USING A LIENARD-WIECHERT SOLVER TO STUDY COHERENT SYNCHROTRON RADIATION EFFECTS

R.D. Ryne*, C.E. Mitchell, J. Qiang, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
B. Carlsten, N. Yampolsky, Los Alamos National Laboratory, Los Alamos, NM, USA

Abstract

We report on coherent synchrotron radiation (CSR) modeling using a massively parallel, first-principles 3D Lienard-Wiechert solver. The solver is able to perform simulations with hundreds of millions to billions of simulation particles, the same as the real-world number of electrons per bunch typically present in modern accelerators. We have recently extended this tool to model a variety of beam transport systems including undulators. In this paper we provide an overview of the tool and present several examples. We also describe the concept of a Lienard-Wiechert particle-mesh (LWPM) code, and how such a code might make it possible to perform parallel, self-consistent modeling using a Lienard-Wiechert approach.

INTRODUCTION

Particle accelerators are among the versatile and important tools of scientific discovery and technology advancement. They are responsible for a wealth of advances in materials science, chemistry, bioscience, high-energy physics, and nuclear physics. They also have important applications to the environment, energy, national security, and medicine. Given the enormous benefit of particle accelerators and their extreme complexity, high performance computing using parallel computers has become an essential tool for their design and optimization to reduce cost and risk, maximize performance, and explore advanced concepts.

Early examples of parallel simulation in the U.S. accelerator community date from the late 1980's and early 1990's [1]. By the mid-1990's parallel beam dynamics codes had been developed to run on the Thinking Machines CM-5 computer at LANL's Advanced Computing Laboratory [2, 3]. In 1997 the U.S. Department of Energy (DOE) approved a "Grand Challenge" project in Computational Accelerator Physics [4]. This later evolved into a DOE Scientific Discovery through Advanced Computing (SciDAC) project [5, 6]. This, as well as other R&D efforts in parallel accelerator simulation worldwide, led to the parallelization of existing beam dynamics codes and the development of new codes. Examples include ASTRA [7], BeamBeam3D [8], CSRtrack [9], elegant/SDDS [10], GENESIS [11], G4Beamline [12], ICOOL [13], IMPACT [14, 15], MaryLie/IMPACT [16], OPAL [17], SPUR [18], Synergia [19], TRACK [20],TREDI [21], and WARP [22], to name a few.

From the mid-1990's to the present, significant attention was devoted to parallel 3D space-charge modeling, multi-physics modeling, and increasingly large-scale simulation. In regard to 3D space-charge modeling, many parallel Poisson solvers were developed to treat a variety of boundary conditions, e.g., [23, 24, 25]. Integrated Green functions (IGFs) were introduced to increase solver performance and address grid aspect ratio issues [26, 27, 28], and are now used in several codes worldwide [7, 8, 14, 16, 17]. IGFs have also been applied to model 1D CSR [29, 30]. In regard to multi-physics modeling, split-operator methods were introduced as a means to combine high-order optics effects with parallel 3D space-charge and other effects [31], and are used in several codes [14, 16, 17, 19]. In general, parallel beam dynamics codes now contain, and are routinely used to model, a variety of phenomena including high-order optics, space-charge effects, wakefield effects, 1-D CSR effects, electron-cloud effects, and beam-material interactions. Regarding the trend toward increasingly large-scale simulation, a start-to-end 2-billion-particle simulation of a future light source, based on the single parallel executable containing IMPACT-T, IMPACT-Z, and GENESIS, requires 10 hours on 2048 cores [32].

Despite these major advances in parallel multi-physics modeling, the simulation of 3D CSR effects has remained a major challenge. A first-principles classical treatment usually involves the Lienard-Wiechert (L-W) formalism. Since this involves quantities when the radiation was emitted (i.e. at retarded times and locations), it requires storing a history of each particle's trajectory. Also, CSR phenomena can exhibit large fluctuations which are physical, not numerical, hence it is often necessary to use a large number of simulation particles if those fluctuations are to be modeled correctly. Storing a large number of particles over a lengthy time history imposes a huge memory requirement. Furthermore the calculation of retarded quantities is iterative and extremely time consuming. Consider that the calculation of an electric field component on a grid in an electrostatic code, e.g., $x/|r|^3$, requires only a small number of floating point operations at each grid point; by contrast the calculation of the L-W field at just a single grid point requires a small simulation code itself to implement the iteration to find the retarded quantities, and furthermore the iteration involves numerical integration of trajectories. In summary a L-W solver involves large memory and many floating point operations, and obviously requires parallel computing. In addition, to embed such a capability in a self-consistent beam dynamics code would greatly compound the computational requirements.

Despite these computational challenges, in the following we will present results that point to the possibility of a mas-

* RDRyne@lbl.gov

ISBN 978-3-95450-126-7

sively parallel L-W-based beam dynamics code. First, we will provide an overview of our parallel L-W solver. Next we will provide some example applications to illustrate the usefulness of this approach. Lastly, we will provide an outline of how standard particle-mesh (PM) techniques used in parallel electrostatic or quasi-electrostatic particle-in-cell codes can be modified to produce a self-consistent Lienard-Wiechert particle-mesh (LWPM) code.

