A Driver Accelerator for an FEL Upgrade

D. Douglas

Abstract

We describe a concept for an accelerator to drive an upgraded IR/UV FEL.

Machine Overview

Assumptions: In this note, we make the following assumptions about accelerator physics issues relevant to FEL drivers.

- We assume the bunch formation process required to support FEL operation is tractable and may be accomplished using an injector similar in performance to the design goals of the IR demo injector.
- We assume space charge effects are not an issue after the bunch is formed.
- We assume that other collective effects, notably BBU, CSR and other wakefield /impedance driven effects are manageable.
- We assume full phase space matching (in six dimensions) will be required at the wiggler(s) to ensure FEL performance goals are met.
- The breadth of the desired spectral reach (IR to UV) suggests at least two FEL systems must be supported. One (the IR) will be based on the available Northrop-Grumman wiggler, the second (the UV) on a wiggler to be specified (but which may be similar to an APS undulator A [1]).

Parameters: A draft parameter set is given in Table 1 [2].

Consequences: The above assumptions and parameter set suggest consideration of a driver accelerator based on an energy recovering SRF linac with FELs in the backleg transport. A similar architecture has been successfully employed in the IR Demo program. Key and unique to this implementation is the assumption that CSR does not hopelessly degrade the accelerated beam, which allows placement of the FEL systems in the machine return transport. The resulting footprint is thereby reduced and the device will readily fit in the available vault.

A straightforward application of the IR Demo design methodology [3] allows an immediate description of the system details.

Table 1: Parameter Set for IR/UV Upgrade

Parameter	IR	UV
Energy (MeV)	100	Best effort, 3 ¼ cryomodules
Normalized emittance (mm-mrad)	<30	<11
Peak current (A)	60-200	270
Charge/bunch (pC); Frequency (MHz)	135; 74.85	135; 74.85
Average Current (mA)	10	5
Longitudinal emittance (keV/deg) Bunch length (psec) Energy spread (keV)	30 0.2 280	30 0.2 280
FEL wavelength (μm)	0.3 – 60	
FEL Power	>10 kW (@ 1 μm)	>1 kW (@ shortest possible wavelength)
Wiggler λ_W (cm); N_{period}; k^2	20; (12,"1",12); 16	3.3; 72; 2.7

Details of System Architecture

The proposed driver is visualized as a two pass (one acceleration, one energy-recovery) recirculating single linac. The FEL(s) will be located parallel to, but displaced from, the accelerating sections in the recirculator backleg. The driver consists of the following subsystems:

- a 10 MeV injector with a source upgraded to meet the performance goals proposed for the IR Demo gun
- injection/extraction systems for 10 MeV injected and energy recovered beams
- a three module linac
- betatron matching to the recirculator
- transport to the machine backleg
- a "wiggler garden" (a really little "undulator farm") in the machine backleg, with any capabilities required for either IR or UV operation.
- transport to the linac axis for reinjection
- reinjection betatron matching

Figure 1 provides a conceptual representation of the accelerator system design. This note assumes that beam dump configurations will remain essentially the same as in the IR Demo and will therefore not address this topic.

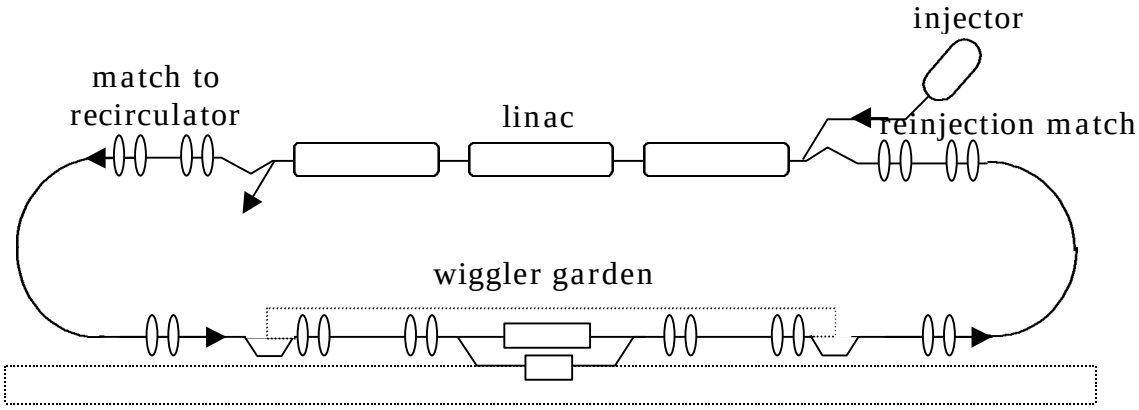

Figure 1: System design concept.

Injector: As it was originally designed to support a 1 kW UV FEL, the existing IR Demo injector design is, in gross detail, capable of supporting the proposed upgrade. Operational experience with the IR Demo and the demand for quite high current (10 mA) to support 20 kW IR operation suggest certain injector improvements:

- Gun performance and reliability must be improved to ensure the required phase space (particularly emittance targets at high bunch charge, as in Table 1) and machine availability can be achieved.
- Cryounit performance should be improved to provide a full 10 MeV (kinetic) injection energy and capability of accelerating 10 mA.
- The sensitivity of beam and machine performance parameters to operational variables should be better understood, controlled, and documented.

