Fermilab Industrial Affiliates Roundtable On

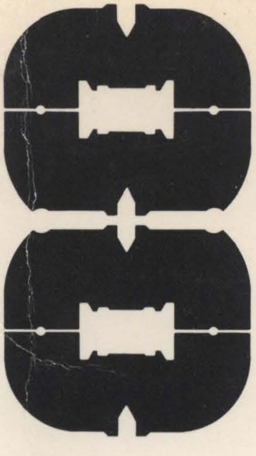

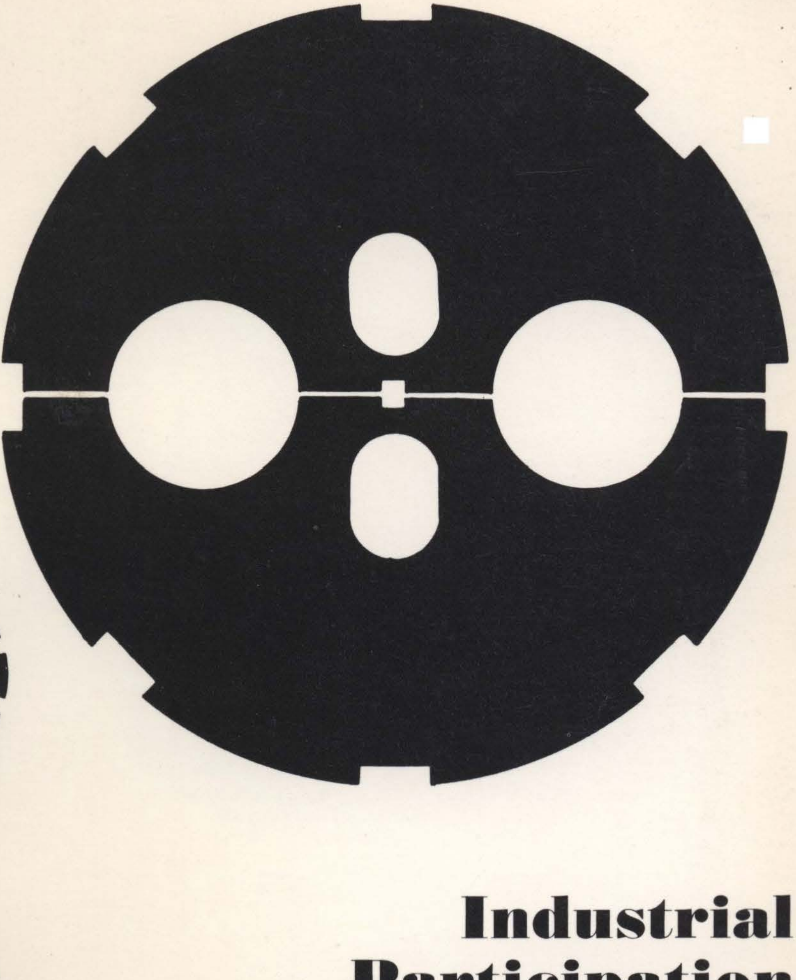

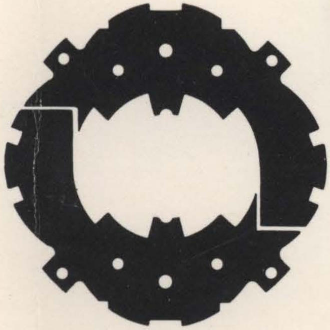

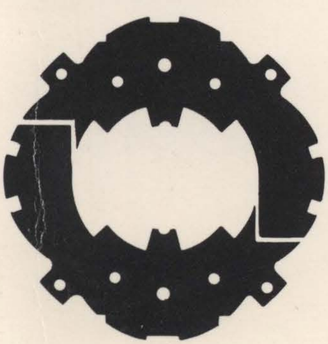

Industrial Participation In Large Science Projects

May, 1984

Fermilab Industrial Affiliates
Fourth Annual Meeting, May 24-25, 1984

Sponsored by Fermilab
And the Fermilab Industrial Affiliates

Fermi National Accelerator Laboratory
Batavia, Illinois

Operated by Universities Research Association Inc.
Under Contract with the United States Department of Energy

INTRODUCTION

Leon M. Lederman, Director
Fermi National Accelerator Laboratory

The Fermilab Industrial Affiliates are a group of more than thirty companies with interests in the research and development work underway at Fermilab. The principal motivation for the Affiliates was to provide a mechanism for ready access to this work, as well as to the work of physicists from seventy or so universities who work at the accelerator. Our experience has been that the Affiliate program's major value is as a forum for communication between the academic and industrial research communities.

Why Affiliates?

The Affiliates do represent an effort on the part of Fermilab to address a larger responsibility and need of science. Basic research in such an exotic subject as particle physics is essentially a cultural activity; however, it is a public trust and the substantial expenditures are justified in terms of long-range benefits to society. The Affiliates aim at exploring ways to hasten and even institutionalize the benefit processes. In addition, modern science is intimately dependent on industrial technology which in turn is beholden to earlier basic science. This interdependency must be understood and fostered. It generates a non-linearity such that advances in the present decade exceed those of the previous three or four decades.

A central feature of the Affiliates annual meeting has been a round table on some important topic. Earlier round tables covered university-industry relations and supercomputers. This year the theme was "Industry and Large Scientific Projects - Particle Accelerators and Projections into the Future: A Super Accelerator." The theme was designed to explore industrial attitudes toward large basic research projects at the leading edge of technology.

Over the last year there has been intense consideration of the possibility of an accelerator twenty times the size of the Fermilab Energy Saver. Serious discussions were initiated in the summer of 1982 and these culminated in a recommendation by the High Energy Physics Advisory Panel in July of 1983. This was to construct, as the highest priority for the field, a proton-proton collider with each superconducting ring having an energy of 20 TeV. Now this did get the attention of the U.S. Department of Energy and it has the enormous enthusiasm of the high-energy physicists. It is called SSC or Superconducting Super Accelerator.

Need For An SSC: Particle Physics Primer

The confluence of several factors served to stimulate interest in the project. First, operation of the Fermilab Energy Saver has emphatically demonstrated that a superconducting accelerator will work. Recent monumental physics discoveries at CERN have shown that experiments can be performed with colliding proton beams at an energy approaching that suggested for a super accelerator. Most important, the veil has begun to lift on the physics of the future and a host of questions lie waiting for a super accelerator.

At first glance, the present picture of basic matter seems almost perfect. The "standard model" of matter has twelve objects divided up into two classes called quarks and leptons. There are six quarks and six leptons that come in three generations. Now the notion is that these twelve objects are the simplest objects that can be found. These particles are supposed to be structureless. They have no insides; they can't be taken apart. So they're literally point objects. That doesn't mean they don't have rich and differentiating properties. They have masses, they have electric charges, and they are subject to forces in different ways. Everything in the universe can be made by combining these objects together. For example, neutrons and protons are made by combining three quarks, while atoms are made by attaching leptons to the protons and neutrons built up from the quarks. Atoms make molecules, and molecules make Industrial Affiliates and all sorts of other things.

There are also four forces: the electromagnetic force, the weak force, the strong force, and gravity. They have different strengths, they have different ranges, and they are enormously different. The forces are described in terms of fields, and the fields are quantized. The quanta of the fields are the force carriers. Characteristically, there's a great mathematical similarity in the description of these forces. The force carrier for electromagnetism is the photon or quantum of light. In the weak force, there are three force carriers: the W^+, W^-, and the Z^0. For the strong force, there are eight carriers, called gluons. There isn't a quantum theory of gravity yet, so not much is known about the carriers, nevertheless they are named; they are called gravitons. A strong motivation exists for trying to unify these forces, i.e., for finding an underlying concept out of which the apparently diverse forces emerge as artifacts of our peculiar situation as observers. Indeed, some success has been achieved in a joining of the weak and electromagnetic forces.

Now this simple picture can almost fit on a T-shirt (soon to be available from Friends of Fermilab). This encapsulates all the published data from all the world's accelerator laboratories for the last 3,000 years. This picture, of course, is highly symbolic, but nevertheless, it tells everything there is to know about all the data.

This picture works! In particular, it predicted the masses of the W^+, W^-, and Z^o to three significant figures. The picture has also made it possible to establish strong links between the origin of the universe and the physics of these fundamental particles.

The Open Questions

But there are open questions when a closer look is taken at the standard picture. Some are intuitive, but some are very disturbing. These questions are a roadblock to progress. The number of publications in the theoretical journals is zooming up because there's no data to limit speculation. What are some of the problems? A dramatic illustration is the Z^o. All the force carriers should have zero mass. Indeed, this is true of the photon and the gluons. However, the Z^o has a very heavy mass. That's been a puzzle. A gentleman named Higgs found a theoretical mechanism for generating that mass. This leads fundamentally to a deeper question, the problem of the origin of mass. All of the theoretical speculations on how the Higgs mechanism might be observed experimentally seem to point to a region of collision energies of the order of 1-2 TeV in the center of mass. There are a lot of Higgs-related speculations that go under the names of supersymmetry and technicolor, which all point to hypothesized objects with masses somewhere in the region of 1 TeV. Another issue is why are there three generations of quarks and leptons? There's a large and seemingly arbitrary set of parameters in the standard model. The quark and lepton masses aren't really understood. Are the quarks and leptons <u>really</u> point objects? There are a lot of speculations about possible substructure of quarks; maybe there are little people running around inside quarks, or something simpler than quarks which would give fewer basic objects. If so, the place to start looking is around 1 TeV. The energy domain at which the SSC will operate is designed to address these questions and any new ones that lie in the future.

The SSC

The possibility of accelerators an order of magnitude larger than the Energy Saver have been considered since the original Fermilab machine went into operation. By 1975 there were serious discussions concerning a Very Big Accelerator or VBA. High-energy physics leaders met in New Orleans that year to map out a ten-year plan to study the possibilities. In 1976, the International Union of Pure and Applied Physics established a committee for future accelerators. International workshops were held at Fermilab in 1978 and Les Diablerets in Europe in 1979. These discussions were capped with the concrete proposal at Snowmass in 1982 to build a superconducting super collider.

By its nature an accelerator twenty times the size of Fermilab will require industrial participation. Industrial involvement is necessary because of the large scale; something on the order of 10,000 magnets have to be built. However, the scope of participation needs discussion. Clearly industry will provide the materials but that's not what we meet about here. Industrial participation on a deep and very technical level is desirable if it serves technology transfer. How to do this without increasing costs and risks is not clear.

In the following discussion, we have assembled a group of industrial experts to address these issues. We have also included representatives from Japan and West Germany in order to explore any differences in attitudes towards these problems.

We at Fermilab have found the round table dramatically illuminating. We hope that publication of the record will help many more in industry, government, the universities, and the national laboratories to understand the factors that influence the character of industrial participation in large-scale science projects in general and the super accelerator project in particular.

If you find this round table interesting, you may want to consider membership in the Fermilab Industrial Affiliates. More details are given on page 127.

[Editor's Note: The round table was organized with the help of Dick Lundy and Dick Carrigan. Dick Carrigan edited the proceedings; Rene Donaldson, Cathy Gianneschi, and Sue Grommes prepared the publication. The cover was designed by Angela Gonzales.]

TABLE OF CONTENTS

BRIEFING CONTRIBUTORS

Dr. Leon Lederman is Director of Fermilab and is spokesman of the group who discovered the Upsilon particle at Fermilab.

Dr. Richard Lundy is Head of the Technical Services Section at Fermilab. This is the group that was responsible for the Doubler magnet construction.

Mr. Claus Rode is Head of the Cryogenic Systems Group at Fermilab. This is the group which did the cryogenic system design for the Energy Doubler.

Mr. Paul Gilbert is Senior Vice President of Parsons, Brinckerhoff, Quade & Douglas, Inc. He is responsible for the preconceptual design and cost estimate on the conventional facilities for the proposed Superconducting Super Collider.

Dr. Dixon Bogert is Head of the Controls Department in the Fermilab Accelerator Division.

Dr. Mel Shochet is Associate Professor of Physics at the University of Chicago. He is now working on the design of the trigger electronics for the Fermilab Colliding Detector Facility.

THE CHALLENGES OF A SUPER ACCELERATOR

Richard Lundy
Fermi National Accelerator Laboratory

In early June the U.S. Department of Energy formally received a reference design for a Superconducting Super Collider (the SSC). The design is summarized in 27 pounds of paper. Without giving away the price, the calibration seems to be a hundred million dollars per pound.

The reference design is an important document. It provides some assurance to DOE that the SSC is not strictly a wild blue sky proposal. There is a technical possibility for building it and for a price that is close to the stated price. The reference design shows that the SSC doesn't require radical new inventions that might never be made. More important, the reference design forces people to think very seriously about the SSC and to begin to optimize features and make compromises.

Probably no one is foolish enough to believe that the high-energy physics community is going to build a machine of this size or this price out of its back pocket. The project is going to take continued industrial help, possibly by a new avenue that hasn't been explored yet.

Very large superconducting accelerators are under construction in several other places in the world. In Germany, the DESY High Energy Physics Laboratory now has approval to bore a tunnel roughly the same size as the Fermilab tunnel under a populated area (Fig. 1). This will be used to collide circulating electrons against protons in a project called HERA. They will use superconducting magnets to put together a proton ring of 820 GeV. With certain design changes that they are now discussing, the energy might even be higher, perhaps 1000 GeV. The work that DESY is doing on magnets is logically connected to the work going on at Fermilab. Naturally, with the passage of time and additional thinking, there have been improvements, and they are still learning and modifying their plans. Fermilab is developing close ties to DESY and will be able to learn from their work. At the same time, DESY will be able to profit from the operating experience at Fermilab.

Another real project is underway at CERN, in Europe. They are boring a tunnel under France and Switzerland through solid rock to hold a large electron-positron collider, LEP (see Fig. 2). The tunnel has a circumference of 27 kilometers or roughly four times the circumference of the Fermilab ring. LEP as it stands will use conventional magnets. Since it's a machine that stores electrons, there is an incredible energy burden due to synchrotron radiation from the electrons that forces them to go to weak magnetic fields. As a result, they have to build a

Fig. 1. Environs of DESY showing the future location of an accelerator under a populated area.

large diameter tunnel. Naturally it offers the future potential for a large proton machine somewhat along the lines of the SSC. There are going to be severe problems in installing a superconducting proton machine in a tunnel that will be chock full of magnets and experimental apparatus. Whether this is a credible alternate to an American SSC or even an International SSC remains to be seen. However, this tunnel is being bored, LEP will be built, and some experience will be gained on a very substantial project.

The SSC is a large undertaking. But just how large? The existing machine at Fermilab, the Tevatron, is a circle a mile and a quarter in diameter (Fig. 3). The experimental areas stretch out for more than a mile. The Tevatron might serve as an injector for the SSC at one twentieth of the main machine energy. So, although Fermilab is a good size piece of real estate, it's really pretty small on this scale. The SSC is going to be a big machine.

Fig. 2. LEP tunnel location at CERN in Geneva, Switzerland. Note the airport in the foreground. France is in the upper portion of the figure, Switzerland in the lower part.

Some people think that the center of the world is Washington, D.C. There is a beltway around it that could just about hold the smallest proposed SSC ring (Fig. 4). No one appears to be seriously considering the beltway as a site. Now consider another scale and a more reasonable location. The smallest proposed ring for the SSC is 20 miles in diameter. This

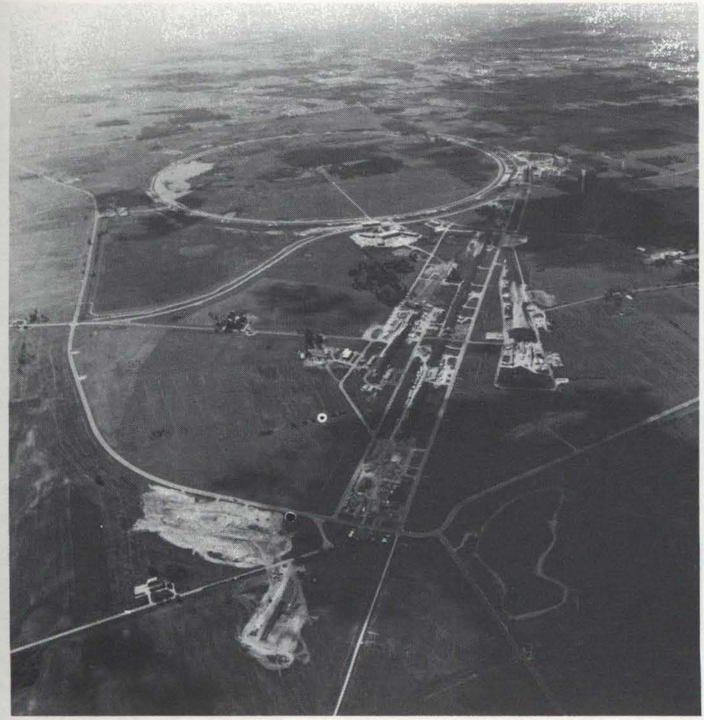

Fig. 3. Aerial view of Fermilab. Note the experimental areas extending away from the ring.

is the ring for the highest field magnet. The ring would fit comfortably into some area west of Fermilab. For a lower magnetic field and correspondingly larger diameter the accelerator ring could go east under the Sears tower in the center of Chicago; however, once the geology is understood it is likely that a proposed Illinois site will be to the west of Fermilab in Kane and DeKalb counties. Figure 5 shows the northern part of Illinois with a typical SSC ring sketched in.

The Reference Design

The reference design treats three fairly different magnet approaches, each with a different field. The different fields lead to different tunnel lengths and the costs vary from design to design. In what follows, reference design B, largely based on Tevatron concepts, will receive the most emphasis. At the moment it happens to have the highest cost. This should be taken with a

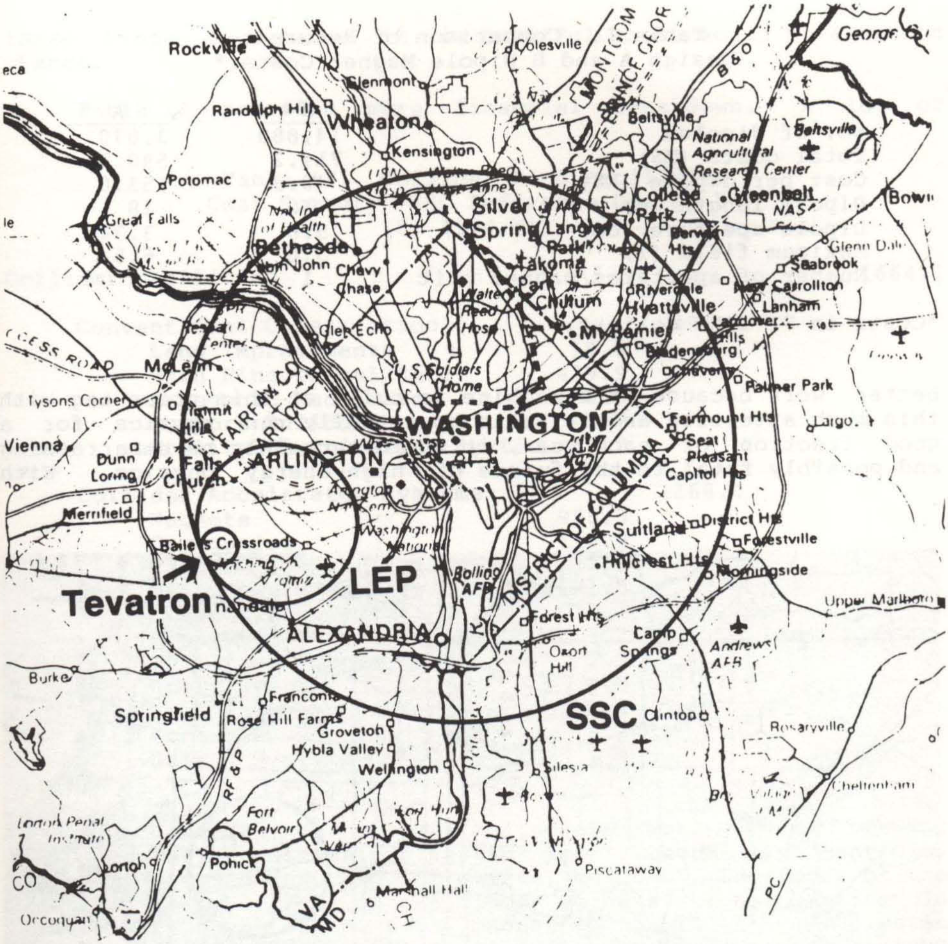

Fig. 4. The smallest diameter version of the SSC along with two smaller colliders, LEP at CERN and the Tevatron at Fermilab, superimposed to scale on the environs of Washington, D.C.

grain of salt. The machine won't be built by following any of these designs, and it won't cost the amount mentioned in the reference design, but these prices are still good guidelines (see Table I). It is a big cost and what is scary about it is that a small overrun on a project this big is an embarrassing sum. Now two things have to be folded into a design like this. The cost must be known, and it must be possible to deliver it for something on that order or preferably less. And, the design had

Table I. Comparison of Reference Design A and B Dipole Magnet Costs.*

	B	A
No. of dipoles	14,880	3,870
Total costs (M$)	719.2	590.9
Cost per dipole (K$)	48.3	153.0
Dipole length (m)	12	18
Dipole aperture (cm)	5.0	3.34
Maximum field (T)	5.0	6.5
Number of apertures/magnet unit	1	2

*Costs in FY 1984 dollars.

better work because if a machine is built at this price tag with this much attention and it fails to actually run physics for a good fraction of each year, it would certainly be embarrassing and possibly fatal to the future of high-energy physics. With

Fig. 5. Map of northern Illinois showing a schematic view of one possible SSC location.

these kinds of stakes it is doubtful there would be a second chance.

Table II provides more financial breakdown. A lot of different parts go into an accelerator. There is the radiofrequency accelerating system (rf), there are controls, there is cryogenics, and, of course, magnets. The focus of the discussion here will be on the dipole magnets which comprise 70 to 80 per cent of the total magnet cost. They'll serve as an adequate illustration of the challenge. The three magnet designs in the reference document represent a diversity of opinion about what is the best magnet to build, what's the most reliable kind, what could go into production on a minimum time scale, and what the cost of the three designs are. These three designs represent the popular magnet concepts at the moment. All were costed and analyzed as to the impact on the whole project. There are also some more radical designs that have not been included. They will receive attention at a later stage. The three designs have been labeled A, B, and C in the reference design report.

Table II. Collider Ring Facilities
Cost Summary--SSC Reference Design B
(FY 1984 M$)

Collider Facilities			1664.1
Conventional Construction		496.1	
Land Improvements	18.8		
Main Ring Tunnel/Encl.	419.9		
Cryogenic Facilities	8.0		
Support Buildings	1.7		
Utility Distribution Syst.	47.7		
Collider Accelerator Systems		1168.0	
Magnets	955.3		
Cryogenics	115.9		
Vacuum	12.2		
Main Power Supplies	21.0		
Correction Element PS	13.4		
RF System	7.7		
Injection System	2.9		
Abort System	8.0		
Beam Instrumentation	7.0		
Controls	19.3		
Safety Systems	5.3		

Design A, the one that is featured in the reference design, pushes niobium-titanium conductors pretty much to the present state of the art. Figure 6 shows the design A cross section. Design A incorporates a 6.5 tesla magnet (65 kilogauss). The

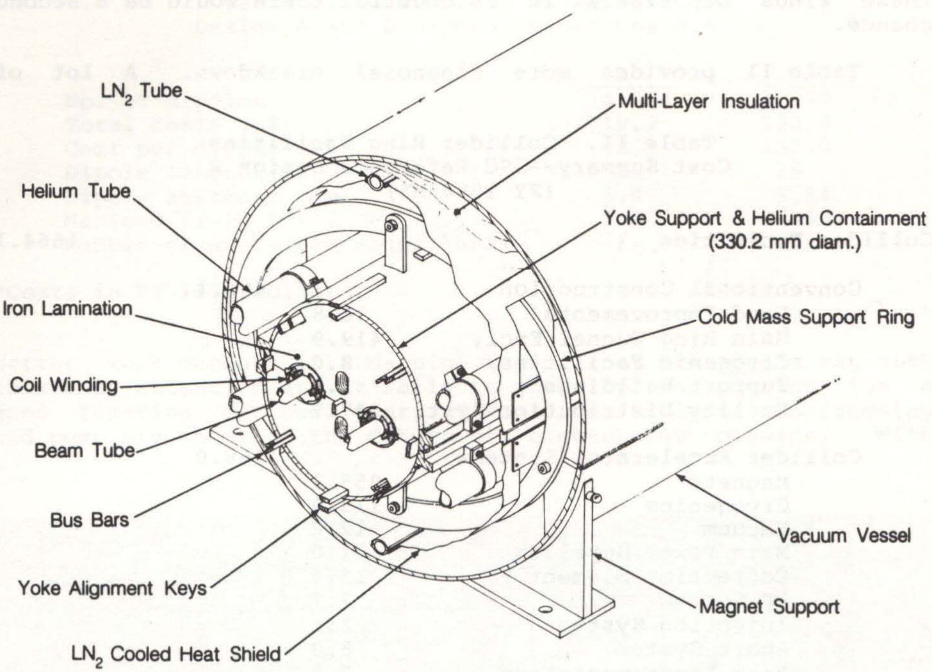

Fig. 6. Cutaway perspective view of the 6.5 T, 2-in-1 SSC dipole magnet.

magnet uses cold iron to generate some portion of the field. It takes advantage of a trick; the iron is not saturated in the central region between the apertures where the two dipoles with fields in opposite directions share the iron. In order to keep cost down the magnets are made long, close to 18 meters. It is only necessary to build about 4000 dipoles that way, but remember each dipole is dual, it has two superconducting coils inside. There may be some problems with that magnet that are not yet anticipated. At the moment probably not more than two or three magnets of this geometry have ever been built and tested. But that will change soon.

This magnet design is the outcome of a collaborative effort on the part of Lawrence-Berkeley Laboratory (LBL) and Brookhaven National Laboratory (BNL). General Dynamics, one of the firms represented on the Round Table, made some of the cost estimates. There are two vacuum pipes for a beam in which protons circulate going in opposite directions. There are two cosine theta superconducting windings close to the vacuum pipes to generate the dipole fields. Iron is wrapped around those windings to return the magnetic flux and to shield the two apertures from each other. As usual there are a lot of cryogenic pipes to make things cold and carry the fluids around the long refrigeration circuits. There are fiber glass epoxy resin rings to support the assembly inside a vacuum vessel using low heat leakage technology.

A strong group based in Berkeley and Brookhaven will be experimenting with this magnet. As an offshoot of this design, a group primarily at Brookhaven is considering even smaller apertures. Instead of an aperture of four centimeters they are considering one something like three centimeters. They would use niobium-tin, which is a much higher field conductor. As a result, they will try for fields of 8-10 tesla. If the BNL group succeeds in that, the tunnel would get smaller, certain civil engineering costs and other length-related costs would go down and to some extent the number of site possibilities might even expand because it's not possible to fit a 5-tesla circle on Long Island. Note that a 10-tesla magnet made with niobium-tin and wire bending radii of an inch or less is a challenge.

Design B, the Fermilab proposal, is closely related to the Energy Doubler because that is what has just been done, and it is known how to do it. It pushes the conductor a little less hard to give a field of 5 tesla. It essentially derives no field from the iron. Each aperture is put into a separate cryostat. That way it's possible to think of building only one ring and going to a proton-antiproton option initially. Figure 7 shows the design B cross section. Design B has a modest length of 14 meters. That may change; it may get longer or shorter. A lot more units are needed, close to 16,000 separate dipole magnets with a correspondingly larger number of quadrupoles. Naturally, the tunnel gets longer going from a 90 kilometer circumference in design A to 113 kilometers in design B. Initially the system was designed with the two apertures one above the other instead of side by side (the figure should be rotated by 90 degrees). This design, which has two independent dipoles with 5 centimeter apertures, uses a cosine theta superconducting winding and a set of aluminum collars to take the Lorentz forces. Again, there is a lot of piping to carry the cryogens around. But now there is just enough iron to make a vacuum vessel and to prevent the magnetic field of one dipole from influencing the beam in the other.

Fermilab Doubler magnets would cost something like a factor of three or four more than design B, because of their more restricted design. They would also have a much higher heat leak than the proposed design B. That would lead to refrigerator costs which would be ten times higher. The Doubler magnets use superconductor less efficient, and they do a lot of other dumb things that seemed smart six years ago.

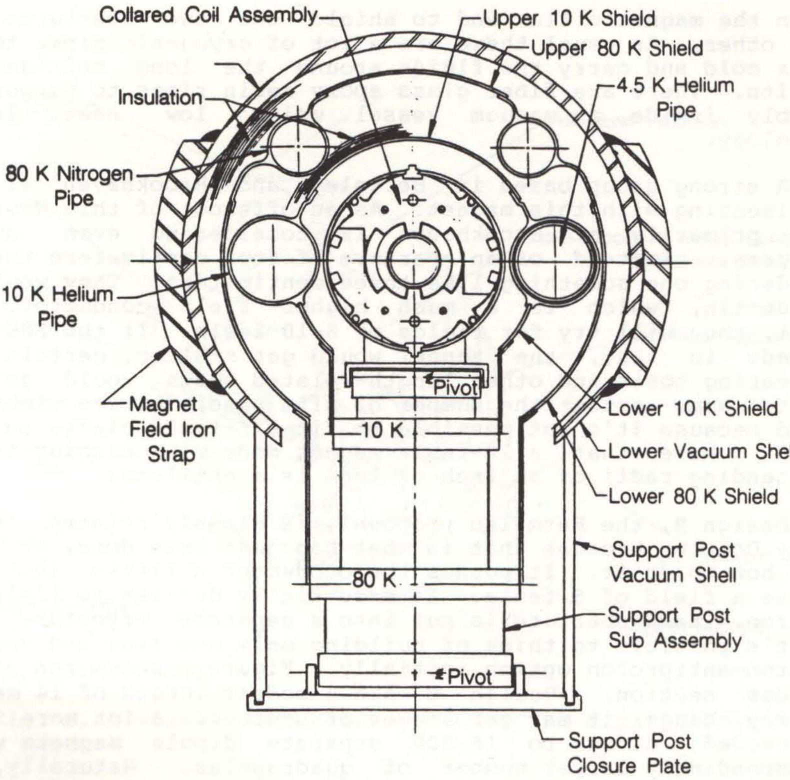

Fig. 7. Reference Design B cryostat cross section showing major components.

Design C was developed by a group from Texas where Russ Huson, Peter McIntyre, and others have picked up on the idea of a superferric magnet. This is a magnet operating in the 2-3-tesla region in which iron forms the main component in the magnetic circuit, and superconductors are chosen as a clever way to excite that iron without a great deal of power investment. Figure 8 illustrates the magnet. What's radical about that magnet? Well the aperture is fairly small. Naturally, since it's an iron magnet, more of that aperture is uniform field than any of the others. The length, 140 meters, is designed to minimize the cost of the ends and of other features. One hundred and forty meters is long, a football field and a half. A lot of light airplanes land and take off in a distance less than that. No one has ever made an accelerator magnet even approaching that length. It would be a tour de force to learn how to build them, but only a thousand would have to be built.

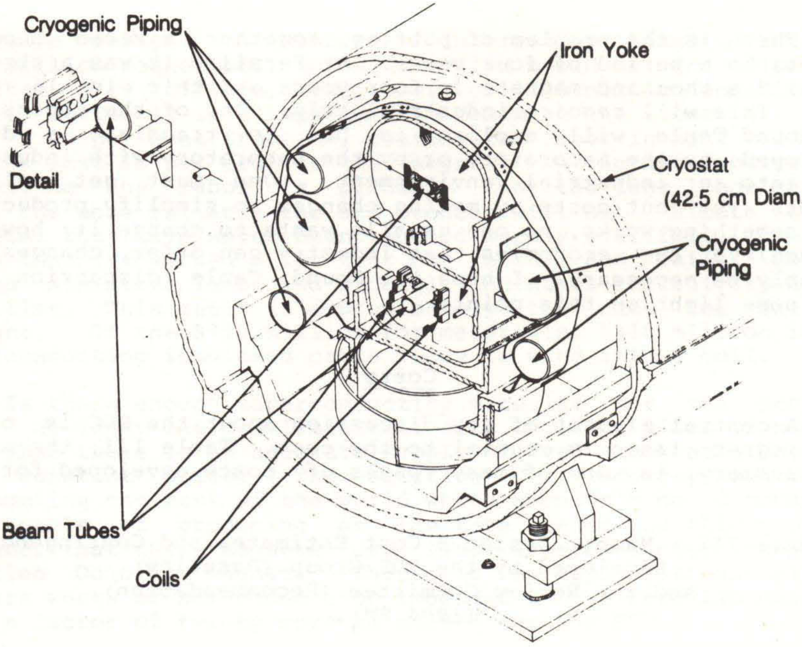

Fig. 8. Cutaway perspective view of the 3T, superferric dipole magnet.

The superferric magnet for design C has a rectangular aperture excited by superconducting windings similar to a classical copper and iron dipole. This leads to a small beam tube. The two apertures are located one above the other. Now a six meter magnet loses about 3/4 of an inch of length when it shrinks. A 140 meter long magnet will have substantial shrinkage. As a result, the ends of this device have to be capable of moving several inches without breaking. There are some substantial engineering challenges involved in that.

In what follows the primary concentration will be on design B. What are some of the technical and logistical challenges with design B? The beam pipe may have to be copper plated inside for low conductivity. Aluminum collars will be used to hold the magnetic forces. Stainless-steel pipe will be employed in conjunction with the beam pipe to form an annular vessel for liquid helium. Heat shields at various temperatures will be needed to help minimize the heat load of the structure. A set of posts made out of epoxy fiberglass will be used to hold the coil mechanically and not introduce a large heat leak. A carbon steel vacuum vessel is part of the thermal insulation. Iron is required to help return the magnetic flux and shield one aperture from another.

