# Recent Progress on FFAGS for Rapid Acceleration

C. Johnstone*
*FNAL, Batavia, IL 60510, USA*†
S. Koscielniak‡
*TRIUMF, Vancouver, B.C., V6T 2A3 Canada*§

Muon acceleration is one of the more difficult stages to develop for a Neutrino Factory or Muon Collider. The large transverse and longitudinal admittances which must be designed into the system and the rapidity with which acceleration must take place because of muon decay preclude the use of conventional synchrotron design. The approach here employs fixed-field architectures for muon acceleration; specifically, a fixed-field alternating gradient or FFAG accelerator. This paper explores the FFAG option, in particular addressing an adjustment in the rf phase which, although characteristic of fixed-field machines, becomes problematic in the context of rapid acceleration.

## 1. Introduction

Because of potentially heavy losses from decay, acceleration must occur rapidly for any application requiring a high-energy, intense muon source; e.g. a Neutrino Factory[2][3] or a Muon Collider[1]. Linear acceleration is the most efficient, but above a GeV it becomes prohibitively expensive. Conventional synchrotrons cannot be used because normal conducting magnets cannot readily cycle in the ramping time required by muon decay. The current baseline approach employs recirculating linacs with separate, fixed-field arcs for each acceleration turn. Given the technical complexity, acceptance limitations, and expense of ultra-rapid cycling synchrotrons and recirculating linacs, the idea of using fixed-field, single-arc accelerators has been revisited in recent work and is the subject of this paper.

In addition to rapid acceleration, another overriding consideration for a Neutrino Factory is an acceleration system which has an exceptionally large acceptance, both transversely and longitudinally, and this acceptance impacts the degree of beam cooling required. Because of its naturally-large longitudinal acceptance, the FFAG is, therefore, an important option to consider. These studies of FFAG accelerators represent an effort to reduce cost and promote an acceptance which is better matched to the performance of the ionisation cooling system and the storage ring. Further, revisiting FFAG lattices in light of present superconducting technology and magnet design has advanced their reach into the multi-GeV regime, making a chain of FFAG accelerators a potential candidate for a complete acceleration scenario. This paper reports significant progress on both a lattice and rf acceleration system for a FFAG

## 2. FFAG Lattices

A circular accelerator system can be designed with fields that remain constant for the duration of the acceleration cycle using an alternating gradient focussing lattice. The arcs of such machines, composed of large- bore superconducting magnets, can be designed to accommodate the large energy range in acceleration. The closed orbits are not fixed as in a ramped machine, but rather move across the magnet aperture during the cycle. Lattices have been developed which can contain an energy change of a factor of four[4]. There are three basic types of alternating

---

*cjj@fnal.gov
†Work supported by the Universities Research Association, Inc., under contract DE-AC02-76CH03000
‡shane@triumf.ca
§TRIUMF receives funding via a contribution agreement through the National Research Council of Canada.

gradient structures used in FFAG lattice design. These include: (i) Traditional scaling FFAG; (ii) Triplet-based scaling FFAG; (iii) Nonscaling FFAG.

The traditional scaling FFAG is comprised of combined-function short FODO cells with edge focussing and magnetic fields which scale with momentum. The consequence of scaling the magnetic fields is that the orbit properties are maintained constant as a function of momentum and the optics are also independent. Such FFAG rings were first designed and studied at MURA[5]. The triplet-based FFAG is a recent innovation based on the scaled-field concept, but formed from a triplet quadrupole structure rather than a FODO one. (It was developed for the KEK Proof of Principle, or POP, machine[6].) Its primary advantage over the previous structure is incorporation of a significantly longer straight section in each cell facilitating injection, extraction, rf insertion, etc.

The nonscaling FFAG is a concept unique to muon acceleration, or when acceleration occurs so rapidly the beam experiences only a few turns in the accelerator. For such rapid acceleration, one does not have to avoid resonances or control lattice parameters as a function of momentum. Instead, one has the freedom to choose parameters optimal for muon acceleration such as minimizing circumference and requiring a large transverse dynamic aperture.