Injection/Extraction: will be accomplished as in the IR Demo. Inasmuch as the injected and extracted energies are the same as in the IR Demo, the existing injection and extraction dipole system could in principle be reused, though the detailed layout must be changed so as to accommodate the change in energy ratios of the new machine. If the machine is to be operated over a broad range of final energies (100 MeV for IR, to 200 MeV for UV), we note that the aperture limitations that are of concern in the IR Demo will become even more aggravated. We therefore will review the design of this subsystem in some detail. A rectangular dipole alternative to the present solution (which uses wedge dipoles) could in principle provide significantly better tolerance to variations in the ratio of injected to final beam energy.

Linac: Significant improvement in available SRF cavity performance has been achieved since the original Jefferson Lab UV FEL proposal [4,5].

Consequently, it is reasonable to propose this driver be a simple single-pass energy recovering linac design based on high gradient cryomodules, rather than a more challenging multipass design. We note that schedule pressures will probably force use of modules with nominal energy gains of order 55, 55, and 60 MeV.

To avoid betatron mismatch with attendant spot growth and beam loss in both accelerated and energy recovered beams, it is desirable (in a short linac with modest front-end to back-end energy ratio) to use "constant gradient" triplet focussing between the modules [6]. The beam transport of the accelerated and energy recovered beams can then be made to be "mirror images" of each other (provide a proper choice of recirculator transport matrix is made). A potential solution is shown in Figure 2, which displays a mildly optimized solution for the linac optics, using the same beam envelope injection conditions as were used in the IR Demo design computations.

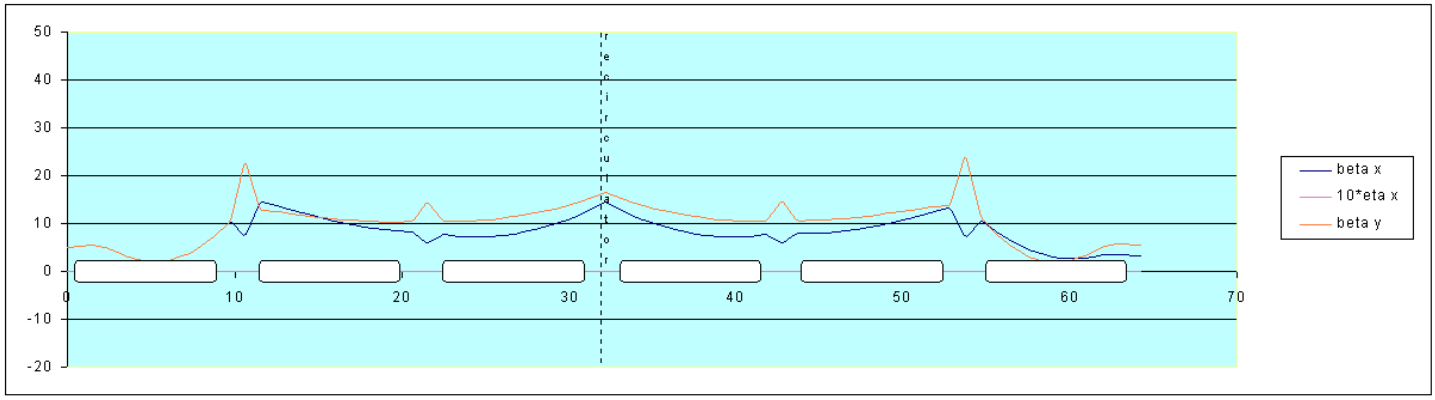

Figure 2: Beam envelope solution for driver linac (55, 60, 55 MeV energy gains assumed).

Preliminary studies suggest that deviations of energy gains on the order of 10% from the ideal do not dramatically degrade beam behavior, provided the quadrupoles are adjusted to compensate. Figure 3 shows the result of altering the linac energy profile from the Figure 2 configuration (of 55, 60, and 55 MeV) to place the highest gradient module last (a 55, 55, 60 MeV pattern). The beam envelopes are not significantly different from the nominal. Schedule and availability can thus probably establish the installation order of cryomodules. We note that the desirability of early beam operations (prior to availability of the third module) would suggest the first two modules be installed in the first and final slots. Initial operation of the linac could then proceed without the need to either transport energy recovered beam through an entire module slot to reach the dump beam line or to move the dump beam line to alleviate low energy beam loss.

