Accelerator Physics issues in eRHIC

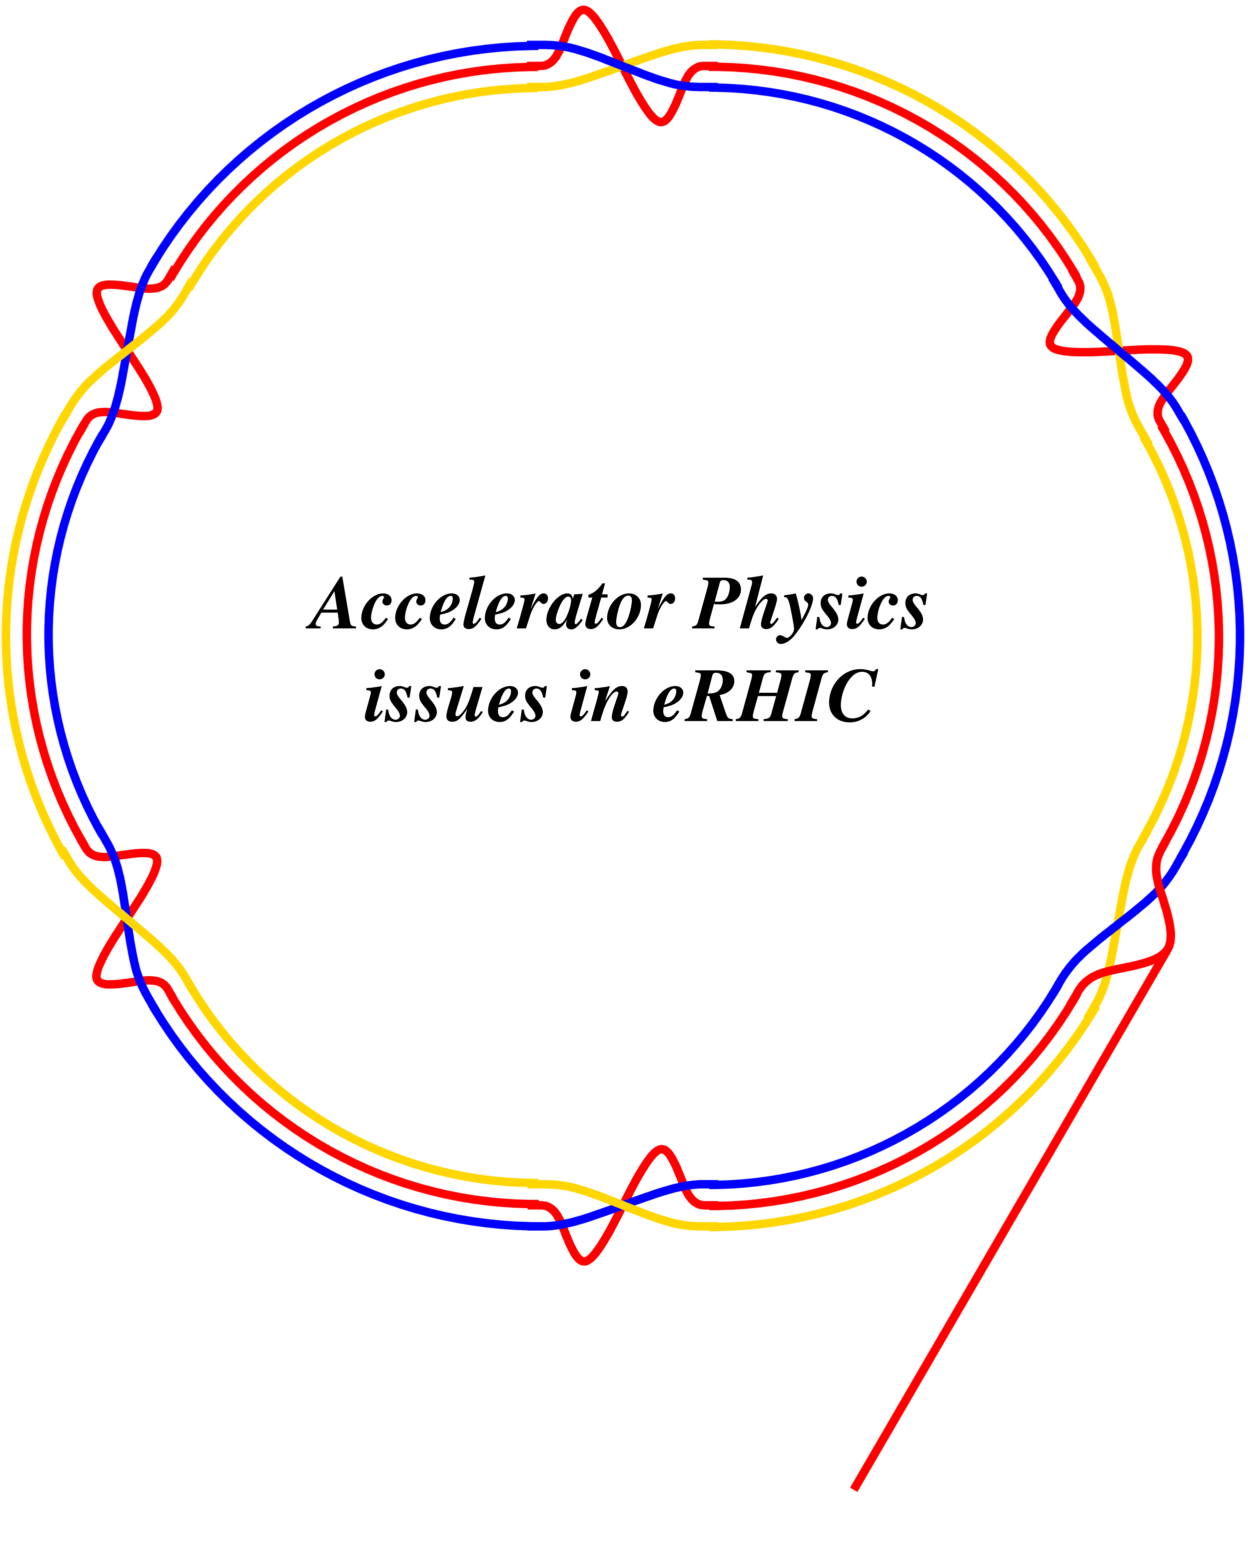

C-A/AP/14
April 3, 2000

Accelerator Physics issues in eRHIC

I. Ben-Zvi, J. Kewisch, J. Murphy, S. Peggs

Abstract

The accelerator physics issues associated with adding electron-gold and polarized electron-proton collisions to the RHIC repertoire are discussed. Two scenarios are considered: *ring-ring* and *linac-ring* collisions. In the first scenario a third electron ring, possibly made from permanent magnets, is added to the RHIC tunnel. In the second scenario the considerable beam power is recuperated by recycling the beam through the full energy linac injector.

Straw man ring optics are presented, including an integrated spin rotator which provides longitudinally polarized electrons at the interaction points. Electron and ion intensity limitations are investigated, for operations with as many as 2520 bunches. Synchrotron radiation, beam-beam, electron cloud, and cryogenic heat load limits are examined. The implications for new and upgraded ion ring systems are discussed, including electron cooling, RF systems, the ion abort dumps, and injection and extraction kickers.

Contents

1 Introduction **4**
- 1.1 Two scenarios: *ring-ring* and *linac-ring* 4
- 1.2 Common parameters and definitions 8

2 Ring-Ring Collisions **9**
- 2.1 Luminosity performance at the beam-beam limit 9
  - 2.1.1 Beam sizes 9
  - 2.1.2 Number of bunches 10
  - 2.1.3 Single bunch intensity 10
  - 2.1.4 Polarization 11
  - 2.1.5 A polarizing booster 12
- 2.2 Electron ring optics 12
  - 2.2.1 An interaction region without spin rotators 15
  - 2.2.2 Spin rotators 15
  - 2.2.3 The utility interaction region, IP4 18
  - 2.2.4 The "high noon" interaction region, IP12 18
- 2.3 Electron beam intensity limitations 21
  - 2.3.1 Synchrotron radiation from the arc dipole magnets 21
  - 2.3.2 Synchrotron radiation from the spin rotator magnets 22
  - 2.3.3 Longitudinal single bunch current limits 22
  - 2.3.4 Multi-bunch current limits 23
  - 2.3.5 Ion trapping 23
- 2.4 Ring energy tunability 23
  - 2.4.1 Energies above the nominal 10 GeV 23
  - 2.4.2 Energies below the nominal 10 GeV 24
  - 2.4.3 Electron beam lifetime 25
- 2.5 Ion beam intensity limitations 25
  - 2.5.1 Radio Frequency system 25
  - 2.5.2 Kickers 26
  - 2.5.3 Abort dump 26
  - 2.5.4 Long range beam-beam interactions 26
  - 2.5.5 Cryogenic beam pipe and BPM signal cable heating 27
  - 2.5.6 Electron cloud effect 28

3 Linac-Ring Collisions **29**
- 3.1 A 10 GeV linac-ring collider 29
- 3.2 Design issues of the high current, energy recuperating linac 30
  - 3.2.1 The linac layout and optics 30
  - 3.2.2 Beam dump 30
  - 3.2.3 The polarized electron source 30
  - 3.2.4 Recuperation issues and beam stability 31
- 3.3 Electron cooling of RHIC 31
- 3.4 Parameters of the linac 32
- 3.5 Interaction region optics 33

4 Summary **34**

1 Introduction

Interest continues to grow among theorists and experimentalists in the physics of collisions between electrons and heavy ions, and between polarized electrons and polarized protons [1, 2, 3].

Physicists at DESY have begun to explore the possibility of upgrading the existing HERA-p ring to store heavy ions, in order to collide them with electrons (or positrons) in the HERA-e ring. Another possible upgrade, to store polarized protons in the HERA-p ring, is also under discussion [1]. This requires polarizing the protons at the source, and then maintaining polarization through the injection chain. Taken together, these two initiatives would upgrade HERA to match RHICs ability to store and collide heavy ions and polarized protons.

At BNL, it is natural to consider adding polarized electrons to the RHIC repertoire. The authors of this paper form an accelerator physics study group which was charged with the mission of answering the following questions [4]:

1. *A 10 GeV electron beam injected into a room temperature ring inside the RHIC tunnel collides with one of the 100 GeV/u gold beams. What luminosity can be achieved with realistic assumptions about the electron beam intensity and phase space? How would the interaction region look? Assume that the electron beam originates in a superconducting linac operating at 1.3 – 1.5 GHz.*
2. *The same electron beam collides with a 250 GeV proton beam in one of the RHIC rings. Discuss the same parameters as above.*
3. *What is the highest possible energy to which electrons could be accelerated in a ring inside the RHIC tunnel, injecting at about 10 GeV? What is the luminosity now?*
4. *Consider what improvement in luminosity could be achieved by using electron cooling or stochastic cooling of the ion or proton beams in either the AGS or the RHIC ring, for gold-gold collisions in the two RHIC rings, or for the electron-gold and electron-proton collisions outlined above.*

This paper attempts to answer these questions. It estimates eRHIC accelerator parameters, and describes some of the physical constraints which would limit its performance.

1.1 Two scenarios: *ring-ring* and *linac-ring*

RHIC provides for the acceleration and storage of bunches of heavy or light ions – typically gold or protons – in two horizontally separated superconducting rings [5]. Head-on collisions occur at six interaction points (IPs). Collisions between like species are usually envisaged: gold-gold or polarized proton-proton collisions. RHIC also enables proton-gold collisions – it is possible to store and collide two beams with very different rigidities. Nonetheless, it is uninteresting to store electrons in one of the two superconducting rings, because of the unacceptable cryogenic heat load which is incurred when even a small electron current is present [6, 7]. If electrons are to be stored in the RHIC tunnel, then a third ring must be constructed.

There is interest among the synchrotron light community in the construction of a superconducting linac in the energy range from 6 GeV to 10 GeV, in order to drive a free electron laser. Except when RHIC is in the (relatively rare) process of cooling down from ambient temperatures, the cryogenic refrigeration plant already installed has enough spare capacity to run such a linac. It is therefore natural to place the linac relatively near to RHIC. The synergy is completed by making spare pulses from the linac available for eRHIC collisions.

In the "*ring-ring collisions*" scenario, 10 GeV electrons circulate in a third storage ring in the RHIC tunnel. This is illustrated in cartoon form in Figure 1. "Pre-polarized" electrons are injected from a full energy linac in the counterclockwise direction at IP4. Collisions are possible with the clockwise rotating ions in the Blue ring at the other five interaction points. The natural place to install the electron ring magnets is just above the floor of the tunnel, in between the stands of the Blue and Yellow ion ring magnets, as shown in Figure 2.

In the "*linac-ring collisions*" scenario there may or may not be a third ring in the RHIC tunnel. The considerable energy which is stored in the electron beam is not thrown away after each single pass collision, but is recovered by recirculating the electron beam back through the linac. The accelerating and decelerating electron beams travel in the same direction along the linac in order to avoid collisions. An arrangement with two linacs in a single straight line is illustrated schematically in Figure 3. This configuration has the problem that the returning 10 GeV electrons must be bent by 180 degrees around a small radius arc. Figure 4 shows an alternate layout with side-by-side linac stages, which uses the large radius RHIC tunnel to return the 10 GeV electrons.

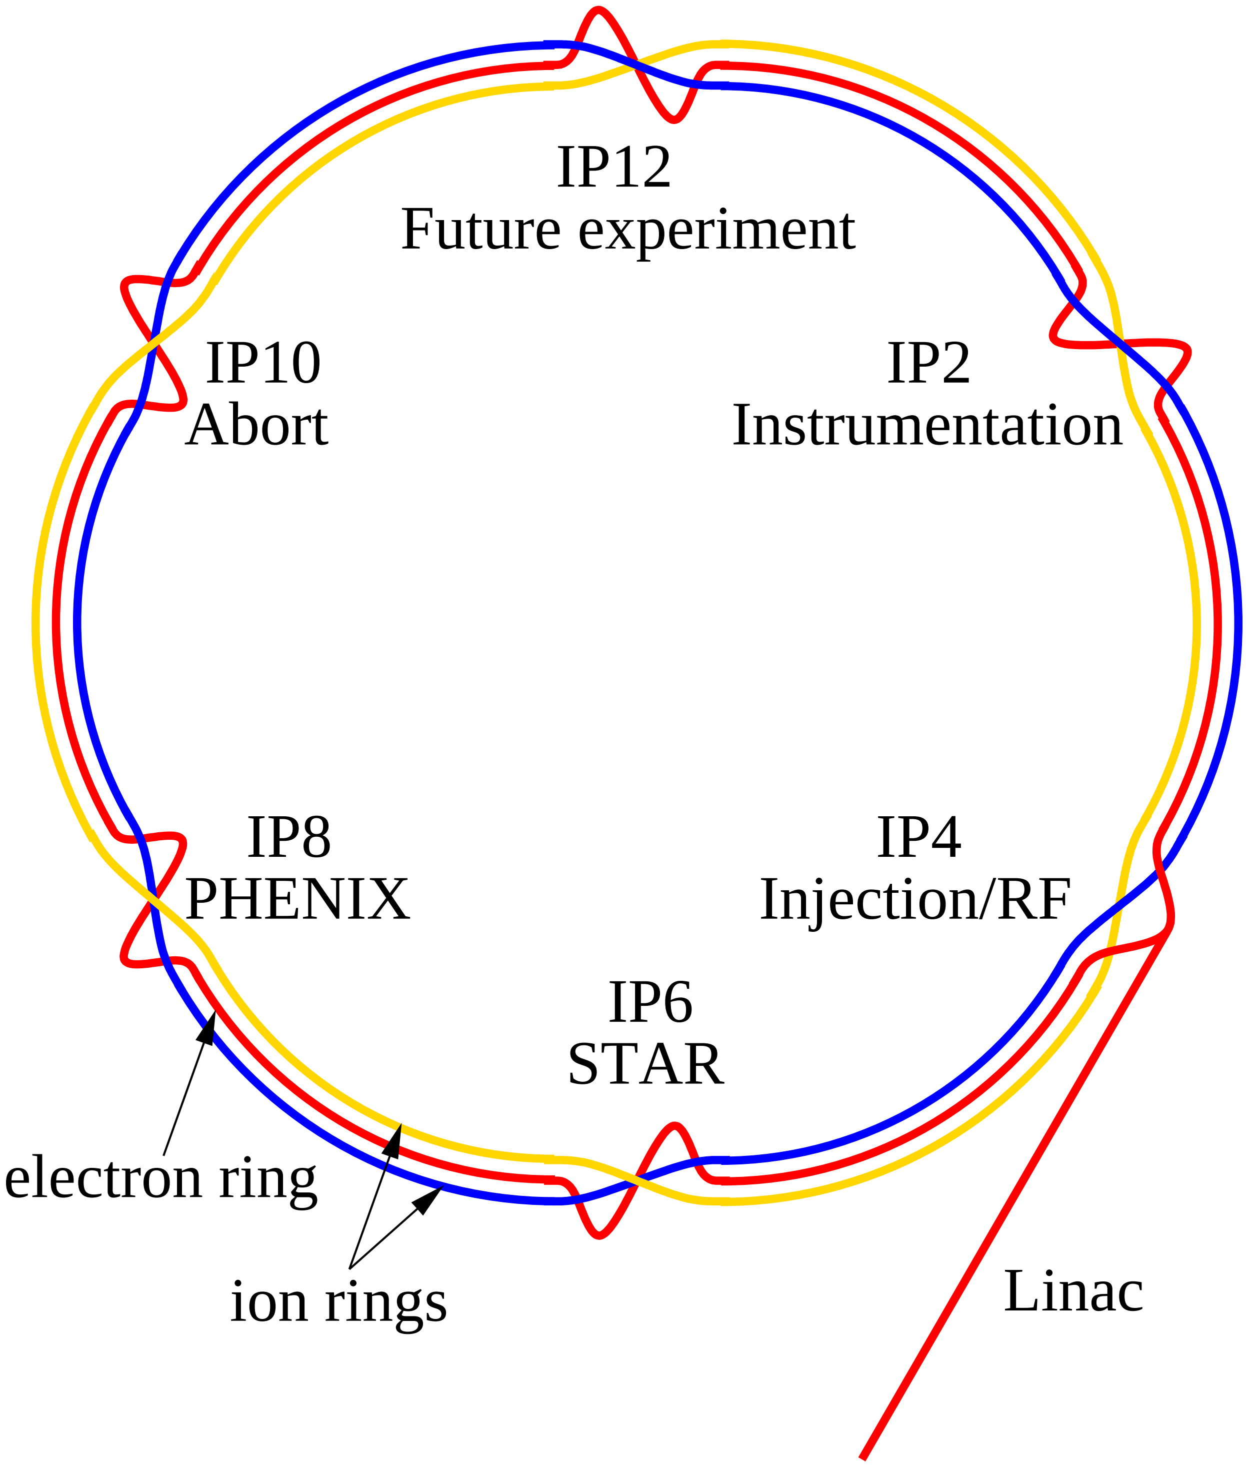

Figure 1: Schematic layout of the injection linac, electron ring, and the two ion rings, in the *ring-ring* collision scenario.

