The submitted manuscript has been authored by a contractor of the U.S. Government under contract No. W-31-109-ENG-38. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

COMPUTING NEEDS OF THE SUPERCONDUCTING SUPER COLLIDER

R. Diebold
Argonne National Laboratory, Argonne, IL 60439

Summary

Following a brief description of the SSC, the computing needs are discussed for both the accelerator design and the experimentation. The computing power required is considerably beyond that being used at present facilities, and parallel processing is expected to play an important role in supplying these needs.

Brief Description of the SSC

The Superconducting Super Collider (SSC) is a proposed high intensity colliding beam machine using two 20-TeV proton beams. It was first seriously discussed as a U.S. machine at a Snowmass workshop held in the summer of 1982. A year later it was identified by a HEPAP subpanel as the highest priority construction project for the American high energy physics (HEP) program.¹ Depending on the magnetic field, the diameter of the machine will be in the neighborhood of 20-30 miles, too big to fit on any of the existing HEP laboratory sites. The project is now in an R&D phase aimed at finding the most economic and reliable design, with concentration on the engineering and prototyping of the superconducting magnets required to guide the beams on their circular paths. Assuming construction approval for FY-88, the machine could be doing physics in 1994. The location of the SSC is completely open, and it is expected that the DOE will ask for site proposals this coming year.

The energy range of the SSC will be well above that presently available, and there is a plethora of theoretical predictions for what might be discovered:

heavy quarks, leptons, gauge bosons;
Higgs/technicolor particles;
supersymmetric particles;
substructure in quarks and leptons.

In addition, cosmic ray experiments have found events which suggest new types of interactions at very high energies. The physics potential of the SSC was explored in detail at the 1984 Snowmass workshop and was briefly reviewed at this conference by Winstein.⁴ Perhaps the most important reason for building the SSC is that we simply do not know what forces and interactions be unleashed at these energies.

The cross sections (or probabilities) for producing massive particles fall as $1/M^2$ and to achieve the highest rates, it is generally assumed that two beams of protons will be collided with one another. Since the two beams have the same charge, the magnetic field must be up for one beam and down for the other, and two rings of magnetic guide field must be provided. The beams travel through the magnets in high-vacuum tubes (about 10^{-9} Torr). The protons are bent through 360° by very strong dipole magnets, a total of 420 Tesla-km being needed for each beam at 20 TeV. The beams are focussed by quadrupole magnets located roughly every 100 meters. An rf system is used both to accelerate the beam (from the 1-TeV injection energy) and to replace the energy which is continually lost to synchrotron radiation. A typical design for the SSC has 10^4 bunches spaced 10 meters apart around a 100 km circumference ring.³ Each of these bunches has roughly 10^{10} protons and is

needle-like in shape with a transverse rms dimension of approximately 0.1 mm and a longitudinal length of perhaps ± 70 mm.

Colliding beams are the latest in a long series of techniques which have been used over the years to increase the energy available for experimentation. Previously, in the so-called fixed target mode, most of the beam's energy was needed just to conserve forward momentum. The useful or center-of-mass energy in this mode scales as

$$E_{cm} = \sqrt{E_{beam} M_{target}}$$

Even with the new superconducting Fermilab Tevatron accelerator, the fixed-target center-of-mass energy is only 44 GeV for a 1000-GeV incident beam.

With colliding beams, all of the beam energy is available to make new particles or cause new types of interactions. At 20 TeV per beam, giving 40 TeV in the center-of-mass system, the equivalent fixed target accelerator would need 10^{18} eV, well beyond most cosmic ray studies. Even with strong 8-Tesla superconducting magnets, such a machine would have a radius of about 400,000 km, the distance from the earth to the moon. Obviously such a machine is not practical, and these energies can only be reached with colliding beams.

Over the years, the accelerator beam energy available for experiments has increased by roughly a factor of two every two years.⁵ Thanks to colliding beams the equivalent accelerator beam energy is still increasing exponentially. It is interesting to note, as pointed out in this session by Pernbach,⁶ that the performance of computers has been increasing at about the same rate with a new generation appearing typically every five years with roughly a factor of five increase in performance.

While I won't predict the ultimate limit on computing speed, it's clear that this exponential growth of accelerator energy cannot continue indefinitely; in the year 3000 the energy would be the equivalent of the entire mass of the universe - back to the big bang!