Even with careful design, the performance limitations of the scaling FFAG generally eliminate it as a candidate for high-energy muon acceleration when a large range in energy is involved. The primary obstacle comes from the large transverse muon beam size which must be transported along with the large change in energy. Keeping the orbit and optical properties consistent require the magnetic field to scale with momentum, but achieving a large transverse dynamic aperture requires that $B'$ must be nearly constant. For a large energy range, this implies the horizontal spread of orbits becomes large. Curtailing the magnet aperture means the field must rise sharply with radius through addition of higher-order field terms. Degradation of dynamic aperture follows with a final acceptance which is inadequate to transport even a muon beam cooled to the extent needed for a collider, much less a beam intended for a Neutrino Factory.

### 2.1. Example of a 6-20 GeV Nonscaling FFAG

As mentioned earlier, muon acceleration occurs so rapidly that resonances are not a consideration. In this case, the beam can be accelerated through an integer, or other resonance-driving “global” tunes if the tune is only valid for a fraction of a turn. With a fast acceleration cycle (the rf systems are assumed to deliver on the order of 1-3 GeV per turn), the lattice’s optical parameters are therefore released from scaling with momentum. Through the further choice of using only linear elements, an FFAG machine can be designed to support a large transverse dynamic aperture in addition to the large longitudinal one, again at the expense of maintaining constant optical properties. One then has the freedom to choose machine parameters which are optimal for muon beam acceleration; i.e. minimizing the circumference to limit intensity loss from decay and maximizing the transverse dynamic aperture to accept a less-cool beam. This approach has been termed a non-scaling FFAG accelerator.

Two steps are important in minimizing the machine circumference. First, the reverse bends required by the criterion to maintain constant optics can be eliminated yielding approximately a 20% decrease in total circumference. Then, also important is to choose the magnet configuration in the basic FODO cell to provide the maximum net bend per cell for a given peak excursion of the closed orbit during acceleration. This is accomplished by favourably positioning the dipole bend field over the defocusing quadrupole element. The cells of the non-scaling ring then contain a horizontally focusing quadrupole followed by a vertically focusing, combined-function bending magnet. The allowed bend is further increased by the choice of focussing strength and cell length: the lower or injection momentum experiences a cell phase advance approaching $\pi$ while the upper momentum approaches zero, depending sensitively on the choice of gradient and relative radial positions of the closed orbits, or total magnet aperture. The non-scaling approach yields the smallest design circumference of any lattice and can approach a factor of two less than that of a scaling lattice.

In designing a non-scaling lattice, the optimal lengths for the magnets are obtained analytically by assuming thin-lens kicks and imposing geometric closure on both off-momentum orbits and transversely-displaced orbits (orbits with the same momentum, but with a nonzero amplitude). In order to solve the set of coupled equations, the maximum off-axis orbit excursion in the F quadrupole must be chosen along with the F quadrupole’s aperture and poletip field. To insure

Table I Parameters of a 6-20 GeV Non-scaling FFAG. Where two sets of values are given, they represent superconducting/normal conducting magnets, respectively.

| | | | | |
|---|---|---|---|---|
| **General** | | | | |
| | FFAG type | nonscaling | Energy range | 6-20 GeV |
| | Central Energy | 16.5 GeV | Circumference | 2041/2355 km |
| | Rigidity | 55 T-m | Poletip Field | 6T/2T |
| **Arc Cell** | | | | |
| | Number | 314 | Length | 6.5/7.5 m |
| | Bend/cell | 0.02 rad | Quad Gradient | 75.9/25.3 T/m |
| | "F" length | 0.15/0.45 m | "F" strength | 1.38/0.46 $m^{-1}$ |
| | "D" length | 0.35/1.05 m | "D" strength | 0.59/0.20 $m^{-1}$ |
| **Cell Tunes** | | | | |
| | 6 GeV | 0.45 (162 deg) | 20 GeV | 0.08 (29 deg) |
| **Maximum Displacements** | | | | |
| | 6 GeV | -7.5 cm | 20 GeV | 7.1 cm |

stability in both planes, the D quadrupole strength is set equal to the F quadrupole strength. The lattice components, parameters, and functions of the non-scaling FFAG cell are given in Table I and plotted in Figures 1 and 2.

