JLAB-TN-99-002
14 January, 1999

Modeling of Longitudinal Phase Space Dynamics In Energy-Recovering FEL Drivers

D. Douglas

Abstract

An Excel 97 based simulation of single particle dynamics in the IR FEL driver longitudinal phase space is documented and used to analyze operational experience from the Summer and Fall 1998 IR-FEL runs. Extension of the simulation to model potential FEL system upgrades is discussed.

Overview

Single-particle longitudinal dynamics in energy-recovering FEL drivers has been previously addressed analytically [1]. Though progress is being made on simulation of such systems [2], there is at present neither a readily available *design* tool accurately modeling both RF and lattice effects in such systems [3] nor a user-friendly operational model that can be used to quickly analyze control room observations.

Given the utility of other Excel-based optics models during IR-FEL commissioning [4], we have developed a spreadsheet-based model of single-particle longitudinal dynamics in the driver. This model uses a moderately detailed (nonlinear) description of particle dynamics in an RF cavity (described below) and includes both time of flight effects in the transport system and RF phase/amplitude sensitivities. Use of Excel allows the model to reside within a user-friendly interface providing features (such as phase, momentum compaction, and path length adjustment sliders) not dissimilar to those used in the control of the actual accelerator.

This model has been used in an analysis of operational experience from the Summer and Fall 1998 runs. Information about the injected longitudinal phase space is inferred from the model, and an almost certainly fatal flaw (with respect to ability to energy recover with lasing) in operational procedures during the summer run is identified. Corrective actions, and results thereof, are described and documented.

Extension of the model to more complex, higher performance FEL driver systems is discussed. A potential upgrade scenario is evaluated.

Excel 97 Simulation of Longitudinal Dynamics

The model under consideration is based on an integration in time of the energy-time phase subspace as the beam is transported through the machine on the acceleration/energy-recovery cycle. A set of "principal rays", which represent the longitudinal subspace of a bunch, is propagated through a sequence of transformations corresponding to individual sub-sections of the accelerator. This characterizes the effect of the accelerator on particular choices of injected phase space.

The choice of coordinates – wall clock arrival time (in nsec) and energy (in MeV) – is driven by the notion that these are the fundamental observables in the system [5]. In addition to variations *around* an origin, the code thus also tracks the dynamics of the origin *itself* [6]. Relevant parameters for each transformation, typically taken to be operationally available control parameters, can be adjusted to determine their influence on the phase space. The code displays an image of the initial phase space at each of several observation points.

The details of the inputs, the algorithm, and the outputs follow. It may prove useful to the reader to activate the spreadsheet, which resides at

http://www.jlab.org/~douglas/FEL/technote/JLABTN99002A.xls

while reading the documentation.

Inputs: Information necessary for the simulation is introduced through a graphical user interface (GUI, Figure 1). The enumeration of the following list corresponds to the enumeration in the figure.

1. *Initial phase space points* in (time, energy) space – These are "extrema points" of phase ellipse that are input by the user, who gives the maximum energy spread ΔE, bunch length δl and bunch tilt r_{12} at injection. These data, as well as the injection energy, are entered within the indicated cells of the spreadsheet shown in Figure 1. Figure 2 illustrates the interpretation of each parameter. These data will be ray-traced to characterize the evolution of the longitudinal phase space.
2. *Design dimensions for RF cavities* are entered as data within spreadsheet cells on a separate worksheet. These are not normally changed during the course of a simulation.
3. *RF phases and amplitudes* for each cavity and the *gang phase* for the module are introduced as user input variables and readily modified. The GUI presents these values in a manner similar to that used in the EPICS control system.
4. *Path length, energy set-point, and momentum compactions (M_{56}, T_{566}) for upstream optical cavity chicane* (compressor) are input as data in spreadsheet cells of the GUI. They can in principle be changed (they are

just spreadsheet cells after all) but are generally considered to assume fixed design values during any simulation.

5. *Data defining the effect of the wiggler on the phase space* are input as cells in the spreadsheet. These data, detailed in a following discussion of algorithms are the energy spread and the energy centroid shift induced by lasing.
6. *Path length, energy set point, and momentum compactions (M_{56}, T_{566}) (all variable) for recirculation transport* (downstream optical cavity chicane, recirculation arc, and reinjection transport to module) are entered through design set-points (introduced as data in spreadsheet cells, 6a). These values may be readily varied using sliders (6b) similar to those used in the EPICS control system to vary the individual trim quadrupole and sextupole families used in momentum compaction management [7].

