SLAC-PUB-3717
June 1985
(M)

Future Plans for HEP Computing in the U.S.*

J. Ballam
Stanford Linear Accelerator Center
Stanford University, Stanford Ca, 94305

Chairman of HEPAP Subpanel on Future Computing Needs

The computing requirements of the U.S. HEP Community are set forth. These will be dominated in the next five years by the $p\bar{p}$ (TEV I) and e^+e^- (SLC and CESR) experiments. The ensuing period will be almost completely driven by the data generated by the superconducting super collider (SSC). Plans for near term computing are presented along with speculations for the SSC.

Brief descriptions of accelerator and theoretical physics plans are also presented.

I. Introduction

Over the past decade the computer has become an increasingly important tool in the field of particle physics. By now it should be considered as a full partner in the triumvirate of accelerator, detector and computer. This is an extremely important development not only for the physicist who designs and uses the apparatus, but for the computer practitioner as well. The immediate consequence is that the nature and extent of the computing demand must be evaluated at the same time as all other aspects of a new accelerator or detector are being considered. This is in sharp contrast to computer planning in previous decades where the impact of computer was suddenly faced after, in most cases, the physics data itself was beginning to emerge from the experiment.

The contrast between then and now is perhaps best illustrated, at least in the U.S. by comparing the Fermilab in its early history, with the efforts now underway in preparing for the proposed Superconducting Super Collider (SSC). The proposal for the 400 GeV proton synchrotron, finally located in Batavia Illinois, did not mention any computing facilities which its construction would eventually cause to surface. Subsequent summer workshops dedicated to investigations of detector technique and possible detectors themselves also did not take into consideration the computing resources needed to realize the physics output from these devices. The laboratory was thus established without a sufficient computing facility which only subsequently was brought up to a reasonable level. In this case the consequence of

---

* Work supported by the Department of Energy, contract DE-AC03-76SF00515.

Presented at the Conference on Computing in High Energy Physics
Amsterdam, The Netherlands, June 25-28, 1985

placing computers last was not devastating because both machine design and detectors were nowhere near at the level of complexity as are the colliding beam facilities of the 1980's.

These reasons, namely the size of the collider itself plus the generation of huge amounts of data from the very complicated and highly granular detectors that are needed, have forced very early consideration of computing needs into the SSC design proposal and attendant workshops. So much so that there is now a pervasive early realization that computing costs, manpower, research and development especially in the early triggering process, have all to be estimated from the beginning for the computer member of the triumvirate. Special workshops have to be set up to study these problems along with the more conventional detector and physics workshops. Consequently it is hoped that before construction starts on the SSC, both the scientific community and the government will have a very good idea of what central computer is needed at the new laboratory, computers needed at each large detector, computers needed at remote locations (universities and other laboratories) and the networking required to link all of these places together.

Along with this burgeoning of computing needs in experimental particle physics has come a corresponding growth in both use and projected need in particle theory. This has been sparked by the proponents of lattice gauge theory, calculations for which can use an enormous amount of computer time and as a result have driven some theorists into designing and building their own processors. Although it is still a bit early to state with certainty, it is reasonable to predict that as more theoretical physicists become familiar with super computers, array processors and special purpose processors, the role played by computers in theoretical physics will become as pervasive as it now is in experimental and accelerator high energy physics.

Finally I would say that there are two new challenges for the physicists in obtaining adequate computing for HEP in the U.S. First to convince the computer industry to play a more active role in scientific computing in general and in particular to help us solve our problems and achieve our goals. Second for ourselves to provide a sensible and useful filter for selecting the new physics from among the incredible amount of data that will be generated by the future $p\bar{p}$ and pp colliders in the U.S.

II. Experimental HEP Physics Needs

I will discuss this matter in two time frames which follow closely the U.S. planning for HEP. The near (next five years) term takes into account the effect of the two new accelerators, the TEV I $p\bar{p}$ ring at Fermilab and the SLAC Linear Collider (SLC) e^+e^- machine as well as the CESR upgrade. During the next five years all machines should be in full operation with all five detectors checked out and in a regular data taking mode. These play the dominant role in computer planning, although data generated at other U.S. accelerators namely TEV

II (fixed target 800 GeV experiments at Fermilab) the AGS at BNL and PEP at SLAC will continue to require considerable computing support.

I note it has become fashionable in the U.S. to state computing hardware requirements in VAX 11/780 units and I will in general use these in the course of this report. Table 1 shows VAX 11/780 equivalents for various computers.

A summary of the TEV I situation is:

- Machine - $p\bar{p}$ collider 2 TeV in CM.

Luminosity	$10^{30}\ cm^{-2} sec^{-1}$.
Instantaneous Event Rate	$10^5\ sec^{-1}$.
Average Event Rate	$\sim 10^4\ sec^{-1}$.
Recorded Event Rate	$1 - 2\ sec^{-1}$.

- Detectors
    - CDF=(Collider Detector Facility)
      a general facility
    - D0 -A Calorimetric Detector.
- Schedule:

Machine operation begins	1986
Initial Data with CDF	1986
Substantial Data with CDF	1987
Initial Data with D0	1988
Substantial Data with both Detectors	1989

- OFFLINE Computing Requirements for CDF

Program size	4 Megabytes
Event size	0.15 Megabytes
Raw Data Tape Generation	10,000/year (6250 p.b.i.)
Total Data Tape Generation	40,000/year
Disk Space Needed	
Rapid Access Data Base Storage	6 Gigabytes
Permanent and Temp. Storage	20 Gigabytes
Overall Mainframe Needs	
Event Reconstruction	50 VAX 11/780 Equiv.
Other, including Data Analysis	50 VAX 11/780 Equiv.

- OFFLINE Computing Requirement for D0

Program Size	3 Megabytes
Event Size	0.1 Megabytes
Raw Data Tape Generation	5000/year
Total Data Tape Generation	20,000/year
Disk Space Needed	
Rapid Access	5 Gigabytes
Permanent and Temp. Storage	10 Gigabytes
Overall Mainframe Needs	
Event Reconstruction	40 VAX 11/780 Equiv.
Other, including Data Analysis	40 VAX 11/780 Equiv.

A summary of the SLC Situation is:

• Machine: e^+e^- Collider	100 GeV in CM
Luminosity	$5 \times 10^{30} cm^{-2}\ sec^{-1}$
Instantaneous Event Rate	0.3 Hz (at the Z)
Average Event Rate	0.15 Hz (at the Z)
Instantaneous Recording Rate	1-2 Hz
Trigger Rate	1-2 Hz

- Detectors

Mk II Upgrade	First Generation Detector
SLD	Second Generation Detector with extensive calorimetry

- Schedule

Machine Operation Begins	1986
Initial Data with Mk II	1987
Substantial Data with Mk II	1987-89
Initial Data with SLD	1988-89
Substantial Data with SLD	1989-

- OFFLINE Computing Requirements for MK II

Program size	3.5 Megabytes
Event size	30 K bytes
Raw Data Tape Generation	800/year (6250 p.b.i.)
Total Data Tape Generation	1200/year

Disk Space Needed	
Rapid Access	500 Megabytes
Permanent and Temp. Storage	5 Gigabytes
Overall Mainframe Needs	
Event Reconstruction	12 VAX 11/780 Equiv.
Other, including Data Analysis	10 VAX 11/780 Equiv.

- OFFLINE Computing Requirement for SLD

Program Size	4 Megabytes
Event Size	50 Kbytes
Raw Data Tapes	2000/year (6250 b.p.i.)
Total Data Tape Generation	2400/year (6250 b.p.i.)

Disk Space Needed	
Rapid Access	1 Gigabyte
Permanent and Temp. Storage	10 Gigabytes

Overall Mainframe Needs	
Event Reconstruction	15 VAX 11/780
Other, including Data Analysis	15 VAX 11/780

Proposed Solutions for the TEV I Detectors.

- CDF - At the present time the computing for this detector is assumed to be done in three stages:

  Stage 1 (1985-86 the first year of operation) - Monte Carlo equipment simulation, track reconstruction and the analysis of a small number of events via the scheme shown in Fig. 1, with the 11/780 replaced by 11/8600's. Two more 8600's will be used for data logging and online analysis.

  Stage 2 (1986-7) - Access to a new mainframe which is now being ordered. No decision has been made on which machine will be selected, but it will certainly be the size of the IBM 3090/200 (45 MIPS) perhaps of 70 VAX 11/780 equiv.

  Stage 3 (1987) - Extensive use of farms of processors. Two are being considered - one developed by the Advanced Computer Project of Fermilab, the other using micro VAX's.

Most physics analysis is planned to be done on the large mainframe. This growth of computer usage coupled with equipment is shown in Figure 2.

- D0 - Plans for offline computing are not definite but they assume that some combination of farms of micro VAX's and a large mainframe central computer will satisfy their requirements. Thus the approach is very similar to CDF. Because of the longer time scale, the decision as to the exact kind of farm to use can be delayed until late 1986.

  Online filtering, to reduce a 200 Hz rate to a 1 Hz rate is proposed to be done by 50 micro VAX's in parallel, each equipped with 3 Megabytes of general memory and 8×64 Kbytes of dual port memory. These will be connected via ethernet to a host (probably a VAX 11/8600). The hope is that early experience with this system will help determine whether it is suitable for some offline event reconstruction.

Proposed Solution for the SLC Detectors.

- Mk II & SLD assume that all production (event reconstruction) will be by a central computer. Present plans are that this will be an IBM 3090/200. Future expansion will probably be with farms of microprocessors. In the SLC machine the real event rate (excluding background events) is a factor of 10 lower than at the TEV I, so that solution is not so heavily dependent on these farms - at least initially.

Networking for Experimental Physics in the U.S.

Existing networks and dedicated lines used are shown in Figures 3, 4, and 5, and a detail list of active network users is shown in Table 2.

As can be readily seen this is a hodgepodge which has no central control or maintenance, BITNET and DECNET are to some extent maintained by the firms, but gateways and cross-connects are not transparent and not very reliable.

