IMPROVED SIMULATIONS OF PHOTOINJECTORS FOR HIGH-AVERAGE-CURRENT FREE-ELECTRON LASERS

C. L. Bohn[1], D. Mihalcea[1], I. V. Pogorelov[2], I. V. Sideris[1], and B. Terzić[1]

[1]Beam Physics and Astrophysics Group, Northern Illinois University, DeKalb, IL 60115

[2]Accelerator and Fusion Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720

Final Submission

Abstract

We report a successful implementation of a three-dimensional wavelet-based solver for Poisson's equation with Dirichlet boundary conditions, optimized for use in particle-in-cell beam dynamics simulations. We explain how the new algorithm works and the advantages it brings to accelerator simulations. The solver is integrated into a full photoinjector-simulation code (IMPACT-T), and the code is then benchmarked by comparing its output against that of other codes (verification) and against laboratory measurements (validation). To enable detailed verification, we developed and applied a new technique that involves quantifying chaos in particle orbits. We also simulated the AES/JLab photoinjector using a suite of codes. This activity revealed certain performance limitations and their causes. Finally, with an eye on eventually doing end-to-end simulations of driver accelerators for high-average-power free-electron lasers, we

sketch a path for future code improvements and applications for which this code is particularly, and possibly even uniquely, useful.

Keywords: photoinjector, space charge, Poisson solver, wavelets, free-electron laser

NOMENCLATURE

E = particle kinetic energy

f = function proportional to charge density

L = matrix representation of the Laplacian operator

N_{gc} = number of grid points along each coordinate axis

N_g = total number of grid points

T = noise threshold

U = true potential

U^i = computed potential after i iterations

u = space-charge potential

X_G = Gaussian-distributed signal

X_P = Poisson-distributed signal

(v_x,v_y,v_z) = particle velocity components

(x,y,z) = particle coordinates

γ = beam energy divided by electron rest mass

Δ = Laplacian operator

ε_n = normalized root-mean-square (rms) transverse emittance of the beam

κ = condition number

λ = wavelength of FEL light

= standard deviation of numerical noise

$|A|_2$ = 2-norm of (arbitrary) argument A

1.0 INTRODUCTION

The work reported herein marks a significant first step toward a faster, more accurate space-charge algorithm for improved simulations of driver accelerators for MW-class free-electron lasers (FELs). Such lasers are envisioned, e.g., for defense of naval vessels against cruise missiles in littoral combat.

The output power of a FEL is proportional to the power of the electron beam, which is in turn proportional to the product of the beam's average current and kinetic energy. For example, suppose the beam has energy 100 MeV and the FEL systems are able to convert 1% of the beam power into light. Achieving a MW-class FEL then requires amp-level average beam current. In turn, the electron bunches that comprise the beam must be of high charge, at the nanocoulomb (nC) level.

Another requirement is that essentially all of the electrons participate in the lasing process. This means the electrons in each bunch must fit within the optical mode of the FEL resonator. This sets a limit on the size of the phase space that the electron bunch can span. Specifically, the normalized root-mean-square (rms) transverse emittance of the beam ε_n needs to be such that $\varepsilon_n \gamma\lambda/(4\pi)$, where γ is the beam energy divided by the electron rest mass, and λ is the wavelength of the FEL light. For example, efficient lasing at ~1 μm wavelength, as is desired for good atmospheric transmission, requires roughly 10 μm emittance.

Coulomb self-forces within the bunch, i.e., space charge, work against achieving the required current and emittance of the electron beam. Nonlinear collective forces act rapidly to redistribute the electrons. Because the beam from the source is generally far from equilibrium, it carries free energy that, as it is redistributed and thermalized, irreversibly expands the phase space; *the emittance grows*. Space charge is thus a key concern regarding electron sources. In the laboratory frame, the space-charge force decreases inversely with the square of the beam energy. For the transverse component, this arises from the partial cancellation between the self-magnetic and self-electrostatic forces; while for the longitudinal component, it is due to Lorentz contraction [1]. Nonetheless, for nC bunch charges, space charge remains *the* important dynamic at beam energies up to ~100 MeV, i.e., through the entire driver accelerator [2].

In addition to degrading the rms properties of the beam's phase space, space charge also generates *beam halo*, a diffuse population of electrons that lie far from the beam axis. This is a special concern for high-average-current machines. Machine protection involves an automated system that will turn off the beam if the loss on the accelerator walls exceeds a certain threshold current. Typically the threshold current is set to ~1 μA/m, a value determined from concerns about accidentally burning through the beamline vacuum chamber or damaging sensitive instrumentation and electronics via x-ray radiation produced from electrons impinging the chamber. For a MW-class FEL, the threshold current is a tiny $\sim 10^{-6}$ of the total beam current per meter. In a military

application in which the machine would normally operate in a continuous-wave mode for relatively short times, the threshold current could perhaps be raised. However, the technology will need to be mature before its operational value can be confidently specified. To reach maturity requires a protracted time of commissioning and optimizing prototypical machines, during which the threshold current must remain at the μA/m-level.

Accelerator designs traditionally derive from procedures that are known to be effective in controlling root-mean-square (rms) properties of the beam. These procedures have proven to be effective for applications involving relatively small currents and/or relatively large phase spaces. However, for large currents and small phase spaces as are required for high-average-power near-infrared FELs, controlling the rms properties of the beam is still necessary, but not sufficient. Thus, the conventional design tools *must* be augmented.

Concerns about space charge are most serious in regions where the beam energy is lowest, i.e., in the electron source and injector. Hence, our objective is to develop an improved space-charge algorithm for simulating electron-beam dynamics in high-brightness, high-average-current photoinjectors, particularly those envisioned for generating average currents in the range 0.1-1 A. Our ultimate goal is to have a fast, accurate code that couples the longitudinal and transverse dynamics, accounts for rapid evolutionary time scales, and quantifies details in the beam's structure. We have used the

improved code to model the 100 mA direct-current (dc) photoinjector that Advanced Energy Systems, Inc. (AES) and Jefferson Laboratory (JLab) are building [3]. Concerning applications to radiofrequency (rf) photoinjectors, Northern Illinois University jointly operates, and thus has ready access to, the Photoinjector Laboratory at Fermilab. This Fermilab/NICADD photoinjector (NICADD stands for Northern Illinois Center for Accelerator and Detector Development) generates low-average-current pulsed beams, but these beams consist of high-brightness, high-charge bunches (1-10 nC with, respectively, ~3-10 μm normalized rms transverse emittance) [4]. We therefore tested and benchmarked the new code against this machine. In turn, we modeled *both* dc and rf photoinjectors.