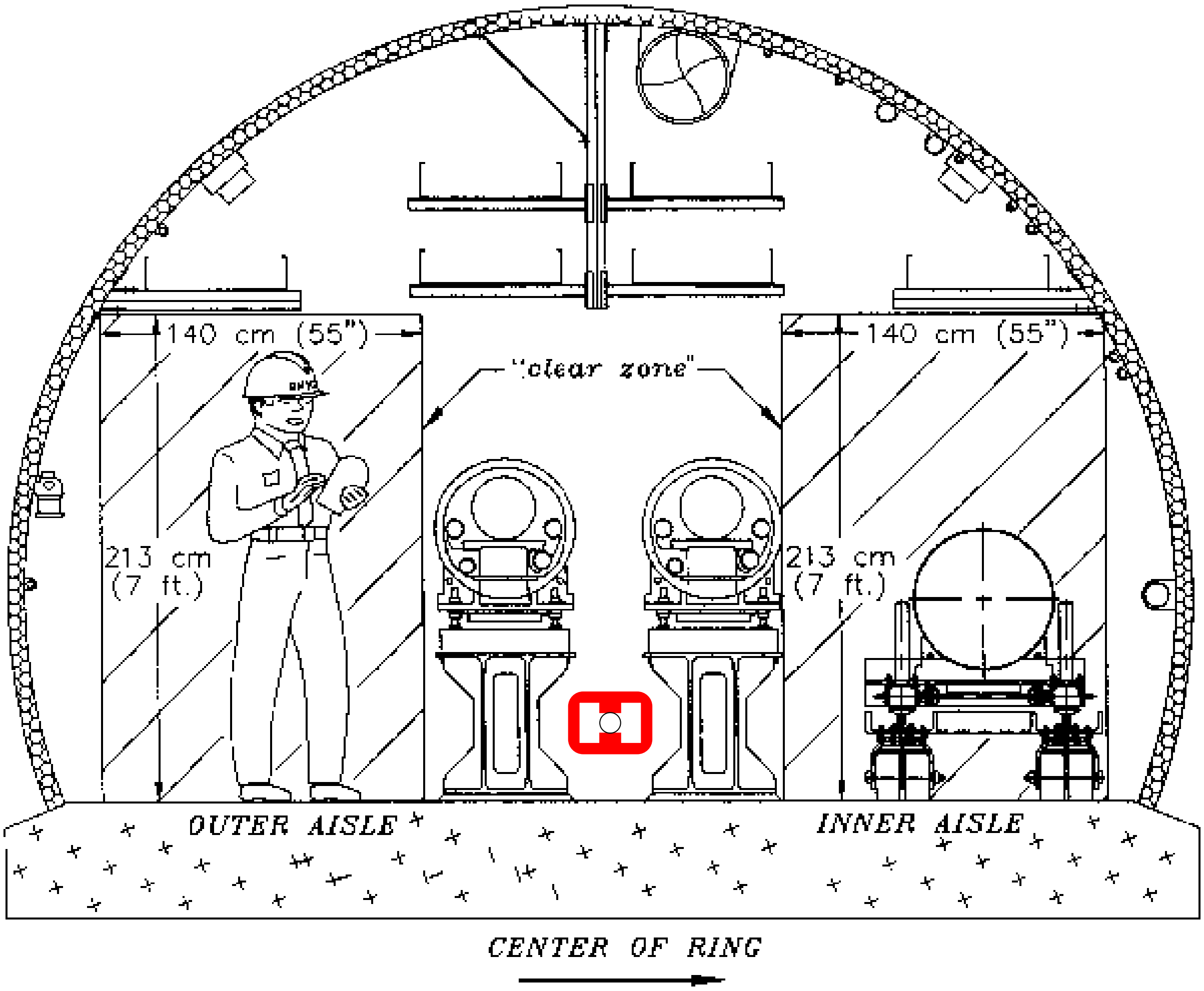

Figure 2: RHIC tunnel cross-section, showing an electron ring magnet between the ion ring stands.

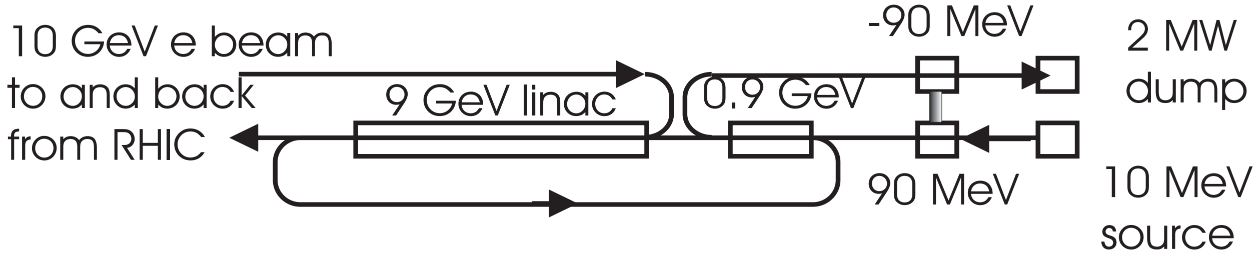

Figure 3: Schematic layout of a two stage recirculating linac structure in the *linac-ring* collision scenario.

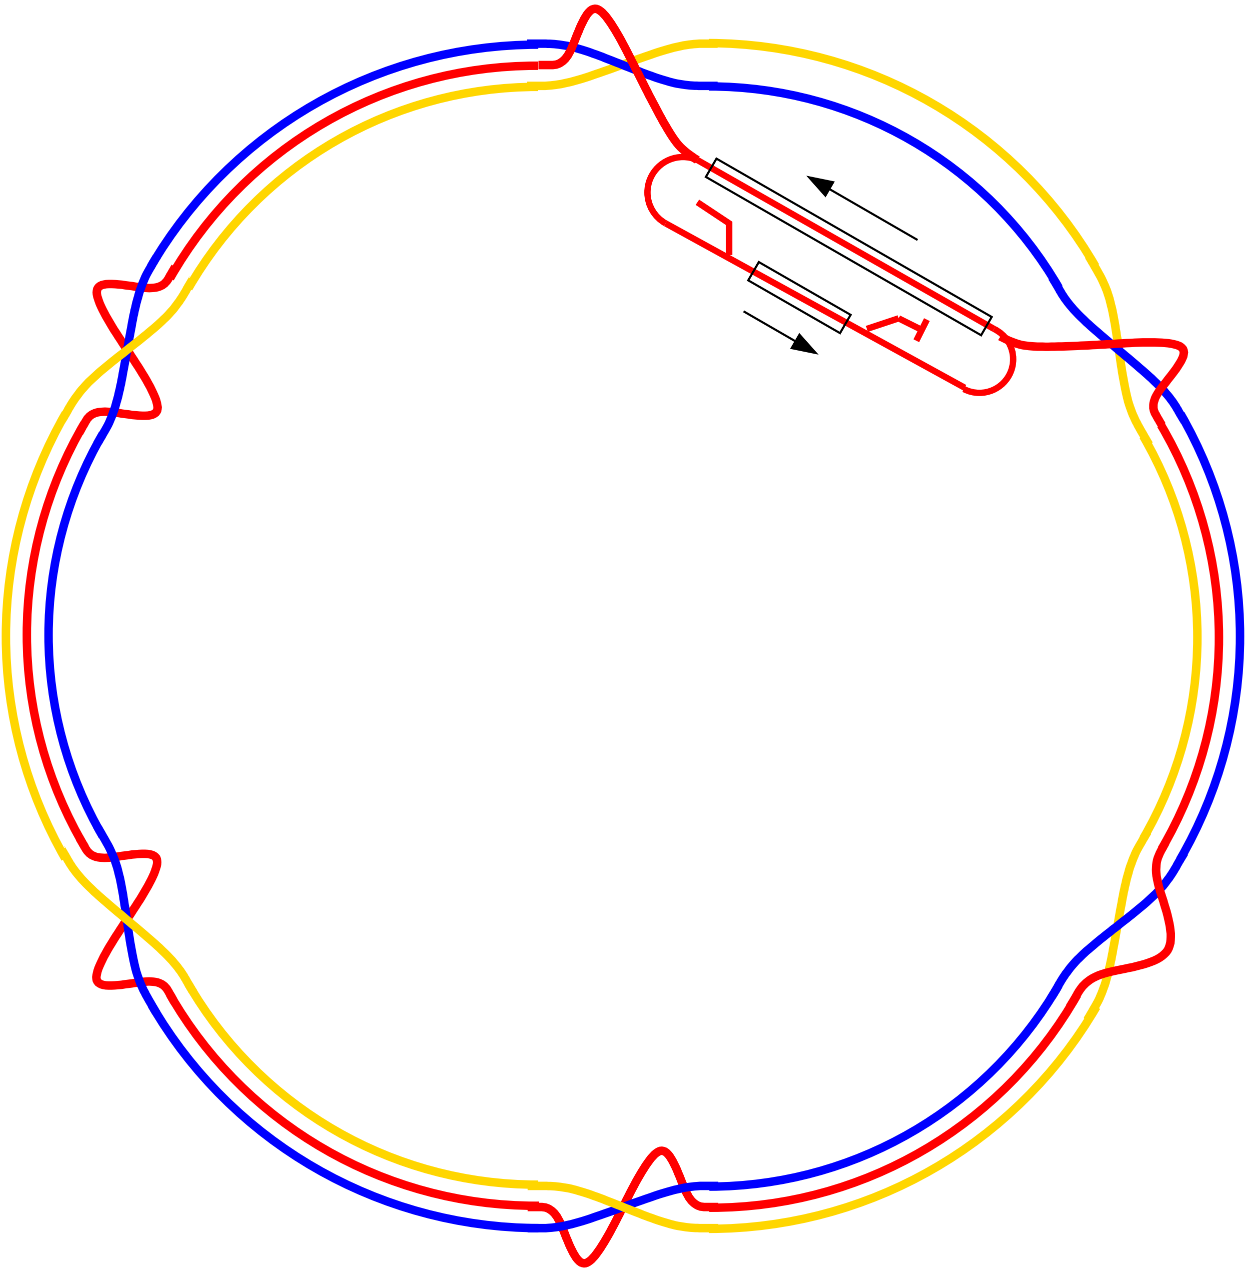

Figure 4: Schematic layout for *linac-ring* collisions, using the RHIC tunnel to return the 10 GeV electrons for energy recovery.

1.2 Common parameters and definitions

The primary parameters listed in Table 1 are common to both eRHIC scenarios. The RMS ion – gold or proton – beam size at the IP is written as

$$\sigma_i^* = \sqrt{\frac{\beta_i^* \epsilon_i}{6\pi(\beta\gamma)_i}} \tag{1}$$

where ϵ_i is the *95% normalized* RMS ion emittance. In contrast, the RMS electron beam size at the IP is written

$$\sigma_e^* = \sqrt{\beta_e^* \epsilon_e} \tag{2}$$

where the electron emittance ϵ_e is *unnormalized.*

Parameter	units	value
Gold top energy per nucleus, E_{Au}	[GeV/u]	100
Proton top energy, E_p	[GeV]	250
Electron energy, E_e	[GeV]	10
Circumference, C	[m]	3833
Revolution frequency, F_{rev}	[kHz]	78.3
Arc dipole bend radius, ρ	[m]	243

Table 1: Primary eRHIC parameters

When the beams are round and identical in size, the electron and ion beam-beam tune shift parameters ξ_e and ξ_i are given by

$$\xi_e = N_i Z \frac{r_e \beta_e^*}{4\pi\gamma_e \sigma^{*2}} \tag{3}$$

$$\xi_i = \frac{N_e}{Z} \frac{r_i \beta_i^*}{4\pi\gamma_i \sigma^{*2}} \tag{4}$$

where N_i and N_e are the number of ions and electrons per bunch, respectively. The classical radius of the ion

$$r_i = \frac{Z^2}{A} \frac{e^2}{4\pi\epsilon_0} \frac{1}{M_0 c^2} \tag{5}$$

depends on the atomic number Z (79 for gold), the atomic mass number A (197 for gold), and the mass per nucleon, M_0. Its value is $r_p = 1.53 \times 10^{-18}$ m for protons, $r_{Au} = 49.0 \times 10^{-18}$ m for gold, and $r_e = 2.82 \times 10^{-15}$ m for electrons.

2 Ring-Ring Collisions

The discussion of ring-ring collisions in this section is organized as follows. First, **Luminosity performance at the beam-beam limit** is discussed, in two sub-scenarios with 360 bunches and with 2520 bunches per ring. The philosophy is that the beam-beam limit sets the ultimate performance scale against which challenges like engineering difficulties, and implementation costs, may be measured.

The sub-section **Electron ring optics** describes a straw man solution for the implementation of the electron ring, going well beyond the cartoon layout of Figure 1. A major challenge is in the provision of spin rotators to arrange for longitudinal polarization of electrons at each IP. **Electron beam intensity limitations** and **Ion beam intensity limitations** discuss a variety of practical difficulties which occur as the single bunch populations and the number of bunches – the peak and average beam currents – are increased. These include the synchrotron radiation heat load in the electron ring, and the electron cloud effect and cryogenic beam pipe heating in the ion rings. Finally, **Ring energy tunability** looks at the costs and advantages in raising or lowering the electron beam energy around its 10 GeV nominal value.

2.1 Luminosity performance at the beam-beam limit

The nominal luminosity is given by the equation

$$L = F_{rev} \frac{N_b N_e N_i}{4\pi\sigma^{*2}} \tag{6}$$

and depends directly on three quantities:

- σ^*, the round beam collision size
- N_b, the number of bunches in each beam, and
- N_e and N_i, the single bunch populations for electrons and ions.

Plausible extreme values for these quantities, and for the maximum achievable luminosity, are derived below.

If the electron and ion single bunch populations N_e and N_i are both limited (for example, by the head-on beam-beam interaction), but the total electron and ion currents are not, then $L \sim N_b(N_e N_i)$, and the luminosity is increased by increasing the number of bunches.

If either the electron or the ion total current, I_e or I_i, is limited, then $L \sim I_i N_e$ or $L \sim N_i I_e$, and the luminosity is at a maximum. The number of bunches is then set by other considerations – most likely to be as small as possible.

