Recirculating Acceleration

J. Scott Berg[*], Alper A. Garren[†], Carol Johnstone[‡]

[*]*Brookhaven National Laboratory; Building 901A; PO Box 5000; Upton, NY 11973-5000*
[†]*University of California Los Angeles, Los Angeles, CA 90095*
[‡]*Fermi National Accelerator Laboratory; PO Box 500; Batavia, IL 60510*

Abstract. This paper compares various types of recirculating accelerators, outlining the advantages and disadvantages of various approaches. The accelerators are characterized according to the types of arcs they use: whether there is a single arc for the entire recirculator or there are multiple arcs, and whether the arc(s) are isochronous or non-isochronous.

INTRODUCTION

A recirculating accelerator (RA henceforth) attempts to make efficient use of RF linacs by passing the beam through the linacs multiple times. The RA is in a racetrack shape (or a shape with more sides), where the straight sides of the racetrack are linacs, and the linacs are connected by one or more arcs.

A ramping synchrotron can be thought of as a RA; there is only a single arc connecting the "linacs" (usually a small number of cavities), and the magnetic fields in the arcs are increased as the momentum of the beam increases. This scheme is impractical for the early stages of a muon collider, since the acceleration must be very rapid, and the magnetic fields in the magnets could not be increased sufficiently rapidly to match the rate at which one would hope to accelerate. The obvious solution to this difficulty, which is used as CEBAF, is to use a separate arc for each pass that the beam makes through the RA. The beam is sent into the appropriate arc either based on its energy by using a bending magnet, or by using a kicker magnet. Another possible scheme is to have a single arc with an extremely large energy acceptance (a fixed field alternating gradient, or FFAG, lattice).

Once the type of arc to be used has been decided, one can decide what type of longitudinal motion one would like and design the arcs accordingly. One may wish to have synchrotron motion, in which case the arcs should be non-isochronous, or one may wish to accelerate on crest, in which case one would like the arcs to be nearly (but not exactly) isochronous.

The primary goals when designing a RA will be to minimize losses (both due to decays and dynamic losses) and emittance growth. Decay losses (and most likely dynamic losses and emittance growth as well) would of course be minimized by

using a single linac accelerating to the desired maximum energy. From the point of view of cost, however, this is not ideal; this is the motivation for using RAs. To keep costs down, one wishes to minimize power requirements for the RF systems, to avoid low frequency RF cavities, and to keep average power efficiencies high, particularly for higher energy machines. At the same time, one wishes to minimize dynamics losses and decays in the machine; this leads to wanting higher RF gradients (and therefore frequencies), and wanting to keep arcs as short as possible. However, in designing the arcs, one often finds that the required energy acceptance of the arcs is very high, the apertures are large, and the required momentum compaction is either very large or has very tight tolerances. These are some of the concerns which will be used to analyze the different types of RA designs.

SINGLE ARC VERSUS MULTIPLE ARCS

The minimum and maximum energies going through a RA may differ by a factor of 4 or more. It is therefore natural to use a separate arc for each turn in cases where it is impractical to ramp the magnets linearly with the reference momentum. For lower energy RAs, since the muons must be accelerated as rapidly as possible to avoid decays, magnets cannot be ramped sufficiently rapidly due to the short ring circumference and rapid energy increase in the machine.

However, one quickly discovers that the energy acceptances in the arcs of a RA need to be very large, particularly if one wishes to use higher frequency RF to reduce costs. Thus, one begins to consider designing an arc that not only accepts the large energy spread required for a beam on a single pass through the RA, but also has sufficient energy acceptance to accept the beam energy for several (even all) passes through the RLA. Machines have been built which can accept extremely large energy ranges with fixed magnetic fields (FFAGs) [1]. The advantages of using a single arc over multiple arcs are that the energy spread in the beam is not a concern (allowing one to go to higher frequencies), and one can have an arbitrary number of turns without the arc complexity changing or the arc cost increasing.