2.0 SOLVING POISSON'S EQUATION WITH WAVELETS

Preserving a hierarchy of scales in the time-dependent space-charge potential is dynamically important. Recent research has revealed that nonlinear, time-dependent forces commonly establish large populations of globally chaotic orbits in beams that are out of equilibrium, and such orbits can even be present in thermal-equilibrium beams [5,6]. When present, these chaotic orbits mix exponentially throughout their accessible phase space with a time scale of only a few orbital periods, i.e., very much faster than via collisional relaxation. The presence of colored noise due to space-charge fluctuations and/or machine imperfections can, when combined with parametric resonance associated

with low-order oscillatory modes, generate large halos over just a few oscillation periods (which are known as 'space-charge-depressed betatron periods') [7]. Thus, all scales are potentially important to the dynamics.

Multiresolution analysis is a mathematical discipline designed to preserve the physically important scales. It hinges on the use of 'wavelets', which are mathematical entities that are defined over a finite range (i.e., they have the property of 'compact support'). One-dimensional and two-dimensional examples of wavelets appear in Fig. 1. Wavelets are also scaleable; a superposition of wavelets distributed over a hierarchy of scales will reproduce very complicated structure, even with a relatively few number of wavelet coefficients. In view of these considerations, we developed an algorithm to solve Poisson's equation linking the charge density in a beam bunch to the space-charge potential. The use of wavelets generates potential-density pairs that correctly preserve the hierarchy of scales in the beam, thereby enabling accurate computations that apply well beyond predictions of conventional rms beam properties. The algorithm is comprehensively documented elsewhere [8]; in what follows we summarize its salient features.

Gaining insight into the dynamics of multiparticle systems, such as charged-particle beams, usually involves *N*-body simulations, and this is certainly true regarding accelerator design. *N*-body codes can be grouped into three main categories: (1) direct summation, (2) tree, and (3) particle-in-cell (PIC). Direct summation codes are

prohibitively expensive for large systems since their computational cost scales as N^2. Tree codes use direct summation for nearby particles and invoke statistical arguments for contributions of particles farther away. PIC codes incorporate a computational grid into which particles are binned, thus resulting in a coarse-grained, discretized particle distribution. The potential associated with such discretized distributions is computed by solving Poisson's equation on the grid. Finally, the forces needed to advance each individual particle are computed by interpolation from the discretized potential. We developed a wavelet-based algorithm for solving Poisson's equation in PIC codes.

To be accurate, algorithms for solving Poisson's equation *must*:

1. *include multiscale dynamics*, because even fluctuations on smallest scales can lead to global instabilities, as exemplified by halo formation and microbunching [7,9];

2. minimize numerical noise due to: (1) the number of macroparticles N used to sample the phase-space distribution function in the N-body simulation being several orders of magnitude smaller than the number of real particles in the physical system that is being modeled (*graininess* of the distribution function), and (2) the continuous physical problem being restricted to a discrete, finite computational grid (*discreteness* of the computational domain); and

3. be *efficient* in terms of computational speed and storage requirements, without compromising accuracy.

With multiresolution analysis (wavelets), all three features are attainable.

An additional and *strong* motivation for developing a wavelet-based Poisson solver is to take advantage of wavelet *compression*. Doing so enables compact storage and easy recall of the beam history. In turn, it facilitates simulations whereby beam history is important. An example is simulating the influence of coherent synchrotron radiation (CSR) on the beam as it transits magnetic bends in the accelerator lattice. Such bends are unavoidable in, for example, recirculating linear accelerators and bunch compressors, both of which are inherent to driver accelerators for high-average-power FELs. To compute the scalar and vector potentials, and hence the force acting on each particle, requires integrating over the history of the evolving charge density to account for retardation arising from the finite speed of light [9]. The end result would be a fully three-dimensional simulation of beam dynamics under the influence of coherent synchrotron radiation (CSR), something that has been notoriously difficult to achieve. It is with such simulations in mind, combined with a desire to preserve accurately the influence of the hierarchy of spatial scales on the space-charge force, and hence on the overarching beam dynamics, that we proceeded. The end result was a fully three-dimensional (3D) wavelet-based solver satisfying all of the above requirements.

To reiterate, the mathematical details of the 3D wavelet-based solver are intricate and can be found elsewhere [8]. Here, we provide a qualitative description of the method. We built on the work of Beylkin and co-workers [10-12] to design and

implement an iterative algorithm subject to general (inhomogeneous) Dirichlet boundary conditions. A combination of three circumstances makes an iterative solver particularly attractive: (1) the discretized Laplacian operator (Δ) remains sparse in a wide variety of wavelet bases; (2) preconditioners for the Laplacian operator exist that are effectively diagonal in a wavelet basis; (3) the fact that Poisson's equation has to be solved at each successive time step, with the source term not changing significantly from one time step to the next. This simplifies the choice of the initial 'guess' (one can, e.g., use the potential computed at the previous time step as the initial guess for the present-time potential).

The flow chart in Fig. 2 summarizes our algorithm. One begins by introducing a rectangular computational grid that tightly envelops the charge-density distribution and has boundaries that may or may not coincide with physical boundaries (e.g., the boundaries of the beamline vacuum chamber). Computing the potential on the surface of the computational grid provides Dirichlet boundary conditions for Poisson's equation in physical (configuration) space. One then applies a Fast Discrete Wavelet Transform (FDWT) [13] to the source term and the Laplacian operator, solves Poisson's equation in wavelet space using the Preconditioned Conjugate Gradient (PCG) method, and then applies the inverse FDWT to recover the potential in configuration space.

We now briefly state how we implement the boundary conditions, and then we present the basic details of the PCG algorithm that is the heart of our solver.

2.1 *Implementation of Boundary Conditions*

A boundary condition (BC) often encountered in practical beam-dynamics simulations is that of zero potential on the walls of an infinite pipe of rectangular cross-section, with open BCs in the longitudinal (z) direction. The Green function for such BCs is known analytically, and the corresponding potential can readily be found [14,15]. In fact, conventional codes often employ a Green function and fast Fourier transform (FFT) to solve Poisson's equation over the whole grid volume. We use this method to compute the potential over the surface of the computational grid, but then employ wavelets to solve Poisson's equation throughout the grid interior. One key advantage for doing so is that the wavelet transform and its inverse are faster to compute than the FFT and its inverse.