A recommendation for a high energy physics network, called not surprisingly, HEPNET has been made. Figure 6 shows stage I of this proposal. Not shown are highly desired 56Kb lines linking the U.S. to Europe and Japan. A desirable goal of this conference would be to recommend a method for establishing these dedicated lines. Important features of HEPNET are central financing of the trunk lines and individual financing of the feeder lines. A full-time staff (may eventually be one person) devoted to maintaining, improving and planning for extension of HEPNET is an essential part of the recommendation. This proposal is under active preparation at this time and, if all goes well, could be in operation in 1986.

Stage I of HEPNET should be adequate for the near term, and experience gained from its use can be embodied in future expansions and modifications, including gateways to other networks.

A summary of the SSC Situation is:

- Machine

pp Collider	40 TeV in CM
Luminosity	$10^{33} cm^{-2} sec^{-1}$
Instantaneous Event Rate	$10^{8} sec^{-1}$
Average Event Rate	$10^{7} sec^{-1}$
Recorded Event Rate	$1 - 2 sec^{-1}$

- Detectors

One for each of the six
interaction regions

- Schedule (Proposed)

R & D	1985-1988
Construction	1988-1994
Preliminary Run	1995
Substantial Data Taking	1996

Typical Numbers* for one Detector Design

Presented at the SSC Workshop, July 1984

Program size	28 Megabytes
Event size	1 Megabyte
Raw Data Tape Generation	50,000/year (6250 p.b.i.)
Total Data Tape Generation	150,000/year (6250 p.b.i.)
Disk Space Needed	
Rapid Access	$\sim$ 20 Gigabytes
Permanent and Temp. Storage	$\sim$ 40 Gigabytes
Overall Mainframe Needs**	
Event Reconstruction	350 VAX 11/780 Equiv.
Other, including Data Analysis	350 VAX 11/780 Equiv.

---

* Most of these are scaled from CDF by the event size ratio of 7/1. The increase is due to the product of the multiplicity of the event times the granularity of the detector

** The Snowmass Study estimated a total of 1000 VAX 11/780 Equiv.

Proposed Solutions for the SSC.

At the present time there is no acceptable overall solution to this enormous problem.

There are several avenues which will be explored over the next five years at the end of which time, if the present schedule is maintained, final decisions on computer specifications will have to be made. These are not meant to be comprehensive, but rather suggestive.

1. Use of very large farms of processors/emulators on the order of 1000/experiment. This system may be made to work since it is assumed that the ratio of CPU time to I/O time is very large. This method would be used for event reconstruction with one event assigned to each processor.
2. Conversion of large ($> 75\%$) of event reconstruction programs to run on a super-vector computer. It is estimated that factors of 2-3 in processing time could be realized for an efficiently written program.
3. The commercial development of multi parallel processors with adequate software support so that the subsequent physics analysis may be run. Then machines would required to run at 2000 to 4000 m.i.p.s. with reasonably short CPU cycle times [$\leq$ 20ns].

   Table 3 shows a recent comparison of various commercial computers, some of which exist and others which have in one way or another been announced as "forthcoming" machines. As can be seen, there is no candidate which could now be taken as satisfying these requirements. The table does show that various companies are moving in the right direction. What is not shown is what is going on in the R and D laboratories in the commercial scene.
4. New forms of computers. Several university groups in the U.S. are working on new architecture. While most of these are directed towards the solution of theoretical particle physics problems, some seem to be of general interest. One example is the hypercube processor being developed at Caltech, but with some interest being shown by industry. The claim is that one processor with 1024 nodes can analyze events each requiring 4 megabytes of program plus event size, in a time taken by 100 VAX 11/780's, so that 4 of those would handle the projected data analysis needs of one of the SSC detectors.

Workshops on computing for triggering, data acquisition, data reduction and data analysis will be organized with the specific goals of making recommendations regarding the online and offline computing equipment needed for the SSC including costs estimates and a schedule for implementation.

Further studies involving specific proposals to modify existing software codes such as EGS or track reconstruction to run on vector machines are to be encouraged.

> Continued and/or new support should be provided to produce multiprocessor systems with the equivalent capacity of 1000-2000 VAX 11/780's by the early 1990's. Cooperation with industry is strongly encouraged.

Nothing has been said heretofore about the most difficult aspect of SSC computing, namely how does one reduce the 10^8/sec real event rate down to a 1-2/sec rate in what the HEP physicist calls a "model independent" way? Since there does not appear to be an answer to this question at the present time it is not logical to include it in a list of "future plans". However since it is considered by many to be at least as difficult a challenge as the proper design of other parts of particle detectors, some discussion of its general features is in order.

First there is general agreement amongst particle physicists that it is not necessary to record every event. A large part of the cross section is assumed to go into quasi-elastic scattering with almost all produced particles continuing on down the beam pipe itself. This however would only remove one-half of the events from consideration. The next largest cross-section of interest may be pair production of B mesons. Some cross-section estimates go as high as 200 μb, leading to an event production of 2×10^5/second. Since there is no known way to record and analyze such a huge rate, and since these events will probably have been produced copiously at LEP and SLC [decay of Z^0] the next step would be to see if there is $B\bar{B}$ mixing and CP violation. This becomes difficult at pp machines unless one goes to exotic decays such as semi leptonic-decays of B $\rightarrow 4K_s$ and $\bar{B} \rightarrow 4K_s$ but these rates are small and are estimated realistically as 1 event every 10^3 seconds. The search for Higgs Bosons is illustrated by considering the process pp $\rightarrow H + X$ where $H \rightarrow t\bar{t}$. For masses of H around 300 GeV the cross section is estimated as 0.5 pb giving an event rate of again 1/1000 seconds where as the rate to 100 MeV Higgs is 1000 times greater.

The problem in all of the above is that if one gets, by virtue of very severe kinematic restraints, a reasonable rate for a given process, other equally or perhaps more important processes may be discarded or distorted.

Obviously the faster that the initial decision to trigger can be made, the better, provided that the decision is calculated by a rather sophisticated process. Therein lies the challenge. No definitive criteria can be set by the physics requirements because basically they are not known. In this case the physicist has to be led by technical developments, before "plans" can be set down.

A summary of the expected growth of computing at National Laboratories is shown in Figure 7.

III. Theoretical Computing Needs

The theoretical community use of large scale computing is at present driven by a widespread interest in lattice gauge theory calculations. Using four dimensional space time the practitioners of this art think in terms of lattices having 16^4 sites. The technique is statistical and the accuracy improves only as the square root of the running time. These lattices already require the full memory of the current generation of supercomputers. Simulations, using already existing algorithms to calculate with quantized quark fields will require thousands of hours on state-of-the-art vector machines.

Another field, which while not currently as dependent on huge machines, but which has enormous potential is that of perturbative studies of quantum field theory. This requires computer assistance both for extensive algebraic manipulation and for numerical evaluation of Feynman integrals. If algorithms can be devised for automatically renormalizing divergent Feynman diagrams dramatic increases in the currently modest amount of computer time being used in these areas will take place. At this time it is not clear whether the greatest impact of this application will be in the area of work stations, conventional mainframes or specially configured "artificial intelligence machines". This is an area to be closely watched by both government and certain parts of industry. There is a developing connection between particle physics theory and theories in other fields. Monte Carlo programs developed for lattice gauge calculations have been adapted from algorithms long used to study statistical mechanical systems. Particle physicists and condensed matter physicists have begun to appreciate the deep connection between critical phenomena and quantum field theory. Nuclear theorists have been attracted to lattice studies of the transition to a quark gluon plasma which could be created in high energy heavy ion collisions. Particle theorists have taken an active interest in computer science and are contributing to it.

This is probably a good place to point out that, while this conference is devoted mainly to computing in HEP, there will be a close connection between this field and high energy heavy ion collision research, both theory and experiment. In the U.S. there are active proposals for new machines and experiments which will be very prominently manned by HEP experimental physicists.

To solve the near term needs of the U.S. theory community it was recommended that the equivalent of two supercomputers of the Class VI capacity [CRAX XMP or CDC CYBER 200 e.g.] are needed. It is also recommended that wide access to these machines be made available to theorists, beginning with an initial allocation of two hours of a class VI machine to any theorist who requests it.

Although supercomputers will continue to be a major source of computer cycles for theoretical physics, alternative approaches have considerable potential for making large-scale computational resources more widely available. For example, the tremendous promise of specially designed processors has already been confirmed in applications to experimental physics and is under exploration for theoretical applications. The most advanced initiative is the machine being developed by N. Christ and A. Terrano at Columbia University, in which special hardware carries out the complex matrix multiplication needed in simulations of SU(3) gauge theory. Other projects at more preliminary stages are underway at MIT, and IBM. One drawback of the home built machine approach is the long lead time for results in comparison with the usually rapid change in theoretical fashion. This lead time should, however, be expected to decrease with experience and the availability of more sophisticated devices from which to fashion a special purpose machine. A more general purpose machine, under development by a collaboration of high energy theorists and computer scientists at Caltech, has already done important lattice gauge theory calculations. Commercially available systems based on the Caltech hypercube have recently been introduced. Farms of microprocessors being developed for track reconstruction may also be suited, with minor modifications, to lattice simulations.

The use of attached array processors as an alternative to vector mainframes is currently a promising economical route to large scale computations. Previously this approach has required considerable dedication of the user in order to overcome programming difficulties due to software shortages. This is becoming less of a problem for users as the needed software is developed. Ambitious projects at Cornell University and Argonne National Laboratory are attempting to show the viability of networking several such devices together in parallel systems. In addition to provide a testing ground for multiprocessing software, the complexes under discussion would, if successful, provide enormous potential computational power.

A synopsis of various activities in this field is given below:

- The Cosmic Cube. Using commercially available microprocessors, with emphasis on simple communications between them, scientists at Caltech have been using a system of 64 processors to study how to formulate algorithms in the multiple processor environment. Having been quite successful for most types of theoretical problems, they are now planning a system of 1024 processors. A version of this processor is under commercial development by a number of different firms.
- The Columbia University project. This project uses commercially available processors but with added hardware to do floating point complex matrix multiplications for QCD lattice gauge calculations at high speed.
- The ANL Multiple Array Processor Project.
  This project plans to ultimately configure 64 Star Technology's ST-100 array processors, a 32-bit 100 mflop machine, with specially constructed shared memory optimized for statistical mechanics and lattice gauge theory calculations.
- FASTBUS VAX Project. NYCB Real-Time Computing is integrating the microVAXII VLSI chip set with a FASTBUS interface into a FASTBUS module. This project is being funded by the DoE Small Business Innovation Research program.
- IBM Project. This presently consists of ten FPS-164 array processors attached to an IBM 4341 host. The host loads each FPS processor with the same program but different data sets. The processors execute independently and return data to the host where it is consolidated.