A crucial question is then: *Does the electron or the ion ring reach its total current limit first, as the number of bunches is increased?* If the total current limit is not a hard physical limit, then a more realistic question becomes: *Where is the optimum balance between the marginal cost (and technical difficulty) of increasing the number of bunches, and the marginal benefit of increasing the luminosity performance?*

2.1.1 Beam sizes

In the absence of a final and rigorous optimization of the electron ring optics, the following four assumptions are made for the sake of simplicity:

1. **Both electron and ion beams are round at the IP.** The horizontal and vertical emittances, beta functions, and beam sizes are assumed to be identical for the electrons and, separately, for the ions. This approximation is quite reasonable for the ion beam, but it is crude for the electron beam. For comparison, in the HERA-e ring the horizontal and vertical emittances differ by a factor of about 4, and the horizontal and vertical beam sizes differ by a factor of about 2 [8].
2. **The normalized ion emittance is $\epsilon_i = 18\ \pi\mu$m.** This is consistent with the nominal parameters used for RHIC, in which the gold emittance grows from 10 $\pi\mu$m to about 40 $\pi\mu$m through Intra-Beam Scattering (IBS) in the course of a 10 hour store [5]. It is also reasonably consistent with expected proton beam emittances.
3. **The unnormalized electron emittance is $\epsilon_e = 40$ nm.** This is roughly consistent with the equilibrium emittance due to synchrotron radiation in the straw man layout of the eRHIC electron optics presented below.

4. **The electron and ion beams sizes are identical, at a constant value of $\sigma_i^* = \sigma_e^* = 0.1$ mm.** This value is consistent with the nominal value of $\sigma_i^* = 0.137$ mm expected for gold-gold collisions at the end of a 10 hour store [5].

Table 2 lists the three numerical values which were introduced (rather arbitrarily) in these assumptions. The table also lists the IP beta functions for electrons, gold, and ions which immediately follow from the the application of Equations 1 and 2. These beta functions are consistent, within a factor of 2, with values for ions and electrons in RHIC ($\beta^* \approx 1.0$ meters at collision) and in HERA [5, 8].

Parameter	units	value
Electron emittance, ϵ_e	[nm]	40.0
Normalized ion emittance, ϵ_i	[$\pi\mu$m]	18.0
RMS beam size at the IP, σ^*	[mm]	0.1
Electron IP beta function, β_e^*	[m]	0.25
Gold IP beta function, β_{Au}^*	[m]	0.36
Proton IP beta function, β_p^*	[m]	0.88

Table 2: Beam size related parameters

2.1.2 Number of bunches

If the total beam current is not limited, the luminosity is directly proportional to N_b, the number of bunches in each beam. What is the maximum number bunches that can be stored?

The RHIC Design Manual [5] quotes a nominal number of 60 ion bunches, but a trivial upgrade to 120 bunches is envisaged, perhaps in the year 2000 production run. A further increase to 360 bunches is possible in a medium term luminosity upgrade, in which every RF bucket in the 28 MHz acceleration system is filled. Although such an upgrade has a non-trivial impact on some of the existing accelerator and experimental hardware, the technical demands of operating with 360 bunches are reasonable.

An electron ring in the RHIC tunnel has an energy and a circumference similar to one of the two rings in the SLAC or KEK B-factory – it has been tellingly observed that *"eRHIC is half a B-factory"* [9]. However, even with 360 bunches, the RHIC bunch spacing of 35.5 ns is modest by comparison with existing B-factories, and by comparison with the projected parameters of the LHC. The SLAC HER has as many as 1658 bunches, with a bunch spacing as small as 4.2 ns. It is therefore natural to consider a much larger number of bunches for eRHIC operations.

After acceleration to storage energy, RHIC ion bunches are "rebucketed" by turning off the 28 MHz acceleration system and turning on the 197 MHz storage system. It is natural to consider what happens when every bucket of the storage system, with a harmonic number of $h = 7 \times 360 = 2520$, is filled. Increasing the number of bunches to 2520 is technically demanding to both electron and ion rings.

Both ends of the spectrum – with $N_b = 360$ and with $N_b = 2520$ bunches – are discussed below, generating potential performance figures (such as luminosity) which range from "reasonable" to "demanding".

2.1.3 Single bunch intensity

Equations 3 and 4 place fundamental limits on the single bunch populations, since ξ_e and ξ_i have maximum critical values which cannot be surpassed. These critical values depend on details such as the number of head-on collisions per turn, on the presence of long range collisions, on the betatron tunes, on the synchrotron radiation damping time (for the electrons), et cetera. Nonetheless, in order to proceed without a full blown investigation of the beam-beam interaction in eRHIC, it is reasonable to assume that

$$\xi_e \leq 0.06 \quad (7)$$
$$\xi_i \leq 0.004 \quad (8)$$

These values are justified by experience at CESR, HERA, SLAC, SPS, and the Tevatron. Note that if the electrons do not circulate, but are "thrown away" in an energy efficient fashion, then the limit on ξ_e can be violated. This is discussed, below, in the section on *Linac-Ring Collisions*.

Table 3 shows the ultimate luminosity performance which is attained at the beam-beam limit. It also lists beam-beam limited values for the circulating electron and ion currents I_e and I_i, for the total synchrotron radiation power P_{rad}, and for the linear power loading per meter of dipole P_{lin}. It is not necessarily possible to reach beam-beam limited performance – for example, the electron current of $I_e = 7.7$ A quoted for electron-gold collisions with $N_b = 2520$ bunches may lead to unacceptable synchrotron radiation linear power loads. Electron and ion intensity limitations are discussed in some detail below.

The beam-beam performance limit sets a scale, against which the key question for the ring-ring scenario can be asked: *Realistically, how close can the $N_b = 2520$ beam-beam performance limit be approached?*

Number of bunches, N_b		360	2520
ELECTRON-GOLD			
Luminosity, L	[$\mathrm{cm}^{-2}\mathrm{s}^{-1}$]	6.4×10^{30}	$45. \times 10^{30}$
Electron bunch population, N_e		2.43×10^{11}	2.43×10^{11}
Gold bunch population, N_i		1.19×10^{9}	1.19×10^{9}
Electron beam current, I_e	[A]	1.1	7.7
Electron stored energy, U_e	[kJ]	140	981
Total radiated power, P_{rad}	[MW]	4.0	27.9
Linear power load, P_{lin}	[kW/m]	2.6	18.3
Gold beam current, I_i	[A]	0.42	2.97
Gold stored energy, U_i	[kJ]	1352	9464
ELECTRON-PROTON			
Luminosity, L	[$\mathrm{cm}^{-2}\mathrm{s}^{-1}$]	2.1×10^{32}	$15. \times 10^{32}$
Electron bunch population, N_e		1.00×10^{11}	1.00×10^{11}
Proton bunch population, N_i		$.93 \times 10^{11}$	$.93 \times 10^{11}$
Electron beam current, I_e	[A]	0.45	3.16
Electron stored energy, U_e	[kJ]	58	404
Total radiated power, P_{rad}	[MW]	1.6	11.5
Linear power load, P_{lin}	[kW/m]	1.1	7.5
Proton beam current, I_i	[A]	0.42	2.92
Proton stored energy, U_i	[kJ]	1335	9346

Table 3: Electron-gold and polarized electron-proton performance parameters, at the beam-beam performance limit, with 360 and 2520 bunches in each beam.

2.1.4 Polarization

The natural polarization time due to the Sokolov-Ternov effect, in an isomagnetic ring with $C = 3,833$ m, $\rho = 243$ m, and $E_e = 10$ GeV, is

$$T_{pol} = \frac{99\, C\rho^2}{2\pi E_e^5} = 3.6 \times 10^4 \text{ [s]} \quad (9)$$

or 9.9 hours! At HERA it is possible to accelerate the electron beam to storage energy in the HERA-e ring, and then to wait for the electrons to self-polarize, in a time which is much smaller than typical storage times. This is clearly not possible in the eRHIC ring. The electrons must be injected pre-polarized.

Intrinsic spin resonances are located at energies given by

$$E_{resonance} \;=\; J\,0.44065 \text{ [GeV]} \tag{10}$$

where J is an integer. It will (probably) not be possible to store beam at energies close to one of these resonances. Nor will it be possible to accelerate or decelerate through one of the resonances, without losing electron polarization. If electron polarization is to be preserved, then the nominal range of energy adjustment *after injection* is something like

$$E_e \;\approx\; 9.92 \pm 0.20 \text{ [GeV]} \tag{11}$$

The eRHIC ring does not accelerate or decelerate polarized electrons more than a fraction of a GeV, and the magnetic field of $B \approx 0.14$ is modest. Thus, it is natural to consider using permanent magnet technology for some or most of the ring lattice magnets. However, this engineering design choice is not critical in any fundamental way.

2.1.5 A polarizing booster

The full energy linac is equipped with a polarized source, in order to be capable of providing electrons polarized in the right orientation at the injection point into the eRHIC ring. Alternatively, a more conventional injection chain using a booster ring could be used. Equation 9 shows that a booster with 1 T dipoles ($\rho = 33.4$ m) and a packing fraction of 0.5 ($C = 420$ m) has a polarization time of only $T_{pol} = 74$ s. Such a booster would accelerate electron bunches from (say) a 1 GeV injection energy to a 10 GeV flat-top, and then hold them there for a couple of minutes, before injection into the eRHIC ring.

At fixed energy the polarization time scales like $T_{pol} \sim C\rho^2$, while the linear power load scales like $P_{lin} \sim 1/(C\rho^2)$. In other words, the price the booster pays for a reduced polarization time is a proportional increase in the heat load per meter of dipole. Fortunately, the booster accelerates only a fraction of the total number of bunches on each acceleration cycle, reducing the heat load per meter. If the linear power load of the booster is made equal to that of the electron ring by adjusting the number of bunches being accelerated, and if each acceleration cycle lasts exactly 74 s (the polarization time in the booster), then the total time to replenish the electron ring is 9.9 hours (the polarization time in the electron ring)!

On the other hand, the *total* average power scales like $P_{rad} \sim 1/(C\rho)$. Thus, a booster with a bend radius of $\rho = 33.4$ consumes only about 14% as much energy as the electron storage ring, in polarizing all the electrons.

2.2 Electron ring optics

The nominal geometry near the RHIC interaction point (for example during gold-gold collisions) is shown in Figures 5 and 6 (top). All six interaction regions are identical. The two beam trajectories which provide head-on collisions share a common beam pipe as they move from the IP. They begin to be separated 9.8 m from the IP in the 180 mm warm bore DX magnet. This is the only magnet through which both beams pass. After further separation in a 7.0 m long warm drift space, the trajectories enter separate side-by-side Blue and Yellow D0 dipoles, which bend the trajectories back to be approximately parallel. The beams then pass through side-by-side triplet quadrupoles. All eight magnets – D0, Q1, Q2, and Q3 in both Blue and Yellow beam lines – are contained in a single massive triplet cryostat. Unlike the DX magnet cryostat, the triplet cryostat cannot be moved.

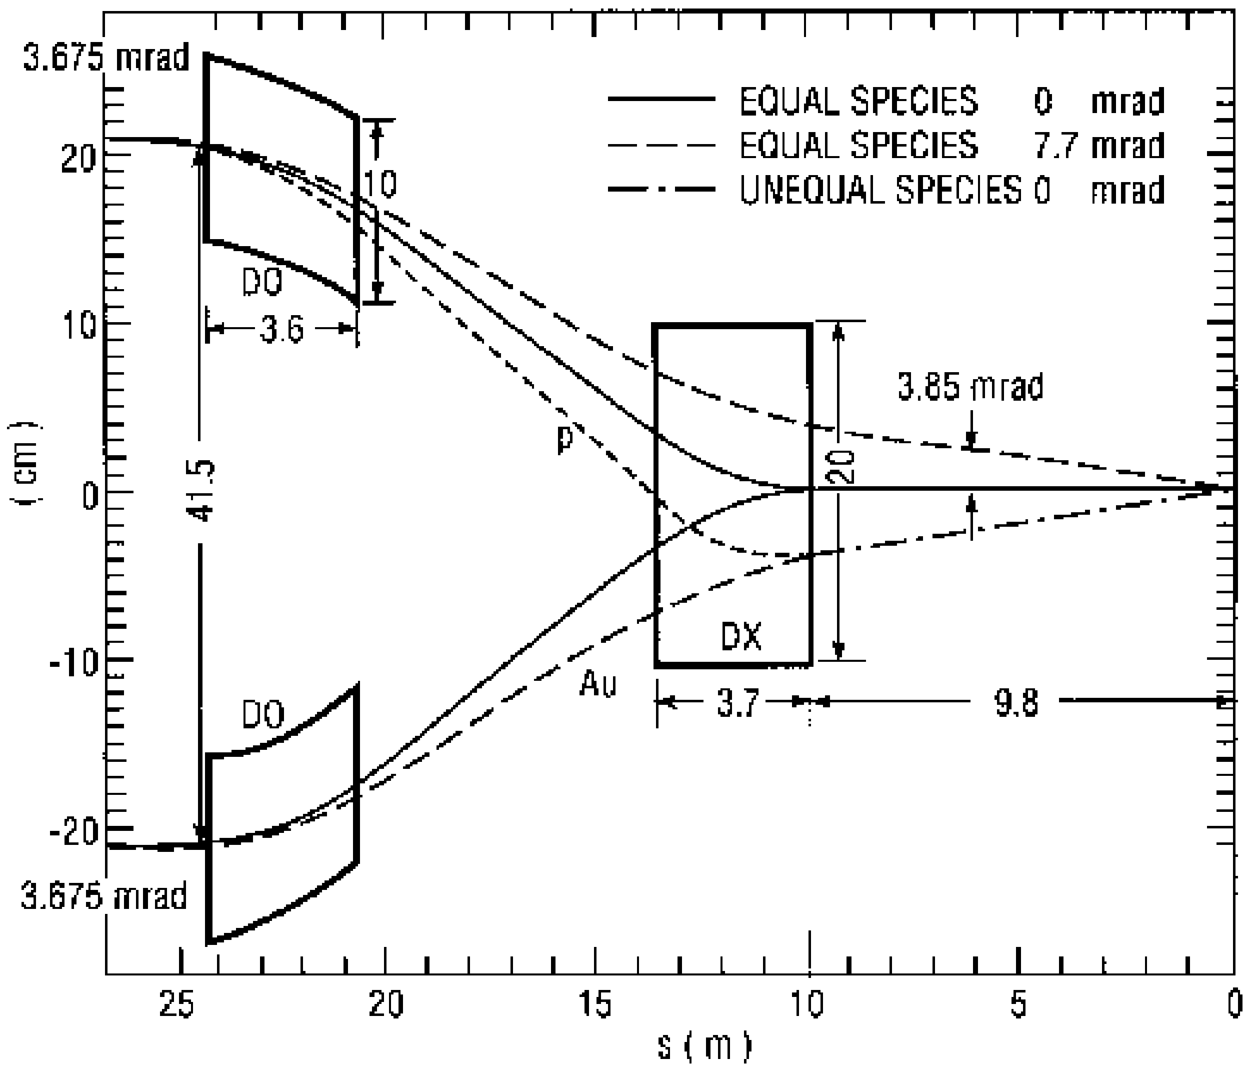

Figure 5: Nominal layout at the interaction region, showing a common crossing angle of 3.85 mrad during proton-gold collisions.