2.2 *Preconditioned Conjugate Gradient*

The three-dimensional (3D) Laplacian Δ is a positive-definite operator. The wavelet-decomposed Laplacian is sparse provided one applies to it a thresholding procedure whereby all coefficients with magnitudes below a user-specified threshold are set equal to zero. In this setting, a conjugate-gradient (CG) algorithm [16] provides an efficient way to invert the operator, with the rate of convergence depending on the 'condition number' κ as

$$\left|U - U^i\right|_2 \le \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^i \left|U\right|_2 .$$

Here, U is the true potential, U^i is the computed potential after i iterations, and $|\ldots|_2$ denotes a 2-norm. A major advantage of working in a wavelet basis is the existence of a diagonal preconditioner that lowers κ from $O(N_{gc}^2)$ to $O(N_{gc})$, with N_{gc} denoting the number of grid points along each coordinate axis. Our PCG algorithm is a combination of the operator formulation of the conjugate-gradient method with a wavelet-basis preconditioner.

2.3 *Denoising by Wavelet Thresholding*

The thresholding operation performed both on operators and potential/density data sets is the simplest possible way to effect simultaneous compression and denoising. After transforming the noisy charge distribution into wavelet space, the signal (physical component) is generally represented by a smaller number of large coefficients, while the numerical noise is largely mapped to many small wavelet coefficients. Wavelet thresholding is a process whereby the contribution of the wavelet coefficients deemed to represent noise is eliminated. Simply put, the wavelet coefficients whose magnitudes are below a certain noise threshold T are set to zero. One must choose the noise threshold carefully, avoiding two extreme perils: on one end, over-smoothing the distribution by overestimating the threshold, and on the other end, not removing enough unphysical small-scale noise by underestimating the threshold.

Studies of wavelet denoising usually concern distributions contaminated with additive (distribution-independent) Gaussian (white) noise [17,18]. The most widely used noise threshold is given in terms of the standard deviation of the noise as $T = (2\log N_g)^{1/2}$, where N_g is the total number of grid points. This is a universal threshold for signals with Gaussian noise, which means it leads to noise removal that is within a small factor of ideal denoising. However, noise in PIC simulations is Poisson-distributed and distribution-dependent. The basic assumption underlying denoising techniques is that, regardless of the details of the noise, small-scale fluctuations due to noise map to small-scale members of the wavelet family. That said, one can transform a Poisson-distributed signal X_P to a Gaussian-distributed signal X_G with unit variance and mean using a variance-stabilizing transformation due to Anscombe: $X_G=2[X_P+(3/8)]^{1/2}$. Applying this 'Anscombe transform' produces a bias in the data; however, it can be removed by ensuring the denoised and noisy data have the same mean (in simulations, this is equivalent to enforcing charge conservation). Details on this procedure and results of tests we did are documented elsewhere [8]. When the number of particles per cell in the PIC simulation is too low, the noise exhibits a departure from the Poisson distribution, in which case the Anscombe transform breaks down. Typically one needs at least three particles per cell in the simulation to avoid this difficulty.

3.0 CODE VERIFICATION AND VALIDATION

To verify a new code means to compare its output to that of established codes. The idea is to ensure the new code correctly incorporates the physics (here, the concern is space-charge physics) while pinpointing differences in code predictions. These differences can, of course, sometimes be subtle.

Since our primary interest herein is photoinjectors, we compare codes in that context. We do so by comparing the detailed predictions concerning the beam's evolution in (1) the three-dimensional configuration space alone, and (2) the full six-dimensional phase space. Here, we concentrate on simulating the Fermilab/NICADD photoinjector [19]. This machine includes a 1.625-cell radiofrequency (rf) electron gun enveloped by three solenoids for emittance compensation, followed by a nine-cell superconducting rf booster cavity, a 'chicane' of four dipole magnets for bunch compression, and a transport line composed of quadrupole magnets. The beam kinetic energy out of the gun is ~4 MeV, and out of the cavity is ~15 MeV. A schematic of this photoinjector appears in Fig. 3.

We begin by comparing beam moments and plots of configuration space computed using several codes, including our new wavelet-based code. We then present a new way of measuring chaos in particle orbits that we developed. This method enables one to construct detailed phase-space plots and to make movies of the beam's evolution in phase space. Finally, we apply our new method to compare phase-space evolution computed from the 'conventional' version of IMPACT-T with our wavelet version. The code

IMPACT-T was developed at Lawrence Berkeley National Laboratory (LBNL) for modeling accelerators of high-brightness beams [14,15]. It is a particle-in-cell code in which the Poisson solver incorporates a 'conventional' algorithm involving a Green function and FFTs. The code is modular, facilitating the replacement of subroutines such as the Poisson solver. Our new code is the same IMPACT-T, but with the conventional Poisson solver replaced by our wavelet-based Poisson solver.

3.1 *Comparing Beam Moments and Beam Spots*

To reiterate, we tested our wavelet-based code in a realistic setting by modeling the Fermilab/NICADD photoinjector. In what follows, we compare simulation results from our wavelet version of IMPACT-T to those from the conventional version. To verify agreement between the space-charge computation of the two codes, we tested them on two highly nonuniform transverse initial distributions of 1 nC bunches: (1) a considerably nonuniform and asymmetric distribution generated from a real laboratory snapshot of the laser-illuminated photocathode in an actual experiment under suboptimal conditions (henceforth called '*Distribution 1*'); and (2) a 5-beamlet quincunx distribution that can be made by masking the photocathode (henceforth called '*Distribution 2*'). We expect that the nonuniformity and asymmetry of the two initial transverse distributions will strongly enhance space-charge effects vis-à-vis a uniform distribution, thereby 'stressing' the Poisson solvers. We compare results computed from these two distributions regarding: (1) rms properties of the beam, (2) phase-space detail, and (3) computational speed.

Figure 4 shows the rms properties of the beam in the Fermilab/NICADD photoinjector for *Distribution 1* computed with conventional IMPACT-T (black lines), and with IMPACT-T with PCG without thresholding (green line), with thresholded after Anscombe transform (blue line), and with thresholding but without Anscombe transform (red line). Figure 5 shows the same for *Distribution 2*. Agreement of rms properties between the versions of IMPACT-T is excellent, to within a few percent.

For *Distribution 1*, the beam size in the laboratory experiment was measured at different positions along the beam line. Figure 6 compares the measurements against values computed using conventional IMPACT-T (red line) and IMPACT-T with PCG (blue line). This represents a first step toward code validation. Taken together, these results clearly demonstrate that simulations using both versions of IMPACT-T are in excellent agreement regarding the computation of beam moments. They also match the measured rms beam sizes reasonably well. We have also done some validation against the longitudinal charge distribution; details are documented elsewhere [20]. Regarding this work, both PARMELA and conventional IMPACT-T were the codes being validated; insofar as we have verified our wavelet version of IMPACT-T against the conventional one, we have also indirectly validated it.

