1.1 FFAGs for Muon Acceleration

J. Scott Berg, Brookhaven National Laboratory, jsberg@bnl.gov

Muons present unique challenges for acceleration. First of all, muons decay: this means that any acceleration must be rapid. A practical minimum is around 1 MV/m. This rules out traditional synchrotron designs, which contain a relatively small amount of RF per turn. One could have more RF per turn in a synchrotron, but then one runs into the challenge of ramping the magnetic fields sufficiently rapidly [1]. One could accelerate very rapidly in a linac, but that becomes very expensive: one would like to pass through the accelerating structures many times to reduce RF costs.

To be able to pass through the same linac many times without having to ramp magnets, previous studies have used a multiple-arc recirculating accelerator [2—3]. These too have their challenges. One must design a switchyard that puts the beam into a different arc corresponding to the beam's energy on that pass. The number of passes through the switchyard is limited by the fact that the highest energy in the beam from one pass must not exceed the lowest energy in the beam at the next pass. With the relatively large energy spreads in our beam, this is a nontrivial restriction. The finite transverse beam size and the necessity of extra space for the magnet coils, cryostats, and other hardware makes this even more restrictive. Furthermore, one has some emittance growth at each matching section from a linac to an arc. Finally, at some point the system cost begins to increase as one makes more turns, because the cost of additional arcs exceeds the cost reduction due to the reduced RF requirement.

An FFAG (**F**ixed **F**ield **A**lternating **G**radient) accelerator is one way to try to address some of these concerns. The idea behind an FFAG is to create a strong-focusing arc lattice with an extremely large energy acceptance (a factor of 2 to 3 in this article). If one can accomplish this, then one avoids the difficulties of the multiple-arc recirculating accelerator. There is no need for a switchyard, so energy overlap from one turn to the next is irrelevant. The matching problem can be avoided if the RF is distributed around the ring, and thus the beam sees an adiabatically changing lattice. Reducing the RF by going to more turns does not require than an additional arc be built. One can therefore hope that an FFAG can provide a cost-effective alternative for accelerating muons.

This article describes the current understanding of how FFAGs can be used for muon acceleration. It presents the work of many individuals and groups which will be listed in the references. In particular, there was recently a workshop on FFAGs for Muon Acceleration at Berkeley, from October 28 through November 8, 2002 (see http://www.cap.bnl.gov/mumu/conf/ffag-021028/), where much progress was made in this area. This article will focus on acceleration to 20 GeV, which was the primary focus of the workshop. Many of the considerations are the same for higher or lower energies, but the dominant problems often change as the energy regime changes.

1.1.1 Technical Constraints

The design of an arc that transports a factor of 2 or 3 in energy is a challenge. First one must deal with the issue of avoiding linear resonances over the entire energy range. One is making relatively few passes through the ring (typically 20 or fewer), so having a tune for the entire ring near a half integer or an integer is not of great concern, since such resonances are generally weak over that time scale. However, one must be concerned with linear resonances over the scale of a single cell, which will lead to catastrophic beam loss. There are three ways that are used to avoid linear resonances over such a large energy range:

1. Make a "scaling FFAG," which has a constant tune over its entire energy range [4]. This is what people have historically referred to as an FFAG.
2. Keep the tune for a single cell below a half integer at the lowest energy, and the tune will then decrease as the energy increases [5].
3. Add sextupoles to set the chromaticity to zero in both planes, minimizing the tune variation over the energy range [6].

Once this basic constraint for transverse dynamics has been satisfied, one considers longitudinal dynamics. A lattice with such a large energy range naturally has a time-of-flight that varies significantly with energy over that energy range. While it is possible to vary the RF frequency as the beam energy increases so as to match the time-of-flight variation, this requires a great deal of RF peak power, and is therefore generally impractical. If one does not vary the RF frequency, there is a minimum amount of RF voltage required to accelerate over a given energy range, irrespective of how many turns over which the beam is accelerated [7]. For a given type of lattice design, this minimum voltage is proportional to the total range in the time-of-flight over the energy range of the accelerator. Thus, achieving cost reductions by reducing RF voltage requires that the time-of-flight variation over the energy range of the accelerator be kept as low as possible.