LIENARD-WIECHERT SOLVER

We begin with the Lienard-Wiechert fields in free space:

$$\begin{aligned} \vec{E} &= \left[\frac{q}{\gamma^2\kappa^3 R^2}\left(\hat{n}-\vec{\beta}\right)+\frac{q}{\kappa^3 Rc}\hat{n}\times\left\{\left(\hat{n}-\vec{\beta}\right)\times\frac{\partial\vec{\beta}}{\partial t}\right\}\right] \\ \vec{B} &= \hat{n}\times\vec{E} \end{aligned} \quad (1)$$

where square brackets denote that the quantity inside is to be evaluated at the retarded time. In the above, $\vec{\beta}$ denotes a particle's velocity vector divided by the speed of light, $\vec{R}$ is a vector pointing from a radiating particle's position at the retarded time to the observation point at time t, $\hat{n}$ is the unit vector $\hat{n} = \vec{R}/|R|$, and $\kappa = 1 - \hat{n}\cdot\vec{\beta}$. For a given observation point (x, y, z) and observation time t, the retarded quantities satisfy,

$$(x-x_r)^2+(y-y_r)^2+(z-z_r)^2-c^2(t-t_r)^2=0, \quad (2)$$

where (x_r, y_r, z_r) denotes a particle's position at the retarded time t_r. Note that other boundary conditions can be treated if the underlying Green function is known. For example, conducting plate boundary conditions can be treated using the fields associated with Eq. 4.2 of [33].

In a previous note we reported on calculations based on an electron bunch in a uniform magnetic field [34]. Here we describe a methodology to treat arbitrary beamlines. Suppose each particle's evolution is described by a six-vector ζ as a function of some independent variable, τ. Suppose also that each particle's history has been stored at a number of locations, i.e., suppose that for each particle we know a sequence, ζ^k, for k values of the independent variable τ.

In our Lienard-Wiechert solver, for a given observation point, we loop over particles and for each particle we first find adjacent quantities $k, k+1$ that bracket t_r. This is done using a bisection search to locate a change of sign in Eq. (2) at the stored values. We then use Brent's method (routine `zbrent` from [35]) to iteratively find the root of Eq. (2). We have found this approach to be more robust to round-off than a simple Newton search. During the iteration we need to perform numerical integration of the equations of motion between two stored history values. For this purpose we use a Dormand-Prince 8-5-3 algorithm with automatic step size adjustment [36]. After the iteration has converged to the retarded quantities we evaluate Eq. (1) to obtain the L-W fields due to each particle. In the following t_r is calculated to a relative accuracy of 10^{-10}. For some problems we have found it useful to use extended precision to ensure that roundoff is not noticeable. For this purpose we use `DDFUN90` package for double-double arithmetic [37]. Our solver is an MPI code that uses particle decomposition. Assuming the observation points are replicated on the processors, the total L-W fields are found by an MPI reduction that adds the contributions from all MPI processes. Later we will discuss an alternative approach that uses domain decomposition.

APPLICATIONS

As mentioned previously, in [34] we focused on steady-state CSR in dipoles. We used a L-W solver to explore limits of the 1D model, the strength of fluctuations due to shot noise, and the enhancement of dipole CSR when a bunch has a microbunched structure. In the following we will present new examples involving dipole CSR, and examples involving undulator radiation. Except where noted, the field plots that follow present just the L-W radiation component. Also, the retarded quantities are computed by integrating trajectories through just the external fields, not including the self-fields.

CSR in Dipoles

Figure 1 shows the results from 6 simulations, each with 6.24 billion particles (corresponding to 1 nC), of a zero emittance, 1 GeV bunch in a dipole with $\rho = 1\ m$. Starting from a 1x1x10 micron Gaussian bunch, we multiplied the longitudinal distribution by a function $a + b\sin(2\pi z/\lambda_{mod})$, where λ_{mod} denotes a modulation wavelength, and with a and b chosen so that the amplitude varied from 0.1 to 1. We randomly sampled the distribution using a rejection method. Figure 1 shows the longitudinal radiation electric field, $E_{z,rad}$ for λ_{mod} = 5, 10, 50, 100, 500 nm. The curve labeled DC has no modulation, i.e., the longitudinal profile is a Gaussian. Since no adequate theory can predict the wavelength dependence of CSR enhancement, except in some regimes, large-scale L-W simulation provides a useful means to explore this phenomenon.