For the two distributions, Figures 7 and 8 show projected transverse cross-sections of the beam at different positions down the beamline. Detailed agreement between the configuration spaces computed with the two versions is clearly very good.

The computational speed of IMPACT-T with PCG is presently comparable to that of conventional IMPACT-T. We found that the fastest simulations are usually achieved with wavelets having smallest compact support. There are rare exceptions, and these pertain to cases wherein convergence of the PCG algorithm at each timestep requires fewer iterations. Further details regarding the computational efficiency of the wavelet-based code can be found elsewhere [8].

3.2 *Measuring Chaos in Particle Orbits*

The historically favored method for characterizing orbits as regular or chaotic is the calculation of Lyapunov exponents [21]. They quantify the local instability of orbits and thus connect directly to sensitivity of initial conditions. Another class of measures relies on information derived by frequency analysis of orbits. For example, a Fourier spectrum provides a picture of the number and strength of frequencies associated with an orbit. Intuitively, one expects the more frequencies of sizable power, the more complicated the orbit should look. Chaotic orbits are characterized by continuous spectra, while regular orbits are characterized by discrete ones. In this context, sophisticated measures have been developed, an example of which is Laskar's frequency-map technique [22]. There are of course other measures not based on either of these two basic ideas: KS-entropy [20] is the most typical example.

There have always been two main concerns related to chaotic measures: (1) the accuracy of the characterization, and (2) the speed, i.e. the length of time one must evolve

an orbit to get a reliable characterization. For time-independent regimes (where energy is conserved, so regular orbits remain always regular, and chaotic orbits remain always chaotic), the longer the evolution time, the more accurate the characterization. However, the shortest evolution time required for confident characterization depends on the measure employed. The traditional Lyapunov exponents may need hundreds, or even thousands, of orbital periods to converge. The most sophisticated measures today have ostensibly lowered this limit to about 30 orbital periods, but this may depend significantly on both the chosen model and the specific orbit.

In view of these considerations, two major questions arise. First, is it possible to do better than 30 orbital periods? In some contexts, the life of systems is very short; researchers do not always have the luxury of long evolution times. Second, is it possible to analyze time-dependent systems, such as an electron beam? In a time-dependent regime, energy is not conserved, and orbits essentially experience a different potential at every different moment. As a result they can experience both regular and chaotic epochs, a phenomenon known as 'transient', or 'intermittent', chaos. Dissecting orbits with the aforementioned established measures to detect short-lived chaotic or regular epochs is either impossible or inefficient.

With these questions and considerations in mind, we developed a new measure based on pattern recognition. It treats a signal (e.g., an orbit) not as one entity but as a *series of distinct epochs*. In this sense, it focuses on *local*, epochal characterization, instead of the

usual approach of global characterization. Pictures of the phase space can be made, including essential details, typically with as little as just 10 orbital periods. Moreover, this measure applies to both time-independent and time-dependent systems without any change in its logic or design. A detailed account of the algorithm, accompanied by numerous example applications, appears elsewhere [23]; here we concentrate on applications of the method to photoinjector simulations.

Perhaps the most attractive feature of our method, and the reason we developed it, is that it enables one to paint detailed phase-space portraits and make movies of the beam as it evolves in phase space. Thus, by comparing movies made from the output of different codes, one accomplishes *detailed* code verification. This is a much more sensitive verification than can be achieved by comparing computed rms properties (like beam moments) alone. A second very attractive feature is that, by making such movies, one can identify and follow the evolution of instabilities, particularly those that start on a localized scale and grow to macroscopic scales. In other words, it serves as a 'phase-space microscope'.

One example illustrating the high resolution of the technique is its application to the Henon-Heiles potential, which is a time-independent potential having the form $U(x,y) = (1/2)(x^2+y^2)+ yx^2-(1/3)y^3$. The number of chaotic orbits this potential admits depends on the choice of the energy. When the energy is small there are no or few chaotic orbits. As the energy increases, the number of chaotic orbits increases and eventually dominates.

The choice of the energy for the example presented here is E=0.125. For this energy about fifty percent of a random set of initial conditions corresponds to regular orbits; the other fifty percent to chaotic.

Corresponding phase-space contour plots appear in Fig. 9. Note the contours in the chaotic sea. They indicate the ability to decipher the *degree* of chaoticity, not just the presence of chaos. This is very important physically. The framework of a real system consists of the regular regions and the sticky chaotic orbits loitering nearby them. The sticky chaotic orbits spend considerable time in the vicinity of regular regions before pulling away and migrating through the chaotic sea (and they can become intermittently stuck and unstuck). The thickness of the sticky regions is an important part of determining the macroscopic properties and subsequent evolution of the system. For systems that contain almost exclusively chaotic orbits, the sticky epochs of chaotic orbits will form the backbone of the system. Producing Fig. 9 required evolving 20,000 orbits for a long time (~300 periods) to get high resolution.

If one were to evolve orbits for much shorter times, the phase-space plots would look like the examples in Fig. 10. Here again, 20,000 orbits are tracked, but only up to ~10, 20, and 30 periods. Clearly the algorithm works quite well given 30 periods. With 15 periods the picture is still clear although small regular islands can no longer be identified. Even with just 10 periods the picture still provides reasonably complete 'global' information in the sense that it clearly identifies many important regions of regularity.

3.3 *Comparing Beam Evolution in Phase Space*

To apply our new measure of chaos, the orbits must be integrated over a sufficiently long time to enable patterns to be recognized. This time is 'long' compared to the transit time of an electron beam through a typical photoinjector, but 'short' compared to integration times required by previous measures. To apply our method toward verifying photoinjector codes, we add a transport line at the end of the photoinjector to enable tracking orbits for a sufficiently long time. By contrast, in a full-scale linac, the integration time may already be sufficiently long; it depends on the actual machine design.

What we did for a first application toward code verification was to 'install' an artificial periodic transport line (quadrupole FODO array) at the exit of the Fermilab/NICADD photoinjector. It provided focusing that was sufficiently strong to contain the beam, and the transport length was sufficiently long to obtain unambiguous results from the measure.