There is the problem of putting together sixteen thousand magnets in a period of four years. At Fermilab it was a struggle to build a thousand magnets in four years and this is 16 times more. This will require industrial help. One of the topics that the Round Table will explore is how to translate a design developed by the Laboratory or by the Laboratory with industrial help into an industrial environment. One must get reliable magnets yet cut costs by making changes to simplify production. Once something works, no one usually wants to change it; however, to achieve the economies that industry can offer, changes will probably be necessary. I hope the Round Table discussion will shed some light on this point.

Costs

A central element of any discussion about the SSC is costs. The magnet issue is crucial to the cost. Table III, the magnet cost summary, is one of many pages of costs developed for the

Table III. Magnet Design B Cost Estimates and Contingency Developed by the RDS Group (Baseline) and the Review Committee (Recommendation) (1984 $M).

	Baseline		Recommendation	
	Cost Estimate	Contingency (25%)	Cost Estimate	Contingency (25%)
1. Tooling (Use 2 shifts/ day, 5 days/ week operation)	21.3	5.3	36.0	9.0
2. Dipoles	719.0	179.8	719.0	179.8
3. Quadrupoles	60.8	15.2	60.8	15.2
4. Special Devices	75.5	18.9	75.5	18.9
5. Special Magnets	20.0	5.0	11.2	2.8
6. Installation	58.5	14.6	58.5	14.6
7. Factory Supervision (Included above)	N.A.	0.0	N.A.	0.0
Subtotal	955.1	238.8	961.0	240.3
8. EDI (WBS 1.6.2)	95.5	28.7 (30%)	96.1	28.8
Total	1050.6	267.5	1057.1	269.1
Grand Total	1318.1		1326.2	
Recommended Change			+8.1	

reference design. A number of interesting things go into the $955 million magnet bill (of which $719M would be for dipoles). Of that $719 million, only 20% or $130 million would be for labor. Obviously it's important to minimize the labor but with only a 20 per cent fraction, cutting the labor in half will only cut the cost 10 per cent on the magnet. A better motive for minimizing the labor is probably that the less labor that is used, the more reliable the end product will be. Simple blocks of metal that have never been touched hardly ever fail. It's the parts that are worked a lot that can break and fail. So the design minimizes labor and probably increases reliability at the same time. This ratio is approximately the same for all the designs. Of the $719 million for materials, $318 million is for superconducting insulated cable ready to wind into a coil.

Is there enough superconducting material for the project? Well, at least one gentleman that sits on large piles of niobium says that there is an adequate supply of feed stock at a more or less constant price. The existing U.S. industry (and that's discounting the rest of the world which shouldn't be discounted) is capable of producing at the rate needed for this project. Remember that these magnets use less superconductor than the Fermilab Doubler magnet. There is perhaps fifteen million dollars worth of superconductor imbedded in the Fermilab magnets. Only a factor of twenty more is needed for the SSC.

The cost of superconductor is almost half the magnet cost, so it's an important fraction of the total costs. In design B this is particularly true; one could say it is wasteful of superconductor since none of the field is from the iron. However, this avoids all of the problems that come with saturating iron. It also opens the window to one area in which advances may occur in the next few years. Nobody expects the saturation field of iron to improve but the ability of superconductor to achieve high current densities could improve dramatically. If it does, the same design might become not a 5-tesla magnet but a 6 or 7-tesla magnet. Steps have been taken in design B to avoid building anything into the magnet that would limit its field performance. Since the iron plays virtually no role, changes in the field level are not influenced by that. The size of the force restraining collars can be made large enough at this time to handle bigger forces. For design B, it would be possible to either cut the cost or improve the performance. Since superconductor cost is such a large fraction it's possible to concentrate attention on a fairly narrow technical area with the hope of producing large yields.

Cable

This cable, so-called Rutherford Cable, shown in Fig. 9, consists of of 23 or 25 individual strands. Each strand is about 30 mils in diameter. The strands are twisted and each strand has

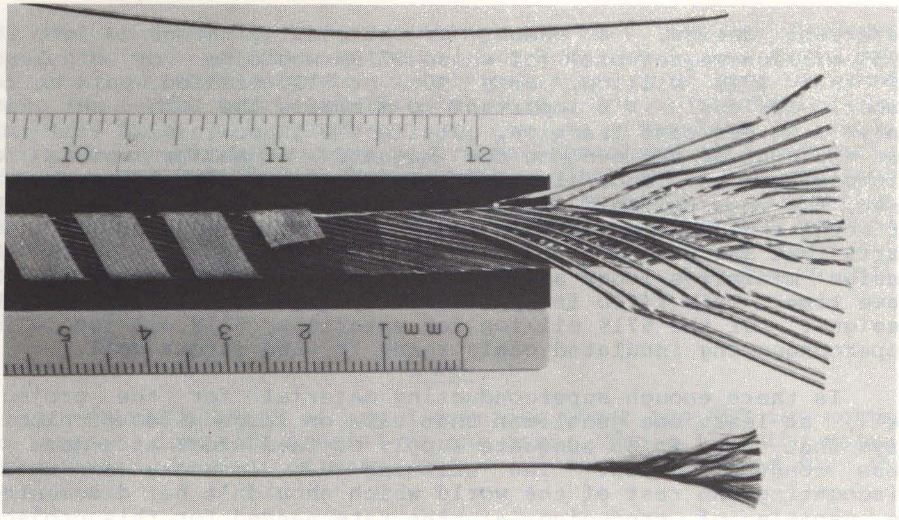

Fig. 9. Rutherford-style superconducting cable.

an internal structure that contains about 500 small filaments of niobium titanium. This is similar to the cable that is used in the Energy Doubler with very small dimensional changes. The quantity is staggering; something close to 100 million feet of this have to be provided. The wire is a simple component, but it's a crucial one. Table IV gives the superconducting cable specifications. There's nothing very important in this table

Table IV. Superconducting Cable Specification.

Conductor type	Nb - Ti (47.5%Ti)
Copper to superconductor ratio,	2.0/1 (inner),
(by volume)	3.0/1 (outer)
Number of filaments	710 (inner),
	533 (outer)
Filament size	16.5μm
Strand size	0.0298 in.
Number of strands in cable	25 (inner)
	23 (outer)
Cable dimensions	
With insulation	$0.053\times0.369\times0.065$ in^3 (inner)
	$0.053\times0.340\times0.065$ in^3 (outer)
Bare metal	$0.048\times0.364\times0.060$ in^3 (inner)
	$0.048\times0.335\times0.060$ in^3 (outer)
Cable short sample current (5T, 4.2K)	6650 amps (inner)
Miits to 500K	15.4 (inner), 12.5 10^6 amp^2-sec

except one number. The superconductor that went into the Energy Doubler had a current density of something like 1800 amps per square millimeter at 5 tesla and 4.2K. This was almost an industry standard for the late 70's and early 80's. All of the designs in the reference design are assuming 2400 amps per square millimeter under the same conditions. There are good reasons to think that could be achieved. There is some evidence that it's possible to go even higher. The more the current density goes up, the fewer dollars it takes to build the machine or the more beam energy that will be available. So it's worth working hard on the current density that can be achieved in a superconductor, and the subject has got to be attacked very vigorously in the next few years.

Magnet Assembly

Once the superconductor is as good as one can make it, it has to be put into cable. The cable has to be wound into long coils of semi-cylindrical cross section. Each cable turn has to be very accurately placed. The cables have to be placed to an accuracy of 2 mils or better and held there against strong magnetic forces. If this can be done, the magnet will have a suitable field quality for the accelerator, tolerances of one part in ten thousand. It's not possible to see the keystoning as it's called in this wire, but to put this package together these cables have to be thinner at one edge and thicker at the other. If that is carried to extremes, the superconductor is damaged and the premium density that has been fought so hard for is lost. A compromise is reached by tilting some of the current elements and filling in with inert wedges. The same trick can also help with the field uniformity. The insulation, winding, compaction, and collaring of this coil are all very tricky steps. A failure to maintain the dimensional tolerances or the pre-load on the cable will lead to a useless magnet.

It will surely require many prototypes to prove out the particular technique. In the Energy Doubler 100 to 200 prototypes were needed. For the SSC it should be possible to get by with something like 100 prototypes. These will test not only the aspects mentioned above but other design features.

Magnetic quality is important. A radius of 2/3 of the aperture is the usual place to specify the field quality. The dipole field will have a strength of something like ten thousand times any of the non-dipole multi poles. These are given in Table V. Since they are high harmonics they increase rapidly at larger radii. At 1 centimeter they are very small. That makes a very uniform magnet, and uniform magnets make accelerators like the Tevatron that are easy to turn on and operate. One would be ill-advised to compromise on the field quality here for economy, since one might end up building a machine that could never really operate.

Table V. Multipole Coefficients
(As Calculated for a Mechanically Perfect Magnet).

	R=1cm	R=1.7cm(2/3 aperture)
B_0	10^4	10^4
B_2	0.07	0.20
B_4	-0.07	-0.55
B_6	0.09	+2.14
B_8	-0.14	-9.49
B_{10}	0.01	+2.78
B_{12}	0.00	-0.56
B_{14}	0.00	-0.03

It would be nice if only geometric placement influenced the field quality in superconducting magnets, but that's not the case. Particularly, at low fields where the beam is injected, superconductivity stabs one in the back with the problem of persistent currents. These are currents produced inside the superconductor filaments by the changing magnetic field. These persistent currents at low excitation lead to error fields, such as sextupoles, that spoil the perfection of the magnet. The bigger the filaments, the worse the persistent current effects. At low fields the persistent current magnitude is set by the short sample performance of the superconductor, just the thing that one has been trying to maximize. If the short sample current distribution in such a conductor is different from the distribution at high field persistent current, terms will be present in the magnets that will vary in an irregular fashion. That would be very bad. Therefore, it's important to learn on a production basis how the very low field short sample valves in these niobium titanium conductors are determined and to control both the temperature of the magnets and all the other factors that influence the error fields. That's true of both designs A and B. Design C side-steps this problem by using the iron to short out these error fields to a large extent.

The superconductor has to be held in place in spite of the Lorentz forces that amount to hundreds of tons on these coils. Fairly strong aluminum collars have been incorporated to take the force.

The design also provides space to put in electrical heaters. These heaters will be used for the rapid warm-up that is desirable. One of the features of design B is the flexibility to handle magnet failure and replacement quickly by warming up a magnet string. In this case a magnet string is 5 kilometers long. It will be possible to warm-up a string in a 24 hour period, change a magnet, and cool the string back down in 24 hours.

At Fermilab the present experience is that it takes 36 hours to warm-up and eighteen hours to cool-down. Those fast rates put large thermal stresses on components. The SSC design has to be proven to be invulnerable to those kinds of thermal cycles. To break something, just warm it and cool it. Big forces are involved. Parenthetically, in design B there are 16,000 magnets. They all have to work all of the time to make the beam go around.

The choice of a high strength aluminum alloy and design of dies to produce enough collars for 60 miles of magnets is a nice challenge. The DESY effort is using aluminum collars, and it seems to work quite well.

As a cost cutting feature, it is proposed to make all of the heat shields out of extruded aluminum. They would be welded together with automated welding techniques to make the longitudinal joints. Aluminum is a good thermal conductor; its mass is low and it has a lot of other good features. In the past Fermilab has always used stainless steel and manual welding. New welding techniques must be learned. The Japanese at KEK have utilized automated aluminum welding.

The suspension system in the design B magnet consists of a G10 post every three or four meters sitting on the vacuum vessel floor. It's a nested set of cylinders with heat intercepts at intermediate temperatures. Different contraction coefficients must be accounted for so that parts can slide back and forth. There is experience on how to do all that; it's just necessary to find the cheapest and most reliable way to do it.

The end detail of the magnet is something that affects installation cost. Every effort should be made to simplify the ends. Figure 10 shows a plan view of a magnet interface. At the end of the magnet, the fins are trimmed away on the extruded aluminum pipes that are part of the heat shields. Aluminum bellows are welded on and welded to the next magnet. There are big bellows on the outside to complete the vacuum vessel. The 100,000 or more welds had best be done by machine to get a high degree of reliability. Note that the Doubler was not welded together but instead put together with mechanical clamp seals. It's probably fair to say that each seal was made at least twice in order to get one leak-tight joint. Welding the joints for the SSC implies a good reliable magnet since taking out a welded joint is a little tougher than unclamping one.

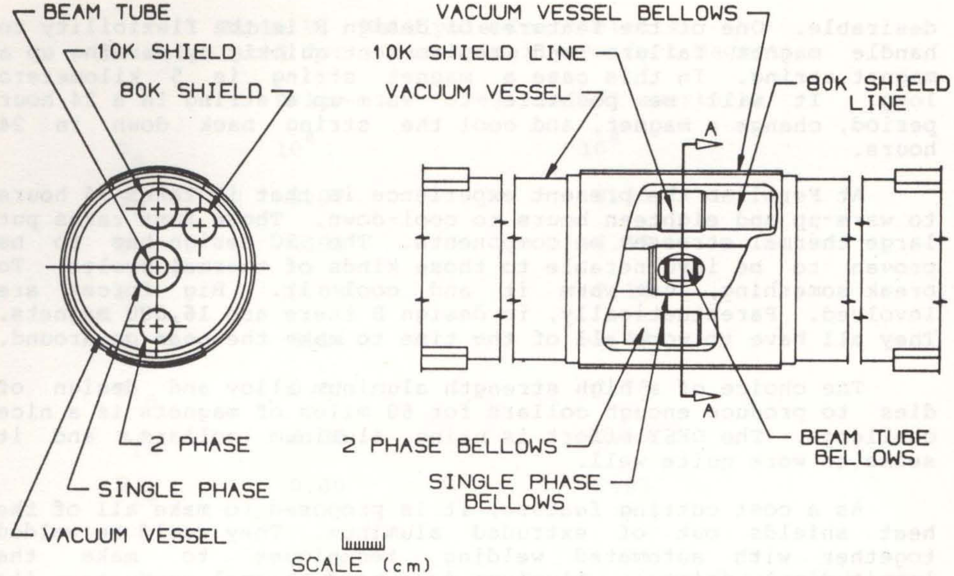

Fig. 10. Plan view of the magnet interface for design B.

R&D Challenge

A nice technical challenge is compensating the persistent field in a reliable way. Measuring all these magnets is another tough problem. It won't be possible to measure 15,000 magnets one by one the way it was done in the Energy Saver where each magnet was individually cooled and measured. The hope is that every magnet can be measured with warm techniques and perhaps one in ten or even one in a larger number will be sampled cold. What that means is that 90 per cent or more of the magnets will be installed in the tunnel, cooled and ramped to full field in the tunnel. Magnets not capable of full-field operation will be removed and replaced.

There are many other R&D challenges. Surveying devices to a fraction of a millimeter in such a large ring is not exactly a proven technique. Quench protection, that is the ability to save the magnet if it inadvertently stops working as a superconductor, has to be made reliable and simple. Table VI contains a complete list of these challenges for design B.

Quality control will be an important part of producing these superconducting magnets. At Fermilab, paper records were kept of who did what to each magnet, what they did it with, and where the parts came from. A computer-based system has been developed with a terminal at every technician work station. This system is now

Table VI. List of R&D Challenges.

- Developing methods to extend to large quantities the industrial production capability of improved Nb-Ti superconductor, and cable with higher critical current density.
- Improving the mechanical cabling process to increase cable performance.
- Developing methods to correct for the effects of persistent currents at low magnetic fields, including both self-energized and externally energized correction coil systems.
- Developing rapid and economical methods for measuring magnetic fields to accuracies of about 10^{-4} in long magnets.
- Developing accurate, labor-efficient methods for magnet alignment.
- Developing a simple and fail-safe quench protection system.
- Developing special superconducting magnets, such as high-gradient quadrupoles and large-bore dipoles for the interaction regions.
- Developing tooling that is appropriate to producing magnets at the required rate.
- Developing coil winding and collaring methods that can eliminate the coil pre-forming step.
- Investigating techniques that allow smaller diameter coils to be fabricated.
- Designing optimized thermal shields to eliminate distortion due to thermal gradients.
- Investigating the behavior of aluminum collars, and the cost/benefit of using more aluminum parts.
- Optimizing the cold-mass support structure.

in use for building some complex assemblies which are part of the TeV I antiproton source. This is just an illustration of the method. The assembly is a one-meter long set of microwave cavities, pick-up loops, nitrogen cooling lines and a delay line combiner board that puts all the rf signals in phase. This is to provide a 1 gigahertz low noise signal indicating the antiproton bunch properties in the machine so that various rf dampers can work on those properties. About 100 of these have to be built to tight tolerances. To do this, a computer terminal is put at each work station and a set of instructions somewhat like a Heathkit manual is loaded in the computer by a supervisor. If someone can read and can recognize part numbers, they can build an assembly by walking into the laboratory and following these instructions.

The point of this is that fairly unskilled labor can be used. When requested, they make measurements on the assembly. They enter those measurements into the computer rather than using pencil and paper. The supervisor sitting in his little cubby hole somewhere can call up any relevant information. He may ask "Where is sub assembly number such and such and how far down the production chain has it got?" The supervisor can also quickly get determinations of the measured properties as a function of serial number or anything else. It doesn't take much computer power to put this work control system in place. A small machine will do it. A system like this will be very useful in the SSC production phase. It provides a quick way of changing the work flow in that new instructions can be inserted without going back to the engineering department and saying "I have to change all these blueprints because I want to change a 10-32 screw to an 8-32."

Summary

Construction and operation of the Fermilab Energy Doubler as well as intensive DOE studies over the last year indicate a Superconducting Super Collider can be built. There will be many challenges. Many people and many institutions will face and solve these challenges. This Round Table was organized to discuss ways whereby the parties could work together to solve the challenges expeditiously and frugally and meanwhile produce a superb accelerator. It is our hope that some progress has been made in that direction.

LARGE CRYOGENIC SYSTEMS

Claus Rode
Fermi National Accelerator Laboratory

The experience with the cryogenic systems at the Fermilab Energy Doubler is an important foundation for extrapolating a design for the SSC. In the material that follows on the SSC, the emphasis will be on design B but the majority of the comments are also directly applicable to designs A and C.

Doubler Cryogenics

A view of the Fermilab tunnel from the Central Laboratory building is shown in Fig. 1, including the Central Helium Liquefier, the satellite refrigeration buildings, and the helium transfer line. The Doubler refrigeration system is described in the schematic in Fig. 2. The Central Helium Liquefier provides liquid helium to 24 satellite refrigerators with their compressor systems. These satellite refrigerators actually are normally running as amplifiers with a liquid helium flow gain of twelve. They each provide one kilowatt of refrigeration to the magnet strings as well as transport 25 liters per hour of liquid helium for lead flow. These flows come back to the compressor system and are returned to the Central Helium Liquefier at 20 atmospheres pressure. This flow is actually ejected into the discharge stream of the central compressors which are running at a lower pressure.

A more detailed schematic is shown in Fig. 3. Above ground there is a compressor discharge header and helium and nitrogen transfer lines. There are two low pressure headers for helium and nitrogen in the tunnel. The helium header serves three purposes: it is the suction header for the compressors, it is the relief header in case of magnet quenches, and it is also the cool-down and lead flow collection header. The nitrogen header has a dual purpose; it is the collection header as well as the relief header for the nitrogen system. While this system worked very well for the Energy Doubler, it cannot be directly scaled to the SSC design. This system provides a very large amount of refrigeration power over a "small" region. If the Doubler system is scaled by a factor of 18, some of these pipe lines would get absurdly long. On the other hand, the refrigeration capacity for the SSC is being scaled by only a factor of three relative to the Doubler.

Figure 4 shows the interior of the Central Helium Liquefier building for the Doubler. The 2000 horsepower reciprocating compressor units are in the front. The Central Helium Liquefier cold box is in the back with the turbines up at the top. The Central Helium Liquifier (CHL) has worked very reliably. The

Fig. 1. View of Fermilab tunnel from Wilson Hall. Note the Central Helium Liquefier at the upper left as well as the satellite refrigerator buildings and the helium transfer line on the berm.

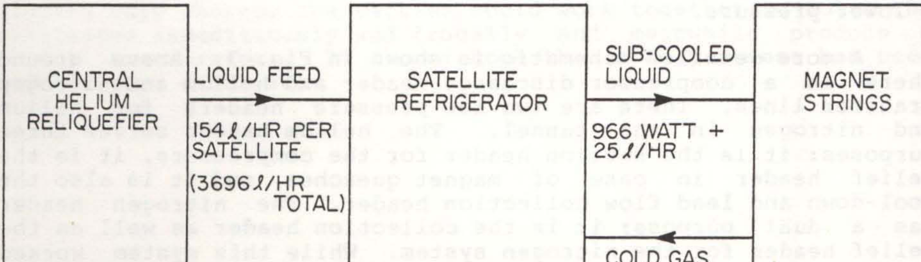

Fig. 2. Energy Doubler helium flow schematic.

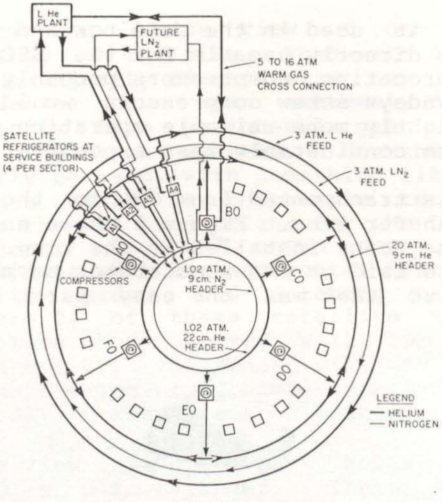

Fig. 3. Layout of the refrigeration system for the Energy Doubler.

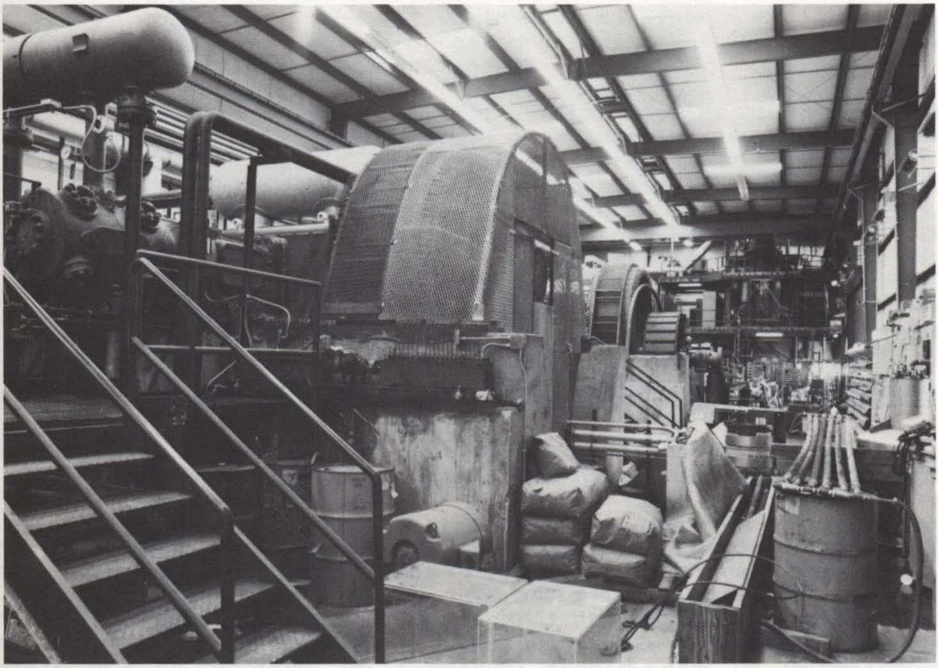

Fig. 4. Interior of Central Helium Liquefier building for the Doubler.

technology that is used in the cold box and in the oil-bearing turbines would be directly useable for the SSC. On the other hand, the reciprocating compressors probably would not be used for the SSC. Nowadays screw compressors would more likely be used due to slightly more reliable operating experience and the fact that they are considerably less expensive.

The liquid is transported from CHL to the ring through a liquid helium transfer line. Figure 5 shows an 80-foot length of this transfer line being installed in the ring. Three miles of transfer line were laid on the accelerator berm in 16 hours using this helicopter, so that was the easy part of this project.

Fig. 5. Installation of eighty foot section of transfer line for the Doubler with a helicopter.

Although the transfer line has worked well for the Doubler, it would not be applicable at all for the SSC. There are two reasons for this. One is the cost of a 70-mile transfer line. More relevant is the fact that the heat load of the magnet system is extremely low so a transfer line would have a heat load very comparable to the magnet system. By putting in a separate transfer line, the total project heat load would be increased by a third. Instead the intent is to use the magnet system as the transfer line rather than having an external one.

Figure 6 shows an inside view of one of the Doubler compressor buildings in the ring. These are 400 horsepower oil flooded screws; the majority of this skid is oil piping. The compressors themselves have been extremely reliable. All the down time is primarily from peripheral gear such as circuit breakers or the interlock systems. For the SSC, much larger units would probably be used, with something like 2000 horsepower compressors. These would be comparable to the compressors delivered both for the CBA project and for the Lawrence Livermore Lab Fusion project. Figure 7 shows one of the satellite refrigeration buildings with the horizontal heat exchanger column protruding from the side. The rotating machinery is in the building. There are 24 of these satellite buildings in the system. Again, they have worked well for the Doubler. The primary hope for the SSC is that there will be fewer of them. From a human engineering standpoint, twenty-four are a few too many. It's difficult to keep track of things.

Figure 8 is a view inside a block house. The satellites have a 30°K and a 6°K expander. These are reciprocating expanders. The expanders have been reliable with a typical mean time between failure of nine months. They would not be used for the SSC because the SSC refrigerators are an order of magnitude larger. Therefore, the SSC would have turbines. Reciprocating liquid helium pumps may end up being used in the SSC, and they would look similar to these expanders.

The actual temperature profile in the Doubler magnets string is illustrated in Fig. 9. Hidden in this graph are 600 channels of information. That is only 2% of the total amount of cryogenic analog data that comes back to the control room. It's a major problem to deal with transducer calibrations, drift and repair, and just keeping track of the units that have died. In this particular version every one of the warm indications is not a hot magnet but just a transducer that is out of calibration. One of the serious problems for the SSC is how to have good reliable thermometry and thermometry that doesn't take continuous maintenance and recalibration.

SSC Magnets

Fermilab actually looked into four different types of posible SSC magnets during the last year. Cold iron magnets were considered first. The other laboratories are now working on those. Because of the tremendous amount of cold mass associated with the cold iron system other options were then considered. The iron generates enormous heat loads during cool-down and warm-up. Typically it would take 1000 liquid nitrogen tankers to cool down a cold iron SSC. Being able to cool down the machine in a reasonable amount of time would need to be a major site criteria. For example, the machine would have to be near a major steel center. While nitrogen is available in the Chicago region,

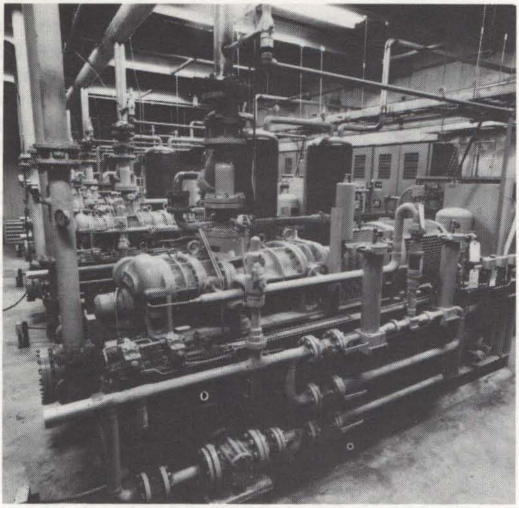

Fig. 6. Inside view of Doubler compressor building in the ring.

Fig. 7. Refrigeration building with horizontal exchanger column.

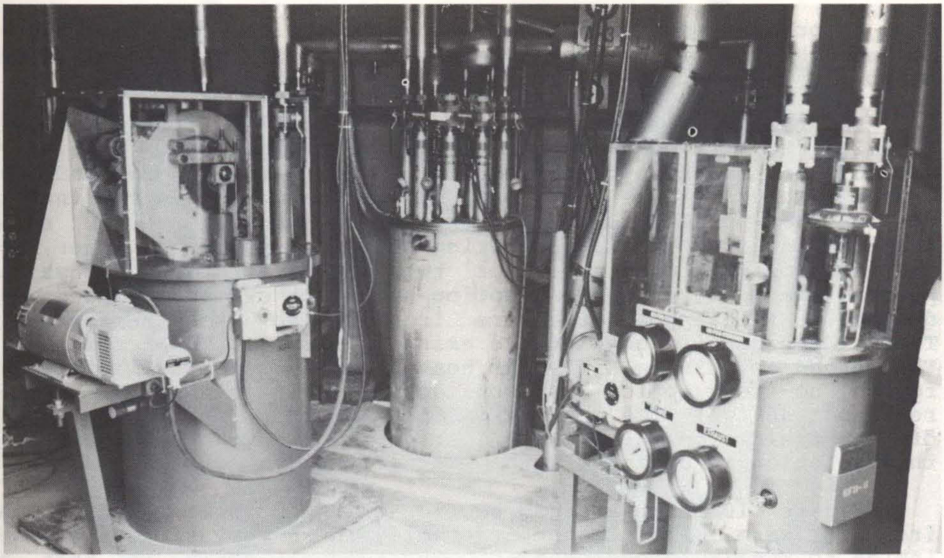

Fig. 8. Inside view of refrigeration building. Notice the expanders.

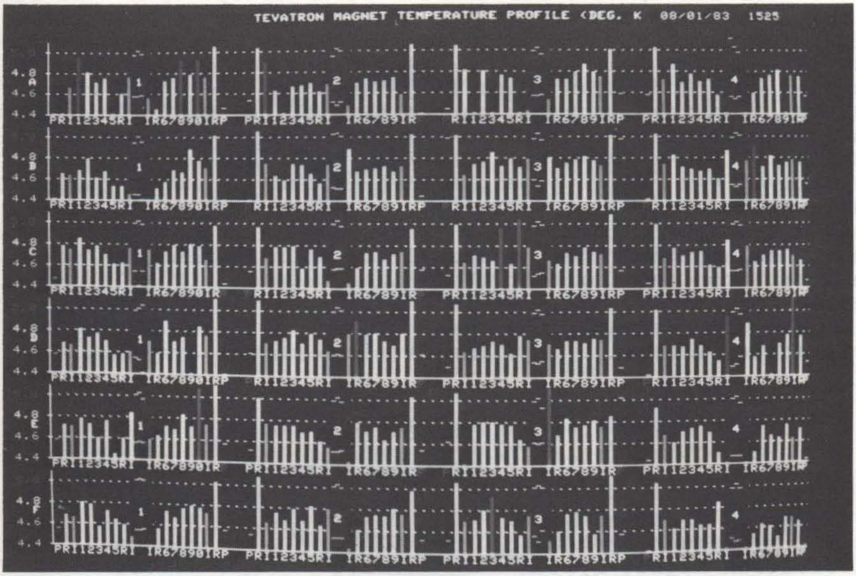

Fig. 9. Temperature profile of the Doubler magnet string. Six sectors of the ring are shown. The scales go from 4.4 to 5.0 degrees kelvin.

this is probably one of the few places in the country where there is enough nitrogen available to properly deal with the cold iron option.

Partially cooled iron was the next option considered at Fermilab to try and get rid of some of the load. Partially cooled iron implies a temperature somewhere between 20 and 80 degrees kelvin. The idea was that the magnet could be running for physics while the iron was still cooling from room temperature. The initial operating heat load would be higher, so the iron would be cool a week or two later after the physics program was underway. That option had a lot of problems and was dropped. At that point the warm iron concept was reexamined. That's the concept used in the existing Fermilab Doubler magnet. That ended up being a difficult heat load problem. Finally a fourth no-iron option was considered. This is design B in the reference design. The penalty is that twenty to forty per cent of the field from the iron is lost. From a cryogenic standpoint this solution is definitely the most attractive.

Figure 10 shows a cross section of the Doubler magnet. The iron is warm. The magnet actually looks very similar on the inside to the SSC cross section. The relevant point is that a great deal of action is jammed into a very small space: a two-phase heat exchange area, a vacuum space, a heat shield, super insulation, and miniature heat intercepts. As the radius of the iron moves out, the field contributed by the iron drops off very rapidly. This is why the no-iron concept is interesting. By the time a reasonable cryogenic magnet is designed, no more than 10 or 15 per cent of the field comes from the iron. So then, why not go the rest of the way? That is what has been done in design B. As a contrast, consider the cross section of the Fermilab transfer line shown in Fig. 11. Basically it's a few pipes, a large amount of super insulation, and large open spaces to work in. As shown in Fig. 12, the SSC design B magnet looks much more like a transfer line than a magnet, except for the intercoil chamber which is basically identical to the current Fermilab magnet. That's important because for the SSC one must have a low heat leak. Table I gives a comparison between the heat loads for the SSC magnet, the Fermilab transfer line, and the Fermilab magnet system.

Table I. Heat Loads Watt/Meter.

	5°K CIRCUITS	80°K CIRCUITS
Tevatron Transfer Line	0.04	0.6
SSC 5T Magnet System	0.19	1.4
Tevatron Magnet System	1.5-2.0	7.5

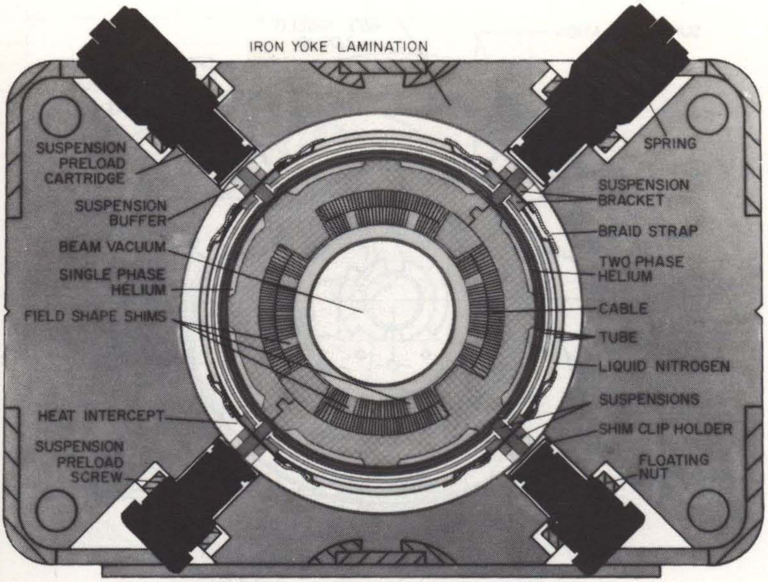

Fig. 10. Doubler magnet cross section.