IV. Accelerator Physics Computing Needs

Two important categories of extensive computer use by accelerator physics have surfaced over the past few years. One is for the design of the SSC and the other for modeling of newly proposed methods of acceleration. In the SSC case there are two components, the first of which has to do with calculation of the dynamic aperture of the machine - a crucial component in the design. These tracking calculations can be very time consuming. For example a single tracking experiment requires a number of particle turns calculated as follows.

$$\begin{aligned} N \;=\;& \text{Number of particles with different transverse initial conditions (10)} \\ &\times\ \text{Number of different synchrotron oscillation amplitudes (3)} \\ &\times\ \text{Number of synchrotron oscillation periods (100)} \\ &\times\ \text{Turns per synch. oscillation period (1000)} \end{aligned}$$

From which $N=3\times10^6$

A typical supercomputer time for one particle turn is 0.091 sec so that the single experiment takes 76 hours. This must be done for different lattice configurations and magnet designs and several times for each lattice. The SSC has random multipoles and these affect the dynamic aperture. Therefore as many as 200 such experiments may have to be done leading to 15000 hours on a supercomputer of the class VI variety. Once the final lattice is chosen these calculations will have to be repeated a number of times depending on how close the locations of magnets and the quality of their magnetic fields come to the numbers put into the original calculations. It also assumed that computers of this class will be needed to design and interpret acceleration experiments on the SSC once it is running.

The second component of the SSC computing is the more mundane task of system and coordinating design. Logistics of construction, storage and eventual installation of thousands of magnets are one example. Another is the design, construction and installation of the cabling plant to say nothing of instrumentation and control. In this sense a computer of the type used in the commercial and industrial world is needed.

The other category, namely modeling of new acceleration schemes, is now beginning to become a significant part of new accelerator design. Experimental work in this field can be very expensive and time consuming especially if conducted in a completely empirical engineering mode. Modeling can eliminate many blind alley approaches and focus in on perhaps one or two schemes that have the greatest promise. If the experimental work begins at this point a great deal of unnecessary work can be avoided. The more extensive the modeling, the more precise the experimental questions can be formulated.

Examples of extensive computer use are in laser plasma accelerator theory where a Class VI machine can on occasion easily be saturated by particle simulation in 2 and 3 dimensions. The work on free electron and two beam accelerators is also being aided by particle simulation on Class VI machines but has not yet reached their full exploitation.

It seems clear that in the U.S. at least two Class VI machines. Each could be fully used by accelerator theory and experiment modeling. If a center for theory is set up around one or more of these processors it is possible that these could be incorporated - or vice versa.

V. Special Topics

Multiple Processors - The following is a list of processors projects within the HEP community which are either in use or far along in planning and implementation. The list is not comprehensive and omits those processors which are primarily for theoretical work and which have already been mentioned previously.

- The 168/E emulators. Already mentioned in the above sections, this processor has been the most used to date since it was started much earlier than others. The processor is equivalent to 1.3-1.5 VAX 11/780s. About 50 processors have been built and used at SLAC, CERN, DESY, Saclay, and University of Toronto for both online and offline processing.
- The 3081/E emulators. A SLAC/CERN collaboration is developing the 3081/E emulator as replacement for the 168/E. Having much more memory space and equivalent to 4-5 VAX 11/780s, the processor is also designed to be much easier to debug and maintain. The 3081/E will be used for both online and offline applications at SLAC, CERN, and Saclay.
- The 370E emulator. Developed at the Weizmann Institute, this processor is much closer to a real computer than the other emulators, thus allowing it to do more easily formatted I/O. Systems are now operating at Weizmann, DESY, Rutherford Laboratory, and Cornell. The processor is 2-3 VAX 11/780 equivalents.
- The ACP project. This Fermilab project emphasizes the use of commercially available microprocessors which are 0.5 to 1 VAX 11/780 equivalent in processing power. A prototype of a 6 processor system exists today with extensive software support. A one hundred VAX 11/780 equivalent farm is planned for implementation in the 1985 calendar year.
- The D0 MicroVAX farm. The D0 detector is planning a large farm of MicroVAXes bought from DEC as the final stage of filtering online. Both DEC supplied hardware and software are used. A small prototype system is already operational at Brown University.

Average cost and manpower estimates for these farms of array processors are given below. For present day systems such as would be used on CDF [in units of VAX 11/780 processing capacity]

Cost per unit/not including host	5-7 K dollars
Manpower to build a farm	20-30 man years
Manpower to interface with a host computer	4-6 man years

For future systems,* such as the 1000 11/780 VAX equivalent for an SSC detector, assuming cost reductions over the next 5-10 years.

Cost per unit 3K dollars	
Total cost	3 M dollars
Host computer and assoc. hardware	1 M dollars
Manpower	60 man years

Mainframes

The greatest increase in raw computing power on the commercial market today lies with the vector architecture supercomputers made by CRAY, CDC, Hitachi and Fujitsu some of which come with as much as 64 Megabytes of memory. It is not clear at this point whether this raw horsepower can be harnessed to the wagon of experimental high energy physics although, as has been pointed out, they are very attractive to users in the other branches of HEP. This is so because experimental data taken from a large detector has a high degree of natural parallelism while the others can be put in a pipeline format. Experiment events have the following characteristics:

- The event structure allows the processing of many events in parallel. This has been discussed already.
- Within a given event many similar calculations occur frequently (e.g. clustering of energies in identical cells of a calorimeter).
- Once tracks are found, track fitting of tracks in the same event is a repetitive procedure.

---

* Of course these must be considered only as estimates but, if achievable, show how dramatic the effect these farms will be on the cost of computing on the SSC era.

It is possible that the last two categories can be vectorized, but real increases in speed can be achieved only when these processes previously constituted $\bar{>}$ 75% of the corresponding scaler program.

However such attempts at Fermilab and at other laboratories as well have not been successful, and enough work has been done to know that this is not an easy task.

The ultimate gain from being able to vectorize 80% of a program is a factor of five, even if the vector hardware is very much faster than the scalar machine. In the case of the SSC a factor of three (a practical goal) is worth considerable effort to obtain and studies of what can be done, beginning from ground zero, with experimental code have begun in the U.S. It has been suggested that vector processors be added to U.S. laboratories - even as front ends to scalar machines in order to provide the hardware for these studies.

Parallel processors of large fraction of giga instructions per second are on the horizon, some purported to arrive in a year or two. There is however little experience with the present day machines. Lack of adequate software support seems to be the greatest drawback. In the tens of millions of instructions per second, there are many working machines, but these are suitable only for near term solutions. Table 3 is a listing of machines that are either on the market, already announced, or rumored to be coming. As can be seen the SSC solution is not yet firmly at hand.

Productivity factors - although raw power is of great importance, as has already been demonstrated, it is interesting to list a number of factors which lead to greater productivity for a machine of given power. Many of these lie in the area of software support which becomes of even greater importance for these future experiments involving hundreds of physicists spread over the world. Some of these are enumerated below:

- Interactive response times should be small and appropriate to the resources required to satisfy a command.
- Large address space available to the user's processes. (Many megawords, perhaps through virtual memory.)
- High quality full screen editing with multiple concurrent files, greater than 40 lines on the screen, and macro capability.
- A modern file system, including hierarchical directories, specialized access methods, and date stamps for access and revision. The system should provide adequate file space for each user.
- A symbolic interactive debugger, including display and the capability to step through source code.
- Flexible command language.
- Extended Fortran compiler (extended beyond Fortran 77), preferably with "standardized" extensions to enhance transportability.

- "Batch" or background capability, with appropriate operations and management tools.
- Networking and electronic mail.
- Graphics capability.
- A broad spectrum of languages.
- A common calling sequence for mixed language programs.
- Remote job execution facilities.
- Easy redirection of I/O.
- Local interprocess communication services.
- Advanced multiprocess services for the individual.
- Advanced applications services, such as parsers, tree builders, queue managers, and real time system hooks. There should be user accessible system building tools, e.g., menu capability and other advanced terminal handling routines for interactive systems.
- A full set of utilities including system management tools (e.g., monitoring of critical computing resources), and user management tools (e.g., individual process monitoring or code management systems for the development of large programs).
- Hardware protection of the individual user from himself (e.g., write protection of code and stable data).
- Advanced scientific document preparation.
- User accessible execution monitor (Proglook).
- Database management system with access from high level languages.
- Good HELP facility with easy user additions to the facility.
- Hierarchical associations of users.
- Spreadsheets (as an integrated utility).
- Screen editing with low bandwidth features for editing over voice grade telephone connections.

Software Manpower

One should not neglect the area of people - both physicists and programmers needed to write the software to run these large detectors and to analyze the data. Table 4 lists, for various U.S. detectors, estimates of the full time equivalent per year needed for at least three of four years of design, construction and initial operation. The extrapolation to a SSC detector is made by assuming a) the number of such systems will increase b) segmentation will increase

c) the selective trigger will require more software d) the data base and code management will be more extensive. Experience with oncoming detectors in the U.S. and elsewhere will of course provide more accurate estimates, but it is clearly a very large enterprise, especially if four or more detectors are contemplated for the SSC. Much more consideration, at least in the U.S., will have to be given to the use of professional programmers than in the past. A similar hierarchy to that of physicist, engineer and technician will probably evolve resulting in physicist, programmer-mathematician, and programmer. At the moment the trend in the U.S. is the conversion of a number of physicists to programmer mathematicians with the greatest number of programmers still in the laboratory computing center.

