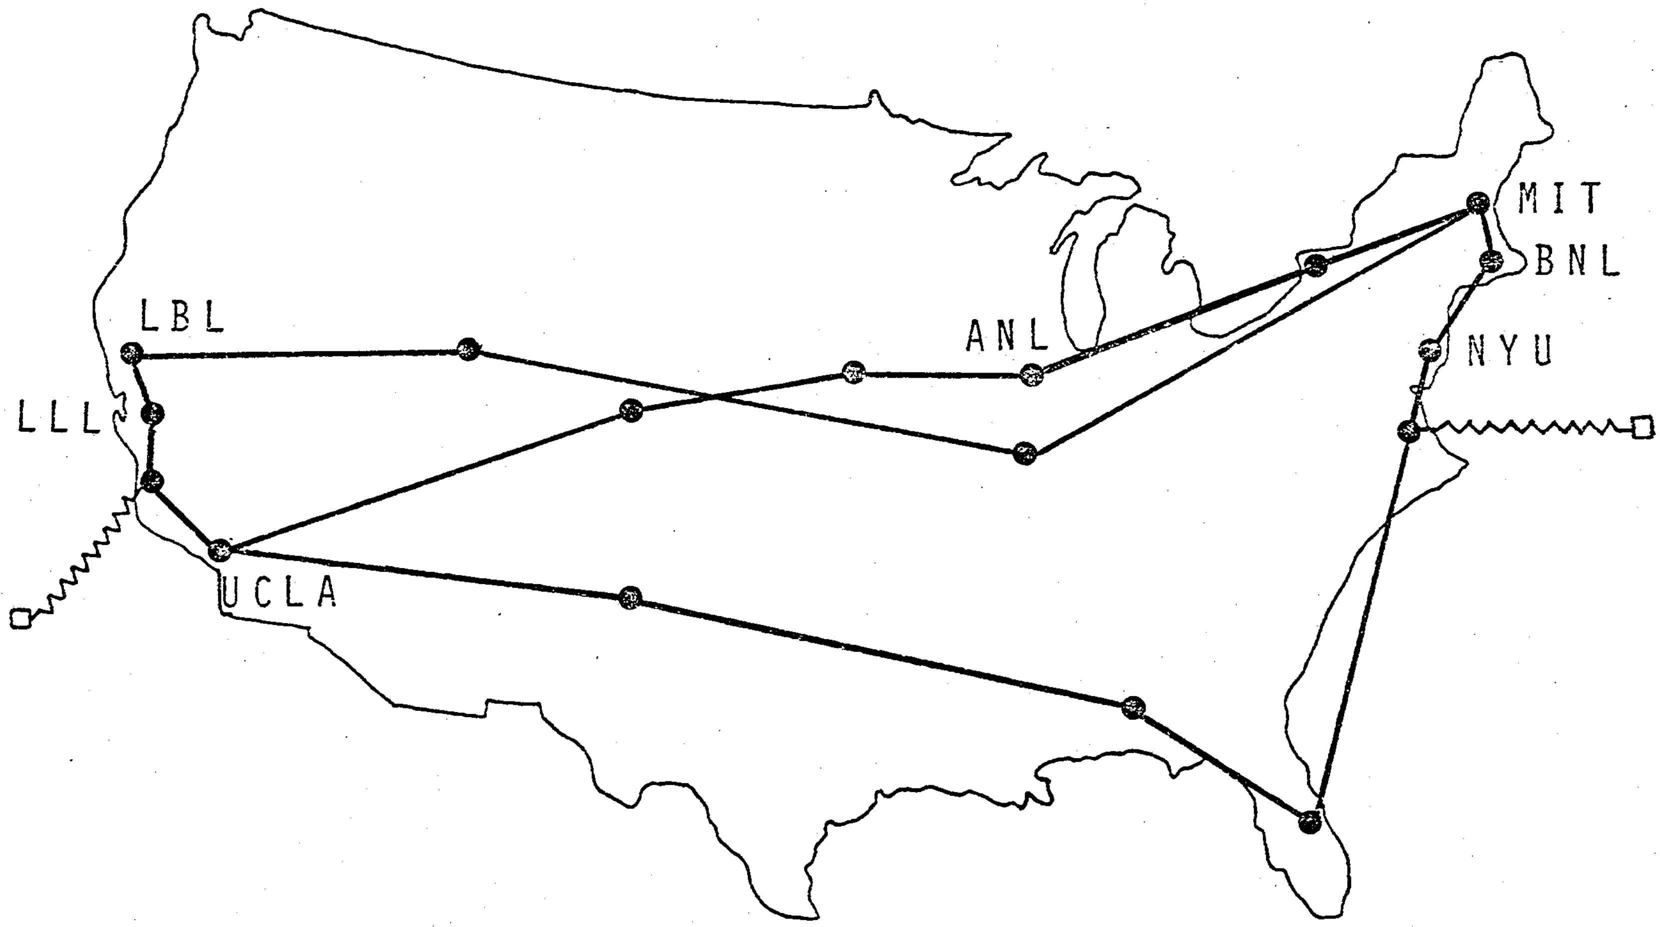

Fig. 3

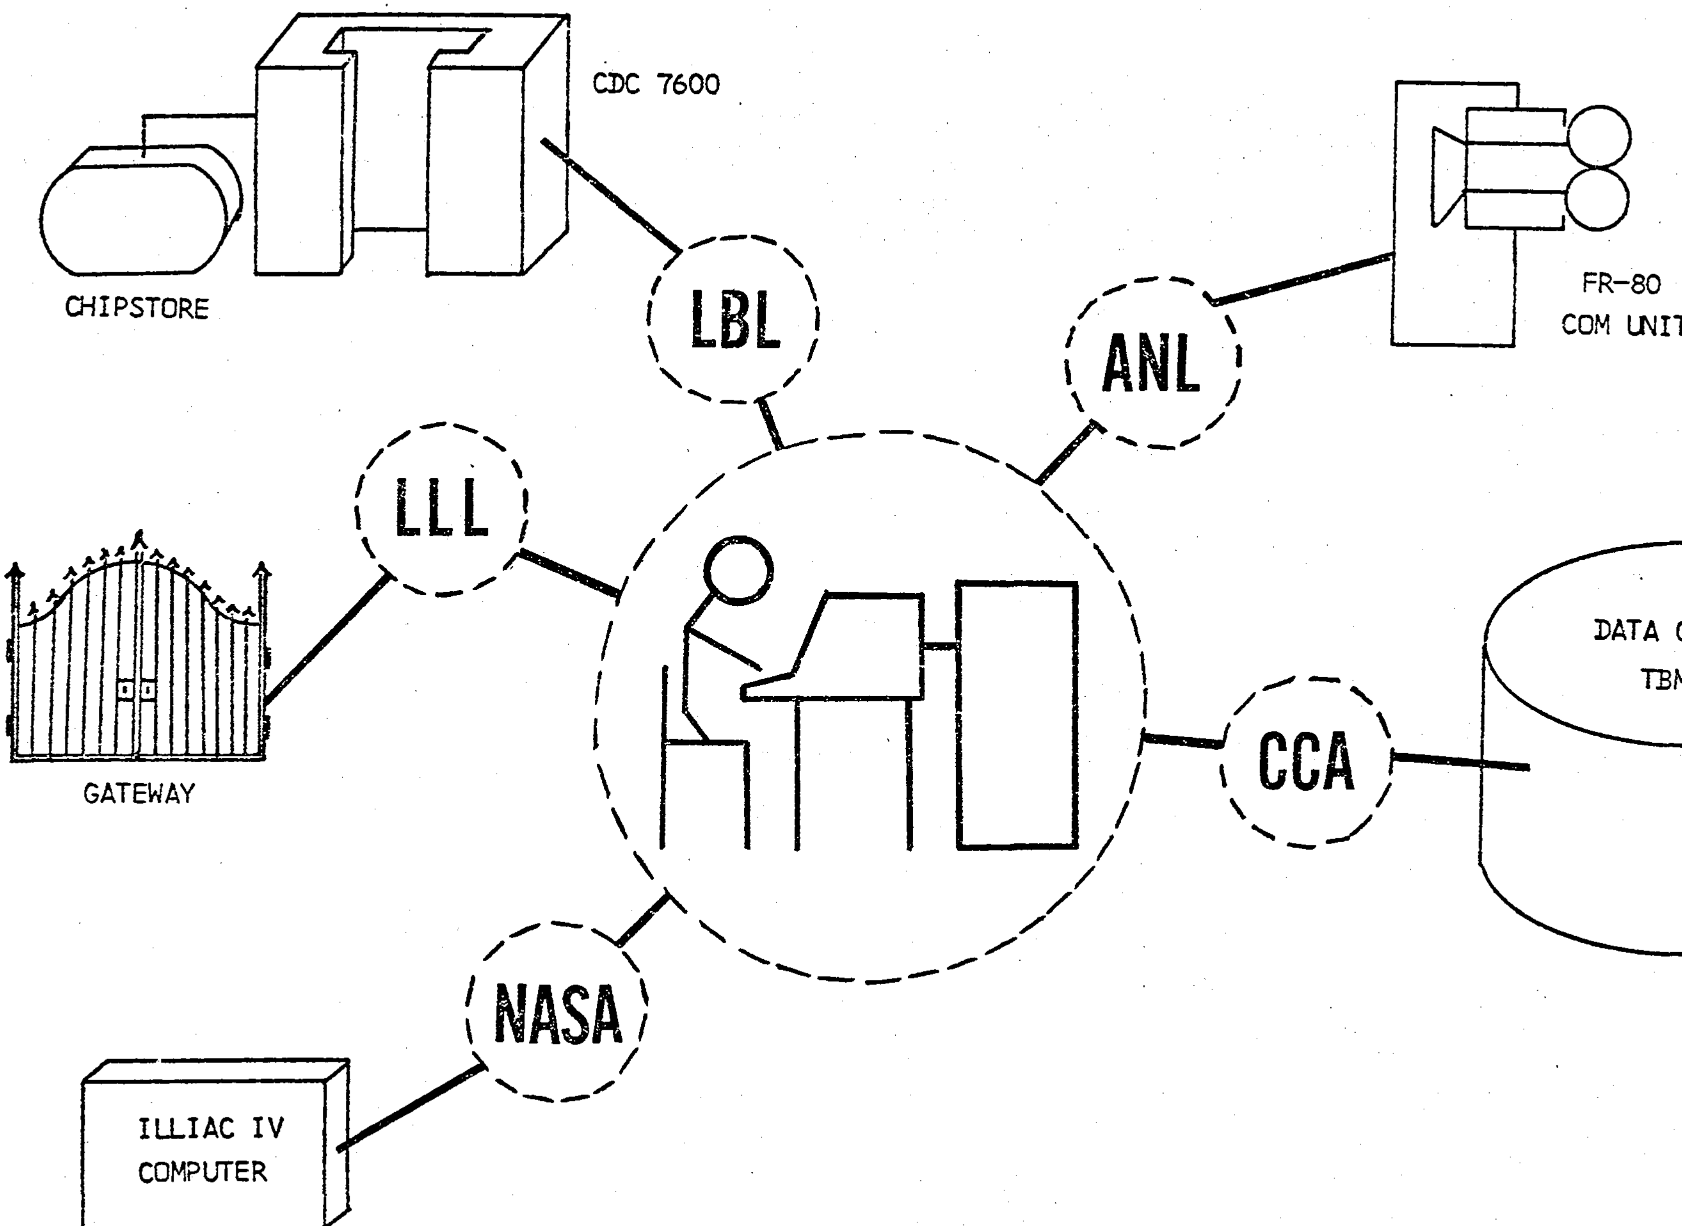

ACCESSING UNIQUE HARDWARE

Fig. 4

-15-

0 0 0 0 4 6 0 0 0 8 4

