

Future Proton and Electron Colliders: Dreams for the 1990's*

Burton Richter

Stanford Linear Accelerator Center,
Stanford University, Stanford, California 94309

I. INTRODUCTION

We are nearing the end of a great cycle of accelerator building. New machines have been recently completed or are nearing completion in China, Japan, the U.S. and Western Europe. These new facilities include (energies are center-of-mass):

- the 5.0 GeV BEPC e^+e^- collider in Beijing
- the 60 GeV Tristan e^+e^- collider at KEK
- the 100 GeV SLC e^+e^- linear collider at SLAC
- the 2.0 TeV P-P collider at Fermilab
- the 100 GeV LEP storage ring at CERN
- the 350 GeV E-P HERA collider at DESY.

With these accelerators, the experimenters and theorists will have a great opportunity to test the limits of the standard model and, we all hope, get a first look at what lies beyond the standard model.

Given the roughly ten year time cycle from the initiation of conceptual design to the completion of the new high energy accelerators, it is clear that the facilities listed above will be the principal new tools for high energy physics research through the mid-90's. My charge from the organizers of this conference is to tell you what might lie beyond that. I will fulfill that charge by giving you a summary of the thinking around the world on new machines for the '90s, restricted to three areas — proton colliders, electron-positron colliders, and B meson factories.

2. PROTON COLLIDERS

The present and the next generation of proton colliding beam machines build on two great pioneering efforts: that of the ISR group at CERN who built and brought into operation the first proton collider, and that of Robert R. Wilson and his colleagues at Fermilab, who made superconducting quality accelerator magnets a practical technology.

We now have two operating proton colliders. These are the $SppS$ at CERN which runs at a center-of-mass energy of 0.6 TeV and the Tevatron Collider at Fermilab running at a center-of-mass energy of two TeV. Both machines run in the luminosity range of $1 - 2 \times 10^{30} \text{ cm}^{-2}\text{s}^{-1}$. The next generation of machines will be titanic engineering tasks involving large extrapolation of a basically known technology. The machines will be hard to use, because of the complexity of the final states, causing a great challenge to the experimenters to design detectors and computer algorithms to extract what we hope is the simplicity hiding in this complexity. We all know the problem. The proton is a composite particle, and what really interests us is the hard collisions of the constituents of the proton. While the cross section for proton-proton collisions is large, the cross section for hard proton-proton collisions is small, and there is a great deal of debris in the final state accompanying the particles of interest. The experimenters problems are illustrated in Fig. 1 which is a Monte Carlo simulation of a two-jet event at 40 TeV in the center-of-mass.

Three new proton colliders are in various stages of the design and approval process. These are the UNK Collider at Serpukhov in the USSR which is an extension to the three TeV proton synchrotron now under construction at Serpukhov; the LHC which might be added to the LEP tunnel at CERN; and the SSC, the largest of them all to be built and at an as yet undetermined site in the U.S. We can compare the capabilities of these three possible new facilities and the two existing proton colliders in terms of a somewhat fuzzy and process dependent concept called "mass reach." This notion combines the given facility's energy and

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




Fig. 1. Monte Carlo Simulations of a two-jet event in a 40 TeV proton-proton collision.

luminosity with theoretical estimates of cross sections and backgrounds to give the maximum mass one might be able to both produce and detect at a particular facility. The table below gives the mass reach for the two existing and three projected facilities.

Table I. A comparison of proton colliders.

Facility	Mass Reach (GeV)
<i>SppS</i>	150
TeV II	300
UNK	600
LHC	1500
SSC	2500

This table is an attempt to crudely compare different machines, and should not be taken terribly seriously except in a relative kind of way. "Mass reach" is a well defined notion for processes like quark jet production (which is where the numbers in the table above come from), and is much more dependent on the experimental assumptions and somewhat questionable background calculations for such things as Higgs production.

I now turn to the three new facilities that are under discussion. I want to take this opportunity to thank Victor Yarba at Serpukhov for the information on UNK, Giorgio Brianti of CERN for the information on LHC, and Chris Quigg of the SSC Central Design Group for the information on SSC.

UNK Collider

The UNK Collider is an as yet unapproved addition to the three TeV Superconducting Proton Synchrotron now under construction at Serpukhov in the USSR. The Synchrotron is scheduled for completion in 1992 or 1993. It uses the existing 70 GeV Serpukhov machine as an injector into a 400 GeV conventional magnet booster which in turn serves as the injector into a three TeV superconducting final accelerator. Both the 400 GeV booster, and the five Tesla

superconducting main ring are contained in a 20.70 kilometer circumference tunnel which has a 5.1 meter bore. The tunnel is about half complete, and the entire project is on schedule for first operation as a fixed target machine in 1993.

The superconducting dipole magnets (five Tesla peak field) are a version of the by now standard HERA modification of the basic Fermilab design. They use a two layer coil compressed with non-magnetic collars which in turn are contained in a magnetic iron tube, the magnetic iron being at a sufficient distance from the high field coil to effectively increase the field obtained per ampere in the coil without at the same time distorting the dipole field from saturation of the iron. The entire coil/collar/iron assembly is operated at liquid helium temperature.

The colliding beam phase of UNK will follow on the completion of the three TeV fixed target facility. Its schedule is not yet firm. The Serpukhov group has considered both proton-proton and proton-antiproton colliders. In their minds, consideration of simplicity, reliability, and luminosity lead to the choice of proton-proton instead of proton-antiproton. Because of the large bore of the UNK tunnel (5.1 meters) there is no difficulty in putting a second superconducting proton ring in the same tunnel with the first.

Figure 2 shows the proposed layout of the collider facility. The beam crossings are in the horizontal plane and four interaction regions are provided for experiments. The necessary cryogenics and power for the second ring are being provided as part of the first phase of the project which will simplify the installation of the second ring when that work begins.




Fig. 2. Schematic of the UNK 6.0 TeV proton-proton collider.

The main parameters of the collider are summarized in Table II.

I would characterize this machine as an extremely conservative design. There is a significant potential for lower interaction point β , smaller crossing angle, and higher beam-beam tune shift; all of which could potentially give a larger luminosity. Experimenters will also appreciate the relatively low number (compared to the other future colliders) of events per beam-beam collision, which will make detector problems somewhat easier. The project requires no new technical developments and should be able to be completed relatively rapidly, particularly if approval is given by the Soviet government to continue magnet production after the magnets for the fixed target machine are completed.

