AN ENHANCED GINGER SIMULATION CODE WITH HARMONIC EMISSION and HDF5 IO CAPABILITIES*

William M. Fawley†, LBNL, Berkeley, CA 94720, USA

Abstract

GINGER [1] is an axisymmetric, polychromatic $(r-z-t)$ FEL simulation code originally developed in the mid-1980's to model the performance of single-pass amplifiers. Over the past 15 years GINGER 's capabilities have been extended to include more complicated configurations such as undulators with drift spaces, dispersive sections, and vacuum chamber wakefield effects; multi-pass oscillators; and multi-stage harmonic cascades. Its coding base has been tuned to permit running effectively on platforms ranging from desktop PC's to massively parallel processors such as the IBM-SP. Recently, we have made significant changes to GINGER by replacing the original predictor-corrector field solver with a new direct implicit algorithm, adding harmonic emission capability, and switching to the HDF5 IO library[2] for output diagnostics. In this paper, we discuss some details regarding these changes and also present simulation results for LCLS SASE emission at λ =0.15 nm and higher harmonics.

INTRODUCTION

Over the past 25 years, there has been a steady advance in the use of numerical simulation codes to explore FEL physics, analyze experimental results, and to help design elaborate and expensive projects such as the LCLS. As the underlying computer hardware has grown both faster and often more complex (*i.e.,* massively parallel environments), FEL codes have similarly become more complex both in terms of the underlying physical phenomena they model (*e.g.,* wakefield losses) and the more realistic forms of simulation they attempt (*e.g.,* importation of massive macroparticle files for start-to-end tracking runs). Moreover, the amount of information the codes utilize and produce has increased by several orders of magnitude with multi-GB output and/or particle restart files becoming necessary for full time-dependent simulations of x-ray FEL's and multi-stage harmonic cascades.

This paper present recent changes to the GINGER simulation code[1] which originally was designed in the mid-1980's to study sideband growth in single-pass FEL amplifiers and which has steadily evolved since to study more and more elaborate configurations such as SASE devices, harmonic cascades and oscillators. The Fortran90 base currently underlying GINGER has proven very useful both in terms of modularization and in the number of useful language features (*e.g.,* array syntax and built-in operators, type definitions, memory management). The code structure has proven very amenable to efficient multiprocessing in which the different longitudinal slices are assigned to different processors with MPI [3] handling the necessary interprocessor communication. In the past year, we have replaced the original predictor-corrector field solver by a new implicit solver and also have extended the code to calculate radiation emission at multiple harmonics; the next section gives details on these changes. We then present results from a GINGER calculation for predicted harmonic emission from the LCLS.

With the addition of calculations of harmonic emission, typical GINGER output file sizes doubled or more. Based upon file compactness, IO speed, and flexibility in data layout, the output file format was changed from simple ASCII to HDF5[2]. In part because we believe that the FEL simulation community should become aware of the usefulness of HDF5 IO, the last section of this paper gives details concerning its implementation into GINGER .

THE NEW FIELD SOLVER AND IMPLEMENTATION OF HARMONICS

Adopting the slowly-varying envelope approximation (SVEA) with $E(r,z,t) \equiv \tilde{E}(r,z,t) \times \exp i(k_o z - \omega_o t)$, the field equation becomes

$$2ik_s \frac{\partial \tilde{E}}{\partial Z} = -\nabla_\perp^2 \tilde{E} - \left(\frac{\omega_0}{c^2} - k_o^2\right)\tilde{E} - \frac{4\pi i}{c^2}\omega_0 J_\perp \quad (1)$$

where $\partial/\partial Z$ is the Lagrangian derivative in the forward (*i.e.,* positive z) direction, and $J_\perp$ is the microbunched transverse current at angular frequency ω_0. The original GINGER field solver used a Gear predictor-corrector scheme [4] which automatically controlled step size to maximize computational speed for a given error allowance, but had negative aspects regarding algorithmic complexity and non-trivial coding difficulties with regards to adding additional dimensions such as harmonics and (eventually) non-axisymmetric modes. Consequently, we developed a completely new field solver based upon a backward-biased, implicit solution of the heat flow equation ("Case 5" in table 8.1 of [5]). This solution is always stable for arbitrary step size ΔZ and the chosen default bias factor makes the scheme equivalent to the Crank and Nicholson algorithm with errors scaling as Δr^2 and ΔZ^2. Tests of simple vacuum Gaussian mode propagation over multiple Rayleigh ranges show normalized energy conservation to better than 10^{-3} for $\Delta r/R_0 \leq 0.1$ where R_0 is the mode waist size.