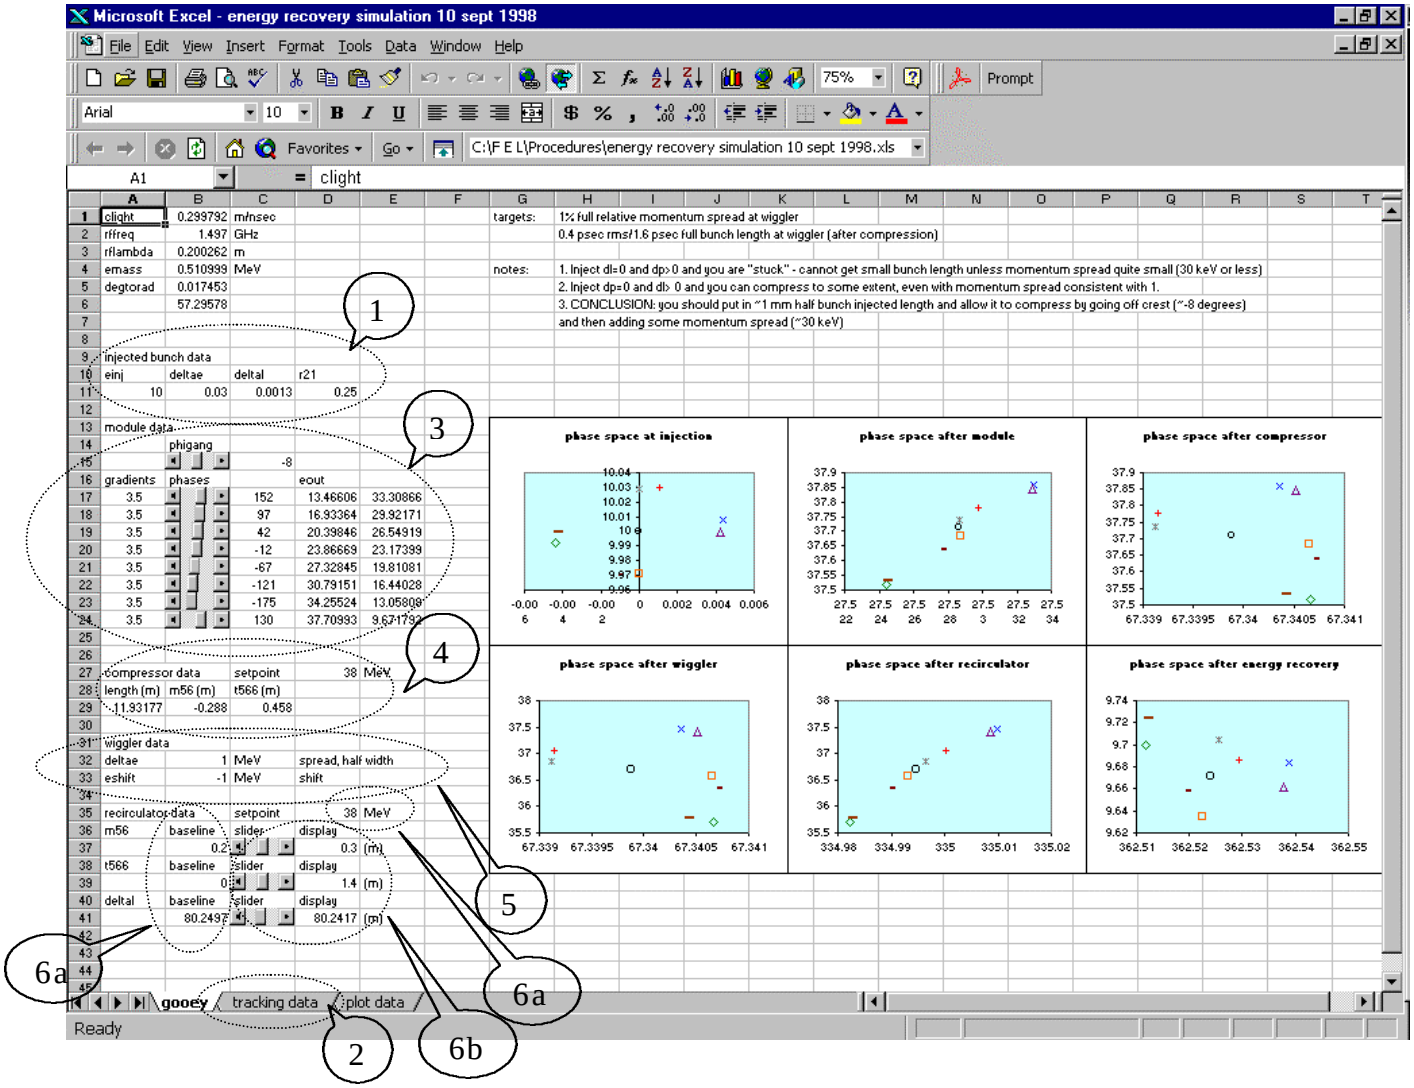

Figure 1: GUI screen shot showing simulation inputs.

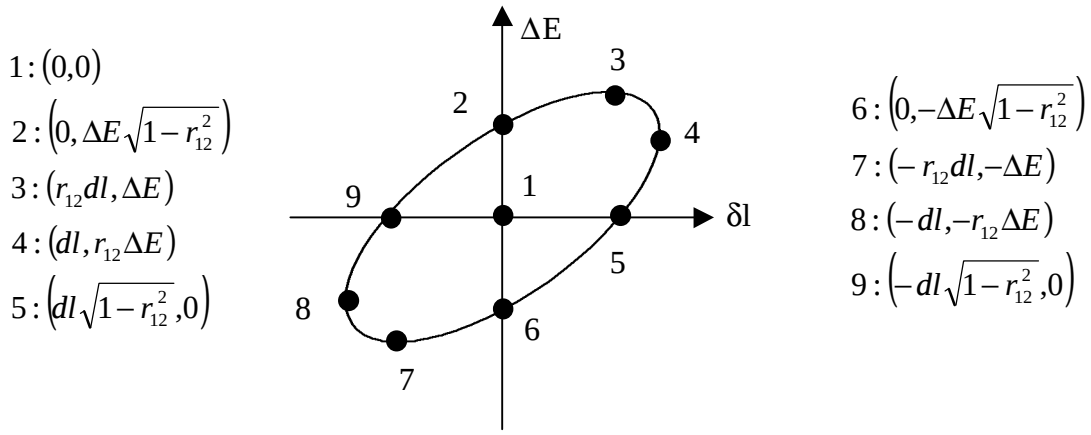

Figure 2: Interpretation of initial ray trace data. The nine data points are ray-traced using the simulation to characterize the downstream longitudinal phase space.

Algorithms – The ray-tracing process for each machine "object" employs algorithms as follow:

1) *Cryomodule* – is treated as eight identical SRF cavities (Figure 3).

2) *SRF Cavity* – is treated as a sequence of finite-length monoenergetic drifts interleaved with impulsive kicks in energy. The energy kicks are intended to simulate the effect on the beam of a single cell of the cavity.
   a) This is a "leap-frog" integration method using 6 drifts and 5 energy kicks per cavity. Figure 3 schematically illustrates the algorithm for the full module.

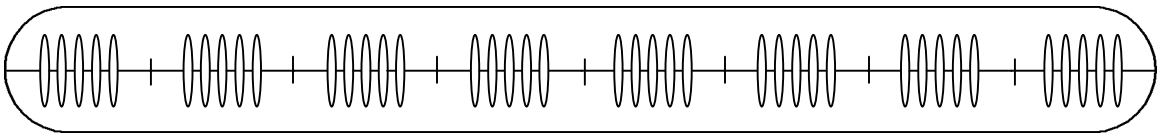

Figure 3: Schematic of integration algorithm. An 8.25 m long cryomodule is modeled as eight identical cavities. Each cavity is modeled as five impulsive (zero length) energy kicks, separated by half-wavelength long drifts; the remainder of the available length is attributed to leading and trailing drifts.

   b) The total cavity length is an $1/8^{th}$ of that of a module (8.25 m) or 1.03125 m.

c) The length of the interior cell-to-cell drifts is $\lambda_{RF}/2$, half an RF wavelength, or 0.100131 m.
d) The length of the exterior (leading and trailing) drifts is half the remainder, or (1.03125 m – 4 x ($\lambda_{RF}/2$))/2 = 0.315363 m
e) The transform for a drift of length *l* is as follows:

$$\bar{t} = t + \frac{l}{v}$$
$$\bar{E} = E$$

Here, *v* is the test ray electron velocity when at energy E; *t* and *E* are the time and energy at incidence, and $\bar{t}$ and $\bar{E}$ are the time and energy at exit.
f) The transform for the nth cell of the cavity is as follows:

$$\bar{t} = t$$
$$\bar{E} = E + \frac{g}{5}\cos(2pt + f_{cavity} + f_{module} + (n-1)p)$$

The dynamical variables are defined as above. The amplitude *g* is the energy gain of a relativistic beam across the full cavity when on crest, ϕ_{cavity} is the individual cavity phase relative to the phase reference, and ϕ_{module} is the module gang phase relative to the phase reference. The $(n\text{-}1)\pi$ summand describes the phase slip from cell to cell.
g) The phases and amplitudes for each cavity and the module gang phase are user-variable; the amplitudes as spreadsheet cell inputs and the phases as sliders (see region 3 of Figure 1 and the associated discussion in the previous section).

3) *Bunch compressor* – the upstream optical cavity chicane, which is used for bunch compression during transport from module to wiggler, is modeled using a second order matrix transformation. The length *l* of the transport from end of module to center of wiggler is 1.93177 m; the compressor is assumed to have the design compaction values of M_{56} = –0.288 m and T_{566} = 0.458 m. The beam line has an energy set-point $E_{central}$; the transformation is thus as follows.
The time and energy variables have the usual meaning; as before, *v* is the

$$\bar{t} = t + \frac{1}{v}\left(l + M_{56}\left(\frac{E - E_{central}}{E_{central}} \right) + T_{566}\left(\frac{E - E_{central}}{E_{central}} \right)^2 \right)$$
$$\bar{E} = E$$

actual test electron velocity at energy *E*.
a) The compressor energy set point, design path length, and compaction values are entered as values in cells of the GUI worksheet (see region 4 of Figure 1). The last three are fixed for the present machine

installation and are typically not changed for any simulation; the energy set point is typically a simulation-specific input and is not changed during the course of the simulation.