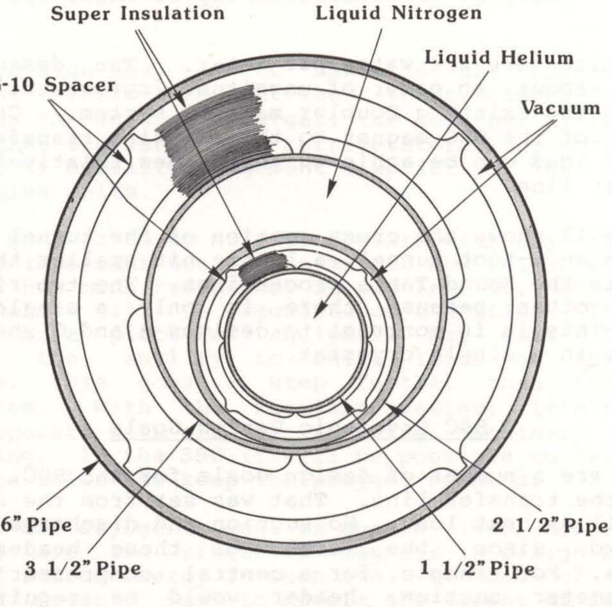

Fig. 11. Doubler transfer line cross section.

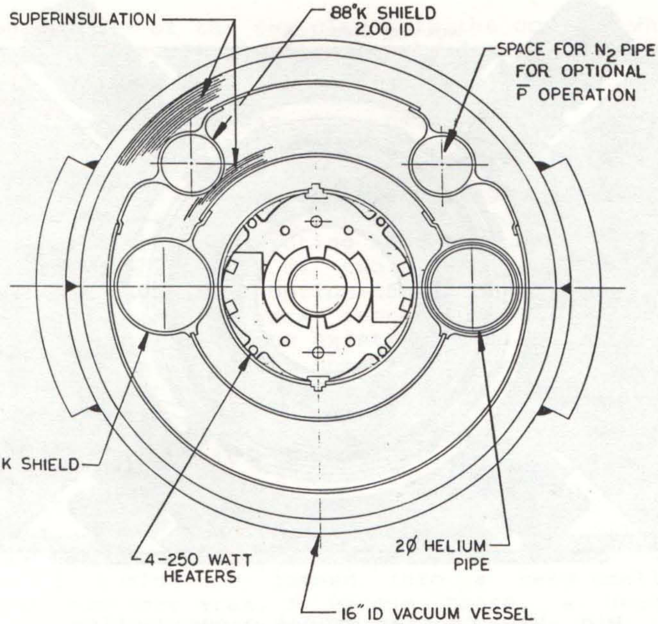

Fig. 12. SSC design B dipole magnet.

The units here are watts per meter. The design B magnet represents about an order of magnitude improvement in heat leak compared to the existing Doubler magnet system. Comparing the heat load of the SSC magnet to the existing transfer line shows that a heat load can be achieved that comes relatively close to the transfer line.

Figure 13 shows the cross section of the tunnel for design B. This is an 8-foot tunnel, a little bit smaller than described elsewhere in the Round Table proceedings. The two rings are one above the other because there is only a single magnet in a cryostat. This is in contrast to designs A and C where there are two magnets in a single cryostat.

SSC Cryogenic Design Goals

There are a number of design goals for the SSC. One is to eliminate the transfer line. That was set from the standpoint of cost as well as heat load. No suction and discharge headers are contemplated since the cost of these headers would be prohibitive. For example, for a central compressor station, a 4-foot diameter suction, header would be required. For one compressor station that would be absurd. There is a small lead flow collection header.

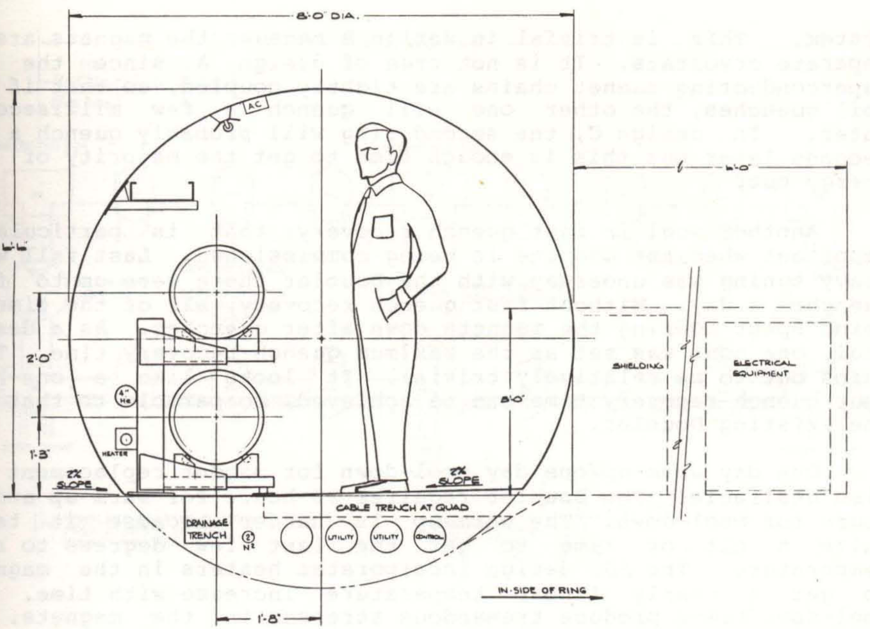

Fig. 13. Design B tunnel cross section.

Another design goal is to maximize the spacing between refrigerators. There are at least three reasons for this: 1) cost, since as the unit gets larger, the cost per kilowatt drops, 2) the fact that the larger the refrigerators, the more efficient they are, and 3) overall management simplicity. The SSC design calls for twelve units.

Another goal is to be able to operate with any refrigerator off. That's important because no matter how reliable the system design is and how much redundancy is put into it, sooner or later some maintenance will be required on a station. There may be contamination or a piece of equipment may die. The design incorporates the ability to shift loads to the adjacent refrigerators. This goes a step further than the Fermilab Doubler system. With the Fermilab design, it's possible to continue to operate with all the rotating machinery off in a given building. In the SSC it will be possible to take the cold box out of operation and keep the project running.

The next SSC cryogenic goal is quench independance of the ring. There are two rings of magnets in design B with many megajoules of stored energy in each one. If one of the rings quenchs, one doesn't want to double the amount of refrigerated energy needed by just dumping the good ring into the cryogenic system. This is trivial in design B because the magnets are in separate cryostats. It is not true of design A since the two superconducting magnet chains are tightly coupled, so that if one coil quenches, the other one will quench a few milliseconds later. In design C, the second ring will probably quench a few seconds later but this is enough time to get the majority of the energy out.

Another goal is fast quench recovery; that is particularly important when the machine is being commissioned. Last fall when heavy tuning was underway with the Doubler there were up to four quenches a day. Without fast quench recovery, all of the time is being spent cooling the magnets down after quenches. As a design goal, one hour was set as the maximum quench recovery time. This turns out to be relatively trivial. It looks like a one-half hour quench recovery time can be achieved, comparable to that for the existing Doubler.

One day warm-up/one day cool-down for magnet replacement is also desirable. The Doubler requires 36 hours for warm-up and 18 hours for cool-down. The warm-up is harder because it takes quite a bit of time to get the last few degrees to room temperature. The SSC design incorporates heaters in the magnets to get a nearly linear temperature increase with time. The cool-down times produce tremendous stresses in the magnets. A one-day cool-down time is really not one day for a magnet. An individual magnet may cool down in thirty minutes. It's much like an electrical pulse going down a transmission line. Doubler magnets have been cooled from room temperature down to 10°K in as fast as ten minutes. The magnet design has to be very carefully done so that it can take these stresses. It's not reasonable to think of a nice tapered wave where the refrigerator is programmed to have a slow cool-down lead edge. That would cost time in the cool-down cycle. Further, if there is a failure of the control system, the magnet system can be damaged. This means the magnets really should be designed for the maximum cool-down rate that they can take. The hardware is then designed to make sure that the rate can't be exceeded. This is done by sizing the valves and pipes so that the cool-down can't go faster.

SSC Cryogenic Design

Figure 14 shows the layout of a sector 1/12 of the ring. There is a refrigerator in the middle and 2-2/3 miles of magnets on either side. Each of these magnet strings in design B have isolation points spaced 2/3 of a mile apart (1 kilometer). With the isolation, if something needs to be repaired, it won't be necessary to warm up an entire building. Instead, a one kilometer space on one of the rings will be warmed up independent of the refrigerator while the other 15 magnet strings continue to operate. This makes it possible to get the one day warm-up and cool-down. In each of these one kilometer sections there are 10

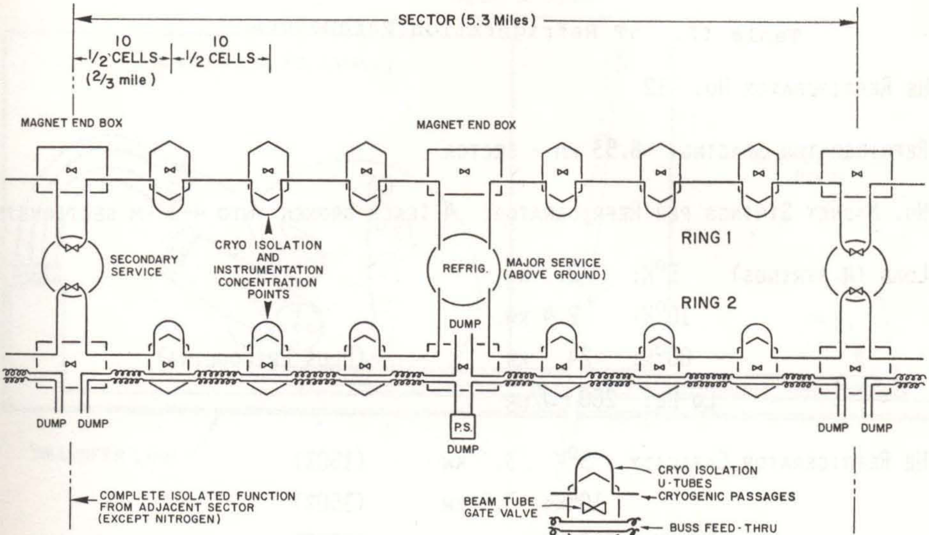

Fig. 14. Sector layout for SSC cryogenics.

half cells. A half cell is 350 feet long and consists of a focusing element and eight bending magnets. These bending magnets may be made longer so there might be only six bending magnets per half cell.

Figure 15 shows a a view of the tunnel cross section at the refrigerator. The tubes of the magnet strings are attached to the refrigerator with a series of bayonets. A bayonet is like a stinger for inserting a transfer line into a dewar. This system allows one to make and break cryogenic connections with no possibility of leaks from one circuit into the other. In this particular case, vacuum-jacketed flexible connections will be used to connect the refrigerator into the magnet strings. There will actually be four of these connections, two upstream and two downstream. At the one kilometer points, there will be vacuum-jacketed U tubes to provide positive isolation between the magnet strings.

Under normal circumstances, a refrigerator would be cooling four strings of magnets. If a refrigerator is off, the two refrigerators on either side would be asked to carry the heat load of five magnet strings, so they would be running at 125 per cent capacity. This means the valving must be available and that there must be very low pressure drops in the magnets to permit

Table II. 5T Refrigeration Parameters.

HE REFRIGERATOR NO: 12

REFRIGERATOR SPACING: 8.53 KM - SECTOR

NO. MAGNET STRINGS PER REFRIGERATOR: 4 (EACH BROKEN INTO 4-1 KM SECTIONS)

LOAD (4 STRINGS) 5°K: 2. KW
10°K: 2.4 KW
80°K: 24 KW (PLUS PRECOOLING)
LQ HE: 260 ℓ/HR

HE REFRIGERATOR CAPACITY 5°K 3. KW (150%)
10°K: 3.6 KW (150%)
LIQ HE: 520 ℓ/HR (200%)

HE COMPRESSOR PER REF NUMBER: 3 INDEPENDENT UNITS
SIZE: 75% FULL LOAD EACH
POWER: 1.3 MW EACH

AIR SEPARATION PLANTS NO: 4
SIZE: 86 TON/DAY (4000 ℓ/HR) EACH (200%)
POWER: 3.0 MW EACH

OPERATING POWER : 26.8 MW
INSTALLED COMPRESSOR POWER : 58.8 MW

STORAGE LIQ N_2: FOUR - 20,000 GAL DEWAR
EIGHT - 10,000 GAL DEWAR (2 DAYS)
GAS HE: TWELVE - 9 X 30,000 GAL "PROPANE" TANKS (15%)
LIQ HE: TWELVE - 40,000 ℓ DEWARS (33%)

Table II summarizes some of the refrigerator parameters.

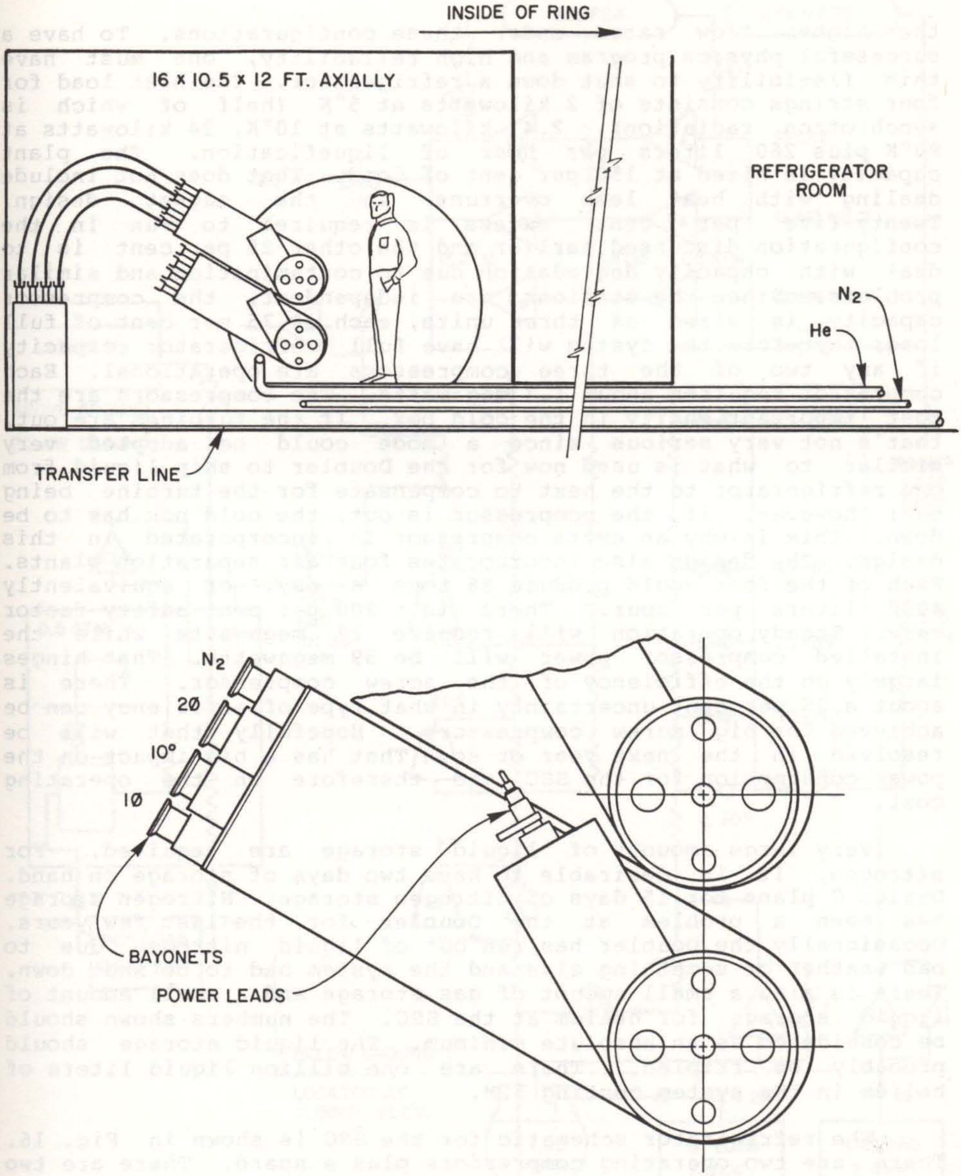

Fig. 15. SSC tunnel cross section at refrigerators. The lower portion shows a detail of the end box.

the higher flow rates under these configurations. To have a successful physics program and high reliability, one must have this flexibility to shut down a refrigerator. The heat load for four strings consists of 2 kilowatts at 5°K (half of which is synchrotron radiation), 2.4 kilowatts at 10°K, 24 kilowatts at 80°K plus 260 liters per hour of liquefication. The plant capacity is sized at 150 per cent of load. That does not include dealing with heat leak overruns in the magnet design. Twenty-five per cent excess is required to run in the configuration discussed earlier and the other 25 per cent is to deal with capacity degradation due to contamination and similar problems. Since the stations are independent, the compressor capacity is sized as three units, each at 75 per cent of full load; therefore the system will have full refrigerator capacity if any two of the three compressors are operational. Each compressor requires about 1.3 megawatts. The compressors are the most important entity in the cold box. If the turbines are out, that's not very serious since a mode could be adopted very similar to what is used now for the Doubler to ship liquid from one refrigerator to the next to compensate for the turbine being out; however, if the compressor is out, the cold box has to be down. This is why an extra compressor is incorporated in this design. The design also incorporates four air separation plants. Each of the four would produce 86 tons a day, or equivalently 4000 liters per hour. There is a 200 per cent safety factor here. Steady operation will require 27 megawatts while the installed compressor power will be 59 megawatts. That hinges largely on the efficiency of the screw compressor. There is about a 25 per cent uncertainty in what type of efficiency can be achieved for big screw compressors. Hopefully that will be resolved in the next year or so. That has a big impact on the power consumption for the SSC, and therefore on the operating cost.

Very large amounts of liquid storage are required. For nitrogen, it is desirable to have two days of storage on hand. Design C plans for 15 days of nitrogen storage. Nitrogen storage has been a problem at the Doubler for the last few years. Occasionally the Doubler has run out of liquid nitrogen due to bad weather or something else and the system had to be shut down. There is also a small amount of gas storage and a small amount of liquid storage for helium at the SSC. The numbers shown should be considered as an absolute minimum. The liquid storage should probably be tripled. There are one million liquid liters of helium in the system costing $2M.

The refrigerator schematic for the SSC is shown in Fig. 16. There are two operating compressors plus a spare. There are two dry turbines, a wet turbine, and two cold compressors. The liquid helium would go through a sub-cooling dewar. The major stream would go through the 5°K load with a small amount going through the 10°K load. There is also a sub-cooled liquid

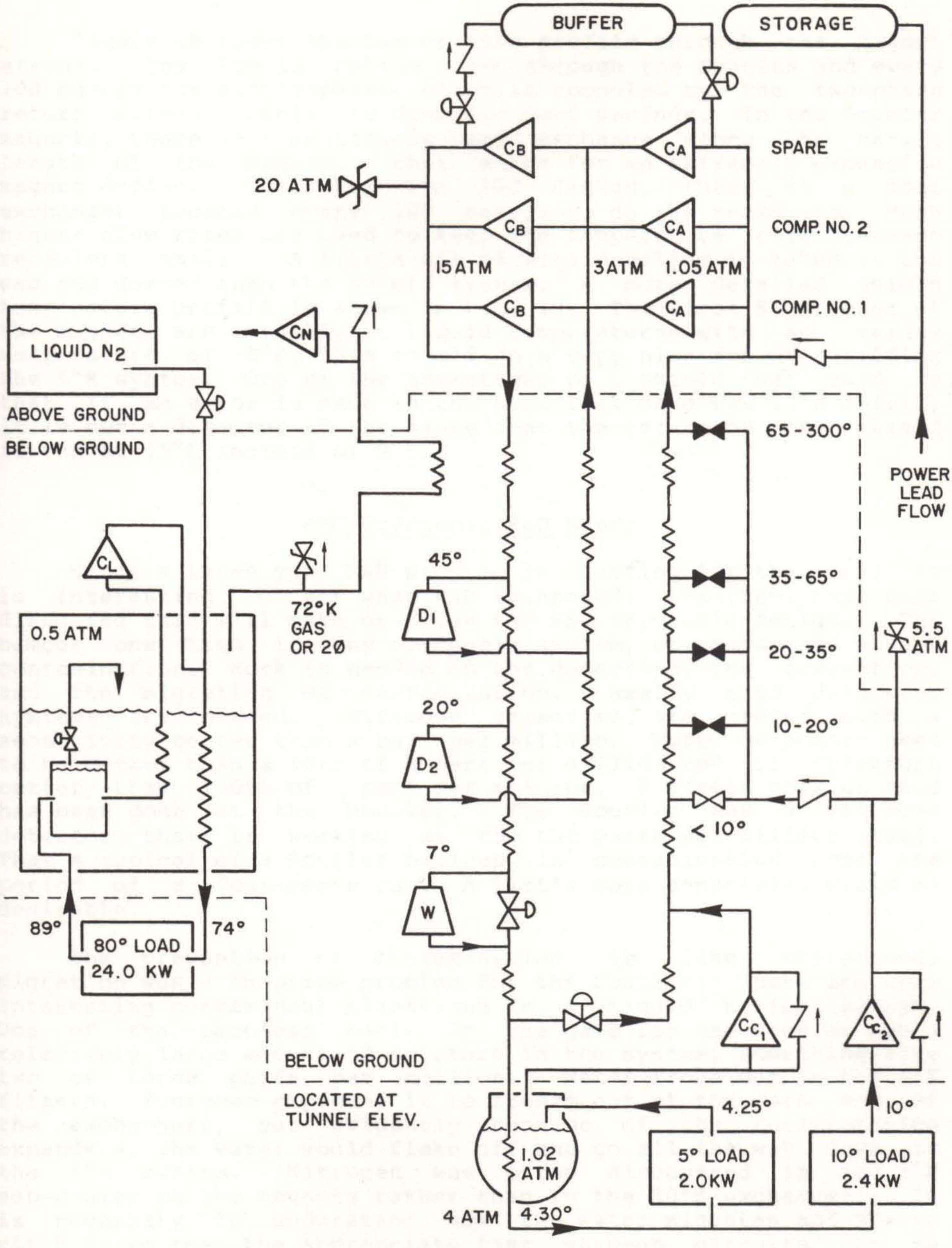

Fig. 16. SSC helium refrigerator schematic.

Figure 18 shows the temperature profile through the magnet string. The liquid helium comes through the magnets and every 100 meters the single-phase helium is recooled by the two-phase return stream. This is done for cost savings. In the Doubler magnets, there is a continuous heat exchange along the entire length of the magnet. That makes for an extremely expensive magnet design. Thus for the SSC design, there is a heat exchanger located every 100 meters to do the recooling. Much higher flow rates are used to keep the temperature rise between recoolers small. A little bit of liquid helium is taken at the end and dumped into the shield system. A more detailed shield temperature profile is shown in Fig. 19. The first 8 per cent of the magnets are actually at liquid temperatures with an average temperature of 8°K. This should do a very nice job of shielding the 5°K system. One of the advantages of a shield that cold is that if an error is made in the heat leak onto the 10°K shield, it is not a disaster in the sense that the return of the shield is run at 15°K instead of 9°K.

SSC Cryogenic R&D Needs

Since a three year R&D program is starting for the SSC, it is interesting to ask what R&D is needed. Six items have been discussed that will make or break the SSC cryogenic design. The number one item in any cryogenic system, of course, is always contamination. Work is needed on the detection, the prevention, and the migration of contamination. Really good detection systems are needed. Nitrogen detectors are needed with a sensitivity better than a part per million. Water detectors need to be better than a 10th of a part per million and oil detectors better than 100th of a part per million. A little work on this has been done at the Doubler. The Doubler has a nitrogen detector that is working at the two parts per million level. That's typical of a Doubler buildup in contamination over the period of a four-month run. A little more senstivity would be desirable.

The prevention of contamination is like motherhood. Migration was a surprise problem for the Doubler. There are some interesting contaminant migrations in a liquid helium system. One of the problems early in the game for the Doubler was a relatively large amount of moisture in the system, something like two or three parts per million. Water froze out in the 5°K filters. Everyone expected it to freeze out at the warm end of the exchangers, but evidently because of the reciprocating expanders, the water would flake off and go all the way down to the 5°K system. Nitrogen was often discovered in the 5°K sub-cooler on the magnets rather than in the 30°K exchanger. It is necessary to understand how the water migrates and how to catch it so that the appropriate fast warm-up circuits can be designed to remove it once it's known where it's going to end up. It usually ends up in the most inconvenient of spots according to Murphy's Law.

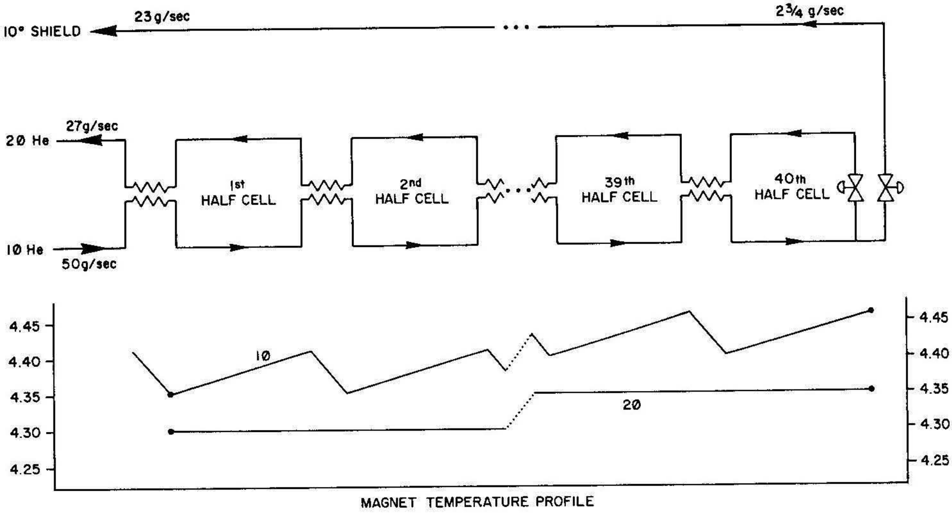

Fig. 18. SSC magnet helium flow schematic and magnet temperature profile.

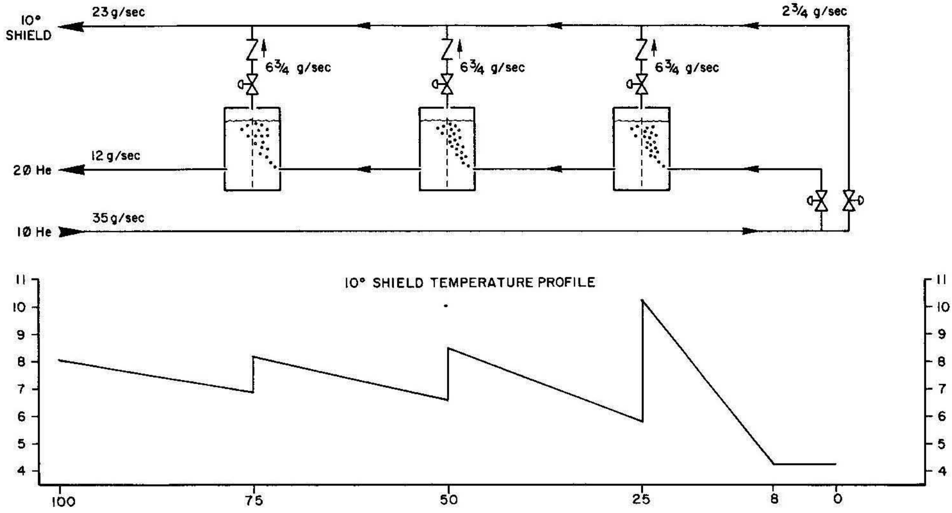

Fig. 19. 10° helium shield temperature profile and helium flow schematic for the SSC.

For design B, a two-phase liquid helium counter-flow is used and it's necessary to know what the stability of this is. This is particularly true if it's an inclined machine. Note that both the DESY and the LEP tunnels are being built on an incline of 0.6-0.7 per cent. It's necessary to know how the two-phase flow will be affected by that type of incline. There is really very little data on two-phase flow for liquid helium or even liquid nitrogen. Control ability and stability are important to have for a reliable system and they also relate to operating cost parameters. For an unstable refrigeration system one often just runs the compressor system harder.

There are three areas (all related) where work on reducing heat leaks is important: super insulation, supports, and heat intercepts. The most important elements of the support design are the heat intercepts on the supports. One designs supports and always assumes that the intercepts for those supports are at the temperature of the cryogens to which those supports are connected. But it's not very difficult to design an intercept on a support where the intercept is actually 30-40°K or even a 100°K above the cryogen temperature. This is an area that needs very close attention and even analytical measurement.

Screw compressor efficiency is important. There is a 25 per cent uncertainty in what efficiency can be achieved in the big compressors. Screw compressors are desirable because of the reliability and lower capital cost. The small Doubler screw compressors have half the efficiency of a very large reciprocating compressor. It's not understood why screw compressors are so inefficient. This is an area that can have a major impact on the operating cost of the machine.

How much helium would the SSC use? The Fermilab machine right now is using over a million dollars a year in helium. Clearly that is not a number that can be afforded on the scale of the SSC. Many man-months have been spent at the Doubler looking for leaks. In general, the leakage problem is not in the cryogenic valves and cryogenic components but in the room temperature valves and room temperature flanges. One of the biggest problems is the shaft seal on a valve that doesn't like the ambient temperature cycling. In one case, $50,000 worth of helium was lost in two nights. This happened in December when it went to -20°F and thirty leaks appeared at midnight. The next morning they sealed again. So, a significant amount of work has to be done in getting components that don't have these types of problems. One would like to use bellows valves for the majority of the equipment, but one needs a valve that can take 50,000 to 100,000 cycles on the bellows. A review was carried out of commercial valves at the start of the Doubler project (which is probably totally obsolete). This found ratings of 1000 cycles for valve bellows. A valve like this on a servo loop moving every ten seconds doesn't last very long.

Conclusion

The knowledge and experience to build the SSC now exists. The three year R&D phase will make it possible to evolve a more reliable design while at the same time decrease the cost.

CONVENTIONAL FACILITY REQUIREMENTS FOR THE SUPERCONDUCTING SUPER COLLIDER

Paul Gilbert
Parsons Brinckerhoff

Background

In February, Parsons-Brinckerhoff was charged with the mission of defining the conventional facilities requirements for the SSC reference design. This was to be done on a generic basis, that is, the location criteria was to be such that it could be assigned to any of a number of suitable sites throughout the United States. Certain anomalies must be faced from the outset when developing a physical plant on a non-physical site. Nevertheless, it was possible to develop a plan for the conventional facilities. This was achieved with a great deal of guidance and patience from Jim Sanford of Brookhaven and Tim Toohig of Fermilab.

A number of important questions were raised. What kind of criteria apply which may relate to the suitability of future site selections? What are some of the loads, some of the services, some of the types of construction, some of the types of materials, some of the labor force needs that must be determined?

The design was developed to a preconceptual level which means in effect, identifying one feasible and economic solution to the conventinal facilities problem in coordination with the technical facilities SSC working group. No significant attempt was made to refine the results, or to optimize them. To start quickly and to build on existing practices, Fermilab, SLAC in California, and Brookhaven on Long Island were used as laboratory models.

The Site

A median site model was developed to deal with the question of a non-physical site. By definition a median site is a location which would represent an attractive location for siting the SSC. It is not a difficult site, nor is it the most favorable site. Instead, it represents those sets of conditions which one would hope to find that constitute the broad middle range of suitability for the siting of an SSC machine.

The accelerator is assumed to be a gravitationally horizontal planar ring. While there are options for the machine to be able to somewhat follow terrain variation, such a machine was not a part of the initial mission. Environmental sensitivity was built into the design since, in all cases, such sensitivity is here today in every way. Another goal was to maximize the operating efficiency. These conditions have been integrated on a just pass basis, into the design that is discussed here.

The median site reflected conditions expected on an attractive site. With the type of capital investment that the facility represents, it would be foolhardy to set out to locate such a machine in an adverse location because of the enormous cost consequences that result. So a favorable site may be taken as synonymous with a realistic site.

The purpose of the preconceptual design is to have a basis to realistically define costs and schedule for the conventional facilities. For that reason a demographic character is ascribed to the site. That character involves such points as establishing the fact that there would be a major regional city within reasonable driving distance. The city would not only provide the labor force, housing, and an international airport but also cultural facilities, university facilities, and other types of research support. These amenities would help make for a suitable location for a laboratory with three thousand or so people involved in this type of intellectual and research activity. Proximity to public services and facilities was important. Electric power supply, natural gas supply, highways, railroads, and other public services had to be given some definition in the median design. Of course there are costs associated with getting the necessary services on to the site. This estimate could be extrapolated to a real site at some time in the future. A climate had to be assigned because climate greatly affects the operating cost and the type of facility required for the SSC. Climate, topography, and geology have a great deal to do with the cost of the facility. For example, a fairly flat site with a lot of rain will have drainage problems and that adds to the cost of operation, the cost of construction, and to the difficulty of construction as well. This meant there had to be a realistic combination of criteria between climate and certain other physical characteristics.

Dick Lundy noted that these facilities consume a considerable amount of power, perhaps 100 megawatts or more. The primary supply thus must not only be subtstantial but also should be backed up by a fairly stiff grid. The vast majority of the power used passes through the system and is then taken off and rejected through cooling towers or other means. Therefore, climate is reflected in the amount of effort and cost that has to be invested in heat rejection.

