L C L S M e m o r a n d u m

Date: November 11, 2005

To: Heinz-Dieter Nuhn, John Galayda, LCLS

From: William M. Fawley, LBNL

Subject: Final Report: FY2005 LCLS-Supported Work on the GINGER FEL Simulation Code

---

This report documents the work done under the FY2005 60K$ contract between SLAC/LCLS and LBNL, whose major purpose was to improve and extend the GINGER free-electron laser numerical simulation code. The agreed-upon specific tasks were as follows:

Task 1: Develop and implement an algorithm for GINGER to compute and propagate higher harmonic emission from a microbunched electron beam including effects of refraction and diffraction. In the initial formulation, ignore effects of harmonic radiation back upon beam electrons --- this approximation should be quite accurate for the nonlinear harmonic emission expected up to nominal saturation from LCLS-like FEL's. Extend the post-processor to analyze the emission in parallel to that done for the fundamental.

Task 2: Modify the I/O capability of GINGER and its postprocessor to use the HDF5 format employed by many advanced simulation codes in the physics and astrophysics community.

Task 3: Extend the wakefield model of ginger to include the so-called AC resistivity component for both Cu and Al vacuum chamber wall material. Examine the effects upon LCLS emission using full time-dependent SASE simulations employing the most recent particle distributions calculated by Emma et al. Also examine the amelioration of wake effects using simple linear undulator strength tapers.

Task 4 (added in mid-contract): Analyze the usefulness of modifying the break lengths in the first few LCLS undulator break sections suggested by Vinokurov and collaborators at ANL. In particular, use GINGER to examine whether the predictions apply to the SASE regime for current LCLS parameters.

We now present some details on the work accomplished under this contract; additional details are included in Appendices A&B.

Task 1 - The GINGER code was originally designed in the mid-1980's to study various non-"time-steady" (i.e. polychromatic) effects common to single pass FEL amplifiers. The code employs the eikonal approximation which examines non-time-steady behavior in a relatively small frequency window centered upon the fundamental FEL resonance wavelength, $\lambda_s \sim \lambda_w (1+ a_w^2)/2\gamma^2$. Another code (NUTMEG) was derived from the original FRED code to study harmonic emission growth; however, this code employed the time-steady (monochromatic) assumption which prevented it from studying harmonic emission in relatively broad band FEL's such as those based on the SASE process. Due the lack of polychromatic ability and the fact that development on the NUMTEG code ceased in the late 1980's, and because to the contrary GINGER has been steadily upgraded to the present (including the ability to run efficiently on massively parallel hardware such as the NERSC IBM-SP), it has been felt that the most effective way to simulate harmonic emission in SASE devices is to add such ability to GINGER.

The FY2005 contract between the LCLS group and LBNL had as its central task this particular upgrade. To do so, we believed with an eye on future upgrades (*e.g.*, adding full 3D field solver capability to the exist full 3D particle mover) that the most robust approach was a complete rewrite of the existing field solver and the nearly complete replacement of the output file format by use of the modern HDF5 library. We now give some details on each.

From its original birth in the mid-1980's, GINGER relied upon an elaborate predictor-corrector field solver based on the so-called Gear scheme. This scheme continually updates the predicted field and its first two *z*-derivatives as the field is advanced in z. While this method is extremely accurate and efficient in terms of asymptotic step size for well-behaved problems such as simple vacuum propagation, the necessary coding to handle time-dependent effects related to slippage was very complex and difficult to extend to higher "dimensions" such as harmonics and full 3D geometries. It also requires a nearly 6-fold increase in required memory size to hold the field and its derivatives. While this increase is not important for the "original" GINGER which typically only has ~1K 64-bit words to hold a given longitudinal slice's field arrays, the necessary memory requirements could grow nearly 100-fold or more for a full 3D problem with multiple harmonics. Since some problems involve 1000 or more slices, this amount of memory begins to become non-trivial, even on modern CPU's.

Due to these reasons and to improve future maintainability, we felt it was more reasonable to completely replace the old field solver with a much simpler one. The new one is currently based upon a simple implicit method described by Richtmeyer and Morton in their classic book. The necessary coding for the new solver is at least 10× smaller in size and does not require any field "history" arrays. Furthermore, the new coding fully embraces the actual representation of the field as being complex; this also halves the number of necessary code lines (although there is probably no actual speedup in terms of CPU operations). The addition of harmonic emission was completely straight-forward; probably the most complex requirement was increasing the dimensionality of the field arrays and making sure such was handled correctly in the non-field solver subroutines in GINGER. In particular, the "message-passing" routines which apply on

massively parallel environments must pass the full field arrays including the harmonic portion.

Initially, the implicit solver is fully "backward-biased" for maximal numerical stability; this leads to some unphysical vacuum energy loss (typically ~1 part per thousand per Rayleigh length for step sizes $\Delta z \sim 0.001\ Z_R$). This loss is more irritating than truly problematic for high gain SASE devices such as the LCLS. Recently, we changed the biasing factor to 0.5 (which makes it equivalent to the Crank-Nickolson algorithm) to reduce this loss. For step sizes ~0.01 Z_R the numerical inaccuracies for vacuum propagation are now dominated by radial grid resolution. we may also examine other solvers (*e.g.*, Runge-Kutta schemes) to see if greater accuracy is possible without any loss of numerical stability.

