SSC-N-579

Report of the Task Force on Computing for the Superconducting Super Collider

December 1988

M. Gilchriese, Editor

Table of Contents

1. INTRODUCTION AND CHARGE TO THE TASK FORCE 1
   - 1.2 Scope of the Task Force Report 1
   - 1.3 Assumptions 2
2. PRESENT STATUS OF COMPUTING RESOURCES FOR THE CDG 2
   - 2.1 SUN Clusters 2
   - 2.2 LBL VAX Cluster 7
   - 2.3 NMFECC Supercomputer Center 8
   - 2.4 Support Services and Personnel 8
3. OVERVIEW OF ACCELERATOR DIVISION COMPUTING 8
   - 3.1 Accelerator Physics Computing Needs for FY89 and FY90 10
     - 3.1.1 VAX cluster usage 10
     - 3.1.2 Workstations 10
     - 3.1.3 Cray 10
     - 3.1.4 Parallel Computing Machine 10
   - 3.2 Accelerator Systems 11
     - 3.2.1 VAX usage 11
     - 3.2.2 Workstations 11
     - 3.2.3 Cray 11
   - 3.3 Magnet Division Computing Needs for FY89 and FY90 11
     - 3.3.1 Workstations 12
     - 3.3.2 VAX 12
     - 3.3.3 Personnel 12
   - 3.4 Computing Needs for the Magnet Test Laboratory (MTL) for FY89 and FY90 12
   - 3.5 Summary of Accelerator Physics Needs for FY89 and FY90 13
     - 3.5.1 VAX 13
     - 3.5.2 Workstations 13
     - 3.5.3 Systems maintenance and support 14
     - 3.5.4 SUN hardware support 14
     - 3.5.5 SDS, database software support 14

3.5.6 Cray 14
3.5.7 VAX and networking support 14
4. PHYSICS AND DETECTOR SIMULATION IN FY89 AND FY90 15
4.1 Status of Physics and Detector Simulation Programs 15
4.1.1 Physics simulation 16
4.1.2 Radiation transport and shower simulation 16
4.1.3 Detector simulation 17
4.1.4 CPU times and data storage needs 18
4.1.5 Vectorization and related issues 20
4.1.5.1 Languages 20
4.1.5.2 Basic HEP software 21
4.1.5.3 Applications 21
4.1.5.4 Production experience at the supercomputer centers 23
4.1.6 Parallel computers 23
4.2 Model for Physics and Detector Simulation Needs 25
4.2.1 CPU requirements 25
4.2.2. Storage requirements 27
4.3 Availability of Computing for Detector Simulation in FY89 and FY90 28
4.3.1 Survey of laboratories 28
4.3.2 Supercomputer centers 28
5. RADIATION TRANSPORT CODES 32
5.1 Types of Codes 32
5.2 CPU Requirements 32
6. NETWORKING 34
6.1 Overview of Wide-Area Networks and their Status 35
6.2 Current SSC Networking at LBL 36
6.2.1 Macintoshes 39
6.2.2 IBM PCs 39
6.2.2.1 Network connections 39
6.2.3 Local area networking at LBL in FY89 and FY90 40
6.3 Networking to SSC Site 40

6.3.1 FY89 40

6.3.2 FY90 42

6.4 Networking at the SSC Site Location 42

6.5 Networking Costs 42

6.6 Hardware/Software 43

6.7 Personnel 43

6.8 Conclusion 43

7. SSC COMPUTING PLAN FOR FY89 AND FY90 44

7.1 Accelerator Physics, Systems, Magnets and MTL 44

7.2 Radiation Transport Codes 46

7.3 Physics and Detector Simulation 47

7.3.1 Plan for FY89 47

7.3.2 Plan for FY90 and beyond 48

7.4 Networking and Communications 49

7.4.1 FY89 49

7.4.2 FY90 50

7.5 Budget & Personnel Summary 50

Appendix A. List of Attendees 53

Appendix B. Computing for High Energy Physics in the 1990's 54

Report of the Task Force on Computing for the Superconducting Super Collider

December 1988

1. INTRODUCTION AND CHARGE TO THE TASK FORCE

The Task Force was given the following charge:

The Task Force should evaluate the large-scale computing needs—those traditionally associated with off-line computing—for the SSC, with emphasis on the near-term (the next 1–2 years) computing requirements for the SSC accelerator complex and for SSC detector design activities. Needs to be considered should include (but are not restricted to):

- accelerator simulation code development and use
- accelerator system and control simulation
- SSC physics simulation development and use by the HEP community
- detector simulation development and use by the HEP community
- radiation transport codes for estimates of radiation levels and shielding requirements
- engineering design of accelerator and detector components
- initial software development for SSC experiments
- networking requirements
- personnel requirements for code development and operation of computing facilities

The Task Force should consider CAD/CAE requirements or the use of personal computers only as they relate to the larger-scale needs indicated above.

In addition to describing SSC computing needs in quantitative detail, the Task Force should examine possible methods of meeting such needs, again emphasizing the possibilities for the next two years. One or more plans for meeting the perceived SSC computing needs should be developed in detail, including preliminary cost estimates for FY89 and FY90.

The members of the Task Force are listed in Appendix A. The Task Force met at the Central Design Group on December 12–14, 1988.

1.2 SCOPE OF THE TASK FORCE REPORT

It is not the intent of this report to cover all aspects of computing for the SSC laboratory and potential users. Furthermore, this report is a planning document for the remainder of FY89 and for FY90 and not a document describing the computing policy of the SSC laboratory. The latter is clearly beyond the authority of the Task Force.

This report describes an aggressive plan for the continuation and necessary expansion of all computing needed for accelerator and accelerator component design for the SSC. In addition, the plan includes provisions for a significant computing capability at the SSC laboratory for physics and detector simulation in support of the design of the initial SSC experiments.

1.3 ASSUMPTIONS

In order to arrive at a plan for providing computing resources for the SSC during the remainder of FY89 and in FY90, it is necessary to make assumptions regarding the location, scope, and nature of its activities during these years. There are obviously considerable uncertainties in any predictions regarding these topics. We have assumed that construction of the SSC will proceed on an aggressive and timely schedule and, therefore, that computing resources must be provided on an equally aggressive schedule. Furthermore, we have assumed that considerable activity requiring computing resources will commence as soon as possible at the SSC site. Our principal assumptions are:

(1) Network communications with the SSC site must be established as soon as it is feasible to do so. This will include the ability to communicate via BITnet and to accomplish remote login to, for example, the LBL VAX cluster.

(2) The ability to access computing resources of the character now (or soon to be) available at the Central Design Group to support SSC accelerator and accelerator component design will be in place and operational at or near the SSC site by the beginning of FY90. A smooth transition will occur during FY89, increasing computing capability at the SSC site while retaining the requisite resources at the LBL location.

(3) Computing resources needed to support some aspects of physics and detector simulation for the design of the initial SSC experiments will be in place and operational at or near the SSC site by the beginning of FY90. However, we assume that the bulk of the computing resources needed for physics and detector simulation during at least FY89 will reside in universities, laboratories and supercomputer centers.

It should be noted that "computing resources" includes networking capability, data storage, graphics capability . . . all aspects of computing, not just CPU availability. Also included in "resources" are the personnel needed to maintain and operate the computing hardware and systems-level software.

2. PRESENT STATUS OF COMPUTING RESOURCES FOR THE CDG

The overall computing environment of the SSC Central Design Group is widely distributed. The SSC/CDG supports two SUN clusters for the development of accelerator operations simulation code and tools and for the analysis of magnet test data and stress analysis. Moreover, the CDG has placed various "satellite" SUN workstations at selected laboratories. The CDG relies heavily on the Lawrence Berkeley Laboratory VAX cluster for moderate to large computations and for general communication with other computing facilities and the scientific community. Large-scale computations are performed on the National Magnetic Fusion Energy Computer Center (NMFECC) Supercomputers located at the Lawrence Livermore National Laboratory.

2.1 SUN CLUSTERS

Since 1986, the Accelerator Physics Division has operated a cluster of SUN workstations for the simulation of SSC operations. The cluster has since grown to ten SUN workstations at the center of which is a SUN 3/160 file server with 8 Mbyte of random access memory (RAM), a

Floating Point Accelerator (FPA), and 2.3 Gbyte of disk storage. The clients include four SUN 3/50s, one color SUN 3/60, one SUN 3/140 with an FPA, two SUN 4/110s, each with its own 327 Mbyte disk, and a SUN 4/260 with 32 Mbyte of RAM (see Fig. 2.1.1). The SUN 3/160 has also served as a file server for the local Macintosh network and is the host for the commercial database SYBASE. Accessories to this cluster include an 8-mm tape backup unit capable of storing 2.2 Gbyte of data on a single 8-mm tape as well as four cartridge tape drives attached to various workstations. The size of the cluster has grown at the rate of approximately four machines per year.

The Magnet Division has acquired a separate SUN cluster for data analysis by its Magnet Test and Data Management Group. The heart of this system is a color SUN 4/280 file server with 64 Mbyte of RAM and 1.8 Gbyte of disk storage. It serves three color SUN 4/110s with 8–16 Mbyte of RAM and a color SUN 4/260 with 32 Mbyte of RAM and 560 Mbyte of local disk storage. The latter machine is dedicated to running the commercial finite element code ANSYS for the Magnet Analysis Group. Hardware to support this cluster includes a 1/2-inch tape unit and an 8-mm tape unit, as well as a 2–4 Gbyte optical write-once-read-many (WORM) device (Fig. 2.1.2).

Several SSC/CDG SUN workstations are located in various national laboratories and universities to support the on-going efforts by those institutions for the SSC Central Design Group. As part of the SSC E778 experiment at the Fermi National Accelerator Laboratory, two SUN workstations—a SUN 3/140 and a SUN 3/60 with a 140 Mbyte disk—have been installed at the Fermilab Accelerator Division. A third SUN workstation, a SUN 3/50, supports the SSC work at the Magnet Test Facility at FNAL. A more powerful CPU, probably a SUN 4/260, will eventually replace this machine. In addition, Brookhaven National Laboratory will obtain an SSC SUN 4/110 with 16 Mbyte of RAM and 327 Mbyte of disk storage to support their efforts in the magnet testing program. A SUN 3/60 with 140 Mbyte of disk storage is located at Cornell's Newman Laboratory for Nuclear Studies to support the accelerator physics program (Fig. 2.1.3). In its present configuration, the various CDG-supported SUN hardware has an estimated worth of $542,000, with the Accelerator Physics Cluster being $181,000, the Magnet Cluster $280,000, and the rest $81,000. These estimates take into account the 33 percent discount which the SSC/CDG currently receives from SUN Microsystems for SUN products. Not included in the estimates are the software purchases and licenses associated with the hardware.

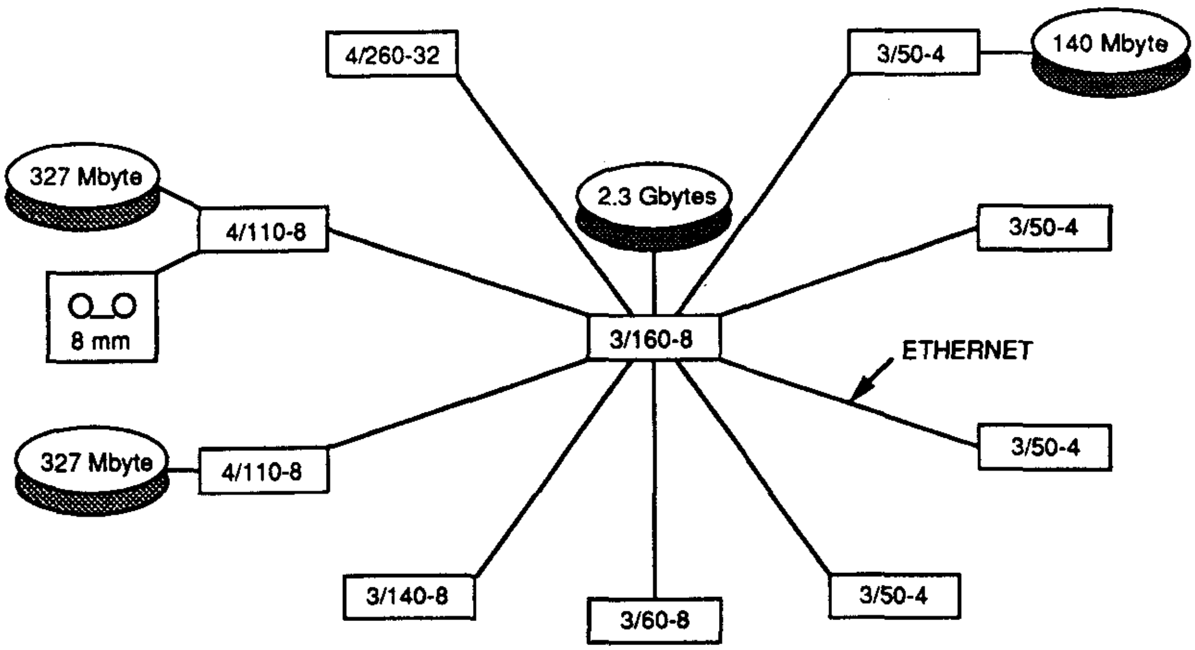

Fig. 2.1.1. Accelerator Physics Division Sun Cluster

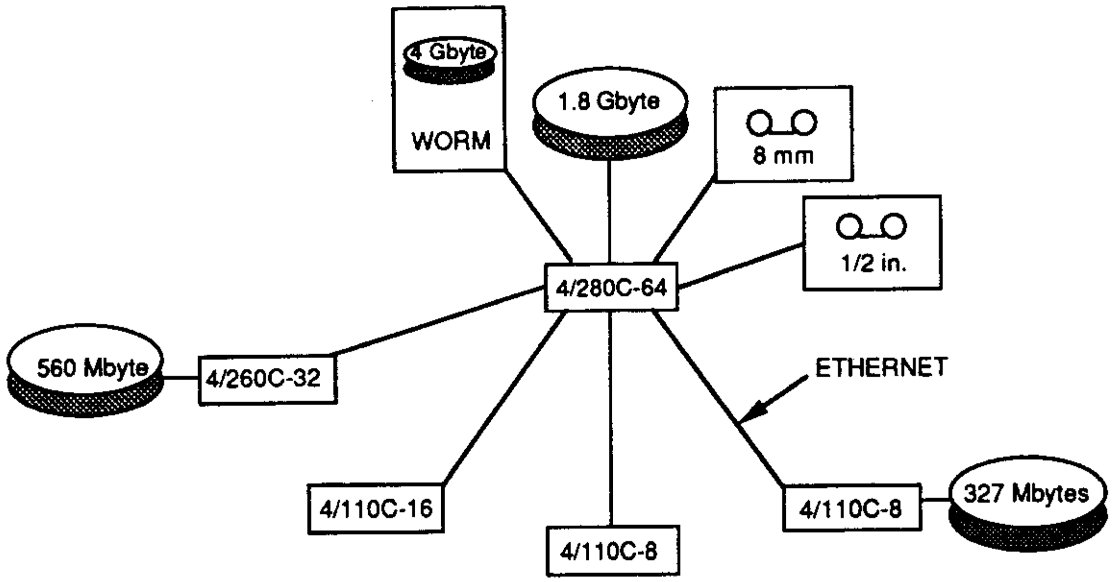

Fig. 2.1.2. Magnet Division Sun Cluster

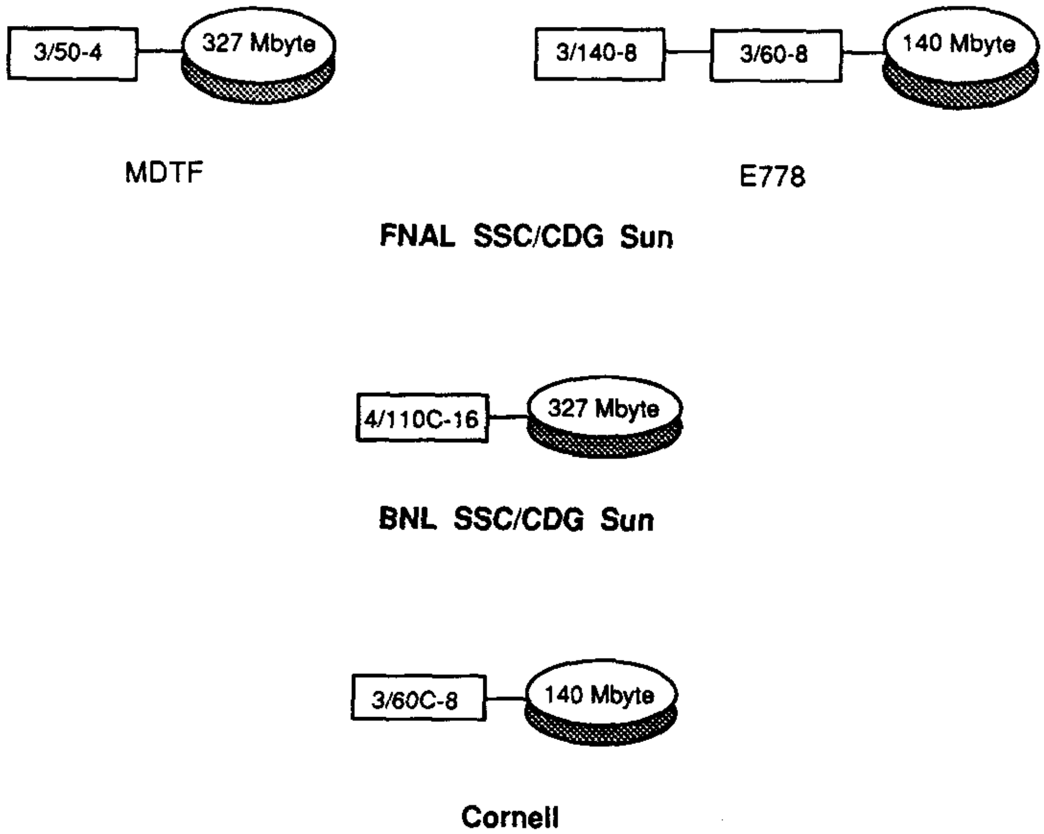

Fig. 2.1.3. SSC/CDG Supported Sun Clusters at Various Laboratories

2.2 LBL VAX CLUSTER

The LBL VAX cluster consists of five Digital Equipment Corporation VAX 8650 CPUs running the VMS operating system. At the writing of this report, there are approximately 105 registered users from the SSC on the VAX cluster, of which about 60 percent are scientific users. The latter group accounts for more than 90 percent of the actual usage of the cluster, which in FY88 was approximately 1,400 hours of CPU time (in VAX 8650 units) with an average disk space usage of 2 Gbyte. The CDG VAX computation budget for FY88 was $322,161, not including the 48 percent overhead. It is expected that VAX usage will increase significantly over the next year. At present, the SSC has committed $432,000 in FY89 for LBL computing ($640,000 including overhead), an increase of 34 percent. Early indications seem to show that the estimate is low. A twelve-month CPU- and disk-use profile is illustrated in Fig. 2.2.

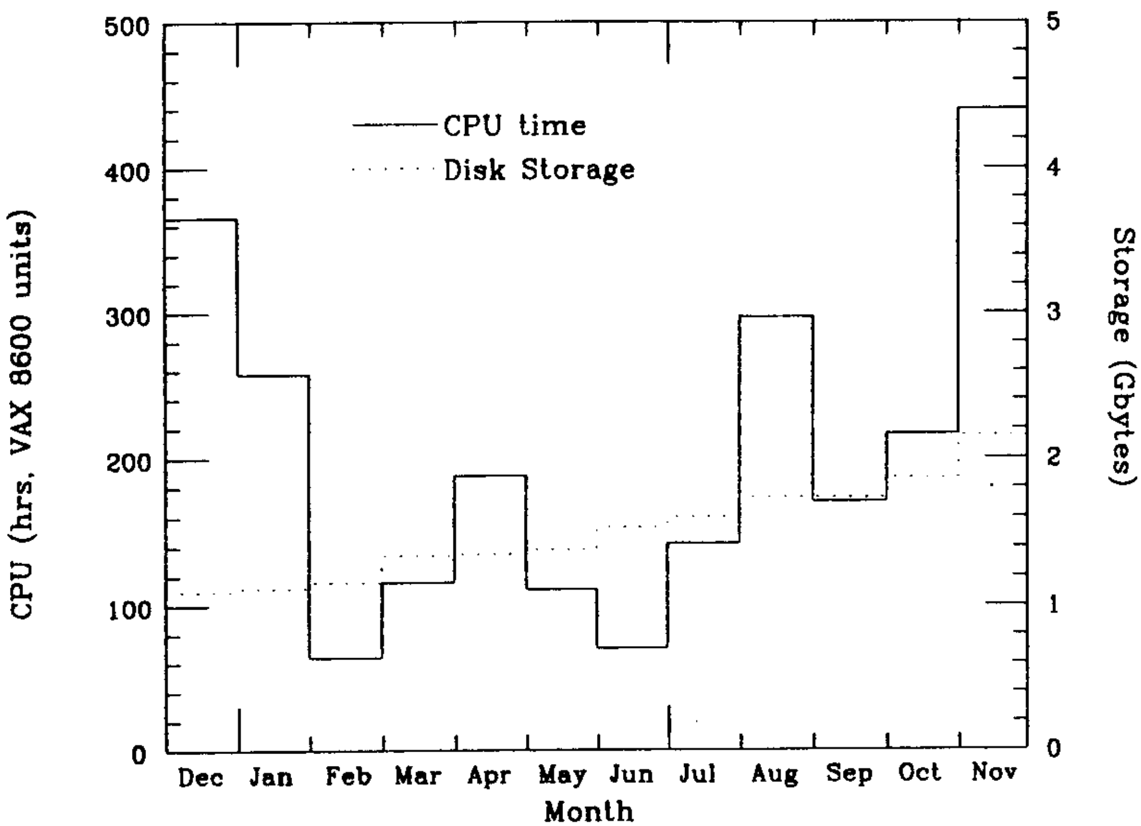

Fig. 2.2 SSC Central Design Group Vax Usage 1987/88. The solid line represents the monthly CPU hours used in VAX 8600 units. The dotted line shows the average diskspace used per month.

Primarily in support of the VAX cluster users, the SSC/CDG has acquired, at an approximate cost of $30,000, an Imagen 3320 Image Server printer capable of producing 20 pages/minute. In addition, there exists a slower Imagen station, along with various Apple LaserWriters, to support the printing needs at the SSC/CDG.