Next consider the physical characteristics that were assumed for the area. As was noted earlier there were three magnet designs, A, B, and C, that reflect in three different physical plant sizes. These would cover an area that may be thirty to fifty miles on a side. However, within that rectangle, the surface-located components of the Laboratory may only occupy seven hundred acres of land and the subsurface facilities may require another seven hundred to twelve hundred acres of subsurface easements. Construction easements in an area where a cut and cover tunnel was used would require an area approximate to one quarter of a mile wide, centered on the main ring alignment. This means a deep tunnel might be more appropriate in a populated area. Dick Lundy's figure showing the ring superimposed on Washington, D.C. illustrates the vast area that would be influenced by the siting of the facility.

Rights-of-way and real estate represent a major consideration for the SSC. There is a natural desire to have a flat site for the SSC. At the same time the median site couldn't be unrealistically flat because Mother Nature just didn't arrange land and climate that way. Some relief has been incorporated and the relief leads to drainage and to drainage ways. Vegetation preservation requirements as well as seismicity were considered. The median site was given a seismic characteristic of zone one which means that there could be some seismic activity but not sufficient to create a significant design consideration. In hindsight, the influence of a more restrictive seismic siting condition would have little influence on the facility cost because of the nature of the facilities themselves. A broad range of geologic types was incorporated in the median site. Air quality and solar access conditions were added since DOE facilities require that both active and passive solar energy conservation be a consideration. Finally, the Los Alamos Laboratory Site Atlas for the SSC was reviewed. They had assembled some six or seven preliminary proposals from different states that have come forward to indicate that they would be willing hosts for the SSC. That also provided a certain amount of basic information.

All of this led to the marvelous median site for the SSC shown in Fig. 1. Now the topography and the highways really don't exist anywhere except on the drawing and in the minds of the creator, but they illustrate the factors that are important in a reasonable site for the SSC. The sort of site illustrated is one where there is a crest in the topography with drainage coming away from the high ground. Notice that there is an interstate highway as well as a public service corridor. This corridor contains not only the interstate but also a 230 KV power line and a railroad right-of-way. Parenthetically, off the map where the branch line railroad joins the main line railroad there is a natural gas pipeline. Again, that is characteristic of what one finds today in public utilities. There is a state highway serving the north-south axis as well as a county road on the site. There are also various farm roads and other graded accesses that are not shown on the site plan but are in fact included in the utilization of this site when the machine is superimposed on it. The spirit is to not pave what one doesn't have to pave when the facility is put into place. Later it will turn out that the access roads that are put in are really minimal in nature and are only to get to those areas where there is a high service demand for access.

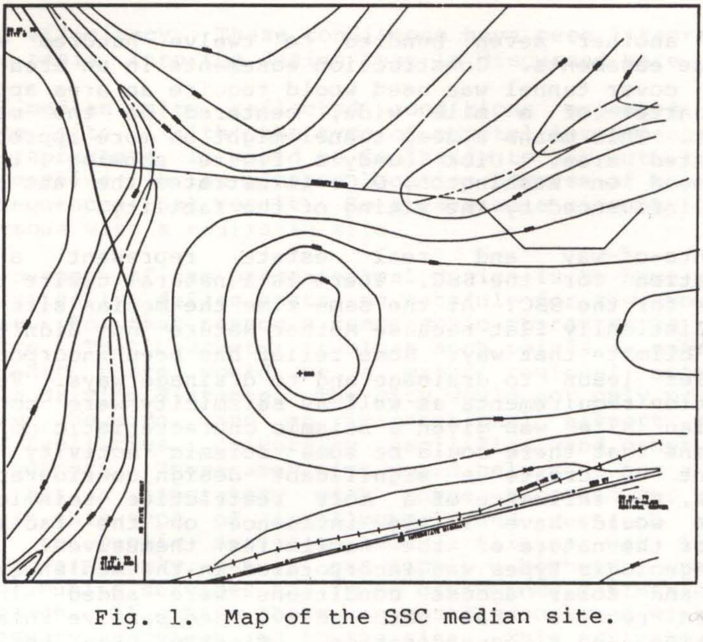

Fig. 1. Map of the SSC median site.

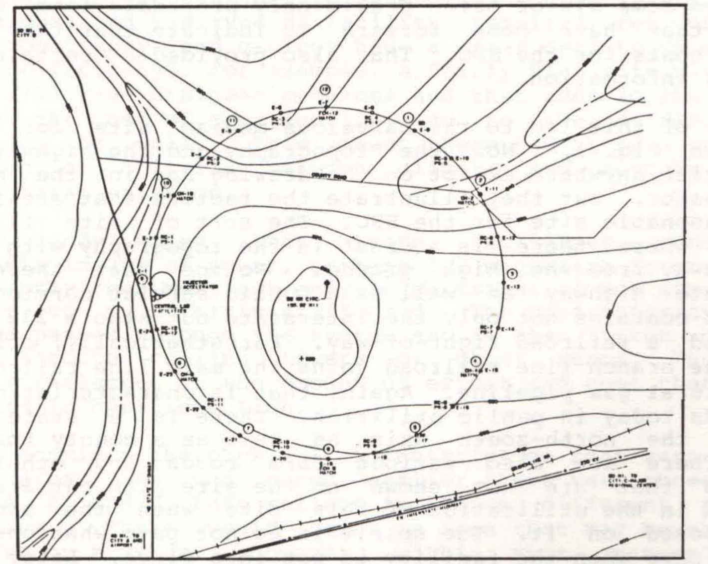

Fig. 2. The median site with the SSC imposed on it.

Looking at this, one asks the question, "How do we go about the matter of locating an SSC in this area?" A number of influences have to be considered. For scale the injector is essentially the same size as the Fermilab Doubler that was described this morning. The high-energy booster is about 20,000 feet around. Figure 2 shows the A design with the 6.5 tesla magnets. This results in a 90 kilometer circumference ring. Design B would have a 113 kilometer ring and the largest one would be 165 kilometers for design C. The injector and the central lab are at nine o'clock (on a clock oriented with twelve toward the north). The beam dump is off at one o'clock while the rf facility is at five o'clock. In addition to that, four developed experimental locations are shown, namely the collision halls at two, four, eight, and at ten o'clock and provision for future collision halls indicated at twelve and six o'clock. There are also refrigeration stations, power stations, accesses and exits, and a variety of other elements that are required to make the system whole. Figure 3 shows a site profile. It is important to realize that this is a much distorted scale profile, the total length on the horizontal is about 56 miles, while the total elevation difference in the vertical is about 180 feet.

For geologic groups, a sufficient amount of each major geologic group was identified to obtain a significant measure of what construction methods, costs, and time could be expected in sites that would be receptive to an SSC. Figure 4 characterizes the geology by sector. The north part of the site has both igneous and sedimentary rock varying in hardness from hard to soft. These could be either granites or basalts on the one hand or shales or limestones in the sedimentary group. At the deep point the tunnels are 100 feet down. There are problems with access shafts and construction logistics in deep tunnels. Typically the access shafts could be 32 feet in diameter and cost $3000 to $4000 a foot. A sufficient amount of such ground was incorporated to get a representative sample of the costs, the difficulties, and the time that it takes to construct in such ground. A collision hall facility was assigned to each one of the geologic groups to get a representative sample of what it would mean to have that type of major underground facility located in each of these groups. The east side of the site consisted of sand, silt, clays, and gravels. This is the type of alluvial out wash that is found in many parts of the country, for example, along the foothills of major mountain ranges where over time the material has been washed down from the mountains and settled out in the valleys. Gradations like this can give a variety of construction problems depending upon how they are mixed. The south end of the site consists of firm clay. This could be the type of material one would find in ancient lake bottom deposits. It can be excellent for construction but it can also have some problems. Finally, there is glacial till on the west side where the injector is. An excellent example can be found by going to the window at Fermilab and looking outside. Glacial till is the result of materials left behind by receding

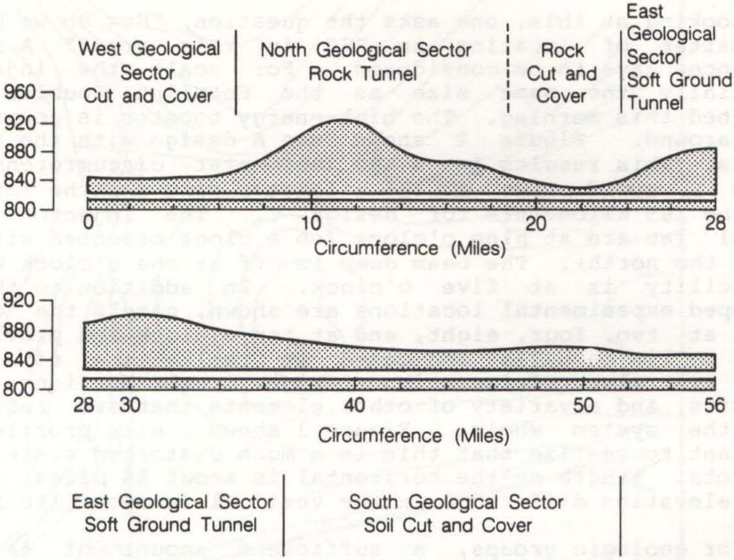

Fig. 3. Distorted profile of site to show angulation of the median site. The elevation difference is on the order of 100 feet which the total horizontal axis extends for 56 miles.

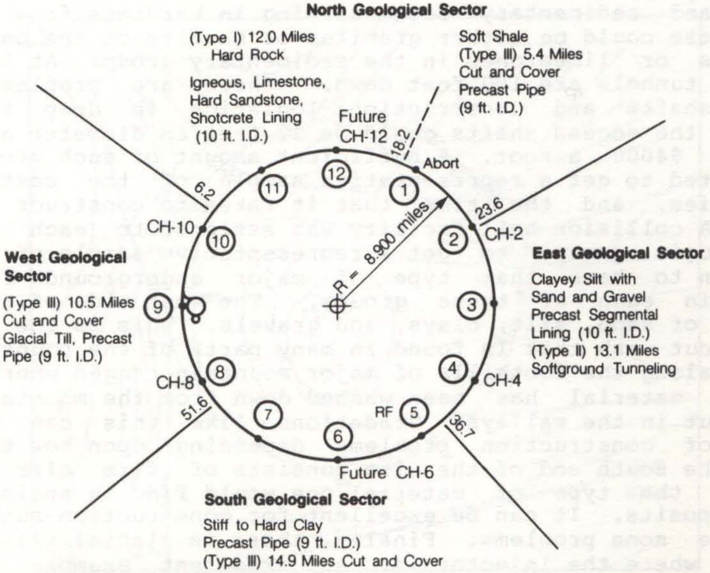

Fig. 4. Geologic character and main ring tunnel types.

glaciers and can be quite variable in conditions. Characteristically it has relatively flat topography and is relatively easy to excavate unless boulder fields, peat areas, or water pockets are encountered.

The Tunnels

To house the main ring, a tunnel must be provided with at least twenty feet of soil cover to provide for the necessary radiation protection. In addition to that, there is an additional radiation requirement for soil embankment to the outside of the ring where a muon shield should normally extend out about 270 feet. Thus, if a fill condition exists below the beam line, a fair amount of earth must be carried along. The most difficult feature to work with here is the scale. A ring 56 miles around has 300,000 feet of tunnels and a 165 kilometer or 103 mile ring has 550,000 feet of tunnels. This is a very major construction undertaking so that when one talks about moving a little bit of dirt per foot of tunnel, one is really talking about moving a lot of dirt overall. With this geologic and topographic mix, roughly 55 per cent of the site would be constructed by the cut and cover tunneling technique. The remaining 45 per cent would be evenly split between rock tunneling and soft ground tunneling. Figure 4 also summarizes the mix of the tunnels. The mix of tunneling methods doesn't have much influence on the cost. The sectors are long enough so they would individually constitute suitable construction projects. Typically the lengths of rock tunnels are such that a new tunnel boring machine would be fully amortized by the end of the construction. The cut and cover tunnel has a 9 foot inside diameter while the mined tunnel sections have a 10 foot inside diameter. When working on this project scale, particularly with underground work, the speed of construction becomes a major parameter because that is where the money is. The 10 foot mined tunnel size provides the driving tolerances to achieve near maximum driving rates. Essentially a 9 foot envelope is required within the tunnel to satisfy the technical facilities requirements to provide the continuity of the alignment space which is required for the beam. This 9 foot parameter is thus obtained at minimum cost by this approach. All of this tunneling information then leads to a construction program, a schedule, and a cost estimate.

The least costly technique for tunneling is cut and cover where, for this project, cut and cover would run $600 to $700 a foot for the ideal burial depths. The most expensive tunneling occurs at the interfaces between different geologic groups in the mined rock tunnels where costs could run to $1600/foot. The main ring housing averaged about $1050/foot without utilities. These estimates are in line with experience at PEP and also recent experience in Europe.

Figure 5 shows what the cross section might look like for the cut and cover case. There are several enclosure materials that could work but a precast pipe has been used here. It is estimated that 400 feet of pipe a day could be installed in a trench and back filled using a forward projection pipe laying technique. That's moving fast, but that's what can be done if a contractor sets up not just a pipeline operation but in fact, a production operation.

Observing the tunnel cross section, the primary device in the enclosure is the beam magnet package which is placed along the inside surface of the tunnel, to the right in this view. The injector, also in a tunnel, is located on the inside of the main ring with a magnet package mounted along the outside surface of the housing. If this were design B, there would be a double magnet package. The utilities are kept on one side. They include a helium gas line, water for heat rejection as well as fire protection, and 480 volt power inside the tunnel. The 480 volt power goes into mini power centers that are set out at 100 meter intervals. At those points it is broken down into 120 volts, a 408 V welding circuit and some other housekeeping circuits. The other side of the tunnel has communications and fire alarm circuits. This cross section provides for both a place for a person in a lay down space to work along the beam line, as well as a way of moving about by an electric-powered cart. Obviously the space inside must be used efficiently. The size of the interior space was determined both by considering the size that would be economic to construct as well as the size necessary on a functional basis to provide for what goes on inside the tunnel. Notice that there is sufficient room for two way traffic.

If a shotcrete lined rock tunnel was used it would not be a water tight tunnel, but it also would not be a wet one. The same general configuration would be used but a 10 foot ID tunnel would be used to provide the 9 foot clearance envelope. The soft ground tunnel would be very similar except the tunnel would be of precast concrete tunnel liners jacked into place and sealed without any further finish. Considering each one of these types of tunnels from the point of view of the supply of the various materials that will be required and doing some extrapolation of lengths and size, it is clear some very significant supply contracts will be required to make this project go. Figure 6 illustrates what the inside of the tunnel might look like when all facilities are in place. This is not too different from the photographs of the existing Fermilab facilities.

Other Underground Features

Around the main ring there's more than just a tunnel, even though the tunnel is the major cost feature. There are also twelve refrigeration compressor stations equally spaced around

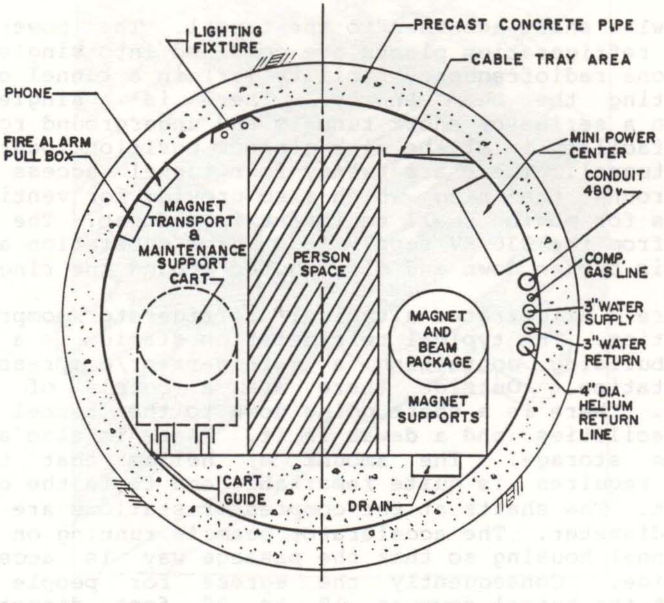

Fig. 5. Tunnel cross section for the cut and cover technique. The inner diameter is 9 feet.

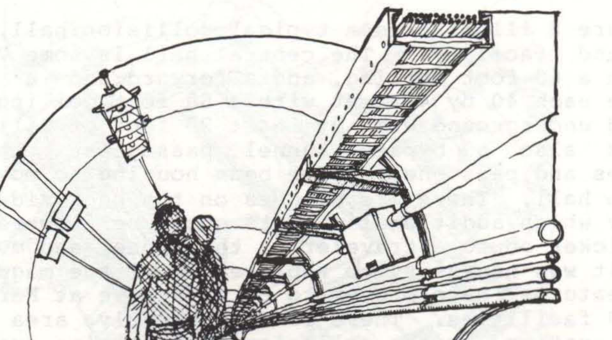

Fig. 6. Perspective view of the main ring tunnel.

the ring with shaft accesses to the tunnel. The power supplies and the refrigeration plants are combined into single stations. There is one radiofrequency facility (rf) in a tunnel enlargement for boosting the beam energy. There is a single beam dump located in a series of short tunnels and underground rooms. The injector facility is at the nine o'clock position located in five miles of tunnel. There are twenty-four tunnel access and exit points around the ring which also provide for ventilation and facilities for moving small equipment in and out. The power is provided from the 230 KV supply to a major substation at the site where it is broken down and distributed around the ring.

Figure 7 illustrates a typical refrigerator/compressor and power station. The typical refrigeration station is a 50 foot by 132 foot building containing a refrigerator/compressor and a power station. Outside there are a couple of additional buildings. There is a shaft going down to the tunnel that has helium facilities and a dewar in it. There is also a tank farm for helium storage. The amount of helium that this whole facility requires is quite remarkable and tests the capacity to produce it. The shafts at the compressor stations are 26 to 30 feet in diameter. The accelerator beam is running on the inside of the tunnel housing so that the passage way is accessible on the outside. Consequently the egress for people is on the outside of the tunnel down an 18 to 20 foot diameter shaft. Double 90 degree dog legs are provided for radiation protection and the cross section is adequate for ventilation of the tunnel.

Figure 8 illustrates a typical collision hall, a very major underground facility. The central hall is some 70-75 feet on a side with a 60 foot ceiling, and a forward and a backward area that are each 40 by 40 feet with a 50 foot ceiling. All of that is buried underground with at least 20 feet of fill on top of it. There is also a bypass tunnel passageway that allows the facilities and personnel in the beam housing to move around the collision hall. There are hatches on the back side of the bypass tunnel by which additional magnets could be lowered down by a cherry picker on to a traveler in the tunnel and moved around the ring if it was necessary to replace one of the magnets. Many of these features have proven to be effective at Fermilab with the B0 and D0 facilities. There is an assembly area at the same floor elevation as the collision hall to give space to lay down components in the course of putting together the experiment. An enormous door, some twenty feet thick, is provided to separate the two areas. At the next level there is a large assembly building for preparation. Access is through a floor opening to the assembly area below. All of this is served by a fifty ton crane. There are a series of offices and floors for the computing area, computer, other equipment, and personnel space on the outside bay of the building. The building is about 120 feet by 300 feet, a fairly large enclosure, but sized to the necessities of one of these research facilities. Recognize that it would probably be at least ten miles from where the

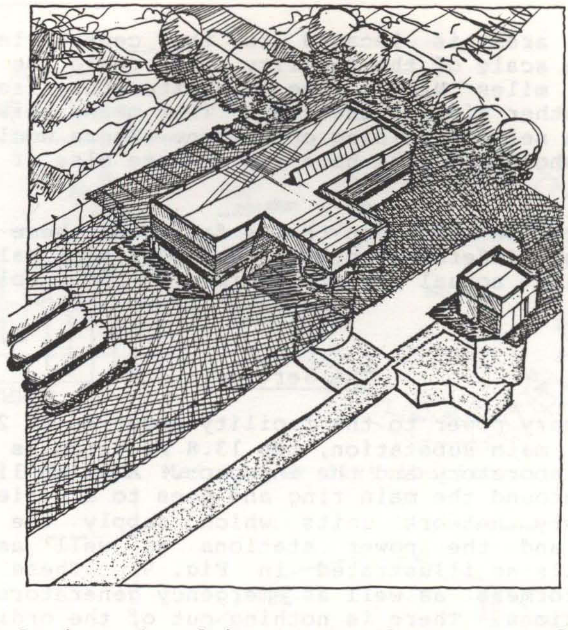

Fig. 7. Aerial view of refrigerator and power supply building.

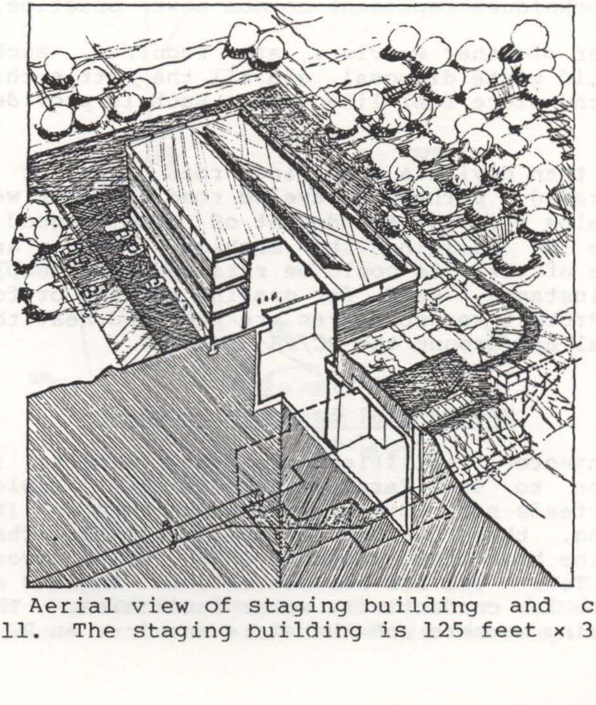

Fig. 8. Aerial view of staging building and cut away of collision hall. The staging building is 125 feet × 300 feet.

experimental area is located to the central laboratory as a result of the scale of the facility. It could be as much as thirty five miles back to the main laboratory from a collision hall on the other side of the ring. That means this has to be essentially a self-sustaining operation. There would probably be people here who would work here and nowhere else if they worked at the site.

Consideration was given to the fact that these buildings may be buried at different depths. Each collision hall was treated in terms of its actual relationship to the geology and the topography.

The Services

The primary power to the facility comes in at 230 KV. From the primary main substation, two 13.8 KV circuits provide power to the main laboratory and the injector. A 69 KV line comes off which loops around the main ring and goes to a series of what are called primary network units which supply the refrigerator compressors and the power stations as well as feeding the collision halls as illustrated in Fig. 9. These contain the power transformers as well as emergency generators for each one of these stations. There is nothing out of the ordinary in any of these systems. The circuit breakers, the switches, and the isolation techniques represent common power practice.

A number of other services are required, such as water, sewage, solid waste disposal, and all the little things that add up to make the place feel like home. Each is provided for in the plan.

Rather than devise a main ring road type of design, the existing graded, perhaps graveled road services were utilized. This also helps from the standpoint of being as good a neighbor as possible by building the machine into an area where the original use of the site could be retained and employed if it was, for instance, used for grazing animals or for other farm purposes. Protective enclosures are provided near the entrances of the tunnel and around the RC/PS areas.

The Central Laboratory Complex

The injector facilities are very similar in size and configuration to the Fermilab accelerator complex. A linear accelerator feeds beam into a low energy booster (LEB) with a "small" ring, three quarters of a mile around. That feeds beam in turn to the high energy booster (HEB) which accelerates the beam to 1 TeV. The LEB has conventional magnets and therefore requires some low conductivity water facilities. The HEB uses superconducting magnets and therefore requires an RC/PS facility.

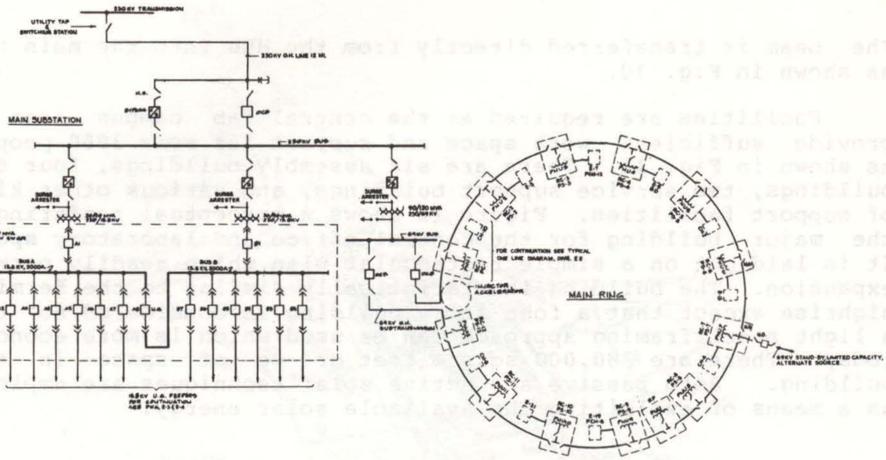

Fig. 9. Main power distribution.

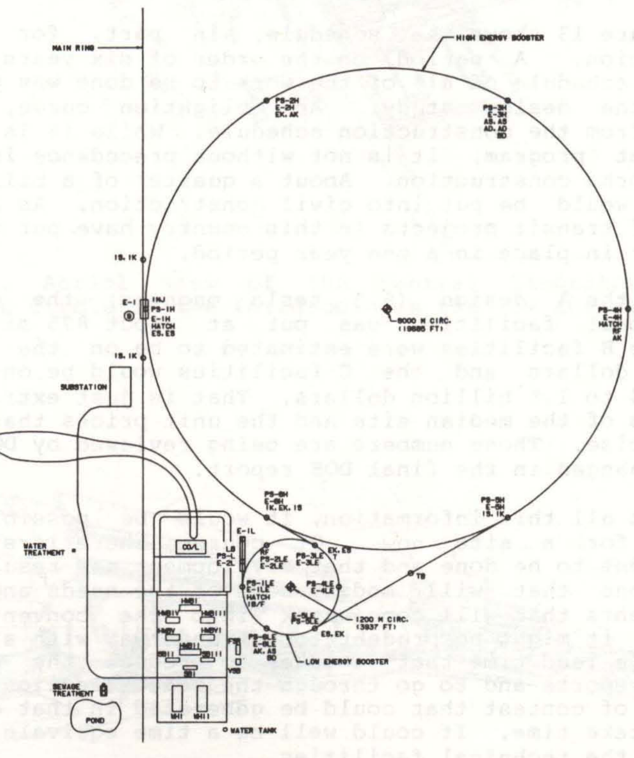

Fig. 10. Central laboratory and injector facilities. The high energy booster is about 6,000 feet in diameter.

The beam is transferred directly from the HEB into the main ring as shown in Fig. 10.

Facilities are required at the central lab campus area to provide sufficient work space and support for some 3000 people, as shown in Fig. 11. There are six assembly buildings, four shop buildings, two service support buildings, and various other kinds of support facilities. Figure 12 shows a conceptual rendering of the major building for the central office and laboratory space. It is laid out on a simple rectangular plan which readily permits expansion. The building is statistically similar to the Fermilab highrise except that a four story building is considered so that a light steel framing approach can be used which is more economic today. There are 380,000 square feet or so of space in this building. Both passive and active solar techniques are employed as a means of exploiting the available solar energy.

The Schedule

Figure 13 shows the schedule, in part, for conventional construction. A period on the order of six years is shown. A detailed schedule of all of the work to be done was put together during the design study. An obligation curve, Fig. 14, was devised from the construction schedule. While it is an ambitious investment program, it is not without precedence in the area of public works construction. About a quarter of a billion dollars a year would be put into civil construction. As an example, a number of transit projects in this country have put that much or more work in place in a one year period.

For the A design (6.5 tesla magnets) the cost of the conventional facilities was put at about 875 million dollars while the B facilities were estimated to be on the order of 1 billion dollars and the C facilities would be on the order of about 1.4 to 1.5 billion dollars. That is just extrapolating on the basis of the median site and the unit prices that came out of the exercise. Those numbers are being reviewed by DOE and they may be changed in the final DOE report.

With all this information, it would be possible to start looking for a site now. Of course, there is still a lot of development to be done and that development may result in further definitions that will modify some of the needs and some of the requirements that will come back into the conventional side. However, it might be prudent to get underway with site selection soon. The lead time that it takes to produce the environmental impact reports and to go through the site selection process with the kind of contest that could be generated in that decision will simply take time. It could well be a time equivalent to the R&D time for the technical facilities.

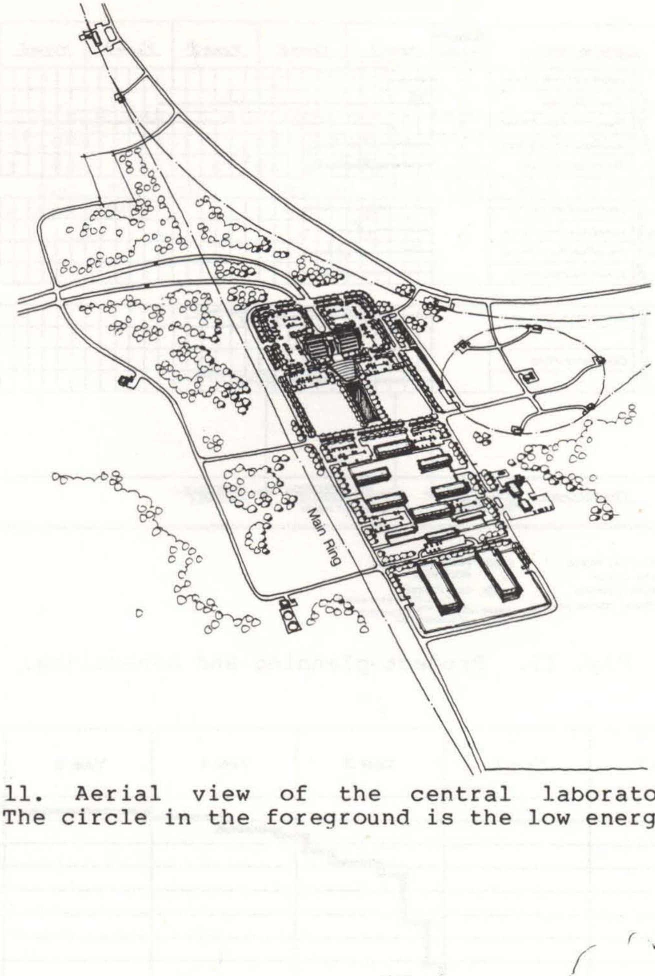

Fig. 11. Aerial view of the central laboratory campus complex. The circle in the foreground is the low energy booster.

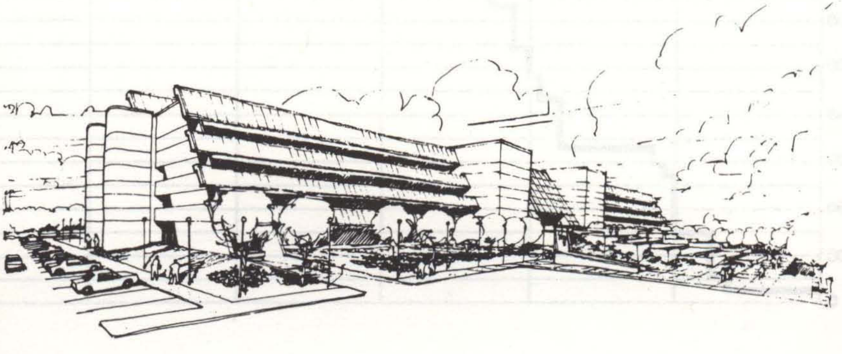

Fig. 12. South facade of the central laboratory.

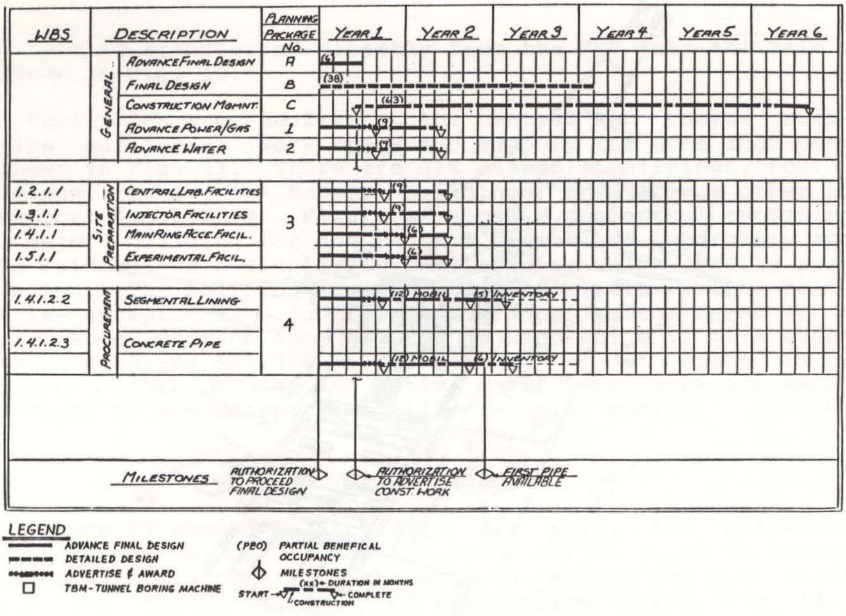

Fig. 13. Project planning and scheduling.

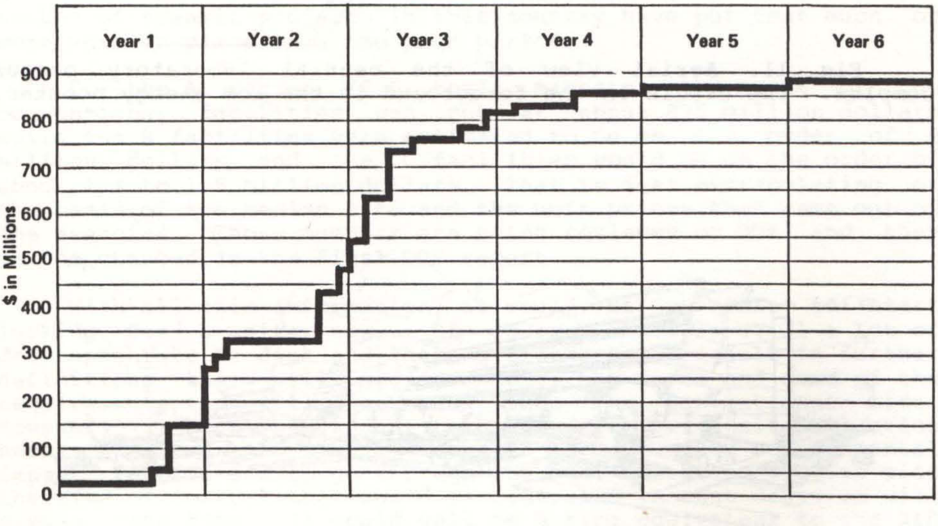

Fig. 14. Obligation curve for the SSC project.