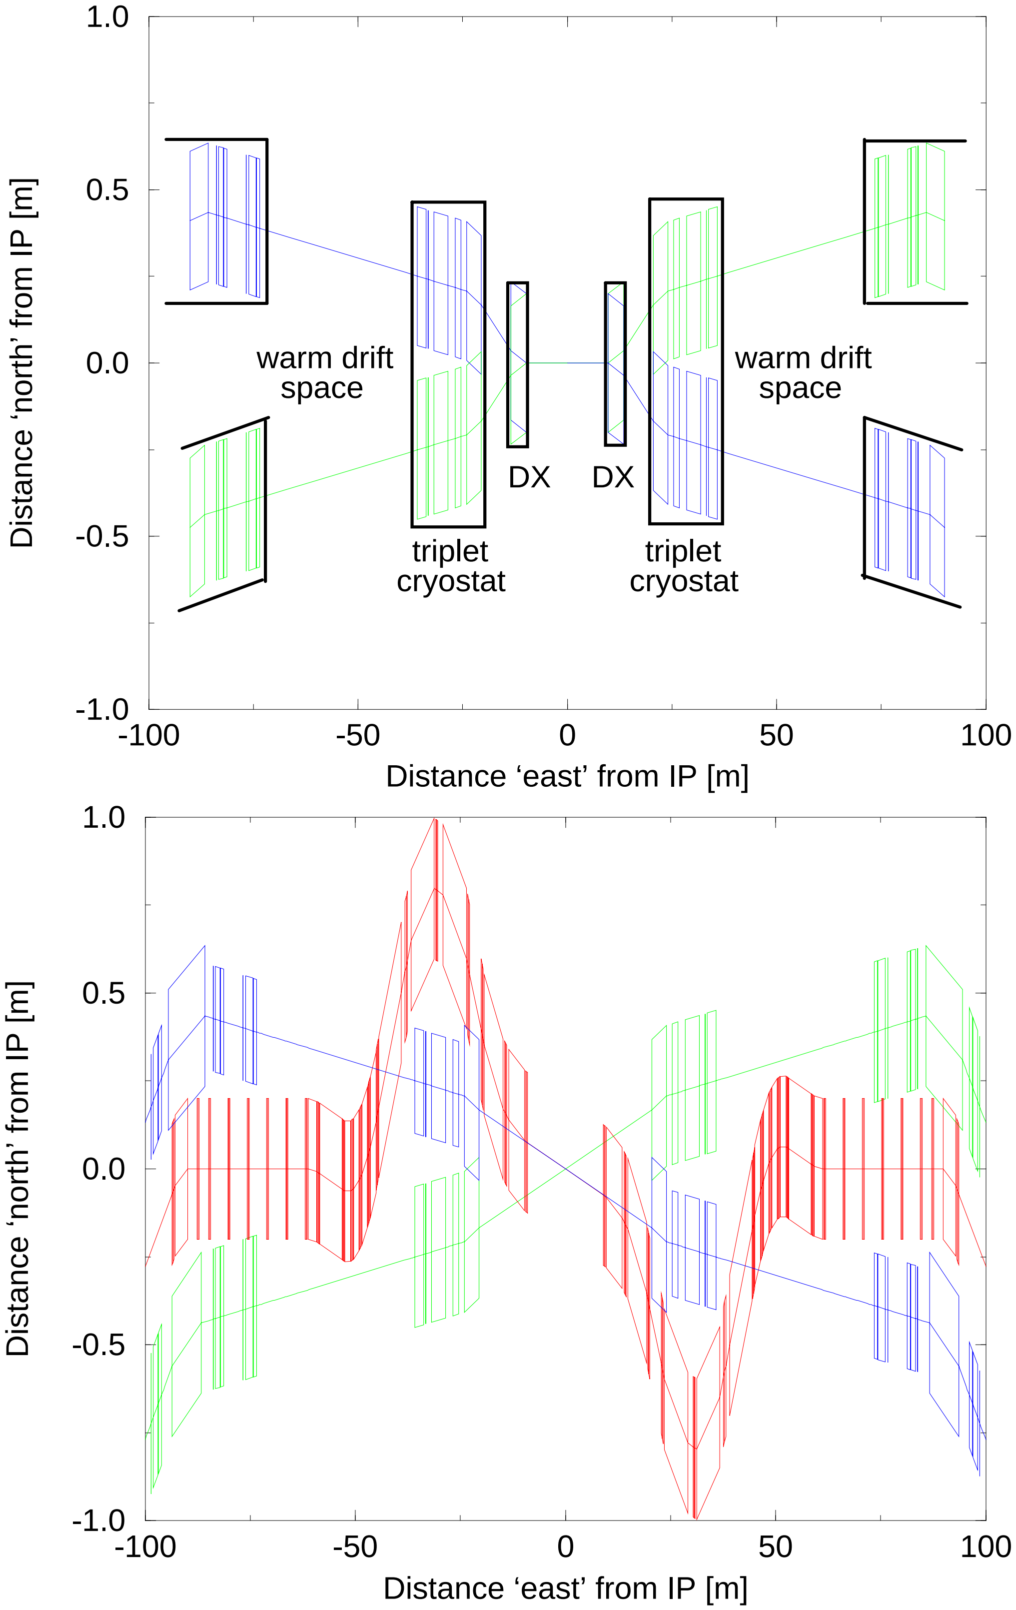

Figure 6: The magnet layout close to the IP in nominal RHIC operation during gold-gold or proton-proton collisions (top), and after modification for eRHIC (bottom).

After the ion trajectories leave the triplet cryostat there is a 35 m long warm drift space before they enter separate cryostats at the Q4 quadrupole. The quadrupoles have a constant FODO cell spacing after Q5, with partial dipole packing until Q9, where the full regular arc begins.

Figure 6 (bottom) shows that the design trajectories still collide head-on in the modified eRHIC straw man interaction region geometry [10]. The beams collide with a common angle of 8.3 mrad with respect to the center line axis, in order to minimize the deflection of the electron beam (and the synchrotron radiation close to the IP). The electrons are horizontally separated from the ions so that they just miss the D0 corner of the triplet cryostat. Beam separation is performed by two magnets. The common BS1 magnet deflects the electron beam by 15.0 mrad with a bending radius of 247 m, shielding the IP from the synchrotron radiation of the stronger BS2 septum magnet.

The magnetic rigidity is 25 times larger for ions than for electrons. Since the BS1 magnet deflects the ion beam by only 60 μrad, the ions essentially go straight through the IP from the D0 magnet on one side to the D0 magnet on the other side. It is probably not possible to use the existing DX magnets as BS1 magnets, since the operating field of 0.136 T is only about 3.2% of the maximum DX field of 4.25 T. Persistent current and other hysteresis effects generally preclude operating superconducting magnets – which usually have a dynamic range of less than 20:1 – at such low fields.

After passing the Q3 end of the triplet cryostat the electron beam trajectory is bent back horizontally to lie midway between the Blue and Yellow rings. The electron trajectory also drops vertically by almost a meter, ending up near the tunnel floor between the stands supporting the Blue and Yellow Q4 cryostats. From this point on the eRHIC lattice mimics that of the Blue and Yellow rings, remaining near the floor between the two existing rings.

2.2.1 An interaction region without spin rotators

In some cases spin rotation is not required. For example, it may not be necessary to provide longitudinally polarized electrons at the IP. Or, the electrons may pass only once through the IP, in the linac-ring collision scenario. Figure 7 shows the relatively simple straw man layout of an interaction region without spin rotators [10]. These optics are attractive because weaker dipole fields are permitted, with lower synchrotron radiation linear power loads.

The BS1 and BS2 beam splitting dipoles are sandwiched by three quadrupoles which restrict the electron β_{max} to 750 m with a β^* of 0.5 m. After the electron trajectory passes the corner of the triplet cryostat, the BH1 magnet bends it back parallel to its original direction. The vertical drop is accomplished in two steps 180 degrees apart in betatron phase, in order to provide local compensation of the vertical dispersion.

2.2.2 Spin rotators

Polarized electron experiments require the spin vector to be longitudinal at the IP, while the spin vector is naturally vertical in the arc of a ring. The transformation from vertical to longitudinal polarization (and back again) is achieved by placing additional constraints on the interaction region optics, in addition to the requirements of geometric and Twiss function matching.

Electron spin rotators use two main principles:

- small angle rotations commute while large angle rotations do not
- a small deflection of the beam causes a large deflection of the spin vector

For example, in order to rotate the spin vector of a 9.92 GeV electron by 90 degrees, the electron beam must be bent by an angle

$$\phi \;=\; 90\,\frac{0.44065}{E_0\,[\mathrm{GeV}]} \;=\; \frac{90}{22.5} \;=\; 4\,\mathrm{degrees} \;=\; 69.8\,\mathrm{mrad} \tag{12}$$

The resulting orbit deflection is about 2.8 times that in HERA-e, which stores 28 GeV electrons. The eRHIC spin rotators are also made more difficult than the HERA-e spin rotators by the availability of only one third as much longitudinal space. Fortunately, the vertical step down which is necessary in eRHIC is convenient for use in the spin rotator design.

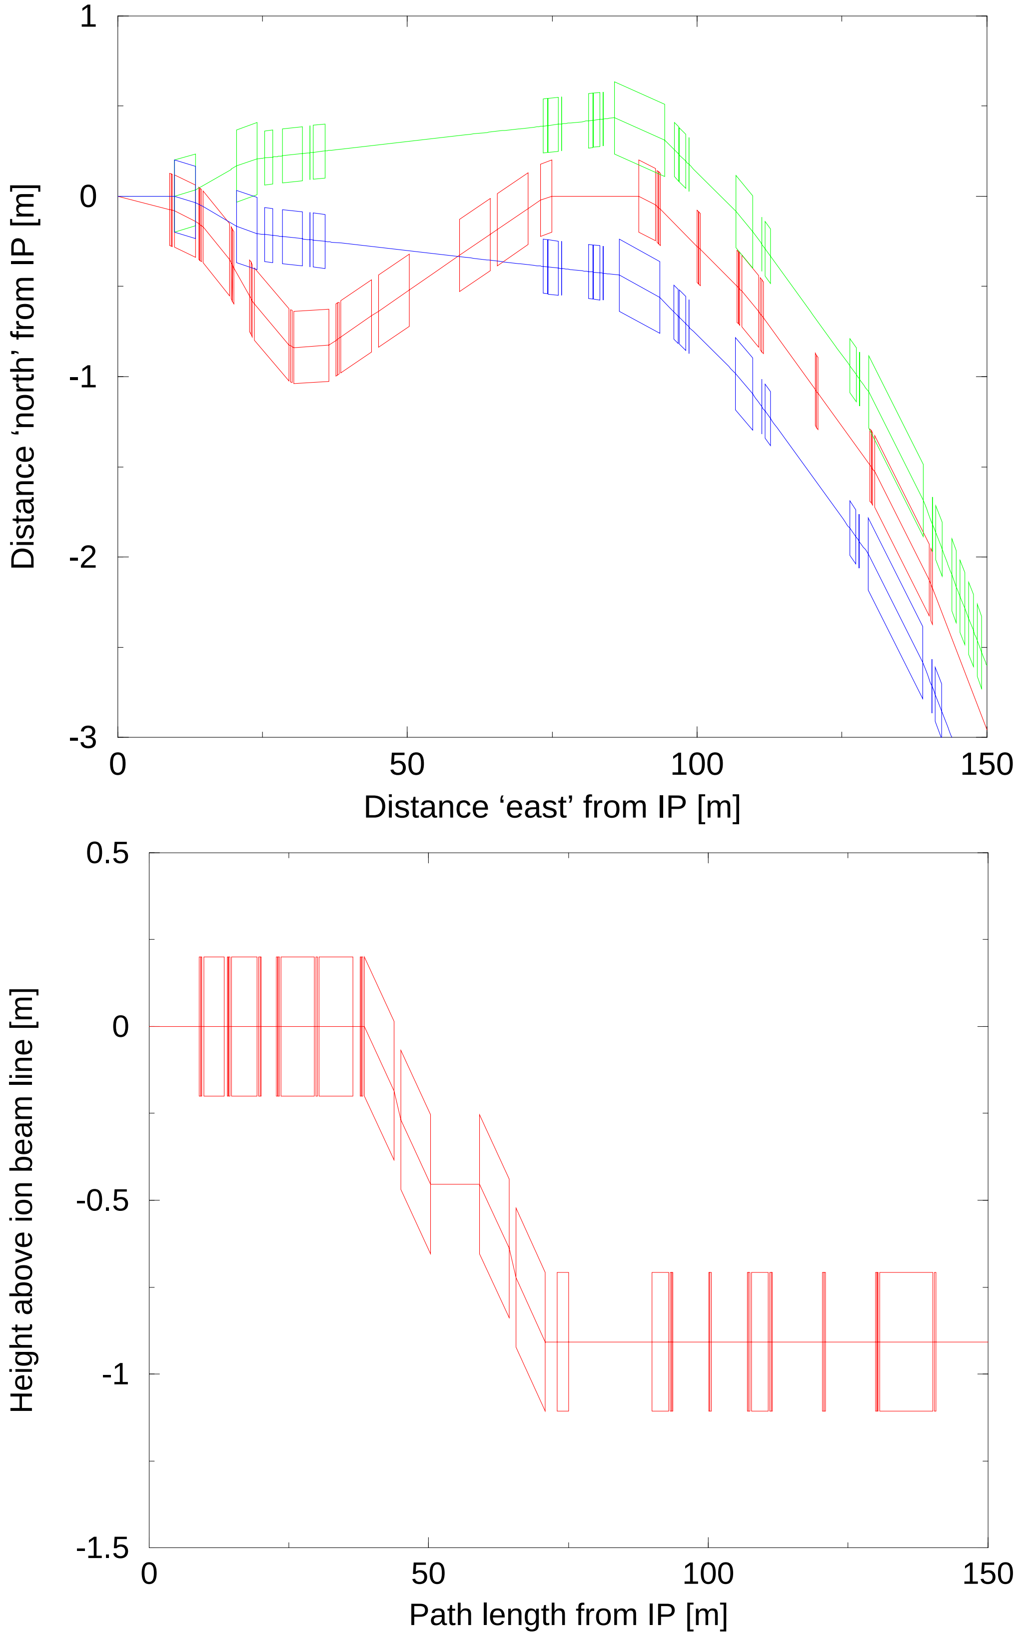

Figure 7: Plan and side views of the eRHIC magnet layout in an interaction region WITHOUT spin rotators.

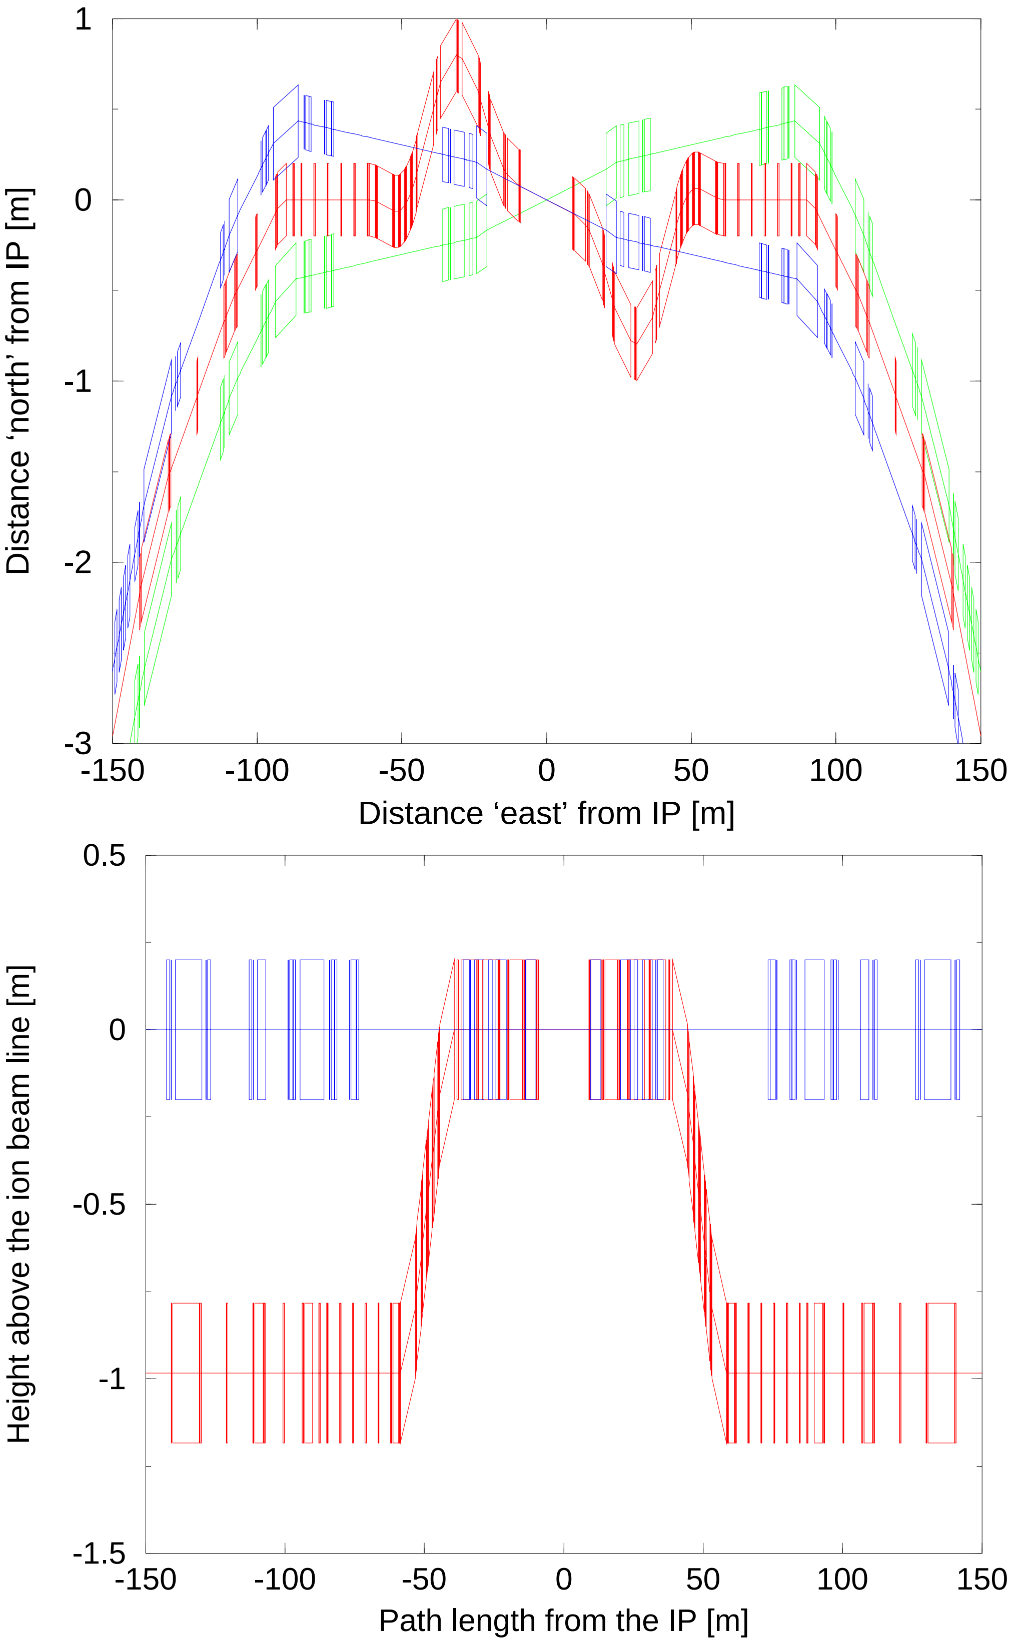

Figure 8: Plan and side views of the eRHIC magnet layout in an interaction region WITH spin rotators.