The prototype for the SSC came into operation about a year and a half ago. This is the Tevatron at Fermilab; it roughly doubles the energy available at that laboratory over the previous 400-500 GeV. Many lessons have been learned in the construction and operation of this superconducting accelerator, lessons which can now be applied to the SSC. Valuable experience with colliding beams has also been gained at CERN on both the ISR and SPS. The present Fermilab complex of accelerators also serves as a model for the SSC injection system. Since a superconducting accelerator can only operate over a factor of roughly 20, SSC injection must occur at about 1 TeV, the energy of the Fermilab Tevatron.

An intensive design effort was mounted this past year under the leadership of M. Tigner (Cornell University).³ This was the Reference Designs Study (RDS), which took place from January to May 1984. It was basically a technical and financial feasibility study for a facility with 20 TeV on 20 TeV pp colliding beams, at a luminosity of $10^{33} \text{ cm}^{-2} \text{ sec}^{-1}$. Three separate magnet designs were examined:

Invited talk at the IEEE 1984 Nuclear Science Symposium, Orlando, FL
(October 31 - November 2, 1984).

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

PS

PARTIONS OF THIS REPORT ARE UNLIMITED.
 IT HAS BEEN REPRODUCED FROM THE BEST
 AVAILABLE COPY TO PERMIT THE BROADEST
 POSSIBLE AVAILABILITY.

B	Circumference	Cost
6.5 Tesla	90 km	\$2.7B
5	114	3.0
3	164	2.7

Computing Needs for SSC Accelerator Design

The high cost of the SSC ring precludes excess conservatism in setting its specifications, yet the machine simply must work. This means that the design must be examined very closely, and this involves a great deal of computational effort, more so than for any machine in the past. In describing this computational effort below, I have drawn heavily on the reports of the Accelerator Computing Group at the recent Snowmass Workshop and of the recent HEPAP Subpanel on Theoretical Computing.⁸ For a machine the size of the SSC, the magnetic aperture and field quality are very expensive, yet are crucial to the successful operation of the machine. Over the next two years a large effort will be made to set the specifications and design of these magnets. This will be followed by an ongoing study to evaluate the fabricated components and to better understand the performance of the machine.

The central issue of these studies is beam stability. In an ideal linear world, there would be no problem. The beam would simply oscillate about its central value, repeating itself endlessly. However, there are many sources of nonlinearity that can limit the regions of stability for the beam. Some of these nonlinearities are introduced on purpose, for example, sextupoles to correct for chromaticity (variation of tune with beam momentum). Other nonlinearities come from magnet construction errors, both random and systematic. Some of the limiting effects are beam-induced: scattering of particles within a single bunch, electric and magnetic fields set up in nearby materials by the beam, and beam-beam interactions at the collision points and elsewhere.

To understand the stability limits, one must calculate various electromagnetic fields, both external fields produced by magnets, rf cavities, and so on, as well as the beam-induced fields: wall currents, wake fields, etc. One then examines particle motion in these fields, both single particles controlled by the external fields, and collective motion of the beam, due to both external and beam-induced fields.

The calculation of external fields is an old classical problem. Usually there is, in principle, a great deal of symmetry, but in practice both errors and end effects must be taken into account. While most calculations have been two-dimensional, in some cases one must have three-dimensional calculations and these are very compute intensive.

In the case of beam-induced fields, the beam acts as a source with boundary conditions given by the vacuum chamber wall with all its geometric and material properties. Although most of the 100-km circumference will be a simple cylindrical beam tube and thus relatively easy to calculate, important effects must be understood from quite irregularly shaped objects such as bellows, vacuum pump-out ports, rf cavities, and special vacuum-chamber geometry at the experimental areas.

In the past, accelerator physicists have relied on expansions of magnetic multipoles or frequency-domain eigenfunctions. While such expansions are well

suited to a physical (intuitive) understanding of the phenomena, direct numerical solutions appear to be more powerful.

Accelerator physicists and engineers have a large repertoire of programs to calculate electromagnetic fields: POISSON, SUPERFISH, BCI, TBCI, TOSCA, URNEL and so on. As an example, one might use a grid 1 mm on a side for an object 50 mm across by 250 mm long.⁸ Just storing the field values for this grid at seven words per grid point would take approximately 4M words. Calculations of the impedance and wake fields, such as are needed for stability calculations, on such a grid takes typically two hours of CRAY 1S CPU time. For the SSC, two hundred such calculations might be needed over the course of a year, i.e., several hundred hours of Class VI CPU time.