4) *The effect of the wiggler* is, in this model, simulated as an induced energy spread and an energy centroid shift.
   a) The user specifies an energy centroid shift ΔE_{shift} and an energy spread half width ΔE_{spread}, both in MeV. These are entered in cells of the GUI worksheet (region 5 of Figure 1).
   b) The model checks through the test rays and finds the ray with the absolute maximum energy deviation from the reference orbit, ΔE_{max}.
   c) For all rays (n=1,2,...,9 in Figure 2) the energy deviations ΔE_n of the nth ray from the reference (n=1) orbit are scaled by $\Delta E_{spread}/\Delta E_{max}$; time offsets are unchanged.
   d) The central energy of all rays are shifted by ΔE_{shift}.
   e) The full transform for all rays may be expressed as follows:

$$\bar{t}_n = t_n$$
$$\bar{E}_n = E_1 + (E_n - E_1)\left(\frac{\Delta E_{spread}}{\Delta E_{\max}}\right) + \Delta E_{shift}$$

for n=1,2,...,9 (as in Figure 2)

This is schematically illustrated in Figure 4.

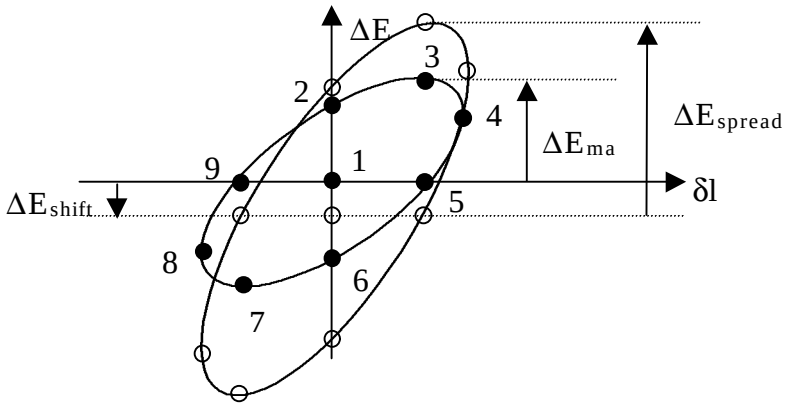

Figure 4: Schematic of wiggler induced energy spread. Spots are original positions, circles are final.

   f) As the longitudinal phase space model used in this simulation is very rudimentary (just a few test rays), no particular details of the actual wiggler dynamics (such as the generation of a bimodal energy

distribution) are included; only the energy positions of the test-rays are modified.

g) Some error trapping is performed to avoid divides by zero. Thus, if all test rays have the same energy, the individual test ray energies are simply shifted by ΔE_{shift}, no scaling is performed.

h) We remark that the energy shift is in fact signed; a negative sign is needed in the input so that the central energy is lower after the wiggler than before.

5) *The recirculator* – is modeled much like the compressor. The length *l* of the transport from wiggler (center) to the reinjection point (module entrance) is 80.24970 m. The recirculator has design compaction values of M_{56} = 0 m and T_{566} = 0 m. These three parameters are input as values in cells of the "gooey" worksheet (region 6a of Figure 1). Unlike in the compressor, these are operationally variable parameters. They may thus be modified in the course of a simulation through use of sliders (region 6b of Figure 1). This is an analog of the EPICS control system, in which these parameters can be modified in real time using, respectively, single families of correctors, trim quads, and trim sextupoles [8]. The total path length and compaction values (baseline and deviation from baseline, as imposed by using the sliders) are then used in the simulation. As with the compressor, the beam line has an energy set-point $E_{central}$ (also entered as a spreadsheet cell value in region 6a of Figure 1). This is normally a case-dependent input to the model, but is not changed during a simulation. The recirculator transformation is thus as follows.

$$\bar{t} = t + \frac{1}{v}\left(l + M_{56}\left(\frac{E - E_{central}}{E_{central}} \right) + T_{566}\left(\frac{E - E_{central}}{E_{central}} \right)^2 \right)$$
$$\bar{E} = E$$

The time and energy variables have the usual meaning; as before, *v* is the actual test electron velocity at energy *E*.

6) *Energy recovery in the module* – is modeled using the cryomodule algorithm described above. The phase of the energy-recovered beam relative to the RF system is defined *via* the recirculator path length.

<u>*Outputs*</u> – the model provides as output images of the longitudinal phase space at various locations along the accelerator. These are phase space plots at

1) injection,

2) after the module during acceleration,

3) at the wiggler center after bunch compression,

4) at the wiggler center after lasing,

5) at the reinjection point following recirculation, and

6) after the module following energy recovery.

The phase space plots for a typical simulation case are clearly visible in Figure 1. The axes are wall clock time in nsec (horizontal) and energy in MeV (vertical).

User interface – has been partially described in the "*Inputs*" section, above. Here we summarize by simply remarking that the model is built in a standard Excel 97 workbook. This contains three worksheets, "gooey", the GUI, "tracking data", on which the actual ray-tracing is performed, and "plot data" where ray-trace data from the tracking data worksheet is reorganized to more conveniently present on the "gooey" worksheet.

As is evident from the previous discussion, the user is primarily interested in the GUI "gooey" worksheet only. It contains both input data cells, ActiveX controls, and output phase space plots presenting all simulation results. If detailed ray-trace data is required, it can be retrieved from the "tracking data" worksheet. For most simulation purposes, the GUI allows the user to interact with the system (at least, with a little imagination) in the way she/he would in the control room.

Using this model, it is possible to study and interpret results of machine operations. We now turn our attention to this activity.

Operational Experience During Summer and Fall 1998

During the summer of 1998, beam was delivered to the wiggler with rms relative momentum spreads of ~¼% and rms bunch lengths of order 0.4 psec [9]. The beam full relative momentum spread and bunch length at the wiggler were therefore about 0.4 MeV and 1.6 psec, respectively. This phase space was assumed to be essentially upright as it was produced by varying the module gang phase to minimize the bunch length as perceived by the Happek device following the wiggler. Acceleration and transport to the wiggler were consequently performed with the module gang phase about 8^{o} off crest. The full relative momentum spread following the wiggler was ~5% when lasing. This corresponds to a 2 MeV momentum spread. Transport from

the wiggler to the 1G dump was performed with "nominal" compactions of M_{56} = 0.2 m and T_{566} = 0 m [10]; the energy recovered beam was roughly 180^o out of phase from the accelerated beam. The resulting spot at the 1G dump was diffuse and poorly focused and transmission through the module was lossy. BLM trips were persistent and beam loss in the module was sufficient to raise He pressure.

Various aspects of this operating scenario are unclear, others are simply incorrect.

1) The bunch length and momentum spread at injection, and any possible correlation thereof (*i.e.* the tilt of the injected phase space), were unknown.

2) During machine set-up it was not considered in detail which side of crest the module was operating on (though this information can in principle be determined from the sign of the gang phase and knowledge of the RF system).

3) It was not entirely clear which side of trough energy recovery was performed on. We do know that the path length offset was rather long, with GASK signals "zeroed" when the path length adjustment string was set to 1400 g-cm. This corresponds to a 40 mm increment to the baseline path length.

4) The magnitude of the recirculator linear momentum compaction was incorrect. It was set to 0.2 m, a value specified assuming acceleration through the module at 12^o off crest (the original design number from *ca*. 1996); for operation 8^o off crest, a 0.3 m value is needed [11].

5) Inasmuch as the location of the beams relative to crest and trough was ambiguous, the sign of the required linear momentum compaction was ambiguous.