Figure 8 shows a straw man eRHIC interaction region layout, including spin rotators [10]. Coming from the IP the spin rotations in the BS1, BS2 and BH1 magnets cancel each other. The BH2 magnet rotates the longitudinal spin by 90 degrees around the vertical axis, leaving it pointing horizontally. The BV1 magnet rotates the horizontal spin vector around the horizontal axis, leaving it unchanged. The electron beam trajectory now points downwards by 4 degrees. The BH3 magnet undoes the rotation of the BH2 magnet, making the spin vector longitudinal again. The BV2 magnet returns the beam trajectory to the horizontal plane, and rotates the spin into the vertical direction. The spin rotation is now complete. Horizontally bending dipoles such as BH4 and BH5 follow the spin rotator, to properly place the electron beam lines between the feet of the Blue and Yellow Q4 cryostats. Integration into the rest of the arc is also shown in Figure 8.

The vertical step down is both a blessing and a curse. Since the spin rotation requires a bending angle of 4 degrees, the distance between the exit of BV1 and the the entrance of BV2 is only about 6.5 m; too little for local dispersion matching. Figure 9 shows how the Twiss functions in the vicinity of an IP with spin rotators match into the neighboring arcs. It confirms that the optical, geometric, and spin constraints can all be simultaneously satisfied.

When an electron emits a photon, it generates betatron and synchrotron oscillations. This can lead to a diffusion of the polarization of a bunch, if the stable spin direction is not maintained for first order perturbations in energy and betatron amplitude. In principle "spin matching" could be performed by the quadrupoles which are interspersed between the dipoles in the spin rotator [11]. This is an 8-by-8 matrix optics problem, which is addressed by codes such as SLIM [12] and SITROS [13]. Spin matching is not performed in the straw man optics presented here, because the 10 hour polarization time makes it unnecessary.

Spin rotators only work for one beam energy because the integrated magnetic field necessary is independent of the energy. A small energy range can be achieved by installing a vertical dipole corrector in front of the BH2 magnet. If this corrector deflects the beam by θ the deflection angle of BV1 and BH3 must be reduced and BH4 must be increased by θ so that the spin rotation and the total orbit deflection is unchanged. This procedure distorts the closed orbit inside the spin rotator. Thus, the energy range of the spin rotator is limited by the aperture of its magnets.

2.2.3 The utility interaction region, IP4

The interaction region at IP4 has no detector but is used for utilities like injection, RF, electron beam cooling, et cetera. There is no reason to raise the eRHIC ring above its (near floor level) horizontal plane. Instead the beam line goes to the outside of the ring to bypass D0 and DX on both sides of the IP, and then returns to the middle of the tunnel. While all other interaction regions are geometrically antisymmetric around the IP this interaction region is symmetric, as illustrated in Figure 10.

The spin rotators described above are actually Siberian snakes of the second kind: they rotate the spin direction by 180 degrees around the radial axis. In order to have longitudinal spin at the other five IPs the optics around IP4 must include a spin rotator which is a Siberian snake of first kind. This snake is located at one end of the straight section and consists of an arrangement of horizontal and vertical dipoles. Since there are no quadrupoles inside the snake it is an achromat in both horizontal and vertical directions [14].

2.2.4 The "high noon" interaction region, IP12

This interaction region is available for a major new detector, which can be tailor built for electron-ion operations. It may be possible, for example to:

- build a vertical step into the Yellow ring to bypass the IP, allowing electron-ion collisions in this IR, and ion-ion collisions in all other interaction regions
- remove the triplet cryostat assembly and replace D0 with a magnet which doubles as a spectrometer for hadron fragments

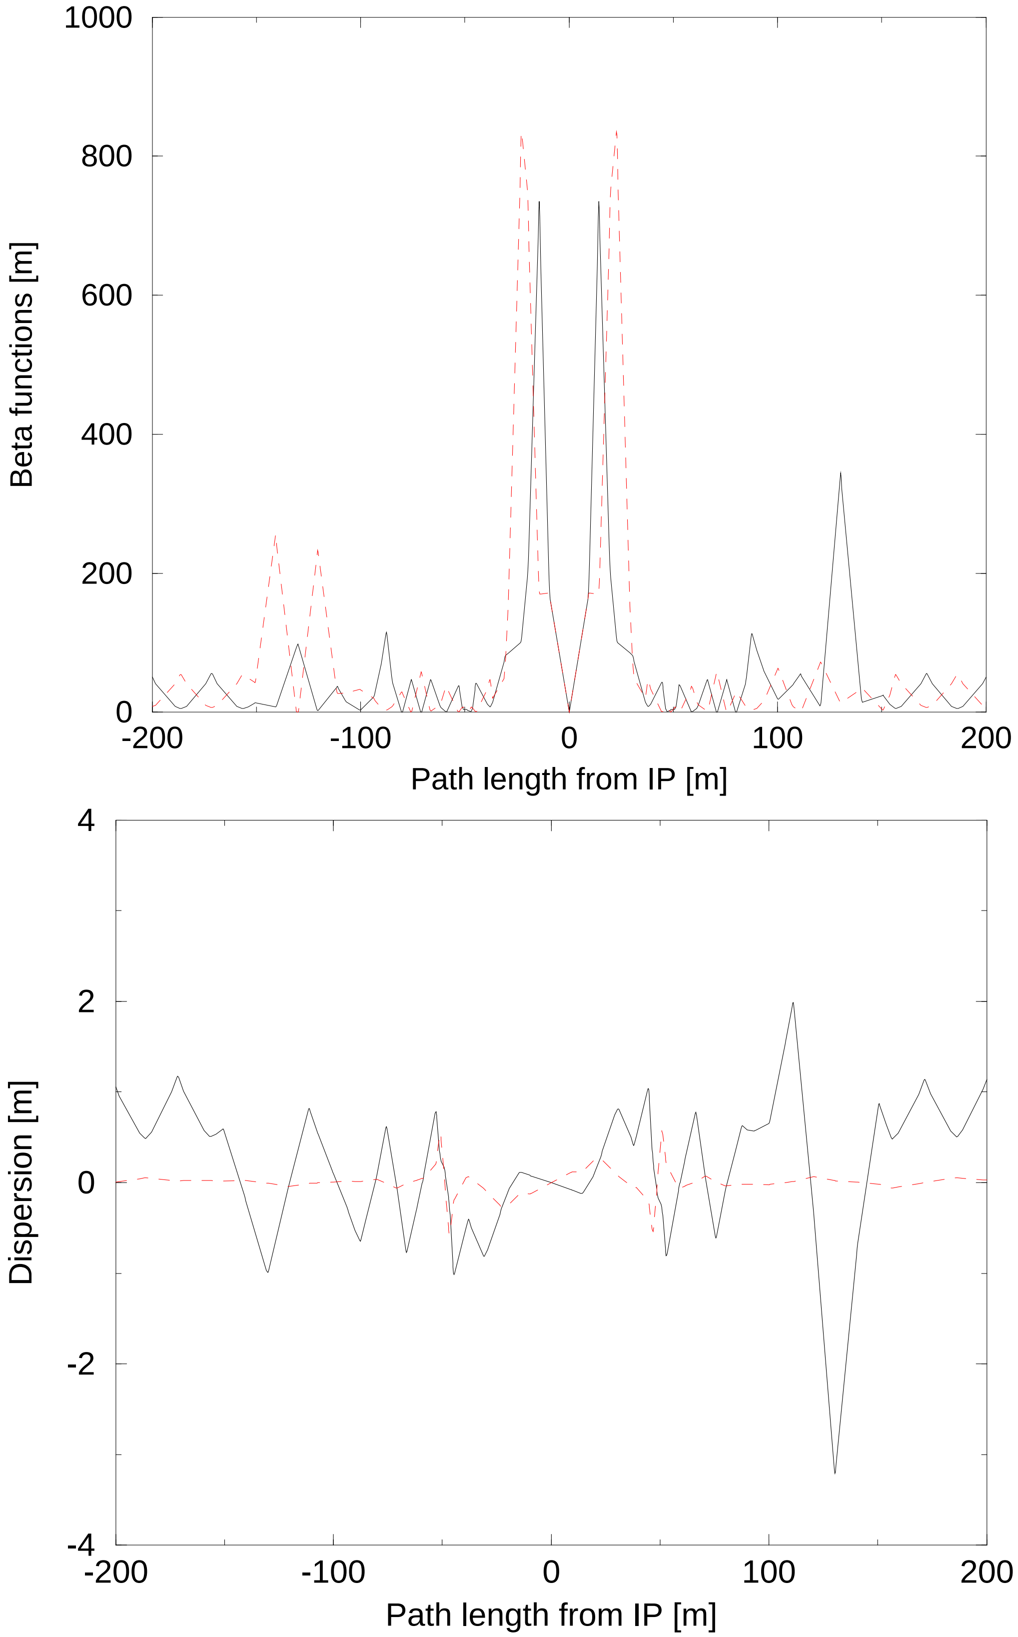

Figure 9: Beta and dispersion functions around an IP with spin rotators.

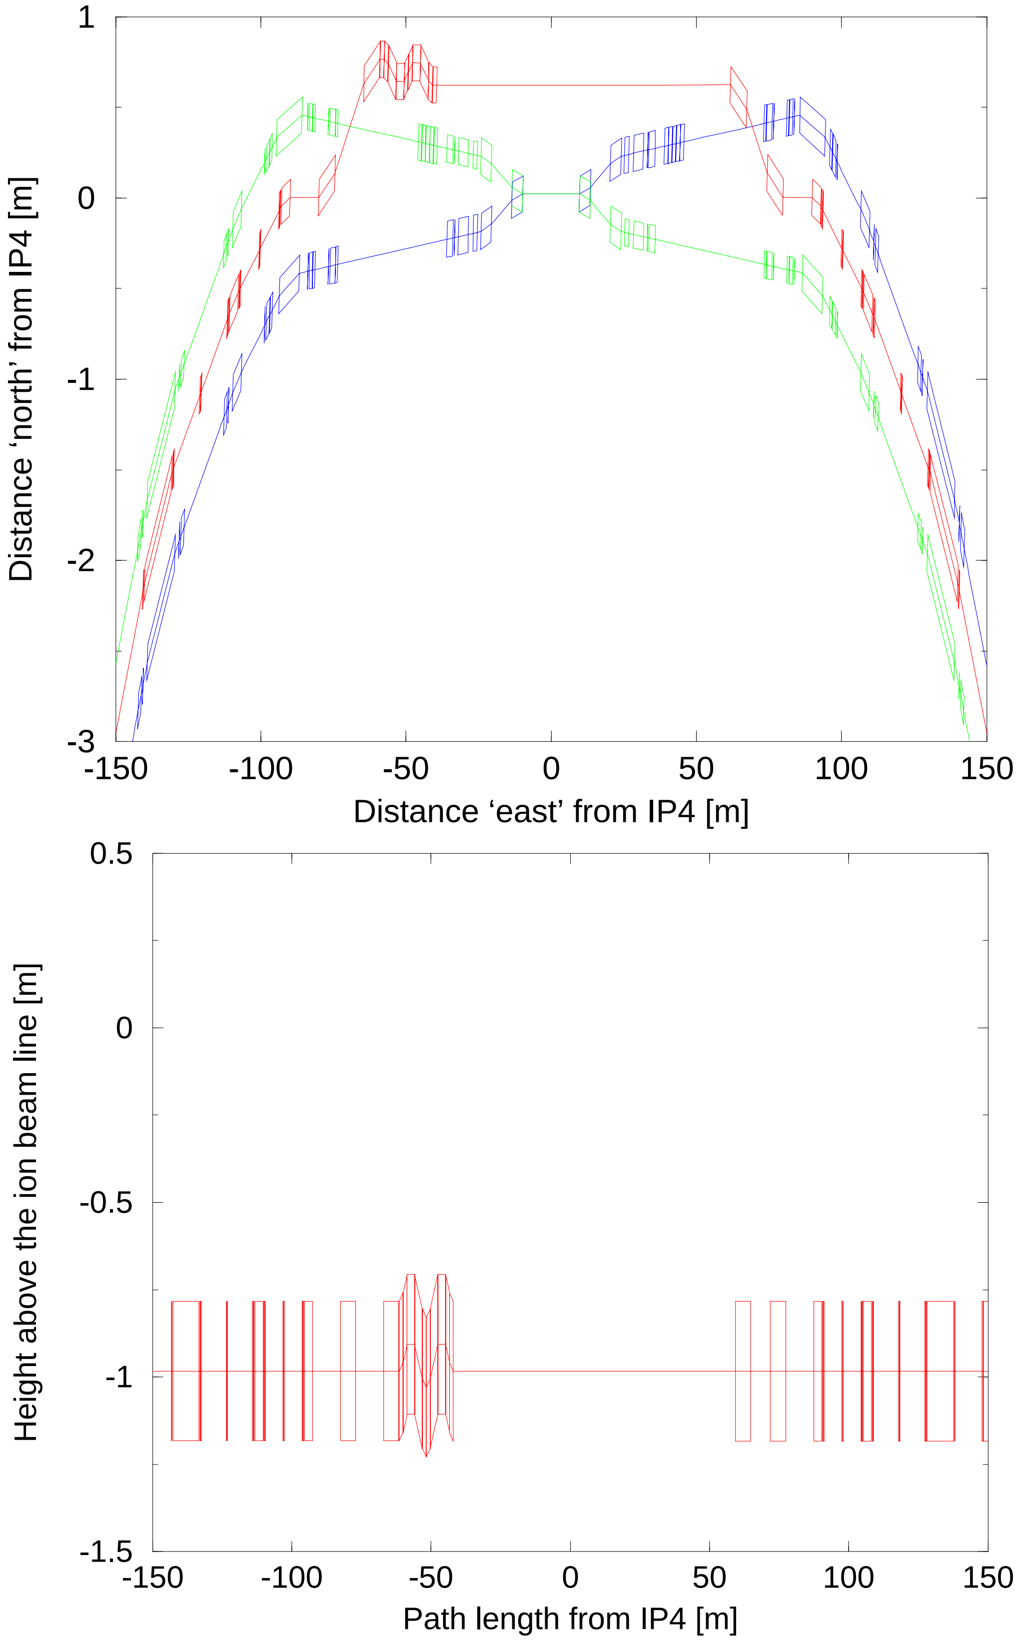

Figure 10: Plan and side views of the magnetic layout around the utility region IP4, including a Siberian snake of the first kind.

2.3 Electron beam intensity limitations

It was suggested at the outset of the electron ring design study that the High Energy Ring (HER) of the B-Factory complex at SLAC [15, 16] could serve as a marvelous prototype for eRHIC [17]. The HER is roughly the same energy as eRHIC, 9 versus 10 GeV, the circumference is within a factor of two, 2200 versus 3833 meters, and the goal of both rings is to achieve high luminosity using many electron bunches. Some of the key parameters of the two rings are listed in Table 4.

This section takes a first pass at addressing the beam current limitations in eRHIC. The HER is used as a "ruler" for comparison with eRHIC parameters.