1.1.1.1 RF Cavities

In general, if the lattice cell is kept shorter, the time-of-flight variation will be kept lower. One of the dominant factors determining the length of the cell is the length of the straight section required for RF cavities. Some of the scenarios described here use 200 MHz cavities, where a single cell would require a length of around 1 m for the cell plus associated hardware. One could shorten this length by using higher frequency RF, but there are two likely problems with doing so. First of all, the minimum required voltage described above is proportional to the RF frequency; thus, it is likely to be more costly to use higher frequency RF, despite the reduction in cell length. Second, because the beam passes through the cavities several times, and it would be prohibitive to put power into the cavities as fast as the beam is extracting it, beam loading may become problematic at higher frequencies, since there is less energy stored in the cavity at higher frequency.

Comparing superconducting RF to normal conducting, superconducting RF appears to be must more cost effective due to its substantially reduced peak power requirement. However, there is a disadvantage to superconducting RF: there must be a substantial separation between the magnets and the RF cavities to shield the cavities from the magnetic field. Previous work has suggested a separation of around 1 m is required [3], leading to a drift length of 3 m. This will substantially increase the time-of-flight range from what it would have been if the drift were only 1 m. A preliminary study by Shlomo Caspi [8] indicated that this magnet-cavity distance could be shortened to 0.5 m. It is clear that further study of how to shield magnets from RF cavities, and in particular obtaining scaling laws for the relationship between shielding distance and magnet type, field, and aperture, is required.

One very interesting solution to this problem is to shield the cavities from the magnets only to around 0.1 T, instead of the approximately 10 mG that the above distances correspond to. One would cool the cavities down to cryogenic temperatures with the magnets off, and then power the magnets. The fields from the magnets would be excluded by the superconducting cavity surface [9]. The main disadvantage to this mode of operation is that if a cavity does quench, the cavity temperatures must be raised, all the surrounding magnets must be powered down, the cavities cooled, and then the magnets powered again; this can take a very long time, potentially having an enormous negative impact on machine availability.

1.1.2 Specific Machine Types

I will classify the machine types based on the lattice cell on which they are constructed. It turns out that the different lattice cells described here all deal with the issue of avoiding linear resonances with a different one of the methods listed above. In addition, they have various advantages and disadvantages which will be described as each lattice is described.

1.1.2.1 Low Emittance Lattice [6]

This lattice is based on a lattice cell that would give a low emittance for an electron ring. Both the dispersion function and the beta functions are small at the bending magnet. The primary appeal of this lattice is that both the closed orbit variation and the variation in the time-of-flight with energy are extremely small, much smaller than in other lattices presented here. The ring is also very short compared to other lattices. This should lead to an extremely inexpensive ring, due to low magnet counts, low magnet apertures, and low RF voltage requirements due to the ability to accelerate for a large number of turns.

Unfortunately, these exemplary characteristics come at a severe cost: the dynamic aperture of this lattice is unacceptably small. The lattice gets its ability to operate over a large energy range (as well as its small range in time-of-flight) through the use of strong sextupoles to control chromaticity. Those sextupoles unfortunately reduce the dynamic aperture significantly. Work is progressing on improving the dynamic aperture without significantly compromising the performance of the lattice. Some significant progress has been made, but the dynamic aperture is still far from what is needed to transport a muon beam.

1.1.2.2 FODO Lattice [5]

A very simple approach to designing an FFAG lattice is to simply make a FODO lattice using gradient dipoles for the two quadrupoles, with drift spaces in between them for RF cavities and other hardware. Such a lattice turns out to be very linear, and therefore has an extremely large dynamic aperture. This lattice avoids the linear resonances by keeping the cell tune well below 0.5 over the entire energy range.