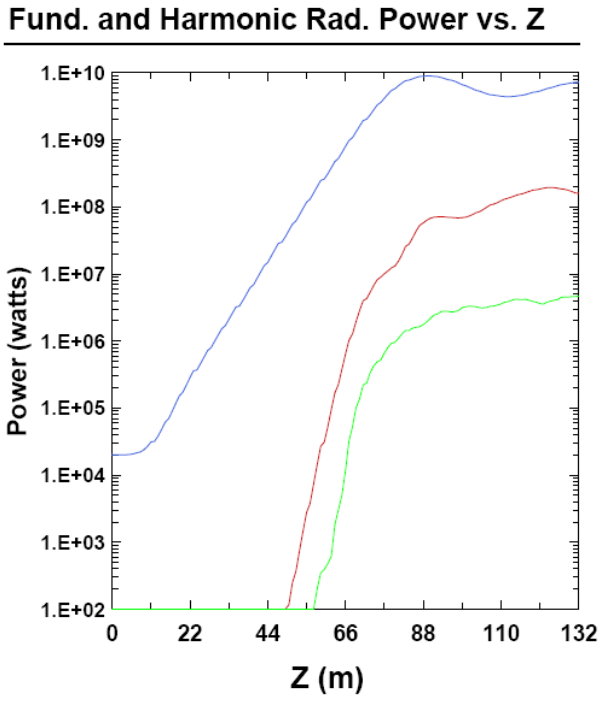

Figure 1: Harmonic GINGER run of a standard LCLS undulator lattice for a 2 kA input beam current (as would correspond to the 0.2-pC "El Chargito" low charge case). The blue curve refers to the .fundamental, the red to the third harmonic, and the green to the fifth harmonic radiation power. The integrator step size was 5.0 mm; the Rayleigh range of the fundamental was of order 33 m.

In addition to changing the field solver, we also completely replaced the longitudinal particle mover that previously was also (unnecessarily for nearly all problems) based upon a Gear-scheme predictor-corrector. To quickly bring up a reasonable mover, we based one on a simple 2nd-order Runge-Kutta algorithm. Moreover, we limited the interaction with the field (via the $d\gamma/dz$ and $d\theta/dz$ equations) to just that involving the fundamental components (and not the harmonics). While this fully captures the growth of so-called "non-linear" harmonic emission due to strong microbunching caused by the fundamental field components, it neglects the "linear" growth components for harmonics. These growth rates are typically 10× or smaller for high gain devices such as the LCLS. At some point in 2006, we expect to add the option of including the full interaction with the harmonic components. This will likely increase the running time by ~2.5× for problems including the 3rd and 5th harmonics. We will also add the option to use a higher order longitudinal particle mover (*e.g.*, 4th order Runge-Kutta) for those users requiring extreme accuracy (although for SASE problems such accuracy in the particle mover cannot make up for the inaccuracies induced by limited macroparticle number in each longitudinal slice).

Figure 1 above displays the predicted power at the fundamental together with that at the 3rd and 5th harmonics for a low charge LCLS case. This was a time-steady case with 2-

kA of beam current. One sees that the 3rd harmonic power reaches about 2.5% of the fundamental. The equivalent number for the 5th harmonic is slightly below 0.1%. This example was run at NERSC and was reduced with the new post-processor which can handle HDF5 format.

Task 2 – A major effort was made to replace the old ASCII-format GINGER output "pltfile" with one based on the modern HDF5 format, in wide use in the advanced simulation community. HDF5 employs a UNIX-like tree file structure in which the user may place a wide variety of data types including multidimensional arrays, string variables, text data, etc.. The user may also define more complicated heterogeneous record structures if needed. Because both ASCII and binary information can be stored in the same file, HDF5 can be very efficient in terms of disk space requirements and in readability. HDF5 is supported by an active group at NCSA and the library has been ported to nearly all UNIX variants including Linux, MacOS-X, and Windows. As of fall 2005, the HDF5-formatted output file was supported in GINGER for both Linux , Sun-OS, and AIX-OS on the NERSC IBM-SP. We also successfully implemented parallel HDF5 I/O which is needed for efficient use of the IBM-SP multiprocessor capability.

The structure looks like the following using the h5dump routine from the HDF5 library:

```
h5dump -n palSDh.hdf5
     HDF5 "palSDh.hdf5" {
     FILE_CONTENTS {
      group      /base_param
      dataset    /base_param/int_param_buf
      dataset    /base_param/real_param_buf
      group      /grids
      dataset    /grids/3Dfld_zgrid
      dataset    /grids/rgrid
      dataset    /grids/scalar_zgrid
      group      /input
      dataset    /input/input_file
      dataset    /input/pltfile
      group      /particles
      dataset    /particles/env_data
      dataset    /particles/info_buffer
      dataset    /particles/scalar_data
      dataset    /particles/spcdata
      group      /radiation
      dataset    /radiation/fund_r-z-t_data
      dataset    /radiation/harm_r-z-t_data
      dataset    /radiation/scalar_data
      }
     }
```

One can see that HDF5 "groups" correspond to UNIX directories and "datasets" correspond to files. There is no difficulty in reading a dataset "deep" in the interior of an HDF5 file --- the read utilities automatically calculate the correct offset in the file. Such is not easy to implement in a binary file (and essentially impossible in an ASCII-formatted file) and one normally needs to read through all the previous information in a file to get at the wanted portion.

We also implemented a very useful suggestion from H.-D. Nuhn that the HDF5 output file also store the input file (and when present, the template file and tapered wiggler "bwfile") which thus allows the interested user in the distant future to see exactly what input was given to the simulation code Upon close inspection of the above structure, one sees that the data is arranged such that particle diagnostic data such as microbunching and 6D phase space coordinates (used for macroparticle scatter plots such as x vs. y) is stored separately from radiation data. Within the radiation data, there are "scalar" quantities such as instantaneous power which are independent of the radial coordinate and vector information (*e.g.*, radiation intensity) that depend upon r. For large multislice time-dependent runs, the HDF5 files is typically 3-4× smaller in disc usage than the old ASCII format. For reading by the post-processor code, there is also a significant speedup since the data is read as large binary arrays as opposed to single line reads; this particularly applies to the 3D (r,z,t) radiation information.