Calculations of stability and particle motion fall into several classes. First one needs to calculate the motion of a single particle (weak beam). This is clearly fundamental to the basic design if the machine is to work at all. The lattice is designed with fitting programs to tune the quadrupoles, sextupoles, and so on. Then, various multipoles are added to further correct the chromaticity, and estimates are made of the strengths of correction magnets needed to overcome random errors in the principle accelerator magnets.

Collective phenomena must also be examined for the intense SSC beams needed for high interaction rates. The evolution of the beam is given by a solution to the Vlasov equation. Linearized approximations are used to give first estimates of important effects. Eventually, however, one must use Monte Carlo methods, tracking test particles through the nonlinear fields over many revolutions in order to see the combined action of various effects, including synchrotron oscillations. As shown by Table 1, it is this tracking, including beam-beam effects, which is the most compute intensive of the accelerator calculations. Many thousands of hours of CRAY 1S CPU are needed for these calculations.

Tracking is so important and time-consuming that it deserves a more detailed discussion. Many programs exist for this problem: PATRICIA, RACETRACK, TEVLAT, MARYLIE, DIMAT, and so on. For a particular machine configuration an "experiment" might consist of ten particles spread in an initial transverse space space, for each of three momenta and each of three synchrotron amplitudes. If these particles are followed for a total of 100 synchrotron oscillations, such an experiment will consist of 3×10^6 particle revolutions. For the Reference Designs Study it was found that using PATRICIA (no vectorization), 90 ms of CRAY 1S time was required per revolution, or a total of 75 hours for a typical experiment. Many of these experiments will be needed. There are three magnet

Table 1. Computing Needs of SSC Accelerator Design in CRAY 1S CPU hours (Ref. 8).

	Basic Estimate	Assuming Vectorization
Lattice	500	500
Tracking	3,750/15,000	550/2,200
Beam-Beam (weak-strong)	4,000/5,000	600/700
Beam-Beam (strong-strong)	5,000	720
EM Fields	400	60
Instabilities	100	15
Miscellaneous	200	100
Total	13,950/26,200	2,545/4,295

designs, and one might examine ten lattices per magnet design (different phase advances per cell and different IR insertions). For each, a minimum of four samples of the error distributions of the magnetic fields should be used. Thus one might have a total of 9,000 hours on the CRAY 1S. The memory requirements for such tracking are nontrivial if one includes, say, 20 multipoles for each of 10^4 magnets.

Once one understands the effects due to magnet errors alone, the combination of beam-beam and magnet effects must be examined. Each of these calculations could take 10^7 particle revolutions or 250 CRAY 1S CPU hours.⁸ Such calculations might be done for two lattice configurations, each with two or three different interaction region insertions; again taking four random magnet samples each, a total of approximately 5,000 hours is needed to study weak-strong stability.

Strong-strong effects are very time-consuming to study, especially if all aspects are taken into account, including the abort gap and long-range interactions between bunches. Here one must involve a substantial number of (super) particles, and 5,000 CPU hours may be required if vectorization is not done properly. If several bunches per beam must be taken into account, the time could increase accordingly. If computing time were not limited, one would want to do many of these strong-strong calculations in order to understand the effects of different lattice and insertion designs. The reason for fewer "experiments" as the calculations get longer is not due to their being unimportant, but rather to the lack of computing power.

The tracking of particles is straightforward to do in parallel. One can simply place one particle in each processor and have many processors working in parallel. A potential problem which must be considered is that of round-off errors. For 10^7 revolutions, each revolution consisting of perhaps 10^4 transfers through as many magnets, it is easy to see that a small amount of round-off error per calculation could eventually lead to nonsensical results.

The conclusions of the HEPAP subpanel on theoretical computing were that there is an immediate need for a Class IV machine for lattice calculations, database manipulations, and code development for accelerator design. There is also a need for a Class VI machine, in particular to do detailed tracking studies.⁸ On the longer term, I believe that additional computing power could play a very cost-effective role in rapidly understanding the SSC once it is operating, so that we can quickly exploit the full potential of the machine. It is clear that we will need a great deal of particle tracking and that a concerted effort should be undertaken to understand the best way to supply adequate and reliable tracking within the limits of computer technology.

SSC Detector - Physics Computing Needs

The experimental physicist has a long list of computational tasks:

- test beam analysis, calibration data bases,
- data acquisition, online filtering,
- algorithm development for filtering, etc.,
- Monte Carlo events for a variety of processes,
- complete detector simulation,
- pattern recognition/reconstruction algorithms,
- production analysis - filtering, reconstruction,
- event scanning and interactive analysis,
- physics analysis.