The same initial distribution of particles was then integrated with the conventional version of IMPACT-T and the new wavelet-based CPG version (with denoising turned off). Then phase-space projections computed with the two codes were compared. The idea is that if the codes are equivalent, then the phase-space projections would be identical, and the time step at which certain orbits become chaotic would likewise be identical. Any deviations from 'identical' are indications of differences between code

algorithms (in this case, differences between the respective Poisson solvers). Now, in practice, not all of the individual orbits calculated with the different codes can be identical; this is especially true during chaotic epochs. The idea is that if the underlying dynamics of the two beams are similar, then the overall phase-space projections and time scales will look the same. This is the first time a test that is completely dynamical in nature has been put to use; usually comparative tests between codes are only statistical in nature, but comparisons based on dynamical evolution are far more subtle and detailed.

Figure 11 shows snapshots of the (y,v_y)-projection of transverse phase space of the same beam (same initial conditions) evolved with the conventional version (top row) and wavelet-based version (bottom row) of IMPACT-T. These are the last four frames of 'movies' of the evolution. A cross-comparison of all frames preceding these shows essentially no differences because the orbits are all regular at these earlier times. This is also true for the two panels corresponding to t=4750-5000 in Fig. 11, wherein the orbits are still regular (blue). Subsequently the dynamics begins to change; in the panels corresponding to t=5000-5250, weak chaos (green) is emerging. It is important that chaos emerges in the *same* frame in both simulations. Later, in the frames t=5250-5500, more weak chaos emerges, and now we see slight differences between the two computations. Later, in the frames t=5500-5750, strong chaos (red) emerges. Again, the strong chaos shows up in the same frame in both simulations, and the phase spaces differ

only slightly. This is a clear indication that these two versions of IMPACT-T are in close detailed agreement; the underlying Poisson solvers and particle pushers yield the same dynamics. It would have been surprising (and unnerving) to see major differences because, since denoising was turned off, the level of numerical noise (PIC and macroparticle noise) is the same for both codes. Thus, we have compared 'apples to apples' and found essentially the same 'applesauce' is produced by the two distinctly different codes.

To reiterate, our primary goal in searching for orbital chaoticity was to verify codes. This is a highly detailed verification in that we look at properties of particle trajectories rather than just beam moments. Because we had to add an artificial transport line to achieve the time scale necessary for the onset of chaoticity in the output beam of the Fermilab/NICADD photoinjector, one could be tempted to conclude that chaos is unimportant in regard to designing such photoinjectors. A conclusion like this is fraught with danger. There is no universal definition, or design, of a 'photoinjector'. The important consideration is time scales. As a rule of thumb, chaos in beams with space charge requires several betatron oscillations to induce mixing and thereby affect the beam properties [24]. If the injector is 'short' (spanning, e.g., 1-2 betatron periods in length), then chaos will likely be unimportant. If it is 'long' (spanning several betatron periods), then chaotic mixing can be important. The same consideration applies to accelerators in general.

4.0 THE AES/JLAB HIGH-AVERAGE-CURRENT PHOTOINJECTOR

The AES/JLab photoinjector is designed ultimately to drive a 100+ kW FEL [25]. First it is to be built, tested, and optimized at Jefferson Lab (JLab). The design has evolved through three iterations, all of which incorporate a dc gun, a solenoid for emittance compensation, and superconducting radiofrequency cavities. The first iteration was for a 100 mA photoinjector comprised of three 750 MHz accelerating cavities; the model is pictured in Fig. 12, and a schematic with dimensions appears in Fig. 13. The second iteration was for an 1100 mA photoinjector comprised of seven 750 MHz accelerating cavities, and a 3rd-harmonic 2250 MHz cavity used to minimize the longitudinal emittance and shape the bunch for better compression. The third, and present, design is for a 1000 mA photoinjector comprised of three 750 MHz accelerating cavities and the 2250 MHz cavity. We simulated all three designs using a suite of codes, including our new wavelet-based code. The process of doing so also provided, as a by-product, additional code verification.

We simulated the first two iterations of the AES/JLab photoinjector design using the codes IMPACT-T (the conventional version) and PARMELA with 200,000 and 50,000 macroparticles, respectively. For the initial (100 mA) design, results for rms beam radius, transverse normalized rms emittance, energy, and longitudinal normalized rms emittance are provided in Fig. 14 and summarized in Table 1. Included in Table 1 are

results computed by Advanced Energy Systems, Inc. (AES) using PARMELA. For the upgraded (1100 mA) design the corresponding results from IMPACT-T and PARMELA simulations are shown in Fig. 15 and summarized in Table 2.

During Spring 2005 AES, in collaboration with JLab, proposed a modified version of the 100 mA photoinjector, one that would operate at higher bunch charge (1.0 nC instead of 0.133 nC) and thus higher average current (approaching 1000 mA). This is the present design and the one slated to be built and tested at JLab. It is based essentially on the same mechanical structure as the original design. The differences consist of an extra ~0.35 m drift between the dc gun and first cavity, and the addition of the 3rd-harmonic cavity. A schematic of this design appears in Fig. 16

Given this new lattice, we optimized the phases and amplitudes of the four cavities (three 750 MHz accelerating cavities and one 2250 MHz 3rd-harmonic cavity) and performed simulations with the conventional version of IMPACT-T, with PARMELA, and with our wavelet-based code (IMPACT-T with PCG). Simulations with the two versions of IMPACT-T involved 200,000 macroparticles, and with PARMELA, 20,000 macroparticles. We chose ten times fewer macroparticles for PARMELA simply to produce a run time comparable to those of the IMPACT-T codes. The input transverse distributions were uniform.

Results of the simulations appear in Fig. 17. The beam radius and bunch length are within specified values. Kinetic energy at the exit of the photoinjector is ~4.8 MeV,

which is lower than the desired value (6 MeV). Longitudinal emittance is only slightly less than 200 keV-ps and this is nearly double the specified value (100 keV-ps). Also, the extra drift and the extra space charge increase the transverse emittance to ~4 mm-mrad at 5 m downstream of the cathode. Using the 3rd-harmonic cavity enables decreasing the longitudinal emittance, but this is done at the expense of increasing the transverse emittance in keeping with Liouville's Theorem, and this inhibits reaching the target value for the latter (3 mm-mrad). Additional optimization is still possible, but based on all of the numerical experiments we did, we believe that the transverse and longitudinal emittances cannot be significantly improved.

The 3rd-harmonic cavity is used not only to minimize the longitudinal emittance, but also to linearize the energy distribution across the electron bunch. This is a necessary prerequisite to successful bunch compression upstream of the FEL system, i.e., prior to lasing. Figure 18 shows the energy distribution as a function of longitudinal position within the bunch before and after the 3rd-harmonic cavity. Blue and red points are those computed with the wavelet and conventional versions of IMPACT-T, respectively, and green points are those computed with PARMELA..