Conclusion

The SSC is a feasible project and its future is possible. Its conventional facilities can be designed and constructed efficiently and cost effectively today, using today's men, materials, and methods. Let us hope that the priorities to be set include the SSC.

COMPUTER CONTROLS AT THE SUPER ACCELERATOR

Dixon Bogert
Fermi National Accelerator Laboratory

Since the SSC controls may have many similarities to those of the Fermilab Doubler, it is useful to review the Doubler controls. On the other hand, there are some significant differences that are worth a little bit of attention and there may even be some areas that will attract some interested industrial participation.

Although the circumference of the tunnel obviously scales linearly for an accelerator that is twenty times more energetic but using the same magnet technology as the Doubler, there are some aspects of the control system which it is now believed will not scale linearly. In fact, there are some instances where fewer devices will be needed, and others where the number of devices will be relatively the same as for the Doubler. For example, for the SSC refrigeration, the plan is to have only twelve very large distributed refrigerator systems rather than scaling the design of the relatively smaller Doubler refrigerator system. The Doubler refrigerator system already includes at least twenty-four distributed refrigerator engines in individual buildings plus about eight compressor buildings leading to a net count that is already in the thirties. Therefore, this is an example of an SSC system projected to have fewer distributed components than presently in the Doubler. An example of a system whose components will remain rather similar in number as in the Doubler is the correction magnet system. There are "cells" in both the Main Ring and the Doubler that consist of four dipoles, a focusing quadrupole, and then a correction coil package. For the SSC, the number of "cells" as defined by the number of dipole correction elements does not scale by a factor of twenty, but is similar to the current Doubler number. Of course, the number of main dipoles per cell will increase, unless very long dipoles are built. Therefore, the number of correction function generators required will remain at the level of a few hundred, rather than maybe five thousand.

One of the necessary features of the superconducting accelerator that exists at Fermilab is an emphasis on distributed processing. The need for distributed processing was driven by considerations which were only beginning to become important for an accelerator of the physical dimensions of the Doubler. These were especially important for the quench protection and refrigeration systems. For a twenty mile diameter machine there are many more instances where a distribution of controls becomes of considerable importance. A quick example is illuminating. The time of flight of a proton around the Doubler is on the order of twenty microseconds. Since the proton is essentially moving at the speed of light, that is also the time of propagation of an electronic signal around the accelerator. Now if someone at a central point wishes to request some piece of information, that implies some sort of round-trip time. (An outbound request for information and an inbound return of data.) This will be true unless a decision was made in advance that the information would always be wanted, so that arrangements were already made to have the data flowing inbound from its source without individual requests. If there is only one serial communication systems (as in the Doubler), it is only possible to make some 50,000 four-mile round-trip communications per second unless those communications are already very complex or, as noted, one-way communications have been previously established and the data is constantly flowing. If the dimensions of the accelerator are expanded by a factor of twenty, then the time of flight of protons around the accelerator is on the order of 400 microseconds and the number of electronic round trips has been correspondingly reduced. The number of requests for information on demand which can be accommodated has been reduced to about 2500 per second. That is a very low number and it absolutely implies that there will be a fair amount of distributed local control. Even intermediate control decisions are going to have to be made in a distributed fashion around the SSC ring.

The completion date for the SSC was given as possibly 1994. It is now about halfway between 1994 and 1974 when some of the early ideas were first considered for the control system for the superconducting Doubler. The differences in fundamental electronic technology which have developed between 1974 and 1984 are probably substantially greater than what one can expect to see in the next decade. Although the rate of development of electronic capabilities will probably continue to accelerate, many of the basic devices that one might need for a large, distributed control system ten years from now are available at present. Undoubtedly these devices will be subject to considerable improvement by 1994. But in 1974 they were not available. The first commercial microprocessors were announced in 1975. The planning documents from 1974 for the Doubler are quite interesting because those involved did not consider the distribution of computer control at all in the way it is now used. This is an example of a "fundamental change" that one does not necessarily expect to see every decade.

The Doubler Control System

As seen in Fig. 1, the Doubler has a basic centralized computer system which consists of two Digital Equipment Corporation VAX 11/780's and about twenty-one DEC PDP-11's which are networked together using DEC-PCL (parallel communications link) hardware. There are undoubtedly aspects of this particular choice of hardware which can be improved. Independent of the location of the SSC these improvements may be carried out at Fermilab during the course of proton-antiproton collider experimentation. Fourteen of the PDP-11/34's shown in Fig. 1 support consoles for operators. Each console is identical. The computers support the consoles on a one-on-one basis. This has proven to be a very useful and friendly sort of organization. It has been accepted as the type of organization that one would continue to support for the SSC. However, the use of something on the order of a DEC microVAX, or the equivalent from another vendor, but something which has considerably more computational power than a PDP-11/34 is probably indicated. The PDP-11/34's drive some hardware (not shown in Fig. 1, but indicated in Fig. 2) which permits serial communication to the actual console location, possibly over some distance.

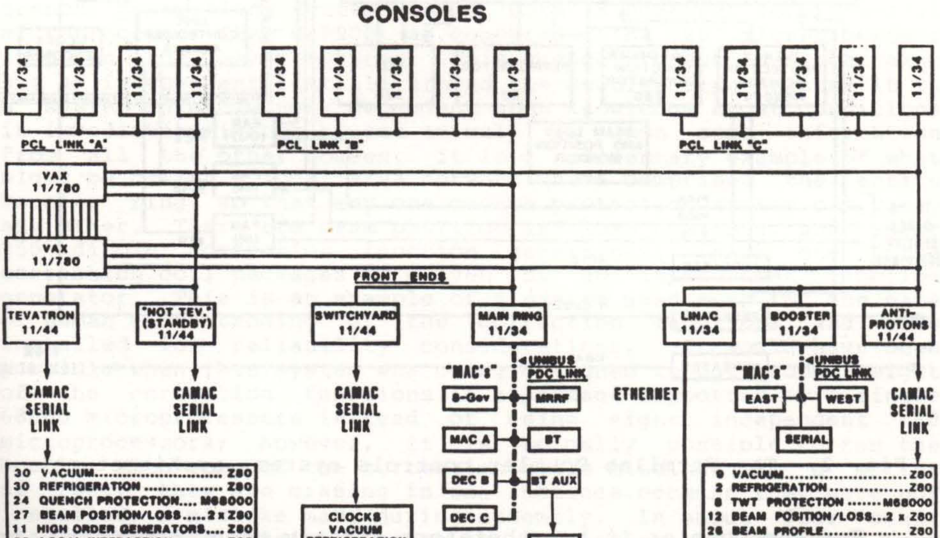

Fig. 1. The Fermilab accelerator controls system "ACNET."

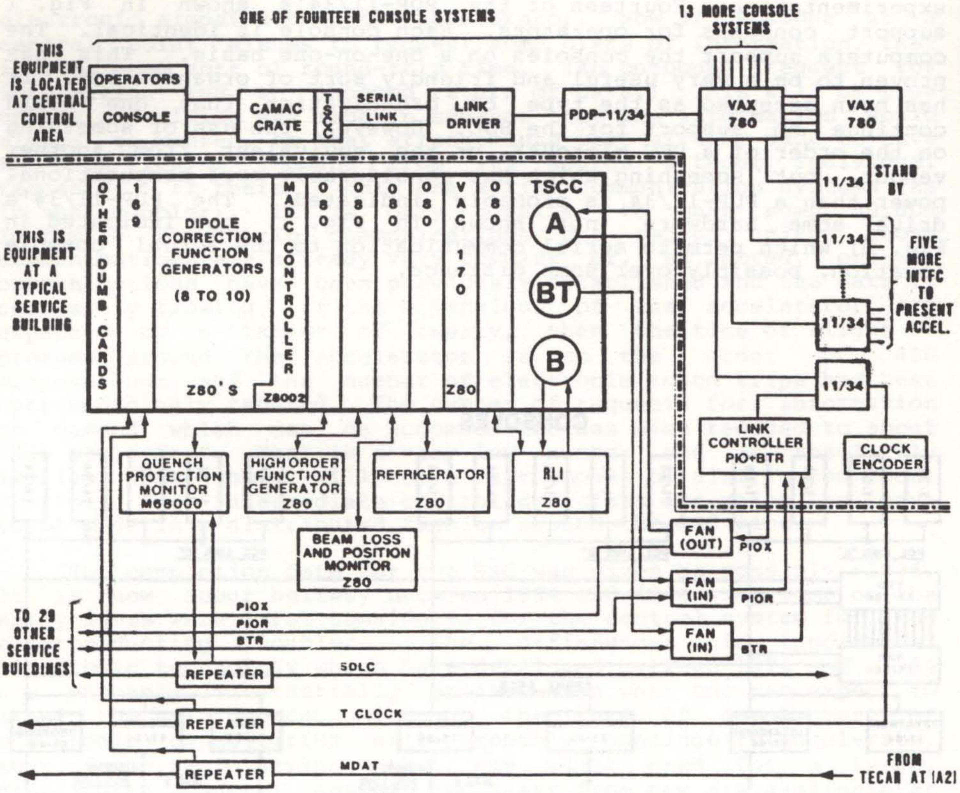

Fig. 2. The Fermilab Doubler controls systems architecture.

For the Doubler it has therefore already been necessary to address the question of console control of the accelerator at distances of several miles from the central control location. The present approach outlined above permits serial communications between the console equipment and its centrally located PDP-11 computer. In the future, this is almost certainly not a reasonable communications design when the distances are not two miles but rather twenty or forty miles. Rather, a greater utilization of long distance computer networking, with the console computers distributed to remote sites in addition to the console equipment, will almost surely be required. This will permit faster local access to data and also will greatly reduce the number of serial communication lines.

In many respects the Fermilab accelerator complex is similar to what is labeled the SSC injector; however, if an SSC is built from scratch in the four corners country of Colorado and New Mexico, the builders will not be faced with some of the historical imperatives that drove the controls group at Fermilab to do some very strange things when preparing to control the Doubler. Once some equipment is in place, it is seldom possible to replace it all at one time. There were some instances in the older systems of the conventional accelerator where very substantial changes were made in order to get a more unified control system going for the Doubler.

There are about 700 microcomputers distributed around the Doubler and Antiproton Source. They generally come in multiples of the service buildings around the rings. There are thirty vacuum controllers, thirty refrigeration systems (each one of which is driving about eleven closed loops), and there is a quench protection system which uses a somewhat more powerful microprocessor (the M68000) as compared tothe Z80 microprocessor used in the other systems. This quench protection system also has an independent link all around the Doubler ring because it is a system which is very dependent upon knowledge about conditions in immediate neighbor houses as well as upon some information from all the other houses. It is a rudimentary example of what might be called a local area network which comprises the entire Doubler ring so that any one quench protection system can query any other. There are beam position and loss monitors, and about 250 dipole correction function generators. Each one of the correction coil packages is driven by an independent function generator. This is an example of a system used both for the ease of human understanding of the correction functions and also installed for reliability considerations. It would have been possible when this system was being designed to put about eight of the correction functions in a package supported by a single 68000 microprocessors instead of using eight independent Z80 microprocessors; however, it is actually possible to run the Doubler with some of the correction coils missing. As a matter of fact, they are missing in one instance because they were not installed, a mistake made during assembly. In about four other places they are missing because the correction coil packages have proven to be one of the physically weak links in the actual superconducting accelerator and four have been destructively damaged. However, if one cannot make a correction at one place, it is possible to rework the the entire ring-wide correction function as distributed around the accelerator to take care of the problem. The extreme independence of this system makes it possible to either lose a single microcomputer (which seldom has happened) or to delete a physical correction element fairly directly. On the other hand, the loss of eight neighboring correction elements, due to the failure of a multiplexed system, would be more serious.

There are several examples of local area networks in the Fermilab accelerator complex. The quench protection system has already been mentioned. The system used in the re-instrumentation of the injection linac is another. In general, a local area network uses a straightforward technology like Ethernet, or a token passing Ethernet, or more generally, any number of computers requesting information from each other in an arbitrary fashion using any networking protocol which is independent of the details of the computers involved. The linac system used an SDLC protocol which is an IBM pre-Ethernet system characterized by a circular serial transmission of messages. A small segment of Ethernet is actually used to connect the linac system to the central host network.

The SSC Control System

As noted earlier, the physical size (rather than the complexity) of the SSC probably indicates that there will be considerably greater utilization of local area networks than was the case at Fermilab to date.

There are other features of an SSC characterized by only twelve major refrigeration centers and access points which are different when compared to the Doubler. At the Doubler, there are 24 service buildings servicing the six 60 degree arcs of the accelerator and six service buildings that are controlling the six straight sections for a total of about thirty service buildings around the four mile ring. At the moment, all of the major electronics is upstairs in the service buildings. In other words, the electronics is out of the tunnel where it is accessible. At the SSC there will be several kilometers between access points and therefore a considerable distance between an access point and an arbitrary controllable device. In such a configuration, the cabling cost quickly becomes rather expensive if all signals were to be brought back to the twelve access points. As a result, people have decided to study the question of actually distributing the electronics in alcoves in the tunnel of the SSC. This has several implications. The tunnel environment will probably be more hostile than the environment found in the Fermilab tunnels. The SSC tunnels are likely to be somewhat damper (since there will be no conventionally powered iron magnets to warm the air) and, of course, it will not be possible to get at the electronics to service it during acceleration or experimentation. Reliability, redundance, and backup will become important considerations. These are all related to the question of "What happens when something doesn't work right?" An important input for these considerations will be dependent upon guesses as to how the SSC will be operated. The operating scenarios will have to come from the accelerator system designers, the accelerator theorists, and the experimental physicists. One scenario might be that a proton-proton or a proton-antiproton fill would be done once per day and that experimenters would be reasonably happy if they got twenty hours of colliding interactions and then had from two to four hours of access to the tunnel for various purposes. At the Doubler, access to the tunnel is not a fast thing to arrange. One basic problem is the exposed electrical buswork in the tunnel. At present, it is a fairly lengthy procedure simply to get the breakers undone in order to allow general access into the Doubler tunnel. Another problem is the oxygen deficiency hazard. This sort of detail must be considered if the scenario includes a plan to service electronics inside the tunnel on a daily basis.

To summarize, there probably will be electronics in the tunnel of the SSC, there probably will be an emphasis on local area networks, and there probably will be an increasing emphasis on reliability. Note that low voltage digital and analog controls are by no means the whole story when considering accelerator reliability. What fails in general is not microcomputers and integrated circuit chips. One recent microcomputer "failure" at the Doubler occurred when the service building roof sprang a leak and a rainstorm drenched the microcomputer. This is an example of the "real" problems a person has to include in one's thoughts.

The SSC accelerator itself will not work if a single major bending dipole is "missing" (out of the circuit). A relatively small loss of bending angle at the one point where a dipole fails will result in substantially less than one centimeter of orbit displacement at the point of failure. However, a "bump" is thereby put into the orbit that results in a ten centimeter displacement at some other point around the accelerator. That would be well outside the aperture of the proposed magnet system. This means that there are certain situations where it is not possible to protect against catastrophic failures.

During the early stages of the SSC preliminary design, some people thought that they might be able to build a truly "passive" quench protection system into the magnets themselves. The SSC reference design does not make that assumption. Magnet builders also feel that it is probably not possible to do so. This means that a very "active" electronic quench protection system for magnet protection will continue to be needed. This was, one recalls, one of the important arguments for the distributed processing utilized in the present Doubler system. At the SSC, the same emphasis only becomes greater. There simply is not time to collect all of the necessary information for evaluating a magnet's superconducting status, and then deliver it to, and process it at, a central location. Typically, a decision as to whether a magnet has gone "normal" has to be made in one or two 60Hz line cycles. At that point, something must be done or the physical integrity of the magnet is at risk. The quench protection system at the SSC will face timing constraints identical to those at the Doubler.

This means that it will be necessary to have some relatively high powered processing capabilities in the distributed locations at the SSC. Some of the SSC reference designers have argued that it may be possible to combine many of the functions, which are separated into individual microcomputers at the Doubler, into larger microVAX or equivalent systems. This is certainly possible but whether it is completely desirable requires some substantial program or failure analysis to evaluate. The SSC reference design, for example, suggests that the same rather high powered miniVAX that would do quench protection of a section of arc would also act as a communications node at a major access point in the network. This could place conflicting demands for processor cycles in a fashion that would be irritating to one or the other of the processes, say quench protection or communication.

One of the features that is very nice for Doubler accelerator operators is the ability to connect a "trackball" (similar to a computer mouse but a little different) directly into the system in real time. This permits one to adjust the numerical value of a variable, have the altered value be sent to the hardware, and then have a new reading of the variable be returned from the hardware at 15Hz. This is fast enough so that from the human (physiological) point of view, there appears to be a real time physical connection. Earlier in our discussion it was noted that the total number of round-trip communications between a remote location at the SSC and some central facility will be reduced to only 2500 or so per second. If, in addition, there were to be too many "layers" of local networks and interfacing computers between the local area networks, there could begin to be some difficulties in trying to pass information along in real time in order to provide some type of response for humans approximating 15Hz.

For the Doubler, the data base for the entire system is centralized in the "Operational VAX." There were some strong arguments in favor of this approach from the people that proposed this, advocated it, and implemented it. With a widely distributed system as at the SSC, it may be desirable to turn on the refrigeration systems, for example, as arcs of the SSC are completed. If the complete central control system is not done at that time, a greater distribution of the data base may be desirable compared to what has been done for the Doubler. However, there are certain problems that arise which are among the reasons that one did not choose to distribute the data base in the Doubler system. The primary information in the data base at the SSC would still be addressing information. In other words, the data base contains the necessary information so that if somebody at any console wants to get at a particular piece of information, one picks up a road map and this map tells one how to get to the information. The road map is handed off further down the line to all other computers involved until one reaches the computer with direct access to the information requested.

Now, one way to avoid the necessity for some of the complexity of the road mapping scheme is to fold into the device-naming-architecture a great deal more directive information than is implicit in a device name at present in the Doubler. There are many people who advocate the practice of a "meaningful mapping" of all device names, and there are some very strong reasons for wanting to do it, avoiding some of the reliance on a centralized mapping data base.

The next three figures illustrate the system proposed in the reference design for the SSC. It shows a system of super-mini's (for example VAX 11/780's or 785's), large mass storage, and twelve operator consoles. The super-mini's do not physically drive the consoles; the consoles are perhaps driven one on one by computers each equivalent to a microVAX. The super-mini's are networking systems and network switch controllers, with perhaps the additional job of providing a redundant system for verification. The SSC in the reference design uses a standard long haul network to distribute local control around the 100 kilometer ring. Figure 3 illustrates a node on a local area network for one of the twelve sectors of the ring, or perhaps for the injector complex. Figure 4 illustrates the proposed architecture. The long haul network is shown, as well as an array of mini computers and some subsidiary local area networks. Microprocessors are located inside the tunnel. Each microprocessor is illustrated managing up to ten "half cells" of magnets. There are switches to permit some redundant paths for communication. Figure 5 shows the architecture at the cell level. This is shown including a 16 bit microprocessor that is managing a number of subsidiary modules in an interface crate. This is the unit that would be repeated most frequently in the tunnel.

The reference design proposal is certainly an example of a system that could be built today. From the technology point of view, there is no overwhelming difficulty with the design. It would probably produce most of the features that have been found necessary in the operation of the Doubler to date, with the possible exception that there might be a little difficulty in making a real time connection between an operator's control device and a piece of accelerator equipment with a 15Hz response.

At the SPS at CERN they do not have a 15Hz operator-to-device connection capability. They do not try to give operators a real time feel of control over particular devices. People who have played with both systems, however, feel that the Doubler system has something to be said for it.

Notice that most of the subsystems mentioned as part of the Doubler are shown in the SSC reference design figures. These include beam position monitors, beam loss monitors, wave form generators, and voltage monitors. The design report does not discuss whether these should involve subsidiary microcomputers or

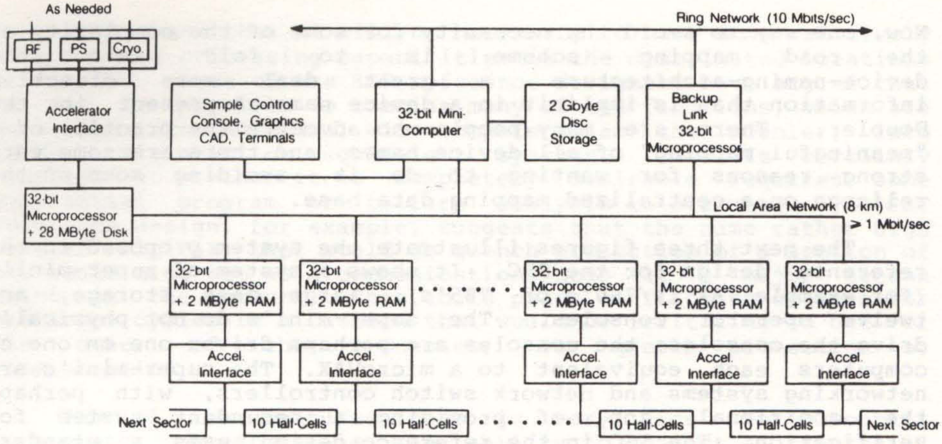

Fig. 3. SSC sector control system block diagram.

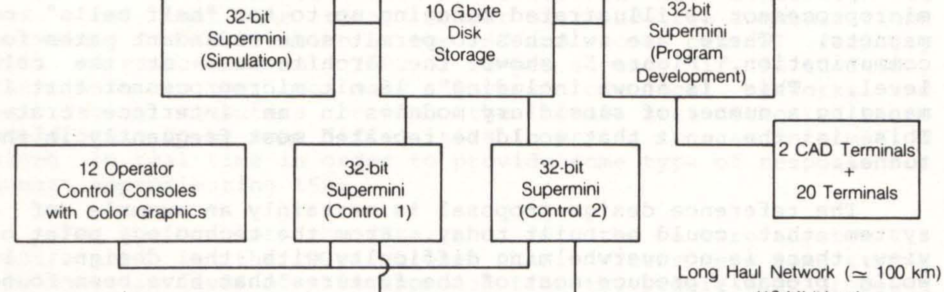

Fig. 4. SSC central control system (architecture overview).

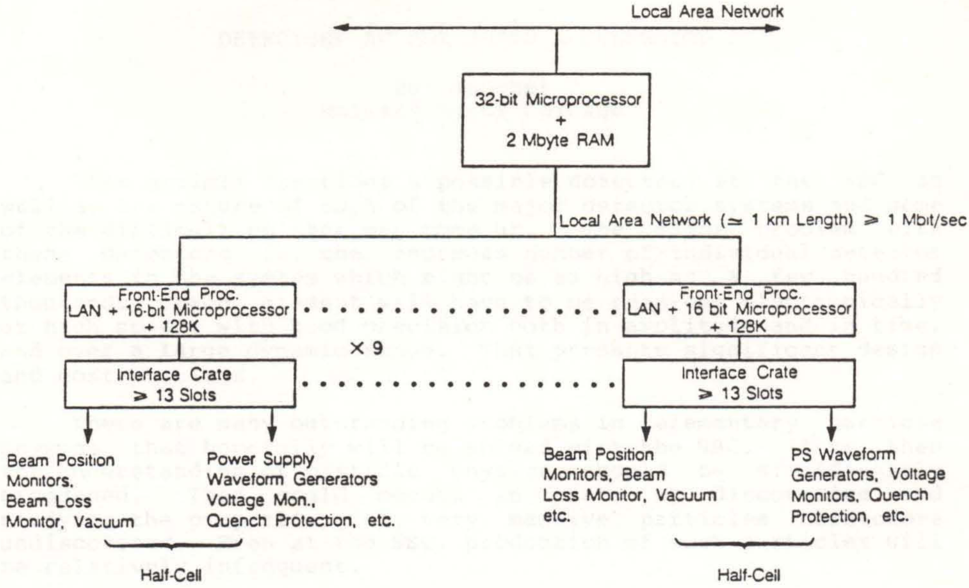

Fig. 5. SSC distributed microprocessors.

be multiplexed out of a single high powered 16 bit microcomputer. That is an example of something the final details of the design would have to address. The similarities to the Doubler system are obviously very large.

Summary

There are a number of challenges with respect to subjects such as the requirements for local area networks, the distribution of the data base, the question of the local distribution of electronic equipment in the SSC tunnel, reliability, and redundancy. Fundamentally, the control system is not a system that will have a significant impact on the basic questions concerning the possibility of constructing the SSC. It is certainly correct to place the major part of the research effort into developing the magnets as well as defining the actual physics goals of the experimentation. The control system can undoubtedly be matched to the requirements so defined with equipment available today, and certainly with equipment to be developed over the next three or four years during the R&D phase.

DETECTORS AT THE SUPER ACCELERATOR

Mel Shochet
University of Chicago

This article describes a possible detector at the SSC as well as the nature of each of the major detector systems and some of the difficulties that may come up. The major problem with these detectors is the enormous number of individual detector elements in the system which might be as high as a few hundred thousand. Each element will have to be read-out electronically at high speed, with good precision both in amplitude and in time, and over a large dynamic range. That presents significant design and cost problems.

There are many outstanding problems in elementary particle physics that hopefully will be solved with the SSC. If so, then the understanding of particle physics should be significantly broadened. This could occur, in part, by discovering and studying the properties of very massive particles heretofore undiscovered. Even at the SSC, production of such particles will be relatively infrequent.

Figure 1 is a schematic of what such a collision might look like. When two protons collide, massive particles may be produced in the process. A massive particle will decay essentially instantaneously into a large number of lighter particles. Some of those particles are metastable. In the figure, the massive particles which decay rather quickly have been indicated by dotted lines. The distances for those decays to occur are something on the order of 100 to 150 microns. The particles from the decay are often produced in clusters. In order to discover and then study the properties of the massive particles, it is necessary to measure the characteristics of each of the final-state particles. The detector will thus have to locate secondary vertices with high efficiency. That means that the detector will have to be able to determine that a metastable particle had been produced and that it decayed 100 to 150 microns away from the primary vertex. The detector will also have to accurately measure the directions of all the particles produced in the interactions and accurately measure the energy of each particle or in the case of very tightly clustered particles, the energy of the cluster.

To give a feeling of the environment that this detector will be working in, one is talking about studying processes which are rather rare so that they might occur at the rate of one event per hour or even one event per day or per week. This is to be compared to the total interaction rate of something on the order of ten million per second. In one hour the detector must pick the single most interesting and unusual event out of the ten billion interactions that occurred in that hour. In each of

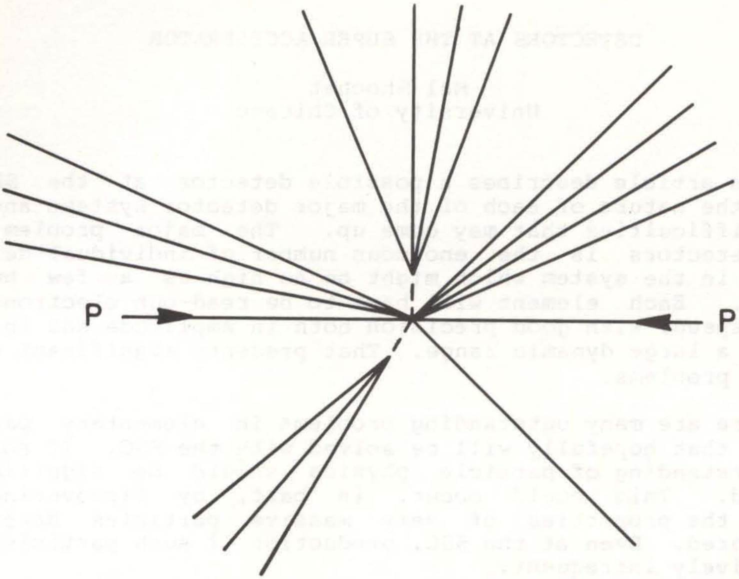

Fig. 1. Schematic of what a typical collision might look like. The incoming protons have arrows on their lines.

these events some 50 or 100 particles will be produced. One will therefore have to build a detector with very good time resolution because of the fact that on an average there will be an interaction every 100 nanoseconds. Very good spatial resolution is needed to detect the secondary vertices and measure the directions of these particles with precision.

Figure 2 is a simulated picture of particles coming out of an interesting interaction superimposed on four other non-interesting events to illustrate the importance of time resolution. Notice that an enormous number of particles come out. Time resolution is really at a premium to reduce the number of events that appear to occur simultaneously.

General Detector

What might such a detector look like? Figure 3 is an illustration. Notice the distance scale of two meters. There is a vacuum pipe bringing the particles into the interaction region. The particles interact in the center. Very close to the interaction region there is a vertex detector to locate any secondary vertices. Following that is a large region filled with the central tracking chamber. The function of that device is to

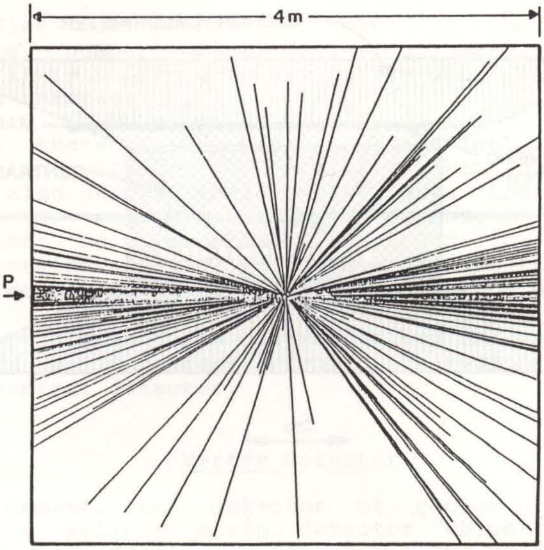

Fig. 2. Simulation of an interesting event superimposed on four dull events.

accurately measure the directions of all the particles produced in the interaction. A superconducting magnetic coil is incorporated to produce a large magnetic field in the center. With the magnetic field the charged particles don't move in straight trajectories but have curved trajectories as they are influenced by the magnetic force. A measurement of the momentum of the individual particles can be made by measuring the radius of curvature of those trajectories. That measurement is also the responsibility of the tracking chamber.

The electromagnetic calorimeter is just outside of the central tracking chamber. The calorimeter is a device which measures the energy of a particle or a very tightly clustered group of particles. The calorimeter will be highly segmented to avoid summing the energy of two particles or clusters that are close together. This segmentation must occur many times in both the axial and azimuthal direction. Individual elements look toward the vertex. The energy of the particle must be measured very well. The general technique is to have the particle interact in a dense material, which is interspersed with detecting layers. In the electromagnetic calorimeter an electron or a photon will interact close to the front end and then produce some additional charged particles. A detector layer then samples how many particles there are. Then there is another layer of very dense material in which additional electrons and photons are produced, thus building up a cascade. As the energy is absorbed

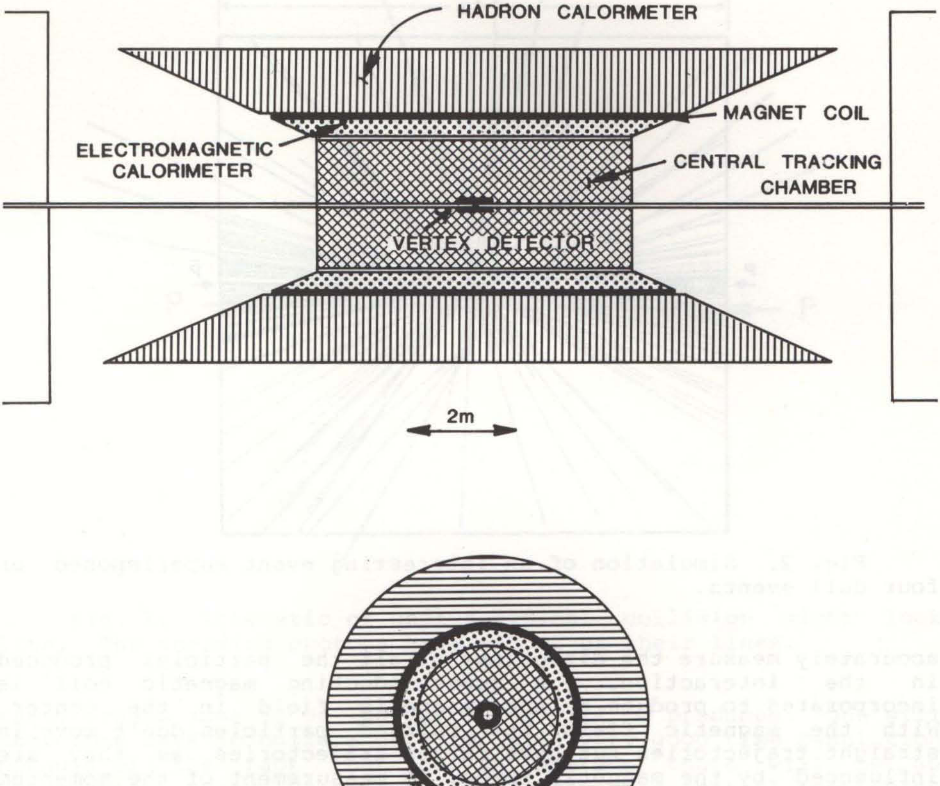

Fig. 3. Schematic of an SSC detector. The upper view is along the direction of the colliding beams, the lower view is transverse to the beams.

the cascade dies out before the back of the detector. By sampling the number of particles at many depths in the calorimeter, it is possible to get a very good estimate of the particle that entered.