Of these tasks, perhaps the most computer-intensive are the Monte Carlo simulations of physics and detector, the production analysis, and the physics analysis including event scanning. As a rule of thumb, the production analysis often takes about half of the total computing power used by an experiment.

The Collider Detector at Fermilab, CDF,⁹ is a good example of a colliding-beam detector. This detector is presently under construction and will have its first run in 1986. Such detectors are quite massive and yet have precision components with wire positions known to within a tenth of a millimeter and calorimeter calibrations good to one percent. This precision necessarily depends on considerable computing to determine and apply corrections.

CDF serves in many ways as a prototype for an SSC detector, although it will run at an energy 20 times below that of the SSC and at a luminosity or event rate down by roughly 10^3 . A typical CDF event is expected to have 100-150 Kbytes of information.¹⁰ This already assumes some preprocessing and suppression of null channels. Approximately 70% of the raw data is expected to come from the tracking devices: drift chambers, vertex TPC and silicon strips. An event rate of 1-3 Hz will fill a tape every 10 to 20 minutes, each tape having between 1000-1500 events. One might thus expect about 10^6 tapes to be generated per calendar year.

The online trigger level for CDF will eventually involve substantial computing power to make the final selection of events written on tape. Present estimates¹⁰ are that the equivalent of 20 VAX-11/780's will be needed online.

Offline there will be several passes through the CDF data. Pass 1 will involve the final calorimeter analysis and clustering of the towers into jets. In addition, event vertices will be found with the vertex TPC. Perhaps 30 seconds per event of VAX-11/780 CPU will be needed for this pass, which will keep 10-15% of the events. Pass 2 will involve a full track reconstruction and take 100-150 seconds per event. All events will be kept at this stage. Pass 3 will correlate data from all detector subsystems and sort the events onto several data summary tapes (DST's). Twenty to sixty seconds per event is anticipated at this level. Altogether it is expected that 25 VAX-11/780 equivalents will be needed for this offline data reduction and reconstruction. Assuming an equal amount of computer power is needed for DST analysis, Monte Carlo studies, and program development, an equivalent of 30 VAX's is required offline.¹⁰

The most cost-effective way of providing this is expected to be a mix of mainframe computers, VAX's, and the modular multi-microprocessor system being developed by the Fermilab Advanced Computer Project (ACP).^{11,12} In many ways, the ACP serves as a prototype for computing at the SSC, both for event reconstruction and for some of the accelerator calculations.

As events pass down the CDF reconstruction chain, it is expected that those events which are retained will grow in length, with new calculated quantities being added to the raw data. This may well lead to a logistics bottleneck and simply be a luxury that can not be long afforded. The conservative experimentalist may have to change his mode of operation and get along without the raw data, instead relying on the output of the reconstruction programs, including results calculated online. This means that such programs will have to be thoroughly debugged, much as the hardware is presently debugged, and finally trusted to give the correct answer. Early and

ruthless compaction of data will become a necessity if high statistics are to be achieved. There will have to be a balance, however, between the systematic and statistical requirements for the study of a particular piece of physics.

An idea of the computer needs for colliding beam experiments can also be gained by looking at the UA-1 computer usage at the CERN computer center.¹⁰ This usage has increased rapidly since their first run, and by the middle of 1983 (their Z run), they were using the equivalent of two IBM 168's or upwards of 30% of the entire CERN computer center. This was far from adequate, however, in that the analysis has gone more slowly than desirable and only a small fraction of the events have been reconstructed.

Estimates are given in Table 2 for the computing requirements of the LEP experiments which will come into operation in 1989. The record length per event, number of tapes per year, and computing power expected are all similar to that for CDF.

For the estimates of the computing needs of SSC detectors, I have drawn heavily on the work of the Snowmass group headed by A. Firestone.¹³ For convenience, I reproduce here their estimates of the relative performance of several computers:

Computer	Relative Compute Power
CDC 7600	10
CDF Cyber 175	7
IBM 3081K (both processors)	19
IBM 370/168	3
DEC VAX-11/780	1

A CRAY 1S is about 100 in these units.