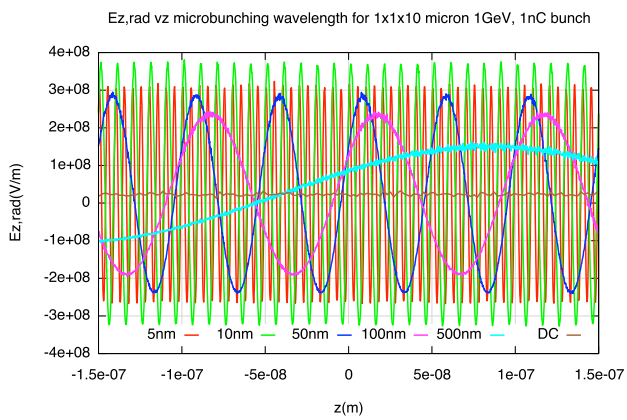

Figure 1: Longitudinal radiation electric field, $E_{z,rad}$ for a 1 GeV bunch with an imposed longitudinal density modulation, simulated with 6.24 billion electrons, for different values of the modulation wavelength.

ISBN 978-3-95450-126-7

In [34] we also showed that the 1D model of CSR was remarkably robust, that only when a Gaussian bunch becomes extremely flattened ("plate-like") does the 1D model break down. However, it is worth pointing out that there are 3D effects that cannot be captured by the 1D model. An example is shown in Fig. 2. In this figure, the bunch is microbunched at $\lambda_{mod} = 100$ nm. The bunch rms dimensions in x-y-z (without microbunching) are shown for 4 cases: 10x10x10 micron, 10x1x10 micron, 1x10x10 micron, and 1x1x10 micron. As is clear from the figure, the microbunching enhancement is sensitive to the vertical bunch size but not to the horizontal bunch size. Such studies could not be carried out with a 1D CSR model.

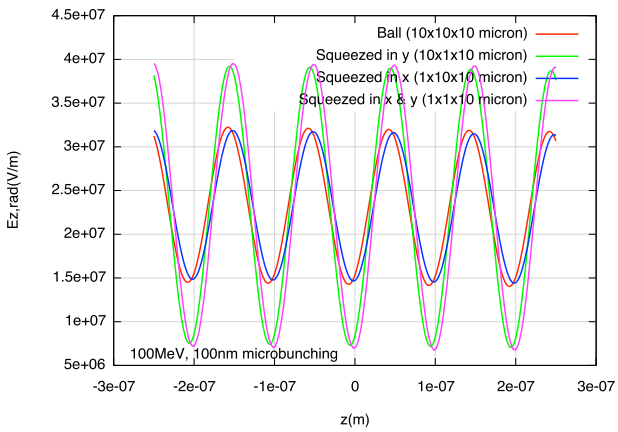

Figure 2: Microbunching field enhancement for different transverse bunch sizes. The microbunching enhancement is seen to be sensitive to the vertical bunch size but not to the horizontal bunch size.

Before leaving the topic of dipole CSR we will present one more example. Consider the evolution of a test particle moving in the CSR field. Figure 3 shows the tangential electric field that would be experienced by a test particle at the center of a 1 GeV, 624M electron, 10x10x10 micron Gaussian bunch as it travels through 3 degrees of transport in a magnet with $\rho = 1\ m$. As can be seen in the inset, the shot noise fluctuations are very well resolved. Using this data we computed its autocorrelation function and examined its dependence on energy. Figure 4 shows the energy change as a function of propagation distance in the dipole that would be experienced by a test particle at the center of the bunch. The bunch is a zero emittance, 1 GeV, 1nC Gaussian with rms sizes $\sigma_x = \sigma_y = \sigma_z = 10\ micron$. The 4 curves correspond to 4 different realizations of the distribution, each modeled with 6.24 billion particles. The time histories corresponding to different realizations of the shot noise seem independent from each other and can be reasonably approximated by a random walk model. The red dashed lines show the boundaries that the energy change should satisfy if the process is diffusive. The diffusion coefficient was calculated based on the autocorrelation function of the energy time history.

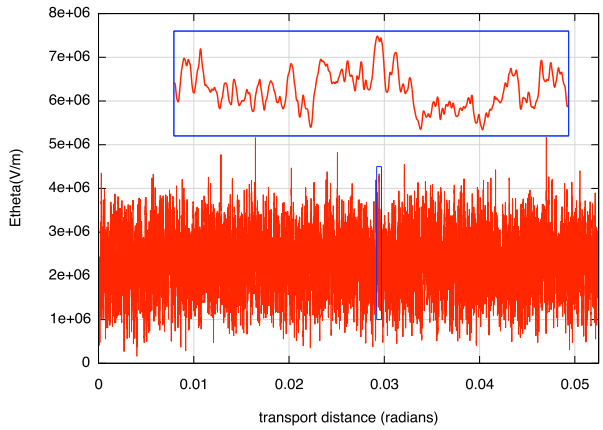

Figure 3: Tangential electric field at the center of a 1 GeV, 624M electron, 10x10x10 micron Gaussian bunch as it travels through 3 degrees of transport in a magnet with $\rho = 1\ m$. The narrow blue rectangle shows the domain of the inset, indicating that the fluctuations are very well resolved.