* Work performed in support of the LCLS project at SLAC under US DOE Contract DE-AC02-76SF00515 (SLAC) and also supported under Contract DE-AC02-05CH11231 (LBNL).

† WMFawley@lbl.gov

Sec. 8.5 of Ref. [5] gives an elegant (and simple to code) solution method based upon a Gaussian elimination procedure where one first sweeps outward in r and then back to the axis Extension of the field solver to include higher harmonics is simple with (k_o, ω_0) being replaced by (hk_o, $h\omega_o$) and use of the appropriate current source term. For the latter, GINGER uses wiggler-period averaging and follows the standard formulation (see, *e.g.*, [6]) for odd harmonic emission that depends only upon the local harmonic microbunching amplitude (and ignores even harmonic emission arising from transverse gradients in the electron beam density). Tests by Z. Huang comparing the predicted 3rd harmonic power from nonlinear microbunching with that computed by GINGER for a simple $\lambda_0 = 0.15$ nm, steady-state amplifier with LCLS parameters show excellent agreement (*e.g.*, 10% or better) for sufficient macroparticle number (*e.g.*, $\geq$ 65K per slice). Use of smaller macroparticle number tends to lead to an overestimation of the harmonic power, perhaps because the longer Rayleigh length (and thus less radial smoothing in Z) at 0.05 nm makes the effective source term "grid noise" (*i.e.*, in higher radial modes) much worse than that corresponding to the fundamental (for which 8K macroparticles appear more to be than sufficient).

The original macroparticle longitudinal mover also used a Gear-scheme (which was extreme computational overkill given the slow evolution of particle phase and energy) and was replaced by a simple second-order Runge-Kutta solver in which only the fundamental radiation field acts back upon the particles. Within this limitation one can study nonlinear harmonic generation (the dominant emission term for most SASE and harmonic cascade devices) but not linear amplification. As shown in Ref. [7], ignoring the harmonic emission feedback is an extremely good approximation for SASE x-ray FELs and essentially all FELs in which gain at the fundamental wavelength is much larger than that at the harmonics, as is the usual case.

In fall 2006 we expect to add the harmonic field feedback terms to the mover and also a fourth-order Runge-Kutta scheme as an option. Lastly, we note that Eq. (1) is separable by azimuthal mode and thus the new field solver can be easily extended (as is hopefully planned for 2007) to solve for $\tilde{E}(r, \varphi, z, t)$.

AN EXAMPLE LCLS HARMONIC EMISSION CALCULATION

Nonlinear harmonic emission from the LCLS is potentially extremely useful, both for experimentalists and as a diagnostic of FEL performance. Adopting design parameters (3.4 kA, $\gamma = 26693$, 1.2 mm-mrad, $B_w = 1.25$ T, $\sigma_\gamma/\gamma = 10^{-4}$), and a quad FODO focusing lattice that produces an average Twiss β-function of 28 m, we simulated a LCLS 9.2-fs segment at 6-attosecond temporal resolution with 32768 macroparticles per slice. The calculation was run on 128-POWER 5 (1.9 GHZ) processors of the "BASSI" IBM p575 system at NERSC and required $\approx$ 40 minutes of wall-clock time. Figure 1 plots the power at the fundamental wavelength of 0.15 nm and also the 3rd and 5th harmonics as functions of z. One sees that the plateau region (on this semi-log plot) for small z where spontaneous emission dominates the coherent FEL component lasts much longer for the harmonics than the fundamental; in reality the plateau region would extend even further because these calculation do not include non-axisymmetric spontaneous emission. While the third harmonic grows more rapidly with z than the fundamental in the region $60 \leq z \leq 85$ m, it saturates at a low ($\sim 1\%$) relative power level.