Relaxing the requirement for consistency in closed orbits at different energies means that the orbits no longer remain parallel. The peak of the closed orbit excursion always occurs at the center of the F quadrupole. Orbit excursions at 6 GeV are plotted in Figure 3. The corresponding orbit excursion at 20 GeV is almost an inversion of the 6 GeV curve as can be seen in Figure 4. The need for a large transverse dynamic aperture is automatically satisfied in this design because only linear elements are used. Because of the short cell structure, variations in the maximum beta functions are not significant enough to instigate a serious beta wave during acceleration due to a transverse optics mismatch.

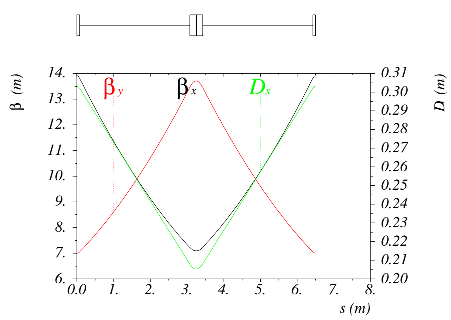


Figure 1: Lattice functions at 16.5 GeV for a 6–20 GeV non-scaling FFAG.

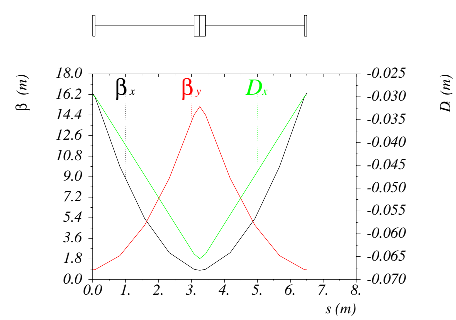


Figure 2: Lattice functions at 6 GeV for a 6–20 GeV non-scaling FFAG.

### 2.2. Pathlength Dependencies in FFAGs

A main drawback to FFAGs in both the scaling and the non-scaling versions is the large changes in pathlength as a function of energy. The pathlength dependence is clearly linear with momentum for radially-staggered, parallel orbits as in the scaling case, but it is parabolic in non-scaling FFAGs. This comes about when the transverse excursion of orbits as a function of momentum (see Figures 3 and 4) is larger than the contribution from the longitudinal pathlength change. Nominally, for small momentum deviations from the central momentum, the fact that the lower momentum stays to the inside of the central orbit and the high momentum to the outside means a smaller total pathlength for low momentum and a larger one for higher momentum. The scaling FFAG follows this norm except the pathlength variations are large because of the large momentum acceptance. For large transverse apertures, and correspondingly large excursions across the

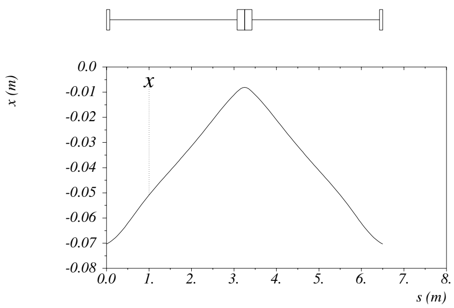


Figure 3: 6-GeV orbit amplitudes for a 6-20 GeV non-scaling FFAG.

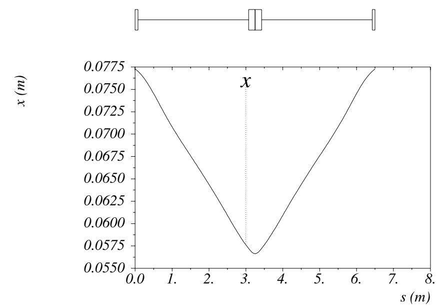


Figure 4: 20-GeV orbit amplitudes for a 6-20 GeV non-scaling FFAG.