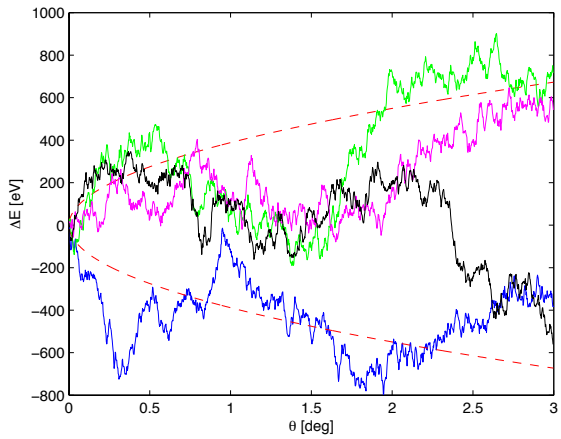

Figure 4: Energy change of a particle at the bunch center vs. propagation distance in a dipole, computed using the tangential electric field produced by the L-W solver. The 4 curves correspond to 4 different realizations of the distribution, each modeled with 6.24 billion particles. The red dashed lines show the boundaries that the energy change should satisfy if the process satisfies a diffusion equation.

Undulator Radiation

The following results are based on an undulator field that is given by,

$$B_y = B_0 \sin(k_w z), \tag{3}$$

where B_0 is the peak magnetic field in the y-direction, and where $k_w = 2\pi/\lambda_w$ is the undulator wave number. The particles travel along the z-direction and wiggle in the $x - z$ plane. This undulator model has no entrance or exit fields. However, since our simulation finds retarded quantities though high-order numerical integration [36], and not

ISBN 978-3-95450-126-7

by simple analytic approximations of trajectories in the undulator, this approach could equally well be applied to any undulator model whose fields are known analytically or are accessible by some numerical procedure.

Before presenting multi-particle simulation results, first consider the single-particle wakefields. Figure 5 shows the radiation component of E_x for a 125 MeV electron in an undulator with $B_0 = .025\ T$ and $\lambda_w = 3\ cm$. For these parameters the undulator K-value is 0.07. Note that the wavelength of the oscillation in Fig. 5 is approximately 0.25 micron, consistent with the expected value $\lambda_u/(2\gamma^2)(1 + K^2/2) = 0.249$ micron. Figure 6 shows the same quantities for a 14 GeV electron in an undulator with $B_0 = 1.3\ T$ and $k_w = 3\ cm$. For these parameters the undulator K-value is comparable to LCLS, $K = 3.64$. As seen in Fig. 6, the oscillation wavelength is approximately 1.53 Angstrom, consistent with the expected value 1.525 Angstrom. The wake for this case contains a lot of high harmonics unlike the previous case with K$\ll$1. The peak at the origin has a value 1.9e10 V/m, while the peak is only 263 V/m in the preceding figure. In both these figures it is understood that the fields are plotted at $t = 0$ when the electron is at $x = y = z = 0$ (and dx/dt chosen to produce a periodic orbit). The time-dependent wakefield is of course oscillatory with a frequency equal to the electron travel time over an undulator period.

Returning to the 125 MeV case, next consider a Gaussian bunch of electrons with rms sizes $\sigma_x = \sigma_y = \sigma_z = 10\ \mu m$ propagating along the undulator. The simulations that follow used only 24 million electrons. Figure 7 shows the radiation component of E_x produced by this bunch. It is interesting to note that the shot noise fluctuations are small, much smaller than those presented in [34] for the steady-state dipole CSR case. This is understood to be due to the fact that all the particles in the undulator radiate inside a cone that is smaller than the wiggle amplitude, in contrast to the dipole case for which the particles radiate along an extended path whose opening angle is much larger than the radiation cone angle. Though the fluctuations are small, they are still present, as seen in Fig. 8 which shows a zoomed-in view. Another observation in Fig. 7 is that the radiation pattern follows the bunch profile, i.e., there is no radiation produced ahead the bunch itself. Next, suppose that the Gaussian bunch is microbunched at a modulation wavelength $\lambda_{mod} = 0.249$ micron. Figure 9 shows the simulation results. Now it is clear that there is a strong radiation field ahead of the bunch. The oscillation is evident in the inset of the figure.

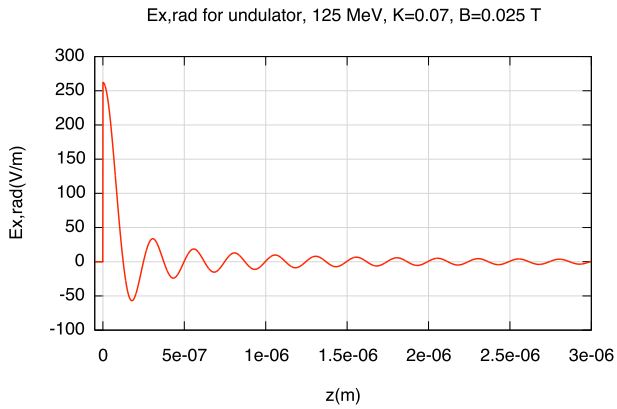

Figure 5: x-component of the single-particle radiation electric field, $E_{x,rad}$, for a 125 MeV electron inside an undulator with $B_0 = 0.025\ T$ and $\lambda_w = 3\ cm$, corresponding to an undulator K value of 0.07 .