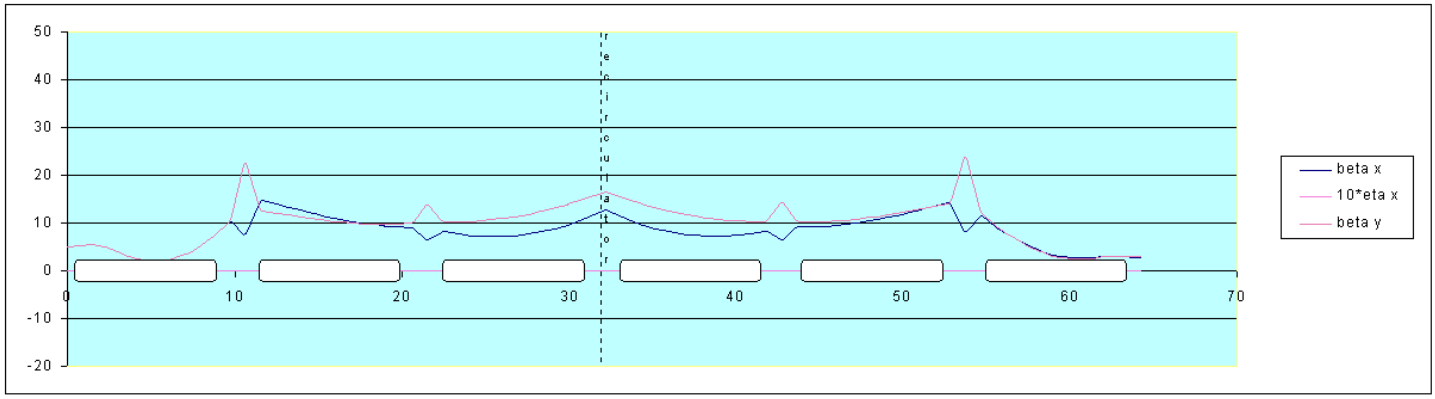

Figure 3: Beam envelope solution for driver linac with asymmetrical energy profile (55, 55, 60 MeV energy gains assumed).

Ongoing beam dynamics studies should lead to better understanding of the sensitivity of beam performance to deviations in energy gain amongst modules. Methods to control beam envelope mismatch induced by this variation in energy gain during acceleration and energy recovery will be developed. It is possible, for example, that the available modules at some point in the upgrade may have gains of order 40, 60 and 80 MeV. In this case, initial results suggest that cavity focussing in the 80 MeV module is so strong that severe mismatch could result if it is used to accelerate or energy recover 10 MeV beam [7]. This is shown in Figure 4, which presents optimized beam envelopes for a linac with 80, 40 and 60 MeV energy gains. The performance appears better for a more "symmetric" energy gain profile, such as 40, 80 and 60 MeV. Beam envelopes using this profile are shown in Figure 5. This case provides distinct schedule advantages as well – the installed module provides ~40 MeV energy gain; installation of a 60 MV module in the third linac slot would allow early operation of the machine at 100 MeV for IR FEL commissioning and studies, and would allow for later installation of a high gradient module at the optimal location.

The issue of compensating HOM-driven skew quad coupling will require resolution prior to finalizing any IR upgrade design. As noted elsewhere [8], DC compensation of this error for one pass will result in a doubling of the effect on the other. Various schemes for compensation must therefore be considered. Simplest is the present approach – correct the accelerated beam (which must be of good quality at the wiggler) and try to survive the error during energy recovery. This may not be possible in a 3-module linac. Alternately, the beam could be "pre-corrected" in the injector and/or in the recirculator. This would meet performance requirements, but would result in a larger, more complex, more difficult to operate machine, inasmuch as either "decoupling" regions would be required or the decoupling functionality would have to be included within modules presently supporting other functions.

The most attractive alternative may be to provide local correction through the use of small RF skew quadrupoles in the linac. This may require the most expensive and complex individual elements, but still give the best cost/performance return inasmuch as it does not impose the complexity associated with nonlocal magnetic decoupling. Local correction of this type may however fail if the energy gains and focussing from individual modules is so high that the phase advance across a module produces betatron phase mixing of the skew coupling within that module. Analysis of and action on this issue will be ongoing during the driver design process.

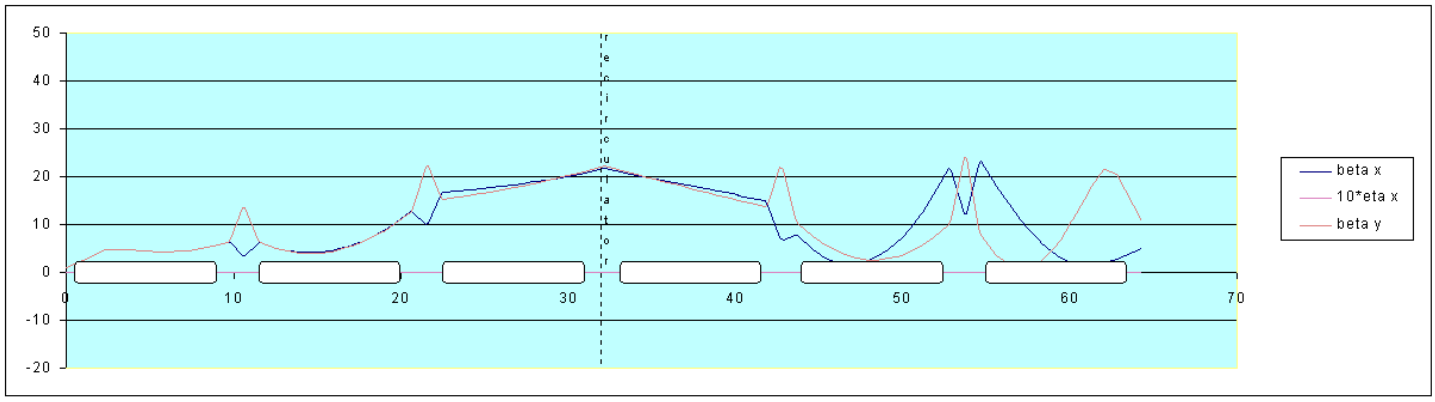

Figure 4: "Optimized" beam envelopes for system with 80, 40 and 60 MeV energy gains. Note the severe focussing in the first module during acceleration and in the final module during energy recovery. To avoid significant mismatch during acceleration we must use very small injected beam envelopes, which leads to concerns about space charge effects.