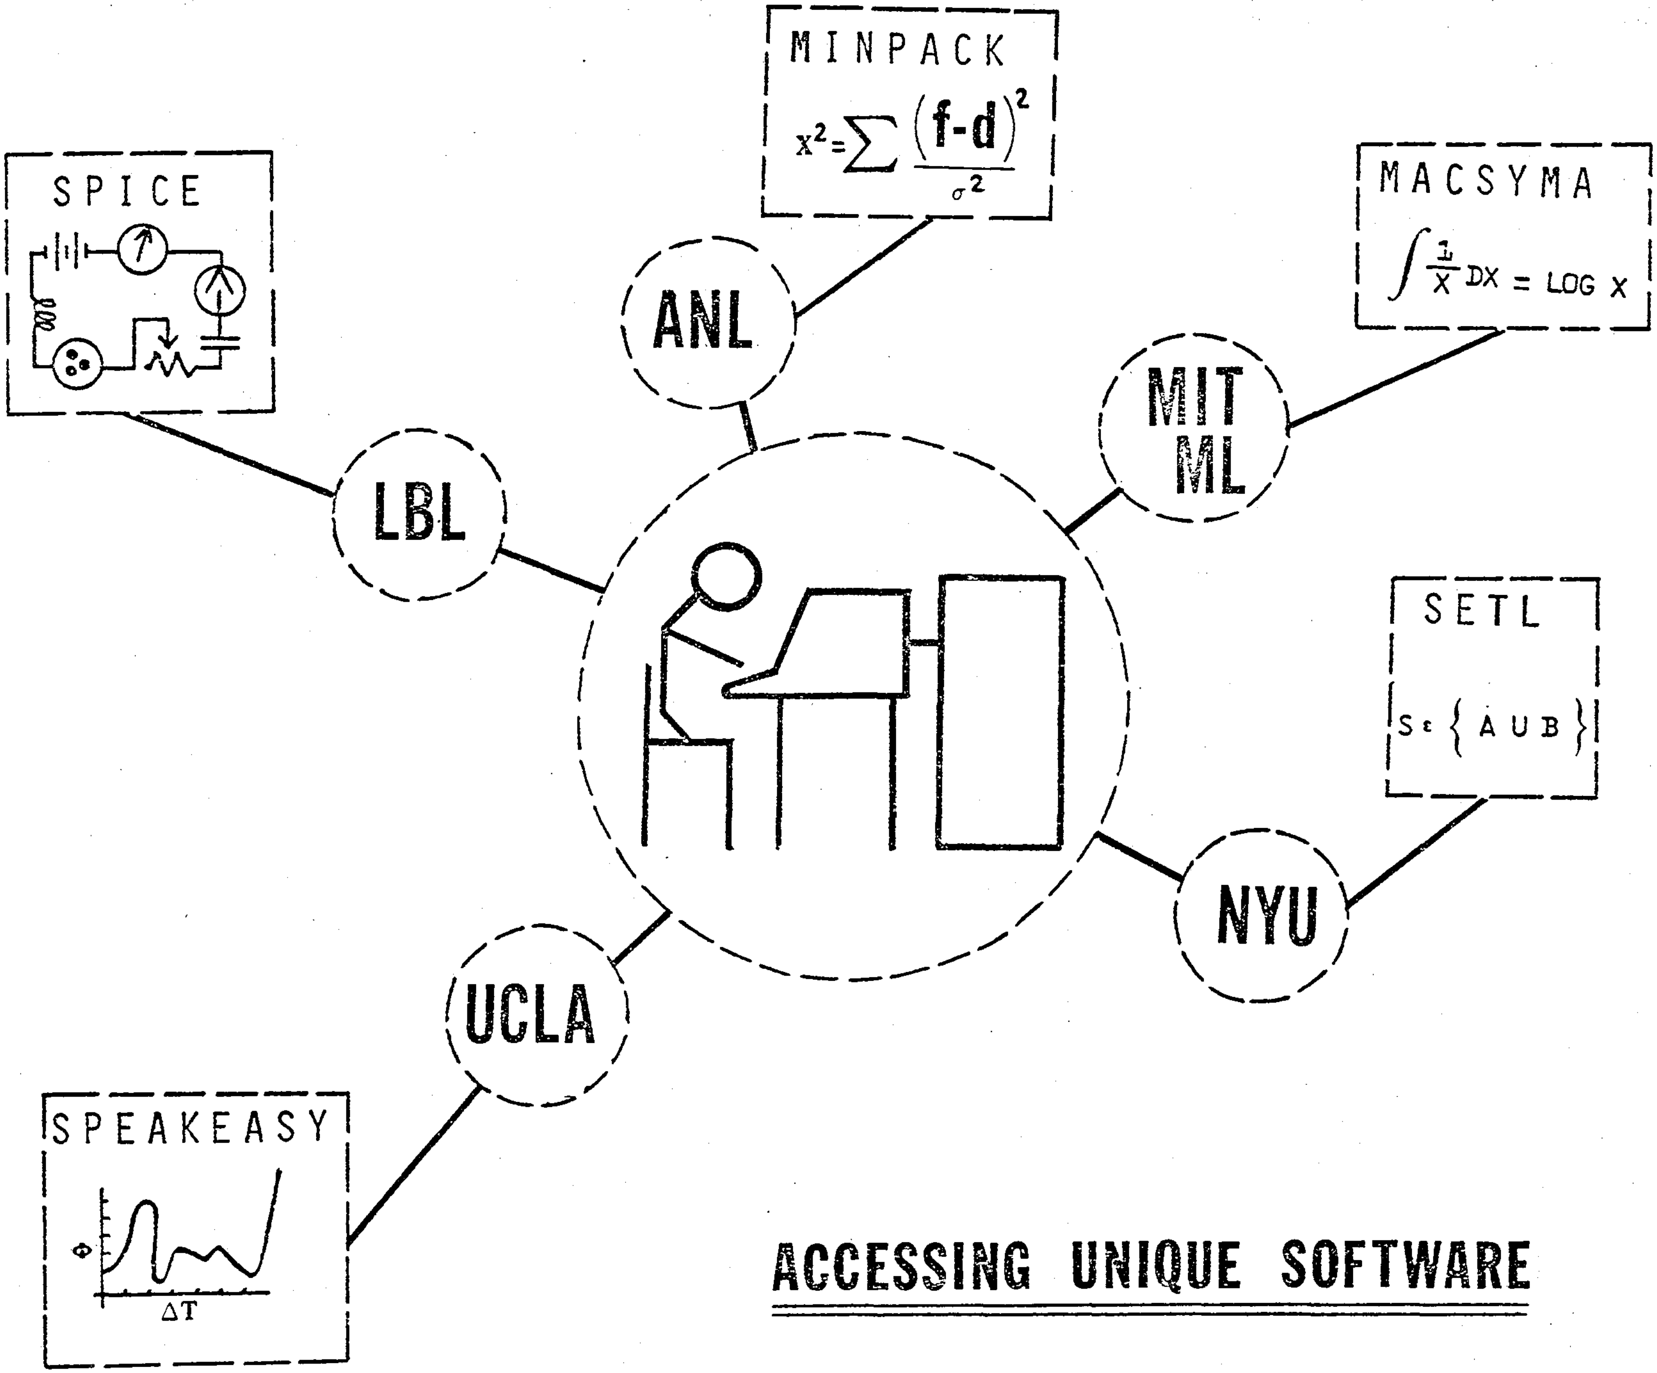

ACCESSING UNIQUE SOFTWARE

Fig. 5

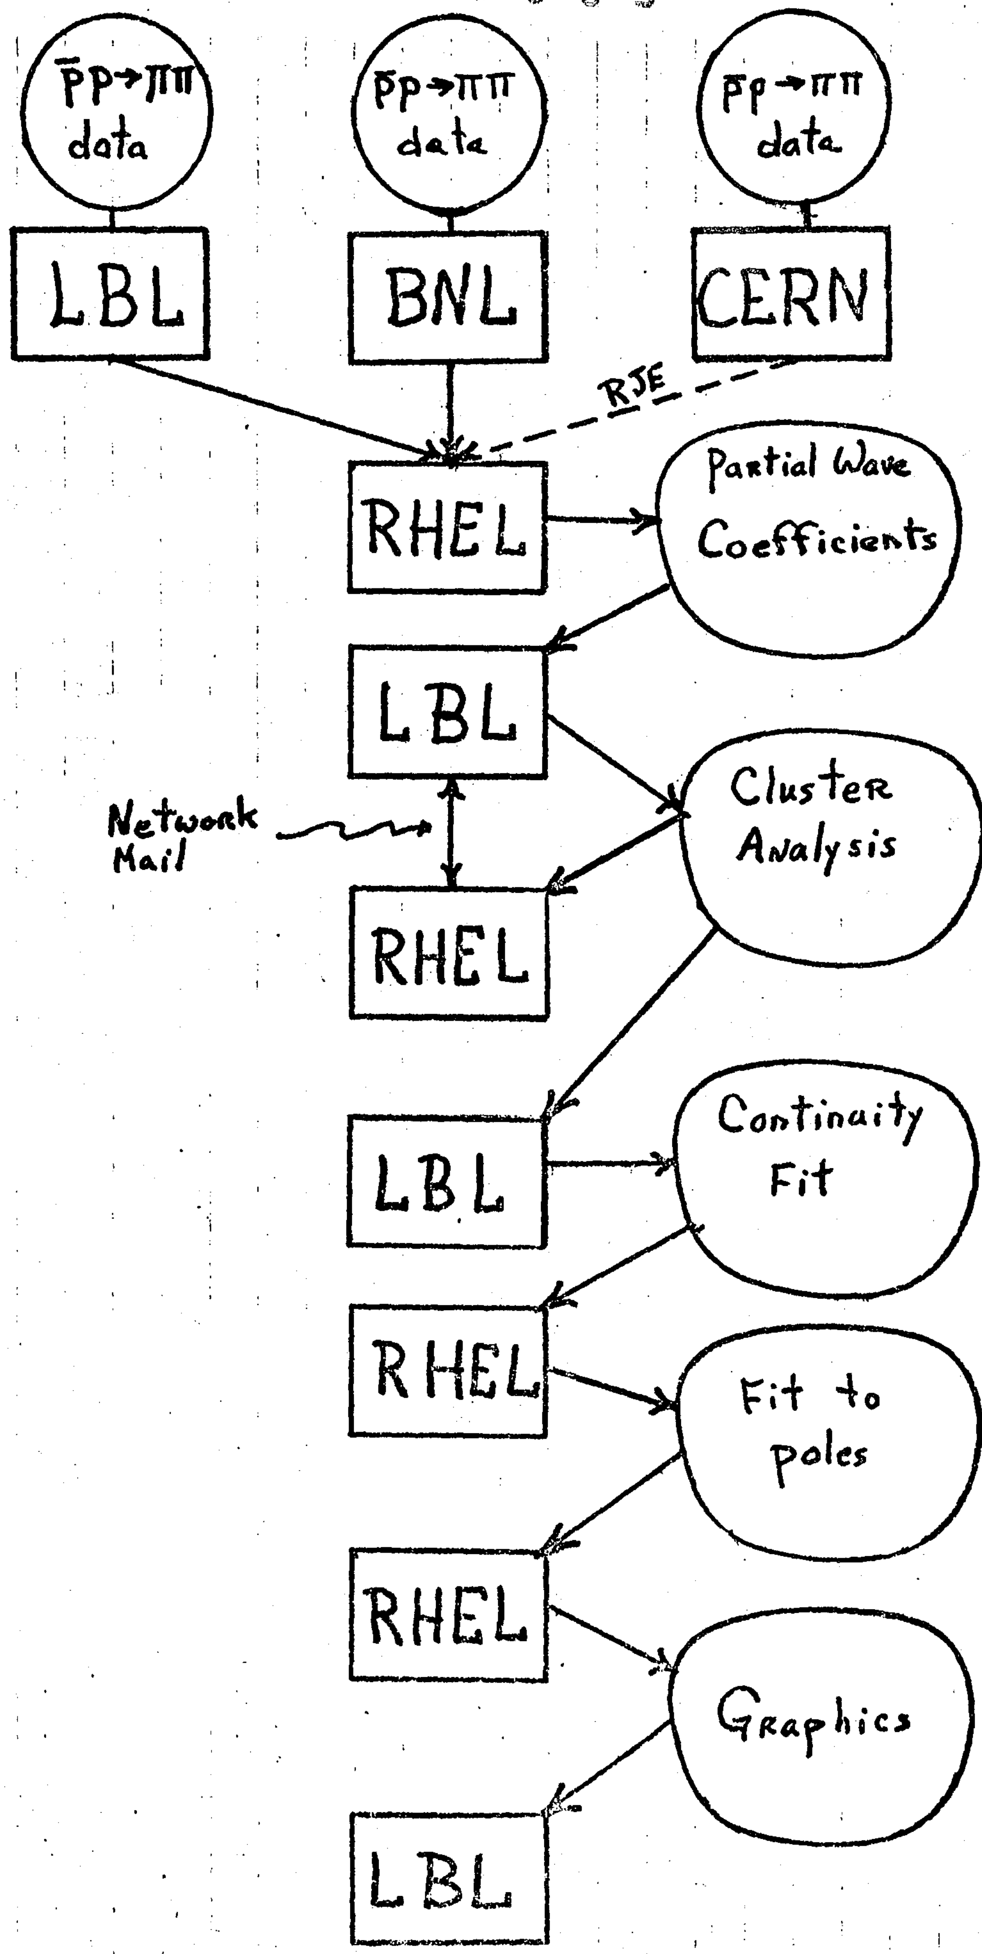
p̄p→ππ data
p̄p→ππ data
p̄p→ππ data
LBL
BNL
CERN
RJE
RHEL
Partial Wave Coefficients
LBL
Network Mail
Cluster Analysis
RHEL
LBL
Continuity Fit
RHEL
Fit to poles
RHEL
Graphics
LBL