We also carefully checked the transverse particle distribution at the exit of this photoinjector to determine whether or not a halo formed. For this purpose we simulated both a uniform and nonuniform input particle distribution using the conventional version of IMPACT-T with 200,000 macroparticles. The transverse particle distributions at the

cathode and at the exit of the photoinjector appear in Fig. 19. Plots of the radial distribution of particles at the cathode and at the photoinjector exit corresponding to the transverse distributions of Fig. 19 are provided in Fig. 20. *For both input distributions, halo does form.* This is an important result; it constitutes a prediction that needs to be confirmed (or refuted) during photoinjector commissioning and optimization. We determined conclusively, by doing simulations with space charge turned 'on' and 'off', that the halo is a consequence of space charge.

To elaborate: based on our simulations, the outermost particles are ~1 cm away (transversely) from the center of the bunch (Fig. 20), regardless whether the initial distribution is uniform or not. Precisely how far from the center seems to depend mainly on the initial size of the beam. For the case shown in Fig. 20 the initial beam has ~4 mm diameter. As a check, we redid the same simulations using IMPACT-T with PCG and found the same quantitative result concerning the halo extent, and very nearly the same quantitative results concerning the halo-density profile. We believe that further optimization of the injector settings can be done to decrease this halo, and we are presently exploring this conjecture.

5.0 DISCUSSION AND CONCLUSIONS

We formulated, implemented, and thoroughly tested a 3D wavelet-based Poisson solver that uses the preconditioned conjugate-gradient method and incorporates general

inhomogeneous Dirichlet boundary conditions. We merged our Poisson solver into IMPACT-T to obtain a fully functional serial PIC code. We found that photoinjector simulations performed using IMPACT-T with its "native" Poisson solver (based on Green functions and fast Fourier transforms) and IMPACT-T with our PCG solver produce essentially equivalent outcomes (in terms of a standard set of rms diagnostics and phase-space detail). To our knowledge, the work reported here constitutes the first application of wavelet-based multiscale methodology to 3D computer simulations of beam dynamics.

We employ wavelet thresholding to remove effects of numerical noise from simulations. Although impossible to quantify precisely, such denoised simulations are *certainly at least as accurate* as simulations performed with solvers based on Green functions and FFTs without denoising, as measured by SNR. Most probably, however, because wavelet thresholding appreciably reduces the level of numerical noise in realistic simulations, it improves accuracy.

Regarding applications to new photoinjector designs, we performed independent simulations for all existing versions of the AES/JLab photoinjector using three distinctly different codes: PARMELA, IMPACT-T (conventional version) and IMPACT-T with PCG (our new wavelet version). Our results agree well with PARMELA simulations done by AES for the older versions of the photoinjector (100 mA and 1100 mA); we do not have recent simulation results from other institutions against which to compare

concerning the latest 1000 mA photoinjector design. Concerning the 1000 mA design, based on our results, we believe that the rms beam radius and rms bunch length can be achieved within specifications. The normalized transverse rms emittance is slightly too large (~4 mm-mrad vis-à-vis the 3 mm-mrad specification). The longitudinal rms emittance is also slightly too large (~190 keV-ps vis-à-vis the 150 keV-ps specification). The output kinetic energy of the electrons is too small (~4.8 MeV vis-à-vis the 6 MeV specification). However, we believe that the performance of this photoinjector can be somewhat improved by further optimization of the beam parameters, though not largely improved. What is perhaps our most important finding, there is beam halo at the exit of this photoinjector, and the halo extent appears to be insensitive to the profile of the input particle distribution. This is something that needs to be understood and minimized via additional simulations and optimization of the injector settings, and measured during commissioning of the real photoinjector. The good agreement between simulations of the AES/JLab photoinjector with IMPACT-T and PARMELA on the one hand and the new wavelet-based code (IMPACT-T with PCG) on the other hand raises our confidence regarding the quality of the new code. Similar agreement not only among the simulation codes, but also between the codes and laboratory measurements, was found in regard to a completely different machine (the Fermilab/NICADD photoinjector). Interestingly, we have in effect investigated two types of photoinjectors, one being based on a dc gun (AES/JLab photoinjector) and the other based on a rf gun (Fermilab/NICADD

photoinjector).

To summarize, the principal advantages of IMPACT-T with the wavelet-based Poisson solver are: (1) denoising, which likely yields simulation results comparable to those of a conventional code employing considerably more macroparticles; (2) compact storage and retrieval of the density and potential, a feature that will enable future comprehensive studies of halo formation and beam self-interaction via coherent synchrotron radiation (CSR); and (3) a slightly faster CPU time (by ~10%) compared to the conventional version of IMPACT-T.

The code comparisons presented herein are with respect to short photoinjectors. We anticipate that noticeable differences would arise between code predictions respecting long accelerators. Nonetheless, our code validations are encouraging, as are our verifications against experimental results.

In view of these successful comparisons, one may be led to wonder whether the choice of code makes any real difference, e.g., why would IMPACT-T with 200,000 macroparticles be any better than PARMELA with 20,000 macroparticles? Perhaps, if all one seeks to control are beam moments in photoinjectors, any of these codes may suffice. However, regarding high-brightness, high-average-current beams, details matter. One example is beam halo and the dynamics of its formation. To do a reliable halo study, especially in the context of a full accelerator driver for a free-electron laser, one needs a huge number of particles to get good halo statistics. None of these codes, in and of

themselves, suffice. However, being able to store and retrieve the bulk potential efficiently, an ability enabled only by the wavelet-based Poisson solver, does permit such halo studies. One can populate the previously computed bulk potential with an enormous number of test charges, charges that do not contribute to the potential but respond to it, and integrate their orbits efficiently. The reason this can be done is that, because the bulk potential is now known, there is no need for further Poisson solves; the problem reduces to particle tracking. The same argument applies to the study of CSR, where now one must integrate over the history of the density to compute the retarded 4-potential. Again, having the density stored compactly at each time step enables this process. These considerations perhaps make advantage (2) the compelling advantage of the wavelet-based Poisson solver.

Our current efforts are focused on several areas that encompass both algorithm optimization and applications. On the optimization side, the top priority is to find new ways to exploit sparsity of operators and data sets that will directly translate into increased computational efficiency. We are also exploring ways to compute more efficiently the potential on the boundary of the computational grid (as distinct from the physical boundaries of the system) so as to reduce the computational overhead at each step of the simulation. Finally, we have yet to parallelize our Poisson solver for use with the parallel version of IMPACT-T on multiprocessor machines, but we intend to do so in the near future.