2.3 NMFECC SUPERCOMPUTER CENTER

The NMFECC supports five Cray supercomputers for the general scientific community. It consists of two Cray-1s, two Cray-2s, and one Cray X-MP. High-speed computations in the areas of particle-ray tracing, beam stability, and magnet design are carried out on these machines by members of the CDG. In FY88, the SSC requested and received a total of 1,100 hours from the Department of Energy to meet its requirement. The actual usage was approximately 600 hours, or 54 percent of the allocation. The Cray usage profile (Fig. 2.3) shows sporadic use of the time, due to the unpredictable nature of projects requiring high-speed computation. The request for NMFECC Cray time for FY89 remained at 1100 hours, and the allocation has been granted. Higher usage of the Cray is anticipated for this year, although the number of active users from the SSC/CDG (now numbering 15 out of a total of 69 accounts) is not expected to increase significantly. Despite the added inconvenience in using the Crays, some computations which would otherwise not be feasible rely heavily on their use.

2.4 SUPPORT SERVICES AND PERSONNEL

The SSC/CDG currently does not have a separate computing services division. The day-to-day administrative tasks of the two local SUN clusters are divided among selected individuals within the CDG. Hardware and software support is provided on an as-needed basis by the Real Time System Section (RTSS) within LBL's Engineering Division. The hourly rate of these services ranges from $40.00 for technical support (mainly hardware repair) to $72.00 for engineering support (software upgrades and development).

The user support for LBL's VAX cluster and LLNL's NMFECC Cray cluster is provided by the respective host institutions.

3. OVERVIEW OF ACCELERATOR DIVISION COMPUTING

The CDG Accelerator Division is composed of three diverse groups: the Magnet Division, the Accelerator Systems Division, and the Accelerator Physics Division. The work of these divisions ranges from magnetic measurements to engineering design of the cryogenics system to calculations of transfer maps using differential algebra techniques. However, all have in common the desire to access, analyze, and display data, whether from simulations or from real measurements. A unified approach to accessing and displaying data is invaluable, if not essential, since it allows the sharing of tools and information among the various parts of the Accelerator Division. Eventually all the diverse views of the machine will coalesce into a real accelerator. At that point (actually well before that) we must be able to merge the various kinds of existing data, and programs must be able to access these data and to interoperate with one another. Only powerful and well-defined standards will make this possible.

The next two years will be crucial to this effort, since many programs will be developed which will eventually be used in the controls system of the machine, and much data will be gathered and stored which will be needed by those programs. At this point the foci of the standardization effort are the SYBASE database, SDS (self-describing datasets, a format for data transfer), UNIX, and the X WINDOW system. SYBASE is being used to create a lattice database, to

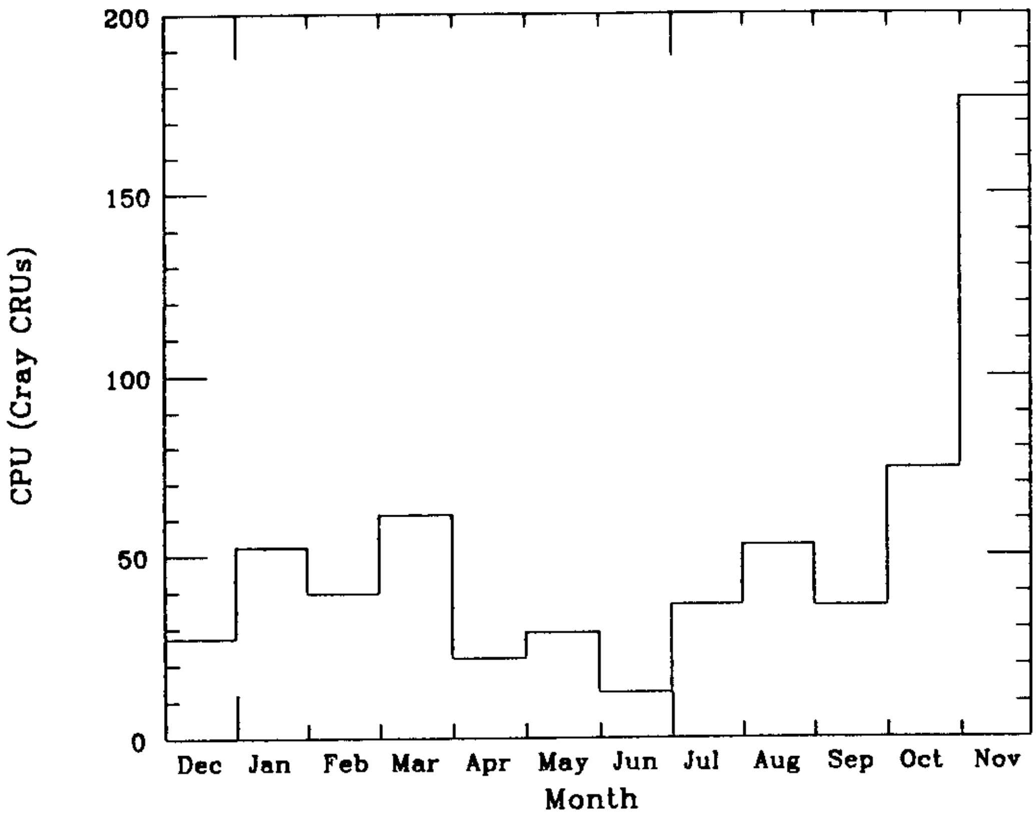

Fig. 2.3. **SSC Central Design Group Cray Usage 1987/88.** The solid line represents the monthly CPU time used in CRU units.

store magnet test data, and to store magnet cable data. Magnetic measurements will also be stored in a SYBASE database, as will the results of various types of simulations. SDS is used in E778 at Fermilab, in the dynamic cryogenics simulator, and will be used in the Magnet Test Laboratory and in modelling work in the Accelerator Physics Division. All the SUN workstations at CDG are UNIX-based and run X WINDOWS. It is extremely important that standards be adopted by the whole of the accelerator division, to allow communication and to avoid duplication of effort.

3.1 ACCELERATOR PHYSICS COMPUTING NEEDS FOR FY89 AND FY90

The Accelerator Physics Division anticipates modest growth in personnel in FY89 and FY90. Currently consisting of 14 members, the division should increase to 18 during FY89, suffer a small decrease at the time of the move to the site, and then add enough personnel to finish FY90 with a total of approximately 19 members. Computing in the Accelerator Physics Division consists mainly of lattice design, analysis, and tracking; analysis of data from E778; accelerator database development; and operations simulation and studies of correction schemes. The latter four tasks primarily use the SUN workstations, while the former, along with various other computing tasks, use the VAX cluster and the Cray.

3.1.1 VAX CLUSTER USAGE

In FY88 the Accelerator Physics Division was responsible for approximately 60 percent of the total 0.15 VAX (8650) years (= 0.09 VAX years) used by the CDG at LBL. We anticipate an increase to 0.12 VAX years for the Accelerator Physics Division by the end of FY89 due to four additional staff members. However, the goal is to reduce the dependence on the MFE Crays by increasing VAX usage by 15 percent per person. This makes the VAX usage estimate for FY89 equal to 0.13 VAX years. Realistically, VAX capability must substantially exceed demand to provide a productive environment. A load greater than 80% causes a significant decrease in system performance. Also, much of the computing is done during daytime hours, increasing the peak load. Therefore, the Accelerator Physics Division alone will need approximately half a VAX 8650 by the end of FY89. The Division's needs for FY90 are anticipated to be the same as those for FY89.

3.1.2 WORKSTATIONS

Accelerator Physics currently has one cluster of nine workstations. In FY89, we will purchase one more disk drive and two workstations, for a total expenditure of $50 K. A second cluster serving six people should be in place early in FY90, and another by the end of that year. The estimated cost of each of these clusters is $240 K, including disk storage and tape drives.

3.1.3 CRAY

The Accelerator Physics Division anticipates no increase in Cray use in the next two years. However, use at the present level must be maintained. Connection with NMFECC, either directly or via LBL, from the SSC site must be established by the end of FY89 or before.

3.1.4 PARALLEL COMPUTING MACHINE

Substantial portions of the accelerator physics simulations are dominated by time spent doing multi-particle tracking. Currently the most CPU-intensive of these simulations are done on the MFE Cray in runs lasting many hours. Multi-particle tracking, however, is extremely well suited for parallel processor architectures. SSC-N-531 discusses a parallel architecture which is projected to deliver from 1/3rd to 3 times the performance of the MFE Cray for a cost of $100 K. Working examples of such systems should be investigated in FY89 and if the needed performance is attainable, one should be acquired. Such a system will add to the system administration needs, though at a level substantially lower than the SUN clusters since the system will see much more restricted use. The added administrative burden is estimated at 20 percent of a SUN cluster, though it would be higher initially.

3.2 ACCELERATOR SYSTEMS

The Accelerator Systems Division will grow substantially in FY89 and FY90, from 9 people at the end of calendar 1988 to approximately 20 by the end of FY90. Most of the computing effort in the Systems division consists of cryogenics simulations, both steady-state and dynamic, which currently run on the VAX cluster. The dynamic simulator also runs on the SUN. The steady-state simulation uses a commercially-available static simulator, ASPEN, and another simulator currently being developed by Air Products. The memory and speed requirements of the static simulators are not substantial. The dynamic simulation of the SSC cryogenics system is done with SSCDYSIM, developed by Air Products. This simulator is a major user of computing power, with each run taking several hours on a VAX 8650. Radiation simulations, which account for a large fraction of the VAX cluster usage in this division, are discussed elsewhere in this report.

3.2.1 VAX USAGE

In FY88 the Accelerator Systems Division was responsible for approximately 10 percent of the total 0.15 VAX (8650) year (= 0.015) used by the CDG at LBL. Personnel growth should cause this to increase to 0.04 VAX years by the end of FY90. This does not include activities related to radiation transport codes.

3.2.2 WORKSTATIONS

Computer users in the Accelerator Systems Division will be moving away from terminals connected to the VAX, and towards workstations. It is anticipated that by the end of FY90, the 20 members of the division will be categorized as follows: 5 VAX users needing terminals (possibly small workstations like the SUN-3/50), 5 users needing small workstations (possibly SUN-3/50s or 386i's which run PC software), and 10 users needing to do substantial computing on a workstation. Two powerful file servers could accommodate 15-20 such individual workstations.

3.2.3 CRAY

The dynamic simulator for the cryogenics system is extremely compute-intensive. Typical CPU times are of the order of hours on either the VAX 8650 or the SUN-4/260. Also, there exists a powerful workstation-based interactive interface to the simulator, which is very helpful for watching the beginnings of long runs to make sure they are well behaved. For these simulations, it would be ideal to have a SUN front-end which could talk to a UNIX Cray. The Cray would do the number-crunching, and the SUN user would set up the input conditions and display and analyze the results. Existing Crays could be used for these simulations, if ones which are not heavily loaded can be found. It would be more convenient eventually to have a local supercomputer.

3.3 MAGNET DIVISION COMPUTING NEEDS FOR FY89 AND FY90

The computing in the Magnet Division can be broken up into four general areas: magnet quench test data analysis and management; finite element analysis; cable database; and prime item specification. Currently, the test data and analysis programs reside on the VAX cluster, with data transfer by DECnet from Fermilab and Brookhaven. The finite element analysis uses ANSYS on the VAX cluster and a dedicated workstation. The cable database resides on PCs and the prime item specification is implemented in Macintosh databases.

3.3.1 WORKSTATIONS

The Magnet Division currently has two SUN-4 servers, one of which is dedicated to ANSYS work, and three SUN-4/110 workstations at the CDG. One 4/110 has been ordered for Brookhaven to allow them to use SYBASE. In the first half of FY89, the test data and analysis programs, which account for approximately 60 percent of the Magnet Division's VAX use, will be moving to the SUN from the VAX . The cable database is in the process of being moved to the SUN, and more workstations will be required for this effort. More workstations will also be necessary for ANSYS users. We estimate that the two existing servers will be sufficient, but that 10 more workstations will be required by the end of FY90. These should have significant computer power, SUN-4/110 or equivalent. This estimate includes two workstations for prime item specification, although it is very difficult to predict needs in this area.

3.3.2 VAX

The Magnet Division currently accounts for 30 percent of the total 0.15 VAX (8650) years (= 0.045) used by the CDG at LBL. It is anticipated that this will grow only slightly to 0.07 VAX years by the end of FY90.

3.3.3 PERSONNEL

By the end of FY89, a full-time database manager will be required, as will a full-time database designer.

3.4 COMPUTING NEEDS FOR THE MAGNET TEST LABORATORY (MTL) FOR FY89 AND FY90

The MTL will have 10 test stands, with five commissioned initially, and one to be ready for data taking by the end of calendar 1990. This implies that much of the computer infrastructure must be in place, running, and maintainable by the end of calendar 1989. The functions to be performed are:

- Hardware control and monitoring through a multidrop bus such as CAMAC, Mil1553, token ring, etc.
- Data management and archiving.
- Online setup—data presentation and analysis.
- Offline analysis—comparison with simulation, trend analysis, etc.

Hardware requirements will therefore be:

- Two large disk servers with WORM devices for raw data storage. One will act as the main file server and database server, the other will be for analysis. Two SUN-4/280s or equivalent will do the job.
- One cryogenics control and monitoring workstation, presumably VME-based. SUN-3/150 or equivalent.
- Five workstations (one per stand) which give control and monitoring access (via multidrop), of the same type as the cryogenics control machine.
- Between 5 and 10 "general access" machines. These can be small workstations—say, SUN-3/60s or equivalent.

3.5 SUMMARY OF ACCELERATOR PHYSICS NEEDS FOR FY89 AND FY90

Below we summarize the VAX-type usage and SUN workstation totals anticipated for FY89 and FY90.

3.5.1 VAX

The VAX usage by FY90 for the three efforts below can be broken down as follows:

	VAX 8650 years
Magnet design and analysis	0.07
Accelerator physics	0.13
Accelerator systems and controls	0.04
Total	0.24

To supply this demand productively will require capability in excess of one VAX 8650. However, for FY89, one VAX 8650 should be sufficient. If the increase in workstations results in a significant decrease in VAX cluster usage, one VAX 8650 may be sufficient for FY90 as well.

A plan for providing VAX-type computing is presented in sections 7.1 and 7.4.

3.5.2 WORKSTATIONS

By the end of FY90, seven clusters of UNIX workstations should be in place for accelerator design. Currently there are three. The breakdown is:

Accelerator Physics — Two new clusters of six workstations two more workstations for the existing cluster, plus one more disk for an estimated total cost of $530 K

Accelerator Systems — Two new clusters. Ten small workstations costing approximately $5 K each, eight workstations at approximately $20 K each, two computer servers at approximately $50 K each, plus two file servers, with storage, at approximately $80 K each. Estimated total cost: $470 K

Magnet Systems — Nine new workstations at approximately $20 K each, plus one at $50 K. It may also be necessary to add more storage for these new users. Estimated total cost: $280 K

Magnet Test Lab — Two large servers at approximately $80 K each. Six workstations at approximately $20 K each, plus 10 small workstations at approximately $5 K each. Estimated total cost: $330 K.

Peripherals These clusters will need printers, and tape drives. Seven printers at approximately $25 K each, plus nine tape drives at approximately $5 K each, gives an estimated total cost of $220 K.

Software Software costs are very hard to estimate; nevertheless, we estimate $300 K. This is meant to include word processors, math libraries, ANSYS, SYBASE, compilers, simulation codes, MACSYMA, etc.

3.5.3 SYSTEMS MAINTENANCE AND SUPPORT

Currently we have approximately one half-time person maintaining the Accelerator Physics workstation cluster. However, this is somewhat misleading because we draw heavily on the networking and UNIX systems expertise of RTSS at LBL. In order to replace this expertise, and in view of planned expansion in the next two years, it will be necessary to hire one full-time System Administrator with considerable UNIX experience, and one UNIX systems expert by the end of FY89. At least two UNIX experts will be needed to help with systems administration and programming during FY90. The MTL will need its own support staff during FY90, which we estimate at one System Administrator and one full-time Programmer.

3.5.4 SUN HARDWARE SUPPORT

One full-time person will be needed in FY90 to handle hardware problems, updates, installations, orders, etc. An estimate of hardware support costs is 5% of total system cost, per year.

3.5.5 SDS, DATABASE SOFTWARE SUPPORT

None of the unification described in the introduction to this section will be possible without programming support. This support should span groups within the Accelerator and Magnet divisions. At least three full-time Programmers will be needed for this effort, two by the end of FY89 and one in FY90. These programmers will be in the Accelerator Division.

3.5.6 CRAY

MFE Cray usage will probably be approximately constant over the next two years. Part of the reason for our relatively small Cray use is that the MFE Crays are inconvenient to use. The operating system is cumbersome, it is difficult for programs to communicate directly between the SUN and the Cray, and the link is unreliable, especially for transfers of large files. If the SSC had its own Cray (or equivalent), running UNIX, it would be straightforward for programs on the SUN and the Cray to communicate data and commands, and the Cray would be more widely used. This would be ideal for the dynamic cryogenics simulator and for long-term tracking, for example. Access to a supercomputer running UNIX should be explored in FY89.

3.5.7 VAX AND NETWORKING SUPPORT

VAX and networking support personnel are outlined in section 7.4 and 7.5.

4. PHYSICS AND DETECTOR SIMULATION IN FY89 AND FY90

Members of the Task Force were convinced that a very strong simulation effort is required for the SSC physics program. The scale of this effort and its significance must be greater in quantity and quality than has been the case for previous large detectors. This necessity is not only the result of the size and cost of the SSC detectors relative to previous detectors, but also, more importantly, due to the complexity and rarity of the physics events at SSC energies. The simulation is motivated not only to determine the most cost-effective solutions to detector design, but to determine if a given design will even work at all.

We are encouraged by the circumstance that the tools required for these simulation efforts have been enhanced in the past few of years. Simulation software, from physics generators to detector response codes, are now available on a variety of computers and with a range of calculational detail that was desirable, but unavailable, when the current generation of detectors were designed. Work is continuing on the enhancement of tools and must be encouraged. In fact, it must be focused on the versions of the SSC simulation tools which will eventually be selected as standard.

The scale of the physics and detector simulation needs for the SSC has been studied at a number of workshops in the past years. This scale is astonishing to those who first encounter it. Two other factors have also influenced the plans developed by this Task Force. The first of these is the distributed nature of the HEP community and the second is the need to provide support for this community to become effective in this area.

Both the lack of computing resources dedicated to the simulation effort and the need to provide, document and consult on the software tools, suggest the development at the earliest possible time of a dedicated SSC computing support group. While the desire to maintain strong user groups is central to the solutions suggested, it is nevertheless required that new large computing capacities be developed in a managed and supportable way. We suggest that the generation and distribution of simulated events be centralized in a manner similar to current models of data generation and distribution in current large detector collaborations.

4.1 STATUS OF PHYSICS AND DETECTOR SIMULATION PROGRAMS

In this section, we summarize existing tools in three categories:

(1) *Physics simulation codes*, which create events as lists of particle types and momentum vectors according to a physical model;

(2) *Radiation transport codes*, which describe the propagation of particles and their interaction with the material of a detector. When applied to simulation of electromagnetic or hadronic showers, where most or all of the energy of an incoming particle is eventually dissipated in the detector through many generations of interactions and secondary particle production, these codes can take very large amounts of CPU time;

(3) *Detector simulation codes*, which combine a description of a detector in terms of geometry and material and simulations of the propagation, decay, and interactions of particles passing through the detector. The physics simulation codes can provide the input particles as whole events, or single particle passages of the detector can be simulated. Radiation transport codes are provided, often with some choice. Parametrizations of shower energy depositions are often provided because of the long execution times of the full shower simulations.

4.1.1 PHYSICS SIMULATION

The standard model provides a theoretical basis both for calculating the production of new particles and for estimating backgrounds for new signatures. Several Monte Carlo programs have been written to simulate complete events based on the standard model. ISAJET[1] and PYTHIA[2] enjoy the widest use among these programs. Both are based on perturbative QCD but attempt to describe complete events by including models for soft interactions and parton fragmentation. Both incorporate $2 \rightarrow 2$ parton processes plus leading-log QCD scaling violations and an approximation for $2 \rightarrow n$ processes from higher order QCD. They also include selected extensions to the standard model which may be of interest at present and future colliders, such as new generations of quarks and leptons, additional gauge bosons, and many scenarios of Higgses.

The primary distinction between ISAJET and PYTHIA lies in the description of fragmentation. ISAJET uses an independent fragmentation model derived from Field and Feynman,[3] in which each quark or gluon undergoes fragmentation independently of the fragmentation of other partons in the final state. PYTHIA takes as the fundamental fragmenting object the color flux "strings" that join pairs of partons, and therefore produces fragmentation that depend on correlations between final state partons. Some evidence from e^+e^- experiments favors the string fragmentation model, but not to such an extent that independent fragmentation has been abandoned.

Both PYTHIA and ISAJET are actively supported by their authors, who have responded to their own and others' interest in different physics processes by frequent additions to the programs. The result is a sometimes unpredictable patchwork of processes, particularly the exotic ones, which are handled by each program. In addition, there are inevitably errors present in different ways in each program, either because of programming "bugs" or because a guess was made to describe a process that had not yet been calculated.

Other physics simulations are also in use. These include FIELDAJET,[4] COJETS,[5] HERWIG,[6] EUROJET,[7] and PAPAGENO.[8] . The latter two provide only parton-level simulation and do not attempt to simulate fragmentation. Each of these may be preferred for more accurate or up-to-date treatment of particular processes, or for treatment of processes that are not available in ISAJET or PYTHIA.

4.1.2 RADIATION TRANSPORT AND SHOWER SIMULATION

Electromagnetic processes and electromagnetic showers are for the most part well understood and simulated. The EGS4 code system[9] generally reproduces test beam data well. Equivalent physics is contained in the electromagnetic part of the GEANT3 detector simulation program (see below).