VI. Summary

- The needs of U.S. experimental HEP computing over the next five years very likely can be served by a combination of the latest 4–80 m.i.p.s. mainframes, coupled with a successful realization of medium sized farms of processors. A notable exception is the data logging both in speed and capacity, which now looms as a real limit to the experiment.
- The U.S. community has come to realize that the productivity component of large processors and their peripherals must be considered along with their raw computing speed.
- Theoretical and accelerator studies are done well on supercomputers (vector machines) and a constellation of four of these, perhaps located in a center, would satisfy most U.S. requirements. However, a great deal of encouragement is being given to groups developing special array processors, which for some purposes are claimed to do the work of tens of supercomputers.
- Networking in the U.S. will be in the form of a special HEPNET with gateways to other systems. No plans exist for extensive ultra high speed links; these await experience with the 56K lines. Linkage of HEPNET to Europe and Japan has been strongly recommended and is deemed necessary for ongoing and approved experiments.
- The period after the next five years will be dominated by the SSC which represents an enormous increase in demand. No solution is now at hand, although promising avenues exist. These must be pursued with the active participation of industry. The physics community should not devote its manpower and resources to designing and manufacturing 10,000 processors plus the software to run them. Parenthetically the same warning applies to future manufacture of special array processors now being developed in the theoretical community. Further, HEP may to rely on the emergence of new mainframes with

adequate software support and which are not yet on the market in order to do the fantastic amount of physics analysis that awaits us.

- Lastly, to repeat the statement at the beginning of this paper, computing in HEP has reached the stage, at least in the U.S., where it should be considered an equal partner with accelerators and detectors.

The author wishes to express his thanks to the Organizing Committee for its support.

TABLE 1

VAX 11/780 Equivalents

Computers[+]	VAX 11/780 Equivalents*	
		Totals
IBM 3081K[2]	21	
Siemens 7890S-MP[2]	48	
CDC Cyber 875 [2]	30	99
DEC 10 [2]	3	3
IBM 3084QX[4]	36	36
CDC Cyber 175-200[1]	9	
CDC Cyber 175-200[1]	9	
CDC Cyber 175-300[1]	10	
CDC Cyber 875[1]	17	45
Hitachi M200H[1]	12	
Hitachi M280D[1]	13	
Hitachi M280H[1]	15	40
IBM 3081K[2]	21	
IBM 3033[1]	8	29

*The normalization used is an approximation ascertained from an arbitrary set of high energy physics benchmarks (see Table V). Thus, all performance values are to be understood to be no better than ±25%.

[+] The number in [] gives the CPU count in each mainframe.

TABLE 2

High Energy Physics Computers on Wide-Area Networks

Compiled By Paul Kunz
Stanford Linear Accelerator Center
<PFKEB at SLACVM>

The following is a survey of computers used by the High Energy Physics community that are attached to wide-area networks. Only computers that support at least mail transfer are considered. Local networks, such as CERNET or DESYNET, are not considered, nor are computers that support only remote logon.

The following networks were considered in the survey:

BITNET	- RSCS network of American university computer centers
EARNET	- RSCS network of European Research Centers.
NETNORTH	- RSCS network of Canadian university computer centers (BITNET, EARNET, NETNORTH are physically one network with different names)
DECNET	- SLAC/LBL DECNET
JANET	- Coloured Books X.25 based network of the U.K.
INFNET	- INFN (Italy) DECNET
PSSN	- Public Packet Switching Networks, e.g. TymNet, TeleNet, only if file/mail transfer is supported.
ARPANET	- DoD network of research centers.
USENET	- UNIX network.

(Network-id in lower case means that connection is not yet made.)

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person,USERID
UNITED STATES				
Argonne:				
HEP	VAX 11/780	ANLHEP	ANL	
HEP	VAX 11/730		ANLCDF	
Physics	VAX 11/780	ANLPHY		
Center	IBM 3033	ANLOS		
Center	IBM 3033	ANLVM		
LBL & U.C. Berkeley:				
TPC VAX	VAX 11/780		TKYO	
Phys Dept	VAX 11/780	lblphys?	PHYS	
Phys Dept	VAX 11/ ?		ICS	
Phys Dept	VAX 11/730		TSTBED	
Phys Dept	VAX 11/730		ECCTST	
Phys Dept	VAX 11/730		MUTEST	
Brandeis University:				
HEP	VAX 11/780		BRND	
Brookhaven:				
Mailserve	PDP 11	BNL		ARPANET:BNL *Ron Reierls* PEIERLS
Phys	VAX		decnet	
Brown University:				
Center	IBM 3081D	BROWNVM		
HEP	VAX	BROWNHEP	decnet	Dave Cutts
CalTech:				
Phys VMS	VAX 11/780	CITHEX	CITHEX	PSSN:311021300219 Harvey Newman NEWMAN
Phys UNIX	VAX 11/780			USENET:CITHEP G. Fox
U. Cincinnati:				
Central	Amdahl V8	UCCCMVS		Brian Meadows (PPHYBZM)
Research	IBM 3033N	UCCCVM1		Brian Meadows (PPHYBZM)
U. Colorado:				
Phys Dept	VAX 11/780	bitnet		Uriel Nauenberg
Colorado State U.:				
Center	CDC 205	CDC205		(Used by Houston for SSC)
Columbia:				
Nevis	VAX 11/780	cunevis		Steve Smith
Cornell:				
Central	IBM 3081D	CORNELLC		
Central	IBM 4341	CORNELLA		
Wilson lab	VAX 11/780	CRNLNS		Ray Hemlke
Theory	VAX 11/780	CRNLTHRY		

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person,USERID
Fermilab:				
Admin.	IBM 4341	FNALVM		Jeff Mack (MAINT)
Front end	VAX 11/780	FNAL	FNAL	Greg Chartrand (GREG)
Front end	VAX 11/780	FNALA	FNALA	"
CD Dev.	VAX 11/780	fnalbsn	BISON	
CDF	VAX 11/780	FNALCDF	CDF	Chris Day (CTDAY)
CDF-Soft	VAX 11/780		CDFSFT	"
CDF-Wigwam	VAX 11/730		CDFHRD	"
CDF-Beam	VAX 11/730		CDFNW	"
CDF-Wedge	VAX 11/730		CDFCRT	"
ACP	VAX 11/780		BSNDBG	Tom Nash (NASH)
ACNET-Dev	VAX		DEVL	
ACNET-Op	VAX		OPER	
Univ. of Houston:				
Center	AS/9000	UHUPVM1		
Physics	VAX	bitnet		
Harvard University:				
HEP	VAX 11/780	HARVHEP	HUHEPL	
IBM Yorktown:				
Center	IBM 3081	YKTVMZ		Mark Bregman (MFB)
J. Hopkins:				
Phys Dept	VAX 11/780	JHUP	JHU (dialin)	Jack Serio (SYSTEM)
U. Ill-Urbana:				
Phys Dept	VAX 11/780	UIUCHEPG	decnet	Gladdings
Indiana U.:				
	VAX 11/780		IND	
Louisiana State Univ.				
Center	IBM 3081	NSNCC		
U.C. Los Angles:				
Phys Dept	VAX 11/780	uclaphys	UCLA	
U. Maryland:				
Center	Univac 1100	UMD2		
HEP	VAX	UMDHEP		
MIT:				
LNS	VAX 11/780	MITLNS		Robin Verdier, CSC
U. Michigan:				Ian Leedom, LEEDOM
Phys Dept	VAX 11/780		MICH	PPSN:3110-3130006266
	VAX 11/		MICH1	
	VAX 11/		MICH2	

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person,USERID
Northeastern Univ.:				
Phys Dept	VAX 11/750	NEUVMS		
U. Pennsylvania:				
Physics	IBM 3081D	PENNDRLS		
HEP	VAX 11/750	PENNHEP1		Richard VanBerg, RICK
Penn. State:				
Central	IBM 3081	PSUVM		
Purdue:				
	VAX		PURDUE	
Princeton University:				
Phys Dept	VAX UNIX			usenet ?
HEP	VAX			PPSN:?
U.C. - Riverside:				
	VAX 11/780		UCR	
	VAX		UCR2	
Rutgers:				
HEP	VAX	bitnet		
Rockefeller Univ:				
Center	VAX 11/780	ROCKVAX		
U.C. San Diego:				
Phys Dept	VAX		SDPH1	
U.C.- Santa Barbara:				
Phys Dept	VAX 11/780	SBHEP	UCSB	Steve Yellin (SPYP09)
Stanford University:				
Central	IBM 3081K	STANFORD		

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person, USERID
SLAC:				
Central	IBM 3081K	SLACVM		Les Cottrell (COTTRELL)
MAC	VAX 11/780	SLACMAC	MAC	Dave Wiser (DEWP06)
CB	VAX 11/780	SLACCB	CB	Didier Bisset (DIDIER)
MarkIII	VAX 11/780	SLACMK3	MK3	Dennis Wycinski
MarkII	VAX 11/780	SLACMKII	MKII	Andy Lankford
ASP	VAX 11/750	SLACASP	ASP	Gabor Bartha
TBF	VAX 11/750	SLACTBF	TBF	Dave Sherden (DJS)
SLD	VAX 11/780	SLACSLD	SLD	Dave Sherden (DJS)
End Stn A	VAX 11/780		ESA	Zen Salata
HRS	VAX 11/780	SLACHRS	HRS	Jim Schereth
TPC	VAX 11/780	slactpcs	TPCS	Ed Whipple
TPC	VAX 11/780		TPCT	Ed Whipple
Two Gamma	VAX 11/780	SLACTWGM	TWGM	Ed Miller
Two Gamma	VAX 11/750	SLACUCSD	UCSD	Ed Miller
Two Gamma	VAX 11/750	SLACNIKH	NIKHEF	Ed Miller
Two Gamma	VAX 11/780	SLACUCD	UCD	Dave Pellett
PEP C.R.	VAX 11/780	SLACPCR	PCR	Tony Gromme
SLC	VAX 11/780	SLACSLC	SLC	Nan Phiney
MCC	VAX 11/780		MCC	Nan Phiney
SPEAR	VAX 11/780		SPEAR	Mike Sullenberger
InterGraph	VAX 11/780	SLACCAD	CAD	John Steffani
StonyBrook:				
HEP	VAX			DECNET: to where?
Syracuse University:				
HEP	VAX 11/782	SUHEP		
U. Tennessee:				
Central	IBM 4341	UTKVM1		Jim Brau (PA87178)
Texas A&M:				
Phys Dept	VAX	TAMVXPHY		
Vanderbilt U.:				
Phys Dept	VAX 11/780	VANDVMS1		Frank Kyle (KYLE)
VPI:				
Center		VPIVAX3		Luke Mo
U. Wisconsin:				
PSL-A	VAX 11/780	WISCPSLA	PSLA	
PSL-B	VAX 11/780	WISCPSLB	PSLB	
Yale University:				
Center	IBM 4341	YALEVM		
Center	VAX	YALEVAX5		
HEP	VAX	YALEHEP		

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person, USERID

NON-U.S.