Ring		SLAC HER	eRHIC
Energy, E_e	[GeV]	9	10
Circumference, C	[m]	2200	3833
Bend radius, ρ	[m]	165	242.8
Energy loss per turn, U_0	[MeV]	3.52	3.64
Emittance, ϵ_e	[nm]	50	40
Number of bunches, N_b		415/831/1658	360/2520
Bunch spacing, τ	[ns]	16.8/8.4/4.2	35.5/5.1
Compaction factor, α		2.4×10^{-3}	0.61×10^{-3}
RF frequency, f_{rf}	[MHz]	476	394
Harmonic number, h		3492	5040
Bunch length, σ_L	[mm]	11	9
Energy spread, σ_E		6×10^{-4}	5.5×10^{-4}
Beam current, I	[A]	3	3
Total synchrotron power, P	[MW]	10.56	10.92
Linear power load, P_{lin}	[kW/m]	10.2	7.2

Table 4: Key parameters for SLAC HER and eRHIC.

2.3.1 Synchrotron radiation from the arc dipole magnets

In contrast to the RHIC heavy ion rings, the properties of the electron beam in the 10 GeV electron storage ring are determined by synchrotron radiation. Synchrotron radiation is responsible for the damping of the electron beam and also for the excitation of the beam through quantum fluctuations caused by the emission of photons. In one turn of an isomagnetic storage ring with identical bending magnets of the same radius ρ, an electron of energy E loses [18]

$$U_0\ [\text{MeV}] \;=\; 0.0885\ \frac{E^4\ [\text{GeV}^4]}{\rho\ [\text{m}]} \tag{13}$$

and the total power radiated by a circulating electron current I is given by

$$P\ [\text{MW}] \;=\; U_0\ [\text{MeV}]\ I\ [\text{A}] \tag{14}$$

The average current limit in eRHIC is based upon the design specifications for the HER at the SLAC B-factory. In the HER the limit is set by the allowable heat load on the copper vacuum chamber [19]. Table 4 lists the relevant parameters for the SLAC HER and eRHIC storage rings. For the purpose of this table we assume an approximate model of eRHIC which is made up of only dipoles of radius 242.8 m. In the next section we discuss the modifications imposed by the spin rotator magnets.

In Table 4 we have chosen the same maximum current of 3 A for both rings. Since the power load specification for the HER was actually set based on the power per meter of dipole, $P_{lin} = P/2\pi\rho$, one could postulate increasing the current in eRHIC to 4.27 A without exceeding the HER power load specification.

2.3.2 Synchrotron radiation from the spin rotator magnets

The above discussion concerned only the gross parameters of eRHIC. In order to provide the desired spin rotation for polarized electron beams it is necessary to include some dipole magnets with much higher fields than those mentioned above. A summary of these magnets is given in Table 5 along with the power load, P_{lin}, for an assumed average ring current of 1 A. For comparison, note that the electron beam currents at the beam-beam limit in Table 3 vary from a minimum of 0.45 A to a maximum of 7.7 A.

Magnet	θ [mrad]	L [m]	ρ [m]	B [T]	P_{lin} [kW/m]
arc dipole	38.9	9.44	242.5	.136	2.4
beam splitters					
BS1 (DX)	15.0	3.70	246.7	.134	2.2
BS2	34.2	4.50	131.8	.250	8.2
no rotator:					
BIH1	39.4	6.00	152.0	.217	6.1
BIH2	34.9	5.50	157.6	.209	5.7
BV1,BV2	61.0	7.50	122.8	.269	9.3
BIH4	21.4	3.00	139.9	.236	7.2
BH5	32.1	7.80	242.5	.136	2.4
BH6,BH9	12.2	2.95	242.5	.136	2.4
BH8	38.9	9.44	242.5	.136	2.4
spin rotator:					
BH2	69.8	5.50	78.8	.418	22.7
BV1,BH3,BV2	17.5	1.33	75.7	.436	26.2
BH4	12.9	2.00	154.8	.213	6.0
BH5	29.0	2.95	101.7	.324	13.6
BH6,BH9	12.2	2.95	242.5	.136	2.4
BH8	38.9	9.44	242.5	.136	2.4

Table 5: Magnet parameters for the eRHIC arc dipole, interaction region dipoles, and for the spin rotator dipoles. The synchrotron radiation linear power load is for a 1 A electron beam current.

The high field bends BH2, BH3, BV1, and BV2 in the spin rotator produce power loads that are a factor of 2 higher than the 10.2 KW/m limit in the HER design. In the straw man optical design presented here, the spin rotator magnets set the eRHIC electron current limit to roughly 0.5 A. More analysis must be done to see if the power load limits can be increased or if a different approach to the spin rotation is possible, because the high current and high luminosity potential of eRHIC is greatly compromised.

2.3.3 Longitudinal single bunch current limits

The "zero current" energy spread and bunch length are given in Table 4. The key question is whether or not these values are preserved as the single bunch current is increased. The first issue to be addressed is potential well distortion which can modify the bunch shape but produces no change in the energy spread. Such an analysis requires a detailed impedance budget for the entire ring which is beyond the scope of the present work. Conventional wisdom says that eRHIC will end up with a combination of inductive and resistive impedances. The magnitude of the inductance determines the amount of bunch lengthening.

The performance of the HER is a useful guideline. The zero current bunch length is 11 mm. Inductive potential well distortion is responsible for increasing this to roughly 13 mm as the single bunch average current is raised from zero to 3 mA [20]. Note that if all of the 1658 bunches are filled this would correspond to a total average current of 5 A with only a 20% increase in the electron bunch length and no change in the energy spread. The low energy ring at the SLAC B-Factory (LER) has stored currents up to 12 mA in a single bunch [20].

Beyond potential well distortion, the dominant effect of concern for longitudinal phase space is the microwave instability. A detailed analysis requires a complete impedance budget for the ring. At the present level of discussion the "Boussard conjecture"

$$eI_p \frac{Z_L}{n} = 2\pi\alpha E\sigma_E^2 \tag{15}$$

is used to gauge the importance of the microwave instability. In this equation I_p is the peak current in amperes, Z_L/n is the longitudinal coupling impedance, α is the ring momentum compaction, and σ_E is the energy spread. As can be seen from Table 4 the relevant parameters in the two rings are similar – the momentum compaction is the most significant difference. If eRHIC employs the same impedance policing used for the HER, the microwave instability should not be a show stopper. This issue needs to be revisited once the lattice is further developed and the current limit imposed by the spin rotator magnets is resolved.

The beam-beam effects for eRHIC have been discussed in an earlier section. As a point of reference, it should be noted that at present, the beam-beam interaction sets the current limit in the HER [20].

2.3.4 Multi-bunch current limits

At the present level of design we do not have the details of the RF system in order to assess the impact of higher order modes on longitudinal and transverse coupled bunch instabilities. We postulate that these effects would be dealt with using broadband feedback systems similar to those used for both the LER and the HER at the SLAC B-Factory.

2.3.5 Ion trapping

Virtually all electron storage rings share the problem of trapping positive ions that are created from the background gas in the vacuum chamber. Fortunately this problem can be ameliorated or eliminated altogether by simply providing a "gap" in the filling of the bunches around the circumference of the ring. For example, the VUV ring at the NSLS fills only 7 of 9 RF buckets, while the X-Ray ring fills only 25 of 30 buckets. These gaps eliminate the problems of ion trapping. In the HER there is a 5% gap of empty buckets to solve this problem [20]. We will pursue a similar strategy for eRHIC.

2.4 Ring energy tunability

There are no resonance phenomena that dictate a specific energy for eRHIC. Since the original eRHIC concept [4] was to have a full energy superconducting linac as the injector, it is desirable to not increase the energy of the ring beyond what is really required. The energy of 10 GeV was selected as a compromise value. However, at least one experimentalist has expressed a desire to be able to reduce the energy from its maximum value. There are a few key areas to address when considering an energy variation of the ring: synchrotron radiation, magnets and RF power, shielding, electron polarization, top off mode, and cost. In what follows we survey the impact of both raising and lowering the energy of eRHIC.

2.4.1 Energies above the nominal 10 GeV

The maximum allowable current in eRHIC is set by the power loading on the vacuum chamber due to synchrotron radiation (as discussed above). It is clear from Equations 13 and 14 that the beam current must be reduced as $I \propto E^{-4}$, if the linear power load is held constant at its maximum value. Thus if the ring energy is pushed to 15 or 20 GeV the maximum beam current would have to be reduced by factors 5 or 16 respectively. The experimentalists must judge if this results in acceptable luminosities. However, we must also assess the effects on other machine parameters before drawing any conclusions.

Two effects not included in the simple power loading analysis above are the changes in the synchrotron radiation spectrum, and in the vertical opening angle. For a fixed lattice the critical photon energy is given by

$$\varepsilon_c\ [\mathrm{keV}] = 2.22\ \frac{E^3\ [\mathrm{GeV}^3]}{\rho\ [\mathrm{m}]} \tag{16}$$

and the vertical opening angle for photons emitted near the critical energy is given by

$$\psi_c = \frac{1}{\gamma} \tag{17}$$

Further analysis must be done to determine if either of these effects has an impact on the power loading of the vacuum chamber and hence the energy tunability of eRHIC. The above discussion of synchrotron radiation only considered the linear power loading, $P/2\pi\rho$.

The electron beam parameters are also functions of the electron beam energy. Perhaps key among the many parameters is the electron beam emittance, which for a given lattice scales as

$$\epsilon_e \propto F(\nu_x)\gamma^2 \tag{18}$$

where $F(\nu_x)$ is a form factor determined mainly by the horizontal betatron tune. If the ring were operated at 20 GeV the beam size would double unless the horizontal tune could be increased to compensate for this. At this stage of the design the flexibility of the lattice has not been fully explored.

The energy acceptance of the RF system is chosen to be 1% at this stage of the design process. The over-voltage is defined as the ratio of the RF voltage to the energy loss per turn, V_{RV}/U_0. As can be seen from Equation 13, U_0 increases as E^4, so the RF voltage will have to be increased considerably if the energy of the ring is raised. This will add to the cost and complexity of eRHIC.

The field in the arc dipole magnets is given by,

$$B \text{ [T]} = \frac{E \text{ [GeV]}}{0.3 \; \rho \text{ [m]}} \tag{19}$$

yielding fields of 0.137 and 0.274 T at 10 and 20 GeV respectively. These field values are not prohibitive but this flexibility must be built in at the outset. Clearly, the use of permanent magnets is not conducive to large energy variations in the ring. The type of electromagnetic dipole, solid versus laminated, could also be effected by the desire to vary the ring energy. The length of all magnets would be chosen to accommodate the maximum desired ring energy.

At this stage of the design we have not performed any estimates of the radiation protection shielding required for eRHIC even for the nominal 10 GeV operation. It is obvious that operation in excess of 10 GeV will have an impact on the amount of shielding required, and this will increase the cost.

The above items could be engineered from the outset to provide for higher energy operation of eRHIC. However, the polarization of the electron beam remains as the key impediment to energy tunability. Each of the four dipoles in the spin rotator rotates the spin by 90 degrees. To do this each dipole needs a fixed integrated field, independent of energy. If the dipoles are ramped to change the ring energy the spin rotation is destroyed. Using some trim dipole compensation schemes it might be possible to vary the energy by $\pm 3\%$. This is an area where the linac scheme may have distinct advantages. The bottom line on energy tunability is that if polarized electrons are required, then the ring energy is more or less fixed!

Another important concern for higher energy operation is that if the linac energy is fixed at 10 GeV, then the ring can not be operated in top-off mode to constantly replenish the electrons lost from the ring.

2.4.2 Energies below the nominal 10 GeV

As is obvious from Equations 13 and 14, the problem of synchrotron radiation power loading disappears rapidly as the energy goes down. However, it must be remembered that the synchrotron radiation not only provides heating, but is also the source of damping. The synchrotron radiation damping time is given by,

$$\tau_i \text{ [ms]} = \frac{C \text{ [m]} \; \rho \text{ [m]}}{13.2 \; J_i E^3 \text{ [GeV}^3\text{]}} \tag{20}$$

where $J_i = x, y, e$ are the damping partition functions.

A reduction in the beam energy by a factor of two increases the damping time by nearly an order of magnitude. Instability thresholds must be reconsidered, and one must check that Intra-Beam Scattering is not an issue for the assumed high intensity beams. If additional damping is required one could consider including damping wigglers.

Since the emittance of the beam scales quadratically with the ring energy (in the absence of IBS), then halving the energy also halves the electron beam size at the IP. As discussed below, the gas scattering and Touschek lifetimes of the electron beam will decrease with the ring energy.

The bottom line consideration in lowering the ring energy is in the loss of polarization – a consideration already observed in the previous section on raising the ring energy.

2.4.3 Electron beam lifetime

A detailed study of the electron beam lifetime is beyond the scope of the present study. However, some general remarks can be made by contrasting eRHIC with the HER. Ignoring for the moment collisions with gold ions or protons, the electron beam lifetime is normally dominated by the gas scattering lifetime and the Touschek lifetime.

The first of these is dominated by the residual vacuum and the energy of the electron beam. If we assume eRHIC can achieve a similar vacuum pressure to that in the HER, we should expect a longer lifetime than in the HER simply because eRHIC operates at a higher energy, 10 GeV versus 9 GeV. The optics of the ring also enters the gas scattering lifetime through the ring acceptance and the maximum betatron functions. The impact of the optics needs to be reconsidered when the lattice details have evolved further.

The Touschek lifetime is dominated by the charge density in a single bunch and is a strong function of the ring energy, because the mechanism for loss is Coulomb scattering of the charged electrons with one another. In longitudinal phase space, eRHIC peak currents comparable to those in HER. In the transverse plane, the eRHIC emittance quoted in Table 4 is approximately the same as in the HER. The higher eRHIC energy is an advantage, and it is also possible to increase the RF acceptance to achieve an acceptable Touschek lifetime. More analysis is needed as the design progresses.

Finally, the electron "burn off" loss rate due to collisions with gold ions or protons must be included in the lifetime analysis, as this effect could dominate.

2.5 Ion beam intensity limitations

The list of ion current limitations addressed in this subsection is not complete. For example, there are significant technical and financial issues associated with increasing the data acquisition rates of the experiments, as the number of bunches (and/or the luminosity) is increased.

2.5.1 Radio Frequency system

In normal RHIC operations ion bunches are accelerated in the $h = 360$ (28 MHz) RF acceleration system, and then rebucketed into the $h = 2520$ (197 MHz) storage system. The nominal rebucketing procedure is begun by suddenly shifting the 28 MHz phase by 180 degrees for a small fraction of a synchrotron period, and then shifting the phase back for about 3/8 of a synchrotron period, to get an "upright" beam ellipse with the correct aspect ratio. Then the 28 MHz system is snapped off at the same time that the 197 MHz system is snapped on.

Rebucketing is done for two reasons. First, the high frequency system applies a much larger voltage gradient dV/ds, so that the bunches can be made much shorter than β^* (avoiding the hourglass effect), with a reasonably large bunch area (to contain longitudinal emittance growth due to IBS). Second, the 197 MHz cavities can not handle the relatively large frequency swing that is necessary to accelerate gold ions from $\gamma \approx 11$ to $\gamma \approx 108$.

If the 28 MHz RF is adiabatically turned off at storage energy, and then the 197 MHz system is adiabatically turned on, then 2520 evenly populated bunches are created. Some of these bunches would have to be removed in order to create a beam abort gap, but this should not be too difficult.

In order to store between 360 and 2520 ion bunches, an alternative to "direct" and "adiabatic" rebucketing must be found. One way is to modify the existing 197 MHz cavities to accept a frequency swing large enough to accelerate gold ions. The feasibility of such a modification has yet to be established, but it would enable the storage of an (essentially) arbitrary number of ion bunches, between 360 and 2520.

Another way is to develop an RF system at an intermediate frequency. The possibility of developing a purpose built superconducting RF system at a frequency of, say, 84 MHz ($h = 3 \times 360 = 1080$) or 113 MHz

($h = 4 \times 360 = 1440$) has recently re-entered discussion, as a way of combating IBS effects which are particularly bad in longitudinal phase space [21]. Adiabatic rebucketing with such an RF system would also fit naturally into an advanced upgrade of the ion-ion luminosity performance.

2.5.2 Kickers

RHIC is currently filled one bunch at a time, from a small number of bunches (for example 4) which have been accelerated in the AGS. The injection kicker therefore has a relatively short rise time, flat top, and fall time, of approximately 90 ns, 45 ns, and 90 ns, respectively.