Finally, there is a hadronic calorimeter outside the electromagnetic calorimeter. Particles other than electrons and photons will also deposit their energy when they strike heavy material in the same characteristic cascade but the cascade is longer so that more material is needed to measure those particles' energies. Again, the same technique is used with many layers of detectors sandwiched between layers of heavy material.

The regions forward and backward along the beam pipe contain similar kinds of detectors to cover the regions which are left open in the central detector.

Vertex Detector

At the moment, the detector of choice for the vertex detector is a silicon strip detector which is made out of a depleted semi-conductor layer with metallic conductor on each side. A particle passes through, ionizes the material, and the charge is collected. If very narrow strips of metal electrode can be used, one can get spatial resolution on the order of five to ten microns and can separate two particles that are as close as fifty or one hundred microns apart. With such resolution, one can do an extremely good job of finding secondary vertices which are no more than 100 to 150 microns away from the primary vertex. Figure 4 shows how such a detector might appear. Typically one might have four layers of these detectors. For a small section of the silicon strip detector, 5 millimeters wide by 50 millimeters long, there would be strips every 50 microns. One hundred detector channels would be required for that 2.5 cm^2 region alone.

What are the problems associated with such a device? First, with strips (essentially detector elements) every 50 microns and with many layers in a cylindrical system, there are between ten thousand and one hundred thousand channels, depending on the longitudinal segmentation, i.e., whether the strips run the whole length of the detector or are broken up into segments with each segment readout. The signals are low level, thus requiring amplification near the device to avoid noise pickup. There will be problems of having to multiplex these channels for readout since it is not possible to bring out 10^4-10^5 cables from the center of the detector. What is needed is a chip which will amplify with very low noise, will integrate the signal, will store it, and will produce a multiplexed readout. If such a chip could actually be imbedded into the silicon detector itself, that would increase reliability because one would not have to have the problems of external mechanical connections. This would result

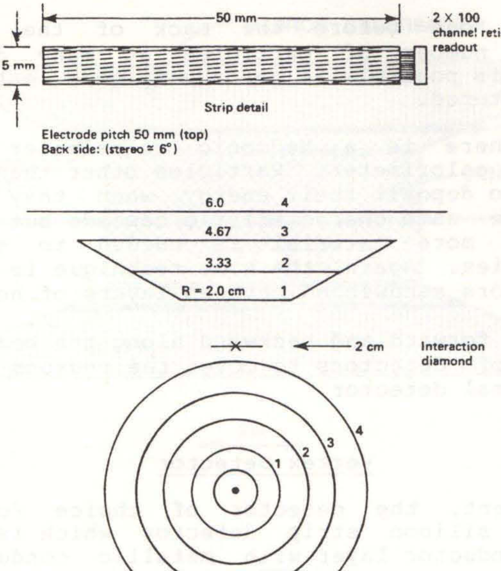

Fig. 4. Silicon strip vertex detector.

in a significant reduction in the number of external connections, and represents both a cost saving and an obvious improvement on reliability.

Second, there is a problem of radiation damage. With this detector as close as it is to the interaction point, the detector will see between 5×10^5-5×10^6 Rads of radiation per year. Normal silicon detectors start deteriorating after about 2×10^5 Rads, so there is a problem. MOS-based amplifiers usually start to run into trouble after about 4×10^3 Rads. Clearly, there is work that has to be done on producing devices which are radiation hardened.

There is one other problem associated with a silicon device of this size and that is the fabrication problem due to the size limitation for a single crystal of silicon. That has led a number of researchers to study the possibility of using non-crystalline semiconductor detectors for such a device. That is research which is in its very early stages.

Central Tracking Chamber

The next detector element is the central tracking chamber. It is required to measure the directions of all the particles produced in the collision. Figure 5 illustrates the situation as the beam would see it with a typical particle coming out and being bent by a magnetic field. The chamber contains many fine sense wires strung in the volume and connected to a high positive voltage. There are other wires not shown in Fig. 5 which are field shaping wires which carry the negative voltage. As a particle passes through the gaseous atmosphere in the chamber, typically, argon/CO_2 or argon/ethane, it ionizes the gas. The electrons that are freed in that ionization drift with a constant velocity toward the nearest sense wire. The length of time that it takes for the electrons to reach the nearest wire measures how far the particle passed from the wire. Thus the time of arrival at a wire measures the angular position of the particle at the radial position of the wire. A third coordinate is also needed, the position along the wire in the direction of the beam. That can be obtained by using wires of finite resistivity so that when charge is deposited at some point along the wire, the resistance from that point to the two ends of the wire are different and the charge which arrives at the two ends will divide accordingly. The charge that is collected is amplified on each end and the ratio of those signals gives the position along the wires. That way, one can get a three-dimensional readout of the particle's position.

There can be a very high local track density due to the fact that there are heavy particles each of which decay into many secondary particles. In addition, the maximum drift time for electrons should be rather small so as not to confuse tracks from one interaction with the tracks from the next interaction which on average occur 100 nanoseconds later. Those two constraints of high spatial density and the need to collect the charge in a short period of time lead to a requirement that the distance between sense wires in the drift direction be less than 1 cm. In addition, a particle should pass at least 100 wires as it moves radially outward from the interaction point. When those numbers are put together with the typical size of the tracking chamber, one finds a requirement of approximately one hundred thousand wires each with amplifiers on both ends.

The electronics must give a time resolution of a few nanoseconds so that the location of the particle can be measured with good precision. Good amplitude information is needed to measure the location of the particle along the wire. Because of high multiplicity, which leads to high track densities, good two particle resolution is also needed. This means that if a signal on the wire appears to show two pulses close together in time, it is desirable to detect the two pulses and to measure their amplitudes and their times of arrival. The net result is a readout system containing flash ADC's which can sample 100 or 200 million times a second. Such a system produces an enormous amount of data. Something like a ring buffer is needed to store all the information, because it takes a finite time to decide that the interaction that occurred was one that should be studied

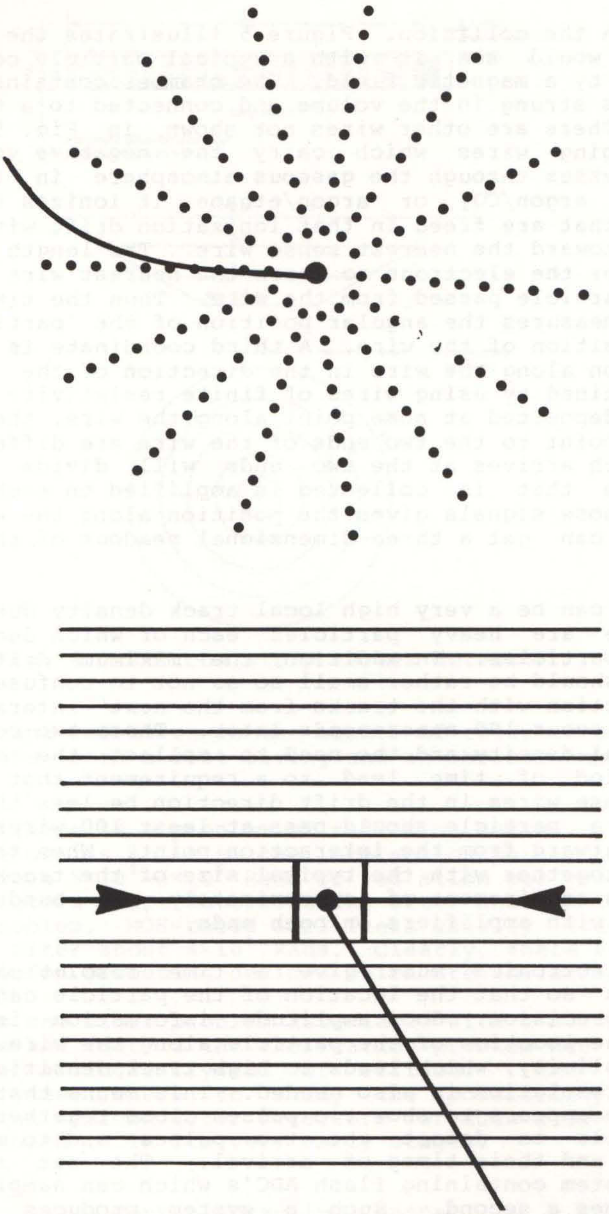

Fig. 5. Wires for central tracking detector. The top view is along the particle beams direction; the bottom view is transverse to the beams.

in more detail. It is not possible to send all of that information up to a computer, because there are one hundred thousand wires each with two amplifiers and each with many samples. A system of data compression is thus needed so that only the channels that have non-zero amplitude are sent up. Again, the main difficulty is that there are approximately one hundred thousand channels of electronics and the requirements on each of those channels is extremely high.

Calorimetry

The next part of the detector is the calorimetry. This is the device which measures the particle energy. A very fine grain division is needed in the electromagnetic calorimeter. The segments which are individually read-out form a geometry which is projective in depth. This arrangement results in tower-like elements each of which points back to the interaction region. This is done so that an electron produced at the interaction region enters one segment and deposits its energy in that one tower as the shower develops in depth. Projective geometry is needed to avoid ambiguities because of the large number of particles produced in each collision.

Electrons play a crucial role in trying to locate and study very massive particles in high-energy collisions. The ability to find electrons near other particles is important. As noted earlier, electrons and photons deposit their energy in a different characteristic depth pattern compared to hadrons; they deposit their energy early. Other particles deposit energy more uniformly throughout the depth of the detector. Since it is necessary to identify an electron that is close to other particles, it is necessary to have a maximum size for each individual tower of roughly ten centimeters on a side. Three or four segmentations in depth is necessary to see the characteristic development of the shower and be able to separate electrons from other particles.

Since it is important that energy information from the event not be lost, it is necessary to minimize the dead space between the calorimeter elements. Also, very good energy resolution is needed. It is thus necessary to be able to associate a voltage that comes out of one of these detector elements with the energy of the entering particle. The response should be uniform to different types of particles so that thirty millivolts on the detector corresponds to five GeV of energy, independent of the kind of particle that went in. Further, in order to make sense out of the information, the response should be uniform over the surface area so all cells will have the same response or at least a response that can be calibrated and is rather constant in time. A scheme is therefore needed for maintaining the calibration of such a system.

When all of those requirements are put together, the detector that at present seems most feasible is made of uranium and liquid argon. It would consist of plates of uranium separated by thin gaps of liquid argon. The ionization in the liquid argon would be read out with copper pads. As particles go through the uranium and interact, the secondary particles that are produced pass through the liquid argon and produce ionization. The ionization is collected by the copper pad electrodes. The problem here is that very small signals are produced so that the detection system and amplifier system must have low noise. A typical system would contain approximately 25,000 channels, in which the amplitude measurement has to have at least a ten bit dynamic range and an eight bit resolution with very low noise. One has to have a reliable and relatively inexpensive way to maintain the calibration so that each of these 25,000 channels can remain calibrated to one half per cent or better. So again, the requirement on the electronics is not trivial.

Triggering

Interactions are occurring at the rate of ten million per second and about one per hour is really interesting. How is that one picked out? It should not be at random or someone will to have wait a long time to see the event of interest. Thus we come to the question of triggering. How are the interesting events selected?

How much data is there for each one of these interactions that occurs? First of all, assume a sparse scanning technique that records only those channels containing non-zero data. If no particle struck a given wire at a given sampling time, nothing is recorded for that wire. But there are still as many as one hundred particles in an event. Each of these tracks passes by one hundred wires in the central tracking chamber. Each one of those wires has amplifiers on two ends. Typically there are five samples out of the flash ADC encoders for a track. That gives one hundred thousand bytes of information for the event. However, since it is a sparse scan, it is necessary to indicate which wire produced each signal. In addition, the time of the digitization must be indicated. When all of that is put together it gives something like 4×10^5 bytes of tracking information, in addition to calorimetry and vertex detector information. Conservatively, there would be approximately a half a million bytes per event. Just to indicate what that means, two or three events per second would saturate a 6250 BPI, 200 inch per second magnetic tape drive. A three month run, recording two or three events per second, would produce something like 5×10^{12} bytes of data to be analyzed. That is a problem all in itself. (Some of these issues were discussed in last year's Fermilab Industrial Affiliate Round Table on Supercomputers.) What is needed then is a system which can accept events at an average rate of one every one hundred nanoseconds and select the one most interesting event out of each ten million so that about one event per second can be written on magnetic tape.

How is that selection done? It can be understood by focusing on one characteristic of these events. That is the fact that in events of interest there is a local concentration of very high energy deposition in the calorimetry. In a typical uninteresting event, there are lots of particles going in all directions. Generally, each one of those particles deposits relatively little energy. This results in a uniform energy deposition over the whole detector. For events of interest there can be a very high-energy deposition in small local regions. What is done is to locate as quickly as possible each one of these clusters of energy and classify them. One finds out where they are in the detector and whether they are likely to be electrons or whether they are likely to be jets of other particles. Once that information is available for all of the clusters and all of the electrons, it can be put together in a list and the overall topology of that event examined to see if it is one that is likely to be of interest.

The strategy that one employs is to reduce the number of interactions, the number of events that one is looking at in a series of steps. Initially, very fast decisions must be made because the number of events is enormous. Then most of the events are removed and relatively few remain so that one can afford to spend more time thinking about each one of the remaining events. Thus decisions that take more time are left to a period when most of the events have been removed from the data sample.

Table I shows an example of how this might be done in a series of steps. One starts out with ten million events per second. Only those elements of the calorimetry which have energy above some reasonably high threshold are examined. All of the elements of the detector which have low energy deposition aren't even looked at. The energy of all the clusters is then added and that total energy has to be above a certain level. If so, it is a potentially interesting event. That decision can probably be made in about three hundred nanoseconds, but that does not cost live time in the detector because storage elements are used that make it possible to make the decision in real time. That is, delay lines are used for the analog signals, and ring buffers are used for digital signals so that the data is slowly percolating through the system while the decisions are being made. No data has been lost. By the time it gets to the end of the percolator, a decision must be made whether or not to keep that event for further analysis. With such a scheme the number of events per second can be reduced to something on the order of a 0.5×10^5.

Table I.

# of Events/Sec	Criteria	Time Required
10^7	ΣE_i for $E_i > E_{TH}$	300 nsec (dc-use delay lines & ring buffers)
$1/2\times10^5$	E of highest energy cluster E of highest e-like cluster Missing energy	1 μsec (5% dead-time)
5000	Find all clusters Determine their energies, locations Identify as e or jet of particles Find μ's Take "interesting" topologies	10 μsec
50	Refine above information	2 msec (read out all data)
10	Analyze in dedicated processor	100 msec - 1 sec

At that point, it is possible to take about one microsecond to look at each one of those events in more detail. That has a cost; about 5% of the detector real time will be lost looking at interactions. This is something that one can live with. At that point the highest energy cluster in the detector can be located very quickly and its energy accurately determined. The highest energy electron cluster in the detector can also be located and its energy determined. All of the energy in the detector can be added up and compared with the energy of the protons that have collided. There is something that hasn't been seen if the energy in the detector is much less than the energy of the initial particles. It has already been noted that electrons are very important for discovering and studying new massive particles. Electrically neutral relatives of electrons called neutrinos are also important. The neutrinos do not interact in the detector. If they carry a lot of energy, that energy will escape and the detector won't see the full energy of the collision but significantly less. That situation can be detected in approximately one microsecond.

One can make some requirements on those parameters and reduce the event sample to something like five thousand per second. Then one can take another ten microseconds to calculate.

For the collider detector that is being built at Fermilab all of the clusters in the system are found. At the University of Chicago, electronics have been designed which can find arbitrarily shaped clusters of energy in such a system without any prejudice on the shape every 150 nanoseconds. Their energies and their locations can be determined with high precision. Likely electrons or jets of other particles can be identified. The same holds for muons, electron-like objects which are identified in another way which won't be discussed here.

This essentially gives the overall topology of the event. The interesting events are finally defined by the software. What the software accepts as interesting can be adjusted as more is learned about the physics at these very high energies. When some new kind of process is found that wasn't anticipated, the software can be modified so that none of the new events are missed.

At that point the number of events has been reduced to a rate of fifty per second and two milliseconds can be taken to read all the data out of all of the subsystems into a large buffer. During that time the information can be refined, and some additional calculations made to reduce the number of events to ten per second. Each of those ten events could be sent to its own dedicated processor which could take 100 milliseconds per second to analyze the event and determine whether it is really of interest and whether or not one wants to keep it on magnetic tape. It is clear that the problem here is basic design. One is trying to make a decision involving a very large number of channels very quickly with high precision.

These are just a few comments that are based on a preliminary analysis of the problems that will face detector designers for the SSC. It is clear that when serious design of these detectors begins, problems will be found that haven't even been thought of yet. That's when the help of industry will be needed to solve them.

ROUND TABLE ON INDUSTRIAL PARTICIPATION IN LARGE-SCALE SCIENCE PROJECTS

Chairman, Dr. Richard Lundy

The Participants

Dr. Ray Beuligman is Program Director of Energy Systems, Convair/General Dynamics. Convair has fabricated many magnets for fusion and isotope separation.

Dr. C. H. Dustmann is from Brown-Boveri. He is now working on superconducting magnets for the HERA colliding-beam system in Germany.

Dr. John Hulm is Acting General Manager, Research and Development Center, Westinghouse. He is one of the developers of modern superconducting wire and a member of the Board of Overseers for the SSC.

Mr. Dick Rhodenizer is Manager of Systems and Products Engineering, Medical Systems, General Electric.

Mr. Carl Rosner is Chairman and Chief Executive Officer, Intermagnetics General. Intermagnetics General was the primary supplier for the Doubler superconducting wire.

Mr. Ryusei Saito is Chief Engineer, Nuclear Fusion, Hitachi. He has been involved in construction of magnets ranging from the KEK accelerator to levitated trains.

Dr. Ed Temple is Head of the U. S. Department of Energy Working Group on the SSC.

ROUND TABLE ON INDUSTRIAL PARTICIPATION IN LARGE-SCALE SCIENCE PROJECTS

Dick Lundy (Fermilab):

Today I hope to elicit from each of the panel members and some of the audience a response to a hard hypothetical question. Let's define an index that runs from one through ten in industrial participation with one being minimal and ten being maximal. The question for each member is: "What do you think the optimum number is on philosophical grounds and what do you think the realistic number is?" Realism has to be taken into account because we have real laboratories, real industries, and real times. I'll also be very interested in examples from the past of cooperation, of participation, and how it worked--a hindsight view of the good and the bad of industrial cooperation.

Let me set the extremes of this participation scale with an imaginary example. On a scale of one, a laboratory would decide to build an SSC and would hardly let anyone know. They would make numerous trips to the hardware store, buy lots of nuts, bolts, bar stock, and steel plate, work furiously night and day on the site and assemble the SSC themselves. They would install the accelerator and pray that it would work. That is the "one" end of the scale, the low end with total laboratory commitment and no real industry involvement except as a basic supplier of materials. The Energy Doubler, probably rates at the two or three level in part due to the high risk nature of the endeavor when it started five years ago. We bought basic commodities, such as steel plate, and we bought the next level up, fabricated subassemblies. We did a lot of drilling and burning and welding and praying here on the site. The other extreme, ten, can be illustrated facetiously with one side of an imaginary conversation: "Yes, this is Big Corporation, incorporated. Glad to be talking to you. My name is Newhart; I'm a sales engineer here. You want 40 TeV in the center of mass, with a luminosity of $10^{33}/cm^2$ sec. Well, we've been selling a lot of those this spring. I'll have to check stock... Yes, you're lucky. We've got two in stock. We've got one with experiments and one without.... Yes, they both have twenty-year warranties--no problem with that. Now, most of our customers take the one with experiments and they get about a Nobel Prize per year with four experimental areas.... You'll take that one? That's fine. Yes, we'll deliver it and set it up next Friday. Only one question, now, will that be cash or on your credit card?" That would be cooperation.

Now I should ask Ed Temple to make some opening remarks. Maybe Ed will tell us what the right numerical index is and whether it's going to be cash or credit. He's in a unique position--he represents the sponsor. Assuming there is an SSC, the participation will have to be played under the ground rules that the sponsor, the Department of Energy, operates under. These are the federal procurement regulations. There are good and bad features of these regulations. We've got to maximize the good and minimize the bad. That may be one of the major problems for the SSC builders.

Ed Temple (Department of Energy):

The first thing I want to discuss is a little bit about the SSC organization. Figure 1 shows the SSC reference designs study organization. Here I want to give full credit to the lab directors and the Reference Design Study Group who have produced the foundation upon which the Department of Energy (DOE) can go forward and upon which this kind of meeting can be held with some real serious paper studies for reference. These paper studies will be available to the world at large sometime in June.

Within the Department of Energy, the Secretary is Paul Hodel and the Director of the Office of Energy Research is Alvin Trivelpiece, so those are two principals in these discussions. The Chicago Operations Office of DOE will be an important contract administration arm for the Department for this effort

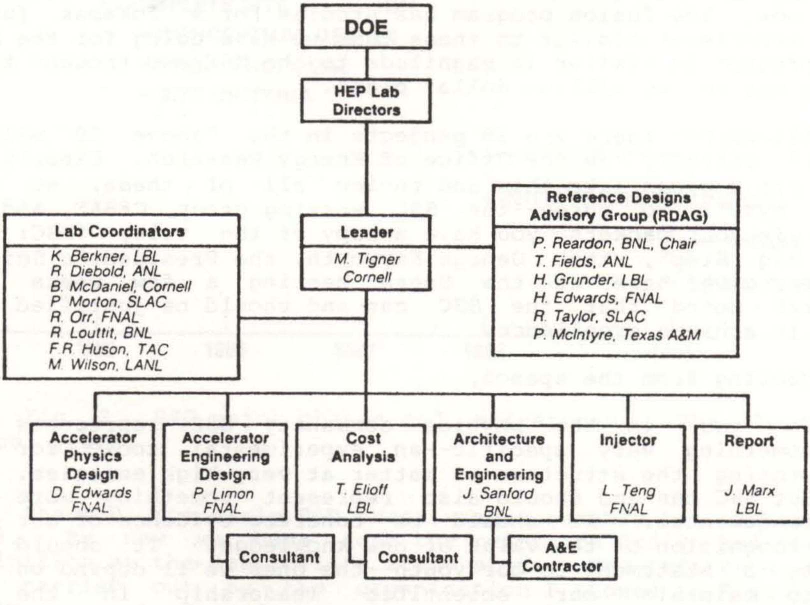

Fig. 1. SSC reference design organization.

for the interim period. In the Office of Energy Research (ER) under Trivelpiece, there is the Office of High Energy Nuclear Physics headed by Jim Kane and the Division of High Energy Physics headed by Bill Wallenmeyer. I head the Division of Construction Management Support in the Office of Energy Research. In that role, I provide office-wide oversight of all projects in ER for the Director or the Deputy Director and then provide construction management support to the various program divisions, high energy and nuclear physics, fusion, and basic energy sciences.

Since this is a round table on industrial participation in large science projects, it may be propitious that I have the opportunity to be here. In the Office of Energy Research, we actually have some semi-large scale science projects right now. In high-energy physics, for example, the Saver has just been completed, TeV I and TeV II are under construction, and the SLAC linear collider is getting underway at SLAC. ISABELLE is just being terminated, and the SSC, we hope, is getting kicked off. The Continuous Electron Beam Accelerator Facility (CEBAF) is a brand new project that the Southeastern Universities Research Association will be building for us in Newport News, Virginia. In fusion, TFTR has recently been completed. I believe that TFTR and the Energy Saver are fantastic successes in big science projects so we do have a record of success to be building on right now. The fusion program has studies for a Tokamak fusion core experiment similar to those that we have going for the SSC. This project is similar in magnitude to the SSC and thought to be in the one to two billion dollar class.

Altogether there are 25 projects in the "above 20 million dollar" category in the Office of Energy Research. Clearly one can't hit a group like this and review all of these, so I'll limit myself to three--the SSC working group, CEBAF, and the TFCX. In your packets, you have a copy of the talk, "SSC: The Next Big Step", that George Keyworth, the President's Science Advisor, gave here at the Users Meeting a few weeks ago. Keyworth noted that the SSC can and should be justified as a means to achieve excellence.

Quoting from the speech,

> To you in the physics community SSC represents something very specific--an experimental tool for probing the structure of matter at very high energies. But SSC can and should also represent something more fundamental. It should be concrete evidence of our recognition of the value of new knowledge. It should be a statement to our youth--the ones we'll depend on to maintain our scientific leadership in the future--that as a nation we value creativity, not just in physics but in all areas of science. And it should be evidence to ourselves and the rest of the world of our commitment to excellence in what we choose to do.

I think it's appropriate that Keyworth talks to the idea of excellence because I think that there's a record of excellence in this field and there's a good record to begin a big project like this on.

Figure 2 shows the staging for the project. These include a phase 0, which we're in right now, a phase 1, where we'll initiate R&D, a design, and complete a site selection process, and phase 2 where we do construction. Operation begins in phase 3.

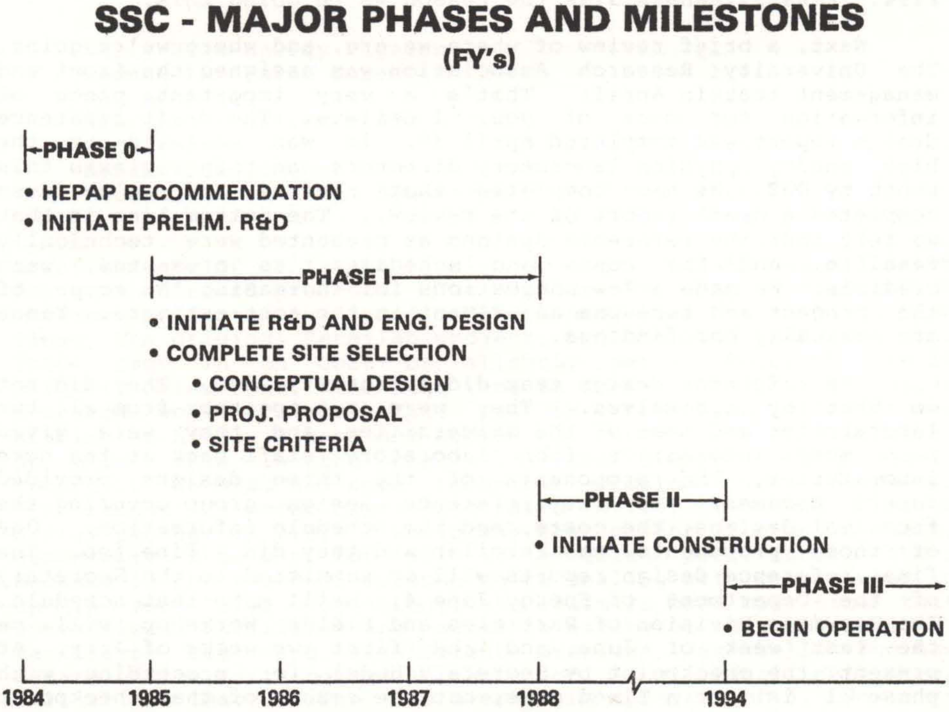

Fig. 2. SSC major phases and milestones. The dates are fiscal years.

Phase 0, the early R&D phase, began roughly in October of 1983. By now we have done the reference design and the cost definition at the feasibility study stage. The reference design was carried out by the organization I showed in Fig. 1. The HEPAP subpanel advised that the R&D be carried out in this phase. Phase 0 will end with the Secretarial checkpoint on proceeding with phase 1 which is scheduled for early August.

The design studies would be done in phase 1. Technical and cost assumptions made in phase 0 would be validated there including the use of supporting R&D. The site criteria document will be developed during phase 1. The conceptual designs will be completed, a proposal made, and systems tests completed. A systems tests here might involve a long string of magnets. Phase 1 will cover a period of three to four years. Funding requirements in that period are in the hundred and fifty to two hundred million dollar range. Phase 2 would be the construction phase. Our goal is to begin that in FY88 and complete it in FY94. Finally, Phase 3 is the reason we're doing this.

Next, a brief review of where we are, and where we're going. The University Research Association was assigned the front end management task in April. That's a very important piece of information for most of you, I believe. The draft reference design report was completed April 30. It was reviewed by the high energy physics laboratory directors and then reviewed this month by DOE. We have completed that review, and we've even completed a draft report of the review,. The bottom line is that we felt that the reference designs as presented were technically feasible, and the costs and schedules, as presented, were credible. We made a few suggestions for increasing the scope of the project and made one adjustment in the cost estimate. Those are basically our findings.

The reference design team did a fantastic job. They did not do that by themselves. They were put together from all the laboratories and some of the universities, and they were given much more information from laboratory staff back at the home laboratories. The proponents of the three designs provided superb documents to the reference design group covering the technical designs, the costs, and the schedule information. One of those proponents was Fermilab and they did a fine job. The final reference design reports will be submitted to the Secretary of the Department of Energy June 4. We'll make that schedule. The Snowmass Division of Particles and Fields Workshop will be the last week of June and the first two weeks of July. At present, the checkpoint by Secretary Hodel for proceeding with phase 1 is not a fixed date, but the result of that checkpoint will be announced on August 6.

For Architecture/Engineering groups, the operational piece of information is that the selection process for AE services and the management of R&D for SSC will be done by the URA integration group. This is the Universities Research Association. There's a member of the Board of Overseers for the administration group here on this panel, John Hulm. Jim Matheson, who's in the audience, is the Vice President of URA. A key part of their assignment is to prepare an R&D plan and a management plan for phase 1 and submit it to the Department by the end of June. The DOE field office for the SSC is the Chicago Operations Office.

In conclusion, there's one other point that I would like to make. It has to do with this idea of excellence and past performance. It is also related to the question Lundy raised, namely, how much industrial participation is desirable and how much is practical in this effort? Table I is a very brief summary of a list of final cost over initial cost estimates for some selected projects. The number for early high-energy physics projects is based on three projects: the original SLAC linac, the original Fermilab project, and the SLAC PEP project. These came in on cost or slightly below. Fermilab came in slightly below cost. The next line is recent Office of Energy Research worst cases. This was in a time when inflation was hitting us hard. I have not backed out the effects of inflation at all. This is the real world that everybody has lived in recently. We also had some technical problems with projects. For recent ER worst cases this ratio is 1.6. The Department of Energy average is 2.5. I would say that either of those numbers, especially the one for ER, is excellent performance. Then, because in the realm of science we're dealing with people wrapped up in very high technology, the claim is made many times that conventional facilities are easier to estimate. But conventional facilities can have significant overruns, as well. The new Senate office building had a ratio of 3 of final cost over initial estimate. The DOD average is a factor of 5. Now just taking the Defense Department and the Energy Department, there must be something different about the way in which they do business. The Alaska pipeline factor was 7.8 and the recent worst case for reactors is a factor of 10. Now a combination of lots of things go into how this number ends--the motivation, the drive, how hard you work, how smart you are. I think that how one does business here is partly what we're discussing when we discuss what is the desirable industrial participation and what is a realistic industrial participation.

Table I. Selected Project Performance Records.

	Final Cost Initial Cost Estimate
Early High-Energy Physics	≲1
Recent ER "Worst Cases"	1.6
DOE Average	2.5
New Senate Office Bldg	3.0
DOD Average	5.0
Alaska Pipeline	7.8
Recent "Worst Case" Reactors	10.0

Lundy:

For the ratio of final cost to initial estimate we want to shoot for numbers like one or less. I just wonder if Temple's ratio has a direct correlation with my 1-10 index. That would be a terrible result if it did. We could just dissolve the panel and go home right now. On my index there should be an optimum value somewhere between 1 and 10 that leads to a one in Temple's ratio.

Now to questions and comments. Would someone like to volunteer a comment about a successful project where the final cost met the original estimate and tell us what the value on my index scale was?

Ray Beuligmann (Convair/General Dynamics):

Convair-General Dynamics has been involved in the fusion program for at least seven years. Indeed a superconducting magnet industry has been producing large-scale superconducting magnets, at least for magnets in the range of 10 to 300 tons, for about seven years. There are companies here, such as Intermagnetics General Corporation (IGC), that have been involved in making smaller superconducting magnets much longer. The superconducting magnet industry exists now. It wasn't there when the Fermilab Tevatron was started. Speaking as someone who has been involved with the evolution of the industry, there have been some good stories and there are some that are not so good. A good example that has come out in discussions and is recognized in the fusion program by the committee that met to look at the role of industry in fusion was the Mirror Fusion Test Facility (MFTF). There Livermore Laboratory developed the magnet technology and did the conceptual design but chose to go out to industry for detail engineering design and analysis because the Laboratory didn't have the necessary skills. It knew that it needed the complementary skills that the aerospace industry possessed. That is they needed industrial skills to dot all the i's and cross all the t's. A competition was set up. We were deeply involved with that project. Bob Tatro, here with us today from Convair/General Dynamics, was the program manager for the project. That engineering job was done under budget and under schedule. We were learning about magnets at the same time we were building large coils. That project was a partnership. When they changed that program to MFTF-B and significantly increased the number of magnets, Livermore was then building the coils for which we had done the detailed design. Without going into detail they changed the physics and added a lot more coils to the machine. About $30M worth of magnets then had to be procured. Livermore came out with the specs and industry competed with its ideas. Convair-General Dynamics was successful in winning the competition. So far we have delivered 12 solenoids. We're winding transition coils and axial axi-cell coils now. I am

In conclusion, there's one other point that I would like to make. It has to do with this idea of excellence and past performance. It is also related to the question Lundy raised, namely, how much industrial participation is desirable and how much is practical in this effort? Table I is a very brief summary of a list of final cost over initial cost estimates for some selected projects. The number for early high-energy physics projects is based on three projects: the original SLAC linac, the original Fermilab project, and the SLAC PEP project. These came in on cost or slightly below. Fermilab came in slightly below cost. The next line is recent Office of Energy Research worst cases. This was in a time when inflation was hitting us hard. I have not backed out the effects of inflation at all. This is the real world that everybody has lived in recently. We also had some technical problems with projects. For recent ER worst cases this ratio is 1.6. The Department of Energy average is 2.5. I would say that either of those numbers, especially the one for ER, is excellent performance. Then, because in the realm of science we're dealing with people wrapped up in very high technology, the claim is made many times that conventional facilities are easier to estimate. But conventional facilities can have significant overruns, as well. The new Senate office building had a ratio of 3 of final cost over initial estimate. The DOD average is a factor of 5. Now just taking the Defense Department and the Energy Department, there must be something different about the way in which they do business. The Alaska pipeline factor was 7.8 and the recent worst case for reactors is a factor of 10. Now a combination of lots of things go into how this number ends--the motivation, the drive, how hard you work, how smart you are. I think that how one does business here is partly what we're discussing when we discuss what is the desirable industrial participation and what is a realistic industrial participation.