6) The quadratic momentum compaction value of 0 m was established using a very rudimentary RF model [12]. It does not take into account any curvature of the RF waveform.

Using the simulation, the "observed" phase space at the wiggler was reproduced by adjusting the momentum spread, bunch length, and tilt of the injected phase space with the injector-to-wiggler transport model configured as an imitation of the machine. The module was modeled as 8^o off crest and the momentum compactions of the compressor were set to design values. When a comparatively long bunch (±1.3 mm or ±4 psec) with a small

momentum spread (±30 keV) and a modest tilt (r_{12}=0.25) is injected and run to the wiggler, the resulting phase space is similar in character to the observed. Variation of the module gang phase away from the 8° set point results in increased bunch length. This is therefore taken as the starting point of the simulation. Results can be seen in the screen shots in Figures 5 and 6.

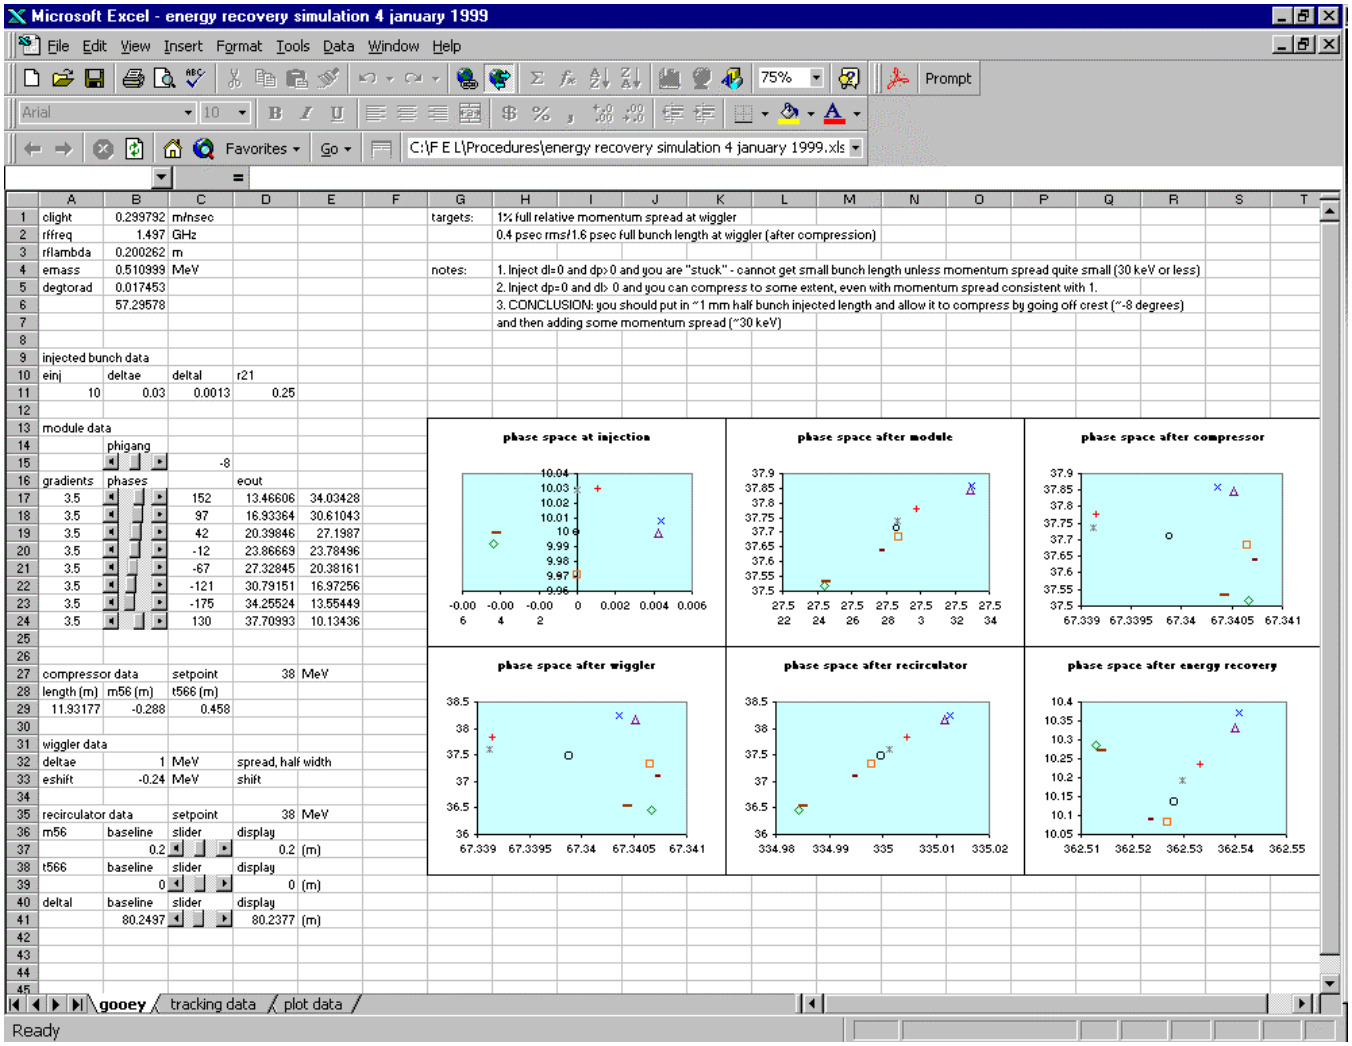

Figure 5: Screen shot of results using Summer 1998 running parameters.

The wiggler is modeled as imposing a 2 MeV (5%) energy spread and a –0.6% (–0.24 MeV) centroid energy shift on the beam [13]. The resulting enlarged phase space is then recirculated and energy recovered. A “best effort” momentum spread at the dump was the obtained by adjusting the recirculator path length. Figure 5 shows the result for the Summer 1998 operating compactions (0.2 m, 0 m). The large final momentum spread of

350 keV (3.5%) is evident. This leads to a horizontal spot size of 3-4 cm at the 1G viewer, where the dispersion is of order 1 m, and would thus be interpreted as providing poor beam quality. Figure 6 shows a result obtained by increasing the linear momentum compaction to 0.3 m and varying not only the recirculator path length but also the second order momentum compaction. The improvement in final momentum spread is evident, with the 2 MeV spread at the wiggler compressed to order 70 keV at the dump.