Large numbers of bunches would naturally be injected into RHIC in "boxcar" fashion – many bunches with the correct spacing at one time. Ideally, for example with protons, it would be possible to inject a "bunch train" – filling the entire circumference of the AGS – into RHIC in one turn of the AGS. A new RHIC injection kicker system would have to sustain an injection pulse with a long smooth flat top. It also means that the AGS needs to be capable of preparing a bunch train with the appropriate parameters.

So long as the rise time and fall time of the injection kicker system are much smaller than the flat top time, their values are not crucial, since the presence of a few "holes" in the bunch structure is inevitable. Some of these holes are necessary and desirable. For example, judiciously placed holes help combat the electron cloud effect (see below) by allowing the clouds to clear. An abort gap approximately 0.5 μs long continues to be necessary, to accommodate the rise time of the abort kicker. It is possible (although perhaps not likely) that the current abort kicker system will be adequate for a much larger number of bunches.

2.5.3 Abort dump

The nominal RHIC beam dump is an internal target, located between the triplet quadrupoles and Q4 on either side of IP10 [22]. As it stands, the dump has a safety factor of about 3 or 4 beyond the "nominal worst case" scenario, in which the carbon-carbon core of the dump is thermally shocked by 60 gold bunches of 10^9 ions, and an emittance of 15π microns. Clearly, the beam dump needs a radical upgrade if very many more than 120 ion bunches are to be stored. Fortunately this is not technically difficult, with the addition of an extraction septum and a conversion to an external dump system. In principle any amount of stored energy can be aborted – if an elaborate enough external dump is constructed.

Even with an external dump, the downstream Q4 remains vulnerable to quenching through accidental spray during an abort [23]. One relatively simple solution for Q4 would be to reconstruct it as a custom magnet with a beam pipe liner [24].

Additional shielding will be necessary around RHIC, and other more formal requirements will need to be met, in order to raise the Operational Safety Limits and to permit much more intense stored beams.

2.5.4 Long range beam-beam interactions

In the nominal RHIC layout the beams begin to be magnetically separated (in the DX magnet) at only 9.80 meters from the IP, *before* entry into the first quadrupole. They enter separate beam pipes at a crotch 15.70 meters from the IP. Similar (or even better) values will apply in eRHIC layouts.

Such early separation has the distinct advantage of immunizing RHIC against the ravages of parasitic long range beam-beam interactions. For example, it has already been recognized that a crossing angle is required with proton-proton collisions only when 180 or more bunches are stored, and that a crossing angle may not be required for gold-gold collisions even with 360 bunches [25, 26]. It is conservative to make the total crossing angle $\alpha = 7\sigma'^*$ when one is required, in which case the tune shift per parasitic crossing is approximately

$$|\Delta Q_{LR}| \;=\; \frac{2}{7^2}\,\xi \;\approx\; .04\,\xi \tag{21}$$

It has been shown that such crossing angles are not difficult with typical RHIC emittances and apertures [26].

Under these conditions it is permissible to have as many as 20 or 30 parasitic collisions per interaction region, since the long range interactions mainly generate tune shifts, not tune spreads, and barely drive nonlinear resonances. Parasitic collisions are spaced by about 0.76 m with $N_b = 2520$ bunches, leading to about 25 parasitic collisions per IR if the beams begin to be magnetically separated ± 10 m from the IP.

Although long range beam-beam interactions are an important concern, they do not rule out scenarios with even 2520 bunches.

2.5.5 Cryogenic beam pipe and BPM signal cable heating

The image current of the beam which flows in the vacuum chamber walls causes resistive heating. This is not a concern in the sections of beam pipe at room temperature, but has the potential to be a serious problem when the heat is deposited at cryogenic temperatures – in the ion rings. A maximum average cryogenic heat load of about 0.5 to 1.0 Watt per meter can be tolerated during continuous running, to stay within the capacity of the RHIC cryogenic refrigerator.

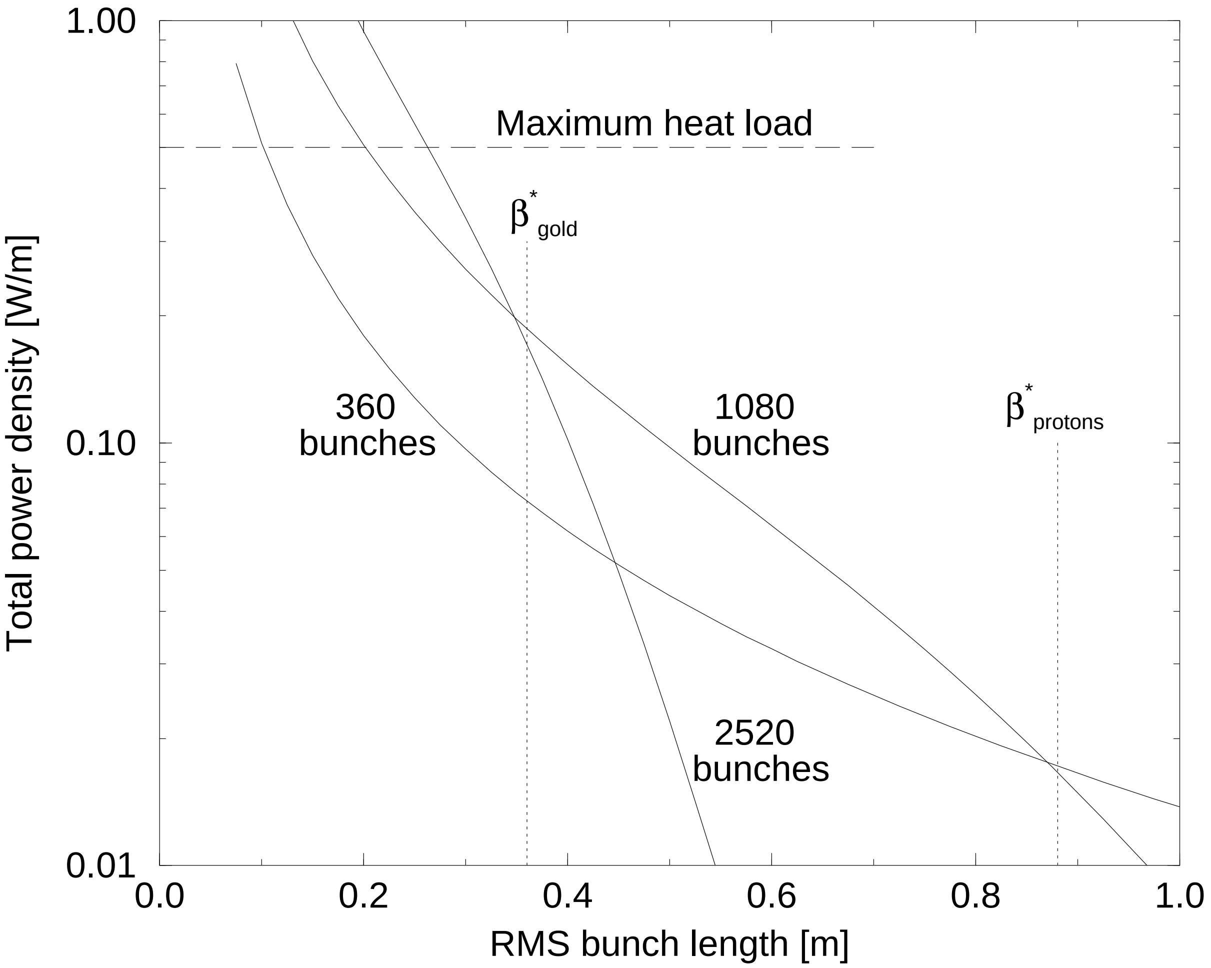

Figure 11: Linear power load deposited at cryogenic temperatures in the stainless steel vacuum chamber of the ion rings, due to beam image currents, for a total charge of 0.93×10^{11} per bunch – at the beam-beam limit for the electrons.

A careful analysis of the issue of "vacuum pipe heating in RHIC" lead to the engineering decision to use stainless steel beam pipes without a copper coating [27]. Figure 11 shows the results of extending that analysis to the parameter range being considered for eRHIC, as recorded in Tables 2 and 3. The linear power load depends strongly on the RMS Gaussian bunch length, and on the number of bunches.

The calculation used to generate the data in Figure 11 naively assumes that all the ion bunches have exactly the same charge, and are spread uniformly around the circumference (without an abort gap). In this case the power spectrum is a series of narrow lines uniformly spaced by $N_b F_{rev}$, under a Gaussian envelope which is the Fourier transform of the bunch shape. The total linear power load is just a sum over all these spectral lines, convoluted with the vacuum chamber resistance at those frequencies – a resistance which is dominated by skin depth effects. As the number of bunches increases, the spacing between spectral lines increases like N_b, but the power in each harmonic increases like N_b^2. Thus *when the bunches are longitudinally spaced by very many bunch lengths* – for example, when $N_b = 360$ – the linear power load is just proportional to N_b, as is intuitively expected.

Figure 11 shows that this scaling breaks down when there are 2520 bunches in an ion ring, and the bunch spacing is only 1.52 meters, except for very short bunch lengths less than, say, 0.25 meters. The suppression of the linear power load which is implied for longer bunch lengths is weakened in more realistic situations – for example, when an abort gap is present and when the bunch populations are not all equal. Nonetheless, it appears possible to store as many as 2520 bunches in the ion rings without violating the maximum heat load limit, and without losing much luminosity to the hourglass effect.

There is also the possibility that the Beam Position Monitor signal cables may suffer unacceptably large head loads, due to resistive heating by the signal current, when the number of bunches becomes large and the bunches are too short [28]. This problem appears to be less serious than the beam pipe heating problem. It is possible to replace all the signal cables with an upgraded design which is more resilient. In an extreme case the inner diameter of the stainless steel outer conductor of the coaxial line would be copper plated, and the dielectric would be made from silicon dioxide.

2.5.6 Electron cloud effect

The electron cloud effect also threatens to overcome the maximum permissible cryogenic heat load in RHIC [29]. The electrons which are produced by ionization of the residual gas when one bunch passes are attracted and accelerated by the electrical field of the next passing bunch. These accelerated electrons eventually hit the vacuum chamber wall and cause the emission of secondary electrons. This process can lead to the runaway accumulation of electrons in the beam pipe, driving a large cryogenic heat load, if the bunches are spaced too closely together.

The effect has been much studied for the LHC, where the nominal bunch spacing is 25 ns and there are nominally about 10^{11} protons per bunch [30, 31, 32]. The recent consensus at CERN appears to be that the electron cloud effect in the LHC is not so dangerous as previously feared [33].

Unfortunately there is a paucity of hard experimental data from existing cryogenic accelerators with closely spaced bunches, although the normal conducting SPS is beginning to generate interesting data with LHC bunch loading parameters [33]. It is noteworthy that in fixed target mode the Tevatron routinely operates with 1008 bunches of approximately 2×10^{10} protons, spaced by about 18.9 ns (53 MHz), without undue cryogenic difficulty. There are 3 missing bunches at 12 locations in the Tevatron bunch train, and an abort gap of about 1 μs. Such gaps act to clear the electron clouds.

With $N_b = 360$ the bunches are spaced by 35.6 ns, and with $N_b = 2520$ by 5.1 ns. It is possible that the electron cloud effect will place a hard limit on the number of ion bunches that can be stored in RHIC. This problem needs more investigation, especially in making careful measurements on RHIC, in other cryogenic storage rings – HERA and the Tevatron – and in the SPS.

3 Linac-Ring Collisions

3.1 A 10 GeV linac-ring collider

The generation of electrons for eRHIC starts with a linac in both major scenarios. The natural advantages of a linac are in the ability to vary the energy over a wide range without affecting the polarization or beam quality, and in its high beam quality (low emittance and energy spread). It is easy to build a linac with energy of 10 or 15 GeV and to provide polarized electrons at up to 80% polarization. A linac injecting into a storage ring can be constructed with copper cavities, say at a SLAC linac frequency of 2.856 GHz (where cavities and RF sources are readily available). Alternately the linac may be constructed of niobium superconducting cavities, for example 1.3 GHz TESLA cavities that may be operated at over 20 MV/m. Refrigerating power for a superconducting linac can be provided by the RHIC liquid helium refrigerator. A full energy linac for eRHIC would be rather short, about 500 meters long. There are no special issues or difficulties with an injector linac except for the polarized electron gun, which is discussed below.

In a storage ring, electrons are injected and stored in the ring for long times, with the RF system making up for the synchrotron radiation losses. The storage of particles in a ring leads to various potential instabilities and resonances, but the advantage is that the high power (many megawatts in average power) of the beam is not lost, by reusing the beam over many turns.

An alternative to using the same electrons over and over in a storage ring is the option to use them only through one turn, and then to decelerate them in the same linac and thus recover their energy. This scheme of a linac which recovers the energy of the beam is called a recuperator. The idea is not new, and there are a few successful examples of such machines, albeit at low electron energies. However, plans have been made to build recuperating linacs and microtrons for multi-GeV energies.

The idea of colliding a linac beam with a storage ring beam as a way of increasing the luminosity of the collider is also not new [34, 35].

Why is a linac attractive for a collider?

- A linac avoids the limitation of beam-beam tune shift inherent in a storage ring. This allows one to reduce the beam size in the ion storage ring, and to increase the charge per bunch considerably.
- The short bunch length of the linac reduces the uncertainty in the location of the interaction by a factor of 2 and simplifies the detector design. The luminosity may be further increased by electron cooling of the RHIC ion beams.
- A linac has a very low emittance. This leads to a small collision point beam size with a relatively large beta function, increasing luminosity and simplifying the optics of the interaction point.
- The fact that the electrons are used only once means that complicated spin rotation conditions are relaxed. Thus a linac-based collider can provide a polarized electron beam at any energy, while a storage ring is limited to a narrow energy range.
- The interaction point optics of the linac are much simpler (since the polarization may be prepared well in advance), thus the bending radii are larger. This significantly ameliorates the synchrotron radiation linear power load restriction to the maximum beam current.
- A linac can operate over a wide energy range without sacrificing performance.
- The polarization of a linac is high and can be alternated rapidly at will.
- A linac produces a naturally round beam, to match well with the RHIC ion beam.
- A linac can be extended to higher energies with a cost that is linear in length whereas a storage ring faces an increase in RF power that goes with the fourth power of its energy.

The beam power of a collider linac would have to be commensurate with the current of a storage ring. This implies that the linac, like a storage ring, would have to recover the beam energy (recuperation) to avoid excessive power budgets. The performance of the Jefferson National Accelerator Facility FEL with energy recovery provides a benchmark for the performance of an energy recovering linac. The energy recovery in that machine is so good that it is difficult to measure. A recent estimate [36] places a lower limit on the recovery efficiency of 99.98%.

3.2 Design issues of the high current, energy recuperating linac

For the purpose of this discussion, we assume that the linac energy is 10 GeV and the current is 0.135 A. The issues to deal with are as follows:

1. Linac layout and optics.
2. Polarized electron source.
3. Recuperation issues and beam instabilities.

In all of these issues the expertise and generous help of many persons at JLAB have been instrumental.

3.2.1 The linac layout and optics

David Douglas has studied the optics of a 10 GeV linac [37]. To avoid collisions of the accelerated and decelerated beam the two beams must propagate in the same direction of the linac. That means that at each end of the linac the transverse optics must deal with the full energy difference, from the injection energy to the final energy. For stability purposes the energy difference should be no more than a factor of ten. On the other hand, the energy recuperation should start at a low energy (see issue 3 below). Thus one arrives at the scheme shown in Figure 3.

As seen in the Figures 3 and 4, all or most of the acceleration is done in a straight line, to avoid emittance growth and synchrotron radiation loss in the accelerated beam. The recuperating beam is bent to return to the appropriate sections. The electron source has an injector linac that accelerates the beam to 10 MeV. The power invested (at 0.135 A) for this purpose is 1.35 MW. This section has no recuperation. Next is a low gradient 90 MeV (energy gain) pre-accelerator. Here energy recuperation is done in a dedicated linac section, and the recovered energy is fed through waveguides to the accelerating section, shown schematically in Figure 3 as a connection between the accelerating and decelerating linacs. The 100 MeV beam from the pre-accelerator is fed into a 0.9 GeV (energy gain) intermediate linac with energy recovery done in the same section. Last is the main linac, with an energy gain of 9 GeV. The 10 GeV beam is taken to the collision point. Figure 3 does not show these details, since the beam may be introduced into a ring-like transport for multiple interaction points, a single IP, or anything in-between.

The beam is returned at 10 GeV to the entrance of the main linac for deceleration and energy recovery. The beam is decelerated to 1 GeV, and then sent to the intermediate energy linac for deceleration to 100 MeV. In recuperating the energy of the beam in the same linac structure we observe the Douglas principle of keeping the energy ratio of the two beams under 10. Deceleration to 10 MeV follows in a dedicated 90 MeV pre-dump linac.