Table I. Selected Project Performance Records.

	Final Cost / Initial Cost Estimate
Early High-Energy Physics	≲1
Recent ER "Worst Cases"	1.6
DOE Average	2.5
New Senate Office Bldg	3.0
DOD Average	5.0
Alaska Pipeline	7.8
Recent "Worst Case" Reactors	10.0

Lundy:

For the ratio of final cost to initial estimate we want to shoot for numbers like one or less. I just wonder if Temple's ratio has a direct correlation with my 1-10 index. That would be a terrible result if it did. We could just dissolve the panel and go home right now. On my index there should be an optimum value somewhere between 1 and 10 that leads to a one in Temple's ratio.

Now to questions and comments. Would someone like to volunteer a comment about a successful project where the final cost met the original estimate and tell us what the value on my index scale was?

Ray Beuligmann (Convair/General Dynamics):

Convair-General Dynamics has been involved in the fusion program for at least seven years. Indeed a superconducting magnet industry has been producing large-scale superconducting magnets, at least for magnets in the range of 10 to 300 tons, for about seven years. There are companies here, such as Intermagnetics General Corporation (IGC), that have been involved in making smaller superconducting magnets much longer. The superconducting magnet industry exists now. It wasn't there when the Fermilab Tevatron was started. Speaking as someone who has been involved with the evolution of the industry, there have been some good stories and there are some that are not so good. A good example that has come out in discussions and is recognized in the fusion program by the committee that met to look at the role of industry in fusion was the Mirror Fusion Test Facility (MFTF). There Livermore Laboratory developed the magnet technology and did the conceptual design but chose to go out to industry for detail engineering design and analysis because the Laboratory didn't have the necessary skills. It knew that it needed the complementary skills that the aerospace industry possessed. That is they needed industrial skills to dot all the i's and cross all the t's. A competition was set up. We were deeply involved with that project. Bob Tatro, here with us today from Convair/General Dynamics, was the program manager for the project. That engineering job was done under budget and under schedule. We were learning about magnets at the same time we were building large coils. That project was a partnership. When they changed that program to MFTF-B and significantly increased the number of magnets, Livermore was then building the coils for which we had done the detailed design. Without going into detail they changed the physics and added a lot more coils to the machine. About $30M worth of magnets then had to be procured. Livermore came out with the specs and industry competed with its ideas. Convair-General Dynamics was successful in winning the competition. So far we have delivered 12 solenoids. We're winding transition coils and axial axi-cell coils now. I am pleased to note that this has been a cost plus incentive program. We have met every incentive milestone to date and collected every incentive fee. We are not going to get rich on those fees because we plow all of them right back into the technology so that we can work on programs like this. We expect to get the remaining incentive awards both on the costs and the schedule milestones.

This is a good example of a laboratory-industry transition. There is an industry now, there are skills. But industry doesn't have all the skills. There's a tremendous amount of skill unique to accelerator magnets within the laboratories. Industry has a different complement of skills. Together we can handle the problem. There will be risks and we are willing to accept those risks.

However, don't ask me to make a fixed price bid on something that you have designed and thrown over the transom to us to build from the print. The upside benefits are not worth the downside risk. We can't even do that successfully within our own companies, design it and do the engineering and then throw it through the transom. This is bound to be unsuccessful when you do it from a laboratory to industry. It must be a partnership from the beginning.

Leon Lederman (Fermilab):

Today we are fortunate to have several people from Japan and Germany here. I'd like to capitalize on that and ask them to give us a rough idea of the equivalent of Ed Temple's table for their countries. In other words, what is the general ratio of final costs over predicted costs for high technology projects?

Cord-Henrich Dustmann (Brown-Boveri):

In Germany, in the field of high-energy physics, the laboratories are also proud of meeting their initial cost estimates. For reactors, we have the same trouble in Germany in that the final costs are much higher than the initial estimated costs. In general, it seems to be the same picture in Germany as in the U.S.

Ryusei Saito (Hitachi):

In Japan we also have some projects suitable to be called Large Scale Science Projects. In most cases, economic conditions are usually not good, especially in the smaller projects or in the R&D projects prior to a big job.

Lundy:

One of the elements that Beuligman implied was an important factor in his success was the key phrase, "cost plus incentive." Now that's something that's almost never used. I think it's accurate to say that DOE discourages us from incentive terms in a contract because positive incentives also usually imply you must put in penalties, and we're never very good at collecting penalties.

John Hulm (Westinghouse):

I'd like to ask a question. I know of some DOE contracts in which there are incentives. For example, GOCO's are operated in that mode, are they not?

Lundy:

Do we have any incentives as a GOCO?

Hulm:

No, but how about contractor operated situations like Stanford or Oak Ridge? Aren't many of those places operated so they have an incentive?

Lundy:

But it's a fixed fee.

Harrison Wroton (Martin-Marietta):

No, at Oak Ridge it's an incentive.

Hulm:

Yes, sure, it depends on performance.

Lundy:

We've got to change our contract with DOE. So you endorse the concept of an incentive?

Hulm:

Yes, I do. Very much so. I think we should do more of it. That's what the free enterprise system is all about.

Lundy:

Basically, it's covering the industry against the risk that makes it work. Not asking them to take a risk in a field that's full of unknowns?

Hulm:

No, DOE could make it so they didn't get anything if they didn't do it right.

Lundy:

That's a fair proposition?

Hulm:

Yes.

Lundy:

Temple defined phase 0, 1, 2, and 3. Do you have a feeling when one can go out for the cost plus incentive or other modes? Can we get industry involved at the end of Phase 0, immediately into the R&D program or would they like a year while we struggle with it?

Hulm:

I believe it's already too late. We saw three designs that were made, one of which had industry participation. I think that's very good. I think industry should be involved from the very earliest possible moment. I would like to see industry involved in the other two designs, to be honest.

Lundy:

You recommend involvement that soon?

Hulm:

In some sense it's too late. But, of course, you can't go back in history. It really isn't too late. At this point, these things are mostly conceptually built. I gather some modules have been built. The next phase would be a great time to involve industry. Let me say from my viewpoint why its so important. There are going to be many industries involved in this machine. It's a very complex system; computer technology and controls technology are also involved. Altogether, a very wide variety of technology is going to be used.

However, I'm only going to say a few words about superconductivity. It's a good example and its the core of the machine. We really couldn't do this project without high field superconductivity. Such a machine could not be built with normal magnets. In my view, and excluding cryogenics, which is a fairly healthy industry in its own right, the superconducting industry is not very healthy at the present time. I make those generalizations, though my colleagues would perhaps disagree. There are a number of small to medium size companies in the industry. For examples, these include the wire suppliers and material suppliers. They are hanging on, although some of them I think are doing all right. Some of these companies are represented on this panel, and they may wish to comment on what I've said.

An important advance was made with superconductivity, the ability to get very high magnetic fields with low expenditure of power. This is a key development which we didn't have prior to 1960. When an important advance is made in a new technology like this, one gets an opportunity. This often happens with such a breakthrough, the first applications are in science. These are primarily by R&D people because they understand, more than anyone else, what can be done with an extension of a variable like the magnetic field. As a result, the industry has been mainly focussed on scientific projects in the past 20 years. The successes are in scientific instrumentation. For example, superconductivity has revolutionized nuclear magnetic spectroscopy. It looks as though it's going to revolutionize medical imaging. In fact, medical imaging may be the first commercial or industrial application in the field. That's going fairly well, but these are fairly small magnets. The other applications are projects like MHD, still basically R&D, fusion, which is R&D, and accelerators, which are clearly a scientific application. Superconductivity simply hasn't found its way into the general world of industry.

But that's typical of brand new technology. In other cases, like lasers, the same situation has occurred. The first applications are scientific and then commercial and industrial needs are identified. We're looking for commercial and industrial applications for superconductivity. They are coming very slowly.

It looked like we had a hot one in power station generators. The same kind of dipoles that were being built for Fermilab can be used as the excitors in thousand megawatt electrical generators. They save a great deal of energy by superconducting excitation. Unfortunately, the electrical industry is an economic disaster at the present time. It's just one of those facts of history that the United States consumption of electricity is on a plateau and hardly anyone is building power stations. We don't expect any new generation capacity for at least the rest of this decade and maybe not even until the mid 90's. So there's very little economic incentive to introduce new technology to the industry. Consequently, I think the application of superconductivity to power station generators is on the back burner at this time. It is being pursued all around the world to some degree--in Japan, the Soviet Union, in Europe, but we don't expect to see a commercial machine introduced for a long time yet.

The only other prospect at this point that I see of a major application is in levitated trains. That is also going on in Japan and nowhere else as far as I am aware. I'm glad they're doing it. It's a very interesting and important development. I'm sorry that we are letting the Japanese do it alone. I wish the Department of Transportation felt enthusiastic about superconducting levitation technology. In short, for large companies, the situation is discouraging. There's not a lot of incentive for our company, Westinghouse, or for General Electric, or even for General Dynamics, perhaps, to build and continue development of superconducting magnets.

So looking at the SSC opportunity, we see this accelerator's going to be the biggest job in superconducting magnets that is coming over the horizon for some time to come. I hope industry will have a major participation in as many phases as possible. If you nuts-and-bolts the job, "one" on Lundy's scale, industry will get nothing. The technology transfer will be zero and even the nuts-and-bolts people won't get any technology transfer. Obviously, at the other end of the scale, Lundy's "ten," there is radically new engineering. New ground is being broken in many fields of engineering. It is difficult to go to consulting industries and say, you guys do everything. I have no difficulty with Fermilab or any other group that is familiar with accelerator design playing a strong role. All the accelerator design knowledge is in the national laboratories and universities; it's not in industry. I have no problem with them providing the engineering leadership needed to put the SSC together since no one else is capable of it. However, it will probably have to be a pooled effort, because it's going to need all the accelerator design knowledge the country has. It may have to be pooled internationally with the entire Western world. However, it's possible to pick out parts of the machine, such as superconductivity, in which industry can have a major part. This would be a place where the engineering participation by industry is essential almost immediately.

I would like to see the SSC R&D group begin to commission magnet development projects soon. A variety of development contracts to build these dipoles put up for competitive bids would be useful. These don't necessarily have to be the three designs which have come out of the national laboratories. Of course, there have to be some specifications on the magnets. The point is to allow industry to innovate. This comes back to what Beuligman said about pushing the drawings over the transom. No technology is transferred by giving industry build-to-print orders.

Lundy:

The design study has shown a very short R&D phase since we want to start construction early in order to finish early. Would you say that a three or four year R&D period is really too short for industry involvement?

Hulm:

It is a little on the short side; however, these magnets are not so difficult. Technologically, they are not radically different from magnets that have already been done in industry. Industry could get on the ball right away and come up with some innovative designs and build some prototypes.

Lederman:

I would like Hulm to clarify the reason why he's anxious for industry to get in at an early stage. If you are discouraged with the pace of industrial applications of superconductivity, is it that you hope the applications will come eventually and that you want to keep industry's hand in?

Hulm:

That's exactly it. Of course, we don't want to do a WPA project. It would be better to lay engineers off than subject them to that. No, I think that because you're going to be building the biggest superconducting project in history, that you owe it to industry to qualify them and to advance the technology for future industrial applications.

Lundy:

We have two representatives from the Nuclear Magnetic Resonance (NMR) industry. Here's a question. How long was it from the time when the light bulb came on and you had the idea magnets should be built and sold until the time a salesman could take orders over the phone? What was the total time span and how does it compare with what you thought it would be when you started? Give us a real time-to-complete over estimated time-to-complete ratio.

Dick Rodenizer (General Electric):

That very factor is germaine to the question whether or not the three or four year R&D time period is appropriate. As far as GE's concerned, we started the initial development work for NMR magnets roughly two to three years ago. In South Carolina, we're currently getting the facility up to speed with very ambitious production targets, starting the end of this year and early next year. I think John Hulm oversimplified the complexity of the NMR magnet. The NMR systems use fairly large magnets. They produce fields which aren't high by your standards--one and a half tesla over large volumes. Undoubtedly, we will be going to higher fields. Uniformity requirements are 10^{-5} over these volumes which is not a simple engineering challenge. That has been done on a timescale similar to what you're talking about here. Another interesting analogy is that the initial work for the magnet was initiated in the R&D center in Schenectady. Their initial role is similar to what you're talking about for the national laboratories. They had the technology base, they started out with the design concepts, and they began to develop the design. However, Medical Systems was involved in the very early stages of the program. We worked with them through the design stages, input on manufacturing, and quality control. The transition from the research and development center to this commercial production business has been extremely successful. I never would have imagined that we could have done it as smoothly as it's now going.

A further point I'd like to make is that this finally is a very substantial commercial product based on superconductivity. General Electric has made a large commitment to the NMR development. There are capabilities and facilities that are now available which just simply hadn't been there in the past. Private industry will do that if there is an appropriate incentive. This can only help the national laboratories.

Lundy:

Carl, IGC is involved in both the nuts-and-bolts side of the business and in making magnets. Seemingly, there is no way you could lose. Would you like to make a comment?

Carl Rosner (IGC):

My association with superconductivity began with the discoveries in 1960. It's been a tremendously exciting technology to be involved in. I spent my first 16 years in superconducting technology with General Electric and was a party to getting superconductivity off the ground within General Electric. But then GE lost interest because there was no apparent industrial application at the time with opportunities to see markets in the hundred millions or even billions of dollars. I maintained my interest and founded Intermagnetics General (IGC) to try to be there when this industry would amount to something. Although some people in the larger companies don't like to acknowledge it, any hundred million or billion dollar industry still starts with the one million dollar industry. IGC chose to commit to that path. It's a bit frustrating to find that when the payoff is there, these large companies jump back in, make new commitments, and rediscover a technology and perhaps relearn something that we have known all along. New money gets wasted in many arenas to try to relearn or re-educate a new generation of participants, ignoring to some extent the accumulated experience that is still there. In fact, there is now a small superconducting industry that's willing to do anything and everything, i. e., both R&D and "nuts-and-bolts."

I remember when Bill Fowler from Fermilab first contacted us in the early stages of the Saver program. Intermagnetics General was quite anxious and willing to build the first magnets in industry. The only thing that kept us apart was the price. We felt that in order not to go into the factor of ten overrun regime, we needed to have sufficient money to do some of the R&D and some of the development work. However, somebody had divided the total cost of magnets that they needed and come up with a price of $10,000. That's what a magnet should cost. And we were asked to build the first magnet for $10,000. That was patently impossible. And so it's to this developmental arena that obviously some thought has to be given.

Now we've moved on to the SSC, and I should really compliment Fermilab for giving me a chance to be here and sound off. I'm grateful for the invitation to be on this panel, and I am excited about the prospect for commercial development of both the NMR and SSC technology. Finally, others have generally recognized the industrial opportunities. And yet, there is a level of frustration as to why couldn't we do this sooner and why aren't we doing it right now here in the U. S. after having learned all these lessons. In particular, John Hulm, who has been a similarly active proponent of getting superconductivity off the ground, also expresses a level of frustration in terms of his experiences in how to do this thing right. There are some real answers out there. I'm proud that, as a small company, Intermagnetics General (IGC) has made key contributions to the success of the Tevatron. Without our ability to produce a conductor of good quality at a reasonable price, the Tevatron might have been a much harder project. So that you don't get any wrong impressions, IGC probably collected less than one per cent of the total cost of the accelerator. Yet our work and accumulated experience was of crucial importance. Last, but not least, the help that we got from Fermilab in making the wire, towards the end of the project, led to a true partnership. These factors made it possible for industrial participation that in the final analysis turned out to be successful for both parties. Bruce Zeitlin in our company was instrumental in holding that effort together and maintaining the contribution and dedication. While this was obviously a corporate effort at IGC, Bruce Zeitlin and his colleagues have been the focal point of this activity, and we've been able to give him the support and the people to make this all possible.

Next, I would like to look at the problem of how to promote industrial, government, and university collaboration. On the basis of the accumulated experience that we have, I think the ideas are there, but the willingness is still the missing link. I think that Fermilab has been particularly successful in putting together a team that had the commitment and the staying power and the willingness to work long and hard hours. I think this has been the case at MMIS as well. We should build on this kind of teaming of partner relationships for the future.

From where I sit, the SSC will be built. We need it from an intellectual point of view, we need it from a national pride point of view, we need it from every conceivable aspect that you can see. The real challenge, however, is to do it constructively in such a way that at the end of the SSC effort there has been a technology transfer to a broader industry. This transfer should allow those participants that have been involved to really be established in a way that gives us a technological edge, if you like, so that we, in turn, can go on to bigger and better opportunities. Unfortunately it's very difficult to find examples at the Tevatron of industrial participation and technology transfer except perhaps for the very limited experience at Intermagnetics. My conclusion from this is that it hardly makes any financial sense to involve industry if indeed the reference design B is the one that has the most to commmend itself. One may ask, why shouldn't it be done at Fermilab? Fermilab already has the facilities, it has the people, it has the experience. How do you transfer that to industry without transferring the people? The way the SSC planning program is going, it's going to require a very hard and conscious effort to jump into industry participation at this point. The next fifty or one hundred magnets could logically be built at Fermilab. But then the opportunity is lost to start this technology transfer right now. That transfer and collaborative commitment literally has to be started from the planning point, because if you lose that first stage and first step, it's too late. People are no longer interested. Industry participation will only be reluctant and may immediately drop off when there either is no profit possibility or when industry is asked to build to print. This is the worst way of getting technology transfer going.

As an aside, I would like to make one more comment in a more speculative vein as to the future course of commercialization of superconductor technology. Would it not be nice if we could find a home use for superconducting magnets. Something like that may be in the offing. This was already a question which I discussed at a visit to the physics department at Stanford about 20 years ago. We were trying to do a market study as part of General Electric to see where superconductivity might be going. I was attached to that marketing effort as a technical advisor. Ultimately that experience led to a transition for me from applied scientist at GE to a businessman. The answer I got from Professor Fairbanks as to where superconductivity might be going was, "you ought to find a way to convince people that they can enormously increase their sexual pleasures if they sleep in a magnetic field." I submit that we may be close to at least testing that hypothesis with the whole-body NMR magnets. Certainly one person can be successfully surrounded by magnetic field and, lo and behold, two people can fit in some of the magnets that we are now building.

But perhaps there really is a potential home use. Whole body magnetic resonance (MR) spectroscopy is around the corner, and I can visualize that some health nuts may want to see what the food they eat in the morning does to their system at night. A handy NMR spectroscopy magnet at home could be used to check out what the food does to their system and how it gets converted. The delightful perspective is that magnetic resonance spectroscopy in addition to imaging will have an ever greater impact on the utilization of superconductivity. Futhermore, in the context of individual participation it is interesting and stimulating to realize that the number of magnets that will be built for applications to NMR or MR in the same time frame that DOE proposes to build 3,000, or 10,000, or 14,000 magnets, may not be so different.

Lundy:

Let me comment on something you said. In the middle and the late stages, the interaction between Fermilab and many industrial firms was very satisfactory and very productive. There has been mutual respect and trust on both sides. I'm personally convinced that with any reputable industry working on the SSC, you would have that same degree of cooperation and warmth. But the problem is in getting started when you don't know who wears the white hat and who wears the black hat. It's like mating porcupines. It's got to be approached delicately because it could go wrong quickly. In the limited time that's available, how can we sort out the pure in heart and the open-minded people that we can put together for the best interaction? How do we filter out the things that are going to end up badly? We really don't have time to work by trial and error on this question. I don't think competitive bidding does it. To my mind, competitive bidding involves a large risk of getting people in who don't understand the job or who have undervalued it, or who have planned to make a profit with change orders. Another option we've discussed is cost plus incentives. Obviously this Laboratory and perhaps others need to become experts in that; this is something we are not at the moment. How do we, during phase 1, get industry to start magnet designs or the analysis of existing designs? Perhaps Dustmann could enlighten us on how DESY and Brown-Bover, in Germany, handled this. Who made the proposition?

<u>Dustmann</u>:

The situation in Europe is different from this country with respect to the tradition of magnet builders. In the last 20-30 years, all the conventional accelerators in Europe have been built with magnets produced in industry. Thus, for conventional machines, there are a couple of companies in Europe which are able to deliver accelerator magnets. Brown-Boveri is one of these companies. On this basis, we came into contact with DESY in connection with the HERA project. This relationship started about 3-1/2 years ago when DESY began to design HERA. They started by contracting an industrial design study. This was contracted to two German companies.

The basis of this design study was, on the one hand, the Tevatron design, which in those days was the basis of plans for superconducting dipole magnets over the world, and, on the other hand, the magnet specification for the field of 4.53 tesla, the length for the magnets of 6 meters, and the harmonic quality which had to be met. On this basis, we started the design study and came to the conclusion that perhaps a cold iron magnet may be better in some respects. This was the basis of the contract between Brown-Boveri and DESY for producing three prototype magnets of our cold iron type. The first of these has been delivered to DESY. Numbers 2 and 3 will be delivered in June or July. The experience here is parallel to that which has been mentioned before--idustry should come into the job as soon as possible; the ideas of industry should be put into the design at an early stage. Finally, the R&D should be done in small steps which can be overseen so that there is interaction before the goal of the final magnet is reached.

<u>Lundy</u>:

I might comment that the relativistic heavy ion collider at Brookhaven (in some sense a replacement for the ISABELLE Colliding Beam Accelerator which was terminated) is probably going to draw heavily on the experience at Brown-Boveri and at DESY because the magnet requirements for the machine that is being discussed will be very similar to the work that's been going on at DESY.

Dustmann:

Let me give you a short impression of what the HERA magnets look like. First, I want to give you an overview of what the HERA project is. Figure 3 is a view of the accelerator enclosure which has a circumference of 3.3 kilometers and is about 15-20

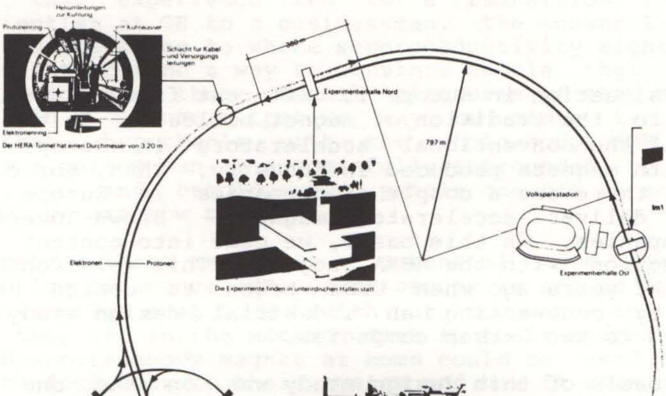

Fig. 3. HERA ring. Note PETRA accelerator at the lower left is used as an injector.

meters below ground level. Notice the PETRA accelerator which will act as an injector. As you can see, there are four interaction regions. Figure 4 shows the tunnel--the same tunnel size has been built for the subway in Hamburg. You see HERA in the tunnel. The young lady was the daughter of Zeus, the boss of the old Greek gods. It is said that all the successful

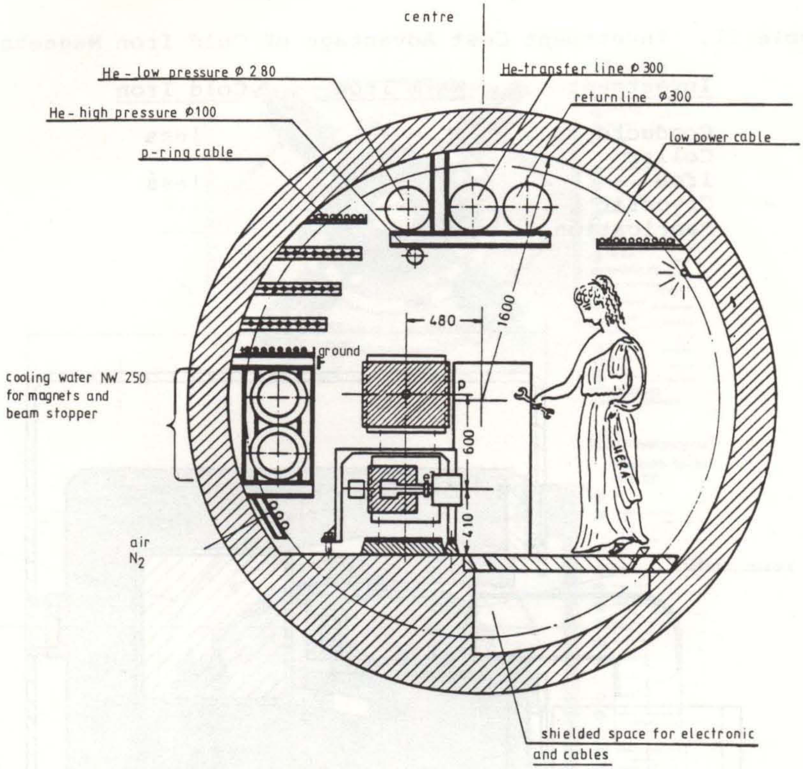

Fig. 4. HERA tunnel cross section.

superconducting projects in the past have had feminine names. As you can see there are two accelerators in one tunnel. Figure 5 shows a cross section of the electron beam magnet. It's a conventional magnet which has been designed on the same principle as the already existing electron accelerator magnet of PETRA.

Figure 6 illustrates the possible superconducting magnets. The upper design, which has been developed by DESY, is very similar to the Tevatron design. That was the basis for the first step of the project where DESY was convinced and knew from the experience at Fermilab that the magnet would work. This evidence was needed to convince the government that superconductivity would work and they could put money into it. This lower picture is our design, based on cold iron. I will not go into the details, but there are some advantages of the cold iron which are summarized in Table II. I will return to this a little bit

Table II. Investment Cost Advantage of Cold Iron Magnets.

Investment	Warm Iron	Cold Iron
Conductor	---	less
Collar	---	---
Iron	---	less
Cryostat	---	---
Fabrication	---	---

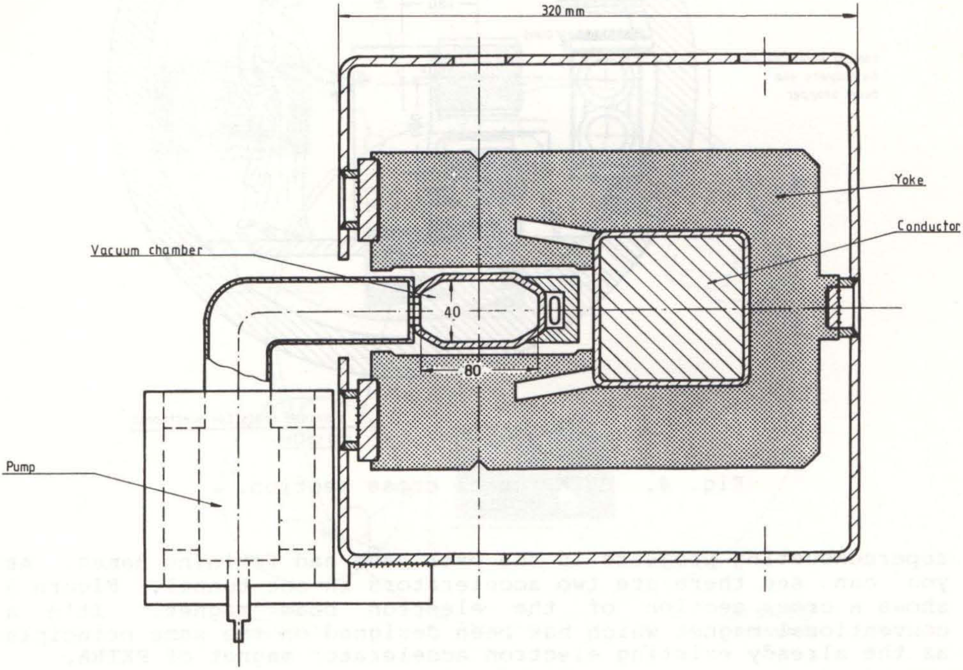

Fig. 5. Cross section of the electron magnet for HERA. This magnet uses normal conductor. Note the vacuum chamber and pump.

later. At some time there has to be a decision between the two designs because, in the end, only one type of magnet can be put into the tunnel. So we have been very lucky that a combination of both of these designs was found which minimizes the drawbacks of each and combines the advantages. This is the so-called hybrid magnet. It may be called hybrid because it comes from two institutions or perhaps it has parts of two different designs.

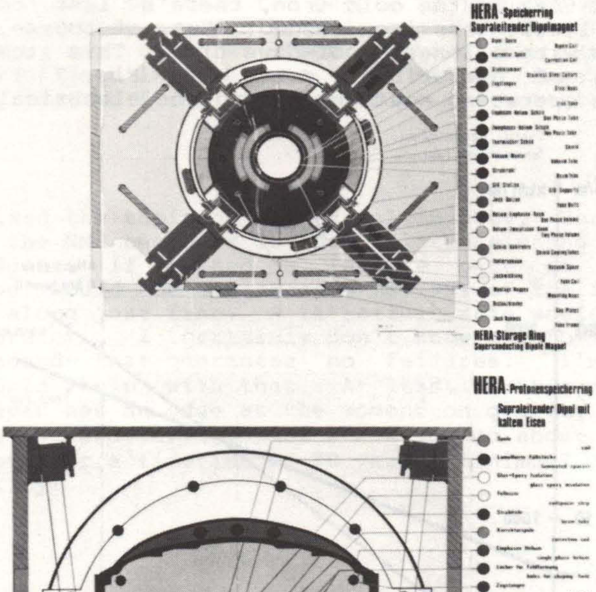

Fig. 6. Two superconducting magnet cross sections considered at HERA. The upper design is similar to the Tevatron. The lower design is based on cold iron.

The winding is inside aluminum collars installed in cold iron. This way the disadvantage of the cold iron is avoided, which is the nonlinearity of the fields. On the other hand, the lower heat losses which are possible with the cold iron are incorporated into the design. This is the magnet that will probably end up being used in HERA. We did some estimates to compare the net costs of the cold and warm iron magnet on the basis of the HERA design. The two designs have been worked out and can be combined for the benefit of the whole project. Now it's obvious that in the cold iron, there's less conductor and there's a little less iron needed. But, of course, you have to pay for this with a longer cool-down time. This item comes under the heading of operational cost. In Fig. 7 I have roughly compared the operating costs by taking the electrical power which has to be put into refrigeration. Of course, much more power is needed for the cooldown of the cold iron magnet whereas, a lower level electrical power is needed for continuous operation. After about nine months, you get the benefit of the cold iron. Now, you may say this depends on how often you have to warm up the ring, but I think the high-energy physics people would like to run the accelerator the entire year. If a magnet fails, one would have to warm up only a section of the ring, not the whole ring. So, I think this argument in principle remains valid. I just put this together to give you an example of how collaboration between industry and a laboratory can come to a solution which is a benefit to both of them.

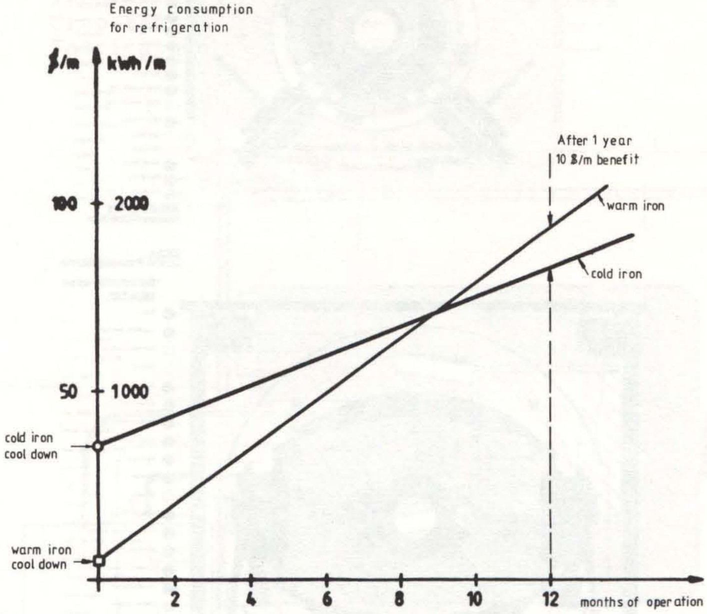

Fig. 7. Comparison of operational cost for warm and cold iron magnets in the case of the HERA magnets.

Let me just make three statements about the relationship between the laboratory and industry. It is very important to involve industry as early as possible. That means even during the design phase so that the people in industry know what the problems are and they can train their own people in the shop. Second, there are benefits from merging the knowledge and capabilities of both institutions, and, third, some competition is necessary just to force everyone to do the best they can do.

Lundy:

You raised the topic of magnet reliability. One of the easy aspects of the NMR business is that all those magnets operate as separate gadgets. If you had a failure rate of 1% for the magnets, you would be embarrassed but 99% of the installations are ticking along just fine. A 1% failure rate would be fatal to this accelerator. I certainly don't know how to write quality control standards that guarantee no failures. I'm sure that industry could help with that. At least in this country, it's felt that Japan has an edge at the moment on quality control, on zero defect manufacturing. How would you go about guaranteeing 15,000 magnets for a lifetime of 20 years minimum? What's the warranty policy?

Saito:

Instead of replying to your question, let me show you the Japanese status relative to Large Scale Science Projects (LSSP). At present, we have several Large Scale Science Projects in Japan. These include the construction of a large accelerator, studies for nuclear fusion systems, and the development of a new transportation system using a superconducting magnetically levitated train. The common feature of these LSSP is that while they are useful for humankind and science in the future, they are too advanced and too large. In the past, the scale for developing such a job was comparatively small. It could be carried out by the research people themselves, and the possibilities for industry to contribute were small. However, the recent trend for LSSP is for the scale to become larger and larger, more costly, and with correspondingly increased requirements on reliability. Under such conditions, the participation of industry is gradually increasing. In Japan, this tendency toward industrial participation was there from a comparatively early stage due to various circumstances in our country.