Hadronic shower simulation programs must describe a more complex array of phenomena, ranging from particle production by high energy collisions with nuclei to the details of nuclear

---

1 F. E. Paige and S. D. Protopescu, in *Proceedings of the 1986 Snomass Summer Study.*

2 H. V. Bengtsson and B. Sjostrand, in *Proceedings of the 1986 Snowmass Summer Study.*

3 R. D. Field and R. P. Feynman, *Nucl. Phys.* B136 1 (1978).

4 R. D. Field, *Proceedings of the 1984 Snowmass DPF Summer Study*, 713.

5 R. Odorico, *Nucl. Phys.* B228, 321 (1983); *Comp. Phys. Comm.* 32, 139 (1984).

6 G. Marchesini and B. Webber, *Nucl. Phys.* B238, 1 (1984).

7 B. van Eijk, *Proc. of 5th Topical Workshop on $\bar{p}p$ Collider Physics.*, 165 (1985).

8 I. Hinchliffe, private communication .

9 W. R. Nelson, H. Himayama, and D. W. O. Rogers, SLAC-265 (1985).

excitation and low energy neutron emission and transport. Two sets of programs attempt to give complete descriptions of hadronic showers with full simulation of all interactions, though of necessity many approximations are employed. These are the CALOR system from Oak Ridge,[10] now being released in a updated form as CALOR89, and GHEISHA, developed at Aachen[11]. Though the details are different, both programs suffer from lack of detail in fundamental nuclear physics. This shortcoming shows up particularly in the model for particle production off nuclei and in the treatment of low-energy phenomena such as nuclear excitation and breakup and neutron production and interaction. Until the release of CALOR89 and GHEISHA version 8, it has been understood that GHEISHA generally simulated the high energy collisions more accurately and that CALOR was better at the low energy physics. The recent releases have attempted to deal with these recognized shortcomings. It remains true that reasonable assumptions have had to be substituted for hard physics knowledge in many places. Some tuning of parameters is necessary in each program to reproduce data. The present state of the art is such that the existence of two programs is useful to map the range of uncertainty in the calculation. Both programs could benefit from more extensive comparison of the calculations with data from actual calorimeters.

FLUKA,[12] developed at Leipzig, and its modification NEUKA,[13] to include evaporation neutron transport, aim for higher speed by a hybrid technique of detailed simulation of the hadronic part of the cascade and parametrization of the electromagnetic energy deposition and (in the case of NEUKA) neutron production and transport.

4.1.3 DETECTOR SIMULATION

The single example of a general-purpose detector simulation program allowing general detector geometry and full simulation of all particle interactions, including showers, is GEANT3,[14] developed at CERN over the last 15 years. The program allows detector geometry to be defined, using a library of elementary shapes, and organized in a hierarchical tree structure. Physical properties such as material constants and magnetic field can be associated with the geometric structures. Particles are traced step-by-step through the detector, with all physical processes such as decay, energy loss, multiple scattering, nuclear interaction, and bremsstrahlung simulated. An interactive graphics package can display the detector geometry, particle tracks, and detector digitizations. It is actively supported by the CERN DD Division, who have recently made the geometry routines run faster and are currently adapting the program for vector computers, in collaboration with SCRI and the University of Michigan.

Several physics simulation programs, particularly ISAJET and PYTHIA, have been interfaced to GEANT3. Electromagnetic interactions equivalent to EGS4 are built into GEANT3. The standard hadronic shower package used with GEANT3 is GHEISHA. Work is now underway to include the Oak Ridge CALOR89 package in GEANT3 as an option for generating hadronic interactions and showers.

While GEANT3 is a very capable program, some drawbacks can be pointed out. At high energies, the long execution time problems of detailed shower simulations are exacerbated by the powerful geometry routines with which it must interact. The general area of user interface could also use some improvement. It is a tedious job to code (or to change) a detailed detector specification.

---

10 T. A. Gabriel, *Proceedings of the Workshop on Detector Simulation for the SSC*,. ANL-HEP-CP-88-51, p. 50 (1988).

11 H. Fesefeldt, ANL Workshop (Ref. 10), p. 214.

12 P. A. Arnio, A. Fassò, H. J. Moehring, J. Ranft, and G. R. Stevenson, CERN TIS-RP/168.

13 H. Kowalski, ANL Workshop (Ref. 10), p. 61

14 R. Brun, F. Bruyant, M. Maire, and A. McPherson, CERN DD/EE84-1.

Two classes of detector simulation program allow faster computation than the standard GEANT3 with full shower simulation. One approach is to modify GEANT3 with the addition of a fast shower algorithm. Several large detectors have taken variants of this route. SLD, D0, and H1 have added parametrized shower calculations (both EM and hadronic) as options within their detector geometries. L3 and D0 have added "frozen shower" options, where a pre-generated library of showers can be kept on disk and replayed when needed to approximate the energy deposition of a particle.

Within the SSC context, an LBL group[15] has developed a parametrized shower addition to GEANT3 for the Large Solenoid Detector discussed in the 1987 Berkeley Summer Study.[16] It has been used by several people starting at the 1988 Snowmass Summer Study. An Argonne group[17] is developing a parametrized shower addition to GEANT3 in connection with an improved detector specification package, intended to allow parametrized shower calculations with easily changed detector geometry for use in conceptual detector design for the SSC.

The second approach is a "super-fast" program using parametrized showers and highly simplified geometry in a standalone program. The prime example of this approach is QFL[18] developed at FNAL. In the first instance, only spherical calorimeters were supported, but a recent upgrade allows more general shapes. The highly simplified geometry produces a factor of order 10 increase in speed compared to parametrized versions of GEANT3.

Although calorimetry has received the most attention, both because of its critical effect on the speed, and because of its central role in much SSC physics, there is at least one example of a detailed tracking simulation for the SSC. This has been done by a SLAC group[19] and has been integrated with the LBL parametrized version of GEANT3.

4.1.4 CPU TIMES AND DATA STORAGE NEEDS

There are various categories of physics and detector simulation that will be useful for the design of experiments at the SSC. These may be grouped roughly in order of increasing CPU time required per event simulated, as given below:

- Physics generation of pp collisions only—that is, stand-alone runs of PYTHIA or ISAJET
- Physics generation with a very simple calorimetric simulation, without any shower parametrizations
- Physics generation with a fast simulation using shower parametrizations or "frozen" showers, magnetic field and reasonably complicated geometries. It should be noted that the speed of this option will depend strongly on the rapidity coverage simulated and the geometry simulated.
- Physics generation with detailed tracking of charged particles to understand tracking requirements

15 E. S. Wang, private communication.

16 R. Donaldson and M. Gilchriese, eds., *Proceedings of 1987 Berkeley Summer Study.*

17 L. Price, private communication.

18 C. Newmn-Holmes and J. Freeman, ANL Workshop (Ref. 10), p. 285.

19 Gail Hanson, private communication.

- Physics generation with full shower simulation using detailed electromagnetic and hadronic shower codes but no detailed tracking simulation
- Same as above, but also including a detailed tracking simulation.

Examples of CPU times (in VAX 8650 units[20]) for various processes generated by PYTHIA v 4.9 are given below:

Physics process	CPU time per event (sec)
Minimum bias	0.75
Two jets, $P_T > 100$ GeV	1.5
Two Jets, $P_T > 1000$ GeV	2.1
Z + jet, P_T of jet > 100 GeV	1.3
Higgs → ZZ, Z → ee or μμ, M_H=400 GeV	1.2
Higgs → all, M_H=800 GeV	1.3

Thus only a small time is needed to generate an adequate sample of signal events. Cross sections for potential backgrounds are often much larger than those for signals, so large numbers of background events may be required.

Estimates of CPU time required for simulations with shower parametrizations, some geometric modeling of calorimetry and magnetic field, but not simulation of particle tracking, are necessarily less precise. Some typical CPU times (in VAX 8650 units) for the packages described above are given below:

(1) QFL. In Ref. 18 it is stated that this program requires 2.5 sec/event to simulate W-pair production (p_T of the W = 500 GeV and Ws decaying to quarks) at the SSC for a rapidity coverage of |4| units.

(2) LBL package. It takes 19.4 sec/event to simulate two jet events with $p_T > 1000$ GeV in a detector geometry representing a large solenoid detector[21] with a 2-T magnetic field and a rapidity coverage of ±6.

(3) For the same program it takes 6 sec/event to simulate a minimum bias event.

(4) Freeman and Newman-Holmes state in Ref. 18 that for the same event types as in (1) but with a full CDF simulation, it takes about 100 times longer, or about 4 min/event.

(5) ANL Package. An early version gave 1 minute per event for $p_T > 2$ TeV/c two-jet events, in a rapidity range of ± 5.

In the study of computing for high energy physics in the 1990s (Appendix B), it was assumed that a typical "fast" simulation required about 1 min/event, with related calculation (event generation of extra events that are cut away and analysis) increasing the time to 3 min/event. It can

20 A VAX 8650 is approximately equivalent to 5.4 VAX 780s.

21 R. Donaldson and M. Gilchriese, eds., *Proceedings of 1987 Berkeley Summer Study.*

be seen that, even for so-called fast simulation packages, there is a wide range of possible CPU requirements, depending upon the level of sophistication desired and required.

A GEANT-based simulation of a large central tracking system for the SSC requires about 4 min/event (VAX 8650 units) for a typical high mass Higgs event and about 30 sec per minimum bias event.[22] These estimates apply to a representative central tracking system covering ± 1.5 units of rapidity. Different systems will, of course, have different time requirements. This time does not include time for pattern recognition or track fitting, which is estimated to add about 10-20% to these estimates.

The time required to simulate with complete shower development depends on the physics process (mostly on the total energy to be dissipated in the calorimeters), on the complexity of the geometry, and on the low-energy cutoffs used in the simulation. This parameter space has not been explored in detail for the SSC. Typical values are 1.3 sec/GeV for EGS4, and 3 sec/GeV for GEANT3, using GHEISHA with a typical mixture of electromagnetic and hadronic energy deposition. The latter number gives a time of 3 hours per event for a high p_T event in which 4 TeV enters calorimeters in the rapidity range considered. As a check, we note that a full simulation of a typical event at LEP at the Z^0 in the OPAL or L3 detector is estimated to take 6 minutes (VAX 8650 units). Since the multiplicity in an SSC event can easily be 10–30 times greater (and the particles have more energy), it is easy to see that a simple extrapolation will yield hours per event at the SSC.

4.1.5 VECTORIZATION AND RELATED ISSUES

In the past, the software environment has been unfavorable for high energy physics computing on supercomputers because of the lack of vectorized applications, transportability problems, lack of support for basic software, and poor network connections. However, this situation has improved substantially over the last year or so.

4.1.5.1 LANGUAGES

For purely scalar applications, languages are not a difficult issue since all of the supercomputers have well-developed FORTRAN 77 compilers.[23] Problems arise, however, for vectorized codes since there is no common method of specifying vector operations. As a result, programs that are vectorized on one machine may not run or may be inefficient on another vector machine. In addition, vector applications are hindered because of the less well developed programming environments on some of the vector machines. Fortunately, vector applications can be written with a relatively modest extension to FORTRAN 77. Such an extension is being produced by a collaboration between SCRI, the FSU Computer Science Department, and LIF Lille.[24] This project will produce a preprocessor implementation (called "AFTRAN") of the array syntax proposed in the FORTRAN 8X standard.[25] Specific translations will be supported for Cray, ETA, and IBM 3090 computers, as well as for scalar machines. Thus, vectorized applications can be developed and debugged on scalar workstations with the resulting code running efficiently on a wide variety of vector machines.

---

22 Gail Hanson, private communication.

23 Note that partial translation from VAX/VMS to FORTRAN 77 is possible. See A. S. Johnson, "FORTRAN Preprocessors," *Comp. Phys. Comm.*, 45 (1987).

24 J. L. Dekeyser, C. Georgiopoulos, F. Hannedouche, G. Riccardi, J. Vagi, and S. Youssef, "The AFTRAN Vector Preprocessor Project," SCRI technical report FSU-SCRI-88-43 (1988).

25 "American National Standard for Information Systems Programming Language FORTRAN," S8(X3.9-198x) Draft 8, version 104, Computer and Business Equipment Manufacturers Association, 311 First Street, N.W. Suite 500, Washington, DC 20001-2178.

4.1.5.2 BASIC HEP SOFTWARE

Even though the supercomputer centers offer a range of scientific and graphical software packages, most programs in High Energy Physics assume a different base of software support. Over the last year or so, the essential software base has been ported to most of the operating systems at the supercomputer centers. Support for CERNLIB, BOS, EPIO, HBOOK, ZEBRA, and GEANT (scalar) is now available on Cray/UNICOS (CERN), Cyber/VSOS (SCRI), ETA/EOS (SCRI), and ETA/UNIX (SCRI). All of these packages are distributed through the CERN DD division.

4.1.5.3 APPLICATIONS

Since the vector supercomputers have scalar speeds ranging from approximately 20 MIPS to 40 MIPS per processor,[26] these machines are attractive for the purpose of ordinary scalar FORTRAN 77 applications. Of course, vector computers are more efficient for vectorized applications where typical factors between 2 and 20 can be gained if most of the processor time is consumed in vector operations.[27] Although scalar detector simulation programs contain no large vectors, it is now well understood that detector simulation Monte Carlos are completely vectorizable after extensive restructuring of the basic algorithms.[28] Here we review some of the scalar and vector application codes that have been run on supercomputers.

The most widely used simulation program in High Energy Physics is the GEANT code[29] from CERN. A joint effort is under way between SCRI, CERN, and LaSC[30] to produce a vectorized version of GEANT for the Cray XMP/48, ETA-10 and IBM 3090vf. A vectorized GEANT will be available in approximately one year with an estimated gain in speed of between a factor of 2 to 3 on the Cray XMP/48 and more than that on the ETA-10.[31,32] For an example of recent progress, Table 4.1 shows some timing benchmarks for the time consuming routine "GNTRAP," which traces a particle from the interior of a trapezoid to the surface. Note the substantial gains in speed from vectorization.

---

26 For reference, we take 1 MIP to be a rough measure of the speed of a VAX 11/780, which is approximately one million typical non-floating point instructions per second.

27 For competely vectorized code, the vector speed/scalar speed ratio is typically 2-3 for the IBM 3090, 5 for the Cray computers and 20 for ETA computers. Note, however, that the scalar speed of the Cray XMP is almost twice that of the ETA-10E.

28 *Proceedings of the SCRI School on Vector Computing in Experimental High Energy Physics*, June 1988, Florida State University.

29 R. Brun, F. Bruyant, M. Maire, A. C. McPherson, and P. Zanarini, "GEANT3 Users Guide," CERN report DD/EE/84-1.

30 LaSC is the Laboratory for Scientific Computations at the University of Michigan, Ann Arbor, Michigan.

31 J. L. Dekyser and C. H. Georgiopoulos, "Technical Report on the Vectorization of GEANT3," FSU-SCRI-87-69 and *Proceedings of the Workshop on Detector Simulation for the SSC*, L.Price, ed., Argonne National Lab., (1987).

32 J. L. Dekeyser, "Managing Data Structures on a Vector Machine—Application to GEANT3," *Proceedings of the SCRI School on Vector Computing in Experimental High Energy Physics*, June, 1988.

Table 4.1

Timing results for the GEANT routine "GNTRAP"

	Scalar Speed	Vector Speed	Vector/Scalar
VAX 11/780	90	—	—
Cyber 205	12	0.65	18.5
IBM 3090vf	5	3.9	1.3
Cray XMP/48	5	0.57	8.8
ETA-10E	8	0.31	25.8

Timing results for the GEANT routine "GNTRAP," which traces the path of a particle from the interior of a trapezoid to the trapezoid surface. 500 particles are traced per call to the subroutine to test the vector speeds. All times are CPU times for single processors given in milliseconds. The results are from Martyn Corden (SCRI).

MC4 is a vectorized detector simulation code intended to take full advantage of vectorization for moderately complex geometric situations such as single calorimeter modules or test beam simulations. The main components of MC4 are vectorized ray tracing[33,34] and a vectorized version of the electromagnetic interaction routines from GEANT3.[35] The code is now being tested in scalar mode comparing results with the EGS code[36] and with experimental results. MC4 is intended to function as a test bed for developing new simulation ideas, including the use of massively parallel computers, variance reduction techniques, Woodcock tracking, correlated sampling, and other methods for improving the speed of simulations. Once benchmarks are established with respect to EGS and experimental results, a group at SCRI intends to also use MC4 for high statistics simulation of the D0 and ALEPH test beam calorimeters and for simulation of a fiber tower calorimeter proposed for the SSC. The MC4 project is closely related to the effort to vectorize GEANT. In fact, most of the vectorized electromagnetic interaction routines can be used in the vectorized version of GEANT. Versions of MC4 will be optimized for the ETA-10, Cray XMP and IBM 3090.

In designing calorimetry for SSC detectors, it is often crucial to have an accurate representation of hadronic interactions and, in particular, the low energy neutron component of hadronic showers. There are several simulation codes which do this, including GEANT 3.12 with the new GHEISHA[37] and the HERMES system which has a version running on the Cray XMP.[38] The HERMES code incorporates HETC, EGS4 and MORSE in a unified system and has been used in designing the ZEUS calorimeter at DESY.

33 S. Youssef, "Vectorized Simulation and Ray Tracing," FSU-SCRI-87-63, and *Proceedings of the Workshop on Detector Simulation for the SSC*, L. Price, ed., Argonne National Laboratory, (1988).

34 S. Youssef, "Generalized Ray Tracing and Point Location," SCRI technical report FSU-SCRI-88-71 (August 1988).

35 MC4 has been produced in a collaboration between SCRI and the Laboratory for Scientific Computations at the University of Michigan.

36 *The EGS Code System*, Ford and Nelson, SLAC.

37 Results presented by several speakers at the International Workshop on Detector Simulation, KFA-Julich, October, 1988.

38 P. Cloth, D. Filges, R. D. Neef, G. Sterzenbach, Ch. Reul, T. W. Armstrong, B. L. Colborn, and H. Bruckmann, "HERMES: A Monte Carlo Program System for Beam-Materials Interaction Studies," KFA-Julich report 2203, 1988.

4.1.5.4 PRODUCTION EXPERIENCE AT THE SUPERCOMPUTER CENTERS

A small amount of experience has been accumulated in large-scale production running of experimental high energy physics programs at the supercomputer centers. The E711 experiment at Fermilab has done all of its Monte Carlo production and data reduction on the SCRI ETA-10E. The E711 data sample consisted of 420 6250 bpi magnetic tapes containing 14 million events. The reconstruction program for E711 took 11 seconds per event on a VAX 11/780, and 18 milliseconds per event on the SCRI ETA-10E with completely vectorized code.[39] The 420 tapes were processed in 34 days consuming 71 hours of single processor time on the ETA-10E. Although promising, the E711 experience represents a case in which a simple detector geometry was ideally suited to a vectorized reconstruction. This example may not apply easily to more complicated geometries and codes. Experience with Monte Carlo simulations includes GEANT studies of a D0 test beam calorimeter on a Cyber 205[40] and current production running of the "GALEPH"—GEANT-based simulation of the ALEPH detector—on the SCRI ETA-10E. Also, note that GEANT 3.12 is running under UNICOS on the CERN Cray XMP/48.

Although the experience with production running of high energy physics applications on supercomputers is limited, there have been no major problems and the supercomputer centers appear to have the ability to handle large data samples if the need arises.

4.1.6 PARALLEL COMPUTERS

The event-oriented nature of certain critical high energy physics computing problems (simulation and reconstruction of experiment event data) has proven to be particularly accessible to loosely-coupled parallel processing techniques.[41] In recent years the main emphasis has been on farms of single board computers based on 32 bit microprocessors supported by adequate if not excellent FORTRAN compilers. Systems with Motorola 68020 based boards, each delivering approximately 1 VAX 11/780 MIPS, developed by Fermilab's Advanced Computer Program (ACP) are now extensively used throughout the high energy physics community with installations in North and South America, Europe and Japan. At Fermilab several systems with a total exceeding 500 "nodes" are in operation (or about to be) and have been essential elements in the success, or anticipated success, of an increasing number of experiments. On a smaller scale, arrays of small processors emulating various IBM systems have also been used successfully in the analysis of data from high energy physics experiments. Also on a smaller scale, plans to use arrays of μVAXes for this purpose are being developed.

The major difficulty encountered in using the "first generation" ACP systems has been in the area of software development for extremely large experiment reconstruction codes while they are being changed rapidly in a large team effort centered on a VMS development and file service environment. The commonly anticipated problem of parallelizing code has proven to be a non-issue, being accomplished with relatively little effort. The difficulties have been associated with moving huge team codes from one FORTRAN/system environment to another. Since simulation efforts are based on relatively stable packages like GEANT with relatively little I/O, they will very likely be more amenable to operation on multiprocessor systems. As an example, one user who had not previously even seen ACP manuals, brought up the Lund Monte Carlo in less than 2 working days. A significant cost of such systems is in the I/O. The relative lack of I/O in simulation activities means that cost-effectiveness of farms is even higher than for reconstruction

---

39 See C. Georgiopoulos, *et al.*, *Nucl. Inst. Meth.* A249 (451) 1986, and C. Georgiopoulos, *et al.*, *Nucl. Inst. Meth.* A261 (493) 1987.

40 S. Linn, "Monte Carlo Studies of CCEM Modules in a Test Beam Configuration," D0 internal report 521, March, 1987.

41 For a a recent review, see I. Gaines and T. Nash, "Use of New Computer Technologies in Elementary Particle Physics," *An. Rev. of Nucl. and Part. Sci.*, 37, 177 (1987).

processing. This assumes, however, that sufficient memory is available to run efficiently the simulation codes.

Several commercial parallel processing systems are available and have been explored at computing research sites such as Argonne's ACRF. Commercially available processors which most clearly fit the needs of HEP Monte Carlo detector design are typically microprocessor-based MIMD type architectures with either shared-memory, bus-based or distributed-memory network-based architectures. Some examples which have been studied are the Alliant FX-8, the Encore MultiMax, the Sequent Balance and Symmetry series, and an ethernet network of Sun 3 workstations. These systems contain typically 8-30 processors of 1-10 MIPS pressing power each. Nearly ideal linear speedup with processor number for typical HEP codes (e.g., EGS4) is obtained in a FORTRAN environment, where the parallelism is implemented with coarse granularity (i.e., events) via user-inserted macro invocations. The resultant codes are then highly portable between MIMD parallel machines. These computers are typically UNIX based and provide an excellent development and debugging environment. They offer a processing solution whose cost lies between purely DEC VAX/VMS machines and ACP-like semi-commercial farms.