Fig. 6

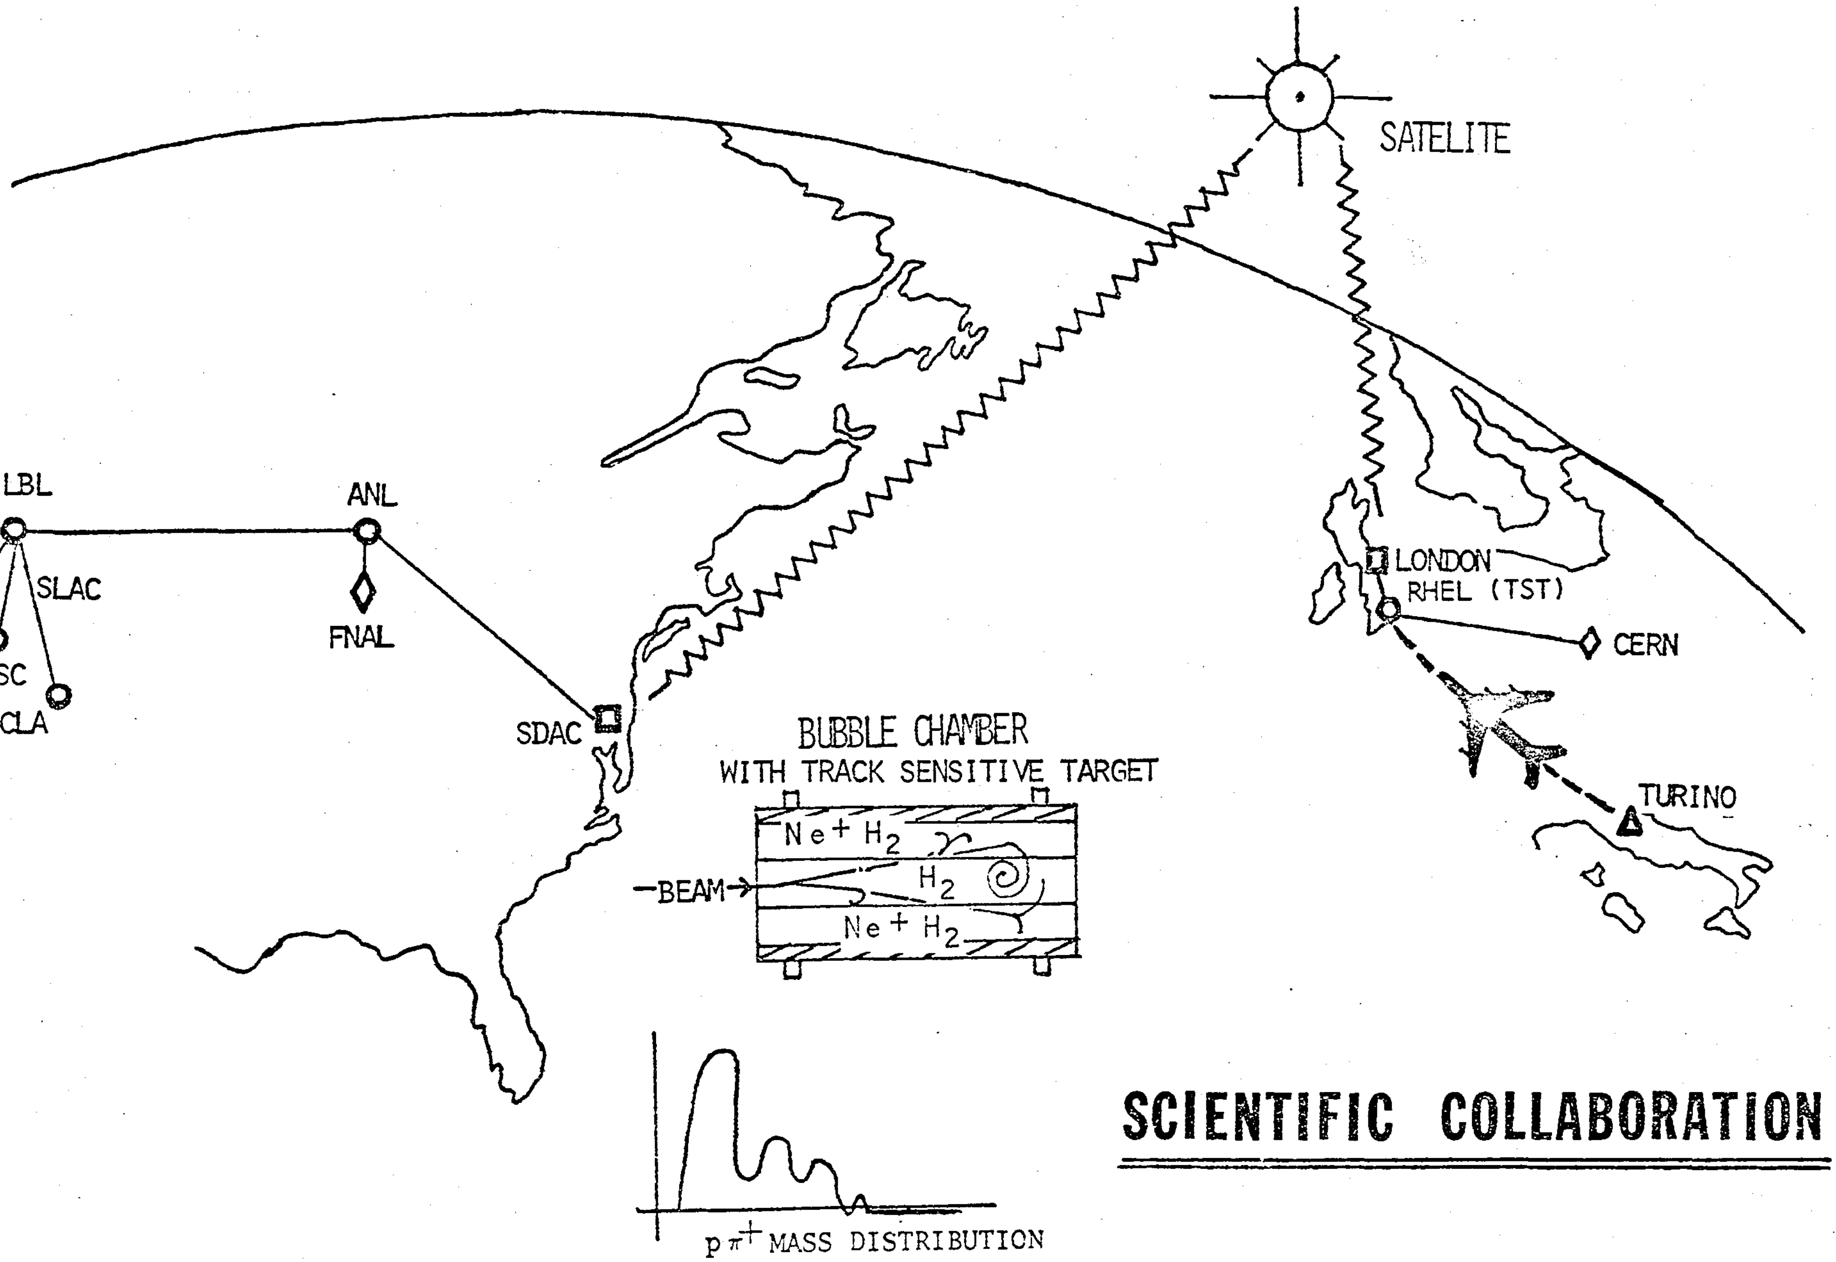
SATELITE
LBL
ANL
SLAC
UCSC
UCLA
FNAL
SDAC
LONDON
RHEL (TST)
CERN
TURINO
BUBBLE CHAMBER
WITH TRACK SENSITIVE TARGET
Ne + H2
H2
Ne + H2
BEAM
pπ+ MASS DISTRIBUTION
SCIENTIFIC COLLABORATION