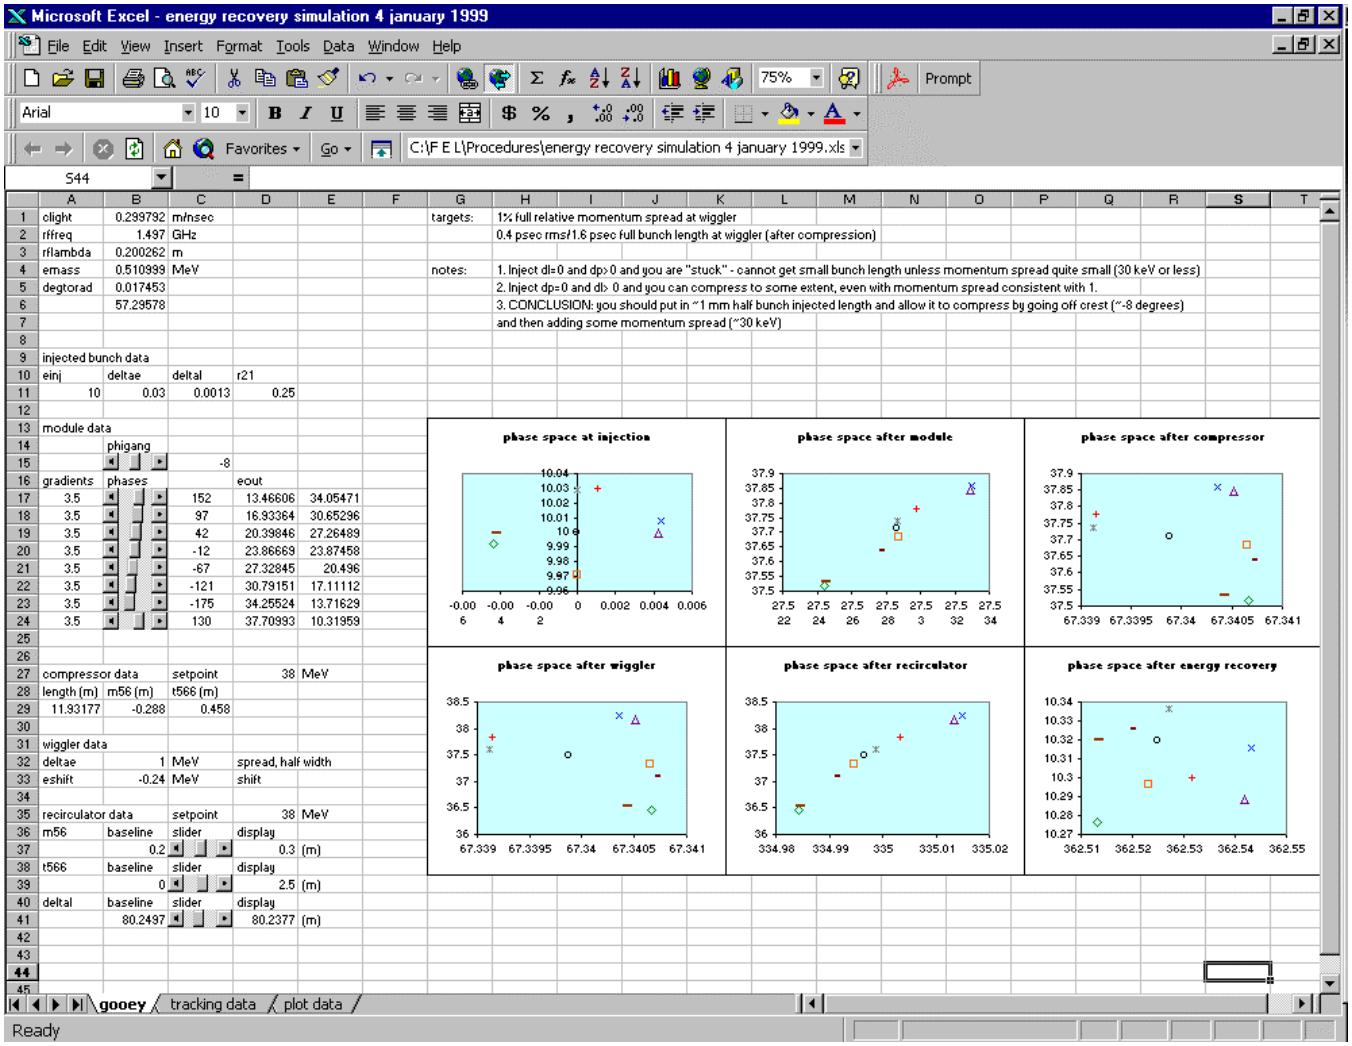

Figure 6: Simulation results with corrected linear momentum compaction and optimized quadratic momentum compaction.

As previously suggested, there is significant quantitative ambiguity in all of this. An assumption was made about the sign of the module gang phase (negative) and the injected beam parameters were varied to fit “observed” results at the wiggler that are not completely firm. We remark that the

choice of negative module gang phase is supported by two factors. First, the module gang phase on the machine tended to run +8^{o} off crest; positive RF system settings historically correspond to modeling by DIMAD and PARMELA with *negative* values. Secondly, with the compressor-generated momentum compactions it was difficult in simulation to produce the observed beam phase space at the wiggler with any even vaguely reasonable injected phase space unless the gang phase was chosen negative. We remark that the compressor compactions are not empirically characterized, but, because they are set by beamline geometry, are probably "close" to design values.

Secondly, the path length set point of the recirculator was far from the expected nominal value – as much as 4 cm or more away from design. It was therefore not entirely clear where deceleration occurred relative to the trough in the RF waveform. The choice of path length was made, operationally, by minimizing four of the GASK signals of the module cavities (#1, 3, 5, and 7). The values of these signals were however not equal to one another during the best machine performance periods (some cavities had higher or lower values than others). Moreover, when readings from the BCM downstream of the module were used to position the recirculated beam 180^{o} out of phase from the accelerated beam, the result suggested that the "desirable" set-point for "best" machine performance was on the other side of trough – the side not giving the best beam performance [14].

These ambiguities led to consequential questions regarding the necessary sign of the recirculator momentum compaction and the location of the recirculated beam relative to the trough. As a result, procedural revisions developed for the Fall 1998 run made allowance for optimization of momentum spread at the dump using the momentum compactions and path lengths as free parameters [15]. The following process was used.

1) Trim quads and sextupoles were set to zero field.

2) The beam was recirculated and the GASK signals minimized by adjusting the recirculator path length.

3) The beam was resteered to the 1G dump

4) With the beam at the dump, one family of trim quadrupoles was excited to give a linear momentum compaction of predicted magnitude 0.3 m. The sign (which turned out to be negative) was selected to give the best horizontal spot size at the dump.

5) With beam at the dump and trim quadrupoles on, one family of sextupoles was activated and adjusted to give the best horizontal spot size at the dump.

6) The path length, trim quad family, and sextupole family were then iteratively adjusted to provide as tight a horizontal spot as possible on the 1G dump viewer.

This process led to a robust machine setup ultimately allowing energy recovery of over 1.5 mA while lasing at nearly 300 Watts. In contrast to the Summer 1998 run, in which the limiting effect was beam loss in the module and at the dump, performance was limited by beam losses in the 5F region of the recirculator. The optimization process in fact examined beam behavior on both sides of trough, for all combinations of sign of quadrupole and sextupole excitation. The resulting trim quad excitations were opposite in sign and about 2/3 the magnitude predicted by the simulation. The sign difference is potentially attributable to running on the "wrong" side of trough, where the slope is opposite in sign to the design value. Recall that BCM observations suggest that we are in fact in this situation. The resulting sextupole excitations were about ¾ the magnitude, and the same in sign, as those predicted by simulation. This is reassuring, as the curvature of the RF waveform *is* the same for both potential operating points even though the slope is opposite in sign.

We note the FELEXT string was running about 6% above the value of FELINJ, indicating higher final beam energy than expected. Conversely, reduction of BLM hits at the dump often required significant horizontal steering with MDB1G00H to the *left*, in the direction of *lower* energy. Moreover, the definition of "best" horizontal spot was decidedly heuristic, as no quantitative measurements of spot size were made. In addition, the dump quadrupole upstream of the viewer was set to somewhat arbitrary values (ranging from –1020 g to 1020 g, depending on the stage of the process, the mood of the operator, and the phase of the moon) to assist in steering to, and keeping the beam visible on, the 1G dump viewer. Thus for example, the dispersion at the 1G viewer was typically not known. Basically, we got a fairly good spot at the dump, but had no idea of how it was *supposed* to look. Given these ambiguities, as well as the various phasing issues, the agreement between simulation and observation at the level of a few tens of percent is rather encouraging.