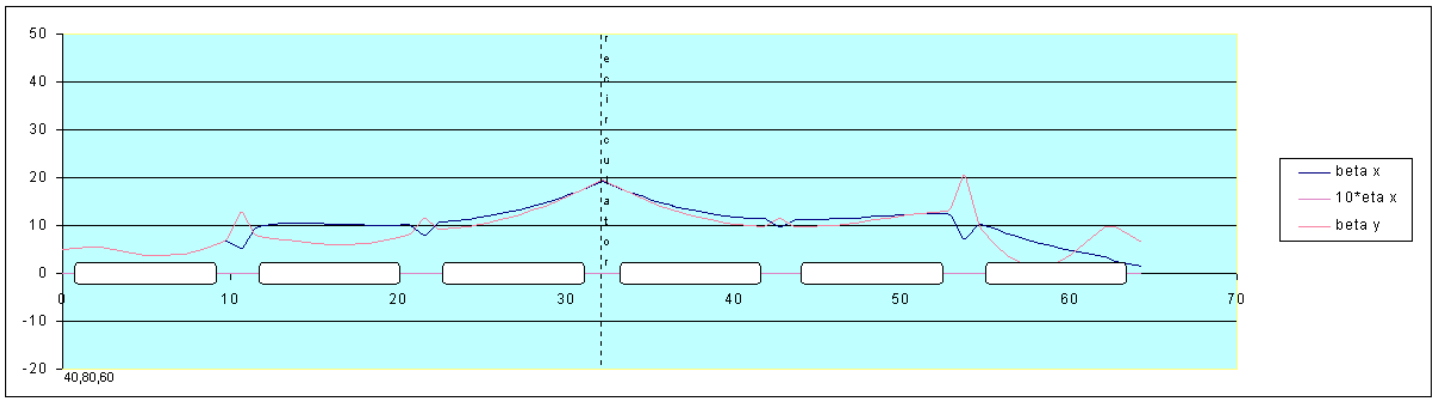

Figure 5: "Optimized" beam envelopes for system with 40, 80 and 60 MeV energy gains. Note the reduction in focussing strength and resultant improvement in beam behavior.

Matching to/from Recirculator: The virtues of matching are well known. The linac transverse optical solution described above is based on the assumption that the recirculator global transport matrix can be configured to provide $\beta_{out} = \beta_{in}$ and $\alpha_{out} = -\alpha_{in}$ in both transverse planes, at arbitrary (but assumed operationally defined and possibly variable) phase advances. As the recirculator transport will in general not *a priori* provide this behavior, matching to and from the recirculator will be required. We will attempt to perform each match (to the recirculator, for reinjection) using an individual quadrupole telescope consisting of a pair of quad triplets. Operational experience with the IR Demo reinjection telescope suggests that it is, at roughly 3 m in length, too short to provide good dynamic range. We will therefore make these telescopes ~6 m long.

If the HOM-driven skew quad effects of the cryomodule cannot be locally corrected, these matching regions may have to provide skew quadrupole systems for decoupling the beam as well.

Transport to Wiggler Garden/Backleg and from the Backleg to Reinjection: Various end loop designs may be considered for the transport of the electron beam from the linac to the wiggler garden, and thence back to the linac for energy recovery. The two end loops must accommodate somewhat different requirements – the first must transport a bright electron beam to the wiggler without phase space degradation (particularly from CSR), the second must accept a far dimmer beam, with very large momentum spread. Both, however, must provide some level of momentum compaction management, the first to provide high peak current at the wiggler, the second to prepare the electron beam for energy compression during energy recovery. At present, it is likely that compaction management through second order will be adequate for the first end loop, but the second end loop must provide compaction management through third order and/or active bunch compression through the use of RF elements [9].

Amongst candidate end loop designs are

- arcs based on multiply periodic FODO structures, with or without dispersion modulation for compaction management,
- systems based on dispersion modulating focussing structures, such as Steffan systems, CEBAF recirculation arcs, and related structures such as the Pelligrini/Amiry design used in the original JLab UV proposal [10], and
- those based on compaction management primarily through reverse bending, such as the Bates design used in the IR Demo or the “Virginia Reel” lattice developed as a preliminary UV design [11].

Each option has its own strengths and deficiencies. FODO based arcs use simple components, but are difficult to modify to provide momentum compaction management. Moreover, the element count for such designs tends to be large, and the apertures have to be rather large to accommodate the large momentum spread [12]. Finally, FODO systems with adequate compaction management typically require significant dispersion modulation, which either dramatically increases their complexity (and size) or breaks their symmetry and periodicity, and degrades chromatic performance.

Steffan systems and CEBAF arcs similarly suffer from either size/complexity issues or degraded chromatic performance, often requiring several sextupole families to provide even a few percent of momentum acceptance. This is true as well of the original JLab UV design, which used many families of quadrupoles and sextupoles to manage off-momentum behavior. The resulting system had good performance over a momentum bite of ~5%, but the system was mechanically complex and existing documentation seems to indicate the performance is tailing off strongly outside the nominal 5% acceptance [13]. Because of the use of multiple sextupole families, nonlinear compaction management is not transparent. However, the elements can often be made with far smaller apertures than in some other options, particularly the Bates system, by virtue of the relatively small dispersions encountered.