We note that the GENESIS code of S. Reiche also has the capability to output information in HDF5 format (thus again proving "great minds think alike" hypothesis). In a more perfect world, it would also be nice if the ELEGANT tracking code implemented an option of HDF5 format and/or an SDDStoHDF5 utility was written.

Task 3 --- A significant amount of effort was spent during the course of the contract (and actually somewhat before) implementing and studying the effects on FEL coherent light production from the "AC resistivity" term on wake losses in the LCLS undulator vacuum chamber. In collaboration with SLAC's H.-D. Nuhn, the AC term was added to the wakefield calculation stand-alone routine whose output is used by GINGER. A great number of full "S2E" SASE-initiated LCLS simulations were done for both copper and aluminum vacuum chamber wall material. Furthermore, we also studied two particular e-beam cases: the standard 1-nC case and also a reduced charge case (~0.2-pC) initially championed by Paul Emma and now by many of us. We have placed the full FEL05 conference paper on this subject as Appendix B to this contract report. Within it, one can examine both the details of the expected phase space entering the undulator and the effects of the wake losses for various compensating tapers.

Task 4 --- In late winter 2005 we were asked to explore with the GINGER code the usefulness of modifying the first three break lengths of the LCLS undulator lattice as suggested by collaborators at ANL. A number of detailed, full SASE runs were done and written up in a note submitted to SLAC in later March 2005. This note is reproduced in Appendix A.

Appendix A

LCLS Output Performance Sensitivity to Drift Length Detuning

Original date: 29 March 2005

R. Dejus and N. Vinokurov of ANL have suggested that it may be possible to increase the LCLS radiation output by increasing, by set amounts, the first three break lengths in the LCLS undulator lattice. Using the steady state code, RON, they found that increments of +45 mm, +20 mm, and +5 mm, respectively, seemed to better couple the exponential growth in the early part of the undulator to the "asymptotic" mode dominating in the last half of the undulator. Inasmuch as RON is a linearized code and cannot examine polychromatic and/or saturation effects, H.-D. Nuhn and Roger Dejus suggested that reexamining this suggestion with the GINGER simulation code would be worthwhile.

To study the problem, GINGER was slightly modified to allow the user via the input file to adjust individual break lengths in the overall periodic lattice. GINGER previously had an input "switch" that allowed the user to specify that the nominal longitudinal phase advance will automatically be an integral multiple of 2π, irrespective of whether the actual break length l_D satisfied the resonance relation $l_D=\lambda_W\,(1+K^2/2)$. Here K is the undulator parameter (presumed linearly polarized) and λ_W is the undulator period. The new input parameter applies an additional shift of $2\pi\,\Delta l_j/\,(\lambda_W\,(1+K^2/2))$ where j refers to the index number of the break length. The ANL-suggested increments thus correspond to 1.32, 0.59, and 0.15 radians, respectively.

We ran two types of problems. The first was a series of purely monochromatic (*i.e.*, steady-state"), single-pass amplifier runs in which 30 kW of input power was introduced at the beginning of the undulator. For a "nominal" undulator with constant and phase-corrected break lengths, the undulator strength parameter K was optimized for peak power at z=60 m before saturation. We then studied the effects on the output power at this point by introducing the break length adjustments suggested by Dejus and Vinokurov. The results indicated that the nominal situation actually performed better.

The second set of runs was more characteristic of a true SASE situation and was done in polychromatic, time-dependent mode. At $z=0$ we introduced a broad-band, white noise input radiation spectrum with $P(\omega)$ constant over the simulation frequency window. This allowed one single run to give a good indication of what would be true when averaged over a broad ensemble of individual shots starting from noise. In the time domain, the resultant $P(t)$ is extremely spiky as there is no phase correlation between the different individual $P(\omega)$ components. GINGER, at present, does not have the same capability to start shot noise with a constant, broadband $b(\omega)$ spectrum. Nonetheless, we are reasonably confident that radiation noise should give a similar result to microbunching shot noise.

We did 3 distinct runs: (a) nominal break lengths with exact phase correction; (b) first three individual break lengths adjusted by the above mentioned values as suggested by

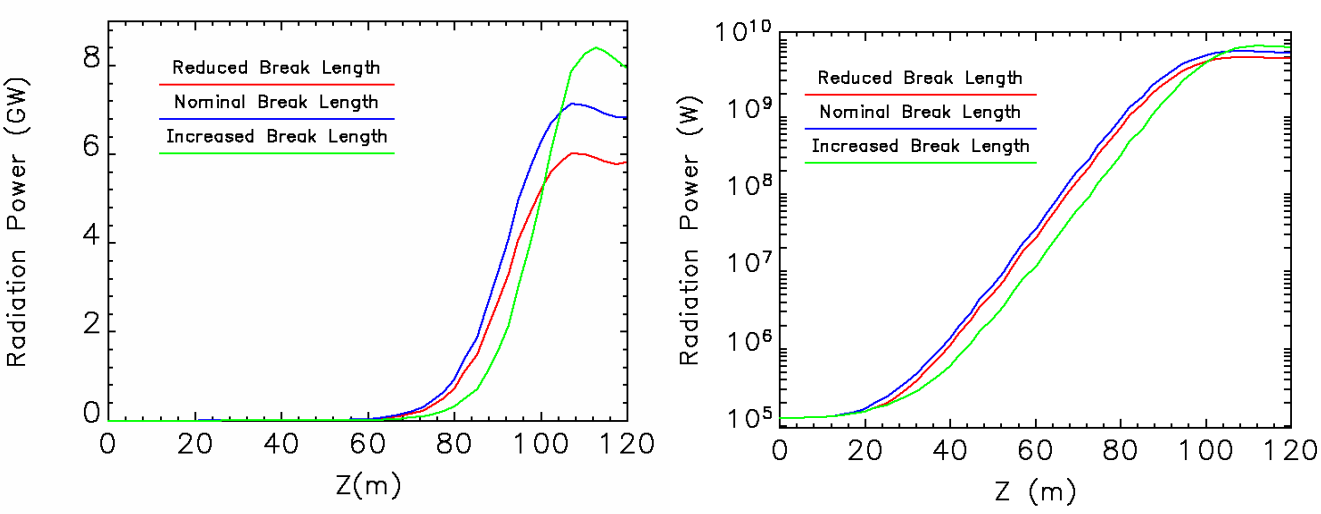

Figure A-1: Linear and semi-log scale plots of time-averaged power vs. z for different break length choices.

Dejus and Vinokurov; (c) Same as case (b) except the break lengths were *decreased* rather than increased by the Dejus and Vinokurov values.

Figure A-1 above shows the results for time-averaged power versus z in the undulator. One sees that, although the nominal break length run shows the greatest power before saturation, the run with increased break lengths eventually catches up and has perhaps 15-20% more power, ~15 m beyond the saturation point. The run with decreased break lengths saturates at the same point as the nominal run but at a power ~15% less. The exponential gain lengths are nearly the same in all three cases. Another sets of runs with the same shot noise seed showed essentially identical results; consequently, we do not believe there is significant sensitivity to the actual phase distribution of the individual components of the initial $P(\omega)$, nor should there be any in the exponential growth regime before saturation.

Figure A-2 below shows near-field spectra at two locations in the undulator. The first at z=80 m is before saturation while the second at z=120 m is well into saturation. It is apparent that increasing the break lengths in the first three drift sections helps amplify the wavelengths components to the red side of the central wavelength, 0.150 nm. Decreasing the break lengths has little effect because there are essentially no particles resonant at wavelengths much shorter than 0.150 nm whereas long wards there are particles whose transverse momenta on the average are much larger than the mean (i.e. the “high” emittance component). The output bandwidth appears to have been increased for the increased break length run. This suggests that, for experiments which require a much narrow bandwidth than that naturally produced by SASE in the LCLS, the increase in power due to the modification of the break lengths will have little direct beneficial effect.

In conclusion, these fully time-dependent runs with GINGER have confirmed Dejus’s and Vinokurov’s suggestion that increasing the first three break lengths in the LCLS undulator lattice can help provide a small increase in power. However, this increase comes at a price of increased saturation length and, more over, the spectral brightness

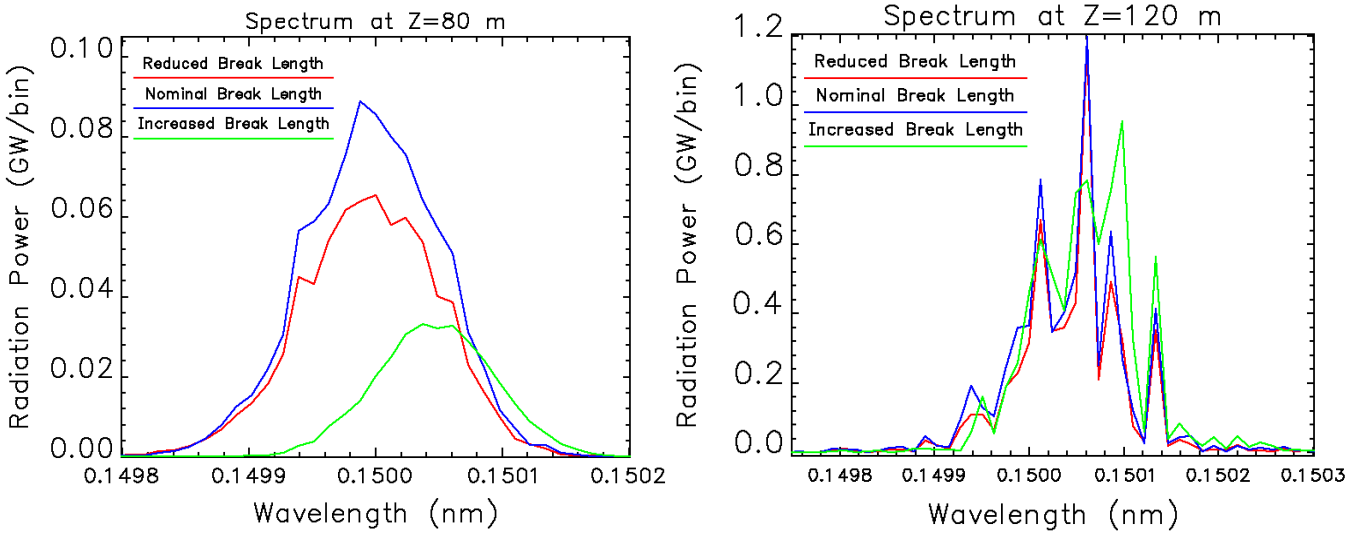

Figure A-2: GINGER predictions for near field spectra at two locations in the undulator for different choices of undulator break lengths configuration.

does not appear to have increased at all. Thus, it is not obvious that this suggestion truly has great utility for LCLS operation.