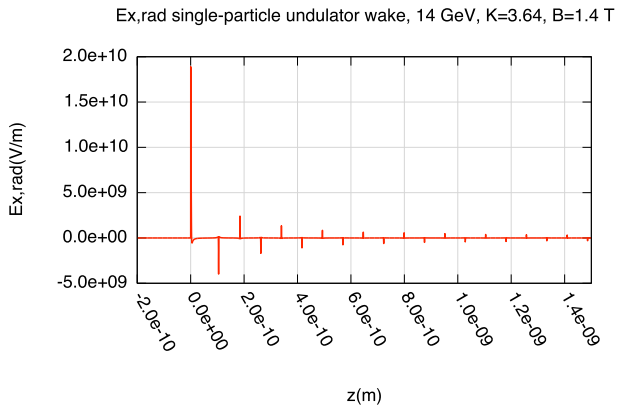

Figure 6: x-component of the single-particle radiation electric field, $E_{x,rad}$, for a 14 GeV electron inside an undulator with $B_0 = 1.3\ T$ and $\lambda_w = 3\ cm$, corresponding to an undulator K value of 3.64 .

CONVOLUTION-BASED SOLVER

The preceding results were all based on computing the exact L-W fields at an observation point by summing the contributions from N particles. In an N-body code this would scale as N^2 which is huge considering that N is of order 10^9. Now we will consider an alternative method to computing the L-W fields from a distribution of particles. Instead of summing N Green functions at a point (i.e. the L-W fields of N point-particles), we use just one Green function, and convolve it with the charge density at the observation time. This method has the advantage that just one Green function needs to be calculated (rather than of order 10^9). Furthermore, by zero-passing the charge density (see, e.g., the appendix in [38]), the convolution can be performing using an FFT-based method, which scales as $M \log M$, where M is the number of grid points. Also, the FFT can be performed using a parallel FFT routine. This approach is analogous to the method in widely used electrostatic or quasi-electrostatic PIC codes: in that approach, a Lorentz-transformation is used to transform to the bunch frame, the electrostatic fields are computed based on a *single* Green function, and the fields are transformed back to the laboratory frame. Such an approach is not always valid, as with some photoinjector simulations where the bunch energy spread is so large that there is no frame of reference

ISBN 978-3-95450-126-7

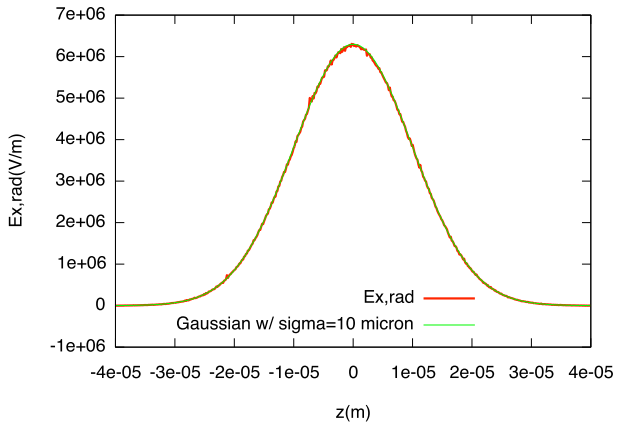

Figure 7: x-component of the radiation electric field, $E_{x,rad}$, inside an undulator with $B_0 = 0.025\ T$ and $\lambda_w = 3\ cm$, produced by a zero emittance, 125 MeV Gaussian bunch with rms sizes $\sigma_x = \sigma_y = \sigma_z = 10\ \mu m$. Also shown (in green) is the longitudinal bunch profile which is a Gaussian with rms size 10 micron. Notice that the radiation field does not extend beyond the front of the bunch, but simply follows the bunch profile.

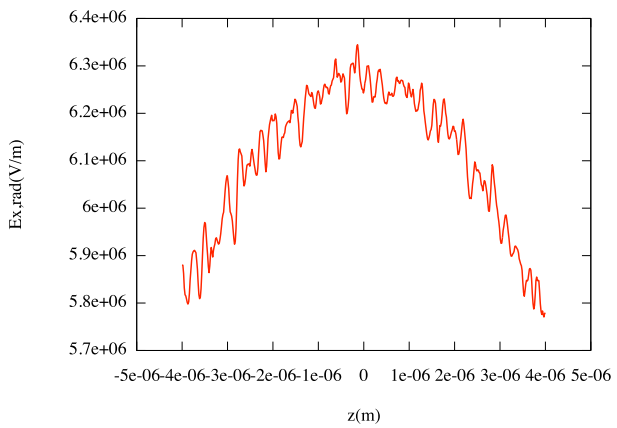

Figure 8: A zoomed-in view of $E_{x,rad}$ from Fig. 7. Notice the small-amplitude microstructure, which is due sampling the distribution at random with 24 million particles.

where longitudinal motion of all particles is nonrelativistic. In that case additional techniques (such as energy binning) are used to address this problem in quasi-electrostatic codes. Similar measures are likely to be found for L-W codes. In this report we do not address the domain of validity of the convolution-based L-W method. Instead we will just provide two examples that illustrate it's applicability.