SSC

The largest of the future projects is the Superconducting Super Collider (SSC) designed to reach an energy in the proton-proton center of mass of 40 TeV with a maximum luminosity of about $10^{33} \text{ cm}^{-2}\text{s}^{-1}$. The machine is to be housed in an 84 kilometer circumference tunnel which will incorporate two 20 TeV proton storage rings in an over and under configuration. The proposed facility is shown schematically in Fig. 3.

Table II. Main parameters of the UNK P-P collider.

Energy (TeV)	3×3
Protons per bunch	3×10^{10}
Number of bunches	8000
β^* (meters)	1
$\mathcal{L}(\text{cm}^{-2}\text{s}^{-1})$	4×10^{32}
Events per collision	0.35
σ_L (meters)	0.10
Crossing Angle (mrad)	1.0
Tune shift	6×10^{-4}




Fig. 3. Schematic of the SSC 40 TeV proton-proton collider.

The configuration of the machine is unusual in that it consists of two long arcs of magnets connecting two clusters of interaction regions, each of which contains four potential crossing points. Since the luminosity in the interaction regions in each cluster are not all the same, the machine actually has a superperiodicity of one. While the machine looks different from existing facilities, there is really not much difference from the beam dynamics point of view between this facility, and, for example, the $SppS$ or the Tevatron Collider. These machines have a high degree of symmetry, but the asymmetric configuration of the interaction regions actually reduces the symmetry to an effective superperiodicity of one.

An intense R&D program has been underway for more than three years. The design of the facility has undergone considerable refinement since the first conceptional design report, and much work has been done on such things as beam dynamics, interaction region configurations, experimental hall design, requirements for experiments, cryogenics, conventional facilities, etc. However, the bulk of the R&D program has concentrated on the development of the superconducting dipole magnets for this facility. These magnets are also of the collared coil-cold iron design with a peak field of 6.5 Tesla, a bore tube diameter of 4 centimeters, and a length per dipole of 17 meters. The magnet length was chosen on an economic basis. The

longer the magnets, the more difficult they are to build and to transport while the cost per unit length of magnet decreases because of fewer complex magnet ends and fewer interconnections between magnets. While no difficulties were experienced in building short magnet models that met specifications, the first few full length magnets had considerable trouble with erratic quench behavior and did not reliably achieve full field. These problems seem to have been solved in the R&D program, and the most recent full length magnets reach full field with very few quenches.

The SSC is the first proton machine to be designed at an energy sufficiently high to have to take into account the effects of synchrotron radiation on machine performance. Figure 4 shows the effects of synchrotron radiation on the beam size and on the luminosity. The lower curve shows the shrinking of the emittance of the beam from synchrotron radiation. The middle curve in shows the decrease in the number of circulating beam particles coming from all the loss mechanisms in the machine. The upper curve shows the luminosity which actually rises during the initial day of the fill, for the decrease in transverse emittance dominates the loss of particles. The synchrotron radiation should also help to stabilize the beam against various slowly growing beam instabilities while at the same time it generates significant amounts of power which will have to be handled by the cryogenic system.




Fig. 4. Effect of synchrotron radiation on the SSC. The lower curve shows the effect of synchrotron radiation damping on the emittance of the stored beam. The middle curve shows the decrease of the circulating beam intensity from all sources. The upper curve shows the combined effect of radiation damping and beam loss on the luminosity.

Table III shows the main parameters of the SSC at a high luminosity collision point at the beginning of the fill before synchrotron radiation affects the beam size. The mean number of events per bunch collision in the SSC is 1.7. This will pose some increased problems for the experimenters beyond those experiencing the existing proton colliders, though it is claimed that these difficulties are not great for a properly designed detector.

Table III. Main parameters of the SSC at a high luminosity collision point.

Energy (TeV)	20×20
Protons per bunch	7.3×10^9
Number of bunches	17,100
β^* (meters)	0.5
$\mathcal{L}(\text{cm}^{-2}\text{s}^{-1})$	1×10^{33}
Events per collision	1.7
Tune shift	0.001

The SSC is not yet a fully approved project. The present administration has requested construction authorization, including a large first-year allocation of construction funding. The

U.S. Congress did not approve the construction request, but did approve \$100 million for R&D in FY89. This allocation of funds specifically prohibited start of construction, but did allow such as things as R&D, industrialization, detailed engineering, and conventional facility design. The decision is really up to the next President and the next Congress. If the project goes as planned, construction will be completed in 1996 at a total cost for the machine of \$3.1 billion (1988 dollars).

Site selection for the SSC is in progress. A large number of site proposals were submitted in the Spring of 1988 and the U.S. Department of Energy turned over the screening of these proposals to a special panel of the National Academy of Sciences and the National Academy of Engineering. This screening procedure resulted in the selection of a subset of proposals that fully met the criteria for the project. The proposals still in contention, going from west to east, are Arizona, Colorado, Texas, Illinois, Tennessee, Michigan and North Carolina. The final site selection is to be made late in 1988. It will not be known until around the summer of 1989 whether the project will or will not proceed on schedule.

LEP Hadron Collider

Ever since the earliest days of the LEP project there has been some discussion at CERN of the possibility of adding a proton machine to the LEP tunnel. While these discussions were very casual at first, in the last two or three years they have become much more serious and have centered on adding a proton-proton collider to the LEP complex. The facility now under discussion is called the LEP Hadron Collider (LHC), and the design effort is focussing on a proton-proton collider that will reach 16 TeV in the center of mass (at a 9.0 tesla bend magnet field) and which includes the possibility of electron-proton collisions at 1.8 TeV in the center of mass (0.1 TeV electrons on 8.0 TeV protons).

The dipole design is of the "two-in-one" type wherein two sets of coils are contained in a common collar system, iron yoke, and cryostat as is shown schematically in Fig. 5. This design has been chosen both because of space limitations in the LEP tunnel (the proton ring goes directly above the LEP magnet ring) and because of the perceived saving in magnet costs and in installation time. However the two-in-one design does pose some new problems, the principal one being the coupling of the fields in the two beam tubes which may add considerable complexity to the design of the necessary correction magnet system. The aperture of each coil is currently specified as 5.0 centimeters and the separation of the centers of the two coils is 18 centimeters.




Fig. 5. Schematic of the LHC two-in-one superconducting dipole magnet.