Successful operation of the IR Demo suggests the use of Bates end loops in this application. The combination of good transverse beam behavior, widely variable momentum compaction, and relatively easy operation offsets the inherently large dispersions and commensurately tight magnet performance specifications. This design choice assumes, as noted above, that CSR will not catastrophically degrade beam quality. An elementary argument suggests that CSR effects will scale inversely with energy in this system. The phenomenon of coherent synchrotron radiation is one in which the radiation field of a bunch moves energy around in the bunch, inducing a momentum shift dp/p on an electron that depends on its longitudinal displacement s from the bunch centroid.

$$\frac{dp}{p} = \left\{ \frac{18000\, Q\, r^{1/3}\, \Delta q}{3^{1/3}\, \sqrt{2p}\, s_s^{4/3}\, E} \right\} Int\left(\frac{s}{s_s} \right)$$

$Int(s/s_s)$ is the wakefield integral, depicted in Figure 6. Of particular interest is the fact that dp/p is proportional to $1/E$. Inasmuch as this momentum shift occurs during bending, and thus at a dispersed location, it produces downstream transverse offsets x and x' that correspond to a growth in the geometric emittance. The geometric emittance is quadratic form in these variables, so its growth scales like $(dp/p)^2$, or, equivalently, $1/E^2$. The CSR

induced growth in normalized emittance thus goes like $1/E$, suggesting that this effect may be less important at higher energies.

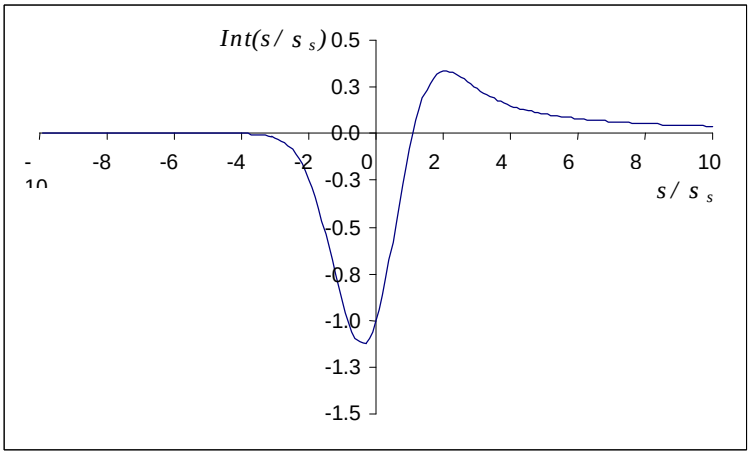

Figure 6: The CSR wakefield integral, giving the magnitude of the momentum shift as a function of position along the bunch

This inference is supported by a simple simulation [14], results of which are shown in Figure 7 and 8. Figure 7 shows portions of the phase space of a bunch with small transverse but finite longitudinal emittance after transport through an IR Demo endloop at 40 MeV. The charge/bunch was 135 pC, and the final geometric horizontal emittance grew from ~1 mm-mrad (normalized) to ~3.5 mm-rad. The structure in the two surfaces of section is indicative of the coherent nature of the phenomenon. This simulation was repeated for various initial geometric emittances over a range of energies. The results are presented in Figure 8. It was found that the growth in normalized emittance fell off at higher energies, so that a 10 mm-mrad (normalized) beam grew by only 10% at 100 MeV, with the increase falling to 5% at 200 MeV. This is within the tolerance of the system, suggesting the use of Bates end loops would be feasible even for the relatively stringent requirements of the UV system.

For design simplicity (and because the Bates design is quite sparse, relatively inexpensive and operationally robust [15]), both end loops will be identical. Though the first end loop will in this machine not need a large momentum acceptance, the Bates system provides several advantages over other designs for the linac to wiggler garden transport. Compaction management is easily performed in the Bates system, it possesses excellent chromatic behavior and IR Demo experience suggests it is simple to reproducibly operate when properly instrumented.

As noted above longitudinal matching will present a greater challenge in this machine than it did in the IR Demo. To this end, we will include in the trim system of the end loops not only trim quadrupoles and sextupoles, but also a

system of octupoles to allow third order compaction management. Further experience with the IR Demo compaction management and phase transfer function measurement systems will give operational guidance on how best to quantify third order compaction management schemes. We note that this is an issue primarily for the 20 kW IR FEL, which induces a much larger momentum spread than that imposed by the UV device. As a fallback position, the backleg will therefore provide a capability to include active RF manipulation of the longitudinal phase space downstream of the IR FEL, following a scheme proposed by Neil [16].