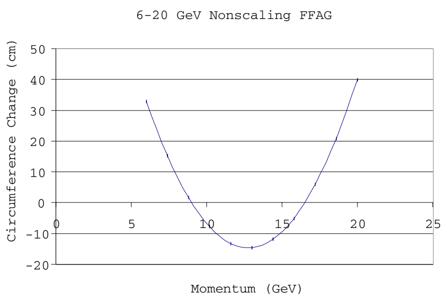


Figure 5: Circumference change as a function of momentum

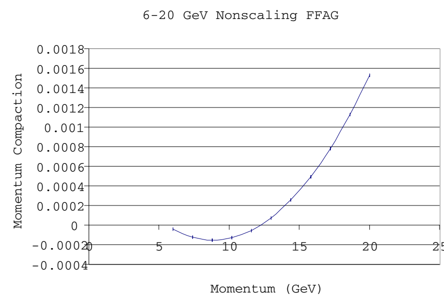


Figure 6: Momentum Compaction as a function of momentum

magnet apertures, as is the case in the non-scaling FFAG, the transverse path changes overtake these longitudinal variations. Because of their transverse offset, both high and low momentum have total pathlengths larger than the central orbit, giving a parabolic shape to the circumference change as a function of momentum (see Figures 5 and 6.

Consequently, path length and cell traversal time change with the reference energy from cell to cell. Of course, the cell traversal times must be synchronized with the waveforms in the RF cavities responsible for acceleration. In a circular machine, particles must make repeated passages through the same cavities; and so on every revolution of the machine the frequency and phasing of each cavity must be readjusted. In the proposed muon FFAG, acceleration is completed in a few turns and the RF adjustments needed are impractically large and fast in the microseconds circulation time characteristic of these machines. For example, if the rf system is 200 MHz, then a pathlength change of 50 cm from the central orbit means the bunch arrives 1/6 of an rf wavelength out of phase. Even if the phase adjustment were possible, the fast variation of path length and cell traversal time means that the notions of synchronous phase and RF bucket cannot be applied to the particle trajectories. Even if one could implement the ideal phases, the longitudinal motion would not be simple; and the usual ideas of synchrotron longitudinal dynamics are not relevant to this type of machine.

Recently workshops have focussed on the phasing problems of FFAGs and a number of solutions are being advanced[8]. These include: (i) momentum-dependent chicanes[9] to correct pathlength differences, (ii) broad-band RF that can be phased quickly but suffers from the disadvantage of comaratively low achievable acceleration voltage (1 MeV/m or less), (iii) a frequency low enough (25 MHz, for example), to make the phase shifts ineffectual, and (iv) mixing of multiple rf frequencies to produce an appropriate waveform to match the beam traversal times. However, the present work investigates the simplest approach: the application and optimization of a single high frequency, high-Q rf system.

## 3. FFAG RF system

To begin, it is important to understand the RF power requirements of such an accelerator. There are designs available for 200 MHz normal conducting (NC) and 400 MHz superconducting cavities (SC) to be used in CERN LHC[10] that provide a starting point for extrapolation. First assume there is 360 MW of wall-plug power is available and that it can be converted to RF power with 50% efficiency. If there are 300 cells and 6 cavities can be installed per cell, then the allowed power consumption is 100 kW per cavity. With a 1.7 MV gap voltage, the shunt resistance becomes 14 Mohm or more. The ratio of resistance to quality factor (R/Q) depends on cavity geometry, and approaches 200 for the CERN cavities. Hence, in this scenario, the quality factor must be at least $7 \times 10^4$. The cavity output faithfully follows its input for variations comparable with (or slower than) the cavity filling time. The filling time is equal to quality factor multiplied by RF period; it is the number of RF cycles for any disturbance to fall to $e^{-\pi} \approx 4\%$ of its initial value. Hence the filling time is 350 $\mu$s, or more, in our high-Q scenario. This should be compared with the revolution period, for a 2 km ring and light-speed particles, which is 6.7 $\mu$s. Evidently, the cavity phase cannot be made to follow the ideal variation on a turn-by-turn basis. One could imagine a vector feedback of the gap voltage so as to reduce the filling time by say a factor of 20, but this would still not guarantee sufficient waveform fidelity and, moreover, the peak power would rise. Given the large power requirement of NC cavities, the muon FFAG would clearly benefit from the adoption of superconducting cavities (SC) with quality factors ranging from $10^6$ to $10^9$. In either case, NC or SC, pure sinusoid operation is the only mode possible.