Commercial high performance technology trends in the areas of microprocessors, workstations, networks, operating systems, and compilers have encouraged a far more open environment in the Second Generation ACP Multiprocessor System now under development.[42] The essence of the Second Generation philosophical goal is to attain a seamless merging of the commercially available workstation/Ethernet-UNIX (or, when required, VMS) world at the physicist's desk with extremely high performance multiprocessor superfarms at the back end. Software tools are being developed to provide the basic high energy physics block transfer capabilities (SEND, GET, BROADCAST, ACCUMULATE) between any UNIX (or VMS) TCP/IP platform, specifically including workstations and individual, or classes of, back end processing nodes.

Even the back end nodes will be open to continuing commercial competition to allow users to weigh considerations of cost effectiveness against the operating system and compiler/debugger environment. The ACP is in fact developing a back end node based on the MIPS R3000 RISC processor on a VME board which will be commercially supported. However, this is not intended to preclude use of other potentially appropriate options on either early or later systems. In particular, products from system vendors like DEC (in both VMS and UNIX) and SUN and single board specialists like Ironics, Heurikon, FORCE, and others, are or will be available for incorporation in the ACP Second Generation Environment with varying amounts of integration effort.

The ACP development of a MIPS-based board is a result of an evaluation that the MIPS compiler and silicon represent the best all-around environment for high energy physics in the near term. The ACP MIPS board will be prototyped early in 1989 and is expected to run at approximately 15 VAX 11/780 MIPS/board on typical high energy physics reconstruction/detector simulation code. This corresponds to approximately three 8650 equivalents per board with a price expected to be in the neighborhood of $5000. (The price will depend on memory costs. The estimate is for 8 Mbyte. Expansion to 32 Mbyte is supported.) Several development environments are available or anticipated that naturally match the powerful MIPS-based multiprocessor farms. These include the MIPS M120, an 8650 class system, that is available at about $20 K, and a

[42] T. Nash, H. Areti, R. Atac, J. Biel, J. Deppe, M. Edel, M. Fischler, I.Gaines, R. Hance, D. Husby, M. Isely, E. Miranda, M. Miranda, T. Pham, T. Zmuda, E. Eichten, G. Hockney, P. Mackenzie, H. B. Thacker, D. Toussaint, "High Performance Parallel Computers for Science: New Developments at the Fermilab Advanced Computer Program," *Proc. Workshop on Computational Atomic and Nuclear Physics at One Gigaflop*, Oak Ridge, TN, April 14-16, 1988, Fermilab Preprint Conf-88/97.

workstation from Silicon Graphics also based on MIPS silicon and compilers. High energy physics code can be developed, compiled, linked, and fully debugged on both of these systems with final executables available for transparent downloading to the superfarm MIPS processor nodes. It is anticipated that DEC will announce shortly an MIPS-based Ultrix workstation. It is not clear yet whether DEC will configure this product so that it can be used in such a seamless fashion for development of code to be used on other MIPS-based systems such as the ACP MIPS node.

In Section 4.2.1, the CPU requirements for SSC detector simulation in FY90 is estimated to be 150 VAX 8650 equivalents. Using our present understanding of costs on this time scale, we believe that this need can be met with a 50-node ACP 2nd Generation System based on MIPS nodes. The cost will be about \$300 K for the basic system, \$100 K for disks, some tape I/O, network interfaces, etc., and \$160 K for 8 development systems (generously assuming only 5 FTEs on each). The total is just about \$600 K. Since the ACP is planning to commission a system of similar size at Fermilab during the summer 1989, it is on a feasible though aggressive schedule to anticipate that a similar system could be made available by the SSC for detector simulation in early FY90. The estimate for FY91 requirements is for an additional 550 8650 equivalents. A factor of 2 improvement in cost effectiveness on that time scale is anticipated (see below). The incremental cost in providing the additional computing power is therefore estimated to be (550/150) $\times$ 0.5 $\times$ \$300 K or about \$550 K plus \$150 K for additional I/O and development systems. The total for FY91 is thus estimated at about \$700 K.

The extraordinary cost-effectiveness of systems like that described above is a result of the great advances being made through the use of Reduced Instruction Set Computers (RISC), like the chips from MIPS, Inc. The impressive performance of RISC-based systems is now widely recognized as is the fact that improvements in RISC microprocessor performance will continue at a rate several times faster than that of conventional computer architectures in terms of both full-size mainframes and microprocessors. Generally accepted projections indicate that the cross-over in performance of VLSI RISC computers versus top end main frames (IBM 370 or DEC VAX series) will come in 1989 for silicon with a delay of a year or so for systems. Recognition of this reality explains DEC's recent moves in this area after more than a decade commitment to the VAX/VMS architecture. Several companies are privately and/or publicly promising RISC silicon in the 500-1000 VAX MIPS range in the 1993-1995 time frame. In addition to MIPS, major actors in this arena include Motorola, AMD, and SUN.[43]

The rapidly changing environment, with huge realized and promised performance improvements, makes it essential for high energy physics activities to be positioned to be as receptive as possible to the technology. This requires a new emphasis on portability of code and openness of systems. UNIX is the operating system which is recognized as meeting these needs as it crosses the boundaries between virtually all vendors, and is the unique operating system supported by the RISC and much of the workstation world. Although VMS will continue to be supported by ACP systems, there is a clear motivation to encourage HEP applications into more open environments.

4.2 MODEL FOR PHYSICS AND DETECTOR SIMULATION NEEDS

4.2.1 CPU REQUIREMENTS

Computing requirements for SSC detector simulation depend on many factors which are incompletely understood, including the physics and background processes to be studied, the

43 T. Nash, "Computing Possibilities in the Mid 1990s," *Proc. Future Directions in Detector R&D for Experiments at pp Colliders*, Snowmass,Colorado, July 5-7,1988 (to be published).

degree of detail required in the simulation, and the number of processes that are considered. Only a crude estimate of the total CPU power needed can be made, but it is clear that the magnitude is large compared to anything now existing and that a study of how to obtain it must begin immediately.

In hadronic collisions the ratio of interesting signal cross sections to cross sections for the raw backgrounds which must be studied in the simulation is often 10^2 to 10^3. For example, the ratios

$$\frac{\sigma(Z + Z)}{\sigma(Z + jet)}$$

and

$$\frac{\sigma(t + \bar{t})}{\sigma(jets)}$$

are both in this range. The ratio can be increased by perhaps another factor of 10 when resolution and other effects are taken into account. Ratios much larger than this probably involve such small tails of the Monte Carlo that they are not worth considering. Hence, to get a final background sample of 10^2 events with a signal/background ratio of order one requires studying a total background event sample of order 10^5 to 10^6 events.

If the background comes from real physics, for example from jets plus real leptons coming from heavy quark decays, then the required sample which must be passed through the detector simulation can be significantly reduced, perhaps by a factor of 10 or more, by filtering the generated events. This may also be possible for some detector-related backgrounds but not for all. We assume that typically 10^5 background events must be run through the detector simulation.

At least three different levels of detector simulation and analysis, with very different CPU requirements, are useful for different sorts of physics and detector questions:

Particle-Level Analysis: This simply uses the generated particles. It is often the first step in developing a strategy to reject backgrounds. The CPU requirements are small, on the order of 1 VAX 8650 sec/event, or 10^5 sec for a background sample. This can be combined with QFL-type calorimeter simulation at the cost of a factor of 3 increase in execution time.

Fast Simulation: This uses some sort of parameterized or other fast calorimeter simulation, and it may also include a tracking simulation. It allows one to study the effects of detector resolution and reconstruction, but not the detailed effects of edges or cracks in the calorimeter. The typical time is 100 VAX 8650 sec/event, or 1/3 of a VAX 8650 year for a background sample.

Full Simulation: Most of the time here is spent in simulating the electromagnetic and hadronic showers in the calorimeter. Such a simulation takes on the order of 10^4 VAX 8650 sec/event, or 30 VAX 8650 years for a background sample.

We think that 10 FTEs are required in FY89 to begin studying the physics and detector issues. This will grow to about 40 FTEs in FY90 for four collaborations preparing detector designs, and to 80 FTEs in FY91. Each FTE might analyze three such background samples per year together with a larger number of much smaller signal samples. The level of detail in the simulation will also increase in time as detector design details are addressed. To model this growth, we use the fast simulation time for FY89 and triple it in each of the next two years. Obviously this is only a crude guess, but it leads to the following (rounded) numbers:

Fiscal Year	FTE	VAX 8650 equivalents
89	10	10
90	40	150
91	80	700

While these CPU estimates are very large, we believe that if anything, they are underestimates: more than the projected 10^2 accepted background events may be needed, and a lot more full detector simulation may be required. A previous attempt to estimate these requirements was made at Snowmass '88 based on experience with simulation of $H \rightarrow b\bar{b}$ searches. That report, included as Appendix B, gave numbers similar to the ones presented here.

4.2.2. STORAGE REQUIREMENTS

Data storage requirements for SSC event simulation depend on the level of detail required. The basic information consists of the output of the event generator and information from the calorimeter out to rapidity $|\eta| \leq 6$ and from tracking out to $|\eta| \leq 3$. A typical $p_T = 1$ TeV/c event requires,[44] for example,

Generator information	30 Kb
Parametrized calorimeter simulation	50 Kb
Tracking simulation information	240 Kb

Vastly more information can be generated. For example, EGS requires about 40 Kb/GeV if all electrons in the cascade are saved, but not all of the basic information may be needed for many analyses. We therefore use 300 Kb/event for our estimates.

By assumption, each FTE generates and passes through detector simulation three runs per year of 10^5 events each. For this information each FTE will make use of two different classes of storage:

Archival Storage: We assume that all events generated for a full year should be saved in archival format. This requires 90 Gb per FTE, or less than 50 eight-millimeter video cartridges, which seems acceptable.

Disk Storage: One model is that the production for a single day must be staged on disk. The processing of a sample will occur over about 30 days, so this requires 1 Gb of disk storage per FTE. Alternatively, one might use the archive medium directly for processing. In this case, one would at least want to hold a DST for the complete sample on disk. This would probably require less than 1000 words/event, or 400 Mb.

Given these assumptions, we have the following storage requirements for SSC detector simulation:

Fiscal Year	FTE	Archive	Disk
89	10	900 Gb	4–10 Gb
90	40	3600 Gb	16–40 Gb
91	80	7200 Gb	32–80 Gb

44 Numbers supplied by G. Hanson, F. Paige, L. Price, and E. Wang.

Memory Requirements: Detailed simulation of complex detectors will require large program and array sizes. We have only anecdotal evidence for how big they will be, but several examples of simulations of present large detectors and of SSC detectors occupy 4 to 16 Mbyte. Over time, simulations will become more complex, so we tend to give more weight to the upper end of this range. An allowance should be made for at least 20 Mb per simulation process. The implementation will be different on virtual and fixed memory computers, but even virtual-memory computers should supply 10 Mb of physical memory for efficient execution.

4.3 AVAILABILITY OF COMPUTING FOR DETECTOR SIMULATION IN FY89 AND FY90

4.3.1 SURVEY OF LABORATORIES

A brief survey of HEP laboratories and research groups was conducted as part of the work of the Task Force. While some relevant computing may be available for the near term needs of SSC simulation, a significant part of the simulation needs cannot be met in this manner even in FY89. For reference, the installed base of computing (mostly VAX and IBM) in the HEP community has been estimated at 150 VAX 8650 units.

Typically, an SSC detector simulation effort is not viewed as a part of the existing HEP laboratory missions. Furthermore, there is not much unused capacity and what there is is expected to be saturated quickly by new rounds of data. The individual research groups, typically at universities, have been expanding their resources with clusters of workstations, when money is available. Nevertheless, it is difficult to see how these can contribute to major event generation tasks on the scale needed for the SSC.

There is no unused VMS-based capacity at the laboratories. There is some short term VM-based capacity at Fermilab, and some soon to be delivered to SLAC. However, these upgrades in capacity are meant to satisfy significant new needs associated with current and next running periods. Furthermore, these upgrades are not to be followed by significant additional upgrades at these or other laboratories in the FY89 or FY90 time frames. There is some small time available on the experimentally-oriented parallel-architecture computers at Argonne and Brookhaven National Laboratories.

Within the context of the above situation, there is some willingness of the multipurpose laboratories to provide some small, but useful, amount of computing for the SSC at their standard rates for in-house programs. Examples of these are specific proposals from Argonne and Lawrence Berkeley Laboratories. However, these standard computing facilities are not viewed as capable of addressing the massive simulation needs described in this document.

In addition to the HEP laboratories, other DOE laboratories were reviewed for new significant computing resources. Some capability (a few hundred hours) will be available at NMFECC as part of the overall SSC request. Beyond this, relevant resources were not found.

4.3.2 SUPERCOMPUTER CENTERS

One of the most difficult problems with simulation of SSC detectors is finding enough computer resources to perform accurate simulations of detector components and full detectors. The problem is particularly acute for simulation results needed to design the SSC detectors since results from a simulation are useless unless they are able to keep up with a schedule of design decisions. Here we discuss the possibility of meeting some of the computing needs for SSC simulation by taking advantage of the vector supercomputers available at the Department of Energy and National Science Foundation funded supercomputer centers.

Table 4.3(a) gives a rough indication of the hardware capabilities of the single Department of Energy funded center[45] and the five National Science Foundation funded centers[46]. The centerpiece of each of these institutions is one or two vector pipeline architecture supercomputers. These machines have scalar speeds ranging from roughly 20 MIPS[47] per processor for the ETA machines to roughly 40 MIPS per processor for the Cray YMP with between 4 and 8 such processors per machine. Of course, the full capacity of the machines is only realized for "vectorized" code where each instruction processes a vector of sufficient length. For vectorized applications, the typical vector-to-scalar speed ratio is 2 for the IBM 3090, 5 for the Cray XMP and 20 for the ETA-10E. To set a scale, note that the proposed need of 700 VAX 8650 units in FY91 (see section 4.2) would require three dedicated, eight processor Cray YMPs, assuming a factor of 3 for vectorization. All of the centers have plentiful mass storage capacity and convenient front end systems for high energy physics users (VM, VMS and Ultrix). Network connections to major centers of high energy physics research are limited to 56 Kb/sec with the possible exception of the San Diego Center which has a 1544 Kb/sec connection to both LBL and Fermilab.

During the late stages of designing SSC detectors, many issues will occur which need answers from Monte Carlo simulations on a time scale much shorter than one year. It is important, therefore, to consider the throughput of computing resources in addition to the average capacity needed per year. To cite a typical example, consider an allocation of 1000 single processor hours on a Cray YMP at one of the supercomputer centers, and suppose that the application is to run purely scalar code. The 1000 YMP hours is equivalent to about 40,000 hours on a VAX 11/780 or approximately 8000 hours on a VAX 8650. Thus, the 1000 Cray YMP hours can be matched with one dedicated VAX 8650s. Note, however, that the throughput for the VAX 8650 solution may be much less than the throughput of the 1000 hours on the Cray provided that the Cray time can be used at a reasonable rate. For example, if a 100 hour job can be run on a single YMP processor (thus using 1/8 of the capacity of an 8 processor machine), the result is obtained in 5 days while the 8650 needs 40 days for the same job.

In addition to the five NSF centers and SCRI, there are vector supercomputer installations at many other university-associated centers (e.g., Purdue University, Colorado State University, the University of Texas at Austin, the University of Mississippi, Ohio State University, and the University of Minnesota). For example, the Minnesota Supercomputer Center has a Cyber 205, a Cray 2 and a four-processor ETA-10E. The possibility of detector simulation for the SSC at such centers is worth further investigation.

Table 4.3(b) summarizes the advantages and disadvantages of using the supercomputer centers for SSC simulation studies. Note that two of the most important advantages have appeared in the last year and that when the vectorized version of GEANT is available, the point about high CPU capacity will become more convincing, especially for the ETA machines. The supercomputer centers appear to have the capability to handle large data samples and to do significant processing on front end systems if needed. The limited network bandwidth could be a serious impediment if data files need to be transferred to and from the centers.

---

45 The Supercomputer Computations Research Institute (SCRI) at Florida State University.

46 The Cornll National Supercomputer Facility (CNSF) at Cornell University; the John von Neuman Center (JvNC) at Princeton; the National Center for Supercomputer Applications (NCSA) at the University of Illinois at Urbana-Champagne; the Pittsburgh Supercomputer Center (PSC) at Carnegie-Mellon University and the University of Pittsburgh and the San Diego Supercomputer Center (SDSC)

47 For reference, we take 1 MIP to be a rough measure of the speed of a VAX 11/780 which is approximately one million typical non-floating point instructions per second.

Table 4.3(a)

	Current Supercomputers	Additions over the next ~1 year	Front Ends	Networks	Mass Storage	Allocation
SCRI Florida State	Cyber 205 [VSOS] ETA10–E(4) [UNIX] ETA10-Q(1) [EOS]	CM–2 (16–32 K) ETA10-G(4) [UNIX] ETA10-Q(2) [UNIX]	VAX8700 VAX 780 [VMS]	HEPNET MFENET SURANET SPAN 64 K to CERN	60 Gbytes	DOE/ FSU
CNSF Cornell	2 IBM 3090-600E [VM] 5 FPS-264 2 FPS-164 iPSC VX/d5			NSFNET	75 Gbytes	NSF/ CNSF
J_vNC Princeton	2 Cyber 205 [VSOS] ETA10-E(4) [EOS] →	ETA10-E(8) [EOS] ETA10-G(8) [EOS]	2 VAX 8600 [VMS] 1 VAX 8600 [ULTRIX]	NSFNET	20 Gbytes	NSF/ J_vNC
NCSA Univ. of Illinois	Cray XMP/48 [CTSS] 2 Alliant FX/8 Cray 2 [UNICOS]	CM-2	2 VAX 785s [VMS]	NSFNET	60 Gbytes	NSF/ NCSA
PSC CMU/Univ. Pittsburgh	Cray XMP/48 [COS] →	Cray YMP-832 [UNICOS] → Cray 3	2 VAX 8650 [VMS]	NSFNET	25 Gbytes 30 Gbytes	NSF/ PSC
SDSC San Diego	Cray XMP/48 [CTSS / UNICOS] SCS-40 [CTSS] S-1/E	→ Cray YMP-832	VAX 780 VAX 785s [VMS]	NSFNET HEPNET SPAN	13 Gbytes front + 60 Gbytes	NSF/ SDSC

Table 4.3(b)

Advantages and disadvantages of SSC simulation at the Supercomputer Centers.

Advantages	Disadvantages
• High CPU capacity, High throughput POTENTIAL.	• System software quality and support are not as good as, e.g., VMS.
• HEP software base is now available on several systems (VM, VSOS, EOS, UNIX, and UNICOS).	• Bandwidth over networks is limited to 56 Kbits/sec except for LBL ↔ SDSC ↔ FNAL.
• AFTRAN version of FORTRAN 8X will soon allow codes to be transportable between supercomputers; develop vector code on workstations.	• A proposal has to be written and accepted for each center.
• Plentiful disk space; large data sets can be handled.	• You must learn a new operating system unless you know VM or UNIX.
• Front end systems are convenient for HEP, i.e., VM, VMS, Ultrix	

The software base assumed by many High Energy applications is now available on many of the systems as demonstrated by production running of GEANT on the Cray XMP/48, ETA-10 and IBM 3090. For new vectorized applications, AFTRAN will allow vector code development and debugging to be done on scalar machines and insure that the resulting code is transportable across several systems.

Another supercomputer-related issue is the question of whether 32-bit floating point calculations are adequate to do simulation of SSC detectors. For example, a drift chamber with 40-micron resolution at a distance of 40 meters from the interaction region has about the same resolution as 32-bit floating point numbers. This question should be carefully investigated since it may affect the hardware able to do the bulk of the simulations for the SSC detectors.

Probably the most serious disadvantage of computing at supercomputer centers is the fact that system software quality and support is not as good as the High Energy Physics community expects. This situation dictates extra care and extra work when using the supercomputers and creates a danger of wasting manpower in using software which does not function correctly. For this reason, groups contemplating simulation at the supercomputer centers would do well to consult the high energy physics group at SCRI, CERN or elsewhere since they have accumulated substantial experience.

The supercomputer centers offer a promising way to provide computing capacity for SSC simulation studies provided that adequate network connections can be established and provided that sufficient time can be allocated at the centers for SSC related work.

5. RADIATION TRANSPORT CODES

Detailed simulations of the radiation environment of the SSC are necessary for proper design of the accelerators and detectors. Increasingly detailed designs for scrapers, injection lines, IR optics, detectors, etc., will create an ongoing need for such calculations. To date essentially all of the these calculations have been done by consultants to the CDG or by visitors to the CDG. Only recently have CDG personnel (one person) become capable of running one of the high energy transport codes (FLUKA87)[48] for productive calculations.

5.1 TYPES OF CODES

There are three types of codes, separated for operational reasons:

(1) *Electromagnetic cascade simulation.* Although other codes exist, EGS4 is the most commonly used code.[49] While it is very consumptive of time, biasing and other time-saving techniques are gradually being introduced. The next release is advertised to be substantially faster.

(2) *High-energy hadronic cascade simulation.* These codes are "knowledgeable" about high-energy processes such as fragmentation, resonances, and hadron-nucleus collisions, but lack detailed nuclear physics information. Even so, they are quite long; FLUKA, for example, consists of about 50,000 lines of FORTRAN. Some are inclusive, in the sense that weighting schemes are used and only selected particles are followed, while others follow every particle to its demise. The former kind are adequate for problems where only an average behavior is required (such as machine design problems), while the exclusive codes are necessary if fluctuations are important. Every high-energy hadronic cascade has an important electromagnetic component, which can be incorporated in an average way or by an explicit call to a code such as EGS4. Typically, these codes transport particles down to energies where nuclear physics effects become important (e. g., 50 MeV).

A FLUKA simulation of a 200 GeV cascade in a simple lead target takes about 4.9 seconds of VAX 8650 time. In the test calculation, average shower behavior was used instead of a full EGS4 simulation. Since a larger fraction of a higher energy shower goes into the electromagnetic channel, the growth of execution time with energy is somewhat less than linear. It will take about 7 minutes for the full simulation of a 20 TeV hadronic cascade.

(3) *Low-energy transport codes.* Because low-energy charged particles quickly come to rest, only neutrons and gamma-rays are of importance. Neutron behavior is of enormous importance for both calorimeter performance and electronics survival in the machine. These codes must acess large files of cross section data and be capable of dealing with time evolution, since pulse integration times in calorimeters are in general shorter than the time scale for energy deposition.