LCLS X-RAY FEL OUTPUT PERFORMANCE IN THE PRESENCE OF HIGHLY TIME-DEPENDENT UNDULATOR WAKEFIELDS*

W.M. Fawley†, LBNL, Berkeley, CA 94720, USA
K.L.F. Bane, P. Emma, Z. Huang, H.-D. Nuhn, G. Stupakov, SLAC, Stanford, CA 94309, USA
S. Reiche, UCLA, Los Angeles, CA 90095, USA

Abstract

Energy loss due to wakefields within a long undulator, if not compensated by an appropriate tapering of the magnetic field strength, can degrade the FEL process by detuning the resonant FEL frequency. The wakefields arise from the vacuum chamber wall resistivity, its surface roughness, and abrupt changes in its aperture. For LCLS parameters, the resistive-wall component is the most critical and depends upon the chamber material (*e.g.*, Cu) and its radius. Of recent interest[1] is the so-called "AC" component of the resistive-wall wake which can lead to strong variations on very short timescales (*e.g.*, ~ 20 fs). To study the expected performance of the LCLS in the presence of these wakefields, we have made an extensive series of start-to-end SASE simulations with tracking codes *PARMELA* and *ELEGANT*, and time-dependent FEL simulation codes *GENESIS1.3* and *GINGER*. We discuss the impact of the wakefield losses upon output energy, spectral bandwidth, and temporal envelope of the output FEL pulse, as well as the benefits of a partial compensation of the time-dependent wake losses obtained with a slight z-dependent taper in the undulator field. We compare the taper results to those predicted analytically[2].

INTRODUCTION

The Linac Coherent Light Source (LCLS)[3] currently under construction at SLAC will operate in the x-ray wavelength range of 0.15 – 1.5 nm. Due to the need for both a large undulator field strength (≈ 1.25 T) and a relatively short period (30 mm), the undulator chamber beam pipe must be quite small with an inner radius of 2.5 mm. The interaction of this chamber with the large instantaneous current of the LCLS electron pulse, $\sim O(1-10\,\mathrm{kA})$ can induce strong electromagnetic wakefields. The longitudinal wakefield can disrupt FEL performance by accelerating electrons off-resonance. Because the wakefields at a given z are not constant in time but depend on the position along the electron bunch, their effects cannot be completely compensated either locally or globally in z by an adjustment of the undulator field (taper). Recently (see, *e.g.*, [1]), there has been strong interest in examining the so-called "AC" component of the resistive-wall longitudinal wake. When excited by high frequency structure on the sub-ps duration electron bunch, the resultant wakefields can vary strongly on timescales as short as 20 fs, with most parts of the pulse suffering net deceleration but other parts net acceleration.

For a SASE FEL device like the LCLS, simple scaling arguments suggest that if by nominal saturation length $L_{sat} \approx 1.5\lambda_u/\rho$ a given portion of the pulse suffers a net acceleration equivalent to a shift of the resonant wavelength by twice the RMS bandwidth (*i.e.*, $\Delta\gamma/\gamma \approx 1.2\rho$) there should be a strong effect upon the instantaneous output power. Here ρ is the FEL parameter and λ_u is the undulator wavelength. A far more detailed analysis [2] shows that the output power has a FWHM in $\Delta\gamma/\gamma \approx 4\rho$ at saturation, and, moreover, using applying net $\Delta\gamma/\gamma \approx 2\rho$ (*i.e.*, including wakes, spontaneous energy losses, and the effects of a linear taper if any) over the saturation length approximately doubles the maximum power extraction as compared with no net $d\gamma/dz$. For LCLS with $\rho \approx 5\times10^{-4}$ and $L_{sat} \approx 90\,\mathrm{m}$ this suggests using an optimum taper equivalent to a net positive $150\,\mathrm{kV/m}$ accelerating field.

In the remainder of this paper, we briefly discuss calculations for the time-dependent wakes for sample predicted LCLS pulses obtained from "start-to-end" simulations upstream of the undulator. To model the expected FEL output radiation for this relatively complex problem, we use two fully time-dependent FEL simulation codes, *GENESIS* and *GINGER*. Of particular interest is the degree to which wakefield effects can be compensated by a simple linear taper in undulator strength (represented in the simulation codes by a constant E_z). We concentrate upon two particular operational modes of the LCLS: (1) the "normal" 1-nC bunch charge case for which there are large head and tail current spikes, each of which couples strongly to the resistive-wall wake (2) a "low" 200-pC bunch charge case [4] in which the current is far more uniform with time. Our results suggest that the latter case should be given strong consideration as the preferred operating mode because wake compensation by a simple undulator strength taper gives a a far more constant output $P(t)$ with little difference in output pulse energy.

START-TO-END SIMULATIONS AND UNDULATOR VACUUM CHAMBER WAKE CALCULATIONS

To produce realistic 6-D phase space distributions as input for the FEL calculations discussed in the following sections, we did detailed "start-to-end" tracking simulations beginning with *PARMELA* for the gun and injector (for which we thank C. Limborg) followed by *ELEGANT* for

* Work supported by the Office of Science, U.S. Dept. of Energy, under Contracts DE-AC03-76SF00515 and DE-AC02-05CH11231. This work was performed in support of the LCLS project at SLAC.
† WMFawley@lbl.gov

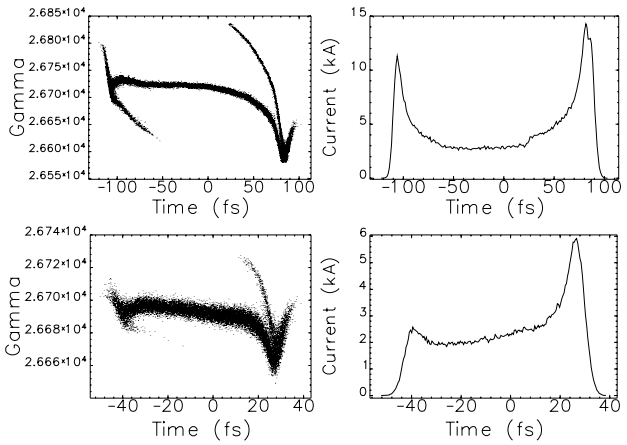

Figure 1: $t - \gamma$ scatterplot and $I(t)$ at undulator entrance for the 1-nC and 200-pC microbunch charge cases.

the remainder of the SLAC linac. The studies (see [4] for more detail) included CSR effects and presumed the existence of a laser-based beam heater [5] used to Landau damp the longitudinal space-charge instability. We modeled both the 1-nC and 200-pC bunch charge cases; Table 1 gives various relevant parameters for each. The low charge case is of particular interest because it is possible to virtually eliminate the high current spikes at the beam head and tail present in the 1-nC case (*i.e.*, compare the plots in Fig. 1). To obtain a 2.1 kA current in the undulator, the 200-pC case requires a significantly shorter bunch length (8μm rms) than that required at 1 nC (22μm). With a total compression factor of 70 (up from 40 at 1-nC) to limit pulse-to-pulse current jitter, the initial bunch length is then 1.5 times smaller. This together with the 5-times less charge drops the peak current in the RF gun to 30 A from 100 A; we believe that a 20% or greater reduction in transverse emittance at the gun is possible. The low bunch charge case has additional important advantages: the micro-bunching instability induced by longitudinal space charge and CSR has 3-times smaller gain; the relative horizontal projected emittance growth due to CSR in the BC2 chicane is reduced by three; finally, transverse wakefields and dispersion errors due to BPM, quadrupole, and RF-structure misalignments are essentially eliminated, due both to the lower charge and also to the shorter average bunch length and the smaller associated chirped energy spread.

At present, both the *GINGER* and *GENESIS* simulation

Table 1: Parameters for 1-nC and 0.2-nC bunch charge.

parameter	**sym.**	**1-nC**	**0.2-nC**	**unit**
init. rms bunch lng.	σ_{z_0}	840	560	μm
init. peak current	I_{pk_0}	100	30	A
init. slice emittance	$\gamma\epsilon_0$	1.0	0.80	μm
final rms bunch lng.	σ_{z_f}	22	8.0	μm
compression factor	C	40	70	
final peak current	I_{pk_f}	3.4	2.1	kA
final slice emittance	$\gamma\epsilon_f$	1.2	0.85	μm
final rms E spread	σ_δ	1.0	1.0	10^{-4}
pred. FEL sat. length	L_{sat}	87	88	m

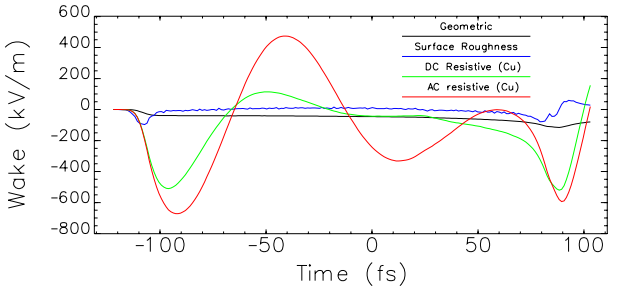

Figure 2: Individual longitudinal wake components for a 1-nC bunch charge propagating in a Cu vacuum chamber. A negative value corresponds to deceleration for electrons.

codes model the effective longitudinal wake as the sum of various components, the most important being the resistive-wall wake, the surface roughness wake, and the "geometric" wake which arises from discrete changes in chamber aperture (*e.g.*, pumping ports). For chamber roughnesses with reasonably large ratios of longitudinal scale length to transverse size, the resistive-wall component should dominate. Figure 2 displays some of the wake components calculated for the current waveform of a 1-nC LCLS pulse (the upper right plot of Fig. 1). For these calculations we presumed a round 2.5-mm inner radius vacuum chamber with a rms surface roughness of 100 nm over a period of 30 μm and an effective geometric wake gap length of 0.18 m over a 4-m period. The curve labeled "DC" refers to the resistive-wall wake calculated from a frequency-independent conductivity model for copper. The "AC" curve shows the predicted wake using a frequency-dependent σ model which for copper used a DC resistivity of 1.725×10^{-8} ohm/m and a time constant τ of 27 fs; for aluminum, the equivalent numbers are 2.733×10^{-8} and 8 fs. The AC and DC conductivity are related by $\sigma_{AC}(\omega) = \sigma_{DC}/(1 - i\omega\tau)$ where ω is the angular frequency.

The most striking difference between the two conductivity models is the nearly sinusoidal shape of the AC wake with a period of order 100-fs, much shorter than the 1-nC LCLS pulse duration. With a peak-to-peak difference of nearly 1 MV/m, one can see that it will be impossible via an undulator strength taper to keep all of the LCLS pulse in optimal resonance. Figure 3 plots the total wake versus time for both the 1-nC and 200-pC bunch charge cases. In contrast to the wake for the 1-nC bunch charge, the 200-pC wake is far more uniform in time for both Cu and Al vacuum chambers with a value between -100 and -200 kV/m. This constancy is due in part to the much shorter duration of the 200-pC bunch and in part to the absence of a high current spike at the beam head.