Our convolution-based code makes use of the same subroutines for computing L-W fields as in our above-mentioned solver, but they are called only once. Also, the code uses domain decomposition, so each MPI process owns only a portion of the computational grid. The grid quantities are computed on a zero-padded domain that is twice the size of the physical domain in each dimension.

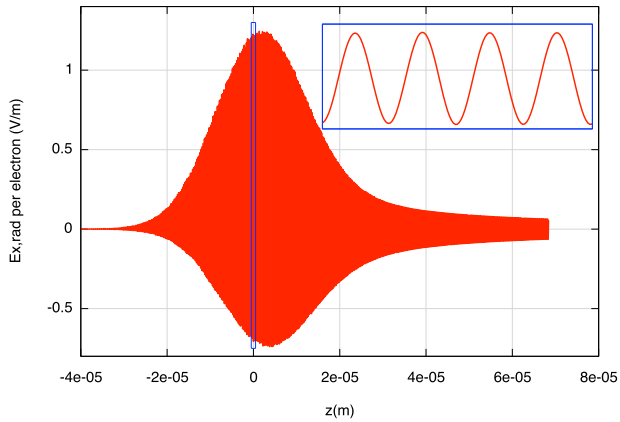

Figure 9: The same as Fig. 7, except that the Gaussian bunch is microbunched at a wavelength of 0.249 micron. Now the radiation field *does* extend beyond the front of the bunch. The narrow blue rectangle shows the domain of the inset.

After the L-W Green function is computed on the grid, a parallel FFT is used to convolve it with the charge density whose values have also been computed on a doubled grid. We use a parallel FFT package [39], which includes subroutines that we also use for domain decomposition. The underlying 1D FFTs are performed with FFTW [40]. A useful feature of [39] is that it allows for several types of decomposition: xyz, xy, yz, xz, x, y, z.

We tested our convolution-based L-W solver on two problems. Figure 10 shows the on-axis z-component of the radiation field produced by a 1x1x10 micron Gaussian ball in a dipole magnet with $\rho = 1\ m$. The plot shows convolution results using 64x64x2048 grid and a 128x128x4096 grid. Also shown is the result from L-W summation over 6.24 billion particles. The convolution method is in good agreement with the L-W sum. Figure 11 shows the x-component of the radiation field produced by a 125 MeV modulated Gaussian bunch in an undulator. The bunch is a $10 \times 10 \times 10$ micron rms modulated Gaussian with $\lambda_{mod} = 0.249\ micron$. The undulator field is given by Eq. 3 with $B_0 = 0.025\ T$ and $\lambda_w = 3\ cm$. The plot shows convolution results using a 64x64x32768 grid, along with results from L-W summation. The convolution-based results and the L-W summation are in excellent agreement.

DISCUSSION AND CONCLUSIONS

In this paper we have described the status of our massively parallel L-W solver. We presented several examples illustrating how such a code can be used to investigate phenomena that would be difficult or impossible to examine in other codes, such as shot-noise effects in a real-world distribution of particles, and 3D effects. We also described a convolution-based L-W approach and compared it with the L-W summation approach in two test problems. It should be noted that those comparisons made use of the "ordi-

ISBN 978-3-95450-126-7

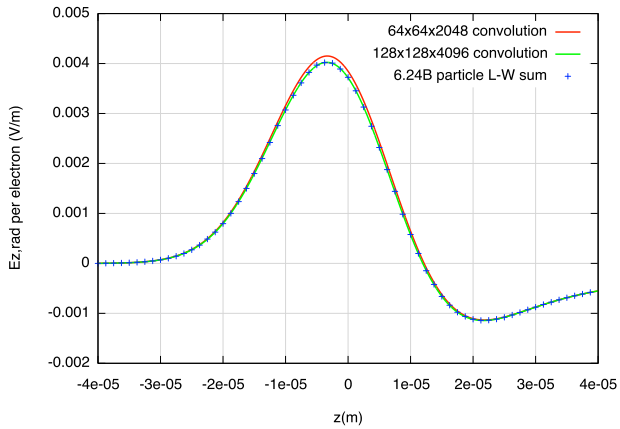

Figure 10: z-component of the radiation electric field, $E_{z,rad}$, for a 1 GeV Gaussian bunch inside a dipole with $\rho = 1\ m$. Comparison of Lienard-Wiechert summation over 6.24 billion particles, and convolution-based methods using a $64 \times 64 \times 2048$ grid and a $128 \times 128 \times 4096$ grid.