ACKNOWLEDGEMENTS

We are grateful to Hans Bluem of Advanced Energy Systems, Inc. for getting us started on the simulations of the AES/JLab photoinjector and for sharing his simulation results. Carlos Hernandez-Garcia from Jefferson Laboratory kept us informed of changes in the layout of this photoinjector, and he shared his simulation results as well. Kevin Jensen of the Naval Research Laboratory kept us abreast of his theoretical developments respecting the physics of photoemission from cathodes, and we incorporated his results where warranted. This work was supported by Air Force contract FA9471-040C-0199 and Department of Energy grant DE-FG02-04ER41323 to NIU, and Department of Energy contract DE-AC02-05CH11231 to LBNL.

REFERENCES

[1] A.W. Chao, *Physics of Collective Beam Instabilities in High Energy Accelerators* (Wiley, New York, 1993), p. 26.

[2] C.L. Bohn, in Proc. Navy FEL Workshop, Newport News, VA, 5-6 June 2001.

[3] A. Todd, *et al*., in Proc. 2003 Part. Accel. Conf., Portland, OR, 12-16 May 2003, Paper MPPB075.

[4] J.P. Carneiro, *et al*., in Proc. 2000 Linear Accel. Conf., Monterey, CA, eConf C000821, pp. 104-106.

[5] R.A. Kishek, *et al*., Phys. Plasmas **10**, 2016 (2003).

[6] C.L. Bohn and I.V. Sideris, Phys. Rev. ST Accel. Beams **6**, 034203 (2003).

[7] C.L. Bohn and I.V. Sideris, Phys. Rev. Lett. **91:** 264801 (2003).

[8] B. Terzić, I. V. Pogorelov, and C. L. Bohn, "Wavelet-Based Poisson Solver for Use in Particle-in-Cell Beam Dynamics Simulations", Phys. Rev. ST Accel. Beams (submitted).

[9] C. L. Bohn, AIP Conf. Proc. **647**, 81 (2002).

[10] G. Beylkin, in *Wavelets: Mathematics and Applications*, eds. J. Benedetto and M. Frazier (CRC Press, Boca Raton, 1993).

[11] A. Averbuch, G. Beylkin, R. Coifman, P. Fischer and M. Israeli, "Adaptive Solution of Multidimensional PDEs via Tensor Product Wavelet Decomposition", Institut de Mathématiques de Bordeaux internal report (2003).

[12] A. Averbuch, G. Beylkin, R. Coifman, P. Fischer, M. Israeli, in *Signal and Image Representation in Combined Spaces*, eds. Y. Zeevi and R. Coifman (Academic Press, San Diego, 1998), pp. 341-359.

[13] I. Daubechies, *Ten Lectures on Wavelets* (SIAM, Philadelphia, 1992).

[14] J. Qiang, R. D. Ryne, S. Habib and V. Decyk, J. Comp. Phys. **163**, 434 (2000).

[15] J. Qiang and R. D. Ryne, Comp. Phys. Comm. **138**, 18(2001).

[16] G. H. Golub and C. F. Van Loan, *Matrix Computation* (Johns Hopkins Press, Baltimore, 1996).

[17] D. Donoho and I. Johnstone, Biometrika **81**, 425 (1994).

[18] R. Coifman and D. Donoho, in *Wavelet and Statistics*, eds. A. Antoniadis and G. Oppenheim (Springer, New York, 1995), pp. 125-150.

[19] http://www.nicadd.niu.edu/fnpl.

[20] D. Mihalcea, C. L. Bohn, U. Happek, and P. Piot, "Longitudinal Electron Bunch Diagnostics Using Coherent Transition Radiation", Phys. Rev. ST Accel. Beams **9**, 082801 (2006).

[21] A.J. Lichtenberg and M.A. Lieberman, *Regular and Chaotic Dynamics* (Springer-Verlag, New York, 1992).

[22] Y. Papaphilippou and J. Laskar, Astron. Astrophys. **329**, 451 (1998).

[23] I. V. Sideris, "A Measure of Orbital Stickiness and Chaos Strength", Phys. Rev. E **73**, 066217 (2006).

[24] C.L. Bohn, in *The Physics of High Brightness Beams*, eds. J. Rosenzweig and L. Serafini (World Scientific, Singapore, 2000), pp. 358-368.

[25] A. Todd, *et al*., in Proc. 2003 Part. Accel. Conf., IEEE Cat. No. **03CH37423**, 977 (2003).

Table 1. Results by AES & NIU for the original AES/JLab 100 mA photoinjector design.

Beam Property	Specifications	AES (PARMELA)	NIU (PARMELA)	NIU (IMPACT-T)
Charge (pC)	> 133	133.3	133.3	133.3
Beam radius (mm)	N/A	0.69	0.65	0.61
Emittance x,n (mm-mrad)	< 3.0	1.20	1.08	1.46
Energy spread (%)	< 3.0	0.5	0.5	0.5
Emittance z,n (keV-ps)	< 100	44	42	39
Bunch length (ps)	N/A	6.30	6.27	6.32
Energy (MeV)	> 7	7.66	7.66	7.66

Table 2. Results by AES & NIU for 1100 mA AES/JLab photoinjector.

Parameter	Specifications	AES (PARMELA)	NIU (PARMELA)	NIU (IMPACT-T)
Charge (pC)	>133	1100	1100	1100
Ave. Current (Amps)	>0.110	0.825	0.825	0.825
Beam Radius (mm rms)	N.A.	2.10	1.95	2.00
Beam Length (ps rms)	N.A.	9.60	9.80	9.80
Trans. Normalized Emittance (mm-mrad rms)	<3.0	5.10	7.13	6.74
Long. Normalized Emittance (keV-ps rms)	<100	42.50	38.05	43.56

Energy Spread (%)	<1.0	1.3	1.3	1.4
Energy (MeV)	>7.0	8.3	8.3	8.3

WAVELET

One-Dimensional

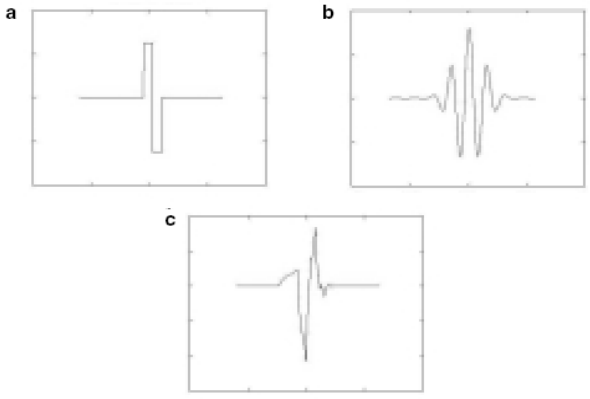

Haar wavelet, (b) Morlet wavelet,
(c) Daubechies wavelet.