Fig. 7

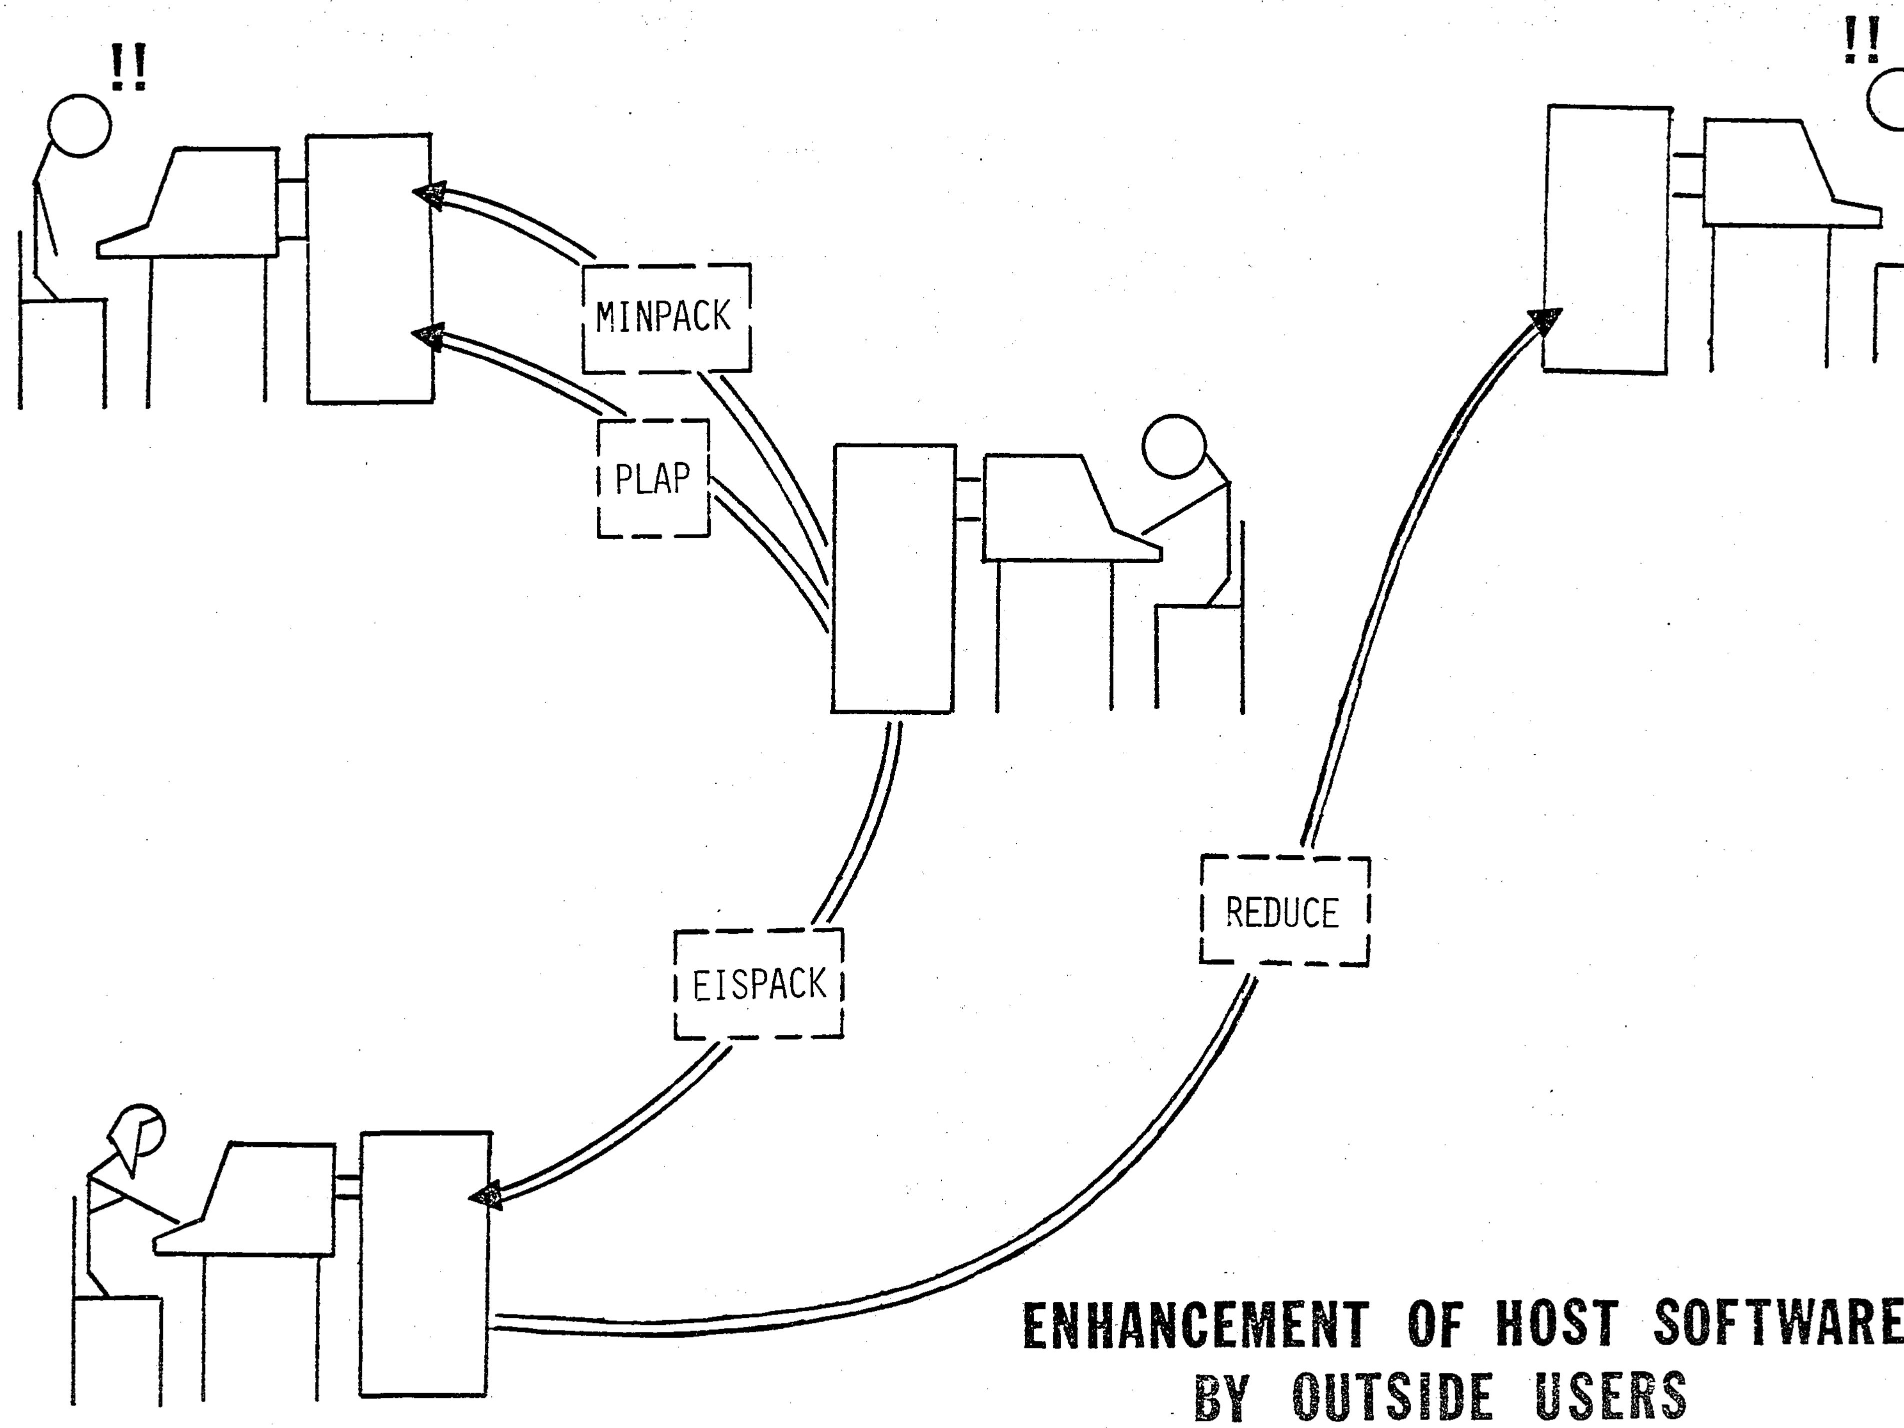
!!
MINPACK
PLAP
!!
EISPACK
REDUCE
ENHANCEMENT OF HOST SOFTWARE
BY OUTSIDE USERS

Fig. 8

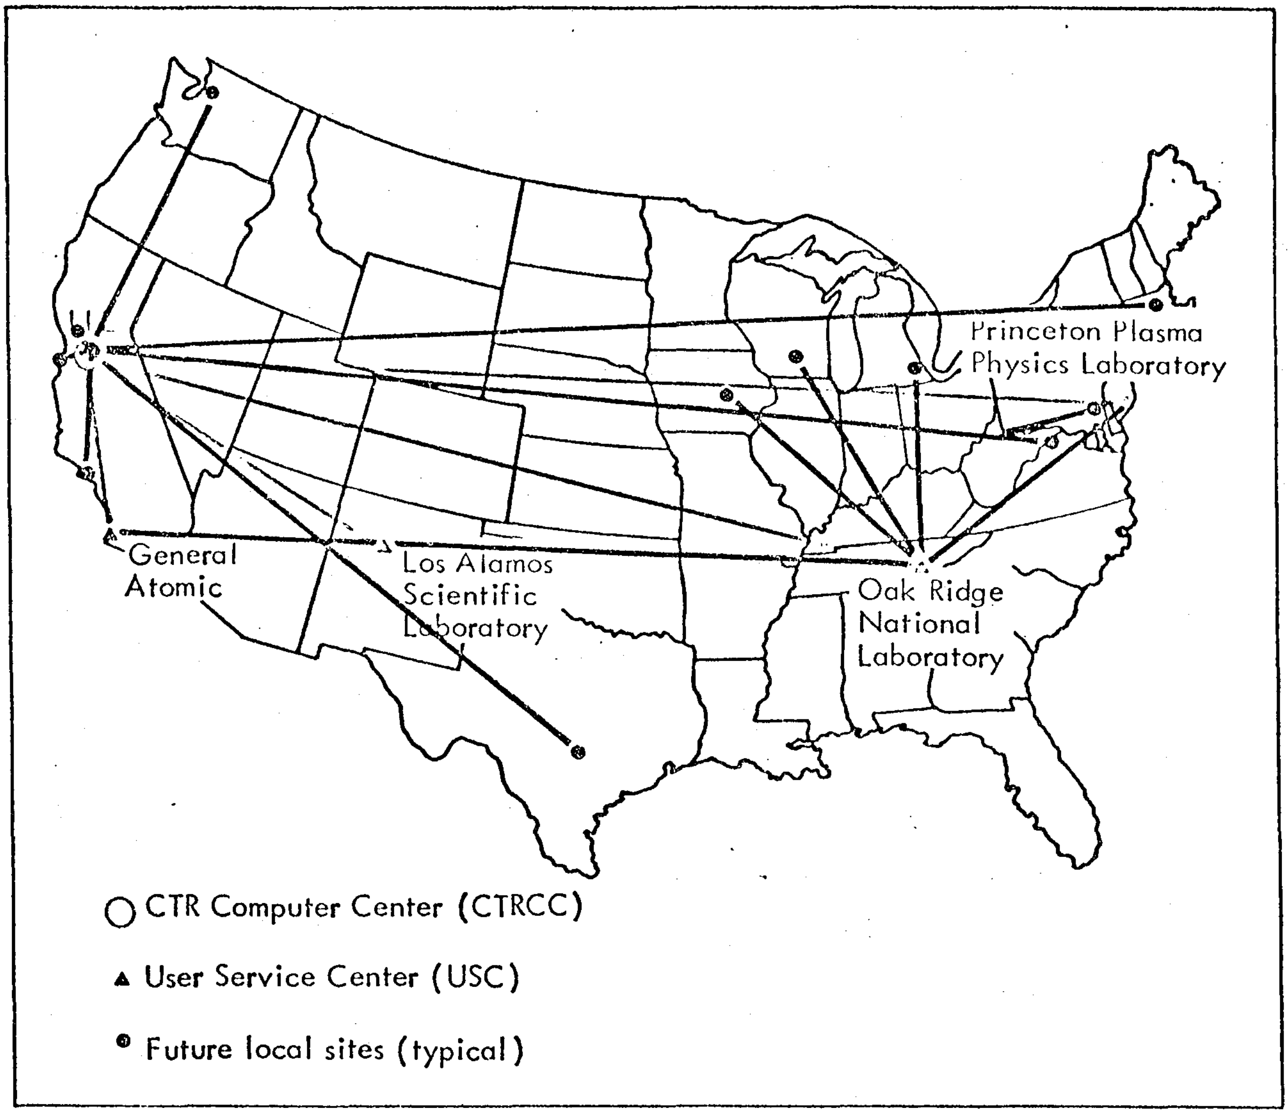
Princeton Plasma
Physics Laboratory
General
Atomic
Los Alamos
Scientific
Laboratory
Oak Ridge
National
Laboratory
CTR Computer Center (CTRCC)
User Service Center (USC)
Future local sites (typical)