3.2.2 Beam dump

The 10 MeV beam is sent to the beam dump, rated at 1.4 MW (a power level of 2 MW was demonstrated in SLAC beam dumps, therefore Figure 3 shows a 2 MW dump). It is possible to decelerate the beam to a lower energy should the beam dump rating be below 1.4 MW, but this RF power will not be used for acceleration. Note that any synchrotron radiation power loss (anywhere in the high energy transport) smaller than 1.35 MW will subtract from the power deposited in the beam dump. The beam dumping, done at 10 MeV or lower (lower energy corresponds to synchrotron radiation losses) and thus the activation of the beam dump is very minimal. An even lower beam energy at the dump and dump power can be obtained by "braking" the beam in a linac structure, by generating RF power that can be disposed off in dummy loads.

3.2.3 The polarized electron source

We follow designs that have been developed at SLAC and JLAB. Charles Sinclair gave a recent summary of the state-of-the-art in this electron gun technology [38]. A strained gallium arsenide DC photo-cathode gun is used. The average current of 0.135 A requires a laser power of approximately 170 watt, which is delivered by a high power CW diode laser at a wavelength of 850 nm. The diodes can be RF pulsed to produce 10 ns bunches out of the photo-cathode. The chopped beam will be accelerated, bunched and further accelerated in special low frequency linac structures, which are part of the 10 MeV injector linac.

The photo-cathode quantum efficiency is expected to be of the order of 0.1% to 0.3%, with a charge half-life of $\sim 10^4$ coulombs/cm^2. Current experience is that the photo-cathode operational lifetime is almost completely dominated by ion back-bombardment. The operational lifetime is expected to be about 150 hours, a reasonable time between cathode changes. The surface area of the cathode is of the order of 7 cm^2, a size that has been built at SLAC [39].

The emittance one may expect from this source needs detailed studies in the parameter regime for this source. The thermal emittance for a 15 mm diameter GaAs cathode is of the order of 5 μm, before emittance growth due to space charge forces, transport, and other effects. Simulation and experimental measurements are reported by Engwall et al [40]. This work, at a cathode diameter of 2 mm and charge of under 100 pC, showed emittances of under 2 μm. We assume a conservative 60 μm for the normalized RMS emittance of the 5 nC bunches until more detailed calculations can be performed.

3.2.4 Recuperation issues and beam stability

A recent estimate of the energy recovery efficiency of the JLAB high power superconducting linac driving the JLAB high power FEL places a lower limit of 99.98% recovery [41]. At this efficiency the power loss from the 1.35 GW beam would be under 0.27 MW. While this is reasonable in terms of RF power, care has to be taken to capture this power at room temperature.

Next to beam loss, beam instabilities are the most serious issue. A recuperating linac is subject to a multi-bunch, multi-pass beam instability, similar to HOM instabilities that can happen in a storage ring. Indeed in a B-factory one has to apply feedback to avoid instabilities, and the same method can be applied in an energy recovery linac. It is instructive to estimate the Beam Break Up (BBU) threshold with no feedback. Lia Merminga ran the code TDBBU [40] for the most difficult section, a 1 GeV linac injected by a 100 MeV beam. The accelerating structures that were used in this numerical simulation were the TESLA cavities, with values for HOMs taken from TESLA data [42]. The result was a BBU threshold of 100 mA with incomplete optimization and no feed-back [43]. With B-Factory style feed-back a much higher BBU limit is expected. If necessary, the linac current may be reduced to under 100 mA with no other change, resulting in just a 35% loss in luminosity.

3.3 Electron cooling of RHIC

It is interesting to note that a high beam quality, high current superconducting linac with energy recovery is the machine of choice to provide electron beams for cooling a relativistic ion beam. The cooling beam energy ranges from about 50 MeV for gold beams at 100 GeV/u to about 130 MeV for 250 GeV protons. Cooling the RHIC ion beam is a worthwhile undertaking in its own right. Electron cooling can reduce the beam emittance and energy spread by orders of magnitude and keep the beam in this condition for as long as the cooling stays on. This will improve the luminosity and luminosity lifetime of the machine very significantly. On top of that, electron cooling potentially allows the collider to operate at larger than usual beam-beam tune shifts [44]. However, in the context of eRHIC with an electron linac, electron cooling of the RHIC ion beam takes on a special significance. Since the linac has an intrinsically low emittance beam, a significant gain in luminosity may be made by reducing the emittance of the ion beam. A large reduction of the ion emittance would be detrimental to an electron storage ring due to the beam-beam tune shift, but may be done when a linac provides the electron beam for the collisions.

Parameter	Units	**GOLD**	**PROTONS**
Cooler energy	[MeV]	54.5	136
Cooler current	[A]	1.5	30
Cooling time	[s]	46	5
Intra Beam Scattering time	[s]	112	768

Table 6: Key parameters for an electron cooler of length 30 m.

For eRHIC purposes, and until further studies of the problem are made, we assume that the beam cooling limit will be determined by the Laslett space charge tune shift, which after further modification by the cooling effect [44] is given by

$$\xi_{sc} = \frac{Z^2}{A}\frac{3N_i r_i}{2\pi\gamma_i^3\epsilon_i}\frac{C}{\sigma_z\sqrt{2\pi}} \tag{22}$$

where σ_z is the ion bunch length in meters. The Laslett space charge tune shift limit is estimated to be as large as 0.04 for gold (under the present parameters) and 0.08 for protons.

A collaboration has been established between BNL and the Budker Institute of Nuclear Physics to investigate electron cooling for RHIC. The initial calculations show that indeed such a scheme should work and the cooling rate for a gold beam in RHIC at high energy is measured in hundreds of seconds.

3.4 Parameters of the linac

Table 7 summarizes the parameters of the linac-ring collider. We take the case of 360 ion bunches, and assume electron cooling of either gold ions or protons. This results in smaller ion beam emittances and allows for larger tune shifts in the RHIC ion beam. Other than that, the RHIC parameters are mostly the same as for the ring-ring case, except that the proton bunch population is reduced to 3×10^{10} in order to keep the electron beam-beam tune shift under 1. In the case of gold the limit is set by the beam-beam tune shift for the ions. We do not take advantage of the large possible increase in the beam-beam tune shift due to the cooling. That may account for a further increase in the luminosity.

Parameter	Units	Value
Electron energy	[GeV]	10
Electron average current	[A]	0.135
Number of ion bunches, N_b		360
Electron bunch population, N_e		3×10^{10}
Electron RMS emittance, ϵ_e	[nm]	3.0
ELECTRON-GOLD		
Luminosity, L	[$\mathrm{cm}^{-2}\mathrm{s}^{-1}$]	5.6×10^{30}
Gold bunch population, N_i		1.9×10^{9}
Gold normalized emittance, ϵ_i	[$\pi\mu$m]	6
RMS beam size at the IP, σ^*	[mm]	0.06
Electron IP beta function, β_e^*	[m]	1.2
Electron beam-beam tune shift ξ_e		0.272
Ion beam-beam tune shift ξ_i		0.0046
Ion space charge tune shift, ξ_{sc}		0.015
ELECTRON-PROTON		
Luminosity, L	[$\mathrm{cm}^{-2}\mathrm{s}^{-1}$]	4.6×10^{32}
Proton bunch population, N_i		3×10^{10}
Proton normalized emittance, ϵ_i	[$\pi\mu$m]	0.8
RMS beam size at the IP, σ^*	[mm]	0.02
Electron IP beta function, β_e^*	[m]	0.147
Electron beam-beam tune shift ξ_e		0.688
Ion beam-beam tune shift ξ_i		0.0046
Ion space charge tune shift, ξ_{sc}		0.006

Table 7: Electron-gold and electron-proton collision parameters, assuming electron cooling of RHIC.

Although the parameters are quoted for an electron energy of 10 GeV, nonetheless the performance would be unaffected by a much lower (or higher) electron energy. Thermal loading is not a limitation for the relatively low electron beam current. In addition, larger radii of curvature in the IP optics are possible due to the removal of spin rotation optics from that area, further reducing the thermal loads.

The bottom line is that good luminosities can be obtained using a linac-ring collider with electron cooling of the ions in the ring. Further increases in the luminosity are possible. In the gold case, the increase would come from pushing the Laslett tune shift to higher values, taking advantage of the cooling. In the proton case the luminosity can be improved by going to a higher current in the electron linac by pushing on the beam-beam tune shift in the ring, once again taking advantage of the cooling.

3.5 Interaction region optics

The interaction region geometry and optics are simpler for linac-ring collisions than for ring-ring collisions, because spin rotator functionality is not required. Instead, the spin direction is easily manipulated close to the source.

Further, it is not necessary to have a small electron β_e^*. The goal is to match the electron beam size $\sigma^* = \sqrt{\epsilon_e \beta_e^*}$ to the ion beam size. At fixed beam size σ^* a storage ring with a relatively large emittance requires a relatively small β^*. The beam-beam effect on the electrons in a storage ring also places a limit on $N_i \beta_e^*$, as shown by Equation 3. Linac-ring collisions permit smaller emittances and larger β_e^* values. They also permit the electron beam to be "destroyed" by the beam-beam effect.

Two scenarios are considered. In the first, only one IP of six provides electron-ion collisions. On the side of the IP where the electron beam is naturally separated towards the radial outside of the ion beam, the electron beam line is taken (almost) tangentially out of the RHIC tunnel. On the other side of the IP the electrons are first bent horizontally around the triplet cryostat to go back across the ion beam lines, before departing tangentially from the tunnel. This beam line uses a set of 4 vertical dipoles to lower the electron beam by about 0.2 m at the ion beam lines. These 4 vertical magnets form an achromat if there are no quadrupoles in between.

The second scenario serves more than one IP. The beam is injected into IP12 as described in the previous paragraph, and as illustrated in Figure 4. Coming out of IP12 the beam goes through the other IPs. By choosing the correct energy the spin will be longitudinal in all interaction regions. Finally, the beam is taken out of the ring at IP2.

4 Summary

1. Electron-gold luminosities in the range in the range 6×10^{30} to 45×10^{30} $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, and polarized electron-proton luminosities in the range in the range 2×10^{32} to 15×10^{32} $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, appear possible.
2. The attainment of luminosities towards the upper end of these ranges depends on the ability to store as many as 2520 bunches in both rings at the beam-beam limit of single bunch intensities.
3. It is probable that further detailed studies and measurements would show that the optimum number of full intensity bunches is less than 2520, but more than 360, set by the point at which one of the rings becomes limited by the total beam current that it can support.
4. It is not clear which ring will reach its total beam current limit first.
5. The electron ring total beam current is probably limited to about 4 A by a synchrotron radiation linear power load limit of about 10 kW/m.
6. Spin rotators are optically challenging. They probably contain the strongest dipoles fields and the highest synchrotron radiation linear power loads, thereby limiting the electron beam current and the luminosity.
7. The ion ring total beam current is possibly limited by the electron cloud effect, and/or by the cryogenic beam pipe heating limit of about 1 W/m. Both of these effects need more study and measurement before a more definite statement can be made.
8. The polarization time at 10 GeV in the RHIC tunnel is about 10 hours, so the electrons must be injected pre-polarized at full energy, and may not be accelerated through depolarizing resonances.
9. A full energy linac or a polarizing booster can provide full energy pre-polarized electron bunches.
10. Permanent magnet technology may be appropriate for eRHIC.
11. Straw man optics have been constructed which meet all constraints except spin matching, which is not necessary in eRHIC.
12. More study is necessary, but eRHIC is feasible.

References

[1] *Polarized protons at high energies – accelerator challenges and physics opportunities*, Editors A. De Roeck, D. Barber, and G. Rädel, DESY-PROC-1999-03, DESY, 1999.

[2] *The first eRHIC workshop*, http://www.star.bnl.gov/afs/rhic/star/doc/www/collab/erhic/erhic.html , BNL, 1999.

[3] *Workshop on eA and polarized ep Collider Physics with eRHIC at BNL*, http://quark.phy.bnl.gov/r̃aju/yale_eRHIC.html , New Haven, 2000.

[4] *Machine studies of an electron-nucleus collider option*, P. Paul, memorandum to S. Peggs, July 15, 1999.

[5] *RHIC Design Manual* http://www.agsrhichome.bnl.gov/NT-share/rhicdm , BNL, 1993.

[6] *Luminosity scaling of electron-gold collisions in the RHIC rings*, S. Peggs, D. Trbojevic, RHIC/AP/168, BNL, 1999.

[7] *10 GeV electrons on 100 GeV/u gold ions at RHIC, with a third ring*, S. Peggs, unpublished note, May 21, 1999.

[8] *Luminosity upgrade of HERA*, M. Seidel et al, PAC99 p. 34, New York, 1999.

[9] *The Fish License Sketch*, M. Python, London, 1972:
Is this wretched demi-bee,
half asleep upon my knee,
some freak from a menagerie?
No! It's Eric the half a bee.
or,
No! It's [eRHIC] the half a [B-factory].

[10] J. Kewisch, owl.pbn.bnl.gov:/rap/lattices/eRHIC/*, BNL, 2000. Files "bv.mad" and "no.mad" describe the straw man optics with and without spin rotators.

[11] *Strong spin matching with and without snakes, a scheme for preserving polarization in large ring accelerators*, K. Steffen, DESY 88-068, DESY, 1988.

[12] *Polarization of a stored electron beam*, A. Chao, SLAC-PUB-2781, SLAC, 1981; A. Chao, NIM 180 29, 1981; and AIP Proc. 87 p. 395, 1981.

[13] *Simulation of electron spin depolarization with the computer code SITROS*, J. Kewisch, DESY 83-032, DESY, 1983; and J. Kewisch et al, PRL 62 No. 4, 419, 1989.

[14] *New practicable Siberian snakes schemes*, K. Steffen, DESY 83-058, DESY, 1983.

[15] *An Asymmetric B Factory Design Report*, SLAC 372, SLAC, 1991.

[16] *Commissioning results of the KEKB and PEP-II B-factories*, J. Seeman, PAC99 p. 1, New York, 1999.

[17] J.B. Murphy, private communication, September 1999.

[18] *The physics of electron storage rings*, M. Sands, SLAC-121, SLAC, 1970.

[19] N. Kurita, private communication, February 2000.

[20] S. Heifets, Presentation at SLAC Impedance Workshop, March 2000.

[21] M. Brennan, private communication, February 2000.

[22] *Conceptual design of the RHIC dump core*, A.J. Stevens, AD/RHIC/RD-94, BNL, 1995.

[23] *Energy deposition downstream of the internal dump*, A.J. Stevens, AD/RHIC/RD-97, BNL, 1995.

[24] A.J. Stevens, private communication, February 2000.

[25] *Parasitic beam-beam collisions and crossing angles in RHIC*, S. Peggs, RHIC/AP/66, BNL, 1995.

[26] *Beam-beam collisions and crossing angles in RHIC*, S. Peggs, Proc. LHC Beam-Beam Workshop, CERN, April 1999; and RHIC/AP/169, BNL, 1999.

[27] *Vacuum Pipe Heating in RHIC*, A.G. Ruggiero, S. Peggs, RHIC/AP/46, BNL, 1994.

[28] , *Thermal Behavior of RHIC BPM Cryogenic Signal Cables*, P. Cameron, M. Morvillo, RHIC/AP/68, BNL, 1996.

[29] *Beam induced electron clouds at RHIC*, K.A. Drees, RHIC/AP/150, BNL, 1998.

[30] *A simulation study of electron-cloud instability and beam-induced multipacting in the LHC*, F. Zimmerman, LHC Project Report 95, CERN, 1997.

[31] *Beam induced multipacting*, O. Gröbner, PAC97, Vancouver, 1997.

[32] *Simulations for the beam-induced electron cloud in the LHC liner*, O. Brüning, LHC Project Report 102, CERN, 1997.

[33] *LHC electron cloud key parameters*, L. Evans, presentation to the LHC Machine Advisory Committee, unpublished, CERN, 1999.

[34] P. Grosse-Wiesmann, SLAC-PUB-4545, SLAC, 1988.

[35] J. Bisognano et al, LINAC88, p. 586, CEBAF, 1988.

[36] G. Krafft and L. Merminga, private communication.

[37] David Douglas, JLAB-TN-98-040, JLAB, 1998.

[38] C.K. Sinclair, PAC99 p.65, New York, 1999.

[39] C.K. Sinclair, private communication, 1999.

[40] D. Engwall et al, PAC97 p. 2693, Vancouver, 1997.

[41] G. Krafft and L. Merminga, private communication, 2000.

[42] J. Sekutowicz, TESLA 94-07, DESY, 1994.

[43] L. Merminga, D. Douglas, JLAB-TN-00-004, JLAB, 2000.

[44] V. Parkhomchuk, private communication, 2000.