These results form part of the basis of planning for the Winter 1998-99 run. We will attempt several tasks at the outset of the run, including the following.

1) Establish if BLM-induced trips are due to an RF instability or scraping:
   a) Reduce faulting BLM head voltage by 100 V and see if more current can be run. If so, limits are probably due to scraping, if not, limits are due to an instability [16]; assuming scraping, proceed as follows; if there is an instability, lower flag, turn out lights and lock door while leaving.

2) Determine if scraping is due to halo or recirculator aperture limitation of the core of the beam. If limit is due to scraping of the core, proceed as follows; if it is due to halo, skip to #10).

3) Resteer the injected beam to improve the match between excitation of the final DU magnet and the beam energy.

4) Improve steering through the 5F region to
   a) center in all trim elements and
   b) improve accuracy of MDY5F02 shunt settings
   and, as a consequence, reduce BLM hits at reinjection.

5) Check/improve dispersion handling in recirculator.

6) Improve steering through the reinjection transport in an effort to better center the recirculated beam through the module during energy recovery:
   a) reacquire the energy recovered beam on the 1F02 viewer and
   b) reduce the impact of HOM-generated and lattice skew quads.

7) Establish improved steering from 1F02 to dump using new viewer near the dump face.

8) Improve spot at dump using dump line quads and new viewer at the dump face.

9) Improve momentum compaction management by quantifying behavior using newly installed BCM cavities.

10) Utilize 3F scraper to reduce horizontal halo/tails; use “halo-suppression optics” [17] to reduce vertical halo/tails, then iterate steps 3)-9).

These tasks individually should improve beam quality and transmission and may collectively allow operation of the FEL at significantly (3 – 4 times) higher power. Even higher powers (10 – 20 kW) will almost certainly require significant modification of the existing machine, including the addition of cryomodules for higher final energy, reconfiguration of FEL and beam transport systems, and extensive source upgrades. Such an upgraded or

rebuilt machine will reside in a different region of parameter space and will have unique longitudinal dynamics issues. We now turn our attention to such topics.

Evolution of the Model for FEL/Driver Upgrade

Upgrades of the IR FEL Demo to higher power and shorter wavelength are presently under consideration [18]. Some upgrade scenarios envision increasing both the machine energy and FEL extraction efficiency, a combination that will result in a much greater potential relative energy spread at the energy recovery dump than is presently encountered.

Recall that in the IR Demo, we use a 40 MeV beam and ½% extraction efficiency, resulting in ~5% relative momentum spread, or an absolute energy spread of 2 MeV after the wiggler. This absolute spread would represent ~20% relative energy spread at the dump, were it not for the 20-to-1 energy compression accomplished during energy recovery. In contrast, a 200 MeV high power IR driver extracting about 1% of the beam power could give an 8% relative momentum spread after the wiggler, or 16 MeV absolute energy spread. This, if energy recovered to an injection energy of 10 MeV, would lead to intractably large momentum spreads and intolerable beam loss. If a 20-to-1 energy compression were achieved, the resulting final energy spread would still be 0.8 MeV, or, at 10 MeV, a rather large 8% relative energy spread.

Management of such large relative momentum spreads may be difficult; implementation of more effective energy compression methods may be even more difficult. A direct solution using an “afterburner” cryounit for bunching immediately after the FEL has been proposed by G. Neil [19]. The desire to evaluate the feasibility of this and other alternatives has led us to extend the Excel longitudinal matching model described above to the case of a three-module energy-recovering machine.

Machine Concept – We consider a direct upgrade of the IR Demo to a fairly similar 200 MeV accelerator. The simulation has been extended to model a three-module linac that accelerates a 10 MeV injected beam to 200 MeV. The beam is transported from the linac to the wiggler through an arc that is envisioned to allow tuning of both first and second order momentum compaction so as to optimize bunch length compression at the wiggler. The wiggler is modeled much as in the previous simulation, specifically, as a device that induces large momentum spread but has no impact on bunch length. The beam, with enlarged momentum spread, is then returned to the linac for energy recovery through an arc that allows for adjustment of first, second and third order momentum compaction. These three variables, together with the path length of the complete recirculation transport, are available as free parameters for control of momentum spread at the energy

recovery dump *via* energy compression. The purpose of the specific exercise described below is to examine the compression efficiency available using bunch length manipulation following the wiggler and off-trough deceleration during energy recovery.

User Interface – The simulation user interface is shown in Figure 7 and is quite similar to that used in the IR Demo model. ActiveX controls for the compactions of the linac-to-wiggler transport have been included, as has a control for the third order compaction of the wiggler-to-linac transport. All these may be operationally useful in the upgraded machine. The location of the controls has been modified, but the worksheet is sufficiently similar to the original interface that the layout should be obvious. The reader may wish to open a copy of this model; it resides at