Designing such an FFAG lattice cell is fairly straightforward: fix the cell length and the total bend angle per cell, allow the bend fields and gradients in the magnets to vary, and fit the tunes at the lowest energy in the range to around 0.3 and the frequency slip factor at the central energy to zero. The latter condition arises because the variation of time-of-flight with energy is well approximated by a parabola, and placing the extremum of the parabola at the central energy minimizes the total height of the parabola over the full energy range.

This procedure allows one to design lattices for given cell lengths and bend angles per cell. Several general conclusions can be made from this procedure. First, the range in the time-of-flight is proportional to the cell length, which should be obvious from scaling considerations. Thus, the RF voltage required is proportional to the cell length as well: there is a clear advantage in reducing the drift length required for the RF cavities. Second, the total range in time-of-flight (per turn) is inversely proportional to the number of cells in the ring. Thus, a tradeoff between arc costs and linac costs occurs which is similar to that which one has for a multiple-arc recirculating accelerator: arc costs increase roughly in proportion to the number of cells, RF costs are inversely proportional.

An additional conclusion that one can draw is that for a given cell length and bend angle, the RF cost is proportional to the cube of the energy gain that one desires. One power is obviously because the voltage needed for a given number of turns is proportional to the energy gain desired. The other two powers are because of the parabolic time-of-flight variation with energy. As a result, it is not clear that fewer accelerating stages is better.

In all likelihood, the above considerations apply to most any FFAG lattice (except that the RF cost goes like the square of the energy gain for the scaling FFAG lattice to be described next). The only difficulty is in finding a method for designing lattices automatically with arbitrary parameters for the purposes of optimization.

If the arc cells did not need to contain RF cavities, they could be made very short, and thus the time-of-flight variation per cell would be small. If the bend angle per cell is very small, the time-of-flight variation is also small. One could try to combine these and get the best of both worlds: make straight (or nearly straight) sections that contain drifts for the RF, and arcs which contain no drifts, forming a racetrack (or oval) shape [10]. This turns out to be a very cost-effective solution. The one challenge is matching the dispersion and beta functions from one type of cell to the other over the large energy range of the accelerator. Recent attempts at doing this by Eberhard Keil [11] seem to be meeting with a great deal of success.

1.1.2.3 Scaling FFAG Lattice [12—13]

The so-called "scaling FFAG" is the original type of FFAG [4]. It is the only type of FFAG that has actually been built [14—16]. The tunes and the momentum compaction of the lattice are independent of energy. The closed orbits at different energies are simply geometrically scaled from one another. To achieve this, the magnets have fields that are proportional to r^k, where r is the distance to the center of the ring. As k increases, the gradient relative to the bending field increases, reducing the required magnet apertures and the momentum compaction. Unfortunately, as k increases, the nonlinearities in the field also increase, resulting in a decrease in the dynamic aperture. Thus, one generally wants the largest k (often several hundred) that will still give an acceptable dynamic aperture. The NufactJ Working Group in Japan has done an extensive design study [13] for a neutrino factory using a sequence of FFAGs for acceleration, and much study has occurred subsequently.

Compared to the previous designs, these accelerators require relatively low frequency RF (24 MHz, as opposed to 200 MHz in the non-scaling designs). The reason is related to the path length variation with energy: first of all, since the momentum compaction is constant, the path length is a monotonic (nearly linear) function of energy, as opposed to being parabolic as in the previous ("non-scaling") designs. This tends to lead to a larger total variation in time-of-flight (for a given maximum slope the parabola has a much smaller difference between maximum and minimum). In addition, the parabolic time-of-flight variation with energy allows the bunch to cross the crest three times [7] whereas it can only cross twice with the monotonic time-of-flight variation; thus, for a given time-of-flight range, the range of motion in time of a bunch is less with the parabolic variation. To accelerate, a stationary RF bucket is created which has a very large energy width, encompassing both the minimum and maximum energy. The bunch is accelerated by undergoing half of a synchrotron oscillation in this bucket [17]. It may be possible to reduce the RF requirements by having two RF systems which create two buckets, one to accelerate from the low energy to an intermediate energy, then a second to accelerate to the final energy. There has been some preliminary success with this scheme, but more work remains to be done.