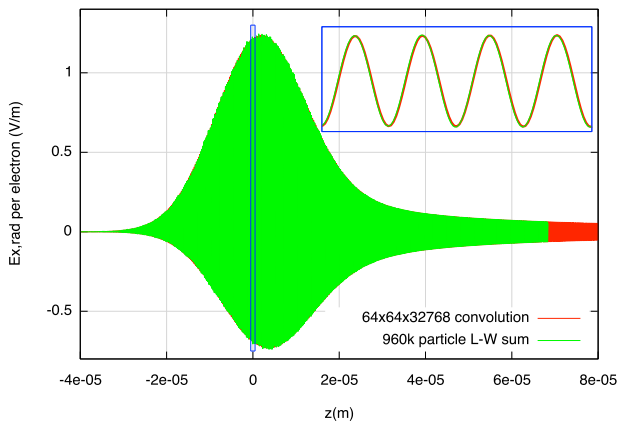

Figure 11: x-component of the radiation electric field, $E_{x,rad}$, for a 125 MeV modulated Gaussian inside an undulator with $B_0 = 0.025\ T$ and $\lambda_w = 3\ cm$,. Comparison of Lienard-Wiechert summation over 960k particles, and convolution-based methods using a $64 \times 64 \times 32768$ grid. The narrow blue rectangle shows the domain of the inset.

nary" L-W Green function, not the integrated Green function (IGF). Up to now IGF techniques have been applied only to problems for which the underlying Green function is known analytically [26, 27, 28, 29, 30, 41]. In the future it might be possible to apply IGF techniques to L-W problems through the use of numerical techniques such as adaptive quadrature.

Using the convolution-based approach, a L-W solver might be used to produce a self-consistent 3D code – a L-W particle-mesh (LWPM) code. In fact, such a code would not be just a code for modeling CSR effects, since the solver could in principle be embedded in any parallel beam dynamics code and be used to treat space-charge and synchrotron radiation effects along with other physical phenomena. The underlying beam dynamics code would need certain modifications, e.g. the ability to store particle histories and the ability to provide absolute particle data (not relative to a reference trajectory) "on the fly."

The approach described is analogous to that used in electrostatic and quasi-electrostatic PM codes. The primary difference is that the Green function on a grid is very much more time consuming to compute compared to the electrostatic case for which the Green function is just 1/r or its gradient. Also, given the nature of CSR fields and FEL radiation, the required grid resolution is much higher than has been dealt with previously. Nevertheless, given the progress in parallel computing resources, it is likely that within 3 years all the major US supercomputer centers will have systems with several hundred thousand cores. In such an environment, "routine" simulations could be done with 10,000 cores, with medium and large-scale runs requiring 100,000 or more. In such an environment, where massive parallelism increased performance by a factor of 10-100 , and where GPUs might be present that are able to increase code performance by a few tens, it seems quite possible that a parallel L-W particle-mesh code would be feasible.

ACKNOWLEDGMENTS

This research was supported by the LDRD program of Los Alamos National Laboratory, and by the Office of Science of the U.S. D.O.E., Office of Basic Energy Sciences, Accelerator and Detector Research and Development program. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

REFERENCES

[1] Robert D. Ryne and Salman Habib, "High Performance Computing for Beam Physics Applications," LA-UR-94-2904, July 1994, http://www.osti.gov/bridge/product.biblio.jsp?osti_id=10180045

[2] Robert D. Ryne and Salman Habib, "Beam Dynamics Calculations and Particle Tracking using Massively Parallel Processors," Proc. LHC'95, Montreux, Switzerland, Oct 16-21, 1995, http://cds.cern.ch/record/1120265/files/p119.pdf

[3] Robert D. Ryne, "Beam Dynamics Simulations using a Parallel Version of PARMILA," Proc. Linac'96, Geneva, Switzerland, August 26-30, 1996, http://accelconf.web.cern.ch/AccelConf/l96/PAPERS/MOP72.PDF

[4] R. D. Ryne et al., "The U.S. DOE Grand Challenge in Computational Accelerator Physics," Proc. Linac'98, August 1998, http://accelconf.web.cern.ch/accelconf/l98/linac98.html

[5] Robert D. Ryne, "Extraordinary tools for extraordinary science: the impact of SciDAC on accelerator science and technology," Journal of Physics: Conference Series vol. 46 issue 1 September 01, 2006. p. 171-189, http://iopscience.iop.org/1742-6596/46/1/025/pdf/1742-6596_46_1_025.pdf

ISBN 978-3-95450-126-7

[6] Ji Qiang et al., "SciDAC Advances in Beam Dynamics Simulation: From Light Sources to Colliders," Journal of Physics: Conference Series vol. 125, SciDAC 2008, 13-17 July 2008, Seattle, WA, USA, http://iopscience.iop.org/1742-6596/125/1/012004/pdf/1742-6596_125_1_012004.pdf

[7] http://tesla.desy.de/~meykopff/

[8] http://amac.lbl.gov/~jiqiang/BeamBeam3D/

[9] http://www.desy.de/xfel-beam/csrtrack/

[10] http://aps.anl.gov/elegant.html

[11] http://genesis.web.psi.ch

[12] http://www.muonsinternal.com/muons3/G4beamline

[13] http://www.cap.bnl.gov/ICOOL/

[14] http://amac.lbl.gov/~jiqiang/IMPACT/

[15] http://amac.lbl.gov/~jiqiang/IMPACT-T/

[16] R. D. Ryne et al., "Recent Progress on the MaryLie/IMPACT Beam Dynamics Code," Proceedings of ICAP 2006, Chamonix, France, http://accelconf.web.cern.ch/accelconf/ICAP06/PAPERS/TUAPMP03.PDF