At a luminosity of 10^{33} cm⁻² sec⁻¹ and an inelastic cross section of 100 mb, there will be 10^8 inelastic interactions per second at the SSC. For a detector covering ± 5 units of rapidity, a typical interaction might give 70 charged tracks, each having 100 drift chamber hits, and a total of 110 particles in the calorimeters, each particle perhaps depositing significant energy in three towers, each tower being read out at five depths.^{14,15} Adding in some miscellaneous information, one then has a total of about 10^4 elements of information per interaction. Each element of information will have minimally three words (12 bytes), giving the address (which wire or tower), time information, and pulse height information. For a 10 m (33 nsec) spacing between the bunches, this luminosity will give on average 3.3 inelastic interactions per crossing. The events of interest will probably have more particles than average and thus more information to be recorded than for a typical event. A crossing of interest may then

Table 2. Offline Computing Estimates for LEP Experiments (Ref. 13).

	ALEPH	DELPHI	OPAL	L3
Record length:				
Hadronic event: (Kbytes)	100	200	240	100
Fraction from central chamber:	80%	30-40%	90%	-
6250 BPI tapes/dav:	100	50-100	25-30	35
Tapes/year:	10^4	10^4	10^4	2500
IBM 168 CPU time/hadronic event: (sec.)	25	20	20	30
Total CPU time/year including reconstruction analysis, MC, etc.): (Khr.)	65	32	80	100
IBM 168 equivalents required:	12	12	10-12	15

have 500 to 1000 Kbytes of information. Since one 6250 bpi tape can contain about 150 Mbytes, this means that between 150-300 events can be recorded per tape. At 1 Hz, this will then fill a tape every few minutes, and one can expect several times 10^4 tapes per year from each large detector. Note that this rate of 1 Hz is chosen not for physics reasons, but is due to logistics and calculational limitations.

As for present-day colliding beam experiments, there will be several online trigger levels sifting through the data to reduce the event rate from 10^8 to a few hertz.¹⁵ Level 1 will use calorimetry and fast electronics to look for large transverse energy and/or imbalance (missing) E_T . This level might reduce the rate from 10^8 to 10^3 Hz. The next several levels will use specialized hardware to count clusters of E_T in the calorimeter (jets) and to determine the characteristics (θ , ϕ , width, etc.) of each cluster. At these levels one will look for electrons and work on the muon information using microprocessors. Calorimeter microprocessors will peak up the resolutions, suppress zeroes, and generally correct and compact the data. These trigger levels will reduce the event rate from 10^8 down to about 10^2 Hz. The final selection level will be made using parallel processors. The exact number and strategy for these processors is not yet well determined. There might be 50 of them, each working on a single event to apply complex acceptance criteria involving all detector subsystems (including tracking), or it might be fewer computers, one each for the vertex, tracking, electromagnetic calorimetry, hadron calorimetry, and muons. These processors will then identify the most interesting of the events and sort them into various classes.

This will result in several data streams being recorded onto different tapes. A possible set of data streams might be:

1. "Express" - all information on highly promising events (e.g., huge E_T) at ≤ 0.1 Hz.
2. Interesting events with all information at ~ 0.5 Hz; these would be of direct physics interest as well as providing systematic studies for the third data stream.
3. Interesting events with minimal information (50 Kbytes) at 5 Hz; for example, only retain calorimeter towers with $E_T > 3$ GeV plus those track segments at large radius that point near the vertex (i.e., high momentum tracks from the interaction of interest).
4. Minimum-bias events at ~ 0.2 Hz, for "ins" physics, normalizations, systematics, etc.

To estimate the calculational time required to reconstruct a typical SSC event, the computing group at Snowmass-84 extrapolated from existing experiments, normalizing to IBM 3081K CPU time:¹³

TPC	20 sec/SSC event
MARK II	44
UA-1	50
Fermilab Fixed Target	60
ISR/SFM	64

The average extrapolated time is 50 seconds. Note that this is several times the estimates for LEP events because of the higher particle multiplicities per event and the overlap of several events.

An effective trigger rate of 1 Hz then requires 50 IBM 3081K equivalents for event reconstruction.

While the SSC won't be producing events every second of the day, additional computing time will be needed for systematic runs and Monte Carlo calculations, as well as the physics analysis; because of this, it was felt that 50 IBM 3081K equivalents will indeed be required for each of the large 4 π detectors at the SSC.

Note that unlike many present experiments, the SSC multilevel trigger will already have rejected nonsense events. The remaining events will require substantial computing time before deciding whether they are truly of interest and worth pursuing.