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person, USERID
Canada				
British Columbia:				
Center	Amdahl			MAILNET: UBCG-MTS%MICH-MTS B. White
McGill University:				
Center	IBM 4341	MCGILLA		
NRC - Ottawa:				
Center	IBM 3033	TSSNRC00		
Univ. Toronto:				
Phys Dept	VAX			PPSN:3020-91600097 R. Orr
TRIUMF:				
Center	VAX			PPSN:3020-83200013 C. Kost
Univ. of Victoria:				
Center		bitnet		
HEP	VAX	bitnet		PPSN:3020-68100058 Lyle Robertson
Denmark				
Niels Bohr Inst:				
Univ Cen.	IBM 4341	DKUCCC11		Bjorn Nilsson (NBIBSN)
RECKU	Sperry 1100	DKCCRE01		
France				
Montpellier:				
Central	IBM 3081D	FRMOP11		
Central	IBM 3033U	FRMOP22		
Germany				
Univ. Bonn:				
Central	IBM 3081	DBNRHRZ2		

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person, USERID
DESY:				
Central	IBM 3084Q	DHHDESY3		
Mark-J	VAX 11/780			PSSN:2624540009306 Coloured Books Richard Mount, MOUNT
Tasso	VAX 11/750		DYVB[*]	K. Gather

[* Restricted to U.K. users only, due to be moved to PPSN]

Heidelberg:				
Center	IBM 3081D	DHDURZ2		
HEP	IBM 4341	DHDIHEP1		
MPI-Munich:				
Central	IBM 4341	DM0MPI11		

Israel

Weizmann Institute:				
Central	IBM 3081D	WEIZMANN		
Tel Aviv:				
Central	IBM 4341	TAUNIVM		
Central	CDC 170	TAUNNOS		
Technion University:				
Central	IBM 3081D	TECHNION		

Italy

Bari:				
CSATA	IBM 4341	IBACSATA		
INFN	VAX			INFNET:VAXBA
Bologna:				
Center	IBM 4341	ICINECA		
INFN	VAX	earnet?		INFNET:VXBO(12), Sarterelli
gateway	PDP11/70	IBOINFN		INFNET:TCCNAF(11),Gateway
Frascati:				
	VAX			INFNET:VAXLNF(5)
Genova				
	VAX			INFNET:VAXGE(19)
Milan:				
	VAX			INFNET:VAXMI(22)
Padova:				
	VAX			INFNET:VXFPD

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person, USERID
Pisa:				
CNUCE	IBM 3033	ICNUCEVM		
CNUCE	IBM 370/168	ICNUCEVS		
INFN	IBM 4341	IPIINFN		
INFN	VAX 11/750	IPIVAXIN		INFNET:VAXPI(20)
Rome:				
	VAX 11/780			INFNET:VXROM(6) Valente
	?			INFNET:CBROM(7)
	?			INFNET:OGROM(8)
Trieste:				
	VAX			INFNET:VAXTS(21), Liello

Japan

KEK:				
Tristan	VAX 11/750			PPSN:44012943104 R. Hayano [Hayano]

Netherlands

Stichting Academ:				
Central	VAX	hasara5		
Kath. Uni Nijmegen:				
Central	NAS9040	HNYURC11		
HEP	VAX	hnykun55		
Univ. Utrecht				
Phys	VAX 11/785	bitnet		

Switzerland

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person,USERID
CERN:				
Central	IBM WYLBUR	GEN		USENET:CERNIBM
Central	IBM VM	CERNVM		
Central	CDC 875	cyber		
3081/E	IBM 4361	CERNVME		
DD Dev	VAX 11/780	CRVXDEV	CPVA[*]	INFNET:VXDEV(60)
OPAL	VAX 11/780		CPVB[*]	N. Gee
Omega	VAX 11/780		CPVC[*]	W. Carena-Bozzoli
Merlin	VAX 11/780		CPVD[*]	D. Myers
Delphi	VAX		CPVE[*]	G. Mornacchi
Gateway	?			INFNET:CERNGW (10)
OPAL	VAX			INFNET:VXHYP(48)
LEP DB2	VAX			INFNET:VXLDB2(37)
Gift Proj	VAX			INFNET:VXGIFT(40)
L3	VAX 11/750			INFNET:VAXL3(41)
Aleph Dev	VAX	CRVXALFB		INFNET:VAXALFB(54)
Aleph beam	VAX		CPVF[*]	INFNET:VAXALBM(55)
Aleph TPC	VAX	CRVXALTP		INFNET:VAXALTP(56)
Unix	VAX	cernvax		USENET:CERNVAX

[* restricted to U.K. users only, due to move CERN X.25 network in January 1985]

United Kingdom

Site:/ Dept.	Computer	BITNET	DECNET	OtherNet Contact Person,USERID
Birmingham U:				
	IBM 4341		BHIA	L. Lowe
Bristol Univ:				
	VAX 11/750		BRVA	J. Alcock
Cambridge Univ:				
	VAX 11/780		CAVA	R Ansorge
Daresbury Laboratory:				
Mail Server	GEC		DLGM	
Edinburgh Univ:				
	VAX 11/750		EDVA	D. Candlin
Glasgow Univ:				
Nat Phil.	IBM 4361		GWIA	A. Conway
Imperial College:				
	VAX 11/780		ZIVA	R. Beuselinck
	IBM 4341		ZIIA	R. Campbell
Lancaster Univ:				
	VAX 11/750		LAVA	R. Henderson
Liverpool Univ:				
	GEC 4085		LLGA	M. Houlden
	IBM 4331		LLIA	
Manchester Univ:				
	GEC 4090		MAGA	R. Hughes-Jones
Oxford University:				
Nucl Phys	VAX 11/780		XXVA	J. Macallister
NP	DEC 10		XXDA	W. Black (BLACK)
Queen Mary College:				
	VAX 11/750		ZMVA	P. Kyberd
Rutherford laboratory:				
Central	IBM 3081(VM)	rlvm370	RLIB	R. Maybury
Central	GEC 4090		RLGB	K. Duffey
Gateway	IBM	ukacrl		
HEP	VAX 11/780		RLVB	M. Waters
Library	Prime		RLPE	
Sheffield Univ:				
	GEC 4085		SHGA	C. Wells.
Southampton:				
	GEC 4070		SNGA	M. Counihan
Surrey Univ:				
	Prime 550		SYPE	
Sussex Univ:				
	Prime 550		SVPA	
University College, London:				
	GEC 4085		ZUGA	J. Conboy

Gateways between Networks:

From: \ To:	BITNET	DECNET	JANET	INFNET	ARPANET
BITNET	-	SLACCB	ukacrl	iboinfn	WISCVM
DECNET	<MANY>	-			
JANET	????		-	gift	ZUXA[*]
INFNET	VAXPI		vxgift	-	
ARPANET	WISCVM		UCL-CS[*]		-

[Note: Lower case indicates not yet operational or planned]
[BITNET and EARNET are considered the same network]
[* restricted to registered ARPANET users]

TABLE 3 - CHARACTERISTICS OF UPPER-END GENERAL-PURPOSE COMPUTERS

Vendor	Model	Options	CPU Cycle (ns)	Memory Max (Mb)	Memory Cycle (ns)	Advertised* Peak Perf.	CPU System+ Price/Mem Size or Availability
AMDAHL	5860		24	64	280	13mips	$4M/64Mb
	5880	2 proc	24	64	280	23mips	$5.5M/64Mb
CDC	Cyber 875	1 proc	25	8	75	19mips	$4M/8Mb
	Cyber 875	2 proc	25	8	75	32mips	$5.5M/8Mb
	Cyber 990	1 proc	16	32	64	32mips	$5M/32Mb
	Cyber 990	2 proc	16	32	64	58mips	$7M/32Mb
	Cyber 205	2 pipes	20	128	80	100-200mfls	$9M/16Mb
	Cyber 205	4 pipes	20	128	80	200-400mfls	$14.5M/32Mb
CRAY	1S		12.5	32	50	160mfls	$7M/8Mb
	2	4 proc	4	2000		1-2gfls	1985
	3	1-8 proc					1986
	X-MP/1	1 proc	9.5	32	70	250mfls	$7M/32Mb
	X-MP/2	2 proc	9.5	32	25	480mfls	$10.5M/32Mb
	X-MP/4	4 proc	9.5	64	25	950mfls	$14M/64Mb
	Y-MP		5	1000		10gfls	1986
DENELCOR	HEP1	1-16 proc	100	1000	50	10-160mips	
	HEP2			2000		250mips	1986
				2000		4000mips	1986
ELXSI	ELXSI64	1-10 proc	50	192	400	4-40mips	$4M/96Mb
ETA(CDC)	GF10	2 proc		256		1gfls	1987
		8 proc		256		10gfls	1987
	GF30	8 proc		4000		30gfls	1990
FUJITSU	VP-200		7.5	256	55	500mfls	$11M/64Mb
HITACHI	S810-20		6	256	40	630mfls	
IBM	3081KX	2 proc	24	64	312	19mips	$5M/64Mb
	3084QX	4 proc	24	128	312	29mips	$8M/64Mb
	3090/200	2 proc	18.5	64+128		45mips	$5m/64Mb
	3090/400	4 proc	18.5	128+256		80mips	1986
NAS	AS9080	2 proc +IAP(VPF)	30	64	320	20mips 100mfls	$6M/64Mb
NEC	SX-2		4.5	256	40	1.3gfls	1985 Japan
SPERRY	1100/91	1 proc	30	64	360	8mips	$4.5M/64Mb
	1100/94	4 proc	30	64	360	27mips	$9.5M/64Mb
DEC	VAX11/780FPA	1 proc	200	32	290	1mips	For Comparison
	VAX 8600	1 proc	80	32	560	4mips	For Comparison
CDC	7600	1 proc	27.5	0.5	275	10mips	For Comparison

NOTES:

*Peak performance figures may be far removed from actual benchmark performance.