A magnet R&D program is now underway aimed at producing magnets with a maximum dipole field of 10 tesla, though the numbers that CERN is now using in specifying the energy of the machine correspond to a practical operating field of 9 tesla. It is planned to achieve this high field with the standard niobium-titanium conductor by operating at a temperature of around 1.8–2.0°K. A 1.3 meter-long model magnet has been built with a modified HERA cable which reached 7.9 tesla after three quenches at 1.8°. Four 2-bore, 10 tesla, short magnets

have been ordered from industry. The first long magnet (9.5 meters) with twin bores will be built using HERA cable and should reach a field above 7.5 tesla at a temperature of 2°.

The LHC layout is shown schematically in Fig. 6; the numbering of the interaction points corresponds to that used for the LEP electron-positron collider. There are four potential interaction regions at IP-1, -3, -5, and -7. One of these (the one deepest in the Jura) will be used for the necessary proton machine beam dumps, making three available initially for some combination of proton-proton and electron-proton experiments. Some sort of bypass will be required to carry the proton machine around the LEP experiments installed at IP-2, -4, -6, and -8.




Fig. 6. Schematic of the LHC indicating the LHC and LEP potential collision points.

In the proton-proton mode, the CERN design study is focusing on providing both a 10^{33} luminosity interaction point for a general purpose detector, and at least one very high luminosity interaction point which will require some sort of special-purpose detector. Some of the parameters now under study are shown in Table IV which includes two variants of the high luminosity interaction region. As one can see, the very large numbers of proton interactions per beam crossing in the high luminosity modes preclude the use of any general-purpose detector both because of the extreme complexity of the analysis of multiple events and because of the extremely high radiation levels around the collision region. The LHC, like the SSC, also has a significant amount of synchrotron radiation emitted by the beams which will have to be caught on higher temperature radiation catchers because of the very low efficiency of refrigerators which run at the nominal magnet temperature.

The LHC parameters in the electron-proton collider mode are shown in Fig. 7. The kinks in the luminosity and electron beam intensity curves are caused by properties of the LEP electron machine. The luminosity curve assumes that if enough RF power is installed on LEP to allow the storage of five milliamperes at 100 GeV, and that this full power is used for all electron energies above about 35 GeV. Below 35 GeV the current in the electron beam is limited by the aperture of the LEP machine and the small kink in the luminosity curve at around 60 GeV is caused by the need to change the focusing structure of the LEP machine at higher energies. In the electron-proton mode there are 540 bunches in each beam, making the peak current in each bunch below that nominally used in the LEP electron-positron collider mode.

If this project is approved, it is planned during the construction and installation phase to operate LEP for around 4000 hours per year and use the rest of the year to build and install the hadron collider. The new experimental areas will all be designed in the "garage and beam enclosure" mode like the experimental areas at the *SppS*. Thus detector fabrication and installation can go on in parallel with normal LEP operation.

Table IV. Main parameters of the LHC.

P-P Parameters	Nominal	High Luminosity	
Luminosity ($\text{cm}^{-2}\text{s}^{-1}$)	1.4×10^{33}	8.2×10^{33}	3.9×10^{34}
Interaction/Crossing	2.6	9.2	44
Bunch Spacing (N_s)	25	15	15
Protons/Bunch (10^{10})	2.6	4.2	12.5
β^* (M)	1	0.25	0.25
Synch. Radiation (KW)	4	7	21
$\Delta\nu$	2.5×10^{-3}	2.2×10^{-3}	3.4×10^{-3}




Fig. 7. Luminosity, electron beam current, and electron-proton center-mass-energy vs. electron beam energy for the LHC in the electron-proton collider mode. The proton beam energy is 8.0 TeV.

In the operational phase, initially both LEP and the LHC are foreseen to be operated each in its own running period, lasting approximately one-half year. In each experimental mode it is planned to have experiments take data when their collider is in operation or be in their garages when the other collider is in operation. If it is desirable at some time in the future to change some of the LEP areas to LHC areas that can be done with very little construction.

It is not completely clear how much downtime will be required of LEP for the construction and installation of the LHC. Present estimates give 24 ± 6 months total downtime, which will be divided into several shorter down periods so as not to keep the LEP machine off for extended periods of time.

The conceptual design of the LHC is still evolving. It is clear that the highest luminosity numbers will be challenging for the machine builders and very challenging for the experimenters. In particular, much more work needs to be done on detectors that can operate at luminosities of 4×10^{34} including studies of radiation problems, the effect of the necessary shielding on mass resolution, and on the probability of more than 40 events per beam crossing generating some odd backgrounds which might mimic or mask the effects that one is looking for. It will probably take another 6-12 months to firm up the LHC design.

3. LINEAR COLLIDERS

It is now generally agreed that the linear collider technique is the only way to reach center-of-mass energies in the electron-positron system much higher than the center-of-mass energy of LEP II. In an electron storage ring intense synchrotron radiation is emitted as the beam circulates with an energy loss to synchrotron radiation proportional to the fourth power of the energy divided by the bending radius. A scaling law for an electron storage ring can be written down which minimizes the cost of a machine for a fixed energy, and this scaling law yields a size and cost for such a machine proportional to the square of the beam energy.¹ Thus, to achieve an energy ten times that of LEP II one would have to increase the circumference by a factor of 100 to 2700 kilometers with a concomitant cost of $1 - 2 \times 10^{11}$ swiss francs. While there are technical problems in building an electron storage of this size, it is clear that the fiscal problems are such that such machines are not really feasible.

In a linear collider, on the other hand, no synchrotron radiation is emitted in the acceleration process, resulting in a more benign scaling law making the cost of high energy linear colliders considerably less than an electron storage ring. Beams in linear colliders can be extremely small, and so high luminosities can be obtained even at the relatively low repetition rate of room temperature linear accelerators.

The first machine of this type, the SLC at the Stanford Linear Accelerator Center, is just now coming into operation. Beams of three or four micron radius can be routinely produced at the collision point, and the stability of these beams are such that they can easily be held in collision by simple feedback systems that hold the beams centered to a small fraction of their size. While the SLC has not yet achieved sufficient operating reliability to begin producing data for physics experiments, it is already clear that the goal of the proof-of-principle has been met. This in turn has led to a great expansion in the R&D devoted to high energy linear collider systems.

This interest in high energy electron-positron colliders comes about because it is possible to do physics with these machines not easily accessible to proton colliders.² There is a kind of democracy in the final states produced in electron-positron collisions in that all partial cross sections are comparable as long as the particles in the final state have either electromagnetic or weak charge. In addition, peripheral processes at large transverse momentum are small and are easy to distinguish from the interesting events with relatively simple cuts. The cleanliness of the final state in electron-positron processes makes the life of the experimenter very much easier than it is at proton colliders, as is shown in Fig. 8 which is a Monte Carlo simulation of W pair production with the W 's decaying into six jets in the final state. The absence of "debris" like that present in hard proton-proton collisions (see Fig. 1) makes detectors much easier to design and build and analysis much simpler to carry out.




Fig. 8. Monte Carlo Simulation at 1.0 TeV of $e^+e^- \rightarrow w^+w^- \rightarrow udb \rightarrow \text{jets}$.

The luminosity required in a high energy electron-positron collider is roughly given by

$$\mathcal{L} = 10^{33} E_{\text{cm}}^2 (\text{TeV}) \text{ cm}^{-2} \text{s}^{-1}$$

With this luminosity, roughly 1000 events per 10^7 operating seconds are produced for each R unit of cross section (one unit of R is the cross section for electromagnetic production of mu pairs). Many of the new kinds of particles which are thought to possibly exist at high mass have cross sections on the order of one unit of R.

Laboratories in the U.S., Europe, the Soviet Union, and Japan are engaged in R&D aimed toward what I will call the Next Linear Collider (NLC) which will be a machine someplace in the energy region 0.5–2.0 TeV in the center-of-mass.

A very qualitative picture of the state of linear collider technology is shown in Fig. 9, which illustrates in the energy/luminosity plane what might be done with small extensions of present technology and the region that will certainly require some kind of new approach. The energy and luminosity requirements of the NLC, which are determined by its physics goals, push us toward the new approaches region.




Fig. 9. Technology requirements of a next-generation linear collider in the luminosity-energy plane.

A numerical example is shown in Table V. Here the parameters of the SLC are compared to two variants of an NLC, one using the well tested linac technology of the SLC, and another using what seems to be plausible parameters which require the development of new technology. The NLC has 10 times the energy of the SLC but requires a luminosity more than 100 times as great. Using the technology of the SLAC linac, which has an accelerating gradient of 20 MeV per meter and runs at an RF frequency of 2.9 gigahertz, results in a machine that is approximately 50 kilometers long and uses a wall plug power of about 500 megawatts. This is to be compared with what might be done with an accelerating gradient of 200 MeV per meter, which has been demonstrated in single cavities, at an RF frequency of 11.4 gigahertz which is four times that of SLAC. The resulting accelerator is five kilometers long and uses a total wall plug power of about 100 megawatts. The new technology machine uses a flat beam at the collision point compared to the round beam of the SLC, and results in one of the beam dimensions being very small indeed — around five nanometers. Perhaps we cannot push quite as far as this new technology example, but it is clear that increased accelerating gradient and higher RF frequency will result in a considerably less costly accelerator facility.

Before going on to discuss the state of linear collider R&D, I want to give a brief introduction to what goes on in the collision region of a linear collider.³ There are problems here that will affect the experimenter's ability to do experiments and questions which require answers that cannot be determined solely by the accelerator physicists.

The beam-beam interaction in linear colliders can be very much stronger than would be allowed in a storage ring. The reason for this is that, since the beam is to be used only once, one can allow the electromagnetic fields of the two beams to disrupt their phase spaces to a much larger extent than is allowable in a storage ring where the beams must continue to circulate in a magnet ring for a very long time. In an electron-positron collider the collective fields of one beam will focus a single particle in the other beam, as illustrated in Fig. 10.

Table V. Some critical parameters of SLC, "old" and "new" technology NLC.

	SLC	NLC	
		SLC Technology	New Technology
Energy (TeV)	0.1	1	1
Rep Rate (Hz)	180	360	90
Luminosity ($\text{cm}^{-2}\text{s}^{-1}$)	6×10^{30}	10^{33}	10^{33}
Accelerating Grad. (kV/M)	20	20	20
Length (km)	3	50	5
RF Frequency (GHz)	2.9	2.9	11.4
Total "Wallplug" Power (MW)	50	500	100
$\sigma_x \times \sigma_y (\mu^2)$	$(1.6) \times (1.6)$	$(0.4) \times (0.4)$	$(1) \times (5 \times 10^{-3})$

The strength of the interaction is measured by a dimensionless parameter (D), the disruption parameter, which is the ratio of the bunch length to the focal length of an equivalent lens. For round trigaussian beams, D is given by

$$D = \frac{\sigma_z}{F} = \frac{r_e \sigma_z N}{\gamma \sigma_r^2}$$

where the bunch has a longitudinal standard deviation σ_z , a radial standard deviation σ_r , a number of particles N , and an energy γ in rest-mass units; r_e is the classic electron radius, and F is the small amplitude focal length of an equivalent thin lens.




Fig. 10. Effect of the intense fields in one bunch of a linear collider on a particle in the other colliding bunch.

The effective fields in a linear collider tend to be very large, and the focal lengths can be very small. Even in the SLC the fields are of the order of megagauss, F is in the order of millimeters, and D is about one. In the higher energy machines being discussed now, the fields are tens of megagauss, F is tens of microns, and D is 5-10.

Large values of D imply that the beam cross section is strongly perturbed during the collision. For values as high as 5-10, the interaction is sufficiently strong so that a kind of mutual pinch occurs reducing the radius of both beams during the collision period and hence enhancing the luminosity.⁴

While synchrotron radiation is no problem in the acceleration process in a linear collider, the very large effective fields in the collision region can generate extremely intense synchrotron radiation. At high luminosities the synchrotron radiation, called "beamstrahlung," dominates the energy spread in the beam. What is important to the experimenters, and hence to the

radiation. At high luminosities the synchrotron radiation, called "beamstrahlung," dominates the energy spread in the beam. What is important to the experimenters, and hence to the machine designers, is the spread in center-of-mass energy generated by this beamstrahlung phenomenon. Particles in one beam lose energy to synchrotron radiation photons as they pass through the other beam and so even for the case of zero energy spread in the incident beams, there can be long tails in the energy distribution of the colliding beams if the synchrotron radiation is sufficiently intense. Naturally it turns out to be easier for the machine designers to make machines with very large values of the center-of-mass energy spread while it turns out to be difficult for the experimenters to do experiments if this energy spread is too big.

Qualitatively, for small values of δ , a parameter approximately equal to the mean loss in energy of a particle in one beam in traveling through the other beam, the center-of-mass collision energy distribution is sharply peaked around the initial center-of-mass energy, while for large values of δ the distribution has long tails stretching out toward low center-of-mass energy. Figure 11 shows the integral distribution of the square of the center-of-mass energy (S). I have plotted three cases which show the fraction of the time that S/S_0 is greater than a given value vs. that value. For $\delta = 0.4$, only around 20% of the time is S within 2.0% of its maximum value, while for $\delta = 0.1$ it is within 2.0% 60% of the time. Studies at SLAC and at CERN indicate that a reasonable compromise would be a value of δ of around 0.25.




Fig. 11. Integral distribution of the energy spread in a 1.0 TeV linear collider for various values of the beamstrahlung energy loss parameter, δ .

Let me now turn to a discussion of the state of accelerator technology. There are four main areas that need considerable research and development before we will be ready to build a machine. These are the electron and positron sources where the beams are born, the accelerators that boost them to the required high energy, the final focus system that squeezes them to an exquisitely small size, and the beam dynamic studies that will tell us how all of these systems interact with each other. The largest and most expensive part of the NLC will be the accelerators, and so I will spend most of my time on that topic. It is, however, worthwhile to say a few words about the other three.

It is easier to make a small beam at the collision point if the beam has been born small at its source. The term "small" in this context means that we require a source of low emittance (the invariant emittance of the beam is proportional to the energy times the transverse size times the transverse angular spread). The NLC will require sources with an invariant emittance no more than about 10% of that used in the SLC. I think we understand how to do that job — we can use existing storage ring technology, but must pay a great deal of attention to the details to make sure that the emittance does come out as small as it can, in principle, be. The damping storage rings will be somewhat different in design from that used now, but it looks like the energy of these damping rings will be in the GeV region.

The final focus system will be difficult. The beam sizes are much smaller than they are in the SLC, while the energies are much higher so that the focusing system requires much stronger elements. The final focusing magnets will surely be superconducting, though some work is going on using plasma lenses which can be made even stronger than superconducting magnets for final focus elements. This focusing system becomes more difficult the larger the

energy spread in the incoming beam will be, and regrettably, the smaller the required energy spread at the end of the accelerator the harder the accelerator is to build. This area needs a great deal of work, which can be theoretical for awhile, but eventually we are going to have to build some prototypes.

More detailed beam dynamic studies are required everywhere. The interaction of the beam with the accelerating structure (wakefields) must be better understood, and much work is required on tolerances, stability requirements, etc. There is more than enough to do to keep the theoreticians in the accelerator community busy for some time.

As mentioned earlier, the accelerators and their power sources will be the most expensive part of the new machine, and it is here that most of the R&D work is now concentrated. The accelerators must be energy-efficient, stable, and able to preserve the small emittance of the beam from the sources through the full acceleration cycle. If one doesn't care about a few billion dollars here or there, one could probably use the SLAC linac technology for the NLC. The machine would be long, expensive and a terrible power hog. New developments in this area will strongly affect not only the construction costs of the machine, but its operating costs as well.

Four main approaches have been under discussion. These are

1. Laser accelerators
2. Plasma accelerators
3. Wakefield accelerators
4. Conventional RF structures with either conventional or exotic power sources.

I think all of us who are active in this field (SLAC, Novosibirsk, KEK, CERN) have come to the conclusion that the NLC can only be built via the fourth method. It is the only one where we can see how, at least in principle, to get the required stability and energy conversion efficiency.

The stability requirement is very severe for we want to make a colliding-beam device, and not a fixed-target device. Beams from two independent accelerators must meet each other reliably and reproducibly within tolerances of a tiny fraction of a micron. The first three methods all have severe problems — intensity fluctuation and mode structure (lasers), laser drivers and plasma uniformity (plasmas), and azimuthal asymmetry of drive beams (wakefield). All of them seem to suffer from serious inefficiency problems as well. I believe they are not for the next generation of linear colliders, though it may well be that new approaches and new technology may make these kinds of systems viable in 15–20 years.

The most promising system appears to be the conventional linear accelerator with some kind of high-power driver, which itself will have to be some new technology. The machines will probably use much higher accelerating gradients than are used now, and will almost surely be considerably shorter RF wavelengths than are used in the SLAC machine. The push toward high accelerating gradients is driven by the costs of the accelerator structure itself. The higher the accelerating gradient, the shorter the machine and its civil construction can be. At SLAC we have shown that for about one microsecond pulses at 3 kHz, copper structures can stand accelerating gradients of more than 100 MV per meter at 3 kHz, and more than 300 MV per meter at 10 kHz. Thus, high accelerating gradients also seem to benefit from higher RF frequency.

The electrical efficiency of the accelerating system also benefits from higher RF frequencies. For a given accelerating gradient, the stored energy per unit length in an accelerator is proportional to the square of the RF wavelength. Thus, for a given charge per bunch, the fraction of the energy stored in the accelerating structure that can be extracted by the bunch increases as the wavelength decreases. Of course, if one had a superconducting accelerator structure, one would not have to worry about the fraction of stored energy extracted, for the leftover energy could be used to accelerate the next bunch. However, superconducting systems cannot attain very high accelerating gradients, and so the cost of a main accelerator done with superconductivity will be very high, as will be the power required to run the compressors of the refrigerator unless the Q of these systems can be significantly increased. Everyone now seems to be talking about systems with frequencies from 10 to 30 GHz.

The power sources for these machines will require something new. Very high accelerating gradients go with high peak power in the accelerating structure. The machines under discussion at various laboratories in the world use peak powers on the order of 1/2 to 1 gigawatt per meter of accelerating structure. Generating these high peak powers will be quite a challenge. Fortunately, the average power is not much higher than we deal with today. These high

peak powers are associated with short pulse lengths (typically 50 nanoseconds, or so) and so the average power required is not much different than that which comes from conventional klystrons.

One method that has been investigated at SLAC to generate high peak power from almost conventional klystrons is pulse compression. By combining multiple power sources through low-loss delay lines, with proper phase manipulation at the power sources, it is possible to get pulse compression ratios of ten or twenty to one. These systems are complicated, delicate, and require an enormous amount of plumbing for the delay lines, but they do seem workable.

Of more interest are the variants of what might be called two-beam accelerator systems. One beam with low energy and high current in one accelerator structure is used to generate RF power which drives a second accelerator structure. Two variants of this are currently under investigation. One being pursued by a Berkeley/Livermore/SLAC collaboration uses induction linacs to produce beams of several kiloamp current at energies of several MeV, with klystron-like bunching and energy extraction cavities. We hope to demonstrate a 500 gigawatt, 50 nanosecond pulse length RF source sometime next year.

A different approach is being pursued at CERN. The CLIC group is investigating the use of superconducting cavities like those already designed and tested to increase the Tristan or LEP energy for the high current, low energy accelerator. A train of short, high-current bunches rides in this low frequency accelerator and interacts with a high frequency cavity structure to produce RF power which is used to charge the high energy accelerator. The CLIC group is interested in frequencies of around 30 GHz for the high energy machine, and are modeling the energy extraction cavities for tests at a lower RF frequency.

This field is moving very fast, and I think in a few year's time there is a very good chance that a practical power source/accelerator combination will be available.

Major R&D programs are either in place or are developing in Japan, Europe, the USSR and the U.S. In Japan a group centered at KEK with contributions from other Japanese universities is aiming toward a machine with about 1.0 TeV in the center-of-mass. This program is growing as people and resources are freed from work on the Tristan colliding beam storage rings. The main emphasis is now on studies of high gradients and appropriate structures for large linear colliders.

In Europe there is the "CLIC" program at CERN which aims toward a 2.0 TeV machine. The R&D work here is concentrating on a superconducting RF generator running at 35 GHz in combination with a room temperature high gradient accelerator. Intense work is underway on the beam dynamics of the driver and on the problem of transferring energy from the superconducting RF generator to the high gradient accelerator.

In addition in Europe there is the Frascati "ARES" project which is aiming toward a recirculating superconducting linac to produce the beams and SLC-like collision geometry. The final phase of the project might be a *B* factory, while the first phase under study is a nuclear physics facility. Fifty four million dollars has been authorized INFN over a five year period for R&D.

At Orsay there is a smaller program aimed at high brightness electron guns and on the use of field emission and lasers to produce high power RF generators.

In the USSR there is a program centered at Novosibirsk which aims at building a 2.0 TeV facility at Serpukhov in stages. This project has a conditional approval from the Soviet authorities. If certain milestones are met, conventional construction might start in 1992 or 1993 on the tunnel to house this facility. The main milestone is the development of a 10-meter-long accelerating section with RF drivers to run at a gradient of 100 MeV per meter.

In the US, the SLAC program is aimed at a 0.5-1.0 TeV machine and is concentrating on the development of new power sources, structures for high frequency accelerators, and on theory. In addition, a final focus test facility is in design which could allow all of the groups in the world interested in this kind of machine to try out new ideas in this very difficult area. There are also programs at LBL, Lawrence Livermore National Laboratory, and UCLA.

No one is ready to proceed as yet with a machine. There is much R&D to do in all areas before a believable design study can be produced with a reasonably reliable cost estimate. The four regions now heavily involved in this kind of work are going to try to carry out the R&D program internationally with a mixture of coordinated and collaborative work. No single group can investigate all the promising alternatives, and all groups will move faster by cooperating. The first international workshop on linear collider R&D will be held at SLAC

in November 1988. If things go well I think we can expect serious proposals in three to four years.

4. B MESON FACTORIES

The results of the Argus group on $B\bar{B}$ mixing⁵ first reported at the Lepton Conference in Hamburg in 1987 have stimulated an enormous amount of interest in facilities that can produce a large number of B mesons. The reason for the interest is that there appears to be a chance to study CP violation in non-kaon systems as well as the intrinsic interest in the B meson system. The interest is great enough, I believe, to lead to the construction of one or more new colliders specifically aimed at experiments on the B meson system in the next few years. In this section I will summarize what is going on in the study of B meson factories and try to compare the large number of different approaches from the perspective of the possibility of studying CP violation. Almost all of these studies are aimed at new kinds of electron-positron colliders or improvements or existing ones.

Before going on to describe the variety of approaches in electron colliders I should mention that proton machines are also copious sources of B mesons. For example:

1. The Tevatron and UNK operating in the fixed-target mode are capable of producing 10^8 to 10^9 B mesons per year, the partial B cross section being a few times 10^{-6} of the total cross section.
2. The Tevatron collider can produce 10^8 B mesons per year, the partial cross section being a few times 10^{-4} of the total cross section.
3. SSC or the LHC can produce 10^{12} to 10^{13} B mesons per year, the partial cross section being a few times 10^{-3} of the total cross section.

While the number of B mesons produced in these proton machines is enormous, certainly enough to study CP violation if a good detection system could be devised, I have seen as yet no credible experiment with sufficient efficiency and resolution to separate out the interesting final states and do the required physics experiments. However, many groups are studying the problem, and perhaps a good experiment can be devised.

Electron-positron colliders seem much more promising as B meson factories. Three machines are running now at the ψ 4S region. They are CESR at Cornell with a luminosity of $10^{32} \text{ cm}^{-2}\text{s}^{-1}$, DORIS II at DESY with a luminosity of 4×10^{31} , and VEPP IV at Novosibirsk with a luminosity of 4×10^{30} . For reference, a machine running with a luminosity of 10^{32} for 10^7 seconds per year at a reconstruction efficiency of 1.0 in the 10 GeV center-of-mass region will produce:

- 9×10^5 tau pairs
- 2.8×10^6 non- B hadrons
- 1.8×10^6 $B\bar{B}$ (4S)
- 2.9×10^5 $B\bar{B}$ (continuum).

Most studies indicate that roughly 10^8 B mesons are needed for the study of CP violation.

Many different approaches are being pursued. Conventional storage rings at around 10 GeV in the center of mass are being studied at Cornell, KEK, Novosibirsk and the Paul Scherrer Institute (SIN). Asymmetric storage rings with a center-of-mass energy of around 10 GeV are being looked at at DESY, KEK, and SLAC/LBL. The possibilities of Z factories as B meson factories are being reviewed at CERN and at SLAC. Linear colliders in the 10 GeV region are being studied at Frascati and UCLA. Hybrid systems involving linacs colliding with the circulating beam in the storage ring are beginning to be looked at as well. All studies are aimed at high luminosities, and some of them may have advantages over others in the detectability of CP violation given the same number of B mesons.

Symmetric Storage Rings

This is the standard electron-positron collider and the limitations of, and technology required to implement this approach are well understood. The CESR machine at Cornell must be our standard of comparison. This is a single storage ring with counter-rotating beams of electrons and positrons traveling in a single vacuum chamber. The Cornell group has cleverly managed to circulate seven bunches of particles in each direction and have these bunches avoid each other except at the designated collision point by a technique where the orbits spiral around

two to three years of hard work. That maximum luminosity would produce approximately 10^7 B mesons per year running on the ϵ 4S.