### 3.1. RF waveform scenarios

Initially all the cavities are assumed to operate at a single frequency. Then “standard” acceleration will be stochastic. However, one may vary the initial individual cavity phases simultaneously optimizing the value of the single fixed frequency; net acceleration is the result. We call this: *single frequency, best fixed phases*. In addition a flat-top can be imposed on the RF waveform by the addition of the 2nd harmonic so as to reduce the rf-beam phase variation.

Alternatively, one can imagine that the phases of the accelerating stations can be changed on a turn-by-turn basis. In essence, the modulation would be achieved by vector addition of an almost constant in-phase term and a varying quadrature term which may be generated either by synthesis of two sinusoids emanating from two high quality-factor systems, or by a single modulated low-quality-factor system. Unfortunately, in either case, at least for the muon-FFAG here, to sweep the phase correctly requires a quadrature modulation term comparable with the in-phase carrier. Thus the “quadrature-subsystems” for the 2-frequency synthesis scheme are no different from the in-phase RF system; and the alternative low-Q quadrature subsystem still has a very high instantaneous power requirement. For the purpose of calculation and comparison, let us set aside for a moment these technological problems, and assume the *ideal phases* can be provided. The question then becomes “what longitudinal phase space can be transmitted?”.

### 3.2. Optimization Strategy

There exist many possible optimization strategies, and we report only that one in which the reference bunch receives equal, maximum possible acceleration on each turn. Our optimization is based on reference particles (one per bunch)–not on complete bunches or ensembles and so it cannot be used directly to maximize input acceptance or minimize output emittance.

Initially, the ideal phases cavity-by-cavity and turn-by-turn are calculated for a single synchronous particle in the reference bunch. One assumes this particle is accelerated perfectly at every cavity, and adjusts the phases artificially to make this possible. The ideal phases and gap-crossing times are recorded.

*Ideal phases:* The ideal phases for a single reference bunch are known; but one must also consider other bunches in the train. We assume that the RF system is not sufficiently agile to make adjustments for individual bunches, but rather that it runs at constant frequency during the train and that phase adjustments are made in the beam gap. Probably the best solution is to have different frequencies on each turn, but we chose a simpler option: to use the "best frequency" found above.

*Fixed frequency and best phases:* To optimally select the fixed frequency and phases, one chooses a frequency and then calculates the phases at the gap-crossing times of the ideal particle. One then forms the mean-square deviation of the fixed-phases from the ideal values, summed over all rf stations and turns. A search is then made to find the "best frequency" which minimizes this square phase deviation. The results of this calculation is a set of "best phases" for the reference bunch; these phases are not-ideal and bear little resemblance to standard acceleration ones.

*Over-voltages:* Because every bunch arrives displaced from the ideal phase, we must find some way of accelerating that is tolerant of poor phasing. Since the time of Veskler and McMillan, it has been recognized as advantageous to use an "over-voltage". Rather than accelerate at the crest of the RF wave, one lets the ideal particle lag or lead the wave and increases the voltage to compensate. In a synchrotron, the advantage of increasing the accelerated phase space area more than compensates for the penalty of higher voltage. In a muon FFAG the benefit is less clear, *a priori*, because of the speed and path length variations. However, *a posteori*, it was found that a modest over-voltage is of enormous benefit to the performance of the machine, even though the effect establishes new limits. The over-voltage is chosen by a numerical optimization procedure so as to minimize the bunch-to-bunch variation of the extraction energy. The merit factor is obtained by tracking a reference particle in every bunch of the train from injection to extraction. There is no need to adjust the phasing; the beam will automatically adjust to the most stable phase configuration.