+The prices listed are estimates for processor and memory system sized as shown. The cost of peripherals is not included.

Table 4

Software Manpower Estimates for some U.S. Detectors

Detector	Number of Full Time Equivalent People per Year*
Mk II Upgrade (SLC)	16
TPC and 2 γ (PEP)	40
CDF (TEV I)	35
D0 (TEV I)	30
L3**	58
One SSC Detector	60-100

---

* These numbers (except for TPC) are estimates for the period of design, construction and initial running. Most experimenters expect that very similar numbers (although not necessarily the same people) will be used during the course of the experiment.

** L3 is listed because of the large U.S. participation.

CDF Offline Computers
(1985-1986)

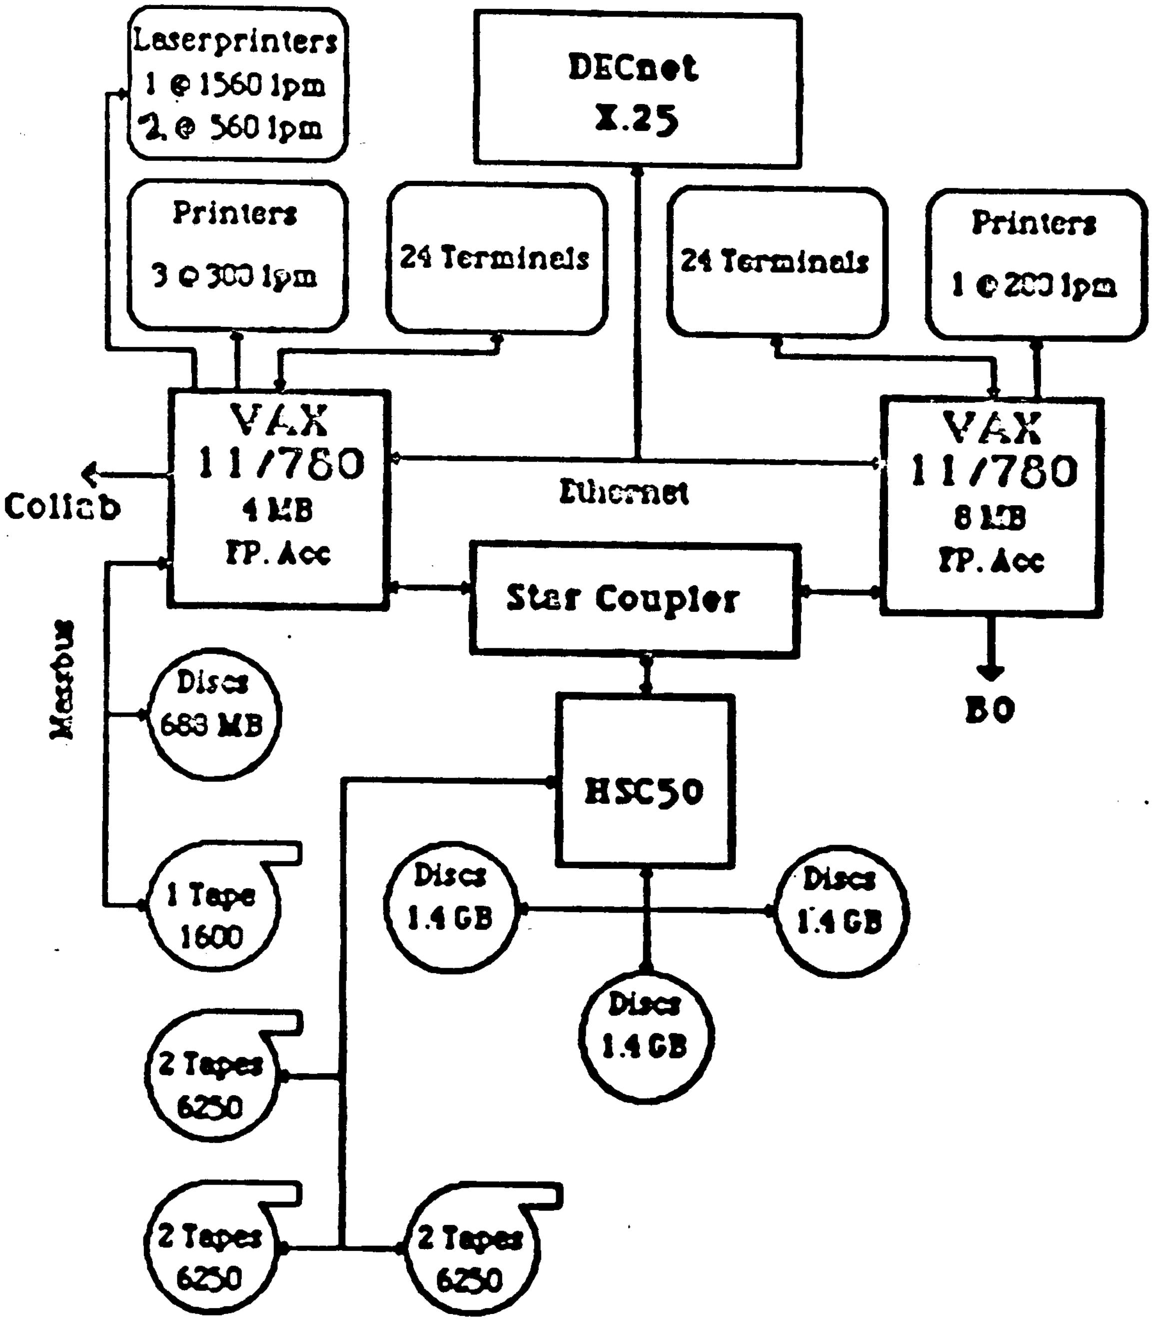

FIGURE 1

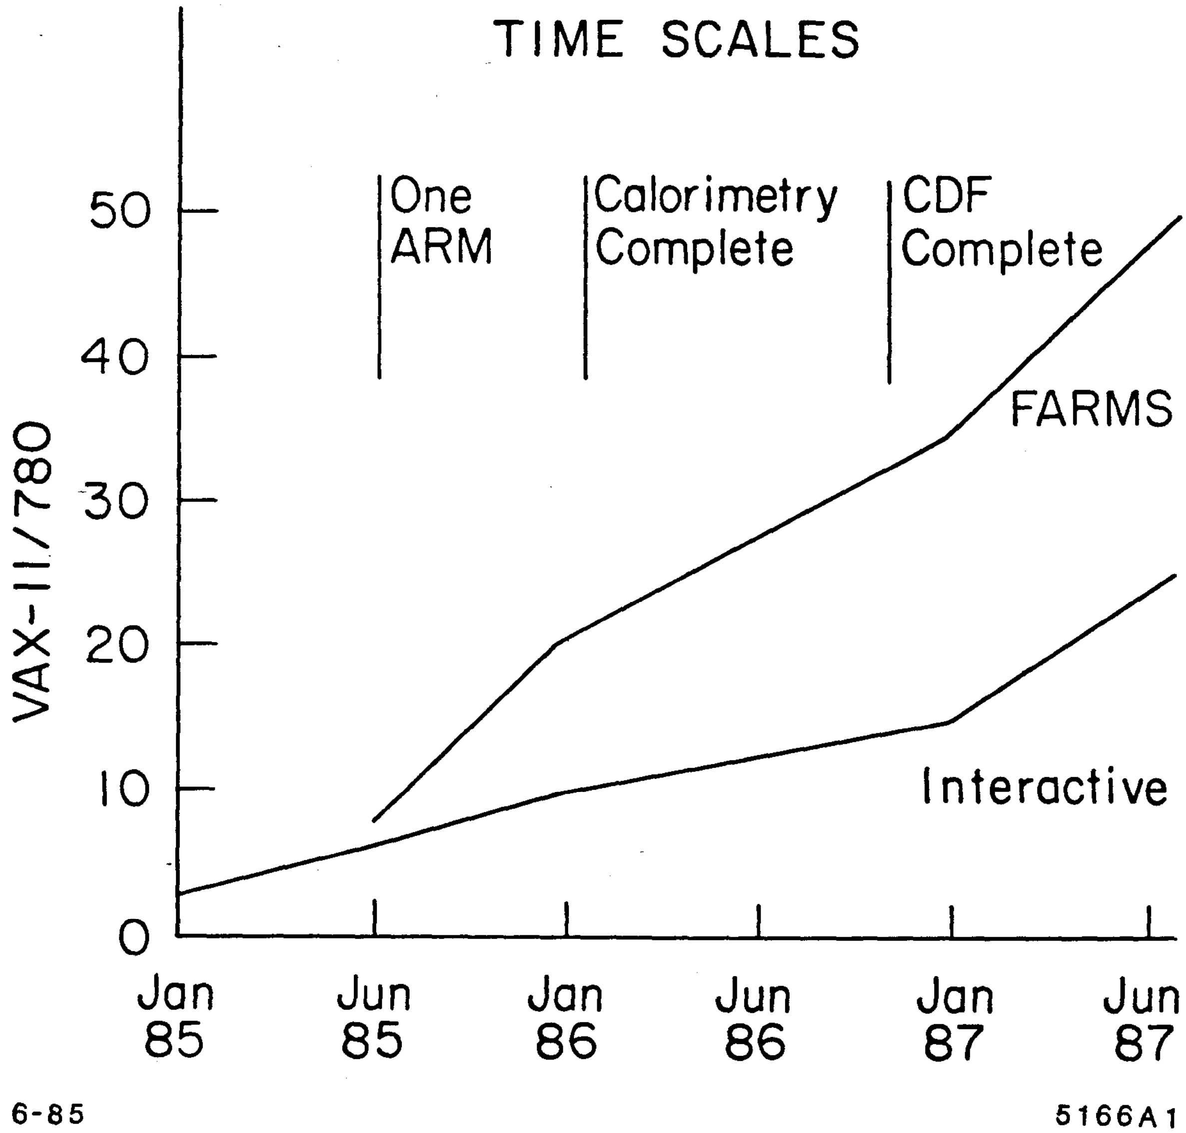
TIME SCALES
One ARM
Calorimetry Complete
CDF Complete
FARMS
Interactive
VAX-11/780
0
10
20
30
40
50
Jan 85
Jun 85
Jan 86
Jun 86
Jan 87
Jun 87
6-85
5166A1