One encounters a great number of difficulties when using a single arc. For non-isochronous arcs, the bunches will not arrive at the correct RF phase if something is not done to correct that RF phase. Correcting the RF phase requires either a large amount of RF peak power to be supplied, some device (ferrite or otherwise) be used to change the frequency of the RF cavities (which most likely absorbs energy), or fast ramping magnets to somehow supply a path length correction (which is probably prohibitive, at least for the low energy stages). Even an isochronous lattice has difficulties, since the lattices must have a small but non-zero momentum compaction which is different from the tendency of an FFAG lattice. In addition, the magnets for a single arc design tend to be very large and complex, and it is therefore not clear whether there is a cost savings of a single arc over multiple arcs. Next, the straight sections containing the linacs necessarily do not fit into

an FFAG scheme, and one therefore has a transverse mismatch in these sections. Finally, a single arc will be longer than even the highest energy arc in a multiple arc design, and therefore there will be more decay losses.

There are different types of FFAG lattices, scaling [1] and non-scaling, and they present different design difficulties. The scaling design is the classic type. It is probably easiest to achieve a large energy acceptance with this type of lattice, but that is not certain. Its primary disadvantage is that it tends to be long compared to a non-scaling design. In addition, when one puts in dispersion-free insertions which try not to disturb the lattice (needed for the linacs), those insertions tend to be short, whereas one would prefer a longer section for the linacs. In principle this can be overcome by using many insertions. Finally, a scaling lattice tends to have constant momentum compaction, whereas one would prefer a momentum compaction increasing roughly linearly with energy to keep the beam matched. The momentum compaction can in principle be adjusted through wigglers in insertion sections, but to what extent its behavior can be modified is unclear.

A non-scaling FFAG lattice in principle overcomes some of these difficulties, but with problems of its own. The natural tendency of a non-scaling lattice is to have zero momentum compaction somewhere in its energy range. To correct this, strong sextupole magnets are required. Furthermore, when it is used in an isochronous lattice, the natural change in the momentum compaction has the wrong sign for what is needed in the lattice.

In a multiple arc design, the major issues to contend with are the energy acceptance of the arc and the complexity and size of the switchyard. Due to the high cost of RF systems and the fact that these costs decrease with increasing frequency, it is generally desirable to use as high a frequency as possible; however, as one increases the frequency, the energy spread in the beam increases also. The energy spreads required for existing muon collider and neutrino factory parameters end up being relatively large, and thus the energy acceptance of the arcs becomes an issue. These energy spreads become an issue for the switchyard as well, limiting the number of turns one can use in a RA, as well as making the magnets in such a switchyard very large. These problems can be ameliorated somewhat by using a kicker in the switchyard, but the required rise time and aperture of such a kicker in lower-energy RAs makes them seem prohibitive.

ISOCHRONOUS VERSUS NON-ISOCHRONOUS

Design considerations for a non-isochronous RA based on longitudinal phase space considerations are described in [2]. In summary, one can choose to run with a matched longitudinal bucket with a synchrotron tune of 0.15 (or possibly less) per half turn and a bucket which just fits around the bunch. This tune is about the maximum value one can use, and it minimizes the energy spread while giving maximal reduction of negative effects from collective effects (instabilities, beam loading

effects due to current offsets and bunch trains). It also minimizes any depolarization of the muons [3]. The one problem with this large synchrotron tune is that the required momentum compaction can become large, particularly at lower frequencies. However, assuming small RF phase angles and small synchrotron tunes, the RMS energy spread and momentum compaction are respectively approximated by

$$\sigma_\Delta \approx \frac{3^{1/6}}{(4\pi)^{1/3}} \left(\frac{\Delta E \omega^2 \epsilon_L^2 k}{n \nu_s} \right)^{1/3} \qquad \alpha_C \approx \frac{\beta^2 p c^2}{L_{\text{arc}}} \left(\frac{128\pi^5}{3} \frac{\nu_s^5 n^2}{\omega^4 (\Delta E)^2 \epsilon_L k^2} \right)^{1/3}. \qquad (1)$$

Here ΔE is the total energy gained in the recirculator, n is the number of turns, and other parameters are as defined in [2]. By reducing the synchrotron tune, a large reduction in the momentum compaction can be accomplished with a much smaller increase in the energy spread. While reducing the synchrotron tune does decrease the beneficial effects of synchrotron motion, the effects that the synchrotron motion is trying to correct are reduced at lower frequencies, where the large α_C was a problem in the first place. Finally, note that increasing the number of turns decreases the energy spread, but increases the required momentum compaction.