SIMULATION RESULTS

We used both the time-dependent FEL simulation codes *GENESIS* and *GINGER* to examine the predicted performance of the LCLS including vacuum chamber wakefields

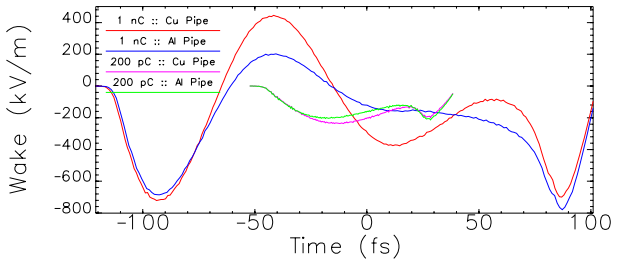

Figure 3: Total longitudinal wake for Cu & Al vacuum chambers for 1-nC and 200-pC bunch charges.

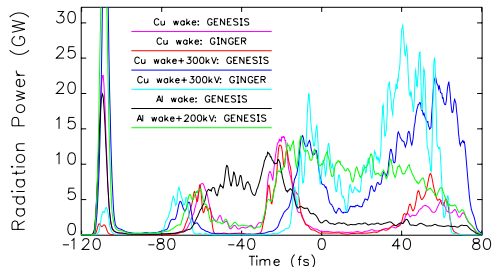

Figure 4: Predicted instantaneous power (artifically smoothed to 1-fs resolution) at $z = 100$ m for a 1-nC bunch charge propagating in a Cu vacuum chamber with and without wake effects.

and possible compensating undulator strength tapers. Both codes imported 6D macroparticle distributions from the *ELEGANT* tracking runs described in the previous section. In order to have sufficient spectral bandwidth around the nominal central radiation wavelength of 0.15 nm, we chose a slice-to-slice spacing of ≈ 12 attoseconds. Since the necessary number of slices for this slice spacing was quite large (~ 20000 for the 1-nC bunch charge case), both codes employed (individually different) algorithms to expand the number of *ELEGANT* macroparticles many-fold while preserving the fine scale details of the complicated 6D phase space (*e.g.*, correlations between x and γ; bimodal γ distributions in the head and tail regions, *etc.*). We also removed temporally-constant transverse offsets and tilts in order to optimize FEL performance. For *GINGER* this was important as its field solver is cylindrically axisymmetric; for *GENESIS* it is less so due to its full $x-y$ solver although uncorrected tilts can still strongly degrade performance. The simulations adopted the "standard" LCLS undulator configuration (as of fall 2004) including break sections and discrete quadrupole focusing magnets. For the wake calculations, we presumed a 2.5-mm inner radius vacuum chamber. No undulator errors were included nor were the effects of spontaneous radiation energy losses.

Figure 4 displays the predicted instantaneous SASE power from a 1-nC bunch charge at the 100-m location in the undulator. The black-colored "No wake" curve shows that in the absence of wakefield losses the radiation has an average level of approximately 14 GW and a duration of 180-fs. Comparison with the current profile in Fig. 1 shows that there is relatively little radiation produced in the head

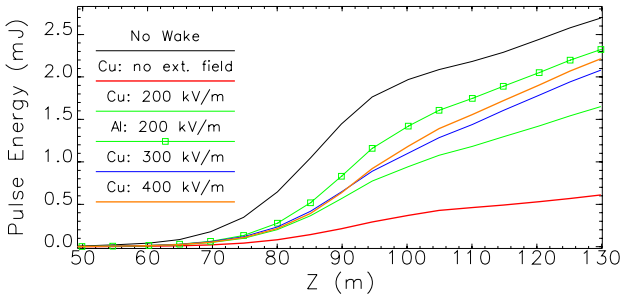

Figure 5: Pulse energy vs. z as predicted by *GINGER* for a 1-nC bunch charge in Al & Cu vacuum chambers for several wake compensation undulator strength tapers.

and tail high current spike regions. When uncompensated wakefields are included (the magenta and red curves), the overall radiation level is strongly suppressed and is mostly confined to three temporal regions around -60, -20, and +60 fs. Examination of the 1-nC Cu wake curve in Fig. 3 shows that these times correspond to when the wakefield lies between 0 and +200kV/m. When a compensating taper equivalent to +300kV/m is applied, the emission comes from temporal regions where the wakefield has a strength -350 to -200 kV/m. Given the computational complexity of these runs, the agreement between the two simulation codes is excellent, both in the level of the output emission and the temporal locations with the major exception being the power levels in the head current spike region (which due to a centroid offset in this region requires another 20 m of undulator for the GINGER simulation to reach the power levels shown by GENESIS at 100 m). Additional comparisons may be examined in Ref. [6]. One can see that time-integrated pulse energy (Fig. 5) drops more than five-fold when the uncompensated wake is compared with the no-wake case. At best, tapering recovers $\sim$80% of the energy by $z = 130$ m for a Cu vacuum chamber; the equivalent at the 100-m point is only 60%. Aluminum chambers result in somewhat better performance although we have not done extensive runs at 1-nC bunch charge for this material.

Simulation runs for the 200-pC bunch charge case show similar results in terms of optimal taper values. With no wake effects (the black curve in Fig. 6), the average power level is about 12 GW over a duration of ≈ 70 fs. Including the effects of an uncompensated wake (typical strength $\approx$-100 to -200 kV/m) drops the power level by nearly an order of magnitude. As increasingly strong compensating tapers are applied, the power is restored to the non-wake level by $\approx$+200 kV/m and nearly doubles for tapers in the +200 kV/m to +300 kV/m region. The optimal taper of +300 kV/m, which corresponds to a net acceleration of $\approx$+150 kV/m, is in good agreement with the analytic prediction of Ref. [2] that one can double the output power over a nominal no wake, no taper case. Greater overcompensation of the wake field with larger tapers steadily reduces the output power; by +600 kV/m the power is down more than four-fold from the no wake case.