FIGURE 2

BITNET

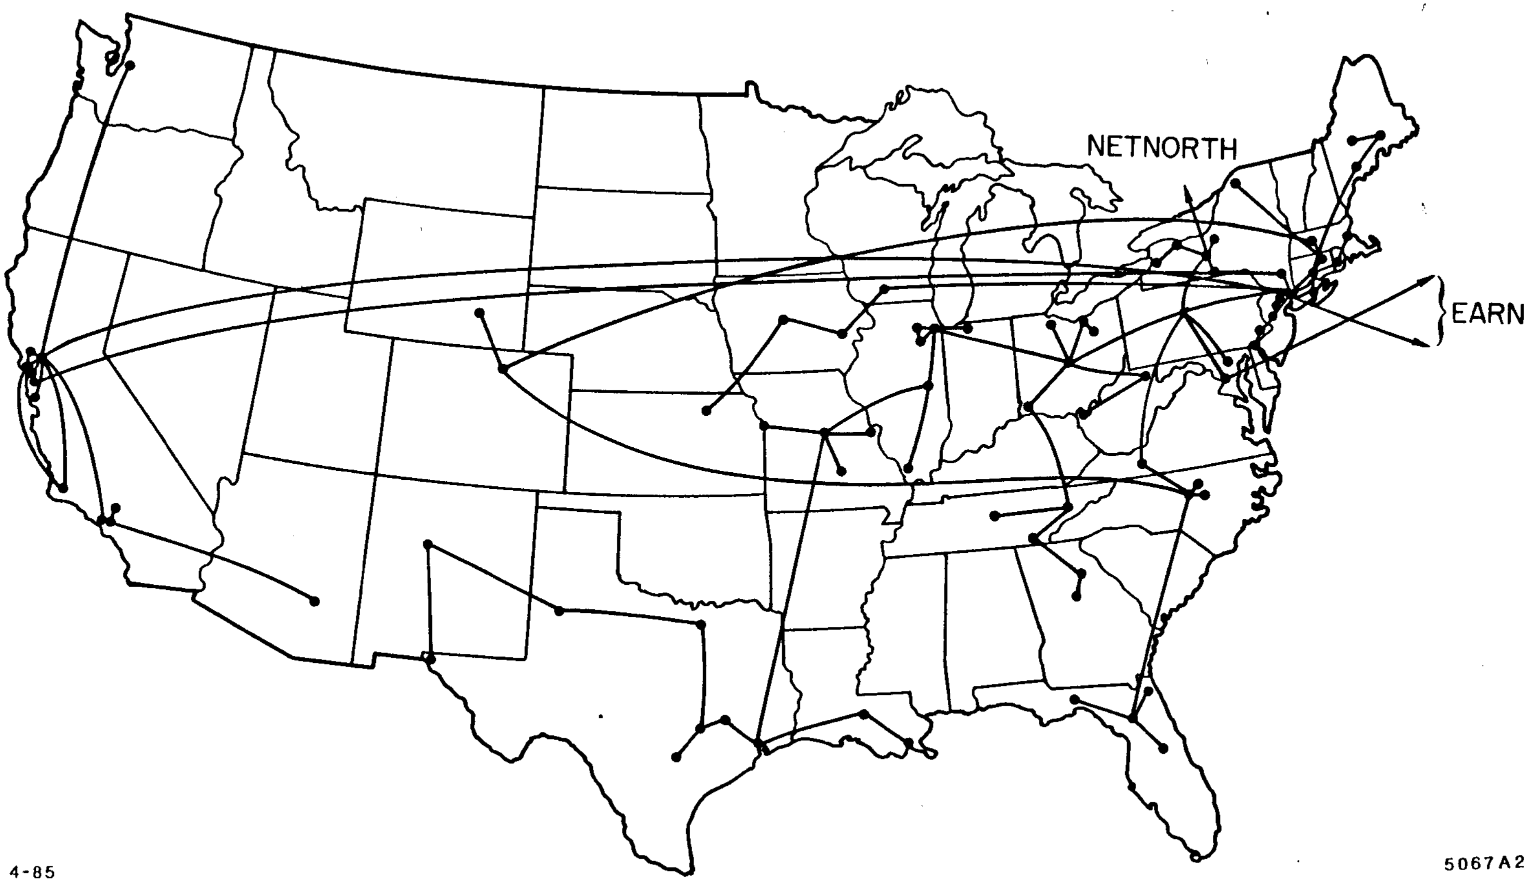

4-85

5067A2

FIGURE 3

HEP LEASED LINES FOR TERMINALS

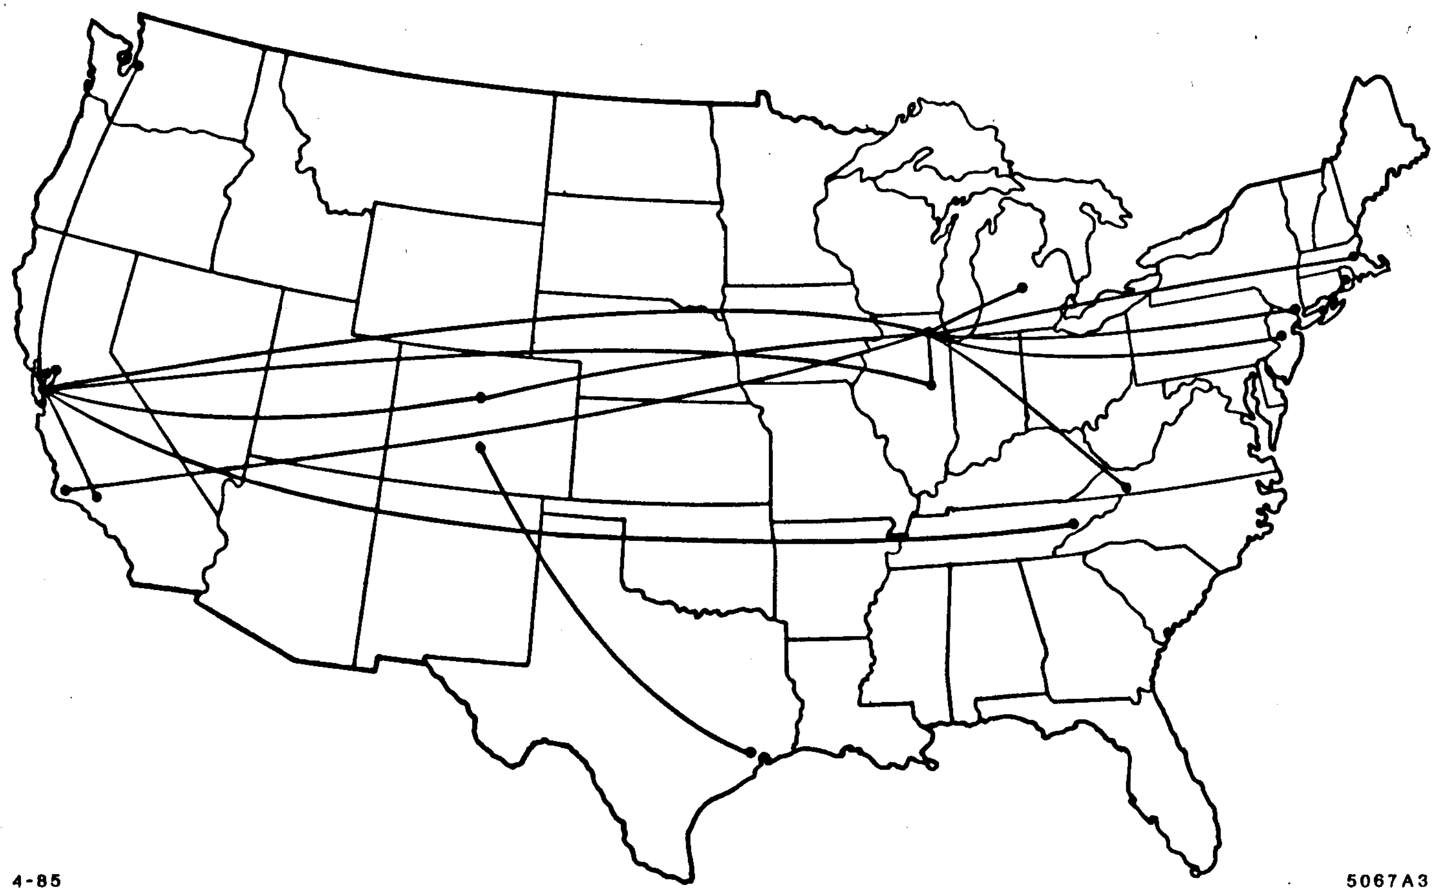

FIGURE 4

SLAC/LBL DECNET

FIGURE 5

HEPNET PHASE I

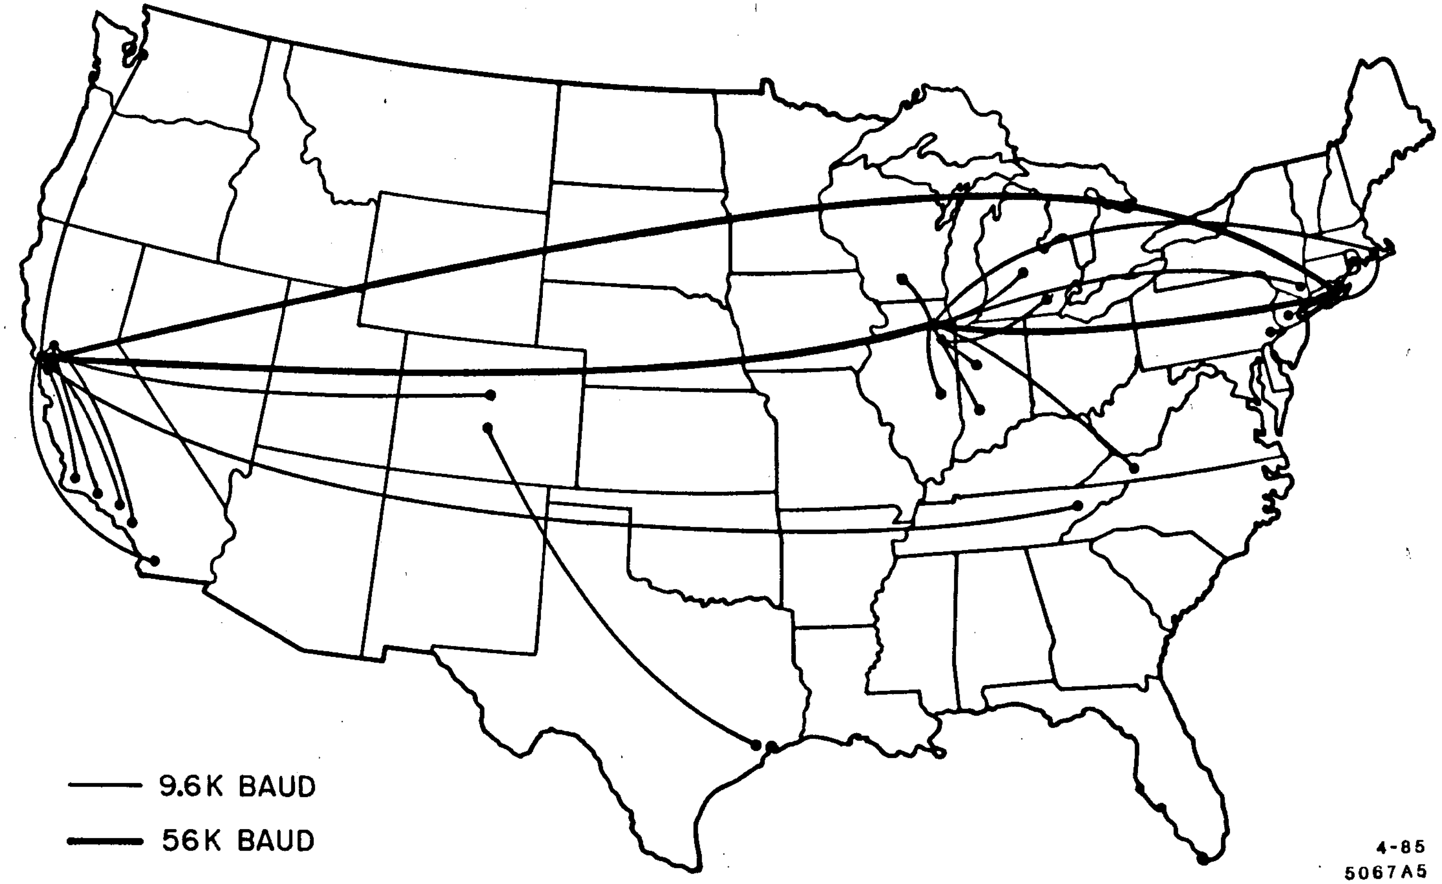