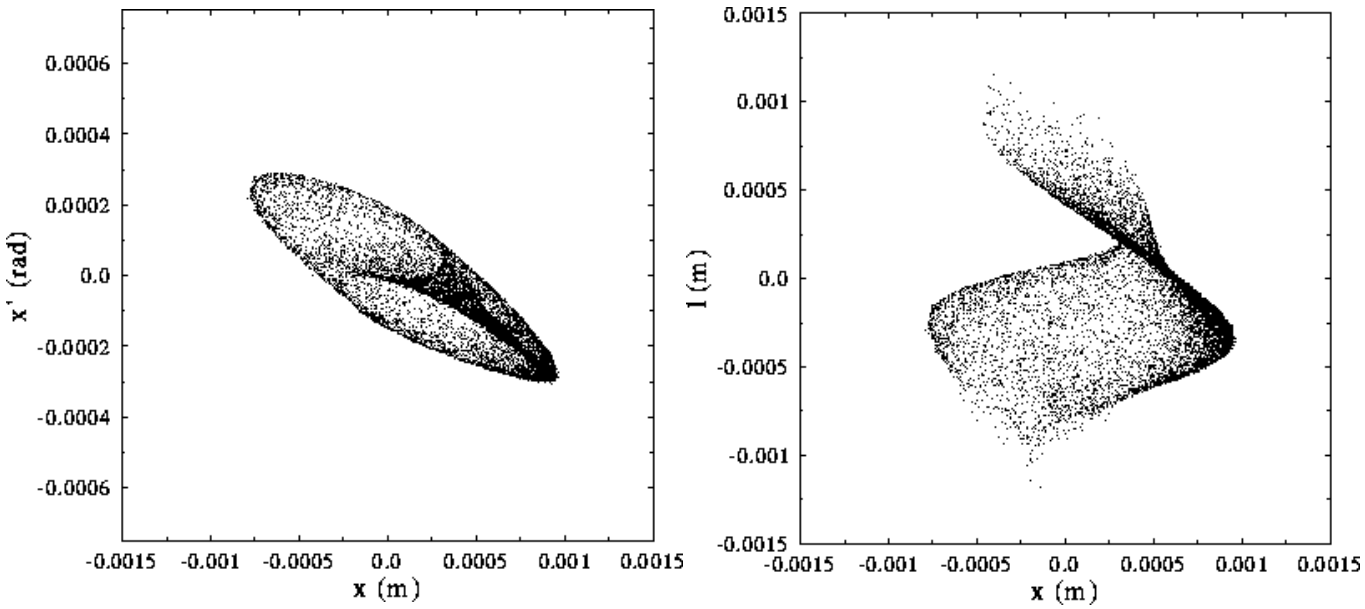

Figure 7: Simulation of CSR effects in Bates end loop.

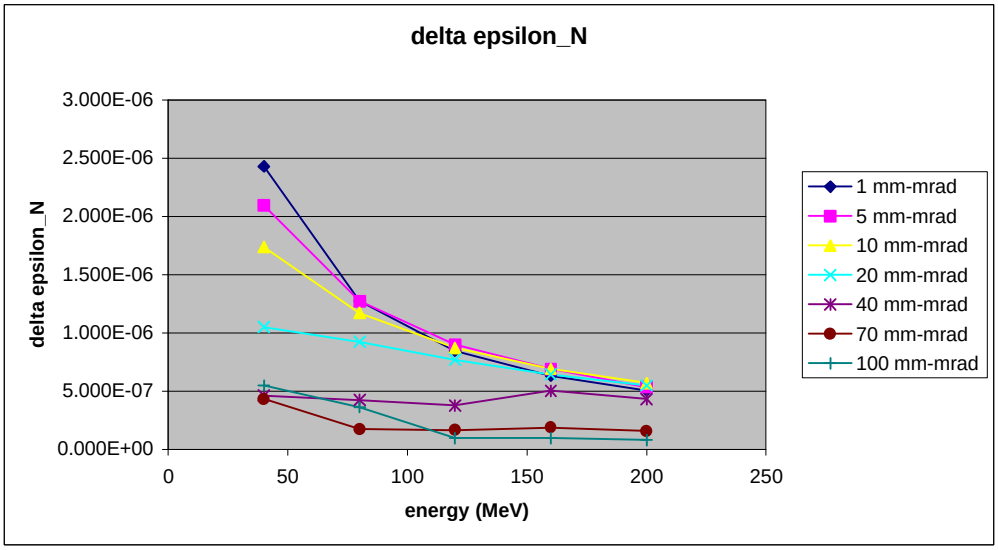

Figure 8: CSR-induced growth of normalized emittance as function of energy and initial emittance.

Backleg/Wiggler Garden: The backleg of the machine must accommodate the wiggler garden containing the IR and UV FELs. It must provide matching to and from the wigglers, allow adequate space for optical components, and, to allow for RF manipulations, provide space for a standard JLab cryounit following a magnetic bunch rotation after the IR wiggler.

A concept providing all the required functionality is shown in Figure 9. The 20 kW far IR FEL, based on the Northrop-Grumman wiggler and requiring a large (10%) momentum acceptance, is embedded on the backleg axis. A 1 kW near IR FEL, which is based on the STI wiggler, requires a smaller momentum acceptance, and has a shorter optical cavity, is parallel the far IR system but displaced to the interior of the recirculator by means of an achromatic staircase translation. A 1 kW UV FEL is similarly parallel to the far IR system, but displaced to the exterior of the machine. This configuration takes advantage of the smaller momentum acceptance required by the UV FEL to provide for the much larger footprint needed by the UV optical cavity. Quadrupole telescopes just after the first and just before the second end loops provide betatron matching across the staircases to and from the near IR and UV wigglers. Additional telescopes (on the backleg axis "under" the staircases, and at the top of the staircases) provide extended dynamic range for matching to and from the various wigglers.