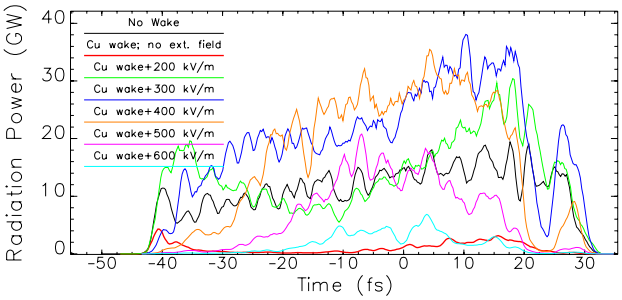

Figure 6: *GINGER* predictions for instantaneous power at $z = 130$ m for a 200-pC bunch charge propagating in a Cu vacuum chamber for various degrees of wake compensation tapers.

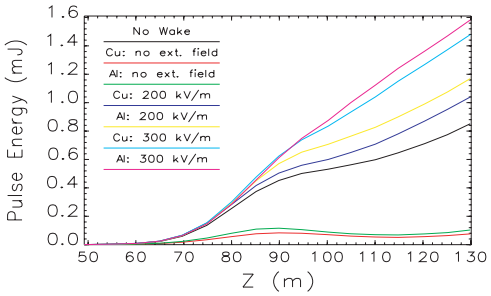

Figure 7: Pulse energy vs. z for 200-pC bunch charge propagating in Al & Cu vacuum chambers for various wake compensation undulator strength tapers.

Examining the total radiation pulse energy evaluated at the undulator exit of $z = 130$ m (Fig. 7), one sees that in accord to the predictions of Ref. [2], one can double the radiation power by overcompensating the effects of the wake. Although the exponential growth rate is only slightly increased by optimally tapering relative to the no-wake case, the power level at "first saturation" ($z \approx 90$ m) is larger by $\approx 50\%$ and then increases in a relative sense even more so in the next 40 meters of undulator.

One important question is whether the increase in output power level at and beyond L_{sat} is accompanied by an increase in spectral brightness. For a constant parameter pulse, Ref. [2] suggested that the spectral bandwidth should not be significantly changed when one uses the optimal net taper. The SASE simulations confirm this prediction as can be seen from examination of Fig. 8 where we plot the normalized inverse RMS bandwidth $\omega_0/\Delta\omega$ versus z. The optimal taper shows an decrease in inverse bandwidth of 20% or less as compared with the no-wake,no-taper case. Consequently, one can in fact achieve a significant increase in spectral brightness by operating with the a taper that leads to a net acceleration of $\Delta\gamma \approx 2\rho\gamma$.

CONCLUSIONS

The inclusion of AC conductivity effects causes strong temporal oscillations ($\tau \approx 100$ fs) in the strength of the predicted resistive-wall wake for a 1-nC LCLS bunch charge propagating in either copper or aluminum vacuum chambers. This oscillation is "shock excited" by the large current spike at the beam head and cannot be completely compensated for by tapering the undulator field strength with z. According to both *GINGER* and *GENESIS* time-dependent simulations, the end result is a temporal fragmentation of the output SASE radiation pulse into multiple sub-pulses. We note that the present LCLS design calls for a rectangular vacuum chamber which will reduce the wake strengths by ≈30%, partially ameliorating the situation.

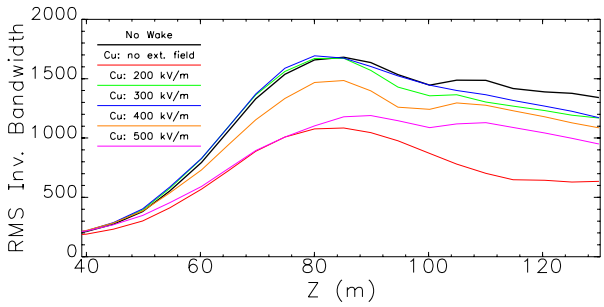

Figure 8: *GINGER* predictions for RMS inverse spectral bandwidth vs. z for a 200-pC bunch charge propagating in a Cu vacuum chamber for various compensation tapers.

Operating the LCLS at a lower charge of 200 pC improves not only the linac stability properties [4] but also strongly reduces peak-to-peak oscillations of the undulator wakefields, both because of the reduced average pulse current and the nearly complete elimination of the current spike at the beam head. Copper and aluminum vacuum chambers produce similar wakes that can be nearly completely compensated by undulator field tapers, resulting in a temporally smooth, relatively constant output radiation output. The simulations confirm the prediction by Ref. [2] that a net field taper equivalent to $\sim +150$ kV/m for LCLS parameters approximately doubles the instantaneous power at and beyond first saturation relative to a no-wake case. Most importantly, the output coherent photon count (1.3 mJ pulse energy $\approx 1.0 \times 10^{12}$ 0.15-nm photons) is reduced by only 25% relative to the 1-nC case. Due to the expectations of more stable operation and less demanding requirements upon the photoinjector, we believe that the low charge case should become the strongly preferred option.

REFERENCES

[1] K. Bane and G. Stupakov, SLAC PUB-10707 (2004).

[2] Z. Huang and G. Stupakov, *Phys. Rev. ST Accel. Beams*, **8**, 040702 (2005).

[3] *LCLS CDR*, SLAC Rpt. SLAC-R-593 (2002).

[4] P. Emma *et al.*, "An Optimized Low-Charge Configuration of the Linac Coherent Light Source", *Proc. 2005 Part. Accel. Conf.*, paper TOAB004 (2005).

[5] Z. Huang *et al.*, *PRST-AB*, **7**, 074401 (2004).

[6] S. Reiche *et al.*, "Optimization of the LCLS X-RAY FEL Performance in the Presence of Strong Undulator Wakefields", *Proc. 2005 Part. Accel. Conf.*, paper RPPT035 (2005).