Fig. 9

GATEWAYS

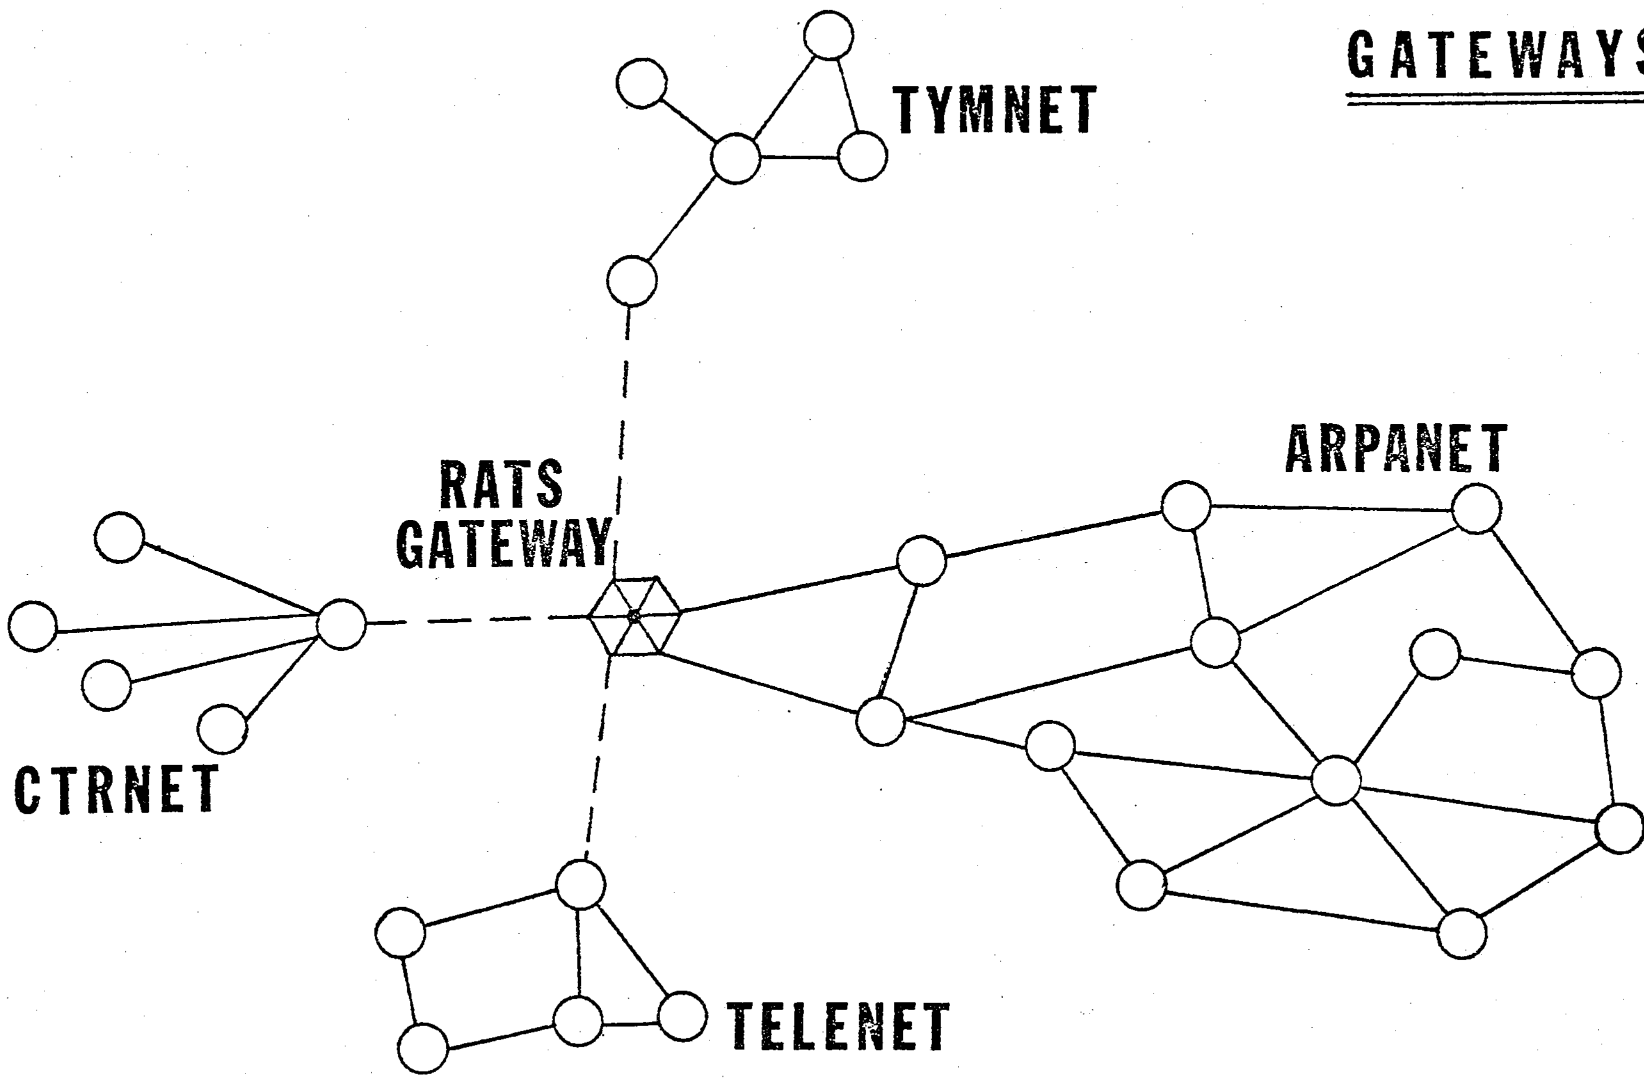

Fig. 10

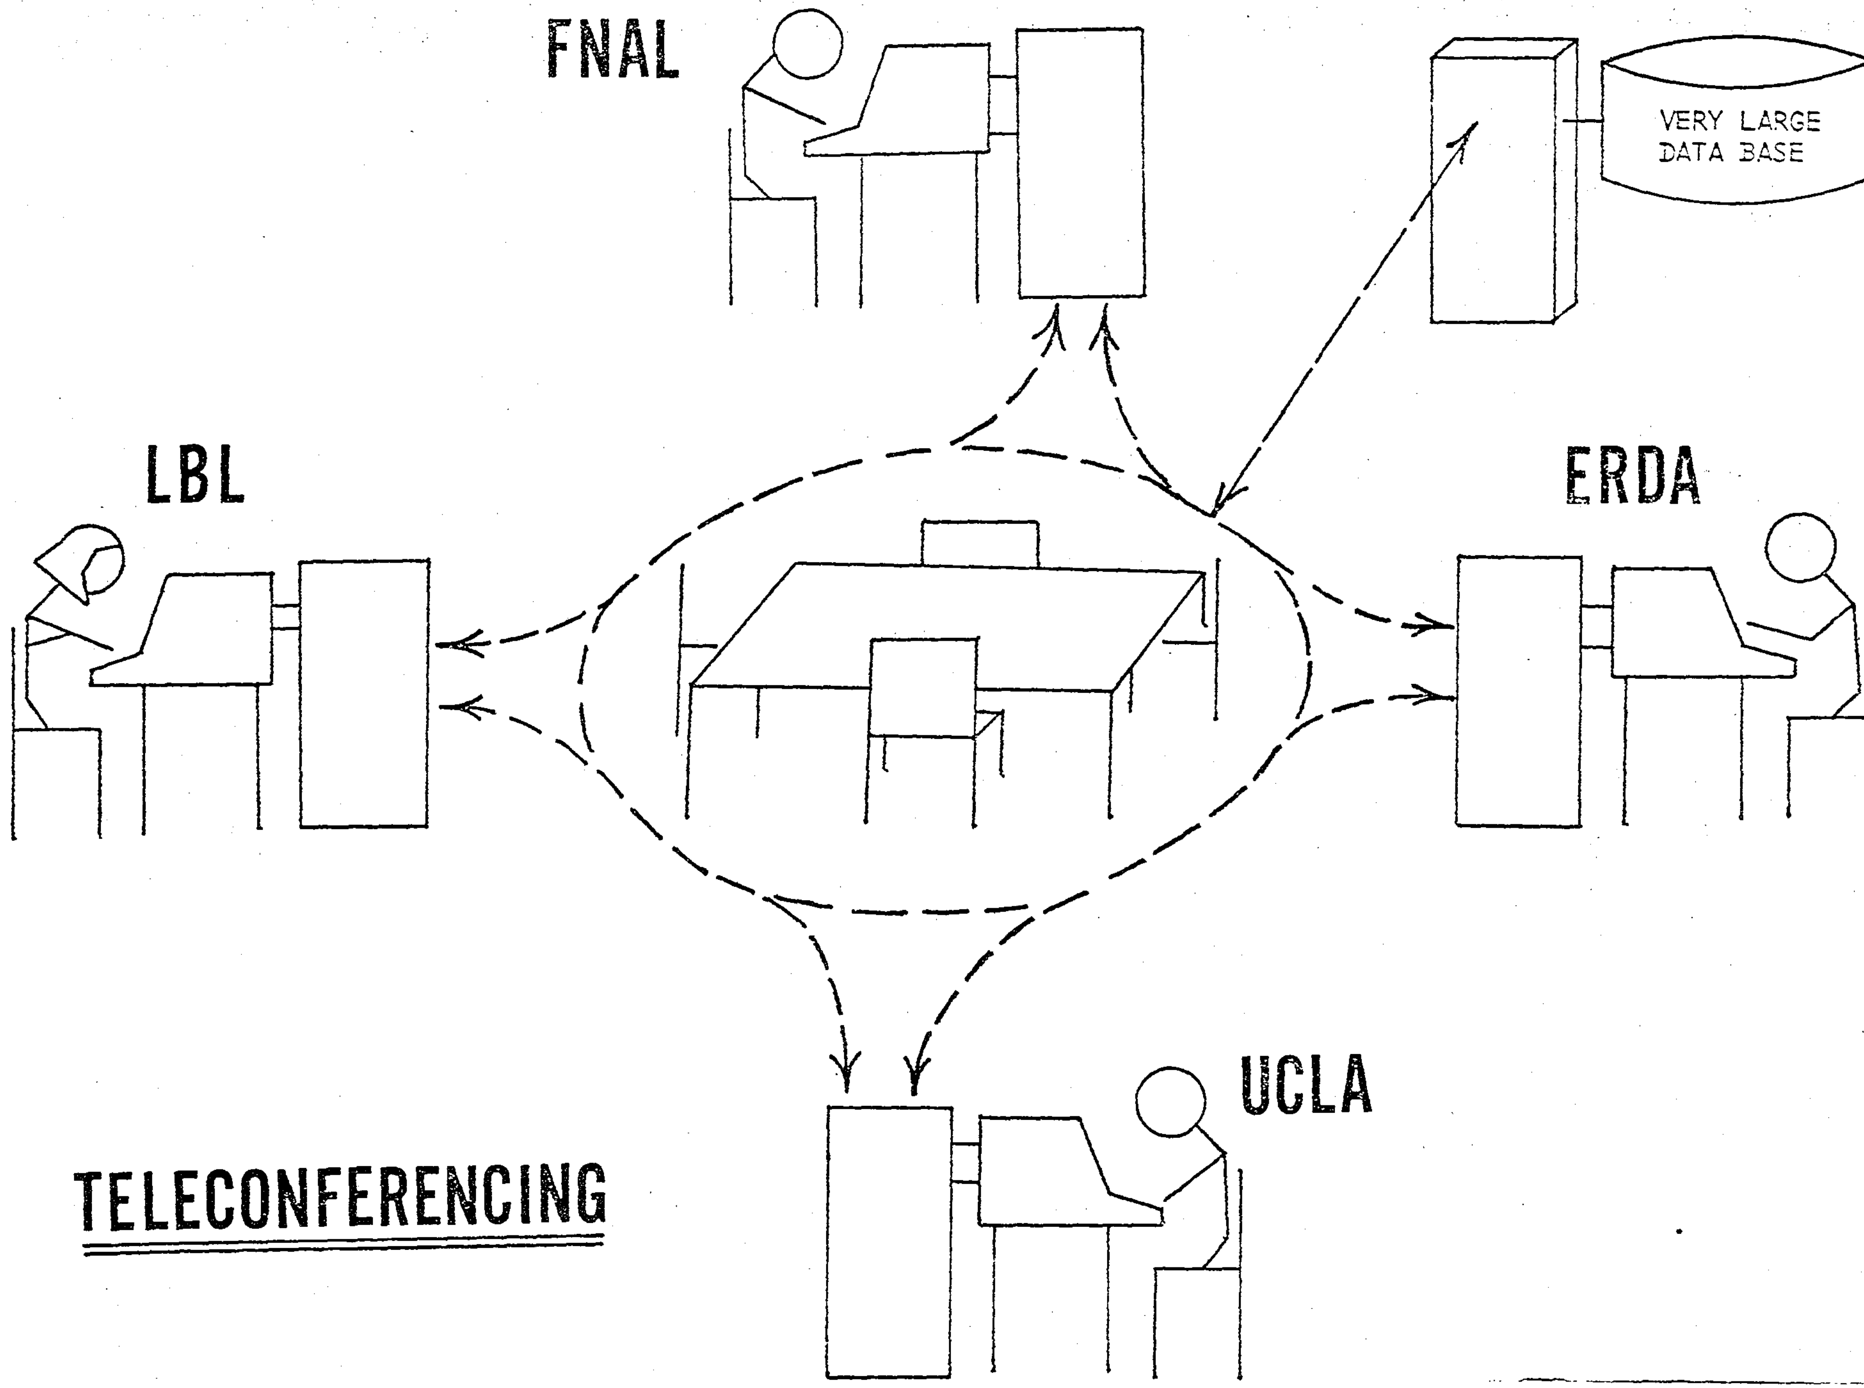

Fig. 11

COMPUTER SCIENCE AND APPLIED MATHEMATICS DEPT.
LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA