Really Large Hadron Collider Working Group Summary*

G. Dugan
Cornell University, Ithaca, NY 14853

P. Limon
Fermilab, P.O. Box 500, Batavia, IL 60510

M. Syphers
Brookhaven National Laboratory, P.O. Box 5000, Upton, NY 11973

Participants in the Really Large Hadron Collider (RLHC) Working Group

M. Albrow, R. Bauer, B. Brown, W. Chou, P. Colestock, D. Denisov, R. Diebold, G. Dugan, S. Feher, G.W. Foster, J. Friant, A. Ghosh, G. Goderre, P. Grant, D. Gross, R. Gupta, M. Harrison, S. Holmes, H. Ishimaru, H. Jostlein, E. Keil, S. Keller, K. Koepke, J. Lach, P. Limon, J. MacLachlan, E. Malamud, M. May, P. Mazur, M. McAshan, A. McInturff, F. Mills, N. Mokhov, T. Murphy, D. Neuffer, S. Peggs, J. Peoples, J. Rogers, G. Snitchler, M. Syphers, J. Tompkins, W. Turner, P. Wanderer, J. Wei, J. Welch, V. Yarba, A. Zlobin

ABSTRACT

A summary is presented of preliminary studies of three 100 TeV center-of-mass hadron colliders made with magnets of different field strengths, 1.8T, 9.5T and 12.6T. Descriptions of the machines, and some of the major and most challenging subsystems, are presented, along with parameter lists and the major issues for future study.

I. INTRODUCTION

Hadron colliders are the "Discovery Machines" for high-energy physics. They reach farther and probe deeper than any other type of accelerator. Any rumor of a possible future hadron collider attracts large crowds of experimentalists and theorists because of the near certainty that exciting new physics will be found. That attraction is well deserved. The W and Z were first observed at the $Sp\bar{p}S$. The surprisingly massive top quark was discovered at the Tevatron. It may be possible to find Higgs or supersymmetric particles there. As we understand more and move to higher mass scales, only hadron colliders can get you there. The LHC will extend our reach another factor of seven. What is after that?

Hadron colliders also attract great crowds of accelerator physicists, not only because of the promise of interesting discoveries, but also because of the challenges of building these large and complex machines. A great deal has been learned in recent years about the design and operation of hadron colliders, and there is no doubt that we can build a working hadron collider. The challenge has become one of building the most effective machine for the smallest possible investment. In fact, it may be that the biggest challenge will be establishing the world-wide cooperation that will be needed to spread the cost of these machines over many nations and regions. Now is the time to accept that challenge.

A. Developments Since the 1994 Indiana Workshop

Since the Workshop on Future Hadron Facilities in the U.S. was held at Bloomington, IN, two years ago [1], there have been a number of significant changes and additions to the hadron collider scene. These changes have influenced this workshop and modified our vision of future hadron colliders in the U.S. Among the most important are:

1. The LHC, a high-luminosity, 14 TeV center-of-mass (CM) proton collider has been approved by the CERN member states to be operational by 2005. The U.S. is an active and significant participant in both the experiments and construction of this collider.
2. The assured existence of the LHC suggests that we reexamine the choice of 60 TeV (CM) used as the energy of the collider studied at the Indiana workshop.
3. An enthusiastic group has started to study a machine called the "Pipetron [2]," based on superferric magnets with a <2 T field.
4. There has continued to be progress in the development of high-temperature superconductor (HTS), with some types reaching commercialization. New processes are being discovered that will improve the performance of HTS and hold the promise of becoming practical production methods.

B. Parameter Lists

Table I is a comprehensive and self-consistent set of machine parameters for the three designs. All parameter sets correspond to 100 TeV (CM) energy and an initial luminosity of 10^{34} cm^{-2} sec^{-1}, and all have similar injectors, with 3 TeV as the energy of the last injector.

*Work supported by the National Science Foundation and the U.S. Department of Energy

Table I : Machine parameters

Parameter	High field-new technology	High field-known technology	Low Field	Units
CM Energy	100	100	100	TeV
Dipole field	12.6	9.5	1.8	T
Circumference	104	138	646	km
Synchrotron radiation damping time (horizontal amplitude)	2.6	4.6	*antidamped*	hr
Initial/peak luminosity	.35/1.2	.35/1.0	1./1.	10^{34} $cm^{-2}sec^{-1}$
Integrated luminosity per day	500	500	700	pb^{-1}
Number of stores per day	2	2	1	
Initial rms normalized emittance	1.	1.	1.	π μm-rad
β^*	20	20	20	cm
Protons/bunch	0.5	0.5	0.94	10^{10}
Number of bunches	20794	27522	129240	
Equilibrium emittance (x)	144.2	62	1.8	10^{-3} π μm-rad
Bunch spacing	16.7	16.7	16.7	nsec
Beam stored energy	.89	1.18	9.73	GJ
Synchrotron radiation power/ring	189	143	48	kW
Total protons/ring	1.1	1.5	12.2	10^{14}
Initial/peak interactions/crossing	7.5/21.5	7.5/21.5	21.5/21.5	
Beam lifetime (pp collisions only)	34	45	130	hr
$\sigma_{inelastic}$	130	130	130	mbarn
Initial beam-beam $\Delta\nu$ (total)	5.1	5.1	11.6	10^{-3}
Revolution frequency	2.89	2.18	.46	kHz
Synchrotron frequency	8.9	5.8	.86	Hz
Rf Voltage	100	100	100	MV
Radio-frequency	360	360	360	MHz
Energy loss/turn	3678	2778	526	keV
Rms relative energy spread(collision)	15.6	18.0	39.0	10^{-6}
Fill time	16.3	16.3	28	min.
Acceleration time	5.8	7.6	35.9	min.
Total time: fill and accelerate	22.1	24	63.9	min.
Longitudinal impedance threshold: $\frac{Z_{\|\|}}{n}$(collision)	3.6	2.7	1.1	Ω
Transverse impedance threshold: $Z_{\perp}$(injection)	731	635	250	MΩ/m
Resistive-wall transverse impedance: $Z_{RW}(\frac{c}{\sigma_s})$ (injection)	0.4	0.5	98	MΩ/m
Resistive-wall multibunch instability growth time	472	310	.36	turns
Total current	.05	.05	.09	Amp
Peak current(inj)	3.6	3.6	4.2	Amp
$<\beta>$	255	255	382	m
Tune	65	86	269	
Half cell length (assumed 90^ocells)	200	200	300	m
Beam pipe radius	1.65	1.65	1.0	cm
Beam pipe	Cold, Cu	Cold, Cu	Warm, Al	

C. Common Parameters: Center-of-Mass Energy and Injection Energy

At the Indiana workshop, 60 TeV (CM) was chosen, more or less arbitrarily, for the energy of the collider. One reason for this was that it was thought to be a reasonable choice for injection from the Tevatron, although a factor of 30 in dynamic range is high for a large synchrotron. Twenty is a more traditional factor.

It was suggested by numerous theorists and experimentalists at this workshop that a factor of four increase over the LHC energy was not sufficient to justify the expense of an RLHC. The choice of a higher energy would increase the reach, and decrease the need for very high luminosity, for at least some interesting processes [3]. Higher energy might help relieve the very serious issue of too many interactions per crossing. Arguments were developed [4] that indicated that 100 TeV (CM) might be the optimum energy if the collider were sited at Fermilab, because the Main Injector, a rapid cycling machine at 150 GeV, could inject directly into a new 3 TeV high energy booster (HEB). It seems unlikely that the Tevatron will be a good injector for any collider energy higher than 30 TeV per beam, and, in any case, it will be a very old machine by the time the RLHC is built. A site-filler at Fermilab could easily reach 3 TeV with magnets operating at 4.2K at 6.5T. The high injection energy also solves many of the beam stability problems [5] that a huge, low-field ring would have at 1 TeV injection.

Applying the rule of twenty brings one to a 120 TeV (CM) collider. A slight decrease from that energy may simplify the collider rings, particularly the correction and feedback systems, and thereby might reduce the cost. Besides, 100 TeV (CM) seems such a nice, round number. An interesting consequence of increasing the beam energy is that the beam emittance damping time is the same with 9.5 T magnets at 50 TeV per beam as it is with 12.2 T magnets at 30 TeV per beam. This opens the possibility of making a 50 TeV ring with synchrotron damping using LHC-like magnets operating at 1.9K.

II. DESCRIPTIONS OF THE DIFFERENT MACHINES

A. Low Field

The concept of using low-field superferric magnets for an RLHC has been around at least since a paper by R. Wilson at the 1982 Snowmass study [6]. Later, a superferric design at 3 T was one of the Reference Designs for the SSC [7]. The ideas used in the present approach have been developed primarily at Fermilab [2], and have been discussed in symposia at the 1996 Indianapolis APS meeting [8]. The essence of the approach is to use a novel, 2-in-1, combined function, superferric magnet, which would operate with a field of less than 2 T. (See Fig. 1).

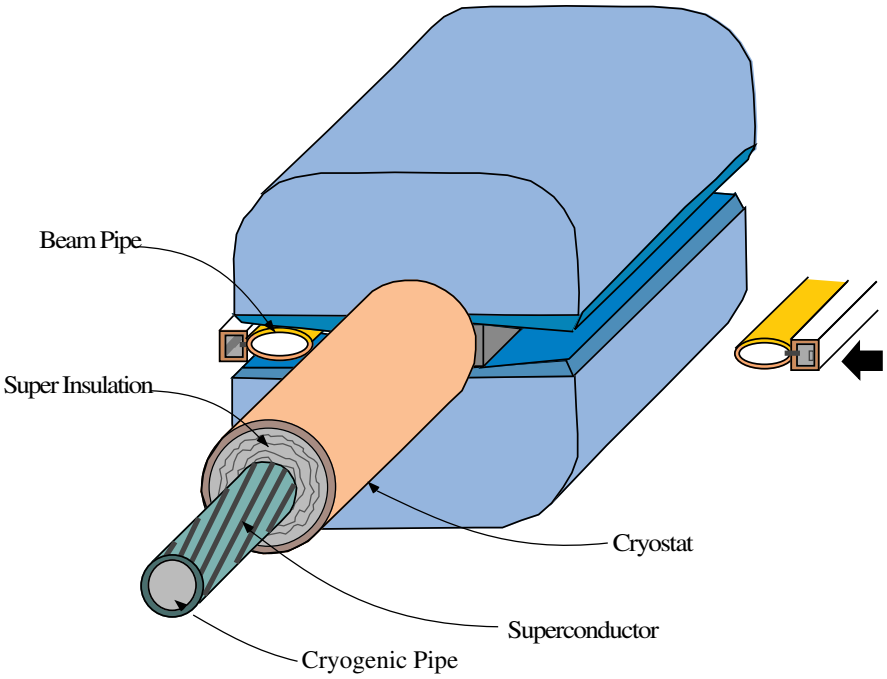

Fig. 1: Double-C twin bore transmission line magnet
The center conductor carries 75 – 100 kA and returns in a separate and nearby cryogenic line. The steel and beam tube are at room temperature. The pole pieces are shaped to provide both bending and alternating gradient focusing.

This magnet, called the “double-C transmission line magnet”, is further described later in this summary. It offers the possibility of dramatically reduced magnet system costs due to its simplicity, ease of construction, and small amount of superconductor relative to a cosine theta magnet. The transmission-line magnet is quite small in cross-section, perhaps 25 cm on a side, with cut-outs on each side to shape the field and receive the beam tube. Both the beam tube and the yoke steel are at ambient temperature. The current that creates the field goes down the center hole in the steel. Depending on the gaps, the peak current will be in the vicinity of 75 kA to 100 kA. The return current comes back in a separate but close-by cryogenic pipe, which will cause some stray field in the tunnel. Because of the simplicity of the magnet, the hope is that it could be made very long, thereby making it even less costly.

Even though this ring is very large, the cryogenics system is quite modest [9]. The conductor is at a force null, so the support structure for it has a low heat leak. Furthermore, the yoke steel is not cryogenic and the synchrotron radiation power is removed at room temperature, permitting a much smaller and simpler cryogenic system than for any high-field design. A preliminary design has been presented using NbTi as the conductor and liquid helium as the coolant [8]. Eight refrigerator plants, each roughly the size and power of an SSC plant are located around the ring. It is interesting to note that the transmission-line design is particularly well-suited to either effective use of pre-reacted Nb_3Sn operating at up to 10 K, presently commercially available high-temperature superconductor at 25 K, or, with some development success, the more speculative HTS deposited on structured substrates, operating at 77 K. In all these cases, operating costs would drop dramatically. The wall-plug power, for example, would be less than 5 MW if the transmission line could be operated at LN_2 temperature.

The choice of low field, however, forces the ring circumference be large (a 50 TeV per beam, 1.8 T ring has a 650 km circumference). For this to be affordable, one must develop and implement low-cost innovative tunneling and installation techniques. The suggestion of simply housing the accelerator in a pipe buried under the ground has led to the name "Pipetron" for this approach. So far, studies [10] have indicated that very small diameter tunnels (about one meter) with no human access will not be the least costly approach because of the high development cost. It appears that an optimum tunnel might be about three meters in diameter. Whether to include an invert, and how much development effort to devote to robotics, are open questions at this time. It seems likely that in this design, robotics could be used effectively to do some of the installation work, and perhaps even some of the maintenance.

The vacuum chamber is extruded aluminum at ambient room temperature, so the synchrotron radiation power does not have to be removed at cryogenic temperatures. The large circumference of the small aperture, warm bore vacuum chamber presents a substantial resistive wall impedance to the beam. Consequently, both multiple and single bunch stability issues must be considered carefully in the machine design. The pipe could be co-extruded with an ultra-pure aluminum coating to help suppress the high frequency single bunch instability.

A distributed pumping system (either ion pumps or getters) is integrated into the vacuum chamber [11,12]. The linear pumping speed needed to cope with the gas desorbed by synchrotron radiation is similar to that required for the PEP II High Energy Ring. If getters are used for the distributed pumping, lumped ion pumps must be installed to handle non-reactive gasses such as methane. This may limit the length of magnets or require a high-conductance antechamber attached to the beam tube.

The large circumference of this ring leads to large values for the number of particles required, and hence to a large value for the stored energy in the beam–almost 10 GJ, 11 times the amount of the high-field design. Beam-abort design, and safely handling single-point beam loss events in the arcs, will be formidable problems. The former was studied at this workshop [13]. It appears that with a sufficiently long straight section, about 1.5 km, reasonable values for septum magnetic fields, kicker strengths, and rise times are possible with a conventional design, such as in the LHC. The most challenging aspect of the design is to prevent the beam from melting the absorber material. Studies with MARS13 code [13] indicate that if graphite is used, a material with excellent properties for this purpose, the beam will have to be blown up to have a spot at the absorber of about 15 mm. This can be done with 50 m of 1 T/m vertically focusing conventional quadrupoles placed about half-way along the 5 km beam line to the absorber. At the same time, the beam will have to be swept back and forth with a 40 cm amplitude, and moved vertically across the absorber face, sort of like the electron beam in a television picture tube, to keep the peak temperature in the absorber to less than 1500° C. A graphite absorber 10 m in length is adequate.

A second type of abort design that was studied involved filling the abort beam line with air to gradually absorb the energy of the proton beam. The fear is that the beam will create a partial vacuum in the air by heating it. Using MARS13 again, it was determined that the beam spot had to be 15 mm in diameter when it was incident on the air-containing pipe, and a 5 m graphite absorber still had to be placed at the end of the 6 km beam line. It does not appear that such a novel design offers any advantage over the graphite absorber by itself.

It does seem that kicking the beam out of the machine in two or more locations would be advantageous.

The low magnetic field of this machine is an advantage for keeping the synchrotron radiation power low, but is a disadvantage for emittance preservation, and places a premium on preserving the emittance during injection and acceleration. This was a major issue in the SSC design, and will be difficult to realize in the low-field RLHC, as well.

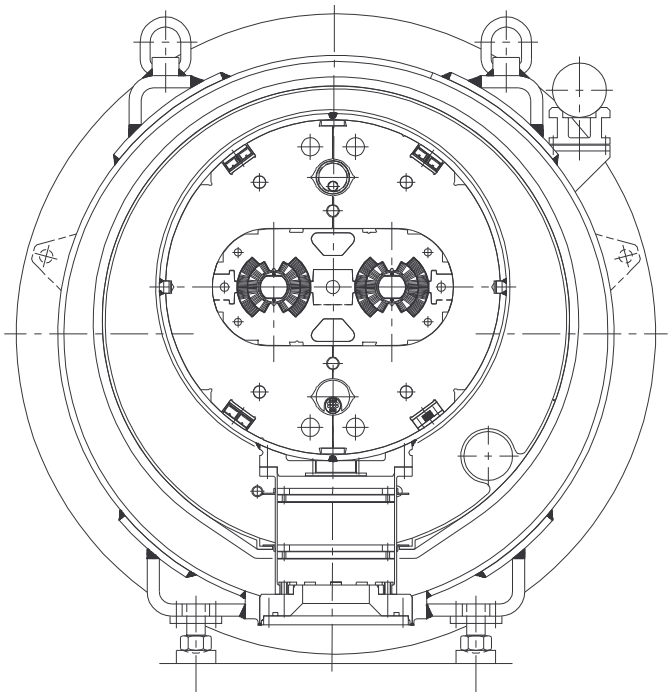

Fig. 2 LHC dipole

B. High Field with Known Technology

It was shown at the 1994 Indiana workshop [1] that transverse damping times of two to four hours improved the integrated luminosity of the collider and essentially erased the impact of beam emittance dilution that occurred during injection and acceleration. Increasing the energy of the rings from 30 TeV to 50 TeV makes it possible to have sufficient transverse damping at magnetic fields as low as 9 T to accomplish the useful functions made possible by synchrotron radiation. This encouraged us to pursue a design based on the mature technology of NbTi conductor and cosine theta coil geometry, similar to the Tevatron, HERA, RHIC, SSC and LHC. The challenges presented by such a machine are completely different from those of either the low-field machine, where they are related to beam instabilities and tunneling costs, or the very high-field machines, where new technology has to be developed and made practical. In the region of B≈10 T the issues are related to perfecting an already reasonably well-understood design and applying the best engineering practices to make it less expensive. The major cost issues will be the magnet systems, of course, and superfluid helium cryogenics. Both the SSC and LHC organizations have studied the machine designs and accelerator physics rather completely at only slightly lower fields.

At 9.5 T, the collider ring is 138 km in circumference, somewhat larger than the SSC ring. The magnet will look similar to the LHC design. (see Fig. 2). For a high-field magnet, a cosine theta coil following the outline of intersecting ellipses is the most efficient use of superconductor, and a 2-in-1 design is certainly less costly than separate cryostats. The design challenges are due to the great forces, approaching 100 MPa, that bear on the conductor, and the initial cost, operating cost and reliability of operating at 1.9 K. The optimization of NbTi strand for high field by-and-large ended when the SSC collapsed. A program has recently been started by Fermilab with commercial superconducting wire manufacturers and the University of Wisconsin to support the Fermilab program for LHC high-gradient insertion quadrupoles.

Most of the accelerator systems for this collider are similar to those of the LHC. It seems that there are no particular problems with instabilities. The beam tube would have a copper-clad liner at about 20 K with holes or slots to pump the desorbed hydrogen onto the 1.9 K cryostat surface. At this temperature, the vapor pressure of hydrogen is not an issue. One of the issues at the LHC is emittance preservation. This is much less of a problem at a high-field RLHC due to synchrotron damping. Quench protection would be accomplished with an active system detecting the quench and firing heaters in the magnet. The current will be bypassed with cold diodes in the magnets.

The RLHC cryogenic system will be significantly larger than any previously built. A preliminary design [14] from this workshop has 20 refrigerators around the ring, each roughly equivalent to an SSC plant. The cryogenics system will be five to 10 times larger than the LHC system, depending on how successfully the heat leak can be reduced from previous designs. This is probably the single most difficult problem in this design, even more important than the magnets. The size and complexity of the cryogenic system will make reliability an issue, and the required redundancy will drive up the capital cost. The operating cost will be high, as well, since using NbTi at 10 T will demand superfluid helium coolant.

On the whole, a design of this type is straightforward. There are few accelerator physics or engineering issues that have not already been rather completely studied. The challenge is bringing down the cost and complexity of the machine.

C. High field with New Technologies

The use of high-field magnets for a 30 TeV per beam collider was promoted at the Indianapolis workshop. Although not necessary for synchrotron damping at 50 TeV per beam, the compact ring size and use of new materials and innovative designs for the magnets make this interesting to pursue. Among the most useful features of such a ring is the possibility of using either Nb_3Sn magnets operating at 4.2 K or HTS materials operating at even higher temperatures. Either of these choices would go a long way to reducing the size, complexity and cost of the cryogenic system compared to that which would have to be used for NbTi high-field magnets.

The major challenge for this design is obviously the magnets. Nb_3Sn has been in use for as long as NbTi, but it has not yet succeeded as a useful engineering material, mostly because it is very brittle when reacted. In addition, its current carrying capability at high field has not been very impressive. It may be that new methods of artificially forming pinning centers (APC) will improve the critical current density. Recently, an accelerator-style magnet has been built and operated successfully at 10 T and 4.2 K by the University of Twente [15], but it is not clear that the techniques used could bring Nb_3Sn magnets to the point of practical production. A magnet designed to attain a field of 13 T is being built at LBNL, and will be tested in a few months.

Even less has been done with high-temperature superconductor. There are engineering materials that can be ordered and delivered in tape form in lengths of about 1 km. These tapes can carry only 100 A to 200 A at 4 K to 20 K, but high-current cables of the type that could be used in accelerator magnets do not yet exist. Very short lengths of tape, about 1 cm to perhaps 20 cm, comprising YBCO deposited as a thin film (about 1 μm) on a 25 μm substrate, have shown remarkable promise. Critical current densities as high as 20,000 A/mm^2 at 75 K, almost independent of magnetic field, have been measured [16]. The engineering current density of these tapes is less than 100A/mm^2, however, due to the thinness of the film. We await developments.

Regardless of the material developments, magnets with field strength in the neighborhood of 15 T and above will be very expensive and complicated. The traditional coil shape of intersecting ellipses will have to be abandoned because of the huge forces. Placing the conductor in blocks that can resist the forces requires much more superconductor than the efficient cosine-theta design. Furthermore, the amount of steel becomes so great that it may actually be more cost-effective to use superconductor to shield the field at the outside of the return yoke.

The development of such strong magnets is exciting, although there is a great deal of progress which needs to be made before they will yield a less costly overall collider design. Even if they never become cost-competitive with other technologies, they may be useful in cases where a finished tunnel already exists, and the goal is to reach the highest energy. They may also be useful as specialty magnets, such as insertion quadrupoles or beam splitters, where strength is often the most important criterion, and the small number of magnets make the cost less relevant.

D. Dynamics

Much has been learned of the particle dynamics issues associated with large hadron colliders from the work performed on the SSC and LHC projects. The issues for RLHC fall into the same categories as for these other projects, though the emphasis may be different. For example, the stored beam energy in the SSC and LHC is quite formidable (440 MJ and 377 MJ respectively) and has required an intense amount of study and new ideas on how to handle this amount of stored energy. For the mid- and high-field RLHC options, the stored energy is comparable (though somewhat larger) to the SSC and LHC designs, and emphasis is placed at reducing the number of protons to help keep this issue under control. On the other hand, the low-field option requires a much larger circumference and much larger number of stored protons, and hence the stored energy is actually many gigajoules, making this issue even more pronounced. Thus, each RLHC option presents its own set of interesting beam dynamics issues and parameter optimization choices. Much of the initial work on the RLHC accelerator physics began at the Indiana Workshop [1], with further progress being undertaken at Fermilab on the low-field option [2] and at Brookhaven on the higher field options [17].

The most prominent feature of the mid- and high-field options is the synchrotron radiation damping. At 50 TeV and with strong magnetic fields, damping times much smaller than the storage time can be realized and hence the luminosity of the collider is enhanced for relatively modest beam intensities. Fig. 3 shows the evolution of the luminosity in the high-field case during the first 10 hours of a store. The model includes radiation damping and quantum excitation, as well as intra-beam scattering and beam-gas scattering. One can see that 10 hours is more or less the useful storage time, and thus this option would require more frequent stores and fairly efficient shot set-up time. On the other hand, the injector requirements for this option would not be nearly as demanding as for the SSC and LHC. The emittance preservation throughout the injector chain would no longer be such an issue. To obtain a luminosity of 10^{33} cm^{-2} sec^{-1}, the SSC injector

system was required to provide an initial normalized rms emittance of 0.8 π mm-mr, 3-4 times smaller than the emittance routinely realized in present day proton collider facilities. The performance of the medium and high-field RLHC options, on the other hand, would not be hindered by larger emittances, so long as the particles remain in the machine, because the emittance quickly damps to its equilibrium value. Fig. 4 illustrates this point, showing the integrated luminosity of the three colliders as a function of initial emittance.

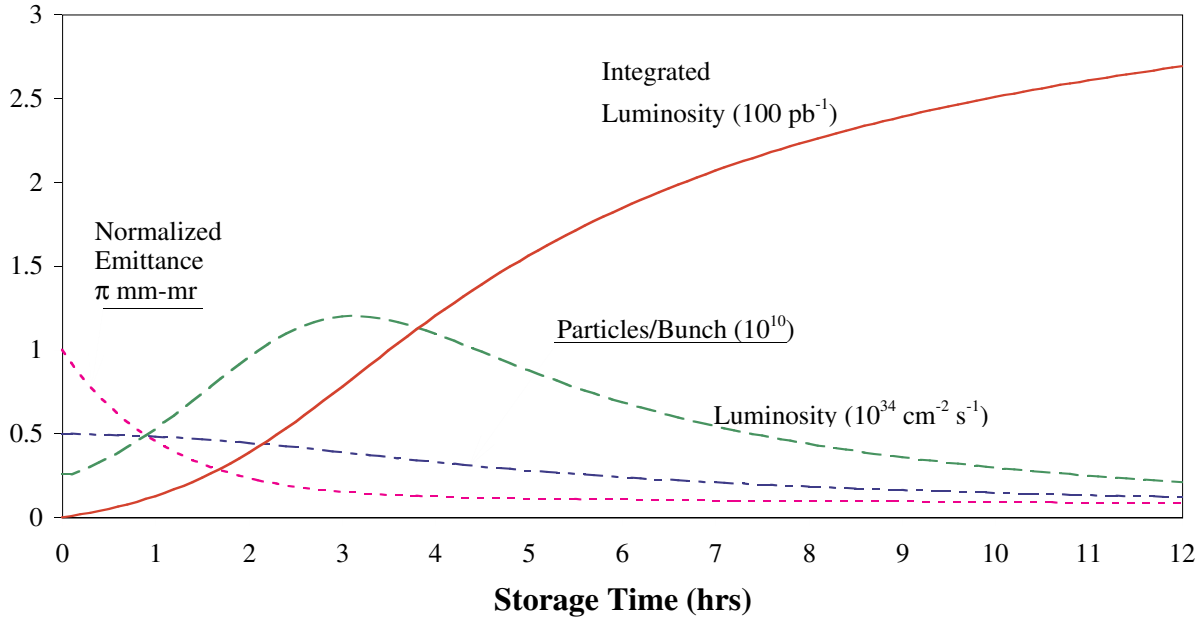

Fig. 3: Beam parameters during a store for high-field RLHC.

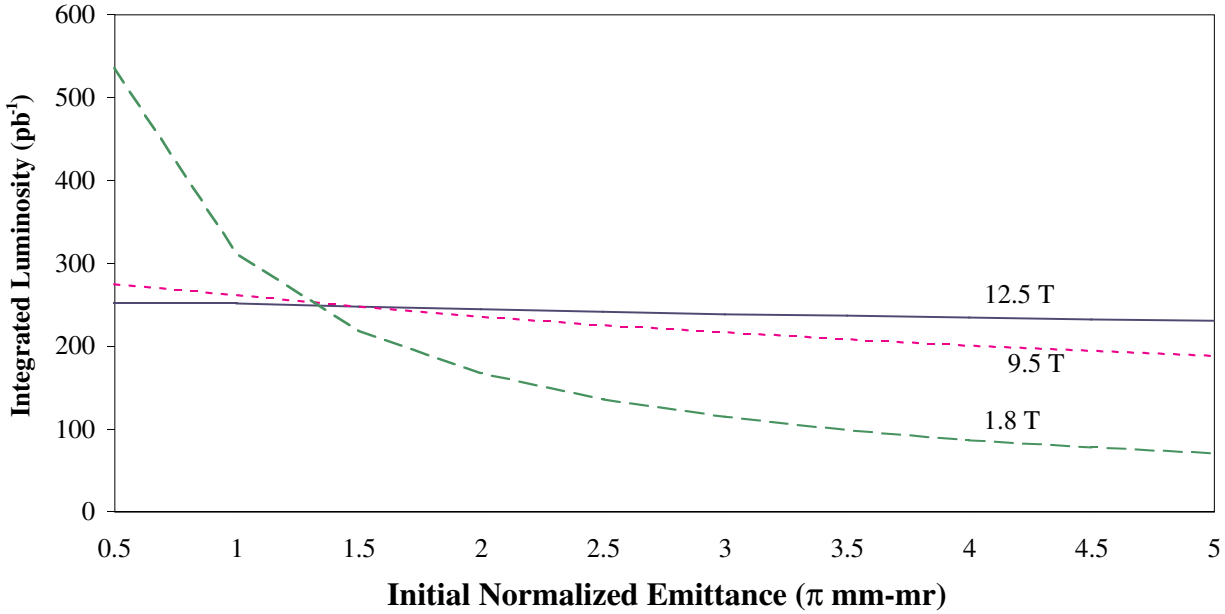

Fig. 4: Integrated luminosity of 10 hour store vs. initial rms emittance for RLHC options. The integrated luminosity of the two high-field cases is almost independent of the initial emittance because of synchrotron radiation damping.

Further enhancements can be made to the damping times and equilibrium emittances by playing various games with the accelerator lattice. Longer cells may be desired from the point of view of cost as well as the desire to reduce the chromatic corrections necessary in the machine; this will be discussed further below. However, longer cells also increase the equilibrium emittance and thus decrease the integrated luminosity. In addition, one can trade off damping partition between the various degrees of freedom, generating a shorter damping time for the horizontal emittance at the expense of the longitudinal damping time. This can be accomplished, for example, by introducing a small gradient in the bending magnets or by offsetting the orbit through the quadrupole magnets in the ring [1].

It should be pointed out that the present design of the low field RLHC lattice uses combined function magnets which inherently leads to anti-damping in the horizontal degree of freedom. Since the damping times in this large circumference ring are on the order of 4 days or so, this may not play a large role in the performance of the machine over the span of a normal store. However, the effects may be noticeable and a final lattice of this machine needs to be optimized with this feature in mind. Neglecting this effect, the low-field RLHC would maintain an average luminosity of roughly 10^{34} cm^{-2} sec^{-1} throughout a 20 hour store, while the mid- and high-field options would require two shots per day to acquire the same level of integrated luminosity.

The half cell length for the RLHC affects many of its features, including the equilibrium emittance, the dynamic aperture and tolerable magnet field errors, chromaticity correction and other machine correction systems, not to mention its cost. In present large hadron accelerators, such as RHIC, HERA, and LHC, systematic errors appear to dominate random errors in the particle dynamics. It was shown by Holmes and by Peggs, et al. [18], that while shorter cells allow larger systematic errors in the bending magnets, quite a bit longer cells (several hundreds of meters) may also be tolerable. Additionally, longer cells generate larger dispersion, which can be used -- so long as the momentum spread is maintained at a tolerably low level -- to reduce the strength of sextupole correctors in the ring. Short cell and long cell lattices have been developed and dynamic apertures computed via particle tracking to show that the longer cells greatly increase the dynamic aperture caused by the correction sextupoles in the machine [19].

As in the SSC and LHC designs, much of the interesting beam dynamics will be dictated by the interaction regions. The field quality requirements of the IR triplet quadrupoles are enhanced by the extremely large amplitude functions (40-69 km in typical designs). In addition, the strong focusing in this region generates tight alignment requirements of the triplet quadrupoles and enhanced sensitivity to ground motion. In designing the triplet layout, one must also be aware of the power delivered to the triplet from the debris generated at the interaction point. In the high field case, this amounts to a peak power of approximately 12 kW emanating from the IP. The triplet magnet design and layout has to accommodate this as well as deliver the proper focusing characteristics. Obviously, there must be much interchange of ideas between magnet designers and beam dynamicists during the design of the IRs.

Studies have also been performed to look at the linear and nonlinear chromatic effects generated by the low β insertions. [19] Long half cells in the arc design would allow the sextupole correctors in the arcs to be used to control the chromatic effects generated by the triplets with β^* values down to 10 cm. However, much shorter half cells may require local control of the chromaticity in the IR regions, especially to handle second order effects.

Coherent instabilities are worrisome in the RLHC, especially for the low field option. The low field machine has a lower single bunch impedance threshold than the high field options. Transverse mode coupling and resistive wall instabilities thresholds of the high field accelerators are typical of SSC/LHC parameters, whereas the low field RLHC has an impedance which is an order of magnitude larger. This is primarily due to the large circumference of the accelerator and small beam pipe aperture relative to the beam size. For the high field option, multibunch stability can be handled by present day feedback technology, while the low field collider will have resistive wall multibunch instability growth times of less than a single turn [5]. This will require a challenging feedback system.

While much initial work has been performed, there are certainly many future accelerator physics issues to study which were identified at Snowmass. For example, further optimization of the half cell length should be performed with the interplay of correction systems, magnet tolerances, and enhancement of the synchrotron radiation effects with the lattice design. Future modeling of the time evolution of the collider's luminosity and performance should include more realistic coupling between the transverse emittances. Crab crossing may be a viable scheme to enhance the useful luminosity, as well as variable β^* schemes which could help to spread out the luminosity and interaction rate over time. Some work has already begun on this issue [22]. Most of the future studies mentioned in this paragraph are more relevant to the high field designs. The issues more relevant to the low field design are in the areas of beam stability and stored energy.

II. ACCELERATOR SUBSYSTEMS

A. Magnets

Magnets are the single most expensive component of any large hadron collider laboratory, accounting for as much as 30% of the cost. Hence, magnets for the RLHC

have to be developed that can be economically and reliably mass produced. The technology of superconducting magnets today naturally breaks the designs into three general types, depending on their field strength:

1. Low-field, usually superferric, in which the field is shaped by steel and the coil is superconducting. These are limited to 2T by saturation of the steel.

2. Moderate field, usually with cosine theta coils, in which the field shape is determined by the location of the conductor. These have generally been built from NbTi material, which is strong and ductile but which limits the field strength to about 10 T.

3. High-field magnets, which require materials with critical fields greater than 15 T and arrangements of conductor to allow for the very large forces present in high-field magnets. The traditional material for these magnets has been Nb_3Sn, but recent progress in high temperature superconductors may show them to be a promising competitor.

1. Low-field magnets, B≤2T

Superferric magnets can be built in a classical H-magnet style similar to the two-in-one low-field design proposed by Huson for the SSC [7]. The low-field design studied at Snowmass was one proposed by Foster [20] for the Pipetron. As shown in Fig. 1, it features a strikingly simple, low cost superferric (1.5-2 T) combined function dipole called the "double-C transmission line" magnet. The magnet drive current is provided by a cylindrical conductor carrying 75-100 kA of supercurrent; this conductor is surrounded by an iron yoke in such a way as naturally to provide a double bore magnet suitable for a proton-proton collider. The conductor is very similar to those used in superconducting transmission lines, and ideally would be fabricated from high temperature superconductor helically wrapped on the cryopipe. It could also be made from conventional NbTi conductor, operating at 5 K or Nb_3Sn operating at 10 K. The location of the conductor at a force null eases many of the mechanical design and heat leak problems. The poles are shaped to provide a gradient as well as a dipole field. The current is returned on a bus located in a separate but nearby cryostat above the double-C magnet in the tunnel.

The field quality in the gaps is determined by the shape and magnetic properties of the pole pieces. Because of the cylindrical symmetry, separate currents for quadrupoles and correction magnets cannot share the space occupied by the transmission line. Hence, it is a great advantage to have a combined function (dipole + quadrupole) magnet. This makes reaching a bend field of 2 T difficult, because one edge of the pole saturates before the other. Nevertheless, it might be possible to reach 2 T bend field by judicious use of crenelations in the steel and by including some high-performance magnetic material in the poles.

This magnet design has many positive features. It is potentially a simple magnet with only one cryogenic and three separate vacuum connections between what could be very long magnets. The only cryogenic part of the magnet is the transmission line itself; both the steel and the beam pipes are at ambient room temperature. Since the steel does not have to be cooled and the power from synchrotron radiation is not absorbed at cryogenic temperatures, the refrigerator and its associated cryogenic system might be much smaller and simpler than in high-field magnets. Furthermore, the transmission line is in a relatively low-field region and experiences small forces, allowing a low heat leak support structure. It might be that cryogenic stability can be obtained even if the system were cooled with gas. This opens the range of superconductors that could be used. In addition to NbTi conductor cooled to 4 K or 5 K, one could use Nb_3Sn tape spiral wrapped around a tube cooled by 10 K helium gas. The gentle bends encountered in such a design might permit the Nb_3Sn to be reacted before fabrication, which would be a great advantage. Existing powder-in-tube HTS tape could be used effectively at 25 K. It might even be possible to plate the more advanced YBCO HTS in a spiral pattern directly on a substrate that is part of a cryogen-carrying tube. Of course, this last option awaits advances in the production of YBCO superconductor. One of the best features of this design is that the superconductor is exactly a DC transmission line, leading one to hope that power companies might help defray the cost of the development.

The major disadvantages of double-C magnet are associated with its low field and resulting large circumference ring – a 50 TeV per beam collider would be about 650 km around. This leads to potential problems in beam stability, difficulty in feedback, very high voltage rf systems and a very large stored energy in the beam. Also, although it is generally thought that the room temperature beam tube is an advantage, it does mean that there must be a significant vacuum system for the beam tube, an added expense and complication. These issues are discussed elsewhere in the summary.

Since no double-C magnet has been built and tested, and very little engineering has, as yet, gone into the details, one should approach the design with caution. Saturation effects might require more steel, complicated shapes, advanced materials and sophisticated correction schemes that could drive up the cost. These items are all part of the estimated costs of higher field accelerators, and could be required in a Pipetron, as well. That will be determined in the coming years as the R&D, design and engineering effort increase.

2. Moderate field magnets, 4T≤B≤10T

Magnets in this field range can be built now using existing and well-understood cosine theta coils of NbTi cable, operating at 4.5 K for fields below about 7 T and at 1.9 K up to 9 T. It is likely that a few more years of

development will yield 10 T magnets of similar designs. It may be that a moderate-field, conductor dominated magnet, similar to the RHIC designs (4 T≤B≤6 T) would be sufficiently inexpensive that the overall cost of an RLHC with such a magnet would be minimized. At that field level, damping by synchrotron radiation would be insignificant, and, hence, a magnet in this moderate-field range was not considered. High-field, for this study, was defined as a field high enough to cause radiation damping times of a few hours.

The major advantages of a cosine-theta magnet are that it is well understood, that it provides for a smaller circumference ring than the superferric design, and that at the highest field strengths, above 9T, it creates enough synchrotron radiation to lead to damping of the beam emittance. Similar magnets have been used in the Tevatron and HERA, and have been engineered and prototyped successfully for SSC and LHC. A typical two-aperture-in-one-cryostat design (for LHC) is shown in Figure 2. Since the magnets are well understood, the cost of building them can be accurately predicted, and the R&D can be accurately focused. The transverse damping time constant for a 10 T field in a 50 TeV ring is 3.8 hr. There appear to be no important instabilities in a 50 TeV ring made from magnets with B≥6 T for an injection energy of 3 TeV. All in all, these magnets are a logical and moderate extension of existing technology.

The major disadvantage of intermediate field magnets is that they are well known, and so are their weaknesses. Relative to superferric magnets they use a lot of superconductor with its high cost and high persistent currents, which need to be compensated. The current-dominated design may lead to lower quality field, but it was shown in the SSC R&D phase that the field quality that was predicted from earlier magnets (like the Tevatron) were at least a factor of three worse than the realized quality. This improvement was a result of better design and quality control of strand, cable and assembly tooling. In addition, new methods of improving the low-field multipoles have been developed recently at BNL, and are used in RHIC IR quadrupoles to decrease random errors [21].

The steel and beam tubes are at cryogenic temperature in these designs, so the refrigerators must be sized to cool down the magnet system in a reasonable time, and absorb the synchrotron power. The beam tubes have a liner that absorbs the synchrotron power (about 140kW/beam for a 50 TeV/beam machine) at a temperature higher than the coil temperature. This saves power, but is a design complication, although it greatly simplifies the beam vacuum system. If the magnets need to be operated in superfluid helium, the cryogenic system will be very expensive both in capital cost and operating cost.

Finally, there are objections to the concept of a design that is basically boring. It is not necessary to do challenging R&D to solve new and interesting problems. Nor does such a machine require new materials that would drive the nascent HTS industry. On the other hand, it is in just such a situation that engineering can concentrate on decreasing costs and increasing reliability.

3. High-field magnets, B≥10T

There are at least three possible technologies for reaching this field level: Nb_3Sn (or Nb_3Al), operating near 4.5 K; Powder-in-tube or dip-coated HTS, operating at 4 K to 25 K; and YBCO (or perhaps other) HTS materials deposited on aligned textured substrates operating at 20 K or possibly higher. The amount of work and the level of development is more advanced in Nb_3Sn, but the greater promise of HTS, particularly the very high critical field and the higher operating temperature makes it very tempting to pursue.

The huge Lorentz forces in the magnet make cosine theta coils unfavored for very high fields because of the buildup of high forces at the midplane, and because the forces are often in a direction that makes the cable mechanically unstable. As a result, very high field magnets often have their cables arranged in blocks, sometimes with inner support rings, so that the forces are easily transmitted to the massive steel yoke. Another type of design uses parallel current walls, which results in a uniform field, but has difficult problems at the ends. A particularly interesting design uses all three types of superconductor, depending on the field that the superconductor is in – HTS at the inside where the field is highest, Nb_3Sn further out, and NbTi on the outside [23]. This magnet, shown in Fig. 5, reduces the amount of return steel by using the NbTi conductor on the outside of the steel to buck out the leakage fringe field. Otherwise, the amount of steel is prohibitive. The conductor is arranged so that the field is always in a direction that produces no torque. The HTS is also oriented in the direction that permits the highest J_c.

Whatever design is used, these high field magnets will use large amounts of conductor compared to cosine-theta coils, because cosine theta coils are very efficient in their use of conductor. Block designs need up to twice the superconductor to reach the same field.

Table II shows comparisons among Nb_3Sn, BSCCO 2223, and YBCO epitaxially grown on a biaxially textured substrate. It is easy to see why the HTS attracts attention. The BSCCO is essentially an engineering material now. Its major problems, which might yield to industrial R&D in the near future, are the strongly directional nature and steep field dependence of the critical current, particularly above 30K. (It is important to note that this is not a problem for superferric magnet designs.) The major problem of YBCO is that it is still very much a laboratory material, with no perfected production techniques. YBCO also has a poor engineering current density, since it is a thin film of the order of 1 μm, deposited on a much thicker substrate. Whereas Nb_3Sn typically will have a copper-to-superconductor ratio of about three, and, hence the engineering current density is

about 25% of the superconductor current density, the substrate of YBCO is 25 times greater in area than the superconductor, so the engineering current density is only four percent of the superconductor critical current density. Another problem of all the HTS materials is that they presently are made in tape form, and can only carry the order of 100 A. In order to make useful accelerator magnets, high-current cables will have to be made.

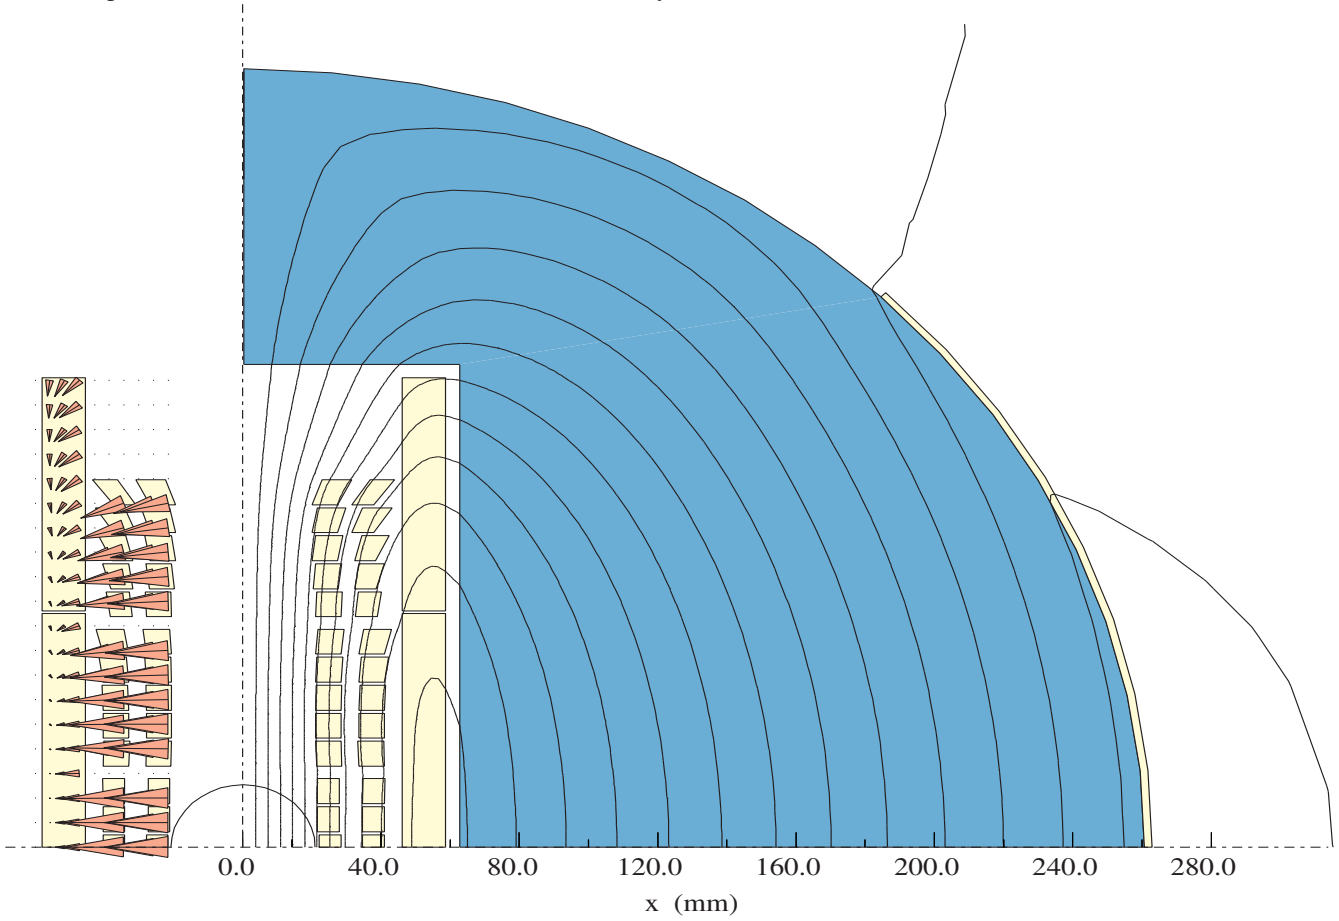

Fig. 5: High-field magnet cross section [23]
The conductor nearest the beam tube is HTS, outside of that is Nb_3Sn, and nearest the yoke is NbTi, which is returned outside the steel, to buck out the fringe field. The arrows on the left are proportional in size to the force on a conductor block, and in the force direction. Note that each conductor block is arranged so there is no torque on it.

Table II: Comparison of superconductors

Property	NbTi	Nb_3Sn	BSCCO-2223	YBCO
Upper critical magnetic field	15 T	25 T	≈ 100 T	≈ 100 T
Critical temperature	9.5 K	18 K	110 K	92 K
Critical current density	2-2.3 kA/mm^2 (7T&4.2K or 10T&1.8K)	1-2.4 kA/mm^2 (10T&4.2K)	<0.9kA/mm^2 (20T&20K)	<2.4 kA/mm^2 (20T&77K)
SC volume fraction	~40-50%	~ 35-40%	~ 35-40%	~ 4%
Conductor type	multifilament wire	multifilament wire	multifilament tape	microbridge
Mechanical property	Ductile	Brittle	Brittle	Brittle
Longest piece made	~10km	> km	≈ km	~ 10mm

B. Tunneling

In the low field case, the large ring circumference implies substantial expenditures on the civil and installation efforts unless innovative, inexpensive tunneling, installation and maintenance technologies are used. In past studies, particularly for the SSC, a shallow minimum in overall cost was found at a field strength of 5 T to 7 T, using conventional tunnel costs. Bringing down the overall cost of the low-field designs will depend on developing techniques to make less expensive tunnels. Progress was made toward understanding prospects for the development of such technologies through discussions at Snowmass with geologists and through a site visit with tunneling experts at the Colorado School of Mines.

The geology of the Illinois site for an RLHC was discussed by David Gross [24]. The Dolomite rock deposits located 100-150 m below the surface can provide an excellent tunneling medium. The rock conditions are predictable and homogeneous, and the area is very stable seismically. At these depths, the environment is virtually vibration-free, and has no settlement problems.

The site visit to the Colorado School of Mines explored various options in tunneling technology: specifically, directional drilling, microtunneling, and the tunnel boring machine (TBM). These options are outlined in the following paragraphs.

Directional drilling uses a small, guided drill head and is generally used in soil for small (<30 cm) diameter holes. Rock penetration rates are slow, and accuracy is limited (30-60 cm).

In standard microtunneling, a cutting head is pushed forward by a pipe jacked by a hydraulic cylinder. The pipe sections are typically 3 m long and up to 3 m in diameter. Jacking stations (shafts) are also required every kilometer. Enhanced microtunneling, which is in the early stages of commercial development, eliminates the pipe jacking by using a gripper to provide forward thrust. With such a system, it would be possible to go 6 km between shafts.

The conventional large-tunnel excavation device is the TBM. These are large, expensive devices manned by a crew of 6-12 people. Rotating cutters, mounted on the front of the machine, bore through the rock. These cutters typically require a great deal of maintenance. Such machines have been built in sizes from 3-12 m in diameter and were used for the SSC tunnels. Typical rock penetration rates are 9 m/hr.

Prior to the site visit to the School of Mines, it was thought that a 1 m diameter tunnel, excavated using microtunneling or directional drilling techniques, would be the most cost-effective approach for a low-field RLHC. However, after the visit, it became clear that the more conventional TBM technology, with improvements to reduce cost, would most likely be the route to take. This would provide a 3 m diameter tunnel, which would be accessible to humans. Some aspects of installation and maintenance might still be performed using robotics.

The present cost of tunnels created with TBMs is about $3000/m. The cost as a function of the tunnel diameter minimizes at about 4 m. Prospects for reducing the cost minimum to $1000/m (with a 3 m shaft diameter) through developments in TBM technology were discussed at the site visit. The developments include automated steering and power thrust control; automated cutter changing, and/or improved cutters requiring less maintenance; optimized cutter positioning; continuous boring without a regrip cycle; improved instrumentation for failure monitoring; and improvements in muck-removal conveyor belt systems. These items provide a basis for a modest R&D program in TBM technology which, if successful, could dramatically reduce the tunneling costs of any accelerator which must be built in underground tunnels.

C. Beam tube vacuum

Beam tube vacuum was discussed for low field (about 2 T) and high field (12.5 T) options for a 50 TeV-on-50 TeV proton collider. The vacuum issues for a 10 T design are similar, but not quite so difficult as in the high field option.

The bound on the average beam tube gas pressure is set by the desire to have the luminosity lifetime dominated by scattering at the interaction points, about 65 hr for the low-field design and 16 hr for the high-field design. If we require that the luminosity lifetime due to beam gas scattering be five times that due to scattering at the interaction points, then the bounds on beam tube pressures are $P(CO) < 3.5 \times 10^{-9}$ Torr for the high-field option and $P(CO) < 0.25 \times 10^{-9}$ Torr for low-field, room temperature equivalent. The distributed pumping speeds necessary to realize these pressures are reasonable to achieve in both cases [11].

For both the high-field and low-field options the product of beam current and ion desorption coefficient is more than an order of magnitude below the threshold for instability. Similarly both options seem to be free of beam induced electron multipactoring.

For both field options the dominant source of gas to be pumped is photodesorption by synchrotron radiation, requiring some type of distributed pumping. Photon stops and discrete pumping do not seem to be viable options. The beam tube vacuum systems for the two field levels differ considerably owing to the different magnet geometries and beam tube temperatures. Firstly, for the low-field option the beam tube is warm and synchrotron radiation is absorbed at room temperature which is an attractive feature. A second attractive feature is that the mechanical assembly of the beam tube vacuum system might be decoupled from magnet assembly. The C type geometry of the low-field magnet allows the distributed pumping antechamber to be located outside the magnet or possibly in the fringe fields if

ion pumping is utilized. In principal, non-evaporable getters (NEGs), distributed ion pumps (DIPs), distributed titanium sublimation pumps (TSPs) and distributed cryopumps could all be utilized in an antechamber configuration. Although detailed cost analyses of these systems were not done it seems obvious that the NEG option would be the most cost effective and that is the one that has been studied in the most detail [12]. Allowance needs to be made for pumping photodesorbed CH_4 which is not pumped by NEGs. One possibility is to allow for a fairly large cross section antechamber outside the magnet iron which has enough axial conductance that the CH_4 can be pumped with discrete ion pumps or cryopumps. A second possibility is to precondition the beam tubes with electron desorption thus removing the bulk of photodesorbable gas before operation. Ishimaru has discussed an aluminum extruded beam tube vacuum system with a NEG pumping strip [12]. An all welded design is proposed with bellows and flanges eliminated to reduce cost.

Discussion of magnet options for the high-field option run the gamut from 1.8 K NbTi , 4.5 K Nb_3Sn to high temperature superconductor operating anywhere from 4 K to 30 K. In all of these cases the magnet bore tube is inside the superconductor and is at the same temperature as the superconductor. Distributed cryopumping is the only high-field option. At the lowest temperatures (1.8 K- 4.5 K) one is led to using a warm (~ 10 - 20 K), perforated beam screen inside the magnet bore tube, similar to the design for the LHC [25], to avoid the high cost of absorbing the synchrotron radiation power at low temperature. Above about 3 K the saturation pressure of the H_2 isotherm is too high for accelerator operation and it is necessary to add cryosorber material (e.g. charcoal) to the magnet bore tube to increase the effective surface area and prolong the time to reach saturation. For higher magnet temperatures (~ 10 - 20 K) it is perhaps reasonable to absorb the synchrotron radiation at the magnet bore tube temperature. Then one can consider a magnet bore tube with co-extruded pumping channels taking the place of the beam screen [25]. This eliminates the complexity of a two temperature system and the cooling lines for the beam screen. Above about 20 K (the precise temperature isn't known) H_2 will cease to be cryopumped effectively. So far there is no known solution to this problem for the highest temperature high-field magnets that have been discussed.

D. Cryogenic Systems

The preliminary studies of the cryogenic systems for RLHC carried out at the Snowmass workshop are particularly interesting in the variety of solutions in both concept and scale that exist for the different machines. [9, 14] The high-field machine of known technology, that is, NbTi in superfluid helium, has been well studied and engineered at LHC. Scaling from that design is completely reasonable, although the result is somewhat frightening. Although magnets made from Nb_3Sn are speculative at this time, from a cryogenic point of view an appropriate system operating at 4.5 K can be reliably scaled from SSC designs and refrigerator experience. In contrast to these well-understood systems, the low-field designs at any of the temperatures considered, 4.5–5 K (NbTi), 4.5-6.5 K (Nb_3Sn), or 20-25 K (HTS), and the high-field machine using HTS at 20 K are much less well studied. One of the major reasons to pursue such designs is the likely possibility that they will require much smaller cryogenic systems and result in much lower operating costs. Nevertheless, the early state of the engineering in all systems for these new and innovative concepts signals us to be cautious in accepting the results of these preliminary studies. Neither a low-field nor a high-field design at temperatures higher than 25 K were considered, although that is surely the most elegant realization of the low-field design, leading to an almost trivial cryogenic plant.

A short description of each system particular to both field levels and each operating temperature is given below, followed by Table III, which describes the high-level parameters of all the systems. In all cases, the flow required to cool the leads assumes leads made of HTS.

1. The low-field designs

The magnet for the low-field machine has the potential for having a very low heat load, because the cold mass comprises only the 75 kA DC transmission line, which is at a force null in the magnet and returns in a force-free cryogenic pipe. These conditions allow a very low heat leak support structure. Furthermore, in this design the synchrotron radiation is absorbed in a room-temperature beam tube, significantly decreasing the operating load, and in the particular design studied at Snowmass, the double-C transmission line magnet, the steel is also at room temperature, which simplifies cool down and allows much smaller refrigeration plants. A calculation of the forces on the conductor indicates that supports can be spaced slightly closer than 1 m. That, along with dimpled superinsulation results in a static heat load for the supply and return transmission line of 0.2 W/m. There may be considerable opportunity to decrease this number with R&D in the coming years.

The low heat load in this model suggests that eight refrigeration plants spaced about 80 km apart, plus one additional plant for the interaction insertions will be sufficient for this design. At each refrigeration plant, the cryogenic fluid is distributed upstream and downstream about 40 km in a supply line and returns in another line both of which share a common vacuum jacket and a thermal shield connected to the return line. These lines are connected to the magnet system transmission line along the string. The supply line also contains the return current transmission line from the refrigeration plant upstream. We will use these parameters for each cryogenic design using low-field magnets, independent of the temperature at which

the magnet operates, although some details are different for each design.

a. NbTi conductor, 4.5 K ≤ T ≤ 5 K

In this case, 300 g/s of helium flow at 3 bar is provided. Of this, 40 g/s is provided by the supply line to the end of the magnet string, and flows back 40 km to the refrigerator. The temperature of the helium rises to 5 K while flowing in the magnets and is recooled to 4.5 K every two cells in recoolers by exchanging heat with helium expanded through valves from the supply helium line. The saturated boil-off at 1 to 1.3 bar flows back to the refrigerator in the return line, keeping the thermal shield cool, and warming up to about 10 K by the time it reaches the refrigerator. Transient operation such as cool-down has not yet been studied, although a preliminary look suggests that 50 g/s would cool the string in about 10 days. For comparison, we note that the total ideal power of about 17 MW is twice that of the LHC, which uses four refrigerators operating at 1.9 K. The total helium inventory is about 3.5 M liters, four times that of the LHC. This handling of this large inventory has not been studied.

b. Nb_3Sn conductor, 4.5 K ≤ T ≤ 6.5 K

Although it might seem that this system should be almost exactly the same as the previous on, since the temperatures are very close, advantage is taken of the high heat capacity of helium near its critical point and the high critical temperature of Nb_3Sn. The helium flow is divided into four parallel paths in the magnet of about 40 g/s each at 4 bar, and all the heat is taken as sensible heat, that is, allowing the helium stream to increase in temperature from 4.5 K to 6.5 K with no recoolers. By this method, the ideal power is reduced by 30%, saving 20 MW of wall-plug power. This saving could be even greater if the Nb_3Sn transmission line could operate at higher temperature, which, because of the low magnetic field, might be possible. A disadvantage of this system is the five hours it takes for the helium to flow through the system.

c. High temperature superconductor, 20 K ≤ T ≤ 25 K

Because of the low magnetic field and low forces on the transmission line, it might be possible to use existing types of commercial HTS for this magnet, allowing operation at 20 K to 25 K. This system is similar to the previous one, taking all the heat as sensible heat, but at higher pressure, near 20 bar. At this pressure, the density of helium gas and the enthalpy to heat it to 25 K from 20 K is about the same as in the 4.5-6.5 K case, resulting in similar flow rates. There is a considerable saving in operating cost and refrigerator size due to the lower power required to cool to 20 K compared to 4.5 K. Operation with more advanced HTS at higher temperatures has not yet been modeled, but probably will be in the coming months, although the cost of the cryogenics at 25 K is already so low that there seems little need to search for higher temperature superconductor.

2. The high-field designs

Two of these designs, NbTi at 1.8 K and Nb_3Sn at 4.5 K, can be scaled from well-engineered cryogenic systems designs–the LHC and SSC, respectively. Hence, their capital and operating power and costs can be fairly accurately predicted. The third high-field design, using HTS at 20 K is completely speculative. In these designs, the field strength of the NbTi magnets is 9.5 T, resulting in a machine circumference of 138 km, and the field strength of the Nb_3Sn and HTS magnets is 12.6 T, resulting in a circumference of 104 km.

a. NbTi, T = 1.8 K

The parameters of this design are scaled from the LHC, with the assumptions that we have very similar magnets that are 10% larger in diameter, 20% heavier with 10% more current. The heat load per magnet is very similar to the LHC "Yellow Book" design [25], except for the intermediate 4.5 K-20 K level, which is roughly three times larger due to the increased synchrotron radiation impinging on the beam-tube liner. The number of dipoles in this RLHC is 6.3 times LHC, and it seems reasonable to have 20 stations each with one two 18 kw refrigerators (rated at 4.5 K), instead of LHC's four stations, each of which also has two 18 kw plants. The RLHC plants actually should be slightly larger than 18 kw to account for necessary redundancy and down-time. The nominal total operating power of 180 MW total is quite impressive, as is the estimated capital cost of between $1.5 billion and $2.0 billion.

b. Nb_3Sn, T = 4.5 K

This design is scaled from SSC because it operates at nearly the same temperature. Since the SSC design effort was terminated before it was complete, and because this magnet design is a complete mystery, the scaling not as accurate as the previous case. We have assumed that the 2-in-1 magnets will be twice the diameter and four times the weight of the LHC magnets. Because of the higher synchrotron radiation load, this ring is divided into short strings, resulting in 16 plants, compared with 10 for the SSC. The plants are about the same size as the 18 kW SSC plants. The total operating power in this case is about 75 MW, only 40% of the 1.8 K case. That old Second Law.

c. HTS, T = 25 K

These magnets are even more of a mystery than the Nb_3Sn, but for the sake of making a guess, we have

assumed that they are the same size and weight, which might be optimistic. At 20 K the synchrotron radiation can be taken directly on the beam-tube walls. In order to pump the desorbed hydrogen gas, a partial beam screen cooled by tubes at 4.5 K is installed in the bore tube, which turns the beam-tube liner solution of the other high-field cases inside out.

Table III: Comparison of Cryogenic Systems for Different RLHC Magnets

Collider Magnet Operating Temperature and Field	Ring Size (km)	Nº of Stations (inc 1 IR)	Total Heat Load at nominal temperature (kilowatts)					Ideal Power (MW)	Wall-Plug Operating Power
			1.8 K	4.5 K	20 K	50 K	Leads		(MW)
Low Field (all 1.8 T)									
NbTi, 4.5-5.0 K	646	9	0	247	0	0	200 g/s	17.2	66
Nb3Sn, 4.5-6.5 K	646	9	0	242	0	0	200 g/s	12.3	47
HTS, 20-25 K	646	9	0	0	242	0	200 g/s	3.7	14
High Field									
NbTi, 1.8 K, 9.5 T	138	20	115	413	0	1644	920 g/s	45	180
Nb3Sn, 4.5 K, 12.5 T	104	18	0	66	420	1080	940 g/s	18	72
HTS, 25 K, 12.5 T	104	18	0	15	590	1080	940 g/s	14	54

E. Power Supplies and Quench Protection

The approaches to power supply systems and quench protection will be quite different depending on whether the RLHC is a low-field or high-field design and whether the magnets are made from low-temperature (LTS) or the new, and as yet speculative, high-temperature superconductors (HTS).

In the low-field design, the inductance is very low, so the power supply and quench protection unit can be quite long. The power supply will be high current, about 75 kA to 100 kA, but can be low voltage, about 5 V. Instead of absorbing the stored energy in the magnets themselves, which is the usual method for high-field designs, the energy will be dissipated in an external resistor which is normally bypassed by a solid-state switch. When a quench is detected, the power supply is turned off and the switch is opened, forcing the current to decay through the series resistor. In order to limit the temperature rise during a quench, the superconductor is paralleled in the magnet by conventional copper or aluminum conductor. To limit the temperature rise to 500 K, 2 cm^2 of copper with a residual resistivity ratio (RRR) of 30 is required. In this case, a quench detection threshold of 1 V is sufficient.

Because HTS has much higher heat capacity than low-temperature superconductor, the quench propagation velocity will be very low, and the detection threshold will have to be very low, of the order of 10 mV. This will be a challenge in a real accelerator environment. Preliminary analysis indicates that 4 cm^2 of ultra-pure aluminum with RRR $\approx$ 1000 for the LTS case, or about 1.5 cm^2 in the HTS case, will suffice to provide cryo-stability of the conductor, at least for a limited time. This means that small quenches will disappear of their own accord. In the case of massive quenches, those caused by large area beam losses, for example, it will be adequate to detect the quench by an increase in the temperature or pressure of the cryogenic fluid, whence the power supply would simply be shut off. This is particularly interesting in the HTS case, where otherwise the quench detection threshold will be prohibitively low.

Power supply systems and quench protection in the high-field designs will be based on the same approach used in the SSC, HERA, RHIC and LHC. The magnet system is divided into long current loops powered by high-voltage (a few kilovolts), moderate-current (10 kA - 15 kA) power supplies. Each power loop consists of many short quench cells bypassed by diodes at cryogenic temperature. The quench is detected by voltage rise, and heaters are fired to spread the quench throughout the quench unit, which could be as small as one magnet. The energy stored in the quench unit is thus safely absorbed in the magnets themselves. The energy stored in the unquenched parts of the machine is bypassed around the quench cell by the diodes, and is absorbed in an external resistor. High quench propagation velocity in the LTS case permits high detection thresholds, about 1 V.