http://www.jlab.org/~douglas/FEL/technote/JLABTN99002B.xls

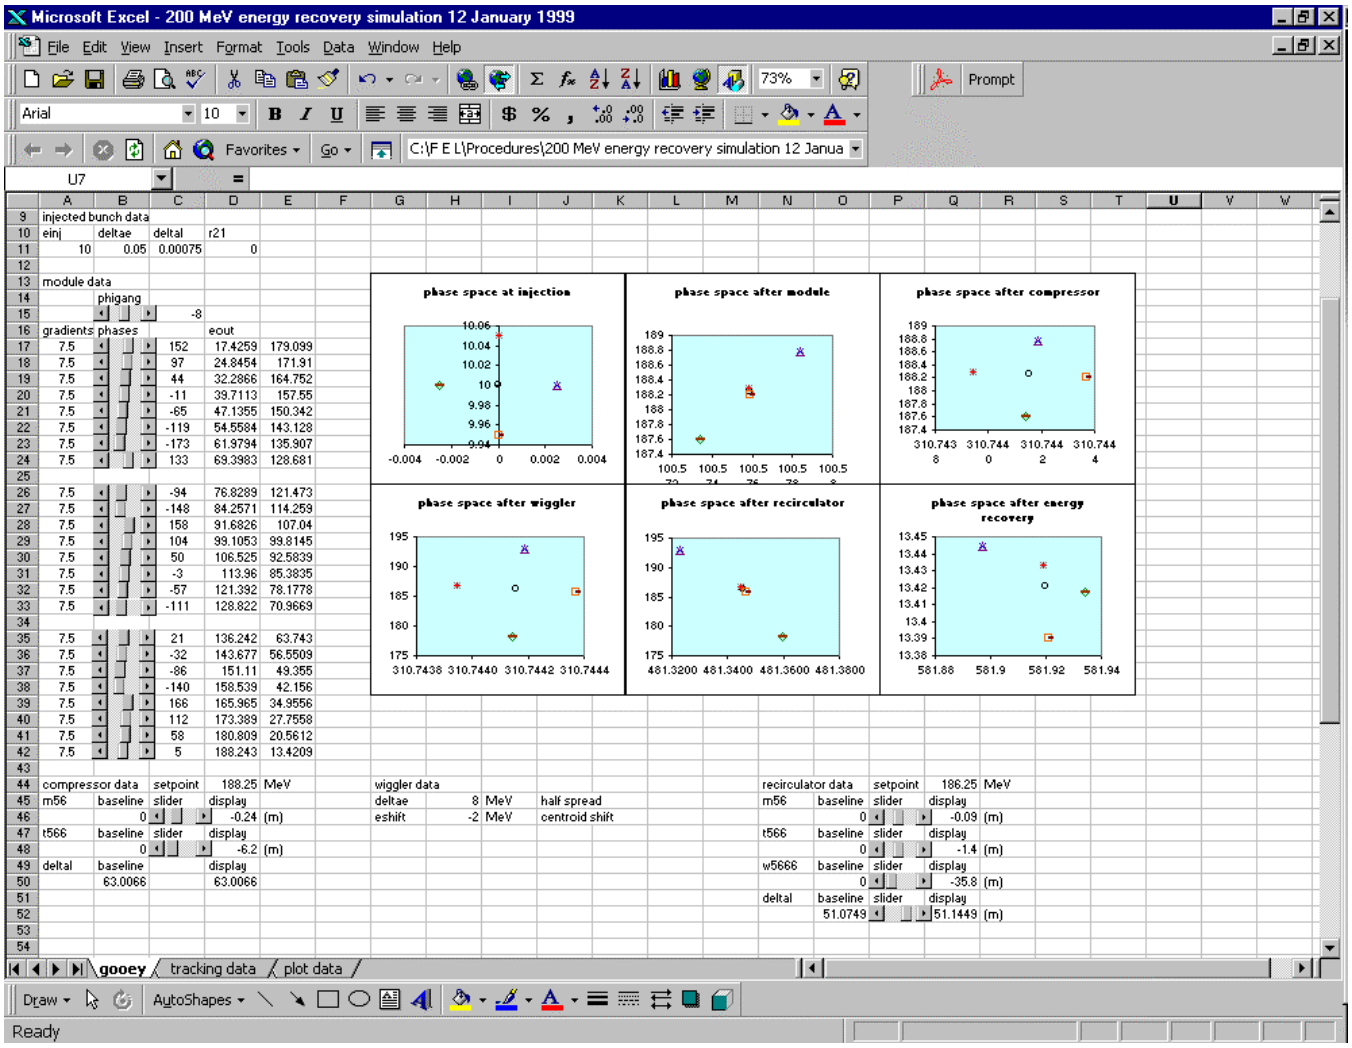

Figure 7: User interface for spreadsheet model of 200 MeV upgrade.

Sample Problem – we have used the extended spreadsheet model to simulate energy recovery in the upgraded machine configuration discussed above. The following comments are relevant.

1) To achieve a "best effort" phase space at the wiggler, some changes from the IR Demo injected parameters were required.
   a) The 10 MeV beam was injected with about the same longitudinal emittance, but with 50 keV energy spread/0.75 mm bunch length, instead of the IR Demo values of 30 kEV/1.3 mm.
2) Acceleration again occurred 8^o off-crest
3) The linac to module transport was used as a bunch length compressor; optimization was possible by adjustment of both linear and quadratic momentum compactions.
4) The wiggler was assumed to impose a 2 MeV (1%) energy shift and 16 MeV (8%) energy spread.
5) The recirculator energy recovery transport (wiggler to linac) had available linear, quadratic and cubic (octupole order) momentum compaction adjustments, which, with path length, were used to optimize the momentum spread at the end of the system.

With these methods and constraints, we find the results presented in Figures 7 and 8. The following comments apply.

1) As noted above, optimization of bunch length and momentum spread at the wiggler required some changes in the injected beam parameters. The longitudinal emittance, however, remained the same.

2) Transport to the wiggler used a linear momentum compaction of –0.24 m (similar but a bit smaller than the IR Demo value) and a quadratic momentum compaction of –6.2 m (quite different from the IR Demo value). The latter change was presumably to accommodate changes in the acceleration dynamics over a larger energy gain, with greater associated curvature.

3) With the aforementioned values, a full bunch length of about 1/2 psec was generated at the wiggler, with a full momentum spread of about 1.4 MeV. See either Figure 7 or 8.

4) The recirculation transport, using only quadrupole and sextupole order compaction control, could compress the 16 MeV momentum spread out of the wiggler to a bit under 1 MeV (Figure 8). This is consistent with the 20-to-1 compression efficiency observed in simulation of the IR Demo and

during IR Demo operations. This is a very large relative momentum spread, which would, as noted above, represent an operational problem.

5) The recirculation transport, when using compaction control through third order, could compress the initial 16 MeV energy spread to a bit under 70 keV. This is probably operationally acceptable.

6) In either case (but particularly in the latter) the configuration providing the "best" compression
   a) was rather difficult to find
   b) resided at slightly over 13 MeV in the cases presented above (unpleasantly close to the neutron production threshold?)
   c) was very sensitively dependent on the detailed values of the compaction parameters (even small changes in compaction or path length would lead to large changes in final momentum spread).