5.2 CPU REQUIREMENTS

As a result, in part, of the limited involvement of CDG personnel in the actual running of transport codes, we do not have a good estimate for the CPU requirements associated with the use of these codes. Furthermore, the number of tasks or problems to be done in the next two years is

---

48 P. A. Arn, A. Fassò, J. Lindgren, J. Ranft and G. R. Stevenson, "Enhancements to the FLUKA86 Program (FLUKA87)," CERN TIS-RP/190 (1987); see also P. A. Aarnio, A. Fassò, H-J. Möhring, J. Ranft and G. R. Stevenson, "FLUKA86 User's Guide," CERN TIS-RP/168 (1986).

49 W.R. Nelson, H. Hirayama and D. W. O. Rogers, "The EGS4 Code System," Stanford Linear Accelerator Center Report SLAC-265 (Dec. 1985).

difficult to specify in detail. Typically, one formulates a problem, debugs it with a series of low-statistics runs, and then makes production runs. The results are then reduced and described. Only the production runs take substantial computer time, while the other phases are labor intensive. As a result, CPU requirements scale as some mixture of the number of tasks and the number of active people.

Several benchmarks have been established as an aid in estimating the near-term CPU time requirements:

- T. A. Gabriel, of Oak Ridge National Laboratory, is a major developer and user of the high-energy transport code HETC[50] and the low-energy neutron transport code MORSE[51]. He estimates his usage at 5 to 7 hr/week (on an IBM 3090 or Cray XMP), reaching 10 to 15 hr/week during periods of maximum effort. Since he has been using evolving versions of the same code for about two decades, these figures are probably higher than would occur in more exploratory situations.
- The Oak Ridge group simulated beam loss in the Tevatron and SSC to provide estimates of the neutron flux in the SSC arcs[52]. In dealing with the problem, a total of about 70 hours of Cray XMP time was required. The problem was done twice to obtain more detailed information that was at first thought necessary, but this is probably typical. Given the time spent with the problem, this is consistent with the 5 to 7 hr/week estimate given above.
- The two leaders of the radiation physics group at SLAC (Jenkins and Nelson) used 10.2 hours/month of IBM 3081 time during FY88, or 1.2 hr/week each. This scales to about 0.6 hr/week of IBM 3090 or Cray XMP time. They were involved in other activities which included writing a book during this period.
- Three visitors to the SSC Central Design Group used SLACVM during the fall of 1987. During a total of 7 person-months they used 143 hr of 3081 time, or 4.7 hr/week each. Again, this scales to 2.4 hr/week of 3090 or XMP time. Although the three people were all doing program development and a sequence of calculations using FLUKA, the range of their usage was a factor of six.

On this basis, we conclude that 2.5 hours of supercomputer time per week would realistically be used by an SSC physicist treating a sequence of radiation problems. An equal amount of time on a slower but more user-friendly computer would be used for program development and debugging.

---

50 G. Alsmiller, Jr. et al., "Modification of the High Energy Transport Code (HETC) and Comparisons with Experimental Results," presented at the ANS Topical Conference on Theory and Practices in Radiation Protection and Shielding, 22–24 April 1987, Knoxville, Tennessee; K. C. Chandler and T. W. Armstrong, "Operating Instructions for the High-Energy Nucleon-Meson Transport Code HETC," Oak Ridge National Laboratory Report ORNL-4744 (1972).
W.R. Nelson, H. Hirayama and D. W. O. Rogers, "The EGS4 Code System," Stanford Linear Accelerator Center Report SLAC-265 (Dec. 1985).

51 M. B. Emmett, Oak Ridge National Laboratory Report ORNL-4972 (1975).

52 T. A. Gabriel, F. S. Alsmiller, R. G. Alsmiller, Jr., B. L. Bishop, O. W. Hermann, and D. E. Groom, "Preliminary Simulations of the Neutron Flux Levels in the Fermilab Tunnel and Proposed SSC Tunnel," SSC Central Design Group Report SSC-110 (1987).

The approach through FY90 will have two main thrusts:

(1) Continuation of collaborative efforts. Many of our consultants (those from CERN and SLAC) use IBM computers operating under VM. In some cases the codes have not been transported to DEC VMS, and in any case it would be exceedingly inefficient for the consultants to spend much of a relatively short visit learning a different system. For this reason, access to an IBM computer using this operating system is essential. The average occupancy of such individuals may be about 1/2, so an average IBM 3090 equivalent usage of 1.2 hr/week should be budgeted for FY89 and FY90.

(2) Continuing development of laboratory capability. A major present handicap is that the one CDG user of the codes is able to make such calculations only on an occasional basis. It is recommended that an additional full-time SSC employee (a physicist/programmer) be added as soon as possible in FY89 to assist this effort, and that another person (more experienced) be added in FY90.

In this scenario we would anticipate only incidental use of a supercomputer during FY89, so that together with visitors the usage would average 2.0 hr/week. "Supercomputer" is not defined, other than for the requirement that the visitor's share be on an IBM machine operating under VM. In FY90 it would rise to 3.0 hr/week. The equivalent of a VAX 8650 would be used 5 hr/week in FY89 and 10 hr/week in FY90.

6. NETWORKING

This section describes an aggressive network plan based on the assumption that it is important to establish (early at the SSC site) high quality communications with other parts of the SSC effort and with the high energy community at large. Because of the need to determine "real" costs for networking which are geographically dependent, one must assume that the SSC will be sited as proposed in the Dallas area.

It is necessary to have some lead time (typically, a minimum of 90 days) to order leased lines and apply for network memberships. Equipment can be ordered in advance as soon as funds are available. A system which can communicate with DECnet, NSFnet, BITnet, and MFEnet should be purchased and prepared for immediate installation. A VAX could be used for this purpose (for example).

The following are the requirements and assumptions upon which we based this work.

- Significant computing will continue to be done at LBL at least into FY90.
- The SSC will provide a significant part of its own computing on-site by the beginning of FY90.
- There will be a local 8XXX class VAX system for detector and accelerator simulation and design. This system will also serve as a repository for data, code, and results.
- There will be multiple SUN clusters used for various purposes at the SSC site.
- Communications between the SSC site and other major HEP facilities should include video teleconferencing, and private phone circuits for local voice PBX systems.

- Communications with the actual SSC site will be required as a physical presence is established.

6.1 OVERVIEW OF WIDE-AREA NETWORKS AND THEIR STATUS

The SSC should use existing network facilities to minimize both the cost and support effort of networking. In other words, the goal should be to exploit existing networks and not to create a "new" network. Because of the diversity of protocols required, no single network connection can provide all of the needs, so multiple network connections will be necessary.

Based upon the current computing resources being used, it appears that TCP/IP, DECnet, MFEnet, and RSCS (BITnet) are the protocols that would be necessary to supply network connectivity. If the SSC were to use a system that was not accessible using the above protocols (and associated networks), special accommodations would have to be made.

Table 6.1

NETWORK	ACCESS	PROTOCOL(S)
HEPnet	All HEP labs and most universities involved in HEP (USA, Europe, Japan)	DECnet/X.25/RSCS
NSFnet	All NSF-funded supercomputer centers and many DOE laboratories	TCP/IP
MFEnet	All DOE-funded supercomputer centers and many DOE laboratories	MFEnet/MFE-IP
BITnet	Thousands of universities and many laboratories (USA, Europe, Asia, Mid-East	RSCS

There are several wide-area networks that should be of interest to the SSC. Each network offers different accesses, protocols, and topologies. The SSC will be interested in the networks described in Table 6.1. Each of these networks is undergoing major up-grades and changes within the time frame of FY89-FY90. These upgrades should negate any major impact SSC utilization may have on these networks. Because of the changes in progress, a brief review of the current state and future plans of each of these networks follows. ESnet provides the foundation for both HEPnet and MFEnet in the U.S.A.

HEPnet HEPnet currently runs over a 56 Kb backbone between SLAC, LBL, FNAL, BNL and MIT. Tail circuits to universities and other labs typically run between 4.8 and 14.4 Kb and attach to the backbone sites. There is a 64 Kb circuit to CERN, and L-3 also has a 19.2 Kb circuit to CERN. There are 9.6 Kb circuits to Italy, Canada and Japan.

UPGRADES

In FY89 the backbone will be upgraded (via ESnet) to about 400 Kb. Several tail circuits are being upgraded to 56 Kb. Internationally, the L-3 link to CERN, Italy, Japan, and Canada will upgrade to 64 Kb or 56 Kb. HEPnet supports DECnet, RSCS (running over DECnet) and X.25 (used for terminal to mainframe traffic and Coloured Books).

NSFnet

NSFnet currently runs over T-1 circuits which connect all of their supercomputer centers. NSF regional networks vary in their robustness and hardware but are generally upgrading to higher speeds (including T-1). The NSFnet provides direct connectivity to ARPAnet and Milnet. The resulting Internet connects over 100,000 systems in university, research and government labs.

UPGRADES

The NSFnet backbone will move to higher speeds (greater than T-1) when it has the hardware to do so. NSFnet only supports TCP/IP but several regional networks are running or considering to run DECnet as well.

MFEnet

MFEnet I supports access to the MFE computer center and runs over a 56 Kb backbone. MFEnet is based on its own protocol which runs on Crays and VAXs. The next generation of MFEnet (MFEnet II) uses IP but to date is not interoperable with INTERNET routing and naming. However, it should work with any TCP/IP system given the ability to locally control IP routing. Tail circuits vary in speed but most are 56 Kb.

UPGRADES

The backbone will move to higher speeds (approximately 500 Kb) this coming year. MFEnet II (soon to be on-line) will support access to all DOE supercomputer centers.

BITnet

BITnet is an RSCS based network which has an extensive distribution in the U.S.A. and Europe plus some connections to Asia (Japan) and the Middle East.

UPGRADES

BITnet will probably merge soon with CSnet thus creating a very large and powerful network. BITnet II (which is RSCS running over IP) is in operation in beta test at several universities. The combined BITnet /CSnet will probably create a new network topology most likely based upon TCP/IP.

6.2 CURRENT SSC NETWORKING AT LBL

CDG personnel with personal computers or workstations have access to large-scale computing via Develcon or the LBL Ethernet backbone. The connection to the Ethernet depends on the user's type of computer. The various SUN clusters and single machines connect directly to the Ethernet as described in Section 2, above. The IBM PCs also connect directly to the Ethernet. Macintoshes, on a local Appletalk network, are attached to the Ethernet through a Kinetics FastPath 4 gateway. (See Figs. 6.2[a] and 6.2[b])

Users also use the LBL VAX cluster, NMFECC or, in a few cases, the SLAC computers via asynchronous terminal connections through the Develcon switch or by remote modem dial-in, either from compatible terminals or from personal computers using terminal emulation software. The Develcon will be replaced in early 1989 by data connections through the new ICS phone system.

The current IBM and Macintosh local area networks and connections to the Ethernet backbone are shown in Figure 6.2(a) (SUN configurations are shown in Section 2, above).

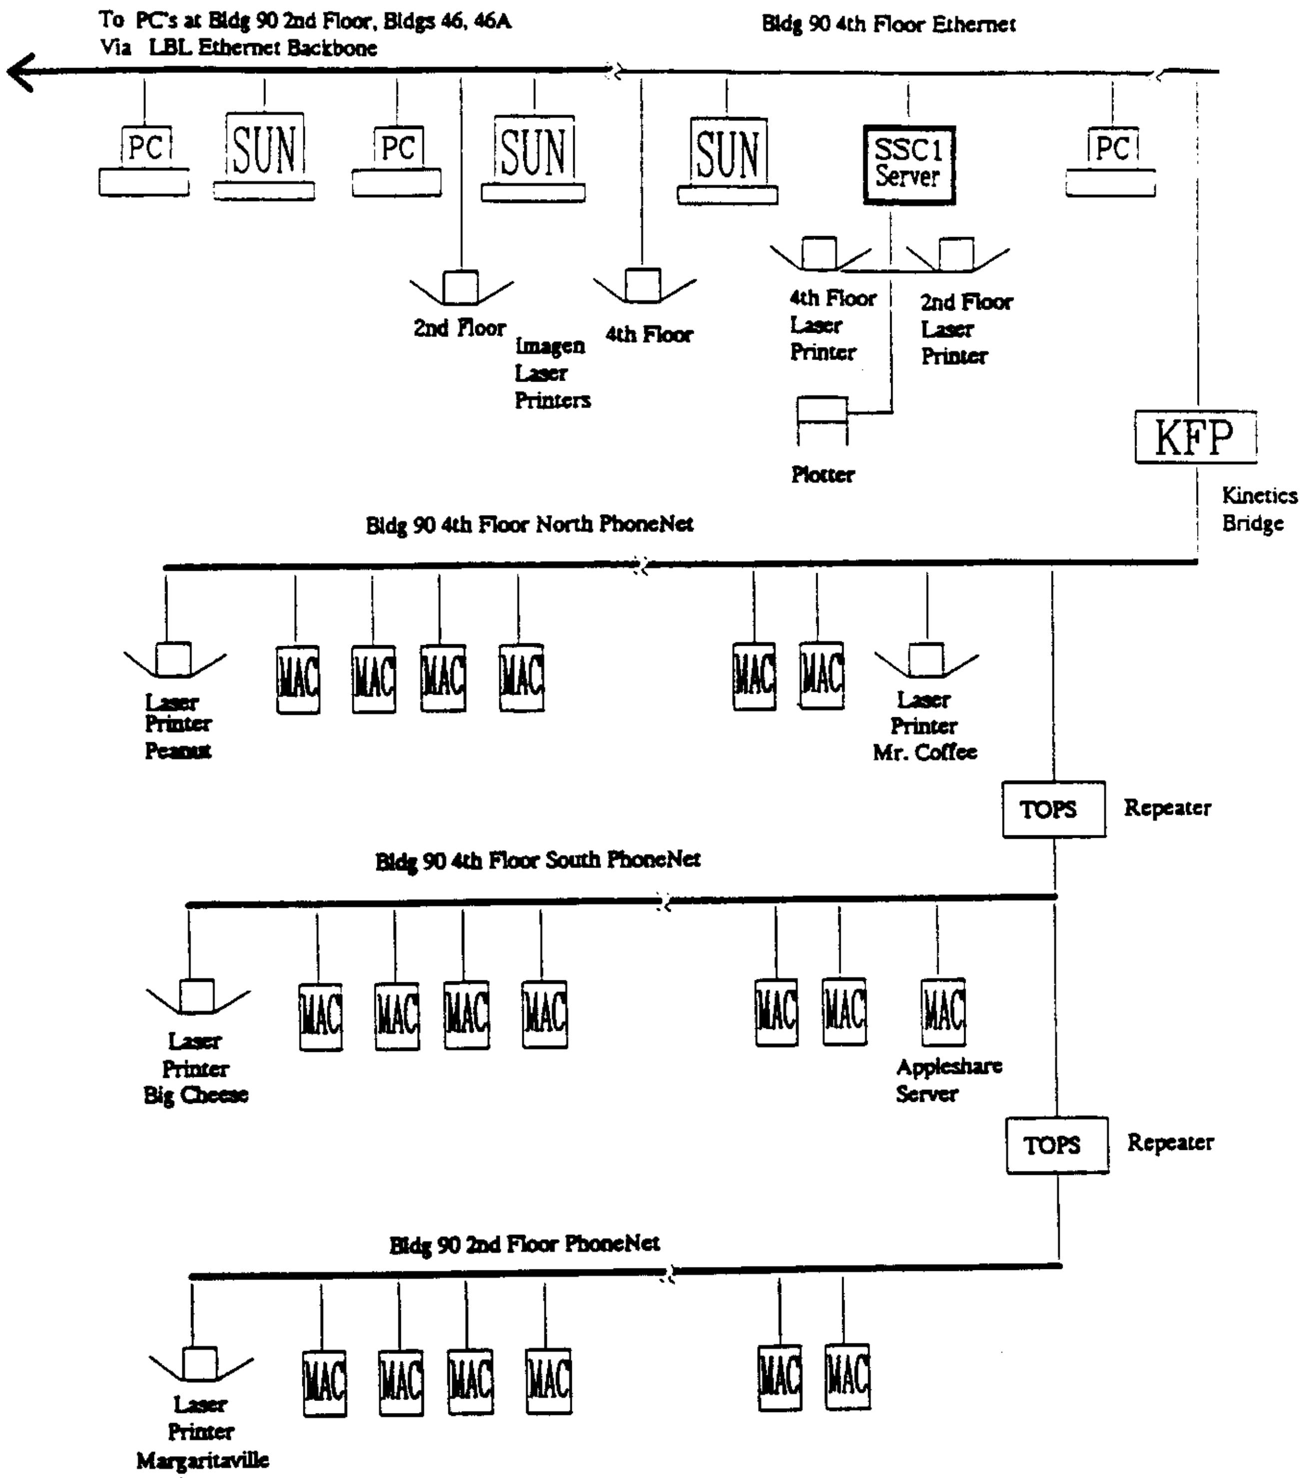

Fig. 6.2(a). SSC Local Area Networks — December, 1988.

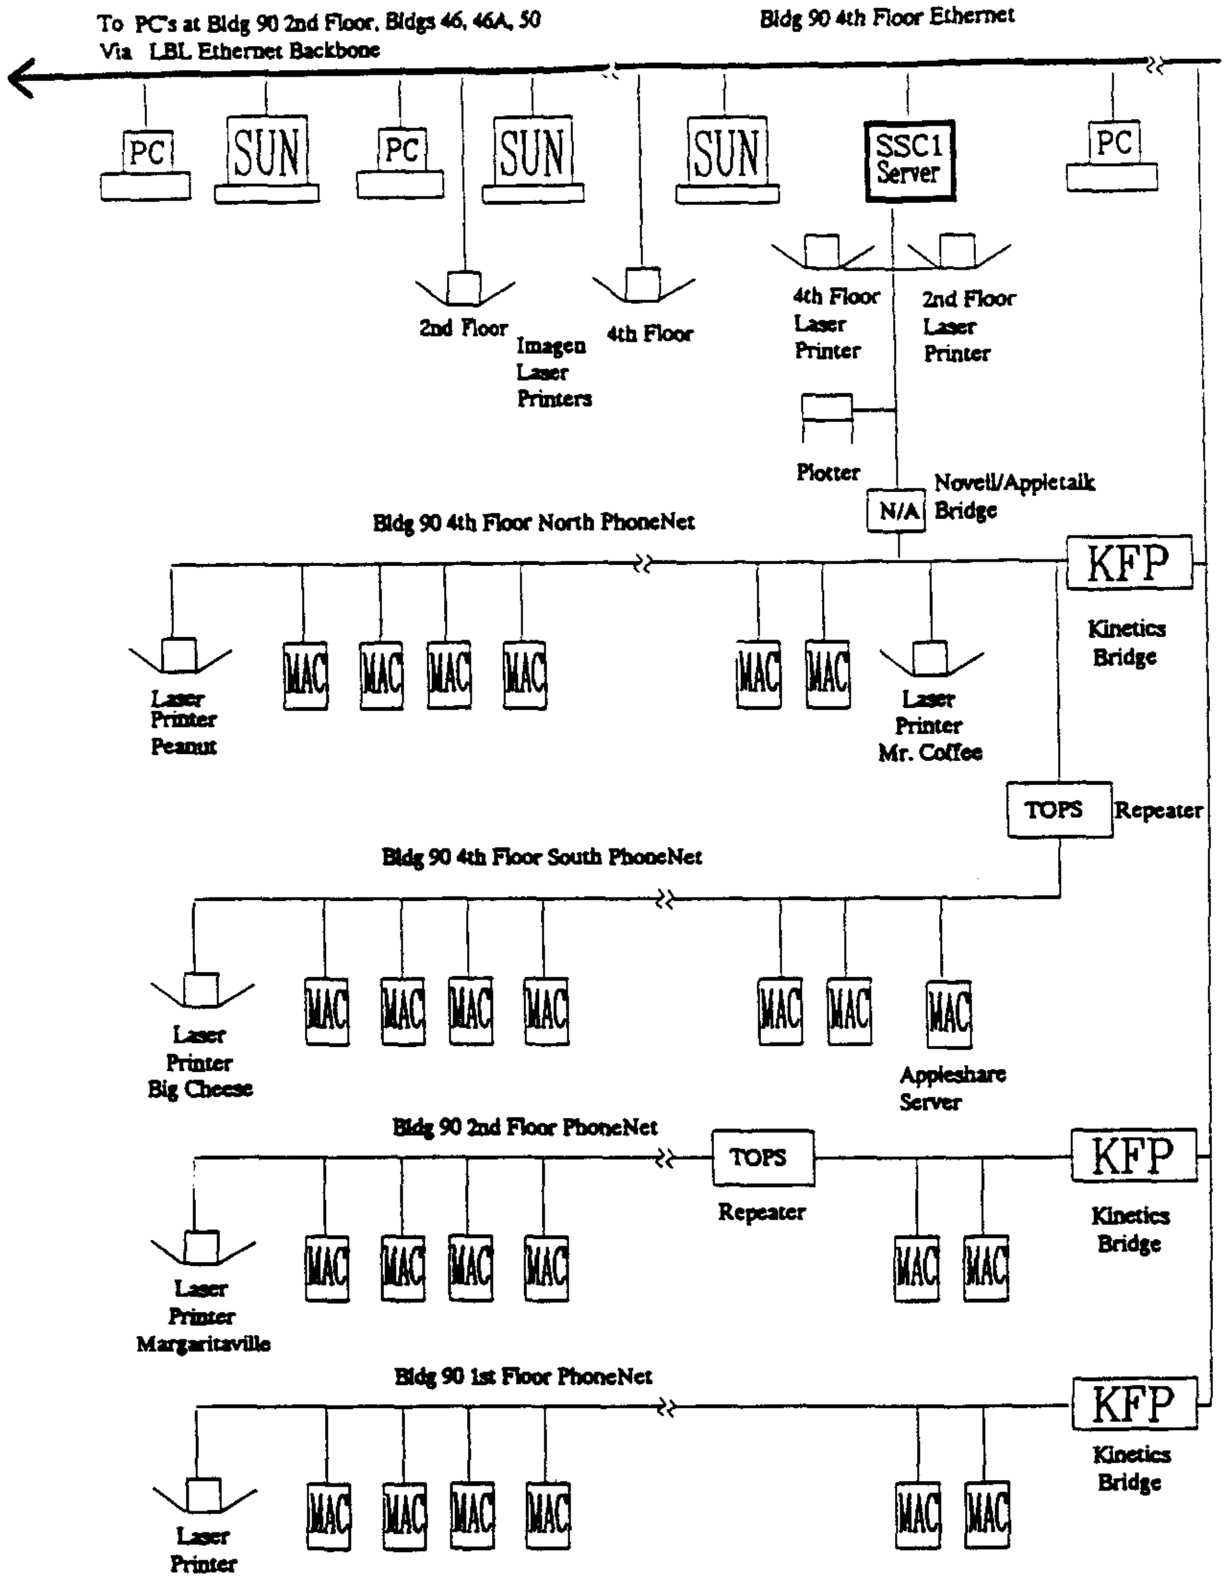

Fig. 6.2(b). SSC Local Area Networks — Early 1989.