In the HTS case, the quench velocity may be very slow, and the high heat capacity (because of the high operating temperature), and large temperature margins will mean that heaters will not effectively spread the quench. Although the margins are very large, making these magnets intrinsically stable, if they do quench, protection will be difficult. This is a clear place for extensive R&D.

E. Detector optimization: interactions per crossing

In this section we will discuss the issue of the mean number of interactions per crossing. This number, $\langle n \rangle$, is given by

$$\langle n \rangle = L \sigma_{int} S_B. \quad (1)$$

in which L is the luminosity, S_B is the bunch spacing, and σ_{int} is the total inelastic cross section.

At Snowmass, a small group studied the impact of the number of interactions per crossing on the performance of a generic general-purpose RLHC detector. They noted that the time resolution of typical detector elements is in the range of 10-20 nsec; hence, bunch spacings smaller than this do not reduce the effective number of interactions per crossing. Consequently, there is a lower limit to the mean number of interactions per crossing which the detector will have to deal with, given from Eq. (1) with S_B in the 10-20 nsec range. For 100 TeV (CM) energy and a luminosity of 10^{34} $cm^{-2}sec^{-1}$, this lower limit is 12-25 interactions per crossing.

The group noted that at 100 TeV (CM) energy, the average charged multiplicity is about 300; so the mean number of charged particles per crossing is 3500-7000. The detector performance was found to degrade quickly as the number of interactions per crossing increased. The central tracker occupancy rises; there is pile-up in the calorimeter with consequent reduction in its energy resolution; and the muon system suffers from combinatoric backgrounds. In general, it becomes much easier for two "soft" events to mimic a rare high-energy process.

In view of the severe consequences for the detector, it is perhaps most prudent for the accelerator design to provide for a bunch spacing in the 10-20 nsec range, which keeps the number of interactions per crossing as low as possible (see Table I above). This is unfortunate, since many of the technical problems with the accelerator are related to the beam stored energy and the synchrotron radiation power. These quantities are proportional to the total number of particles in the machine, N_T. The luminosity

$$L \propto N_T^2 S_B \quad (2)$$

can remain constant as the number of particles is reduced only if the bunch spacing increases, which of course increases $\langle n \rangle$. A crucial machine/detector interface issue, one deserving of substantial additional R&D effort, is the proper tradeoff between detector problems related to large numbers of interactions per crossing and the accelerator systems issues required to cope with large beam stored energy and large amounts of synchrotron radiation power.

III. CONCLUSIONS

A. The major issues

Colliders of such immense size and complexity share a number of possible problems independent of the details of their designs and magnetic field strength. First among these is reliability, because of the large number of components in these large rings, and their relative inaccessibility. The simple magnets of the low field ring mitigate this issue somewhat for the magnet, cryogenic and power supply systems, but exacerbate the problems for tunneling, controls, alignment and some other length-dependent systems. In any case, reliability engineering at a level presently unknown in high-energy accelerators will be necessary. Another common issue is the large number of interactions per bunch crossing and the large charged particle multiplicity at 100 TeV (CM). Here there is a trade-off between the number of bunches, the total stored beam energy, and the number of interactions per crossing. This issue is more difficult to solve in the low-field design. The R&D on detectors being done now for LHC will certainly help find solutions to this problem. More will need to be done.

Each machine design has its own specific problems. In the high-field ring, the major issue is inventing a magnet that can be mass produced for a reasonable cost, and handling the very large synchrotron power. In the moderate field ring, the challenge will be to value engineer a low-cost magnet, and to design and operate reliably a very large superfluid helium system. In both the high and moderate-field rings the accelerator physics issues have been well studied and appear straightforward. The existence of synchrotron damping in these two designs forgives a host of sins, making the challenge of emittance preservation far less daunting, and significantly reducing the amount of stored beam.

The issues in the low-field ring are less concerned with the magnet, and more with accelerator physics issues, such as beam stability and feedback, and methods to decrease the cost of tunneling and length dependent systems, both driven by the large circumference of the machine. In addition, the lack of synchrotron damping in the low-field design drastically increases the amount of stored beam in the collider, and makes the consequences of emittance blow-up during injection and acceleration more severe.

In all three of these designs there are complicated trade-offs between the magnet aperture and field quality, and among the collider cell length, correction system complexity and injection energy. Choosing a 3 TeV injector has eliminated or at least moderated a number of injection field quality and beam stability issues.

B. A list of possible R&D activities

The discussion of major issues leads to the list of R&D activities in Table IV that we might start to work on, and that could help us decide which sort of design would result in the best and most cost effective collider. We believe that many, if not most, of these R&D activities are of general intellectual and practical interest right now, and should be actively pursued even if political and fiscal reality puts the RLHC in the distant future. In the list below, the suggested R&D will affect all the designs, but often affects one more than the others.

Table IV: R&D Items

R&D ITEM	DESIGN MOST AFFECTED
1. Accelerator physics paper studies	
a. Single and multibunch stability	low field
b. Aperture/cell length trade-off	all
c. Field quality requirements	all
d. Vibration & ripple sensitivity	low field
2. Tunneling development	low field
3. Magnet R&D	
a. Value engineering of existing NbTi magnet designs	mid field
b. Inventing and perfecting new design	high field
c. Transmission line/C-magnet R&D	low field
4. Superconducting materials	
a. High-field NbTi	mid field
b. APC techniques for NbTi	mid field
c. Transmission lines with HTS	low field
d. Nb_3Sn & thin film HTS	high field
5. Systems development	
a. Superfluid refrigeration (paper study)	mid field
b. HTS quench development	high field
c. Vibration feedback	all
d. beam stability feedback	low field
e. H_2 absorption at 20 K	high & mid
f. Beam loss studies	all
g. Beam absorber studies	low field

IV. REFERENCES

[1] Proceedings of the Workshop on Future Hadron Facilities in the US, (Bloomington, 1994), Fermilab TM-1907

[2] G. W. Foster, E. Malamud, "Low-cost Hadron Colliders at Fermilab", Fermilab TM-1976 (1996)

[3] Stephane Keller, "Interim Report from the Pipetron Physics/Detector Team," Snowmass 96.

[4] P. J. Limon, "The Optimum Energy of the Really Large Hadron Collider Sited at Fermilab," contribution to the Snowmass 96 Proceedings.

[5] J. T. Rogers, "Collective Effects and Impedences in the RLHC(s)," contribution to the Snowmass 96 Proceedings.

[6] R. R. Wilson, "Superferric Magnets for 20 TeV," 1982 Snowmass Proceedings, p. 330 ff

[7] Donaldson, et al, (ed), "Reference Design C," Report of the Reference Designs Study Group on the SSC, Berkeley, 1984

[8] "New Low-Cost Approaches to High Energy Hadron Colliders at Fermilab", mini-symposium at the spring APS meeting, Indianapolis (1996)

[9] M. McAshan and P. Mazur, "Cryogenic System for the Low-Field RLHC Study Case," contribution to the Snowmass 96 Proceedings.

[10] J. Friant, et al, "Pipetron Tunnel Construction Issues," contribution to the Snowmass 96 Proceedings.

[11] W. Turner, "Beam Tube Vacuum in Very Large Hadron Colliders", contribution to the Snowmass 96 Proceedings.

[12] H. Ishimaru, "Conceptual Design Study for a Very Low-Cost Aluminum Alloy Vacuum Chamber in a High-Energy Low-Field Collider", contribution to the Snowmass 96 Proceedings.

[13] N. Mokhov, T. Murphy and S. Pruss, "Extracted Beamlines and Absorber for the RLHC," contribution to the Snowmass 96 Proceedings.

[14] M. McAshan and P. Mazur, "Cryogenic System for the High-Field RLHC Study Case," contribution to the Snowmass 96 Proceedings.

[15] A. Den Ouden and H.H.J. Ten Kate, "Application of Nb_3Sn Superconductors in High-Field Accelerator Magnets," Invited talk at the Applied Superconductivity Conference, Pittsburgh, August, 1996

[16] D. Wu et. al., Appl. Phys. Lett. **67**, 2397 (1995)

[17] M. Syphers, M. Harrison, S. Peggs, "Beyond the LHC: A Conceptual Approach to a Future High Energy Hadron Collider", in Proceedings of the 1995 Particle Accelerator Conference and International Conference on High Energy Accelerators (Dallas, 1995) p. 431

[18] S. Peggs and F. Dell, "Tolerable Systematic Errors in Really Large Hadron Collider Dipoles," contribution to the Snowmass 96 Proceedings.

[19] S. Peggs, et al., "Lattice Optimization for a Really Large Hadron Collider (RLHC)", contribution to the EPAC 96 Proceedings.

[20] G. W. Foster, "Double-C Transmission Line Magnet," contribution to the Snowmass 96 Proceedings.

[21] R. Gupta, et al, "Shimming Magnets for Good Field Quality", contribution to the Snowmass 96 Proceedings.

[22] J. Wei, contribution to the Snowmass 96 Proceedings.

[23] A. Koski, S. L. Wipf, "Computational Design Study for an Accelerator Dipole in the Range of 15-20 T", DESY report

[24] R. Bauer and D. Gross, "Geology of the Greater Fermilab Region," contribution to the Snowmass 96 Proceedings.

[25]The Large Hadron Collider Conceptual Design, (The Yellow Book), October, 1995, CERN/AC/95-05 (LHC)