A new double storage ring is being proposed by the Paul Scherrer Institute. This double ring is something like the original DORIS design at DESY and keeps the beams completely separated except at the two interaction regions where they collide head-on. The circumference of the machine would be 650 meters, and the operating range would be from 2–6.5 GeV per beam. The design is quite conservative, and the proposers, I believe, considerably understate the potential of the machine. In what the proposers call their standard optics ($\Delta\nu = 0.03$, $\beta = 3$ cm) the luminosity at 10 GeV in the center-of-mass is specified as around 5×10^{32} with one megawatt of RF power and one-half ampere of current circulating in each ring. The high luminosity optics ($\Delta\nu = 0.04$, $\beta = 2.0$ cm) reaches 1.2×10^{33} with the same power and circulating beam current. However, other storage rings have achieved values of $\Delta\nu = 0.05$ and run at $\beta = 1.5$ cm, and assuming those parameters for the PSI machine the luminosity would be 2.5×10^{33} . The cost of the PSI machine is estimated to be 129 million Swiss francs plus manpower and detectors, and the group hopes for approval to start construction in 1990 with beam-on time being 1994.

Novosibirsk is working on the conceptual design of a two-ring $B\bar{B}$ factory as well. The design is not as far along as the PSI project. They assume a $\Delta\nu$ of 0.05, a β of 1.0 cm, and 10 megawatts of RF power available yielding a luminosity of 10^{34} . This is an aggressive design.

At KEK modifications to the 8-GeV Tristram booster are being considered. There is no design as yet, but this is a single ring and would be something like the Cornell machine.

Asymmetric Rings

This is an idea put forward a few years ago by Oddone of LBL⁶ and further developed by him and Feldman of SLAC. They envisioned a small ring colliding with a larger ring (a new 2-GeV ring, for example, colliding with PEP running at 12 GeV) to give a center-of-mass energy of around 10 GeV, but with the center of mass in motion. The advantage of this system over the conventional symmetric storage rings lies in this center-of-mass motion which separates the decay vertices of the B and $\bar{B}$ mesons. For example, at $\beta\gamma = 1.0$ the mean separation between the two decay vertices is 300 microns. Using vertex detectors there would be less combinatoric background, and one can see a clean CP signature by studying the time evolution of certain final states which cannot be studied with symmetric machines at the 4S where the two decay vertices are essentially on top of each other because of the low Q value in the decay of the 4S to a B and a $\bar{B}$.

This system is being studied at DESY, KEK and SLAC/LBL. This is no accident, for these are the labs which have the high energy ring and where the low energy ring could be added at moderate cost. The studies are still in an early state, but it looks like the luminosity might be in the 5×10^{32} to 2×10^{33} range. However, there are new beam dynamics problems in these asymmetric machines with which we have not had much experience with before, and they need considerably more study.

Z Factories

Z factories are B factories as well, for the cross section for $B\bar{B}$ production from Z decay is about five times the cross section on ϵ 4S resonance. Both LEP and the SLC could be used, though the luminosity would certainly have to be considerably above the initial design values at both machines to make them effective.

The LEP luminosity could be increased by using the RF to be installed as part of the LEP II, while running at the Z . If a larger number of bunches could be circulated in LEP than it was originally designed to hold (some sort of separation scheme would be required), it is estimated that LEP might run at a luminosity of 10^{32} .

The SLC machine with polarized beams has the potential for "self-tagging" through the polarization asymmetry. Any CP violation experiment must use some method to determine whether the particle which decays was originally a B or a $\bar{B}$. However, with longitudinal polarization life is much easier since the left-handed nature of the weak interaction means that b quarks are produced in the forward direction while $\bar{b}$ quarks are produced in a backward direction. The effective analyzing power is 0.75 times the polarization.

Linear Colliders and Hybrids

Two linear collider schemes are under discussion as possible *B* factories. These are the ARES project at Frascati which is a superconducting machine, and a room temperature linear collider at UCLA. The conceptual design of the ARES project is much further advanced, and so I will concentrate on it, as shown schematically in Fig. 12.




Fig. 12. Schematic of the ARES project.

The first stage of the machine is an electron gun and a 150 MeV superconducting linac whose sole purpose is to produce positrons for the next stage of the machine. This is technically a very difficult problem, for the positron yield per incident electron at 150 MeV is probably about 1/2%, and so several megawatts of beam power are required to get the necessary positron yield. This is a difficult problem for the accelerator and an even more difficult problem for the design of the positron production and collection system.

The next stage of the project is a superconducting recirculating linac where the beams make four passes through the linac reaching a final energy of 1.5 GeV. In concept this is much like the CEBAF nuclear physics facility in the U.S. After acceleration to 1.5 GeV the beams are inserted into two small storage rings where their emittance decreases by synchrotron radiation. Here, too, there will be problems since the time between bunches in the linac is short compared to the damping time in the storage rings. Either the storage rings must have a very large circumference or there must be several of them and the injection and extraction pulsed magnets must be very fast.

Upon extraction from the damping rings the beams go through the next stage which is again a recirculating linear accelerator, giving a final energy at the output after four passes of up to 5.5 GeV. At this point the beams go into a transport system which is conceptually like that used in the SLC, and are brought to the collision point.

The luminosity goals of ARES are 10^{33} to 10^{34} . It is a difficult project and a great deal of R&D will be required to bring it to fruition. The INFN in Italy has authorized a \$54 million R&D program for '89-'93. The first step in the program will be a four-cavity superconducting linear accelerator with an energy of about 30 GeV.

Paul Grosse-Weissman⁷ has revived an old idea of Csonka and Rees⁸ to collide a beam from a linear accelerator with a beam circulating in a storage ring. The goal is a luminosity of around 10^{34} . The use of the storage ring solves the positron problem which is present in the ARES design, by building up a recirculating positron beam in the storage ring which is in this concept to operate at a much higher energy than the linac. Thus, the disruption of the lower energy electron beam by the beam-beam collision can be very severe, while the positron disruption is small, allowing that beam to continue to circulate in the storage ring. This system can only reach high luminosity with a high repetition rate linear accelerator, for there must be an electron bunch to collide with each of the positron bunches circulating in the storage ring. In practice this means that the linac must be superconducting. This idea gives a way to reuse older high energy storage rings. For example, at PEP a 2.0 GeV linac would

be required to collide with a 12 GeV positron beam in the storage ring; at KEK a 1.0 GeV linac could collide with a 25 GeV beam in the storage ring; while at LEP a 0.5 GeV linac is required to collide with a 50 GeV circulating beam. There has been as yet very little detailed work on this concept, and much remains to be done before one can conclude that it is viable.