Comparing the behavior of the on-axis far-field power (not displayed) with the total near-field power (Fig. 1) reveals that the plateau region is much smaller for the fundamental but there is essentially no change for the harmonics. This is likely related both to the short gain length for the

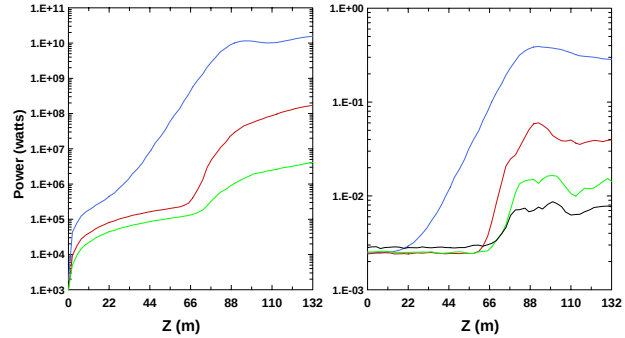

Figure 1: Time-averaged SASE power (left) and microbunching fraction (right) vs. z for an LCLS pulse at the 0.15-nm fundamental (blue curve) and also the 3rd, 5th, and 7th harmonics (bunching only) (red, green, black curves).

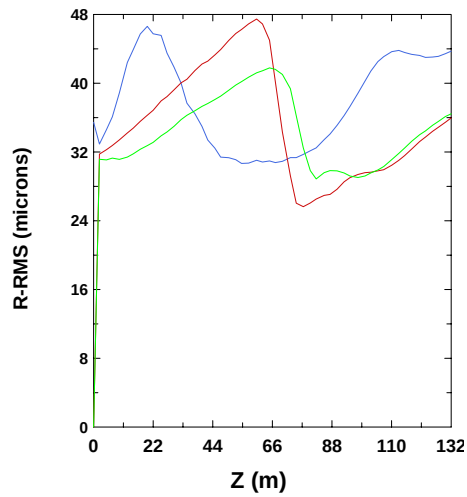

Figure 2: RMS transverse size vs. z for SASE emission from an LCLS pulse at the 0.15-nm fundamental wavelength (blue curve) and also the 3rd and 5th harmonics (red and green curves).

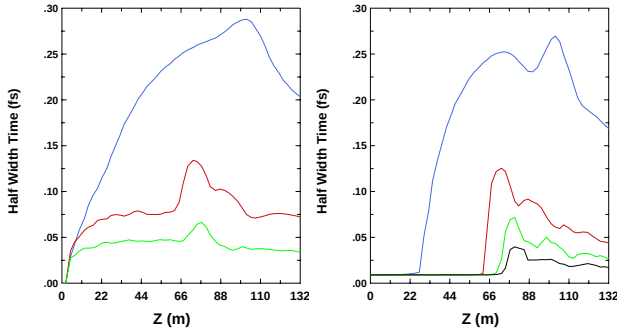

Figure 3: Plots of $C_{1/2}(\tau)$ for the on-axis far field emission (left) and microbunching (right) for the fundamental and harmonics (same color scheme as Fig. 1).

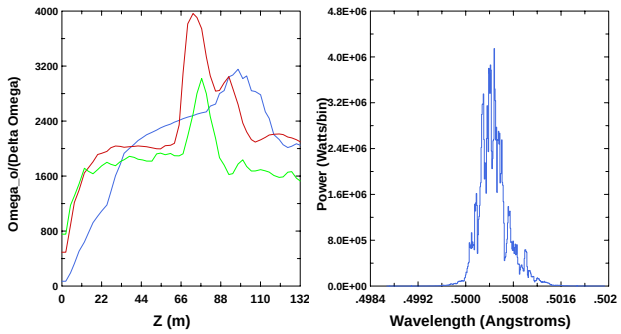

Figure 4: Plots of RMS spectral width for on-axis far field emission at the fundamental and harmonics (left) and the 3rd harmonic output (near-field) spectrum (smoothed with a 5-bin average) (right).

harmonics $L_{g,h} \approx L_{g,1}/h$ and the smaller mode size (see Fig. 2) for the coherent harmonic emission which reduces some of the far-field contrast *vis-a-vis* the spontaneous harmonic emission. Near power saturation, the RMS transverse sizes of the 3rd and 5th harmonic emission are ∼20% smaller than that of the fundamental (Fig. 2). However, after saturation, they also grow with z despite a greater Rayleigh range, indicating perhaps that optical guiding effects are far more important for the fundamental.

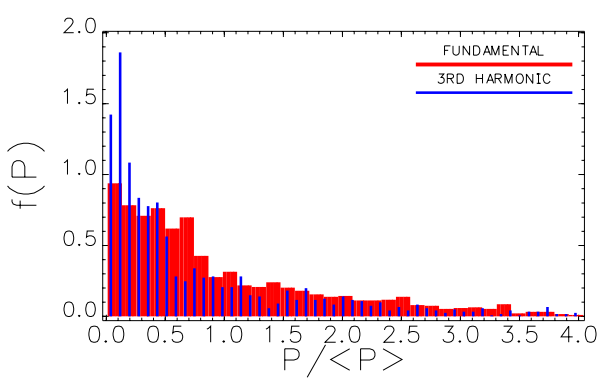

Figure 5: Histogram of the probability distribution of $P(t)/<P>$ at the fundamental and 3rd harmonic at the undulator exit.