Preliminary designs for superconducting magnets for the highest energy accelerator (10—20 GeV) have been made [18]. They use a cos θ style of design (with an elliptical vacuum chamber), but the coils are distributed highly asymmetrically to give the r^k field dependence. In addition, a trim coil has been introduced to allow the adjustment of k over a limited range.

1.1.3 Cost Estimation [10]

Robert Palmer has created a model for magnet and RF costs [10] and used it to estimate the cost of several FODO-based accelerators and the scaling FFAG accelerator. The results are summarized in the following table:

	Magnets (M$)	RF (M$)	Other (M$)	Total (M$)
FODO, 3 m Drift, 6-20 GeV	105	89	36	230
FODO, 1 m Drift, 6-20 GeV	45	117	19	181
FODO, Racetrack, 6-20 GeV	46	34	14	94
FODO, 3 m Drift, 10-20 GeV	19	37	17	73
Scaling, 10-20 GeV	89	89	25	203

The numbers should not be taken as absolute numbers, but should be taken relative to each other. The "Other" costs are for vacuum, diagnostics, and civil construction. These designs are not cost optimized but are optimized to have roughly the same decay (corresponding to an average accelerating gradient of around 1 MV/m).

The high cost of the FODO lattice with 3 m drifts comes from the larger magnet apertures (due to the longer cell length) and the large amount of RF needed (because the range of time-of-flights is relatively long). Shortening the drifts to 1 m decreases the costs substantially for the reasons just mentioned. The RF cost is increased since one is forced to use normal conducting RF, but the shorter cell length reduces the RF voltage requirement, so the additional cost of superconducting over normal conducting RF is partially compensated. The racetrack, as expected, gives the best of both worlds: the time-of-flight range is kept under control, allowing a small amount of RF voltage to be installed, the magnet apertures are kept reasonably small, and it is still possible to use superconducting RF. This lattice has not been analyzed completely self-consistently, so one should be careful in this comparison, but these results indicated that the racetrack design is likely to be optimal if the matching can be done properly.

Reducing the energy range from 6-20 GeV to 10-20 GeV indeed results in a substantial cost reduction for the accelerator. Note that in this example, cost reduction techniques applied to the 6-20 GeV FODO design have not even been applied, and so the optimal design will probably cost far less than this. Of course, one must add in the costs associated with making more stages, including in particular transfer lines between the accelerators, so it is not clear to what extent going to more stages will be beneficial.

Using these costs estimates, the scaling FFAG is coming out substantially more expensive. The magnets are more expensive because a substantially larger aperture is required in the defocusing quadrupoles (the non-scaling designs have a smaller orbit swing in the defocusing quadrupoles, whereas the scaling designs do not) and the ring is substantially longer (due at least partly to the smaller time-of-flight range in the non-scaling design). The RF costs are higher because of the low frequency and large voltage required, and the fact that the RF must be normal conducting.

The comparison of the scaling FFAG design to the FODO-based FFAG designs is not completely fair. The FODO-based designs have not been analyzed nearly as extensively as the scaling FFAG design. In fact, an examination of the dynamic aperture has not even been done on most of them (although based on tracking done for one of them, it is expected that their dynamic aperture will be high). The nature of the magnets in the two designs is likely to be very different. However, this comparison does indicate where improvements in the scaling FFAG design should be made in order to lower its cost.

1.1.4 Conclusions

FFAGs appear to be an effective way of reducing the cost of accelerating muons. A great deal of research is being done to verify and improve their performance and cost. Individuals are constantly coming up with new and better ideas for how to design these systems (in particular, Carol Johnstone has proposed using triplets instead of FODO cells, and initial results indicate that this improves the time-of-flight range [19]). This work has the potential to be used for many other types of acceleration applications as well, such as high-intensity proton sources.