At the present price of about \$5M each, fifty 3081K's would cost approximately \$250M per 4 π experiment. Clearly it is of interest to find a more cost-effective solution and the Snowmass computing group considered three options:¹³

1. Micro-VAX farm of 1000 micro-VAX's at \$20K each for a total of \$20M. In addition, there would be costs for controls, peripherals, and a mainframe computer for physics analysis and other tasks.
2. Emulator farm of 350 3081E's at \$30K each for a total of \$10M. Again one must add host computers, controls, communications, and so on.
3. Supercomputer mainframes. The Snowmass group assumed that five years from now one might achieve the equivalent compute power of a 3081K for a factor of 8 reduction in price, or a total cost of \$35M.

To this cost must be added the online and data acquisition computers and microprocessors necessary to collect, correct, compact and record the data. Also, the collaborating institutions will need to upgrade their local computers with workstations, graphics terminals, networking and so on. The total cost per large detector is then expected to be in the neighborhood of \$50M for computers.¹³ This is roughly 20% of the cost estimated for such a detector (about \$250M).¹⁴

Conclusions

Computational limitations are expected to be a major factor restricting the accelerator design and physics output of the SSC. Given the huge capital investment in the SSC and its detectors, it seems foolish not to spend adequate resources in an attempt to ameliorate this limitation. It appears that parallel processing will be a key element in this.

In order to achieve the highest useful event rates, the attitudes of the experimentalist will have to undergo a radical change. Computers must be considered an integral part of the experiment and taken into account from the earliest design phases. The output from the online computers must be highly compacted, and this process must be thoroughly debugged and eventually trusted in much the same way as the hardware is debugged and trusted.

In closing, I quote the conclusions of the Snowmass-84 group:¹⁵

"This group recommends that an intensive R&D program be mounted very soon to work out a reasonable solution which would involve computer manufacturers as much as possible as well as the software expertise now resident in the high energy physics community."

"The compute power available has a crucial impact on the scientific productivity of an experiment. Too often in the past the physics output has been limited because of: (i) long processing times to get reasonable statistics, (ii) wasted efforts on translations, reformatting, and reprocessing because of poor communications, and (iii) limitations on physics analysis due to inadequate graphics and poor access to data. Intelligent planning, begun now, can reduce these problems at the SSC."

Acknowledgements

I would like to thank the many people with whom I discussed this subject, and especially A. Firestone, T. Nash, D. Quarrie, and J. Simpson.

References

1. R. Diebold et al., Proc. 1982 DPF Summer Study on Elementary Particle Physics and Future Facilities, R. Donaldson, R. Gustafson, and F. Paige, Eds. (Fermilab, 1982), p. 307; R. Huson et al., *ibid.*, p. 315.
2. S. Wojcicki et al., Report on the 1983 Subpanel on New Facilities for the U.S. High Energy Physics Program (DOE/ER Report, Department of Energy, Washington, D.C., 1983).
3. M. Tigner et al., Superconducting Super Collider Reference Design Study for U.S. Department of Energy (May 8, 1984).
4. B. Winstein, "The Nature of Experimentation at the SSC," this conference.
5. See, for example, Ref. 4 or R. Diebold, Science 222, p. 13 (1983).
6. S. Fernbach, "Computers: the First Six Generations," this conference.
7. A. Dragt et al., "Computing Requirements for SSC Accelerator Design and Studies," Proceedings of the 1984 Summer Study on the Design and Utilization of the Superconducting Super Collider, R. Donaldson and J. Morfin, Eds. (Fermilab, 1984).
8. C. Quigg et al., "Report of the Subpanel on Theoretical Computing of the High Energy Physics Advisory Panel" (September 1984).
9. D. Ayres et al., "Design Report of the Fermilab Collider Detector Facility," CDF-111 (Fermilab, 1981).
10. J. Ballam et al., "Future Computing Needs for Fermilab" (December 1983).
11. I. Gaines et al., Proc. of the Symposium on Recent Developments in Computing, Processor, and Software Research for High-Energy Physics, Eds. R. Donaldson and M. Kreisler, p. 183 (Fermilab, 1984).
12. D. Husby et al., "The Fermilab ACP Multi-Microprocessor Project," this conference.
13. J. Ballam et al., "Computing Requirements at the SSC," Proceedings of the 1984 Summer Study on the Design and Utilization of the Superconducting Super Collider, R. Donaldson and J. Morfin, Eds. (Fermilab-1984).
14. G. J. Feldman et al., "4 π Detectors," *ibid.*
15. L. E. Price et al., "Triggering at a High Luminosity Hadron Collider," *ibid.*; R. Diebold and R. Wagner, "Physica at 10³⁴ cm⁻² s⁻¹," *ibid.*

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.