## 4. Simulations of longitudinal dynamics

The model used is very simple. One assumes *complete* decoupling from the transverse motion. The 2 km ring is divided into roughly 300 identical cells each with an RF station. Although the station could comprise up to 6 cavities, the energy gain is lumped in a single element. Stations all run at the same frequency, but have initial individual phases which may differ. The arrival time depends on $\beta = v/c$ and path length. A parabola has been fitted through the circumference-change versus energy curve, figure 5, and is used for the path length computation. The machine's energy acceptance is assumed to be 6-to-20 GeV $\pm 10\%$ of injection and extraction energies. Typically we track 100 bunches, with roughly 1600 particles per bunch. Initially, the longitudinal phase plane is uniformly flooded with trial particles; and one attempts to accelerate them. Particles which survive the complete acceleration to 20 GeV are recorded and used to map out both the input admittance and the output emittance. To increase accuracy, the procedure is repeated using the admittance of the previous trial as a basis for populating the input ensemble of the subsequent trial. A wide variety of cases has been considered and is reported elsewhere[11]. Here we shall describe indicative cases.

### 4.1. Five-turn, 200 MHz acceleration

*Single frequency:* We shall consider acceleration in five (5) turns of the FFAG. If one uses a single (but optimal) frequency, fixed initial cavity phases and the nominal RF voltage of 9.33 MV/cell, then the desired acceleration is not achieved. If one adds second harmonic, a small phase space area of 0.15 eV.s is transported to 18.2 GeV and there is no emittance growth.

*With over-voltage:* If one allows a 40% over-voltage (case #8), then an input phase space area (admittance) of 1.37 eV.s is transported from 6 to 20 GeV by which time the emittance has increased to 1.375 eV.s.

*Dual harmonic:* If one adds 2nd harmonic (case #9), then one needs $1.333 \times 1.249$ of the nominal voltage. The admittance rises to 1.94 eV.s and the output emittance is 1.91 eV.s. As will be noted from Figures 7 and 8, transport is non-linear and the occupied phase space has a bizarre shape. The input admittances for cases #8 and #9 are similar to that in Figure 11 except that their time-width is halved.

*Ideal phases:* When the ideal phasing is used, and no other measures are taken (case #6), an input admittance of 1.21 eV.s is successfully accelerated to 20 GeV with an output emittance of 1.18 eV.s. Notice, Figures 9 and 10, that despite the ideal phasing, the transport is non-linear.

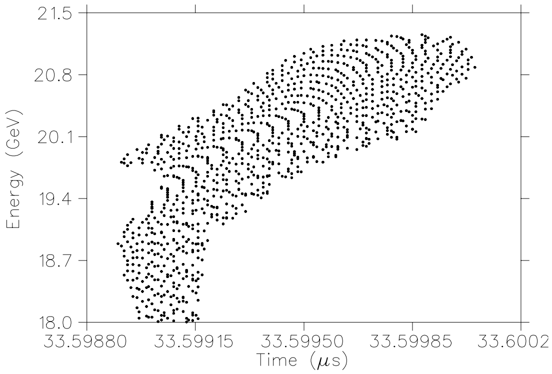


Figure 7: Output Emittance, Case #8. Single harmonic.

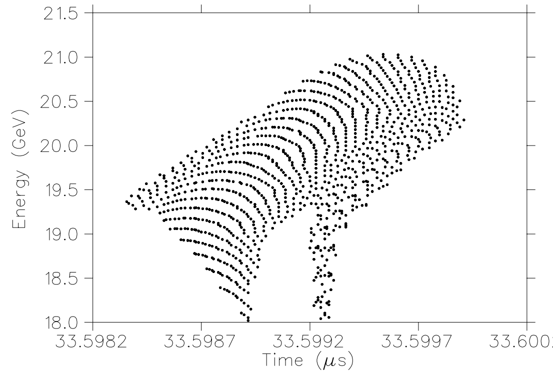


Figure 8: Output Emittance, Case #9. Dual harmonic.

If one adds second harmonic to the basic scheme, case #5, then the input admittance rises to 2.22 eV.s and there is no output emittance growth. However, the transport is so non-linear that the useful phase space area is probably only one half of these values. Perhaps the situation could be improved by adjusting the phase of the second harmonic.