1.1.5 References

[1] Don Summers, "Accelerating muons to 2400 GeV/c with dogbones followed by interleaved fast ramping iron and fixed superconducting magnets," in *Proc. of the*

APS/DPF/DPB Summer Study on the Future of Particle Physics (Snowmass 2001), ed. N. Graf, arXiv:hep-ex/208010.

[2] N. Holtkamp and D. Finley, eds., "A Feasibility Study of a Neutrino Source Based on a Muon Storage Ring," Fermilab-Pub-00/180-E (2000).

[3] S. Ozaki, R. Palmer, M. Zisman, and J. Gallardo, eds., "Feasibility Study-II of a Muon-Based Neutrino Source, BNL-52623 (2001).

[4] K. R. Symon *et al.*, "Fixed-Field Alternating-Gradient Particle Accelerators," Phys. Rev. **103**, 1837 (1956).

[5] C. Johnstone and S. Koscielniak, "Recent progress on FFAGs for rapid acceleration," in *Proc. of the APS/DPF/DPB Summer Study on the Future of Particle Physics (Snowmass 2001)*, ed. N. Graf, SNOWMASS-2001-T508.

[6] D. Trbojevic, M. Blaskiewicz, E. D. Courant, and A. Garren, "Fixed field alternating gradient lattice design without opposite bend," in *Proceedings of EPAC 2002, Paris, France* (EPS-IGA/CERN, Geneva, 2002), p. 1199.

[7] J. S. Berg, "Dynamics in Imperfectly-Isochronous FFAG Accelerators," in *Proceedings of EPAC 2002, Paris, France* (EPS-IGA/CERN, Geneva, 2002), p. 1124.

[8] Shlomo Caspi, presentation given at the Neutrino Factory and Muon Collider Collaboration workshop on FFAG Acceleration, Berkeley, CA, Oct. 28—Nov. 8, 2002.

[9] M. Ono *et al.*, "Magnetic field effects on superconducting cavity," in *9^th^ Workshop on RF Superconductivity 1999* (Los Alamos, NM, 2000), Los Alamos National Laboratory report LA-13782-C.

[10] Robert Palmer, presentation given at the Neutrino Factory and Muon Collider Collaboration workshop on FFAG Acceleration, Berkeley, CA, Oct. 28—Nov. 8, 2002.

[11] Eberhard Keil, unpublished note, http://keil.home.cern.ch/keil/MuMu/Doc/FFAG02/adiabatic.pdf (2002).

[12] Y. Mori, "Neutrino Factory in Japan: Based on FFAG Accelerators," in *Proceedings of EPAC 2002, Paris, France* (EPS-IGA/CERN, Geneva, 2002), p. 278.

[13] NufactJ Working Group, "A Feasibility Study of a Neutrino Factory in Japan," http://www-prism.kek.jp/nufactj/ (2001).

[14] F. T. Cole *et al.*, Rev. Sci. Instrum. **28**, 403 (1957).

[15] D. W. Kerst *et al.*, Rev. Sci. Instrum. **31**, 1076 (1960).

[16] M. Aiba *et al.*, "Development of a FFAG Proton Synchrotron," in *Proceedings of EPAC 2000, Vienna, Austria*, p. 581.

[17] T. Uesugi and C. Ohmori, presentations given at the Neutrino Factory and Muon Collider Collaboration workshop on FFAG Acceleration, Berkeley, CA, Oct. 28—Nov. 8, 2002.

[18] T. Ogitsu, presentation given at the Neutrino Factory and Muon Collider Collaboration workshop on FFAG Acceleration, Berkeley, CA, Oct. 28—Nov. 8, 2002.

[19] Shinji Machida and Carol Johnstone, work presented at the Neutrino Factory and Muon Collider Collaboration workshop on FFAG Acceleration, Berkeley, CA, Oct. 28—Nov. 8, 2002.