Standard SASE theory (*e.g.*, Ref. [8]) predicts that the inverse spectral bandwidth $\omega/\Delta\omega$ and autocorrelation time $C_{1/2}(\tau)$ (the point at which the temporal autocorrelation function drops to 0.5) for FEL radiation increase as $z^{1/2}$. Such behavior is demonstrated by radiation at the fundamental (blue curves in Fig. 3. For the harmonics, the coherent component of $C_{1/2}(\tau)$ suddenly dominates over the spontaneous emission (for the radiation) and shot noise (for the microbunching) at two-thirds of a saturation length, reaches a maximum within a gain length, and then begins declining with z, in contrast to the fundamental which reaches a maximum somewhat later in z (right plot in Fig. 1) and only drops in the last 20-m or so. Presumably, debunching due to the development of a large energy spread is responsible for the greater sensitivity of the harmonics.

An autocorrelation function and power spectrum can be similarly calculated for microbunching by using the complex average computed over the macroparticles corresponding to each longitudinal electron beam slice. In contrast to the radiation quantities, the microbunching $C_{1/2}(\tau)$ (right plot in Fig. 3) has a slow but steady linear decay from its maximum which is perhaps associated with the temporal variation of Ω_{syn} over the individual microspikes. Ideally, in a near time-steady situation with essentially no variation in radiation eikonal phase, $C_{1/2}(\tau)$ would be constant with harmonic. However, in the SASE regime, Saldin *et al.* [7] have predicted from 1-D numerical studies that $C_{1/2}^{h=3}(\tau)/C_{1/2}^{fund.}(\tau) \approx 0.65$; the GINGER LCLS results show a ratio $\simeq 0.5$, a reduction attributable perhaps to nonzero emittance effects not included in the 1-D study.

Figure 4 plots the inverse spectra width versus z together with the output spectrum at the third harmonic. Another diagnostic of radiation field correlation, the ratio $[\omega/\Delta\omega]/[1.18\omega C_{1/2}(\tau)]$ ($\equiv 1$ for a Gaussian power spectrum), is 0.76 for the fundamental and drops to to 0.69 and 0.65 for the third and fifth harmonics at the position of maximum inverse bandwidth.

Another apparent difference with the Saldin *et al.* 1-D results concerns the probability distribution for $P(t)$. In Fig. 5 we plot the normalized histogram for the fundamental and third harmonic ofthe output (near-field) radiation. One sees that in accord with the earlier study, the third harmonic emission is more peaked toward $P = 0$. However, the peak is more extreme and the first order moment is much smaller, evidently because of large temporal regions of relatively low harmonic power. The probability distribution for the fundamental is also more peaked toward $P = 0$ whereas Ref. [7] found a peak near $P =< P >$. Because our data is taken somewhat deeper in saturation than the z point chosen by Saldin *et al.*, it is possible (albeit unlikely) that the distribution function reverts back to a negative exponential corresponding to the exponential gain regime. or-that this particular SASE run does not have sufficient statistics in terms of slice number to clearly resolve the distribution shape. However, it may also be possible that inclusion of true transverse effects actually change the nature of

the distribution function in the saturation regime. Consequently, this topic may warrant some additional study.

IMPLEMENTATION OF HDF5 IO

As is true for most large-scale simulation tools, an FEL code must have numerous IO capabilities. On the input side, the user must be able to specify a number of beam and run parameters. Depending upon the run complexity, additional input might be required, *e.g.,* 6D macroparticle phase space information from a tracking particle output file or a "restart" file from a previous FEL code run (as might be true for a harmonic cascade); a time-dependent input radiation field; an undulator "lattice" including $a_w(z)$, possible pole error field and dipole corrector information, and possible vacuum chamber longitudinal wake field input. Output information can be very extensive: spatially- and temporally-resolved radiation field information including harmonics; 5D macroparticle phase space snapshots at different $z-$ and $t-$ locations for later use in scatterplots; $z-$ and $t-$ resolved "scalar" information such as total radiation power and macroparticle microbunching (at multiple harmonics), energy spread, on-axis far field complex amplitude; and possible full 6D macroparticle and/or 2D (r, t) radiation field information for a subsequent downstream undulator run. As the raw binary size of this information can exceed 100 MB for fields and 1 GB for particles, simple ASCII formats are unattractive both due to their size ($\sim$3X larger than pure binary) and slow IO transfer speed. Moreover, unless one designs a very clever and robust self-describing data structure and/or divides the heterogeneous output into many separate files per run, it is difficult to have an output file system that is reasonably easy to analyze (and which maintains good backward compatibility as the simulation code inevitably evolves in complexity).