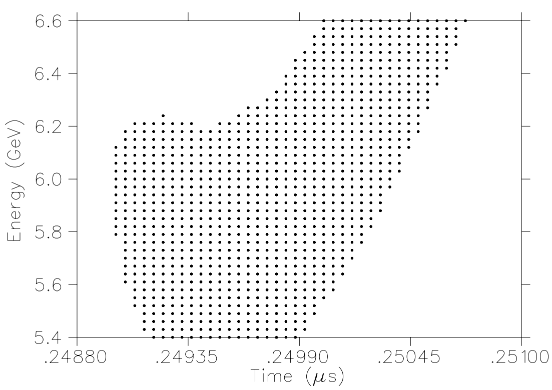


Figure 9: ±10% band from Input Acceptance

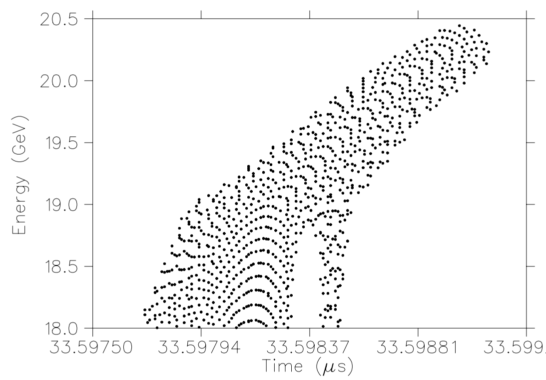


Figure 10: Maps to the Output Emittance, Case #6. No over voltage.

*Conclusion:* For 5-turn acceleration, there is little difference in the output emittance between the use of "ideal phases" versus using a combination of "best phases" and an over-voltage. When second harmonic is employed, the useful acceptance is greater when the "single-frequency" scenario is adopted.

### 4.2. Ten-turn, 200 MHz acceleration

To first order, the phase slippage will increase linearly with time; and so one expects the transmission to *fall* in a roughly parabolic manner as the number of turns is increased. We have not achieved successful acceleration over ten (10) turns with 200 MHz RF and 4.7 MV/cell unless the ideal phases are used. When ideal-phases and single-harmonic is used for 10-turn acceleration, the admittance falls by about 40% compared with 5-turns (case #6). The addition of second harmonic restores the admittance to within 12% of the 5-turn value (case #5).

### 4.3. Acceleration with 100 MHz RF

The phase slips accumulate half as quickly when the RF is halved, and so one anticipates that 100 MHz acceleration will be less compromised by using a larger number of turns. We have considered only the use of a single frequency and "best phases". The input admittance rises almost linearly from 3.0 to 3.7 eV.s as the number of turns is reduced from nine to five and the cell-voltage is raised from 5.2 to 9.3 MV.

*10-turns:* With the nominal voltage of 4.7 MV/cell, and a small 4% over-voltage, the transmission falls to zero for 10-turn acceleration. However, a modest over-voltage of 30% will restore transmission, case #31, and yields a 1.8 eV.s acceptance. If one adds second harmonic and an over-voltage of 27%, the admittance rises to almost 4.0 eV.s.