[17] A. Adelmann et al., "The Object Oriented Parallel Accelerator Library (OPAL), Design, Implementation and Application," Proc. ICAP09, San Francisco, CA, USA

[18] N. C. Ryder, D. J. Scott, and S. Reiche, "SPUR: A New Code for the Calculation of Synchrotron Radiation from Very Long Undulator Systems," Proc. EPAC08, Genoa, Italy

[19] J. Amundson et al., "Synergia: A 3D Accelerator Modelling Tool with 3D Space Charge. Journal of Computational Physics," Volume 211, Issue 1 , 1 January 2006, Pages 229-248.

[20] P. Ostroumov at al., "TRACK, The Beam Dynamics Code", Proc. PAC2005

[21] L. Giannessi and M. Quattromini, "TREDI Simulations for High-brilliance Photoinjectors and Magnetic Chicanes," Phys.Rev. ST-AB 6 120101 (2003)

[22] D.P. Grote, A. Friedman, I. Haber, Methods used in WARP3d, a Three-Dimensional PIC/Accelerator Code, Proc. of the 1996 Comp. Accel. Physics Conf., AIP Conference Proceedings 391, p . 51 (1996).

[23] Ji Qiang and Robert D. Ryne, "Parallel 3D Poisson solver for a charged beam in a conducting pipe," Comp. Phys. Comm. 138, pp 18-28, 2001.

[24] G. Poplau, U. van Rienen, K. Flottmann, "The Performance of 3D Space Charge Models for High Brightness Electron Bunches," TUPP103, Proc. EPAC08, Genoa, Italy

[25] A. Adelmann, P. Arbenz, and Y. Ineichen, "A Fast Parallel Poisson Solver on Irregular Domains Applied to Beam Dynamic Simulations," arXiv:0907.4863, February 2010

[26] Robert D. Ryne, "A new technique for solving Poissons equation with high accuracy on domains of any aspect ratio, ICFA Beam Dynamics Workshop on Space-Charge Simulation, Oxford, April 2-4, 2003.

[27] J. Qiang, M. A. Furman, R. D. Ryne, W. Fischer, K. Ohmi, "Recent advances of strong-strong beam-beam simulation," Nuclear Instruments & Methods in Physics Research A558, 351, (2006).

[28] D.T. Abell et al.,"Three-Dimensional Integrated Green Functions for the Poisson Equation," Proc. PAC2007, 3546-3548.

[29] R. D. Ryne, B. Carlsten, J. Qiang, and N. Yampolsky, "A model for one-dimensional coherent synchrotron radiation including short-range effects," http://arxiv.org/abs/1202.2409

[30] C. Mitchell, J. Qiang, and R. Ryne, "A Fast Method for Computing 1-D Wakefields due to Cohereent Synchrotron Radiation," Nucl. Instrum. Methods Phys. Res., Sect. A 715, 119 (2013)

[31] Robert D. Ryne, Ji Qiang, and Salman Habib, "Computational Challenges in High Intensity Ion Beam Physics," Proc. 2nd ICFA Advanced Accelerator Workshop on the Physics of High Brightness Beams, World Scientific, James Rosenzweig & Luca Serafini, Eds., Nov. 9-12, 1999, Los Angeles, CA, USA

[32] J. Qiang, "Start-to-End Simulation of a Next Generation Light Source Using the Real Number of Electrons," this proceedings

[33] J. B. Murphy, S. Krinsky, and R. L. Gluckstern, Longitudinal Wakefield for an Electron Moving on a Circular Orbit, Particle Accelerators, Vol 57, pp. 9-64, 1997.

[34] R. D. Ryne et al., "Large-Scale Simulation of Synchrotron Radiation using a Lienard-Wiechert Approach," Proc. IPAC 2012, New Orleans, LA, USA, May 2012

[35] William H. Press, Saul A. Teukolsky, William T. Vetterling, Brian P. Flannery, "Numerical Recipes in Fortran 77: The Art of Scientific Computing," Cambridge University Press

[36] http://www.unige.ch/~hairer/prog/nonstiff/dop853.f

[37] http://crd-legacy.lbl.gov/~dhbailey/mpdist/

[38] Robert D. Ryne, "On FFT-based convolutions and correlations, with application to solving Poisson's equation in an open rectangular pipe," arxiv.org/pdf/1111.4971

[39] http://www.sandia.gov/~sjplimp/docs/fft/README.html

[40] http://http://www.fftw.org

[41] Robert D. Ryne, "An Integrated Green Function Poisson Solver for Rectangular Waveguides," Proc. IPAC 2012, New Orleans, LA, USA, May 2012

ISBN 978-3-95450-126-7