The far IR system is asymmetrically located on its optical cavity [17], allowing positioning of one end of the far IR optical cavity upstream of the backleg, beyond the end loop, with the second end embedded in an optical cavity chicane downstream of the wiggler. This chicane also introduces momentum compaction, which can be used with RF supplied by a downstream cryounit to provide active bunching prior to transport for energy recovery, as discussed above. Engineering designs for the transport elements must, in this scenario, allow adequate aperture to accommodate the large angular divergence of the far IR optical mode. Dipole gaps and quadrupole bores of as much as 3 inches (7.5 cm) may be required.

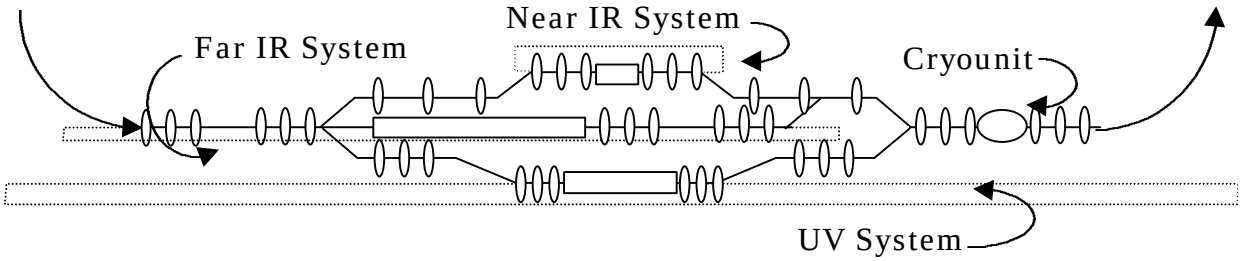

Figure 9: Concept for Backleg

Decoupled operation of the various UV and IR systems is readily achieved in this scenario. If driving one FEL, one need only turn on or off the staircase dipole string, change the machine energy, resteer as needed, and operation of the other FEL can commence. (This activation/deactivation of the dipole string can in fact be used as a flag in the machine protection system to limit current to the UV system, for example, by switching the electron gun drive laser repetition rate. This is akin to the present operating limits imposed by dipole string excitations or the presence of insertion devices in the IR Demo). As noted, the proposed layout also provides opportunity for RF-based longitudinal matching following the IR FEL (which may induce a very large electron beam momentum spread). The far IR downstream optical cavity chicane provides momentum compaction, which, together with the indicated downstream cryounit, can be used for energy compression of the beam before transport through the second end loop and energy recovery.

A simpler configuration, perhaps more appropriate to initial machine operation in the far IR only, is shown in Figure 10. This layout embeds only the far IR system in the backleg, relegating other FEL systems to future upgrades. It may aggravate potential pitfalls in the operation of the machine to recover large momentum spreads, inasmuch as it seeks to minimize the transport distance of the waste beam by locating the far IR wiggler as close as possible to the second end loop. It is essentially a mirror image of the backleg axis shown in Figure 9. Near IR and UV systems may be added at will in the future; the principle loss of capability comes with the absence of the downstream optical cavity chicane (and consequential loss of momentum compaction) and the lack of available space for a cryounit. This configuration, though operationally simpler than the one discussed above, precludes the use of active RF compression of the longitudinal phase space prior to energy recovery, and thus relies completely on transport-induced bunch length manipulations to provide energy compression during energy recovery.

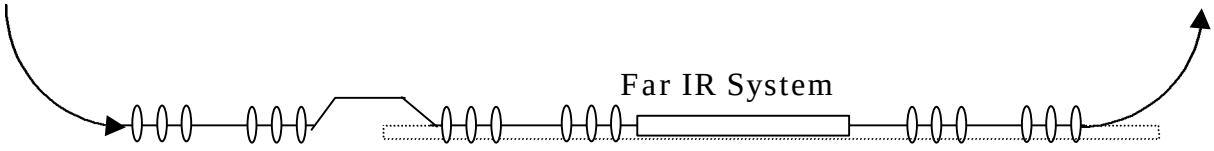

Figure 10: Alternate Concept for Backleg

Accelerator System Component Requirements

Table 2 gives a preliminary breakdown of accelerator system components required for this machine. The parts count assumes the “minimalist” upgrade of Figure 10 is to be implemented. If the Figure 9 configuration is desired, an additional cryounit, 12 ~30^{o} dipoles, 24 quadrupoles, and perhaps 8 or so each of sextupoles and octupoles will be needed.

Table 2: Major Accelerator System Components

System	*#*	*Utilization/Comments*
RF		
Injector		
buncher	1	may reuse existing system
cryounit	1	new unit for higher injector energy
Linac		
cryomodules	3	potentially 2 new and 1 reprocessed
RF		
RF skew quads	6?	requirements to be set by design study*
Beam Transport		
Injector		
~20^{o} dipoles	3	injection line dipoles (one may be reused)
solenoids	2?	requirements known after source redesign, may reuse existing components
quadrupoles	4-6?	requirements known after source redesign, may reuse existing QJs
skew quads	2?	requirements to be set by design study*
Linac		
quadrupoles	6	reuse existing QGs
skew quads	6?	requirements to be set by design study*
Recirculator		
Upstream/reinjection matches		
quadrupoles	12	probably QCs
skew quads	8?	requirements to be set by design study*
End loops		
~30^{o} bends	8	large far IR optical mode may demand 3” gap for (some) reverse bends
180^{o} bends	2	
trim quads	8	
sextupoles	8	possibly recoiled SCs?
octupoles	8	
Backleg transport – matching		
~20^{o} bends	4	optical cavity chicane – large far IR optical mode may demand 3” gap
quadrupoles	18	possibly QC, but large far IR optical mode may demand larger (3”) bore
Reinjection/Extraction		
dipoles	7	all assumed identical, required in 4-dipole reinjection/extraction chicanes