6.2.1 MACINTOSHES

The Macintosh net has 66 Macs, four Apple LaserWriters, a Macintosh II Appleshare server, and one Kinetics KFP2 gateway to the Ethernet. This network, spread over the fourth and second floors of Building 90, has two TOPS repeaters to boost signal strength over the long lengths. The Mac net also uses one of the SUNs as an Appleshare file server, running Columbia's AUFS software. Currently this network is used for TCP/IP communications, for LaserWriter printing, for file exchange, and for shared database applications.

6.2.2 IBM PCs

There are approximately 15 IBM and equivalent systems currently used at the CDG location. Most are 8-MHz AT systems configured with 30 Mbyte hard disks, Enhanced Graphics Adapters, two floppy disks (1.2 Mbyte and 360 K), and 1.5-3 Mbyte of RAM. Several clones are in use with more appearing soon. One IBM PS-2/60 and at least one XT are also in use.

6.2.2.1 NETWORK CONNECTIONS

Develcon and modem connections are widely used for communications with other systems. Most of the systems have 3COM Ethernet cards connected directly to LBL Ethernet backbone. This allows access to the SSC1 IBM server and network TCP/IP operation. Four IBMs have assigned IP addresses on LBLnet and use NCSA Telnet and FTP for network access and file transfers.

The SSC1 server is an enhanced IBM AT with Novell LAN server software and a 150 Mbyte hard disk. The software will support 100 users with configurable individual and group file and directory access privileges, passworded user logins, and several levels of usage accounting and security.

Two HP LaserJet printers and one Bruning size E plotter are connected to the server. Each is assigned a print queue which may be managed by any user with printer console privilege. (Currently assigned to all).

The server communicates with its clients via the LBL Ethernet backbone directly using its 3COM Etherlink plus Ethernet card. Communication packet protocol is APX, a subset of the Xerox XNS protocol suite.

Current services performed are:

- shared database for accounting
- laser printer sharing
- plotter sharing

The system is used by the following CDG Divisions:

- Project Planning
- Conventional Facilities
- Administration
- Magnet Division

The server is located on the 4th floor of building 90, but the client IBM systems currently configured for server logins are widely scattered over LBLnet:

- 5 are on the 2nd floor of building 90
- 6 are on the 4th floor of 90
- 2 are in building 46A
- 1 is in building 46, and
- 1 is in building 50.

6.2.3 LOCAL AREA NETWORKING AT LBL IN FY89 AND FY90

Not shown on the diagram but under consideration is the installation of additional IBM servers. In particular, the CDG Accounting Group may require the additional computing power, printer and disk service, and isolation that a second server would provide. Addition of another Novell server to the LAN is accomplished simply by connecting the server to the Ethernet. The necessary client workstation software is already incorporated into the current systems.

6.3 NETWORKING TO SSC SITE

The plan to provide communications services to the SSC site is phased relative to the FY89, FY90 boundary. In FY89, most of the computing resources will come from the LBL VAX cluster, thus requiring high performance networking between the two. Costs for leased lines are given in Table 6.3. An overall schematic of SSC networking is given in Figure 6.3.

6.3.1 FY89

In FY89, most of the computing resources will come from the LBL VAX cluster, thus requiring high bandwidth networking between the two locations. In FY90, there will most likely be a major shift from using resources at LBL to using those at the SSC site. This transition is not understood in detail, but the continued need for a T-1 circuit to LBL will be weighed relative to the networking costs plus computing costs versus establishing a larger VAX computing system at the SSC site. In Table 6.5, we have assumed that there will continue to be a need for an LBL T-1 circuit through the end of FY90.

The T-1 line connecting the SSC site will provide access to/from the LBL computing facility, HEPnet, MFEnet, telephone systems, and video conferencing facilities. A T-1 multiplexer will be used to separate these (incompatible) functions. The T-1 multiplexer will have three channels, one for data (DECnet, IP, and X.25), one for telephone (inter-PBX systems), and one for slow-scan video teleconferencing. If ESnet operates as a "normal" IP network, then there will be no problem using the data channel to gain access to the ESnet gateway at LBL. If ESnet will not support this, then there will have to be an ESnet-supplied line installed from the SSC site to Austin (an ESnet terminus).

Sesquinet (the Texas NSFnet regional network) and BITnet access should be established through a connection to the University of Texas - Dallas campus. We have contacted the Sesquinet manager, who has informed us that there will be a T-1 trunk running to the University of Texas in Dallas from Austin and Houston this coming spring. The Austin location provides a T-1 link to the NSFnet backbone. The UT Dallas connection also provides connectivity to the Texas Higher Educational network (THEnet). THEnet is a multi-protocol Texas wide network which supports both TCP/IP and DECnet. Due to addressing conflicts between THEnet and HEPnet, total inter-operability may not be possible with DECnet; however, THEnet personnel are working on

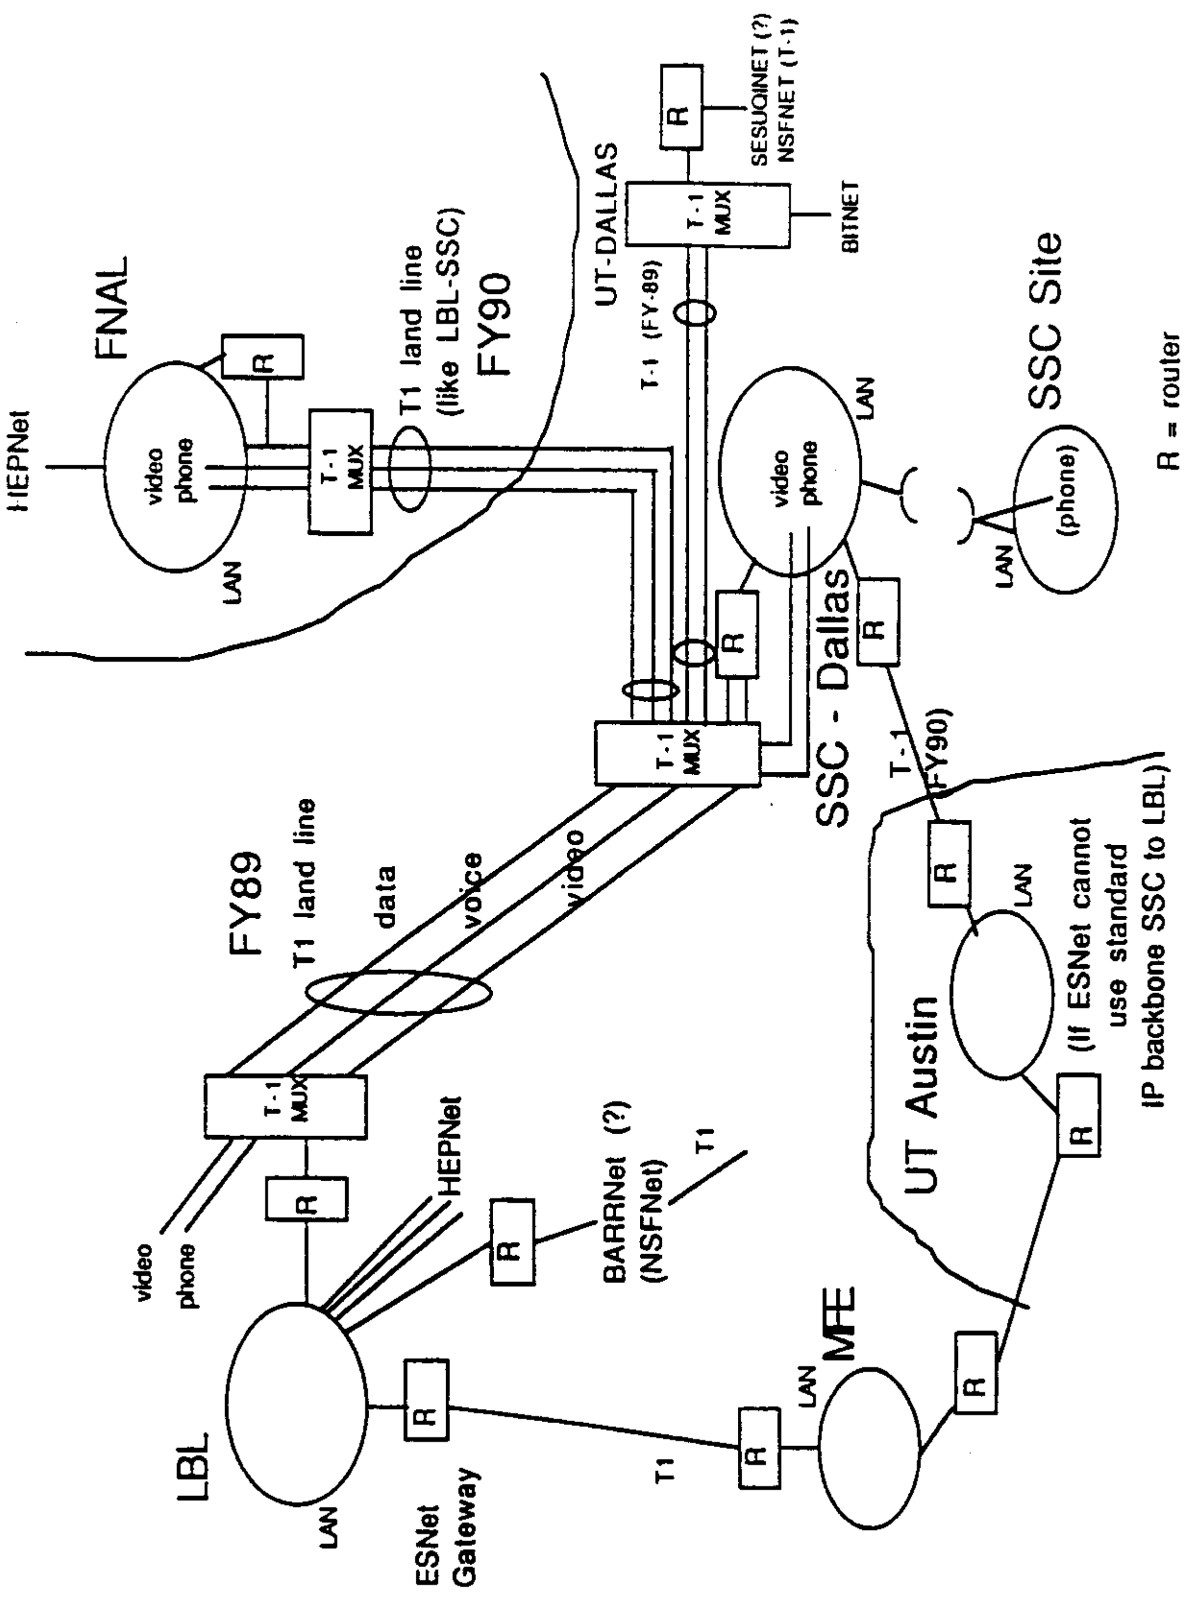

Fig. 6.3. SSC Networking FY89/90.

software which should be in place by the time the site is to be occupied. THEnet is an ideal means by which the SSC personnel can communicate with most Texas universities. Alternatively, the data portion of the LBL T-1 circuit could be used to access BARRnet through its on-campus T-1 connection to NSFnet. In any case, IP traffic bound to LBL will be directed through the LBL T-1 link.

Table 6.3

SPEED	FROM/TO	COST FY-89 (K$)	COST FY-90 (K$)
56 KB	DALLAS TO LBL	24	47
T-1	DALLAS TO LBL	104	207
56 KB	DALLAS TO FERMILAB	20	39
T-1	DALLAS TO FERMILAB	86	172

*For six months in FY89

6.3.2 FY90

In FY90 we assume the expanded program and personnel will require that a T-1 connection be established between the SSC site and Fermilab for support of detector design/development and magnet design, construction, and measurement. This T-1 arrangement would be identical to the one proposed for LBL and used for the same purposes. MFE access through the ESnet gateway at LBL (soon to be operational) could be used through the LBL T-1 link (if allowed by ESnet) or through the Internet (if MFE continues to maintain a functional Internet connection to the Crays). MFE access could be obtained from a direct ESnet connection with a line to Austin (as previously stated). It is expected that telephone (PBX) and video conferencing would be established in FY90.

6.4 NETWORKING AT THE SSC SITE LOCATION

We assume that the initial operations at or near the SSC site will be established in an office-type environment similar to that now occupied by the CDG at LBL. If this environment is off-site and if a need arises to communicate with the actual site, leased lines or a microwave link would be used for that purpose.

An Ethernet Local Area Network (LAN) will have to be established in the office facilities occupied by SSC personnel. The LAN will provide the means by which communications services will be delivered to the personnel's workstation, terminal, PC, or Macintosh. If there is to be a population of 200 to 300 persons, there is a good possibility that several floors of a building or even multiple buildings may have to be wired. An appropriate wiring plant should be established if the same office complex is to be used for several years. An organized wiring arrangement would make creating NFS, LAVC, PC, and Appletalk sub-LANs easier to create and modify. The LAN (in conjunction with multi-protocol terminal servers) could provide "dumb" terminal access to systems at the SSC site for dial-in access.

6.5 NETWORKING COSTS

We have assumed that the requirements for networking have been identified as an essential component for SSC development in its early stages. Unlike other situations where networking was an add-on, after-the-fact detail, our assumption is that networking is a prerequisite to effective productivity. The plan as outlined establishes the SSC site with a full array of networks by the end of FY89. (Table 6.5)

Table 6.5

ITEM	FY89 (K $)	FY90 (K$)
NETWORK HARDWARE/SOFTWARE	160	70
LEASED LINES	104	379
TOTALS	264	449
PERSONNEL (FTEs)	3	5

6.6 HARDWARE/SOFTWARE

At least two routers will be needed to connect the SSC site's LAN to the wide-area networks. In FY89 (for THEnet and NSFnet), a Cisco router should be used. In FY90, communications access to Fermilab may require an additional DEC router 2000.

As stated previously, a VAX system used as a communications hub would be very desirable. To handle the potential load, this system should be a μVAX III with sufficient disk to buffer BITnet mail and file activity. This system's software should include Joiner's JNET (for BITnet access), SRI's Multinet (for TCP/IP access), DEC's PSI (for X.25 access), and possibly PC link or Alisa Talk for Macintosh and PC wide area access. A SUN workstation might also be purchased to provide TCP/IP dedicated name service for the establishment of an SSC domain on the Internet. All of these systems should be purchased in FY89.

Two T-1 multiplexers will be needed in FY89 between LBL and the SSC site, and two more in FY90 for T-1 connections to Fermilab and UT Dallas. Of course, this assumes that T-1 speeds are desired. A wiring plant for the SSC office facilities should be established if the same facility is to be used for more than a few years. This wiring plant includes Ethernet, wire distribution, individual office drops, media conversion devices and equipment, Ethernet backbone isolation devices (routers and/or bridges), and wiring frames.

6.7 PERSONNEL

There should be a person designated to design and coordinate all of the local and wide-area networking for the SSC group. He or she will need at least two additional FTEs to install the described facilities in FY89, and an additional FTE to expand and maintain the networks in FY90. These three FTEs would also be used to support computing facilities at the SSC site and/or workstation NFS/LAVC/Appletalk file servers.

6.8 CONCLUSION

Providing networking services for the SSC/CDG is straightforward. However, the networking associated with the personnel involved in the SSC design will be more diverse than the HEP community in general, and thus is more demanding. Planning will be the key for providing a productive network environment for SSC personnel in both the short and long terms.

7. SSC COMPUTING PLAN FOR FY89 AND FY90

7.1 ACCELERATOR PHYSICS, SYSTEMS, MAGNETS AND MTL

The Accelerator Division anticipates having access to the LBL cluster and subsequently to a leased VAX 8820 in FY89 and FY90. Several clusters of SUN workstations will also be purchased during this time frame. Below we split the costs associated with these systems into totals for the two fiscal years, FY89 and FY90. We have been very conservative in new workstation purchases for FY89, in line with realistic budget constraints. However, before moving to the SSC site, it may be desirable to purchase some of the workstations planned for FY90, shifting people away from the VAX at that time.

Systems support costs, hardware support, software support, and networking support are listed in Section 7.4.

VAX REQUIREMENTS
(in VAX 8650 equivalent years)

FY89:	Current LBL Cluster and Leased VAX	0.24
FY90	Leased VAX	0.24

TOTALS:

FY89:	0.24 VAX 8650 years
FY90:	0.24 VAX 8650 years

CRAY REQUIREMENTS

FY89:	1100 hours at MFE
FY90	1100 hours at MFE

If a Cray running UNICOS were available over a high-speed link, our Cray usage would increase considerably, perhaps doubling over the next two fiscal years. Only the inconvenience of use of the MFE Cray keeps our request constant.

UNIX WORKSTATION REQUIREMENTS

HARDWARE COSTS:

Accelerator Physics

FY89:	1 disk drive @ $10 K	$10 K	
	2 workstations @ $20 K	40 K	
			$50 K
FY90:	2 clusters of 6 workstations each @ $240 K		$480 K

Accelerator Systems

FY89:	1 file server @ $80 K	$80 K	
	1 workstation @ $20 K	20 K	
			$100 K
FY90:	1 file server @ $80 K	$80 K	
	2 computer servers @ $50 K	100 K	
	7 workstations @ $20 K	140 K	
	10 workstations @ $5 K	50 K	
			$370 K

Accelerator Totals

	FY89:		**$150** K
	FY90:		**$850** K

Magnet Division

FY89:	1 workstation @ $50 K	$50 K	
	1 workstation @ $20 K	20 K	
			$70 K
FY90:	8 workstations @ $20 K	$160 K	
	Disk storage	50 K	
			$210 K

Magnet Division Totals

	FY89:		**$70** K
	FY90:		**$210** K

Peripherals:

FY89:	2 printers @ $25 K	$50 K	
	4 Exabyte tape drives @ $5 K	$20 K	
			$70 K
FY90:	2 printers @ $25 K	$50 K	
	3 Exabyte tape drives @ $5 K	$15 K	
			$65 K

Peripheral Sub-Totals

	FY89:		**$70** K
	FY90:		**$65** K

Magnet Test Lab

FY89:	1 file server @ $80 K	$80 K	
	3 workstations @ $20 K	$60 K	
	3 workstations @ $5 K	$15 K	
			$155 K
FY90:	1 file server @ $80 K	$80 K	
	3 workstations @ $20 K	$60 K	
	7 workstations @ $5 K	$35 K	
			$175 K

Magnet Test Lab (MTL) Totals

	FY89:		**$155** K
	FY90:		**$175** K

Peripherals for MTL

FY89	2 printers @ $25 K	$50 K	
	1 tape drive @ $5 K	$ 5 K	
			$55 K
FY90	1 printer @ $25 K	$25 K	
	1 tape drive @ $5 K	$ 5 K	
			$30 K

Peripheral Sub-Totals for MTL

FY89	$55 K
FY90	$30 K

TOTAL WORKSTATION COSTS:

FY89:	**$ 500 K**
FY90:	**1330 K**
	$1830 K

SOFTWARE COSTS

Software costs in FY89 are expected to be approximately $100 K. This includes purchase and/or leasing of programs such as MACSYMA, ANSYS, and SYBASE, and word processors, math libraries, compilers, simulation codes, etc. We expect that in FY90 these costs will increase to $200 K.

Software support personnel will be needed to maintain and enhance tools such as the database, SDS, and simulation interfaces. For FY89, two such people will be needed, and one more should be hired in FY90. One database manager and one database designer will also be needed in FY89.

7.2 RADIATION TRANSPORT CODES

Goals for the FY89 and FY90 period include:

- implementing software within the SSC laboratory for radiation transport calculations and to further develop capabilities to productively interface with consultants in this area
- application of this capability to urgent problems in the design of the SSC accelerators and detectors: dose in the distributed correction coils, neutron fluences in electronics enclosures, radiation field near a scraper, etc.
- support of detector simulation activities that require radiation transport codes

Access to a substantial IBM mainframe running VM will be necessary during FY89 and FY90 for productive use by consultants. To this end, access from the SSC site to a large IBM mainframe running VM would be highly desirable. We expect that VM system time usage will average 3 hr/week/consultant (in IBM 3090 single processor units) when such consultants are in residence (perhaps 1/3 of the time). Other supercomputer usage will rise from occasional use in FY89 to about 10 hr/week by late FY90. VAX 8650 (or equivalent) usage will total about 5 hr/week in late FY89 and 10 hr/week in late FY90.

Support of radiation transport calculations will require additional personnel to be added to the SSC laboratory in FY89 and FY90. These are:

- a physicist/programmer with demonstrated skills in FORTRAN and experience with large codes.
- a physicist with experience in using radiation transport codes.

7.3 PHYSICS AND DETECTOR SIMULATION

7.3.1 PLAN FOR FY89

The model presented above in Section 4.2 shows computing for detector design at the rate of 10 VAX 8650 equivalents by the end of FY89. We plan for the use of seven VAX 8650 years during the remaining 75 percent of FY89 and the provision by the end of FY89 for computing at the rate of ten 8650 equivalents. Because it is important to provide some computing for general use by the community as soon as possible and to provide computing at the SSC site by the end of the 1989 fiscal year, it is necessary to make use of computing from a variety of sources. We recommend the following mix of computing be provided in FY89:

- Existing HEP computers. A small amount of this computing can be done on computers at the universities and labs, as local physicists spend a portion of their research time on SSC detector problems. We estimate that this may account for about one 8650 year in FY89.
- Allocations at supercomputer centers. The possibility of using time at the DOE and NSF supercomputer centers should be explored vigorously. It will be necessary to make proposals and to provide some coordination and guidance to new users on the appropriate versions of standard codes and instruction on the different operating systems. It is thought possible that two 8650 years (each 8650 year is estimated as about 1000 hours on a supercomputer) could be obtained.
- Purchase of time on non-HEP computers. This appears to be a useful way to get substantial amounts of computer time to the community quickly and is important if the time at supercomputer centers is limited as expected. We consider that early availability of "SSC" computer time is important, both to encourage an increased level of work on detector simulation and design questions and to begin to measure the level of interest and need for this computing. Two approaches have been presented to this task force:
  - The ANL proposal makes use of time on an existing CRAY XMP computer supported by the Argonne Computer Center. A rough equivalence is given by 1000 CRAY hours per VAX 8650 year.
  - The proposal from LBL envisions an approach based on farms of μVAX III workstations in a local area VAX cluster arrangement. Money is saved by not supplying software (such as compilers) on each workstation. The result is quite low-cost computing, but will not supply all needs, e.g., the interactive data analysis will not be practical on this platform but can be done on the LBL cluster or on VAX computers at other institutions..