How does industry view LSSP? Strictly speaking, it seems not only attractive but risky. The plan itself is very beautiful. The personnel associated with it are wonderful. Often there is a great deal of money provided for the budget of the LSSP. So an LSSP should be attractive, but there is another element in the LSSP for industry. Industry earns profit by getting high productivity. The LSSP has some problem from this point of view. Typically, the specification of the LSSP is unique and usually difficult. The schedule is often demanding, but at the same time, trial and error is needed before the start of real manufacturing. Often an LSSP costs much more than was expected beforehand, not only by the planner, but also by industry itself. These costs are usually difficult to reevaluate. There is also a problem in that repeated production is rarely expected for the LSSP. However, industry does have a passion for work on an LSSP. There are rewards from the viewpoint of the status of the company and the spinoffs.

Industry has several needs that must be fulfilled to make it easy to participate in LSSP. Consider the case in Japan. There are two ways for industries to join such a project. The first possibility is for industry to act solely as a manufacturer. In this case, the scope of the responsibility and the specifications must be clear and acceptable to the industry. A good plan and design are needed to be sure that the industrial participation will be productive. The price must be reasonable and allow for necessary R&D and contingencies. The other possibility is for a more extended scope for industry in the LSSP. Industry itself has the abilities to carry out planning, engineering design, cost estimation, and scheduling. For some of these items, industry is rather professional. In order to have industrial contributions of industry in much more fundamental ways, the future of the project must be assured, to a certain degree, including the budget. The technical proposal on the cost estimate from industry must be well understood and reflected in the engineering or the budget. If not so, the latter case is not interesting for industry at all. In any case, communication between the planner and industry is very important from the early stage of the project. This early communication makes it easy for industry to participate in the project.

Finally, I'd like to discuss international collaboration. When industry participates in a project in a foreign country, there are some problems, especially for the LSSP. These relate to the status of each country. It is very desirable that the agreement be confirmed between the governments that are involved and that the division of work for each country is well established. The existence of a national research organization providing the appropriate coordination and advising (our own) industry is also desirable for us not only in the lead country but also in collaborating countries. Figure 8 illustrates schematically how such projects could be organized.

Finally, I hope each LSSP of the world, as well as the SSC in the United States, overcomes various barriers and blooms with beautiful flowers, and then gets fruitful results. I expect that Japanese industry can make many kinds of contributions to the LSSP as far as possible.

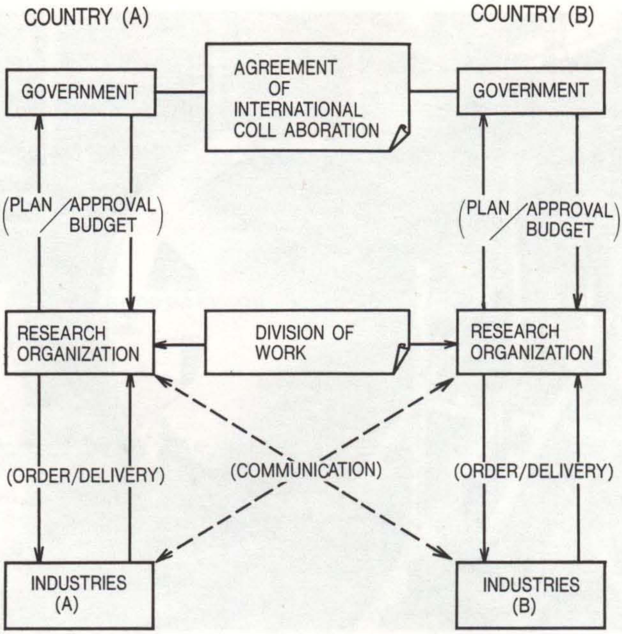

Fig. 8. Typical schema of organization for international collaboration.

As you know, international participation in particle physics has proved quite fruitful. Figure 9 is one such illustration, a picture of the Collider Detector solenoid being prepared by Hitachi for Fermilab. Another example is LCP coil Japan now already installed at Oak Ridge National Laboratory to develop the technology of nuclear fusion.

Lundy:

It comes as no real surprise that industry wants to get involved early in the SSC. However, it can't be that black and white. Some of you must have been buyers, not sellers. Who in the audience has had a bad experience with industry getting involved early?

Bob Tatro (Convair/GD):

Industry wants to get involved. But I raise the following question for the buyer. How do you let industry participate in an equal way so there is competition and do it in a time frame such that when the die is cast, people know what the game's going

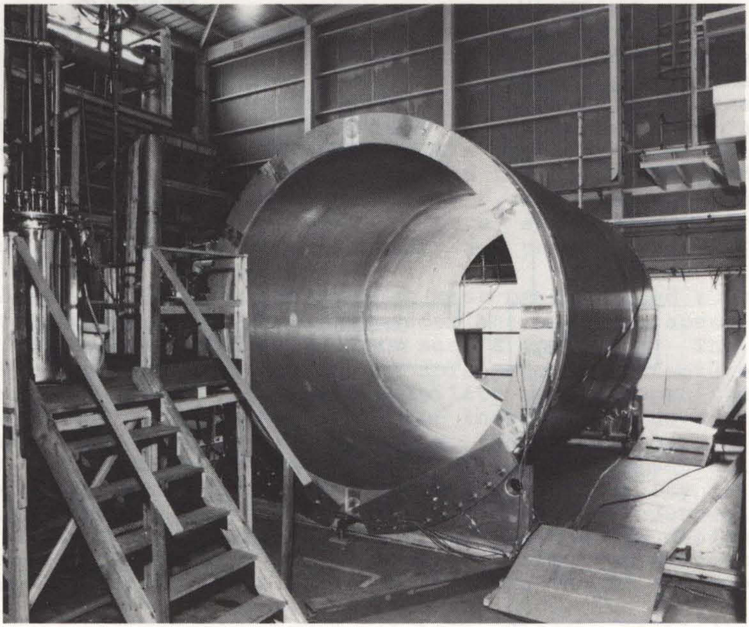

Fig. 9. Superconducting solenoid prepared by Hitachi for CDF at Fermilab.

to be, and all of industry is not strung along for years because everybody's not going to be a winner. There are going to be significant investments, contributions, and commitments by that industry and that company. How can the competition be structured so that we can get involved early, all of us that are interested, and yet the sands will sift and it will sift down to how it comes out. I've gone through this on the MFTF program now for four and one half years. I saw how that program evolved and how we were successful. I haven't heard all the stories on how the competitors felt it evolved. That does concern me because, although we have a very valid strong interest in the SSC, we want to know how we will get in and get out if we're not successful on the SSC.

Lundy:

Of course, you don't mean to imply that only one industry can be successful.

Tatro:

No, I'm not implying that. That's the point. I was going to ask Hitachi, Brown-Boveri, and the other people, how many other companies within their agencies or their countries are interested? How many will be brought along? Is it just one company from this point on in both cases?

Lundy:

Do you have a monopoly on magnets for HERA? A potential monopoly?

Dustmann:

I would hope so. The actual situation for HERA is this. HERA was funded only on the condition that there is European participation on the project. That means that different countries contribute to HERA and this contribution is in hardware. Different countries will deliver parts such as magnets, vacuum components, or other elements. These discussions have gone in the following direction. There is participation by France. They have developed the quadrupole magnets through their own knowledge and resources. There is discussion about Dutch participation. They are talking about making the correction coils. What we at Brown-Boveri don't like so much is that the Italians will also participate in dipole construction. It looks like there will be a division of suppliers, partly from Italy and partly from Germany. So, if you ask for the number of companies that may be involved in this business, it's between three and five in Europe, I would say.

Lundy:

My own thinking on this (and this doesn't represent DOE or even Fermilab policy) is that we have to take advantage of some of the facilities that are at the laboratories in order to compress the R&D timing. Say that at some point we finish a design that represents a laboratory's best shot at what would be a successful final magnet. The national labs commence to build these, somewhere between 10 and 100 units. At the very beginning of that process, you invite in teams from firms that are likely to be suppliers. I don't know how you'd limit it to a few or even what the right number is. Those people come to the laboratory probably at their own expense. After all, they've got to take some risk in this. They work alongside laboratory people and help build magnets. They learn all the good parts and all the bad parts. They keep their own counsel. They go home and make propositions for design changes, new methods, radical deviations, and somehow you evaluate and you select one, two, or three, and they get the job. The tricky part is selecting the necessarily limited number of initial participants. You can't entertain a hundred teams; you'll never make a magnet.

David Vroom (Raychem):

The question of how many companies can be involved goes back to how soon industry gets invited. The sooner you get industry involved, the more people you can have participating with relatively limited risk. A number of companies can be asked to submit early designs of their concepts. This will begin qualifying companies and perhaps uncover other design possibilities. On the other hand, if you wait until you get right down to the final bidding, of course it's going to be extremely difficult to choose who's going to build the magnets.

Lundy:

In fact, I think it's impossible because if we take the risk for that, that's over the transom engineering.

Patrick Stone (UOP):

SSC is probably five or maybe ten years late in involving industry in this arena. The time to involve industry in this arena is when the first idea is broached. I'm basing this on what has been recently called the justification for the federal role in anything (of course administrations change the way they look at this). But in science that's not the situation. In science, you're dealing with a non-proprietary but totally monopsonistic market. There is only one buyer for the SSC. And the last time I checked, the king never made mistakes. In fact, according to the government rules, he's not allowed to because that's the taxpayer's money and that would involve fraud. Anytime you enter a program where you're not allowed to make mistakes, you've already made one.

Now this is not meant to chide our leadership in any form. I've been in the aerospace market myself for the better part of 30 years, both as a technician and a marketer, and I've bled over it. But, if someone feels they have the vast knowledge necessary to provide leadership in a project as expensive, as involved, and as obvious as this particular one is, we ought to go back and let him throw the first stone. You can't do this by committee. Ferdinand Porsche had a great quotation. He said that there's never been a winning race car designed by committee. So the first thing you have to do is pick the leadership and this time you picked Fermilab. Now, the less Fermilab does other than lead, the greater the likelihood they'll come up with the best solution. If you want to involve industry, you start with the bottom of the pyramid--lots of folks--and you do it as early as possible. And in the early parts of the program, when you're talking about ideas, you very rapidly establish those people with whom you can deal in both a trustworthy and a competent way at very little expense. By the time you get up to where you are today, you have, in fact, established the club. And those who are competent to propose know who they are. You'd better put it on the street and you better put it on the street early and you better put it on the street when it's just ideas. This should happen even if it's only holding a conference, to say, hey, we're thinking about building a huge, new accelerator and get industry involved right there because the presumption of perfection is in itself the basic error.

Bob Remsbottom (Wisconsin):

One question that needs to be addressed is how do we stay away from a repeat of the Large Coil Project (LCP), where there's a large number of coils being made by a bunch of different people. Industry is involved in it from the start, but will it ever go on line? I don't think that what we're doing on the SSC could ever survive something like that.

Beuligmann:

In answer to that, the LCP will go on line. It may not go on line with all the magnets. There are some questions about one or so, at least at full current. LCP is a technology program. That's different than what we're talking about here. First of its kind, industry had never built anything like that and so there were problems. It should not be put in the same context as other magnets that have been built that are slightly different. The PMS-F magnet has been running three years now without a hitch. The MFTF Ying Yang's have been tested. Large coil is a different program, and I would be glad to go into the constraints that drove the technology and some of the excuses (some of them not so good).

Sure, we were naive at the start. There are some other problems within the industry of lack of commitment. That can be embarrassing to the whole superconducting magnet industry. But there won't be a repeat. I don't know if I've answered all those questions. I think it's pretty obvious that there is an industry out there now. We have been working on SSC for a year helping one of the labs do some work leading to production. Still, we are not working in the heart of the system. That's inherent to the work here that Bob Remsbottom, Dick Lundy, and the others have been doing at Fermilab and Clyde Taylor and others at other laboratories. We have been working on that problem for over a year now. Anyone who thinks that they're going to come on line and reinvent a wheel just doesn't understand where SSC is today and where it can go. We'll build those magnets for SSC. By we, I mean the whole business and the industry. I have no doubt that SSC can be made with the concepts that are being talked about here. It is not an infant industry anymore. I wouldn't say it's mature, but it sure isn't an infant.

Convair-General Dynamics is up in the order of 60 to 70 million dollars worth of business just in superconducting magnets. That's a lot of broad skills and experience and some of them were built on the magnetic skills that existed in the laboratories. We complement the set of skills at the laboratories, and we don't intend to replace the skill that exists in the laboratories. That is a very inherent ingredient to the success of SSC. I disagree with John Hulm. I don't think you allow industry to go out and reinvent some wheels and then propose. There are many man-years of knowledge and skill existing right here in the team at Fermilab. The same holds true at some other laboratories. I've got a lot of respect for industry, but I don't propose that even with our background we're going to go out and reinvent that wheel and get a better idea. We can find improvements to it. We can complement it, but we cannot reinvent it. Have I answered your question?

Remsbottom:

Basically, yes. You can look at General Dynamics as being successful on LCP. Here you have a very broad industry, many people involved and so forth. If we have a thousand magnets sitting out here it wouldn't be very good if one of them blew.

Lundy:

There are two comments that I haven't heard, and I'm surprised and I want to throw them out to see why I haven't heard them. One was triggered by the mention of Porsche. One of the reasons that Porsche goes in for auto racing, besides the sheer fun of it and the advertising value, is that they believe by participating in racing, they're able to attract engineers who also do the passenger car work that are much better than they would be able to hire otherwise. I would think that the project we're talking about here has enough sex appeal that it's a recruitment aid or a morale builder or a source of adrenalin for a firm. Is it a project that will get a company's adrenalin going?

Mike Morgan (Meyer Tool):

The laboratories have the expertise. They have to define the magnet and the magnetic field properties. And oftentimes, from the past experiences I've had in working at the labs, the requirements are not always realistic. Industry is now in the position to point that out to you. It makes sense for the laboratory to design the magnet and to go to industry, who have the background and experience and capabilities for putting things like this into production. The laboratories will have to show industry how they've done it, and explain what they will accept. They can't go to industry with requirements that are not attainable.

Lundy:

I understand that. If you, in fact, elect the route--build to print, you've got to prove those prints are good by building some.

Hulm:

Or come back later and change the design, piece by piece, as you go along through the manufacture.

Lundy:

Of course, the classical reason for getting industry involved is that they will economize, find cleverer ways to do the job.

Morgan:

There's another aspect that I think is important. I've been on both sides of the fence. The laboratories must come to a cognizance of the cost of making changes. This is in light of an earlier comment about unscrupulous companies taking the job on a low bid and making it up on change orders. I'm sure that that does happen. But by the same token, changes are extremely expensive. If you have good documentation and look at what it costs to do something, you sometimes scratch your head in bewilderment. If you had to estimate what that change order would cost, you wouldn't believe it. And you don't believe it after you find out what it did cost. And the people at the laboratories look at it and say, hey, you guys really stick it to us. But we haven't.

Lundy:

I couldn't agree with you more. One of my own biggest problems during the Energy Doubler was to prevent changes. Naturally we had lots of ideas, but we knew that if we incorporated them that we would have never produced magnets.

There would be 990 different dipoles. Of course, that takes some of the fun out of it. If you're going to have to make 15,000 magnets all looking alike, it's going to get boring. Be prepared to face that.

Rosner:

I'd like to make a point that is now a little bit on the other side of the fence from my earlier observations. There was some good judgment and rationale why Fermilab did as much as it did in-house. There were many changes in the early periods, because a superconducting accelerator was an evolving concept using new and evolving technology. As an old magnet designer and builder, I really appreciate what Fermilab has done in designing and building these magnets. It is a fantastic achievement to have every single magnet in a ring of 800 or a 1,000 magnets work. Most superconducting magnets operate in the dc mode, but Fermilab encounters the most demanding application, namely, pulsed operation where you have to worry about cycle fatigue and shorts and who knows what. I don't want to detract from what Fermilab has done and the way it's gone about it. The fact of the matter is that it was successful and that's a real tribute to the way they went about it. The question that I was trying to address is how can that experience and that accomplishment be translated and transferred to industry thus allowing us to go on to bigger and more productive projects. HERA couldn't have done what they're planning to do without the Fermilab experience. That's where the benefit from the Saver experience has gone. When you try to bring it closer to home, unfortunately I have a hard time seeing the benefits to U.S. industry at this juncture. For the future, that's what I hope will come out of the SSC experience.

Lundy:

Thank you for your compliments. At Fermilab we also believe that it was the right way to build the Doubler. That doesn't mean it's the right way to repeat the experience, and that's part of the reason for this discussion. I was going to ask Leon Lederman to put his hands over his ears so I could say how to transfer the technology--hire all the smart people.

Lederman:

I saw John Hulm nodding at the last comment and yet he has been rightly pointing out that collaboration in the construction of such a machine would be worthy for a lot of other reasons. Do you see a possible conflict between these two points of view?

Hulm:

No, but I certainly would agree with what Rosner said. The Fermilab Energy Doubler is an amazing engineering achievement. Fermilab deserves to be congratulated, but I also agree with Rosner in that if you organized the SSC work along the same lines and continue in the same way, the technology will remain buried in the national laboratories.

Lederman:

I'm sorry. I wasn't clear. My question has to do with international collaboration. I'm raising a very delicate point. I'd love to have a full and frank discussion. You have raised the issue and others have raised the issue of the importance of international collaboration in constructing the machine. I resonate very strongly with this idea. But I'm now asking you how this is consistent with the other virtues of SSC as direct benefit to U.S. industry? You don't have to answer that if you don't want to.

Hulm:

It's a very difficult question, of course. I assume that the other countries that might be involved in such a collaboration--Japan, Germany, Switzerland, France--would all want to get their industries the same kind of benefits that we would hope to get for U.S. industry. I would hope they would act in some kind of competitive mode in the procurement. The Japanese do this all the time. Almost all of their major projects involve several companies, at least in the first stage, and they try to get the best ideas from the companies and then somebody wins the follow-on project. Of course, we do this in many other areas. I think it's very dangerous if any kind of monopolistic situation results from the SSC. Something would have to be built into the agreements that we would have with our international friends.

Lederman:

You don't see any problems with Mr. Saito's model?

Hulm:

Not basically.

Wroton:

I'm so far down on the learning curve on SSC that I hesitate to speak, but I got a strong impression from some of the earlier comments that industry cannot support changes, that build-to-print is even a conceivable point of view in this kind of project. I'm not at all certain that it is. In fact, there are many ways to contract R&D, some of them very suitable for rational changes. As was noted already, the ability to hold changes to a desirable minimum is part of that rationality.

An example is the Viking Lander spacecraft sent to Mars, with a couple dozen experiments, an unknown planet, a year in transit, an unknown atmosphere, an unknown surface structure. Finally, the spacecraft had to be sterilized for 48 hours before takeoff with almost no testing following the sterilization. You're all familiar with the pictures from cameras and the negative results concerning life on Mars from the biologists.

That was a performance contract. The contract was very simple. It said, go to Mars and take data successfully. It was an incentivized contract which had as its final carrot some fifteen million dollars of incentive that would interest almost any corporate president or anybody under him. That amounted to about 3 or 3-1/2 per cent of the value of the contract. The total contract was about 400 million for Martin-Marietta's part which was to provide the lander. And that was totally successful. We got 100% of that award in the end. It took seven years to do that job, during which time the project operated against a countdown schedule which says 1,022 days to launch or 368 days to launch or what have you. We had an absolutely definite window during which we had to complete that job and get it off, as well as providing the entire design of the mission and the support of the spacecraft and the scientific team associated with it. Now, of course, there were lots of changes in a program like that. When we began, we didn't know what you do to accomplish that kind of a task except put in large contingencies, which was not an acceptable option. All those disciplines that we built into the Viking lander and which have been developed in the aerospace industry and in other high-tech industries are usable in the SSC application. Martin-Marietta is now the operator of the Oak Ridge Laboratories. There is a major incentive on that contract, although I'm not personally very familiar with it. My point is I don't believe that the fact that there will be developments after the initiation of the build process needs to be an overwhelming concern as to whether these types of products can be built in industry or must be built in a government laboratory.

Lundy:

Thank you for your comment, because in addition to being worthwhile, it reminded me of a question that I wanted to ask. In a recent article in Fortune, the thesis was put forward that the biggest challenge right now for U.S. industry is to learn to respond to change more rapidly. This is not because change in itself is a desirable end, but because, we find that the life cycle from product introduction to the profitability phase to obsolescence is getting shorter and shorter. It's approaching months instead of years. This is particularly true for computers and electronics and some consumer goods. A company has to be able to get with it and produce and get out and get on to the next boom, whatever it is. The times have to be shortened, the flexibility has to be increased. Some companies do that by building a skunk works so that the smaller operation can be more dynamic. Presumably, whatever you do to learn to respond to these accelerated product life cycles, also enhances your ability to respond to changes with minimum friction. It may be that corporations may have to learn to respond more quickly to the customer, be it a monopsonistic customer like a laboratory or to the millions of consumers out there. I think the U.S. auto industry, properly goaded by Japan and by Western Europe, has been able to respond more rapidly and had to in order to survive. They had lost touch not only with the customer but also with the rate of change of the customer's wishes.

Beuligmann:

I think we're talking about two kinds of changes. Some changes are inevitable. You can have something going in production and you can no longer get that semiconductor or part. This happens all the time. There are 50 changes a month for a program that I know about in the Pomona division making missiles, typically because of items no longer being available. You can't get rid of changes but you have to minimize them. However, when you've got not one Viking lander that you're making, but thousands of magnets coming down the pipeline, you've got to watch those changes. In particular, you can't make changes readily in a contract where you say, hey, industry take this fixed price or put these tough incentives on and then see changes flowing freely. You're going to end up with big teams in negotiations all the time trying to figure out how much the change impacted. Of course there will be changes. I think the changes Lundy's talking about are in the rapidity of getting marketing sense of what the consumer wants and implementing the change.

To reiterate, there are two different kinds of changes here. In the R&D phase you need changes, flexibility within the government, and fluidity between the laboratory-and-industry teams. But when you hit production, the outlook should be that the system is good and really cost effective, and that we're not going to play with it willy-nilly. At Convair/General Dynamics, we've had troubles with some of the laboratory programs in superconductivity. We've had to sit hard on people who wanted to make changes who came into the program late at Livermore. Essentially they've wanted to just do it their way. It wasn't better, it was just a different way. That is disaster for a program. Luckily the program managers up there understood that and took care of the problem. But there are always some engineers who want to do it a different way because it's their way, and they're more comfortable with it. They don't pay the bill.

Tatro:

Since January 1981 when we started the contractual effort (at Convair/General Dynamics for the MFTF-B), there have only been five contractual changes to the program, every one of them initiated by configuration changes at Livermore. We've had a cost plus incentive contract. We recognized at the beginning from our discussions with Livermore that they wanted us to be flexible, they knew things were going to change, and they didn't want to be dogged to death by contractual changes on our part. We scoped and priced the project initially, knowing we'd have to handle some changes. I assure you the magnets were not defined with performance specifications. In fact, they were not well defined. We are within the contracted budget through the entire program on every single type of magnet, the solenoids, the axicell, the transition coils, the high field insert coils, and over 3,000 thermal shields with 800 different configurations. In each of those areas many changes have been made, but we have done the design and fabrication and we are on budget and on schedule. That's the kind of situation that must be accommodated. The contract has to be that way but there is a strong program requirement to recognize there are people within the laboratories and the government, as well as within industry, that are going to want to do something that is just different. That attitude has to be stopped and stopped quickly when it develops.

Lundy:

To summarize, I want to ask Ed Temple if he's profited from this discussion and if he's figured out what the R&D director's going to do about all this. Perhaps his answer will be, "You haven't quite solved all the problems, but don't worry because I'm sure there are going to be many more meetings like this." It sounds like there should be. I apologize that this one is seven or even ten years too late. I didn't even think about superconductors at that time, thank goodness.

Temple:

The SSC as we see it now, is four to five times larger than the effort for the creation of this laboratory (when compared in comparable dollars). The number of superconducting magnets that we're going to build is something like 10 times the number that we did on the Doubler. Given that situation, it's just a fact that we have to use industry in a very big way to help us do this. Now the best way to do that is not known by me. I showed some cost figures early on that are the very foremost consideration in our mind. We have a situation where we can't make errors.

On Lundy's scale of industrial participation, the Doubler was something like 2.5. Remember the Doubler wasn't just buying up some bolts, but did involve contracting some pieces. A nine or ten on Lundy's scale is something like buying a Van de Graaf from the High Voltage Engineering Company. There are some other recent projects, such as TFTR where industry was involved in a big way building devices that had been designed by either Princeton or industry. I would put TFTR around four or five on the Lundy scale. MFTF took a different tack and in very many of their large systems, put out performance specs. The most outstanding performance in this area was in their vacuum chamber system, the huge tank that housed the Ying-Yang coils. It included the tank plus all of the normal vacuum pumping system for that tank plus all the cryogenic system. (By the way, it was another quote on an MFTF refrigerator that we used in the SSC cost estimate, so at least within the Department of Energy, we do have some transfer of knowledge and information. That's hard, as well, sometimes.) That contract was put out at about 30 million dollars and they finished it with less than a 5% increase. It was a performance spec and they had some changes along the way. I think that's fantastic performance. We've heard how they've done a major set of the solenoidal superconducting coils and some of the axicell coils. I would put MFTF maybe on a scale of six to seven.

Now, for those of you here, we have two projects that I think you might be semi-interested in at this time. They are CEBAF and the SSC. For different reasons, I think they're both going to be significantly higher on the Lundy scale than the Doubler. CEBAF is not going to be able to put together an organization fast enough to do all of their own work and they're going to have to do some awfully big pieces by putting them out to industry. For the SSC, as I noted, we don't have the manpower in our labs to do it, and we couldn't keep them there if we did. In any case, it's probably not a very effective way to do things. In my estimation, a well-managed project will probably end up in the four-six range. It would be helpful to have your thoughts on what you would like to do, what you have done well, and how you think we can sort out some of these questions. This should be helpful to the labs, to URA, and to us at DOE in actually getting our plans together.

I would like to note that the schedules that I showed earlier were basically those put together by the not so infallible planning bureaucracy about four months ago. Based on the feasibility study effort, the reference design study effort, and especially the conventional facilities work that the architects from Parson Brinckerhoff did, we have to get moving on the SSC now if we really plan to do it in the overall time frame that I showed. That means that if we really do this project in the time scale that has been outlined, we're going to have to get some changes in our planning way up front. From that standpoint, this meeting happened none too soon. I hope that there will be exciting and continued interaction amongst you all and the participants in the lead contracting groups in the SSC over the next year.

Lundy:

Thank you, Ed and thanks to the rest of the panelists. As the chairman, I declare this session formally closed.

FERMILAB INDUSTRIAL AFFILIATES

The Fermilab Industrial Affiliates organization was established in 1980 to improve university-industry research communications and to foster technology transfer from Fermilab. By now the Affiliates number more than 30 institutions including many research-oriented companies in the Fortune 500 list as well as several companies formed by Fermilab staff members and users.

Direct activities of the Affiliates include visits of company representatives to Fermilab and Fermilab personnel to Affiliates. The annual meeting is one of the principal opportunities for such visits. This Round Table was presented at the fourth annual meeting. At the meeting, the visitors are given a comprehensive presentation of the activities underway at the Laboratory. Tours and individual conferences present an opportunity to see the Fermilab work in detail. Affiliate members have direct access to Fermilab staff for information on the work at the Laboratory. They receive copies of significant Fermilab technical publications and are kept abreast of important seminars on technical matters at the Laboratory.

Specific technology innovations are only one facet of the work of the Laboratory that is emphasized. The "scientific culture" related to particle physics is given heavy weight as well as the long-range potential of activities such as the development of superconductivity technology. The participation of more than a hundred universities in all phases of the Laboratory is also important to Affiliate members. Often, an Affiliate's interests in the Laboratory are hard to gauge. A major farm equipment manufacturer turned out to be one of the heaviest users of large computers in the United States.

Fermi National Accelerator Laboratory, by its nature, amalgamates a wide array of engineering and physics disciplines of interest to Affiliate members and industry in general. The acceleration of particles requires a working together of systems of high voltage electrostatics, high power radiofrequency signals, and rapidly pulsed magnets all under rigid and precise computer control. Beam optics, high vacuum techniques, ion sourcery are also involved. Particle detection adds new areas in terms of spatial and temporal resolution, fast logic circuitry and decision making, techniques of multi-dimensional pattern recognition, signal processing, and efficient number crunching.

A seven year R&D effort in superconductivity has culminated in the construction and operation of a four-mile ring of superconducting magnets with associated cryogenic systems. A substantial fraction of the world helium refrigeration capacity is at Fermilab. Advanced R&D looks to new materials, refrigeration, and understanding which will lead to pulsed magnets operating with magnetic fields of greater than 100 kilogauss.

Since the scale of Fermilab is so large, four miles of tunnel filled with sophisticated magnets and a 6800 acre site, an important ingredient in much of the R&D has been a search for innovative, cost-conscious designs. Special fabrication techniques such as laminar tooling have been invented in pursuit of precision coupled with economy. Remote and autonomous control is important for the same reason. This has led to important developments in large-scale distributed control and data-collection systems.

Technology-related programs at Fermilab include holography, solar energy, and neutron cancer therapy. A brochure is available (Technology Development at Fermilab) with capsule descriptions of all these activities.

For information on the Affiliates contact:

Dr. Richard A. Carrigan, Jr.
Assistant Head of the Research Division

or

Dr. Leon M. Lederman
Director

Fermi National Accelerator Laboratory
P. O. Box 500
Batavia, IL 60510
(312) 840-3333

LIST OF FERMILAB INDUSTRIAL AFFILIATES

Arco Petroleum Products
Bell Laboratories
Borg-Warner Corporation
Brunswick Corporation
Cherry Electrical Products Corporation
CMD Development
Commonwealth Edison
Deere & Company
Digital Equipment Corporation
Digital Pathways, Inc.
Eaton Corporation
FMC Corporation
General Electrical Corporation
The Goodyear Tire & Rubber Company
W.W. Grainger, Inc.
The Harshaw Chemical Company
Hewlett Packard
State of Illinois
International Business Machines Corporation
Johnson & Johnson
Kinetic Systems Corporation
Litton Industries
McGraw-Edison Company
Nalco Chemical Company
New England Electric Wire
Nuclear Data, Inc.
NYCB Real-Time Computing, Inc.
Omnibyte Corporation
Raychem Corporation
Richardson Electronics
Shell Development Company
Standard Oil Company (Indiana)
Sunbeam Appliance Company
Union Carbide Corporation
UOP, Inc.
Varian Associates
Westinghouse Electric Corporation

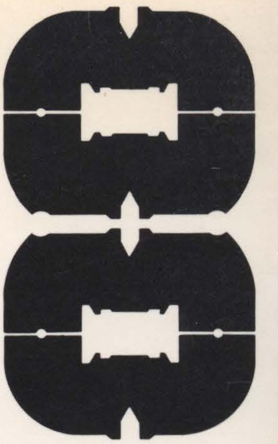

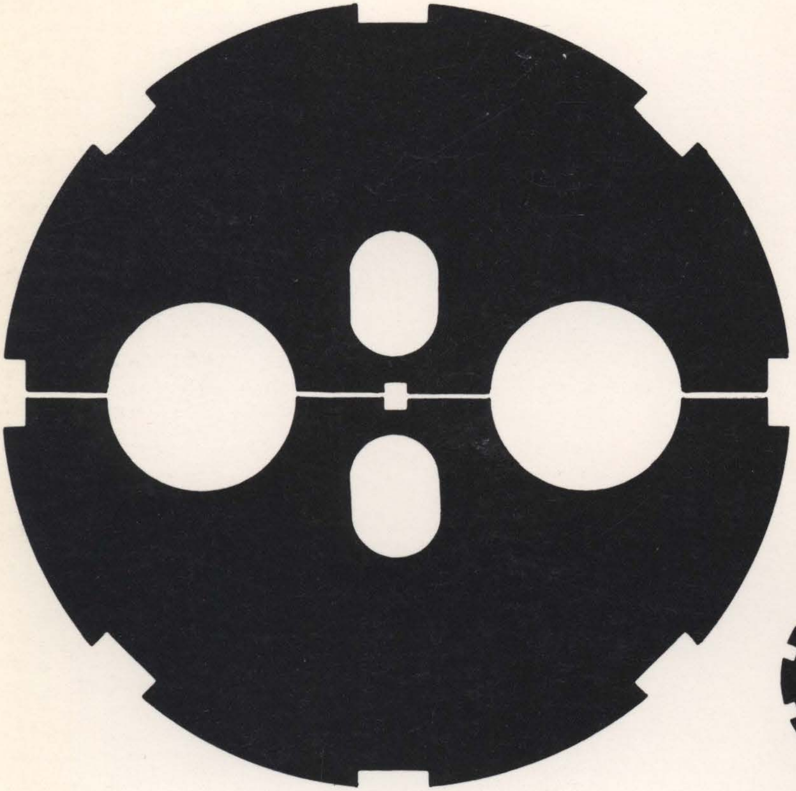

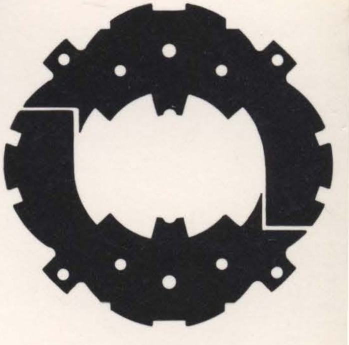

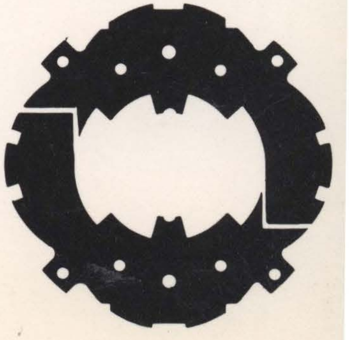

Fermilab
Industrial Affiliates
P.O. Box 500
Batavia, Illinois 60510