FIGURE 6

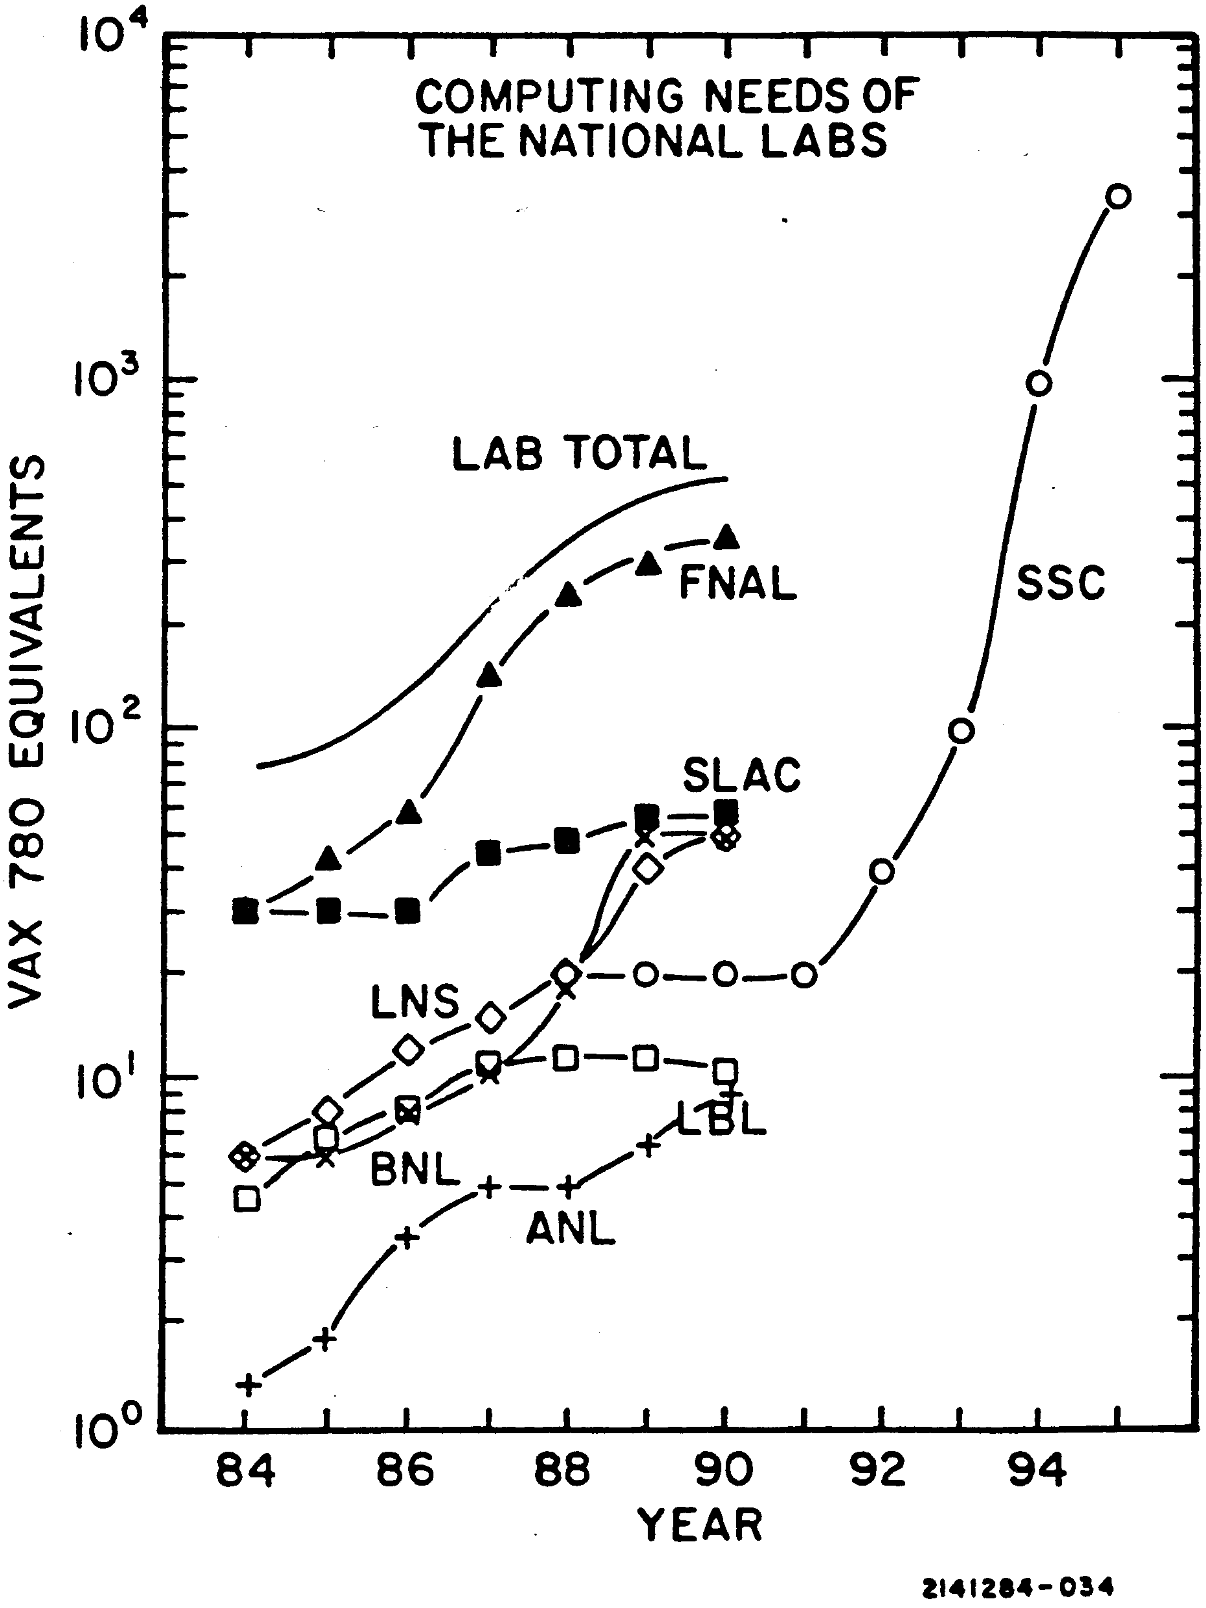
COMPUTING NEEDS OF THE NATIONAL LABS
VAX 780 EQUIVALENTS
YEAR
10^0
10^1
10^2
10^3
10^4
84
86
88
90
92
94
LAB TOTAL
FNAL
SLAC
LNS
BNL
ANL
LBL
SSC
2141284-034

FIGURE 7