We recommend exploring both of these options in order to provide computing resources immediately and to gain experience with the possible use of farms of microprocessors in a computing center environment. A possible scenario is to buy an initial 2000 CRAY hours at ANL and begin installation of the LBL MICROVAXES. Later in the year, the usefulness of the MICROVAX system can be assessed, along with the need (if any) for more computing at ANL.

- Acquisition of computers by the SSC Laboratory. To come up to the desired level of ten VAX 8650 units by the end of the fiscal year, the equivalent of four VAX 8650s should be leased or purchased. If good networking is available to the SSC site, all of this acquired computing could be installed there and made available for outside use. Otherwise, at least two 8650 units should be located at the SSC site, while the other two could be located at existing labs or universities. The VAX 8820 appears to be a good choice, providing approximately the equivalent of two VAX 8650s initially and being expandable to up to four 8650s equivalent.

  In addition to the computer resources described above, there will be a need for computer systems people and software-oriented physicists to provide support and software to make the computers usable to the physicists carrying out the detector simulation at many different institutions. Coordination of the software, documentation, fixing of bugs, helping users solve problems, and keeping the computer systems running are examples of the functions these people will perform.

  Such people are needed as soon as computer time becomes available for general use in FY89. We estimate four FTEs will be required.

The cost of the purchased time and hardware for the plan described above would be as follows:

Initial two VAX 8650 equivalent years at ANL	$600 K
Initial two VAX 8650 equivalent years at LBL, including storage and use of existing central facility	$225 K
Lease two VAX 8820s and peripherals for six months	$300 K
TOTAL	$1125 K

7.3.2 PLAN FOR FY90 AND BEYOND

The requirement of 150 VAX 8650 equivalents available by the end of FY90, and *a fortiori*, the requirement of 700 in FY91, cannot be met at an attainable cost by solutions that are presently familiar in the HEP community. The clearest conclusion we can draw is that work must begin immediately to identify and gain experience with a new computer architecture.

The Task Force considered the use of the present generation of computers (e.g., CRAY YMP) of scalar mainframes and mini-supercomputers, but found that all were unaffordable at the required level of computing. The only architecture we are aware of that offers the possibility of providing enough computing within attainable budgets is a parallel one making use of the new generation of Reduced Instruction Set Computer (RISC) microprocessors. These may involve multiple computers on a bus, as in the Silicon Graphics computer server or networked single processors in workstations based on new RISC chips such as Apollo DN10000. We note that this approach is being used in the second generation ACP under development at Fermilab. It is important to recognize the considerable uncertainty of this approach, since it has not previously been tried in production HEP settings on the scale required. It requires early study for this reason.

As a specific example, we consider the Silicon Graphics solution, with four processors in a package giving a total performance of 12 VAX 8650 equivalents. One of these systems, with additional memory and disk space, costs about $150 K. Thus, the FY90 level of 150 VAX 8650 equivalents could be met with 12 systems, costing $2 M. The FY91 level of 700 8650s requires 60 such systems, costing about $10 M. We emphasize that the feasibility of this or similar systems has not yet been demonstrated.

The major portion of this new computing should be located at the SSC site. Depending on the state of networking, however, it may be appropriate to locate the 1/3 of the detector simulation computing that is devoted to analysis at selected labs and universities throughout the community so it will be close to the users during the stage where responsive computing is most needed. In this model, the event generation and detector simulation would be done on the concentrated computers at the SSC site, and the data sent, e.g., on 8-mm helical scan tapes, to distributed locations for analysis.

In order to identify the appropriate approach to computing to have in place at the end of FY90, it will be necessary to hire people to study the problem, make the choice, and begin the process of installing standard software and supporting the HEP users in this unfamiliar mode of computing. At least one such person should start working in FY89, and more early in FY90. By the end of FY90, when the computers are installed in quantity, the support group will require at least five FTEs. Sample quantities of a tentatively-chosen computer should be in trial use by the second quarter of FY90.

Provision of computing during FY90 will need to represent a transition from the interim provisions of FY89 to the new architecture chosen during FY90. Thus, at the beginning of the year, support could be continued for computing at the LBL VAX cluster, at the SSC site in association with the 8820s, or elsewhere. By the second half of FY90, expansion should come by beginning the acquisition of the new-style computers.

7.4 NETWORKING AND COMMUNICATIONS

7.4.1 FY89

In FY89, most of the computing resources will come from the LBL VAX cluster, thus requiring high-performance networking between the two locations. This would be provided by a T-1 circuit with T-1 muliplexers which would allow data, voice and video communications. The LBL link would provide HEPnet and MFEnet networking to the SSC site.

A T-1 connection to THEnet at UT Dallas would provide Sesquinet (and Internet) access as well as access to BITnet and the Texas-wide DECnet.

An Ethernet Local Area Network (LAN) will have to be established in the office facilities occupied by SSC personnel. We assume that the initial operations at or near the SSC site will be established in an office-type environment similar to that now occupied by the CDG at LBL. The LAN will provide the means by which communications services will be delivered to each workstation, terminal, PC, or Macintosh. Costs are summarized in Table 7.4.

Table 7.4

ITEM	FY89 (K$)	FY90 (K$)
NETWORK HARDWARE/SOFTWARE	160	70
LEASED LINES	104	379
TOTALS	264	449
PERSONNEL (FTEs)	3	5

7.4.2 FY90

In FY90 the expanded program and personnel will require that a T-1 connection be established between the SSC site and Fermilab for support of detector design/development and magnet design, construction, and measurement. This T-1 arrangement would be identical to that at LBL and used for the same purposes. A T-1 connection to ESnet at Austin should be established for high-performance MFE access.

7.5 BUDGET & PERSONNEL SUMMARY

This section summarizes a possible computing plan for FY89 and FY90 in terms of money and staffing necessary to support the requirements of (1) numeric-intensive calculations for design and simulations, (2) network and telecommunications, and (3) supporting services such as mail, text and word processing, and computer conferencing, in a transition from LBL to the SSC site and still satisfy the requirements for the continued growth of SSC R&D. There are substantial uncertainties in the area of detector simulation that may only be resolved by further study and by decisions regarding the development of the SSC experimental program.

The capacity proposed to be leased at or near the site is to provide an anchor for the startup of computing at the SSC site and is intended as transition capacity through 1990. This will allow a smooth continuation of capability for accelerator-related design and a significant start on capability for detector design that may be readily accessible to a wide community of users. We specifically have chosen the lease option to avoid long-term commitments to "old" architectures and to allow for future policy decisions.

We believe that a review for the next phase of computing is an important issue which should be addressed in early 1990. The real startup capacity for 1990–1994 is suggested in more detail in the report of the Offline Computing Advisory Panel in Report SSC-SR-1023 (June 1986).

Costs for:	FY89 ($ K)	FY90 ($ K)
Accelerator Physics, Systems & Magnets		
Work Stations		
H/W Maintenance	50	200
S/W Maintenance	20	40
Equipment Purchases	445	1,230
Software Purchases	100	200
Purchased Computer Time		
LBL VAX Cluster	750	550*
Cray MFE	—	—
Physics and Detector Simulation		
Purchased Computer Time†	825	825
Acquisition of Hardware for Simulation	100	2000‡
Networking		
Network & Telecommunications		
Leased Lines	104	380
Local Area Network @ SSC Site	40	60
Hardware/Software in Support of Local and Wide-Area Network	160	70
General-Purpose Computing Resources at the SSC Site		
Hardware		
Dec VAX 8820 (lease w/option) (2)	280	550
20 Gb disk (lease w/option)	50	100
Exabyte Tape (purchase)	15	7
Printer (purchase)	15	15
Software (purchase)		
Compilers/Edit/Code MGT	135	60
Math/Stat	16	—
Finite Element	10	30
Maintenance		
Hardware	3	64
Software	16	24
Personnel (FTE's) (Assumes 12-hr operation in FY89 and 24-hr operation in FY90)		
Systems	1	2
Network/Telecommunications	3	5
Operations	2	5
Product Support	2	3

*This assumes six-month usage of the LBL cluster in FY90 for accelerator studies.

†Purchased, for example, from LBL via μVAX farms and the LBL cluster and/or on the Cray X-MP at ANL

‡ This is a very crude estimate and should be refined during the next year.

Note: SUN hardware maintenance is assumed at 10% purchase price, VAX hardware maintenance at 3% purchase price. Software maintenance is assumed at 10% purchase price.

Although it is not the mandate of this Task Force to consider all the long-range issues of SSC-era computing, we feel we need to mention the particularly important matter of software engineering issues as they apply to the huge on- and off-line detector codes which will be required at SSC detectors. Development of such codes in HEP has to date been notoriously undisciplined. A few early computer-aided software engineering (CASE) tools are starting to appear in HEP groups. However, the most basic and standard software project management procedures, like requirements and design documents and reviews, are virtually non-existent. The problem in the HEP environment has many similarities with large software elsewhere. However, it has one essential difference. Because of the many unexpected problems that frequently occur in the experimental physics environment on the one hand, and because of the large scale of detectors and software in HEP on the other, there is a unique combination of extremely large yet rapidly changing software. This is an especially difficult and interesting case for the professional software engineering community and the computer-aided software engineering tool industry. It will be important for the SSC organization to address this area early on.

Appendix A

Task Force Members

Dr. Jeffrey A. Appel
Fermi National Accel. Laboratory

Mr. Marvin Atchley
Lawrence Berkeley Laboratory

Dr. Ruben Carcagno
SSC Central Design Group

Mr. Gregory A. Chartrand
Fermi National Accel. Laboratory

Dr. Thomas Dombeck
U.S. Dept. of Energy/GTN

Dr. Murdock G. D. Gilchriese
SSC Central Design Group

Dr. Donald E. Groom
SSC Central Design Group

Dr. Gail Hanson
Stanford Linear Accelerator
Center

Mr. William E. Johnston
Lawrence Berkeley Laboratory

Dr. Stewart Loken
Lawrence Berkeley Laboratory

Dr. Thomas Nash
Fermi National Accel. Laboratory

Mr. Vern Paxson
SSC Central Design Group

Dr. Steve Peggs
SSC Central Design Group

Mr. Jack O. Pfister
Fermi National Accel. Laboratory

Mr. Lev Pope
SSC Central Design Group

Dr. Lawrence Price
Argonne National Laboratory

Dr. Chris Quigg
SSC Central Design Group

Dr. Chris Saltmarsh
SSC Central Design Group

Dr. Lindsay Schachinger
SSC Central Design Group

Dr. Robert Schermer
SSC Central Design Group

Dr. Tjet Sun
SSC Central Design Group

Dr. John C. Tompkins
SSC Central Design Group

Dr. Edward Wang
Lawrence Berkeley Laboratory

Dr. Saul Youssef
Florida State University

Dr. Jack Zelver
SSC Central Design Group

Appendix B

Appendix B is a partial draft of the report of the working group on Computing for High Energy Physics in the 1990s from the recent Snowmass Summer Study on High Energy Physics in the 1990s.

This report discusses some aspects of SSC computing not covered by the Task Force and does not necessarily represent the conclusions of the Task Force.

COMPUTING FOR HIGH ENERGY PHYSICS IN THE 1990's

S. C. Loken
Lawrence Berkeley Laboratory
1 Cyclotron Rd., Berkeley, CA 94720

Partial Draft

Abstract

The working group on computing for high energy physics in the 1990's has reviewed a wide range of issues related to computing and software development. Many of these are specific to the problems of the SSC. Many, however are more general. Foremost among these is the need to improve the quality of the software used in high energy physics and to improve the productivity of physicists and programmers working in the field.

1 Introduction

In this report we consider the computing problems of high energy physics in the 1990's. It is clear that the computing needs of the field are increasing significantly and will continue to do so. We are building larger and more complex experiments. The detectors consist of more individual elements and the experiments wil collect more events.

The size and composition of experimental collaborations also present new problems. The collaborations are large and include members from many countries. There are more programmers involved in the computing effort and they are distributed through all institutions in the collaboration.

The Superconducting Super Collider (SSC) is just one example of the computing problem for the next decade. The reserch program of the 1990's will present similar challenges in many other areas.

In this report, we discuss the specific problems of the SSC computing including those of detector simulation. We also review the software development issues that will affect computing in the 1990's.

2 Computer Planning for the SSC Era

In the near future, the effort to design detectors for the SSC will increase significantly. This will result in a need to provide computer time and programs for detector simulation.

Once the site is selected and the laboratory is formed we will need to establish a computing department. In this section, we consider these issues.

2.1 Computing for Detector Design

As the deadlines for letters of intent and for proposals approach it will be increasingly important to have sufficient computing capacity to satisfy the needs of collaborations to simulate SSC detectors and to optimize their design. We have made an estimate of the needed capacity based on recent simulation work. In doing this we make the following assumptions:

- Typical run of 20,000 events
- Using fast shower Monte Carlo
- Time/event of 15 minutes (VAX-11/780)
- Total of 5000 VAX CPU hours per run

There is a great deal of effort going into improving the performance of the detector simulation program and the time to simulate a single event is expected to improve. The assumptions above are, however, useful to set the scale of the problem. For the next three fiscal years, we get the following:

- FY1989
  - Ten simulation efforts
  - Each does ten runs
  - Total requirement is 70 VAX-780 equivalents
- FY1990
  - Four collaborations with fifteen people working on simulation
  - Each does some full simulation runs which double time
  - Total requirement is 840 VAX equivalents
- FY1991
  - Four collaborations with thirty people working on simulation
  - Each does many full simulation runs
  - Total requirement is 5000 VAX equivalents

The total computing requirement is large even in the next year. Even with progress in speeding up the simulation or in making the code run efficiently on vector supercomputers, additional facilities must be provided to satisfy the needs of detector designers. We recommend that a center for detector simulation be established as soon as possible to provide the computing needed in FY1989. In addition to providing the needed computing, this center should manage program libraries and databases and save simulated data so that it can be shared by other groups. To be useful to the community, the center must be available on HEPNET and other networks. As the use increases, higher speed networks will also be necessary.[1]

As soon as the SSC laboratory is established, the computing facilities and support group should be moved there. We discuss below the need for an SSC Computing Division which will manage SSC computing facilities and act as a resource to the national high energy physics program.

2.2 Programs for Detector Simulation

There has been considerable progress in developing software for detector simulation. This was the focus of a workshop in August, 1987 [2] and was reviewed at Snowmass by Larry Price. Some of the important areas of progress are the following:

- Agreement to focus development on GEANT3 and ZEBRA
- Plan to make Oak Ridge calorimeter code CALOR available in GEANT3
- Outline of standard interface for event generators
- Plans for utilization of vector and parallel computers
- Standard particle numbering system coordinated with PDG
- Ideas for improving the ease of use of GEANT3

While there has been significant progress, much remains to be done to develop a simulation that is easy to use and which runs fast enough to generate the large numbers of events needed to study subtle backgrounds in SSC detectors. In addition, there needs to be a coordinated effort to compare simulations with test beam data to ensure that the simulations are correct.

3 SSC Computing Division

Our working group has concluded that a strong SSC Computing Division is necessary and that it should be formed as soon as possible, hopefully within the next year or two. Although there is a serious immediate need for this division it is also extremely important that it be formed with the right charter.

This division would have charge of computing, networking, software, and related facilities for SSC users. Because of the size and centrality of the SSC program this division would have a position of leadership for computing activities for the entire U.S. high energy physics community.

It is **essential** that this division be responsive to the needs and desires of the community. This must be reinforced structurally, perhaps by a policy committee of users, both SSC employees and outside users. **This division should lead the community by persuasion, not arbitrary fiat.**

The duties of the division would be:

- Obtain and operate SSC computing facilities
- Provide and maintain SSC and community (HEPNET) networking facilities
- Provide software utilities for the community
- Suggest community standards to aid the development, maintenance, and portability of programs
- Investigate and develop advanced hardware and software facilities for HEP use

We will describe each of these duties in turn.

3.1 Obtain and Operate SSC Computing Facilities

1. **Mainframe.** Although the computing environment is continually changing we expect that a mainframe computer to serve as a center around which the other activities are grouped will still be essential. It will need very large mass storage capabilities, and fast dataways for routing information internally and externally. The division may well decide that a whole set of multi-user computers is needed.

2. **Workstations.** The laboratory will need many workstations of widely varying power. Some of these will be very small computers, others will have signifigant computing power or be very fast graphics engines in their own right. The division will have many choices here. The division will have to decide quite early what level of commonality to establish between the various stations. They could be all members of a single computing family, or they could be entirely disparate, connected only by a data link to the central mainframe(s) or to other workstations. X-windows may be quite advantageous here since it is non-proprietary.

3. **Terminals.** The laboratory will have to provide terminals. It may well be that terminals will be largely replaced by various levels of workstations.

4. **Specialized off-line computing facilities.** At present it appears that many specialized often non-commercial arrays of processors or other facilities will provide a significant fraction of the computing power needed for SSC data analysis. Some of these will be designed and built by the computing division and some will be obtained from outside the division. A present day example of such facilities is the set of very successful Fermilab ACP systems, some of which are now operated as part of the Fermilab computing center.

5. **On-line facilities.** The division will have the very difficult problem of having to draw the line between the responsibilities of the division and that of individual groups. At the outer trigger levels the hardware is probably rather standard, but it gets more and

more specialized for the inner trigger levels. Policy decisions will have to made about the level of laboratory equipment support for these facilities.

6. **Operating systems.** This is another gray area. The level of support that must go into non-standard (or standard) operating systems for on-line and off-line systems will depend on the state of computing at the time.

7. **Programming assistance.** It is clear that assistance must be available for specialized operating systems and specialized programming packages maintained by the division. However, programming assistance can also mean helping students with trivial errors, suggesting programming strategies to groups, or even developing and coding various packages for groups. We will discuss generic packages later, but it is clear that the activities listed here could take a great deal of time, and limits will have to be set.

8. **Evaluation of various commercial computing systems and options.** In the course of making evaluations for its own purchasing the division will have to develop good batch and interactive benchmarks and apply them to many systems. It will have to examine the differing advantages of scalar, vector, and parallel processing. It would be extremely valuable to the community if the division could make available to the high energy physics community results from good benchmarks involving high energy physics type calculations, and could serve as a resource for groups trying to make decisions for their own computing.

Items 2, 4 and 8 will help standardize HEP facilities nationwide, improving computing efficiency, cost, and protability.

3.2 Provide and maintain networking facilities.

1. **Install and maintain networking within the laboratory including to the experiments.** Networks will have to be physically installed throughout the laboratory and switches and gateways maintained for efficient transfer of data between the various sites.

2. **Install and maintain networking for the community.** The SSC is the natural center of the network operations for high energy physics. This may mean nothing more than connecting to NSFNET or it may mean much more depending on the way networking evolves. Effective networking within the lab, between the lab and users, and between users is absolutely crucial for physics during the next decade. Networking outside of the U.S. will be very important for the large multi-national collaborations.

3. **Maintain and develop network software.** Good networking protocols, translation between various lower level protocols, and higher level software (e.g. coloured books or X.400) are all vitally necessary. They may all be available as a single commercial package or they may require considerable effort. The division should strive for uniform procedures for the high energy community. Special problems for users outside the U.S. will have to be considered.

3.3 Provide software utilities for the community.

For almost two decades the U.S. high energy community has depended to a large extent on CERN for development of software utilities. The working group feels it is necessary for the U.S. community to do a greater share of development than in the past. Utilities should be developed as needed. Sometimes this may involve collaboration with industry or with other labs. When feasible, utilities should be purchased from vendors. If utilities are to be purchased the division should attempt to arrange for discounts for the high energy physics community and should take into account the total cost to the community of decisions made by the division. Any general utilities developed or purchased must be available for a range of hardware, not just the machines at the SSC. They should work on the machines used by the bulk of the users. Some examples of possible utility packages follow:

1. Program development facility. This would perform the functions of PATCHY, HISTORIAN, etc. It should include a number of modern facilities developed to assist and organize large tasks being worked on by a number of people.

2. Graphics standard system.

3. Data base management system.

4. On-line software aids. This is a gray area. The old programs such as the Fermilab MULTI package were extremely useful. However, the modern trigger systems involve much more complicated structures. Some of it can be standardized. Much of it is still under development, but the line between standard and specialized utilities will be a constantly moving one as the techniques develop.

5. Off-line workstation and graphics aids. There are a number of utilities developed or being developed at CERN in this area (PAW, ...). This area is in intense development now and packages will be changing quickly.

6. Standard large programs. An example of this is GEANT. Coordination with CERN will be very useful on some of these programs.

7. Program library. CERN has built a large library of standard subroutines. The division should coordinate with CERN on further developments.

8. I/O utilities. I/O utilities are needed to aid the portability of programs and data within and perhaps even between experimental groups when the individual members have diverse facilities.

3.4 Suggest community programming standards

This has a strong educational aspect and also a strong relation to the preceeding section. New standards are needed in a number of areas as indicated below. These standards should be developed in conjunction with the community. Seminars, working groups, colloquia, schools, meetings, etc. are needed to define community desires and to obtain community agreement with and adherence to standards suggested. Some of the areas needing standards are:

1. Portability standards. These are standards of programming styles, of I/O, etc. which will make it easier to port programs from one hardware system to another.

2. Standards for designing new programs. Structured analysis, structured design (SASD) is one example of such standards. An educational effort here could result in programs which are better documented and easier to maintain then is the case for most of the present generation of high energy physics programs.

3. Programming styles, languages. It may well be that it is time to consider the possibility of doing a part of particle physics programming in languages other than FORTRAN. Such a major shift in style would require a considerable effort and investment by the community. The SSC division would be an appropriate place to have the community consider the problem and to lead an educational effort after a decision is reached.