While with synchrotron oscillations, one can in principle inject a "matched" beam and avoid any emittance growth, one will necessarily have emittance growth with isochronous operation. Thus, one puts bounds on the emittance growth. The relative longitudinal emittance growth in an isochronous RA is approximately

$$\frac{\delta \epsilon_L}{\epsilon_L} = \left(\frac{\Delta E}{2 \epsilon_L \omega} \right)^2 [\omega^2 (\sigma_\tau^2 + \mathcal{P}^2 \sigma_\Delta^2)]^3 \qquad \mathcal{P} = \sum_{k=0}^{n-1} \frac{n-k}{n} \frac{L_k \Delta \alpha_k}{\beta_k^2 p_k c^2}, \qquad (2)$$

where again ΔE is the total energy gained in the RA, k is an index referring to each individual arc of length L_k, reference momentum p_k and corresponding reference velocity $\beta_k c$, with an error in the momentum compaction of $\Delta \alpha_k$. The first arc is $k = 0$ and there are n arcs. Notice the extremely strong dependence on the bunch length and frequency. In particular, the momentum compaction tolerances get tighter when you go to higher frequency or more turns, the very things one would like to do to improve RF efficiency and costs.

EXAMPLES

Table 1 shows some sample parameters for a non-isochronous RA accelerating from 4 GeV to 16 GeV. The systems are compared at several frequencies. For the isochronous lattice, a 5% emittance blowup was allowed from the bunch length alone, and an additional 5% from a momentum compaction error. The table shows the momentum compaction for the highest energy arc for the non-isochronous RA and the momentum compaction tolerance for the lowest energy arc.

The energy spread for the isochronous RA is substantially larger than the energy spread for an equivalent non-isochronous RA. This not only potentially allows for

4

TABLE 1. Example parameters for 4–16 GeV RA.

f	n	Non-isochronous					Isochronous				
		L_{lin}	σ_τ	σ_Δ	loss	α_C	L_{lin}	σ_τ	σ_Δ	loss	$\Delta\alpha_C$
MHz		m	ps	MeV	%		m	ps	MeV	%	10^{-5}
100	4						182	111	108	3.9	82
100	7						104	111	108	4.8	40
100	11						66	111	108	5.9	23
200	4	145	121	112	3.4	0.051	129	70	171	3.1	33
200	7	88	144	94	4.3	0.076	74	70	171	4.0	16
200	17	41	188	72	7.5	0.133					
400	4	110	75	180	2.8	0.020	91	44	272	2.5	13
400	8	62	92	146	4.2	0.031					
800	4	87	46	292	2.5	0.008					

more turns in a non-isochronous RA, but it also means that the arcs will have fewer problems due to the smaller required energy acceptances. In addition, one can in principle go to higher frequency RF with a non-isochronous design. One only gains slightly in the length of the linac for an isochronous design, since when one is near the crest, going off crest by a small amount has little effect on the gradient.

CONCLUSIONS AND ACKNOWLEDGEMENTS

We have outlined some of the issues related to the design of recirculating accelerators, in particular considering issues related to arcs. If one uses a multiple arc design, it seems that there is a clear preference for a non-isochronous design, primarily because of the smaller energy acceptance required. However, if one uses a single arc design, and isochronous lattice is probably preferable, since it avoids issues related to keeping the particles at the correct phase. Also, in a single arc design the energy spread is not of so much concern.

REFERENCES

[1] K. R. Symon *et al.*, "Fixed-Field Alternating-Gradient Particle Accelerators," Phys. Rev. **103**(6), 1837–1859 (1956).
[2] J. Scott Berg, "Acceleration for a High Energy Muon Collider," to appear in the proceedings of *Studies on Colliders and Collider Physics at the Highest Energies: Muon Colliders at 10 TeV to 100 TeV—HEMC '99*, Montauk, NY, 27 September–1 October 1999. Also BNL Report BNL-67115.
[3] Alain Blondel, "Muon Polarization in the Neutrino Factory," to appear in the proceedings of *Neutrino Factories based on Muon Storage Rings*, Lyon, France, 5–9 July 1999.