*See discussion H/V coupling correction in the previous section, under the heading of “Linac”

Table 2 suggests contemplating some component reuse. The following tabulation reviews some of the elements that may be salvageable:

- *Injector* – solenoids, buncher, cryounit (though it should be reprocessed to raise gradient and accommodate higher forward powers and beam currents), correction dipoles, matching quadrupoles, central injection line dipole
- Many or even all beam line quadrupoles and dump line quads
- Some or perhaps all correction dipoles
- Beamline instrumentation
- Control components
- SC sextupoles, with new coils

Virtually all dipoles must be replaced. This is necessary in the injection/extraction chicanes to increase the energy bandwidth of the recirculated beam and in the arcs to provide higher fields and larger momentum acceptance. Many new quadrupoles will be needed, either to reach the required gradients or to provide adequate bore to accommodate the larger dispersed beam size (trim quads) and IR optical mode (backleg matching quads).

References

[1] S. Benson, private communication

[2] G. Neil, e-mail of 15 March 1999.

[3] See, for example, the "Guide to the Design" links available from http://www.jlab.org/~douglas/FEL/FELmasterindex.html .

[4] I. E. Campisi, J. R. Delayen, L. R. Doolittle, P. Kneisel, J. Mammosser, L. Phillips, "Superconducting Cavity Development for the CEBAF Upgrade", Proc. 1999 I.E.E.E. Part. Accel. Conf., New York, 1999.

[5] Laser Processing Consortium, "Free Electron Lasers For Industry", May 1995.

[6] For a contrasting discussion in the context of high energy accelerators, see D. Douglas, "Lattice Design Principles for a Recirculated, High Energy, SRF Electron Accelerator", Proc. 1993 I.E.E.E. Part. Accel. Conf., Washington, DC, May, 1993.

[7] This, of course, has implications for the proposed CEBAF upgrade, which may use 80 MeV cryomodules to provide higher injection energy.

[8] D. Douglas, "Beam Transport Issues in the Winter/Spring 1999 FEL Run", JLAB-TN-98-008, 9 April 1999; for a complete description of the source of this error, see Z. Li, "Beam Dynamics in the CEBAF Superconducting Cavities", Ph.D. thesis, Dept. of Physics, College of William and Mary, March, 1995.

[9] This issue was addressed in D. Douglas, "Modeling of Longitudinal Phase Space Dynamics in Energy-Recovering FEL Drivers", JLAB-TN-99-002, 14 January 1999. G. Neil has proposed active bunching through the use of RF components in an e-mail message of 10 November 1998, in which he details an active energy compression system using an "afterburner" cryounit.

[10] Laser Processing Consortium, *op. cit.*

[11] D. Douglas, S. Chen, and J. van Zeijts, "A Lattice Design for a 200 MeV FEL Driver Linac", CEBAF-TN-94-019, 14 March 1994.

[12] For parameters typical of systems at this energy scale (4 period FODO structures with 1 meter bend radii and 50% packing fraction), matched dispersions of order 1 m are typical, meaning quadrupole apertures must be in excess of 10 cm to allow 10% momentum acceptance. Moreover, the resulting FODO structure momentum compaction will be on the order of p m, requiring external correction through the use of either multiple chicanes or dispersion modulation. The first option consumes considerable space and can introduce significant chromatic aberration in focussing; the second option may be more compact, but typically suffer from even larger chromatic effects than the first.

[13] See, for example, Figures 5-17 and 5-20 of Laser Processing Consortium, *op. cit.*

[14] D. Douglas, "Suppression and Enhancement of CSR-Driven Emittance Degradation in the IR-FEL Driver", JLAB-TN-98-012, 24 March 1998.

[15] Bates end loops consist of only five dipoles. Only one is particularly large – the 180 degree bend – though two of the four smaller bends require large aperture and all demand excellent field quality. Though individual magnet costs may be high, the low element count leads to modest system costs. The design possesses, moreover, symmetries that allow straightforward compaction management through second or even

third order. This is to be contrasted with the cost/performance advantages of other designs. Steffan systems require fewer dipoles (three) but also need multiple quadrupoles. It is difficult to manage compaction beyond first order in this system without introduction of severe geometric aberrations. FODO systems with sufficient symmetry to provide good higher order compaction management (such as second order achromats) require several dipoles (at least four and more typically eight to twelve) and a large number of quadrupoles (eight or more). The loss of operational flexibility in the Steffan system and the large number of elements required for the FODO, offset naïve notions of potentially lower cost.

[16] D. Douglas, "Modeling of Longitudinal Phase Space Dynamics in Energy-Recovering FEL Drivers", *op. cit.*; G. Neil e-mail of 10 November 1998, *op. cit.*

[17] *Per* a requirement of G. Neil and S. Benson.