4. Graphics. Graphics standards are in flux, but graphics is an increasingly large area of our programming.

5. Networking. It is clearly in the community interest to standardize networking protocols, gateways, higher level protocols, etc. as much as possible. Lines of the highest speed obtainable ($\geq$ T1) should be provided for the communication needs of the community.

6. Operating systems. This is at present a gray area. Is it possible to arrange to interface all systems with UNIX?

3.5 Investigate and develop advanced HEP facilities

The division will have to examine possible new hardware and software facilities. This may involve collaboration with industry, with other laboratories, or with Universities. Some examples of possible projects are given below. Even more than with the other subsections this is an area of rapid development and the items listed are only a few of the many projects that will undoubtedly develop.

1. Large processing arrays for on-line, off-line use. The Fermilab ACP is an example of a successful present day project in this area.

2. Specialized processors for on-line and off-line work. For trigger processors neural network or systolic processors or other special purpose processors might prove extremely useful.

3. New programming strategies. The computing community has developed a number of new programming strategies such as object oriented programming which may prove to be useful for some high energy physics programming. It would be very worthwhile to explore their usefulness.

4. File servers, mass storage devices. This field is in great flux. The division should examine the field, attempt to select some appropriate technologies for high energy physics use and should attempt to set standards for these devices. (For example the VHS based tape cartridge systems do not have inter-vendor standards for how the data is stored. Portability requires some decisions on standardization.)

4 Software

It is traditional that software problems receive less attention than hardware, especially in the early phases of an experiment. This has been a serious problem for the current generation of high energy physics experiments and the problem will be even worse in the future.

Each of the current generation of colliding beam experiments is developing a software base of $3 - 5 \times 10^5$ lines of code. This development will take an effort of 300–500 man-years. The developers are typically not professional programmers and they have little, if any, training in computer science. They are not located in one place and, for the most part, do not devote full time to their software responsibilities. When a physicist who is working on the design of a hardware component needs additional help, he often turns to and engineer or a technician. It is extremely rare that physicists will enlist professional help during any phase of the software development effort.

This situation is now changing. Recognizing the effort that software development will require, many groups are beginning to use software development methods that have been standard in the computing industry for many years. Many are using specialized systems to maintain data structures within FORTRAN programs and to use database management tools to maintain constants. While it is still too early to assess the success of these efforts, the indications are that the software will be better and that it will be easier to maintain.

4.1 Traditional Development Methodology

Up to now, most experiments have followed a "bottom-up" strategy for software development; the software is put together from components which were written for isolated studies during the design phase of the experiment. The result is a system with little design documentation. There is agreement on a general outline but individuals work on components with little intercommunication. Little attention is paid to integration of modules until after they are complete.

4.2 The Software Life Cycle

A modern approach to the problem of software development for experiments begins by recognizing that the software project will have a long and complex life and that each phase must be planned in detail. In fact, the approach is very similar to the engineering of any hardware component of an experiment. The phases of the software life cycle are the following:

- Requirements analysis

- Define what the software is to do without consideration of how it is to do it or what the hardware configuration will be.
- Build a logical model of the software system.

- Design
  - Define details of the software.
  - Create a physical model which describes the implementation.
  - Subdivide the problem into processing units, then into tasks, and finally into modules within tasks.
- Coding
  - Write code, test and document modules.
- System Integration
  - Assemble modules and test.
- Acceptance Test
  - Demonstrate system performance
- Operation and Maintenance
  - Install in sites.
  - Update as necessary.

The formal structure emphasizes the early phases of the development cycle. Many studies have showed that the cost to fix an error increases significantly in later phases of the project (see Figure 1). In the early phases, the group working on the project is smaller. The effort to create a new model is much less than the effort to recode, compile and retest a module.

It is important to recognise that there are real costs associated with software errors. These are not measured by the same criteria as in the commercial world. The cost, however, can be measured in terms of lost beam-time, the expense of rerunning data tapes, missed discoveries, or incorrect results. It is very much worth the effort to get it right.

Figure 1: The relative cost to fix an error or to make a change in software as a function of the time in the project.

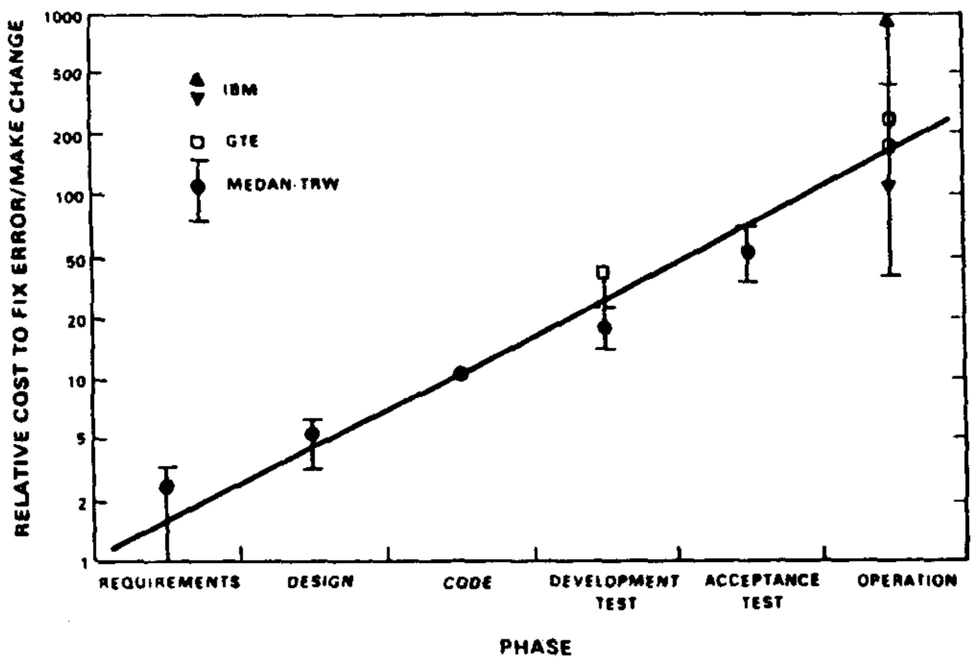

4.3 Structured Analysis/Structured Design

One of the formal methodologies for the first two phases of the project is Structured Analysis/Structured Design or SA/SD [8]. This methodology has been adopted by the ALEPH collaboration at CERN [9] and the D0 experiment at Fermilab[10]. While it is too early to determine whether the software for these experiments will be significantly better than that developed using traditional methods, both groups agree that they are developing programs that are different from what they would have done without SA/SD.

In the analysis phase, SA/SD uses three graphical tools to model the software system. These tools describe the software in much the same way that engineering drawings model a hardware object. A Data Flow Diagram models the flow and transformation of data in the system (see Figure 2). A State Transition Diagram (Figure 3) describes the time dependence and is especially useful for online programs or control systems. The Entity Relationship Diagram (Figure 4) defines the data elements of the system and the relationships between them. Together, these three tools describe the character of what is to be built. The graphical tools are augmented with textual specifications such as the Data Dictionary which specifies the characteristics of the stored data and Minispecs which describe the transformations in the diagrams.

The diagrams developed in the analysis phase are useful for describing the software system to people outside the development team. They should be the subject of a formal

Figure 2: An example of a Data Flow Diagram taken from the ALEPH software [9].

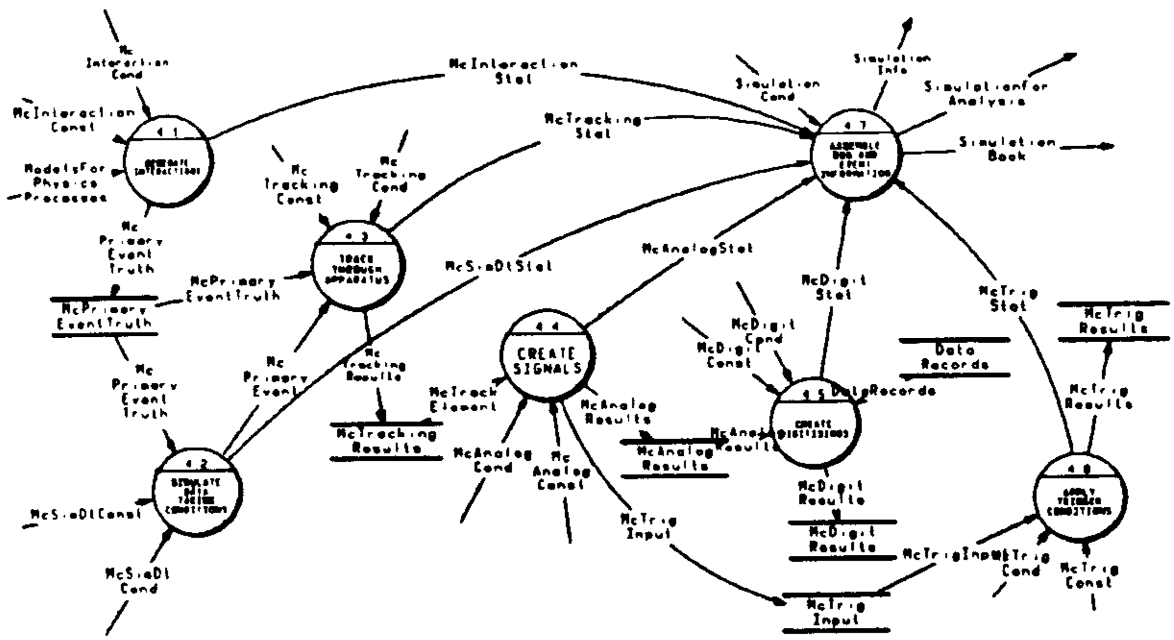

review, often called a Walkthrough. An error found by a review at this stage is usually fixed much more easily than later in the development cycle. There may be many ways to describe a system with the graphical tools. To choose the best model, or to refine a model, there are a number of criteria that may be applied. The most important of these is correctness. In addition, it is useful to simplify the interconnections in the system. The result of the analysis phase is a logical model of the system.

In the design phase, the constraints on the software are added. Pieces of the logical model will be allocated to various computers. The details of the human interaction will be incorporated. The elements of the logical model are assigned to modules and a physical model for the software is developed. The primary tool for the design phase is the Structure Chart. This is again a graphical tool and models the hierarchy, the partitioning and the interfaces of modules within a single program (see Figure 5). The Structure Chart should also be reviewed in a Walkthrough. There are also techniques to evaluate and refine this physical model of the software.

4.4 Computer-Aided Software Engineering

The commercial interest in software development methodologies has spurred the development of new tools to support what is now called Computer-Aided Software Engineering (CASE). Many of these products support SA/SD. The tools include programs to draw the diagrams of SA/SD. There are routines to check the consistency of various diagrams and

Figure 3: A State Transition Diagram taken from the D0 calibration software [10].

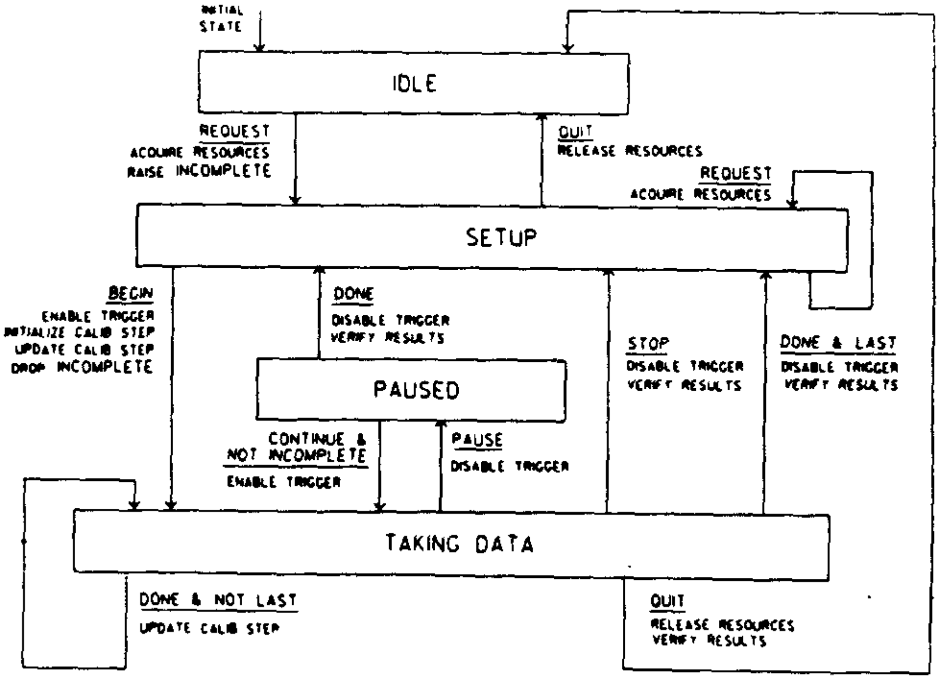

Figure 4: An Entity Relationship Diagram taken from the ADAMO data management system for ALEPH [9].

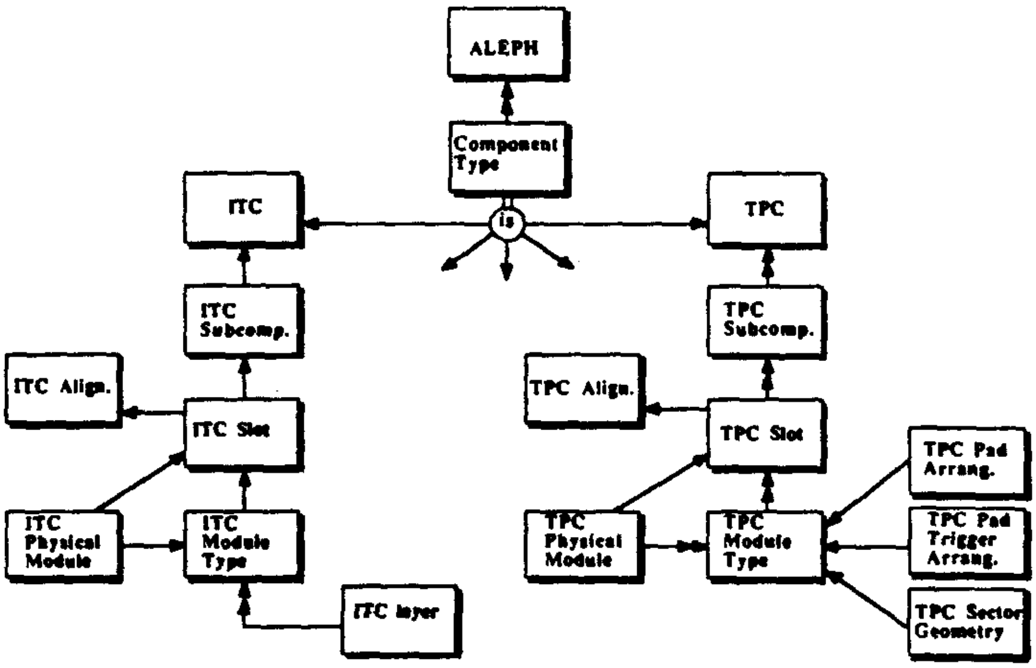

Figure 5: An example of a structure chart taken from ALEPH software [9].

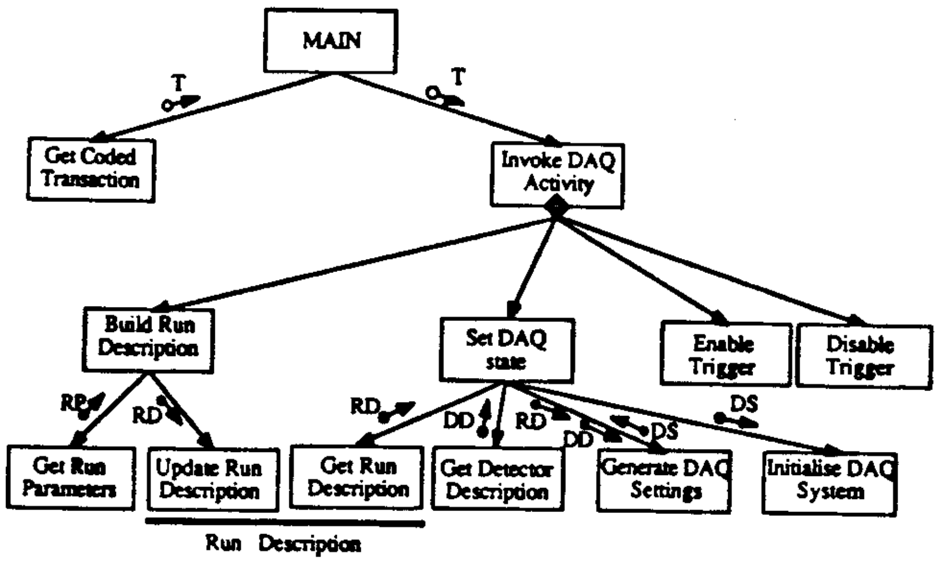

to ensure that there is a complete specification of each data element or transformation. The tools store the diagrams and specifications in databases which are available over a network from remote computers.

These commercial have improved substantially in the last year. The two experiments which have committed to SA/SD were not able to find products that would satisfy their requirements even two years ago and have developed some tools themselves. D0 has recently undertaken an evaluation of available products and will try to put existing documentation into a new system. ALEPH has developed a set of tools which use the MacIntosh to manipulate Entity Relationship Diagrams in their ALEPH DAta MOdel (ADAMO)[11] . The ADAMO package goes beyond many commercial products and provides FORTRAN tools to manipulate data within the analysis code.

4.5 Code Management

During the development cycle, the software library must be carefully managed. Stable versions must be distributed to all developers so that new software can be tested with the rest of the system. There must be a full record of changes and there must be set of rules to ensure that new versions in the library are fully tested and documented. In fact, the problems of management begin earlier in the project. The same considerations apply to the requirements documents and the design documents.

The development of software for an experiment presents some challenges which are often not found in a commercial organization. The software continues to evolve over the history of the experiment as the apparatus is upgraded, our knowledge of the detector improves, or our physics interests change. There are many people involved in the project and they are at many sites. There are often many different computers.

The problem of code management is a difficult one and there is no simple solution. The PATCHY system developed at CERN can run on many computer systems but is not adapted to interactive access and is not widely used in the United States. Its replacement at CERN is a commercial product, HISTORIAN, but it is expensive for groups outside CERN and has not been widely adopted, even by LEP experiments. Many groups in the United States have decided to manage their software only on the VAX computer systems. The VAX products, CMS (Code Management System) and MMS (Module Management System) provide most of the functionality needed for experiments and support library management over a network with DECnet. There is a significant problem, however, for groups that have other computer systems at their home institutions.

5 Software Testing

No one developing software would dispute the need to test new software. The problem comes as the program is changed. Each change is tested to ensure that it does what it is expected to do. Seldom is the program subjected to all the tests that were run at earlier stages to ensure that old problems have not been reintroduced or that the change did not have some unexpected impact. To do this requires maintaining a series of test scripts that are run on each new version before it is distributed for general use. A set of standard output files must be maintained with the scripts so that a new file can be verified. Such a product is available as part of the VAX software system, DEC Test Manager (DTM) but again it works only for programs which run on the VAX computers.

Another aspect of testing that seldom gets the attention it deserves is the need to follow all possible paths through the program. A single program may have many thousands of possible paths depending, for example, on the nature of a complicated event. Often, these paths are not tried until a real event provides the set of parameters. If the program does not handle it properly, it is necessary to rerun the tape and chase the problem. It is more efficient to minimize the number of possible paths by the use of structured programming techniques, and then design the testing to ensure that all paths have been tried.

Various studies in the computer industry have indicated that up to half of the effort in a large project may be spent in testing programs. This statistic by itself should may it a high priority to improve the efficiency of software testing. The book by Myers [12] provides useful guidance for all aspects of software testing.

References

[1] *High Energy Physics Computer Networking: Report of the HEPNET Review Committee.* DOE/ER-0372, Department of Energy, 1988.

[2] L.E. Price, editor *Proceedings of the Workshop on Detector Simulation for the SSC*, ANL-HEP-CP-88-51

[3] K. Miura. Vectorization of EGS4 on vector processor system. In P. Kunz and T. Schalk, editors, *Computing in High Energy Physics: Proceedings of the International Conference on Computing in High Energy Physics, Asilomar, 2-6 February 1987*, page 127, 1987.

[4] D. Levinthal. Experimental HEP supercomputing at FSU. In P. Kunz and T. Schalk, editors, *Computing in High Energy Physics: Proceedings of the International Conference on Computing in High Energy Physics, Asilomar, 2-6 February 1987*, page 137, 1987.

[5] I. Glendinning and A. Hey. Transputer arrays as FORTRAN farms for particle physics. In P. Kunz and T. Schalk, editors, *Computing in High Energy Physics: Proceedings of the International Conference on Computing in High Energy Physics, Asilomar, 2-6 February 1987*, page 367, 1987.

[6] *Computing for Particle Physics: Report of the HEPAP Subpanel on Computer Needs for the Next Decade.* DOE/ER-0234, Department of Energy, 1985.

[7] R. Bock et al. PAW-towards a physics analysis workstation. In P. Kunz and T. Schalk, editors, *Computing in High Energy Physics: Proceedings of the International Conference on Computing in High Energy Physics, Asilomar, 2-6 February 1987*, page 181, 1987.

[8] S.J. Mellor and P.T. Ward. *Structured Development for Real-Time Systems.* Yourdon Press, New York, 1986.

[9] G. Kellner. Develpoment of software for ALEPH using structured techniques. In P. Kunz and T. Schalk, editors, *Computing in High Energy Physics: Proceedings of the International Conference on Computing in High Energy Physics, Asilomar, 2-6 February 1987*, page 229, 1987.

[10] J.J. Linnemann et al. The use of SA/SD methods in D0 software development. In P. Kunz and T. Schalk, editors, *Computing in High Energy Physics: Proceedings of the International Conference on Computing in High Energy Physics, Asilomar, 2-6 February 1987*, page 245, 1987.

[11] Z. Qian et al. Use of the ADAMO data management system within ALEPH. In P. Kunz and T. Schalk, editors, *Computing in High Energy Physics: Proceedings of the International Conference on Computing in High Energy Physics, Asilomar, 2-6 February 1987*, page 283, 1987.

[12] Glenford J. Myers. *The Art of Software Testing.* Wiley-Interscience, New York, 1979.