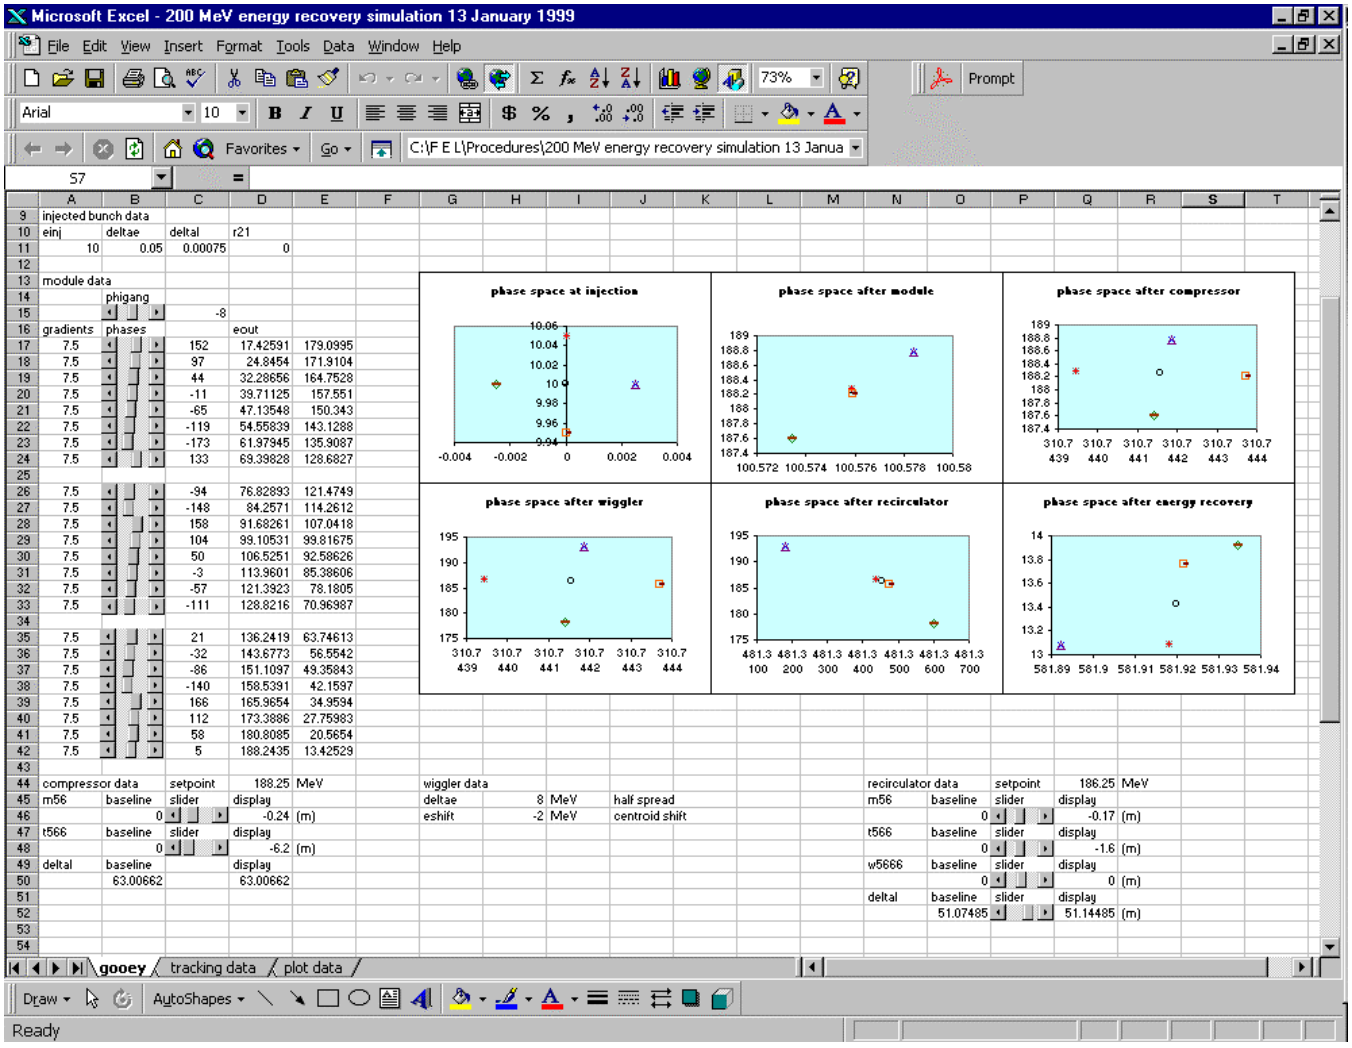

Figure 8: Upgrade simulation using only M_{56} and T_{566} control.

We therefore conclude that energy compression using compaction manipulations and off-trough energy recovery is possible and may be used to produce operationally acceptable results. We do however strongly suspect that it is a parametrically sensitive process that will have attendant tight tolerances. It may be operationally difficult to achieve and may not, without some effort, reproduce well. It will, at the very least, require careful characterization of the accelerator. Any design should allow for detailed measurement and correction through octupole order, suggesting that early consideration be given to robust diagnostic and tuning schemes such as that developed by Krafft [20]. Furthermore, careful consideration should be given to providing alternative momentum management technologies, such as that proposed by Neil [21].

Conclusions

A simulation of single-particle longitudinal dynamics in energy recovering FEL drivers has been documented. This simulation uses a moderately detailed cavity model and includes nonlinear chromatic effects due to beam transport. It is embedded in a spreadsheet providing an intuitive, user-friendly interface.

Based on this simulation, we infer that the phase space following the FEL injector during the summer 1998 run likely had small momentum spread and long bunch length so as to provide the observed (rms) 0.4 psec bunch length and ¼% (rms) momentum spread at the wiggler. Given this observed phase space at the wiggler, use of the simulation to develop an energy compression/energy recovery scenario is described. From this scenario we conclude that Summer 1998 attempts to operate the machine in energy recovery mode were almost certainly made using the wrong recirculator momentum compaction, resulting in very large momentum spread during energy recovery. Revised values for linear and quadratic momentum compactions were evaluated, and a tuning scenario developed and implemented [22]. The implementation was successful inasmuch as cw FEL power and energy-recovered current were raised, respectively, to about 300 W and 1.6 mA during the Fall 1998 run.

The simulation can be readily extended to model upgrades of the IR FEL to a multi-module UV/high power IR capability. Results from a study of a 200 MeV UV upgrade indicate that energy compression is possible during energy recovery. However, performance targets achieved in the IR Demo are difficult to meet in preliminary studies using possible parameters for a UV/high-power IR machine. Successful machine operation will probably require matching to quite high (octupole) order [23] and/or may not be possible

without additional phase space manipulations of the type recently proposed by George Neil [24].

Acknowledgments

I would like to thank Jean Delayen for useful instruction on the description of RF cavities used in this simulation; his insight made this work possible. I would like to thank the FEL group for gracefully suffering though a presentation on this work in the Fall of 1998, and for much valuable feedback that led to the development of an improved high power setup procedure.

Notes and References

[1] D. Douglas, A. Hutton, C. Leemann, and D. Neuffer, "Analytic Modeling and Lattice Scaling Relations for FEL Driver Accelerators", CEBAF-TN-95-015, 21 March 1995.

[2] Byung Yunn has put some second-order effects into PARMELA and is evaluating both the code evolution and the significance of the physics in the IR FEL driver.

[3] Most optics design programs do not model acceleration at all. DIMAD models linear acceleration effects and second (or higher) order lattice effects in all phase space dimensions, but cannot at present directly simulate effects due to the curvature of the RF wave form.

[4] D. Douglas, "Matching Solutions for Various Wiggler Operating Configurations", JLAB-TN-97-040, 27 October 1997; D. Douglas, "An Alternate Paradigm for High-Level Application Development", JLAB-TN-97-047, 12 December 1997; D. Douglas, "Suppression and Enhancement of CSR-Driven Emittance Degradation in the IR-FEL Driver", JLAB-TN-98-012, 24 March 1998; D. Douglas, "How to Use Trim Quads to Modify Dispersion", JLAB-TN-98-035, 28 August, 1998.

[5] Thanks to Jean Delayen for much general instruction on appropriate description of RF cavities and for specifically pointing out that you can always establish a phase (time) reference using a master oscillator.

[6] This is a feature of the model distinguishing it from nearly all standard optical modeling codes, which address behavior using expansions around a “design orbit” of assumed ideal behavior and which assume this origin of phase space is mapped onto itself. This common technique simplifies and speeds up the computation of lattice and beam properties, but precludes easy modeling of acceleration, which, by its very nature, is a translation in phase space!

[7] D. Douglas, "Lattice Issues Affecting Longitudinal Phase Space Management During Energy Recovery, Or, 'Why the FEL Needs Sextupoles'", JLAB-TN-98-025, 9 July 1998.

[8] *ibid*.

[9] These very approximate numbers evolved from conversations with Geoff Krafft and Steve Benson.

[10] D. Douglas, JLAB-TN-98-025, *op. cit.*

[11] D. Douglas *et al.*, CEBAF-TN-95-015, *op. cit.*

[12] D. Douglas, JLAB-TN-98-025, *op. cit.*

[13] Steve Benson, private communication.

[14] Phillipe Piot, FLOG entry #3540, 10 December 1998.

[15] D. Douglas, “High-Power Setup Procedure” Revision 2.1, 21 September 1998; this is available on-line at http://www.jlab.org/~douglas/FEL/procedures/onekwsetuprev2_1.pdf

[16] G. Neil, private communication.

[17] D. Douglas and G. Biallas, unpublished lattice solution, 8 January 1999.

[18] See, for example, published viewgraphs from the “FEL Upgrade Brainstorming Meeting”, October 22-23, 1998.

[19] G. Neil, e-mail message of 10 November 1998 describing an active energy compression scheme using an “afterburner” cryounit.

[20] G. A. Krafft, “Analyzing Measurements of Nonlinear Transfer Functions With Tschebyshev Polynomials”, to appear in the Proc. of the 1997 IEEE Particle Accelerator Conference, Vancouver, B. C. (May, 1997).

[21] G. Neil, email *op. cit.*

[22] D. Douglas, “High-Power Setup Procedure”, *op. cit.*

[23] G. A. Krafft, *op. cit.*

[24] G. Neil, email *op. cit.*