In answer to the above requirements and problems, the HDF5[2] IO system offers a strong set of attractive features: (1) self-describing data format, including optional dataset attributes (2) Unix-like directory trees including soft links (3) compact data storage with little additional overhead relative to binary (4) multiplatform data file portability (*e.g.,* little- to big-endian, 32- to 64-bit native) (5) a rich class of native data types range from simple strings to 1-D integer arrays to many dimensional real arrays (6) user-definable data types (7) MPI-based parallel IO on platforms such as the IBM-SP. The HDF5 development group also provides a set of useful tools for HDF5 file data visualization and data dumping of individual structures. Many projects in the high performance computing community have embraced HDF5 as well as a number of 3rd party visualization tools venders (*e.g.,* IDL). I note that the GENESIS code [9] also has a growing HDF5 IO capability and would suggest that other FEL and tracking code developers consider this option.

Currently, the GINGER HDF5 output file has 5 main top directories: /base_param to pass simulation run variables such as the central wavelength, number of slices, etc.; /grids which have the r- and z- locations of grid points; /input which has the complete ASCII input and template files for the run; /radiation which contains all scalar and vector radiation field information; and /particles which has all particle-derived information including scatterplot dumps, if any. Separately, particle restart files can also be written in HDF5 format and the extension to radiation restart files (as is needed for multistage harmonic cascades) will be done in the near future. Knowledge of the exact "path" of a given data set (*e.g.,* /radiation/scalar_data/...) is completely sufficient for a HDF5 utility (and GINGER post-processor) to access the data set; the exact disk address within the HDF5 file is not needed by external programs but is provided by a low-level look-up table in the HDF5 file. This massively simplifies data access With the migration to HDF5 format, we believe it will be relatively straight-forward and painless to extend the GINGER output file to handle fully "3D" field information (*i.e.,* $[r, z, t] \rightarrow [r, \varphi, z, t]$) and variable spatial resolution (*e.g.,* relative finely z-resolved scalar quantities such a radiation power and microbunching but coarser resolution of vector quantities such as $E(r, z, t)$) as needed by future code upgrades.

ACKNOWLEDGEMENTS

I am pleased to acknowledge many useful discussions with numerous individuals including A. Adelmann, E. Allaria, M. Borland, W. Graves, Z. Huang, H.-D. Nuhn, G. Penn, E. Pourmal (HDF5 group/NCSA), and S. Reiche. I also thank the LCLS project office at SLAC for financially supporting much of the work presented here and also the US DOE National Energy Supercomputer Center for providing Computational resources.

REFERENCES

[1] W.M. Fawley, "A User Manual for GINGER ...", LBNL-49625-Rev. 1 (2004); see also http://www-ssrl.slac.stanford.edu/lcls/lcls_tech_notes.html/LCLS-TN-04-3.pdf .

[2] National Center for Supercomputing Applications (U. Ill.), "HDF5 Home Page", http://hdf.ncsa.uiuc.edu/HDF5/

[3] Message Passing Interface Forum, "MPI: A Message-Passing Interface standard", Int. J. Supercomputer Applications and High Perf. Comp. **8** (1994); see also: http://www.mpi-forum.org/docs/docs.html .

[4] C.W. Gear, Comm. ACM **14**, 176 (1971).

[5] R.D. Richtmeyer and K.W, Morton, "Difference Methods for Initial-Value Problems", (Interscience, New York), 2nd. Ed., 1967.

[6] Z. Huang and K.-J. Kim, Phys. Rev. E **62**, 7295 (2000).

[7] E.L. Saldin, E.A. Schneidmiller, and M.V. Yurkov, *Proc. 27th Int. FEL Conf.* (2005, Stanford, CA), 51.

[8] K.-J. Kim, Phys. Rev. Lett. **57**, 1871 (1986).

[9] S. Reiche, Nucl. Inst. Meth. Phys. Res. A **429**, 243 (1999).