Comparisons

The best comparison of the potential of all of these different methods has been done recently by the B factory working group of the 1988 Snowmass Summer Study.⁹ This group (chaired by D. G. Cassel and G. J. Feldman) chose a particular process to use in the study of CP violation, was realistic about detection problems, included uncertainties in the CP violation parameters, and computed the luminosity required for a three standard deviation effect in a collection of scenarios using different kinds of B factories. The process chosen for the study was B or $\bar{B}$ decaying to ψK_s . The required luminosity is determined from

$$\int \mathcal{L} dt = \sigma(e^+e^- \rightarrow b\bar{b}) f^0 2B\epsilon_r\epsilon_t(1-2w)^2 d^2 \delta^2 (\sin \phi)^{-1}$$

where f^0 is the fraction of B^0 in the b quark fragmentation; B is the branching fraction assumed to be 5×10^{-4} for the ψK_s model, and 0.14 for $\psi \rightarrow l^+ l^-$; ϵ_r is the ψK_s reconstruction efficiency; ϵ_t is the tagging efficiency; w is the fraction of incorrect tags; d is a dilution factor having to do with fitting, integrating over time, and the mixing of the tagged state; and $\delta(\sin \phi)$ is the required accuracy on the CP asymmetry parameter $\sin \phi$.

Table VI. Comparison of B -factory techniques.

Case Factor	Asym. Y(4S)	Sym. Y(4S) ⁺	$\sqrt{s} =$ 16 GeV	Z ($P = 0$)	Z $P = 0.9$ ($P = 0.45$)
$b\bar{b}$ cross section, $\sigma(\text{nb})$	1.2	0.3	0.11	6.3	6.3
Fraction of B^0 , f^0	0.43	0.34	0.35	0.35	0.35
ψK_s reconstruction efficiency, ϵ_r	0.61	0.61	0.61	0.46	0.46
Tag Efficiency, ϵ_t (& method)	0.48 (1,K)	0.48 (1,K)	0.30 (1,D)	0.18 (1,D)	0.61 (A_{FB})
Wrong tag fraction, w	0.08	0.08	0.08	0.08	0.125 (0.27)
Asymmetry dilution, d	0.71	0.63	0.52	0.52	0.71
$\int L dt$ needed for 3σ effect (10^{40} cm^{-2}) *	0.3-12.	2.2-78.	14.-490.	0.5-19.	0.1-3.6 (0.3 - 9.6)
Relative $\int L dt$ needed	1.0	6.4	40.	1.5	0.3 (0.8)

* = the peak luminosity needed in units of $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ for 10^7 seconds of fully efficient running at peak luminosity.

Table VI shows the results of their study. The large uncertainty in the integrated luminosity required to measure the effect comes almost entirely from our present state of knowledge of the CP violating phase in the KM matrix. The Z factories and the asymmetric ring at the upsilon 4S are the clear winners in this analysis.

There are of course other ways of looking for CP violation than by studying the time evolution of the B^0 system, and it would be interesting to see the results of a similar analysis

to conclude that this is going to be a difficult problem to solve. The lower edge of the luminosity band shown in Table VI. is within reach, while the upper edge of the band will require something new in colliders. Thus, if nature is not too unkind, and CP violation is near the maximum of what it might be, we can expect to find it with a collider like that under consideration at the Paul Scherrer Institute (perhaps even with CESR) or with a new kind of machine which is not much of a stretch with existing technology. The asymmetric machines or the polarized beam Z factory looks to be the only reasonable way to get to the upper end.

SUMMARY

In this paper I have reviewed the possibilities for new colliders that might be available in the 1990's. One or more new proton colliders should be available in the late-90s based on plans of Europe, the U.S. and the U.S.S.R. The two very high energy machines, LHC and SSC, are quite expensive, and their construction will be more decided by the politicians' views on the availability of resources than by the physicists' views of the need for new machines. Certainly something will be built, but the question is when.

New electron colliders beyond LEP II could be available in the late 1990's as well. Most of the people who have looked at this problem believe that at a minimum three years of R&D are required before a proposal can be made, two years will be required to convince the authorities to go ahead, and five years will be required to build such a machine. Thus the earliest time a new electron collider at high energy could be available is around 1998. A strong international R&D program will be required to meet that schedule.

In the field of B factories, PSI's proposal is the first serious step beyond the capabilities of CESR. There are other promising techniques but these need more R&D. The least R&D would be required for the asymmetric storage ring systems, while the most would be required for high luminosity linear colliders.

For the next decade, high energy physics will be doing its work at the high energy frontier with Tevatron I and II, UNK, SLC, LEP I and II, and HERA. The opportunities for science presented by experiments at these facilities are very great, and it is to be hoped that the pressure for funding to construct the next generation facilities will not badly affect the operating budgets of the ones we now have or which will soon be turning on. There are great things to do with what we have or will soon have.

To go beyond these machines, a reasonable share of the high energy physics resources *must* go into accelerator R&D. Accelerator R&D represents the seeds from which future accelerators will flower, and though budgets may be tight, we must reserve an appropriate fraction to build for the future.

REFERENCES

1. B. Richter, Nuclear Instruments and Methods **136**, 47 (1976).
2. J. H. Mulvey, ed. *Proceedings of the Workshop on Physics at Future Accelerators*, (CERN 1987); C. Ahn et al., *Opportunities and Requirements for Experimentation at High Energy e^+e^- Colliders*, SLAC Report-329 (May 1988).
3. A more extensive primer on linear colliders can be found in B. Richter, *Proceedings of the 7th International Conference on Physics in Collision*, p. 423 (Tsukuba, Japan, 1987).
4. P. Chen and K. Yokoya, Phys. Rev. D **38**, 987 (1988).
5. H. Albrecht et al. (ARGUS), Phys. Lett. **192B**, 245 (1987).
6. P. Oddone, Detector Considerations, *Proceedings of UCLA Workshop on Linear Collider $B\bar{B}$ Factory Conceptual Design*, January 1987.
7. P. Grosse-Wiesmann, SLAC-PUB-4545 (Feb. 1988), submitted to Nuclear Instruments and Methods.
8. P. L. Csonka and J. Rees, Nuclear Instruments and Methods **96**, 149 (1971).
9. Report of the B Factory Working Group, *Proceedings of the 1988 Snowmass Summer Study*, (to be published).