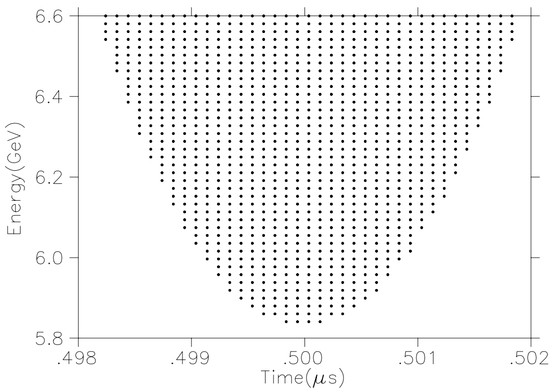


Figure 11: ±10% band from Input Acceptance

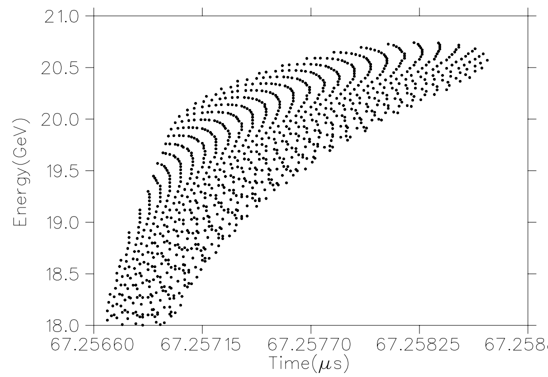


Figure 12: Maps to the Output Emittance, Case #31. 10 turns, 30% over voltage.

*Conclusion:* For acceleration with 100 MHz RF using a single frequency and fixed "best phases", the optimum number of turns appears to be seven or eight. However, if it is desirable to reduce the gradient, acceptable results can be achieved in 9,10 or even 11 turns.

## 5. Summary

Scaling FFAGs for the most part have not been found to be applicable to muon acceleration in the multi-GeV regime. This is due to their poor transverse dynamic aperture resulting from strong nonlinear field profiles. Although superconducting magnets with large horizontal apertures are required for efficient acceleration in a non-scaling FFAG, this approach does provide the necessary transverse and longitudinal acceptance match to high-energy muon beams. Initially it was felt that the circumference change, or phase-slip, posed a serious problem, but since then numerous solutions have been proposed.

When 200 MHz RF is utilized, useful admittances ($\simeq$ 1 eV.s) can be achieved with acceleration in 6 or less turns using either either a phase-agile or a high-Q fixed-frequency RF system and a modest over-voltage. For the case of 100 MHz RF, study of a fixed frequency system, shows that an output emittance of ($\simeq$ 2 eV.s) is realized for acceleration in 10 or less turns. In all cases, addition of second harmonic roughly doubles the admittance but the non-linear effect is augmented.

In conclusion, it looks promising to build a chain of muon accelerators from FFAGs and replace the costly and somewhat restrictive RLAs, which so far have been the baseline accelerator for the feasibility studies of a Neutrino Factory in the U.S [2, 3].

## References

[1] The $\mu+\mu-$ Collaboration, $\mu+\mu-$ Collider: A Feasibility Study, BNL-52503, Fermi-Conf-96-/092, LBNL-38946, July 1996.

[2] N. Holtkamp and D. Finley, eds., A Feasibility Study of a Neutrino Source Based on a Muon Storage Ring, Fermilab-Pub-00/108-E (2000).

[3] S. Ozaki, R. Palmer, M. Zisman, and J. Gallardo, eds., "Feasibility Study-II of a Muon-based Neutrino Source", BNL-52623, June, 2001, available at http://www.cap.bnl.gov/mumu/studii/FS2-report.html.

[4] C. Johnstone, W. Wan, and A. Garren, "Fixed Field Circular Acceleration Design", Proceedings of the 1999 Particle Accelerator Conference, New York, NY, Mar. 29-Apr. 2, 1999, pp. 3068.

[5] K. Simon, MURA-KRS-6, (MURA-43) (1954); K. Simon, et al, Phys. ReV. 103 p. 1837, (1956)..

[6] S. Machida, et al, “Beam Optics Design of an FFAG Synchrotron”, (MOP1B20), submitted to EPAC 2000, and Y. Sato, et al, “Development of a FFAG Proton Synchrotron”, (MOP1B21), submitted to EPAC 2000.
[7] W. Wan, unpublished, 1998.
[8] FFAG Workshop, CERN, Geneva, Switzerland, July3-4, 2000 and FFAG2000, KEK, Tsukuba, Japan, Oct. 11-13, 2000.
[9] A. Garren, presented at the CERN FFAG Workshop, Geneva, Switzerland, July3-4, 2000.
[10] D. Boussard et al: “Design Considerations for the LHC 200 MHz RF System”, LHC Project Report 368, January 2000.
[11] S. Koscielniak: TRIUMF design note TRI-DN-01-15, September 2001.