Two-Dimensional

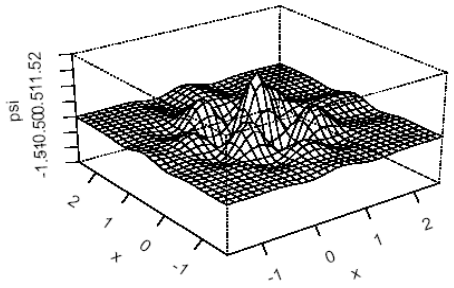

Daubechies least symmetric
N = 10 two-dimensional

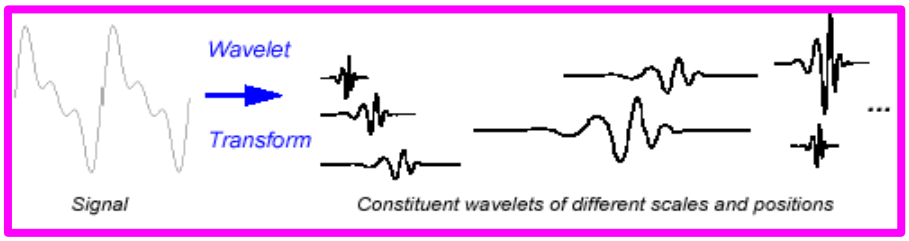

FIGURE 1

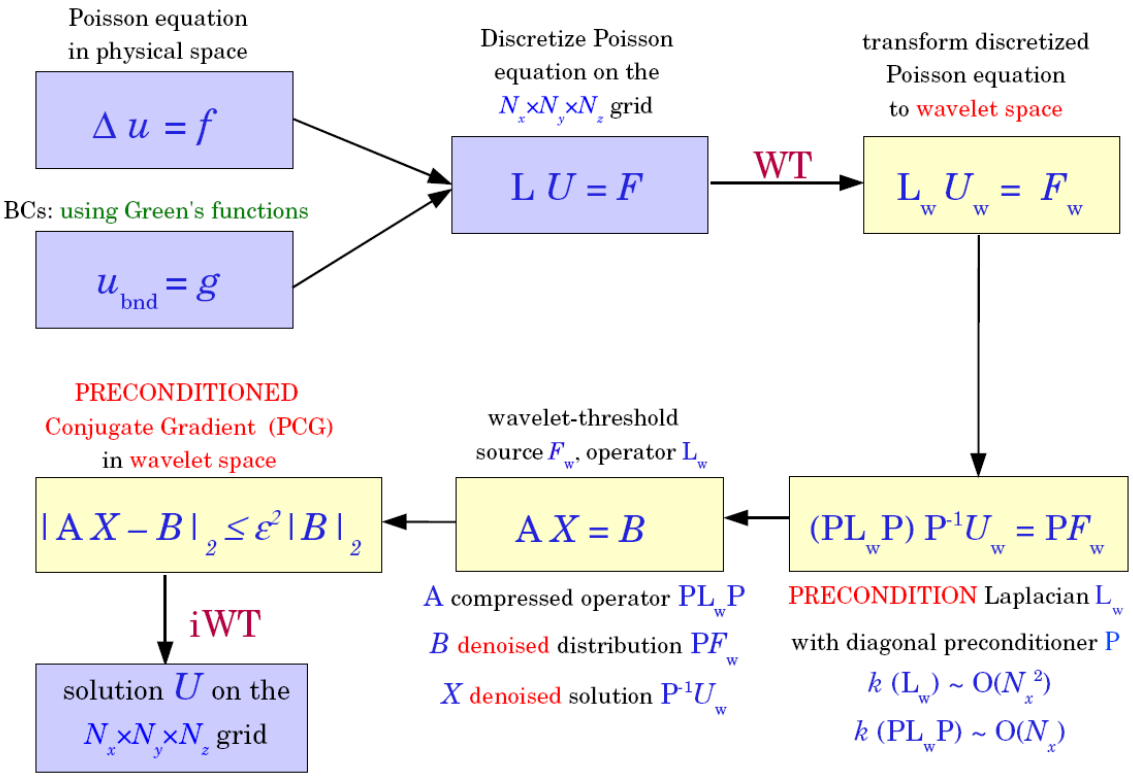
Poisson equation
in physical space
$\Delta u = f$
BCs: using Green's functions
$u_{bnd} = g$
Discretize Poisson
equation on the
$N_x \times N_y \times N_z$ grid
$L\,U = F$
WT
transform discretized
Poisson equation
to wavelet space
$L_w\,U_w = F_w$
$(PL_wP)\,P^{-1}U_w = PF_w$
PRECONDITION Laplacian L_w
with diagonal preconditioner P
$k\,(L_w) \sim O(N_x^2)$
$k\,(PL_wP) \sim O(N_x)$
wavelet-threshold
source F_w, operator L_w
$A\,X = B$
A compressed operator PL_wP
B denoised distribution PF_w
X denoised solution $P^{-1}U_w$
PRECONDITIONED
Conjugate Gradient (PCG)
in wavelet space
$|A\,X - B|_2 \le \varepsilon^2 |B|_2$
iWT
solution U on the
$N_x \times N_y \times N_z$ grid

FIGURE 2

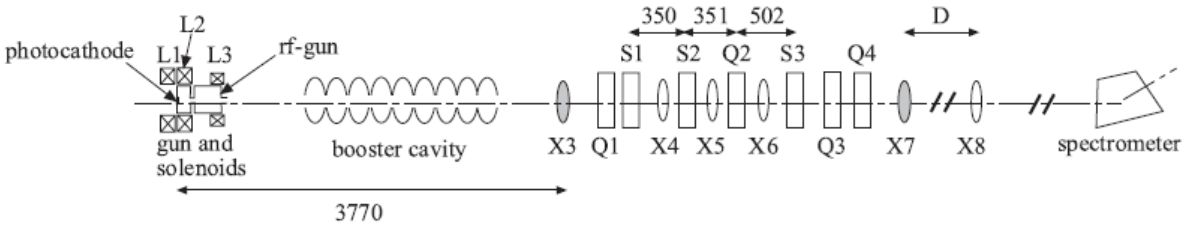
L2
photocathode
L1
L3
rf-gun
gun and
solenoids
booster cavity
3770
350
351
502
D
S1
S2
Q2
S3
Q4
X3
Q1
X4
X5
X6
Q3
X7
X8
spectrometer

FIGURE 3

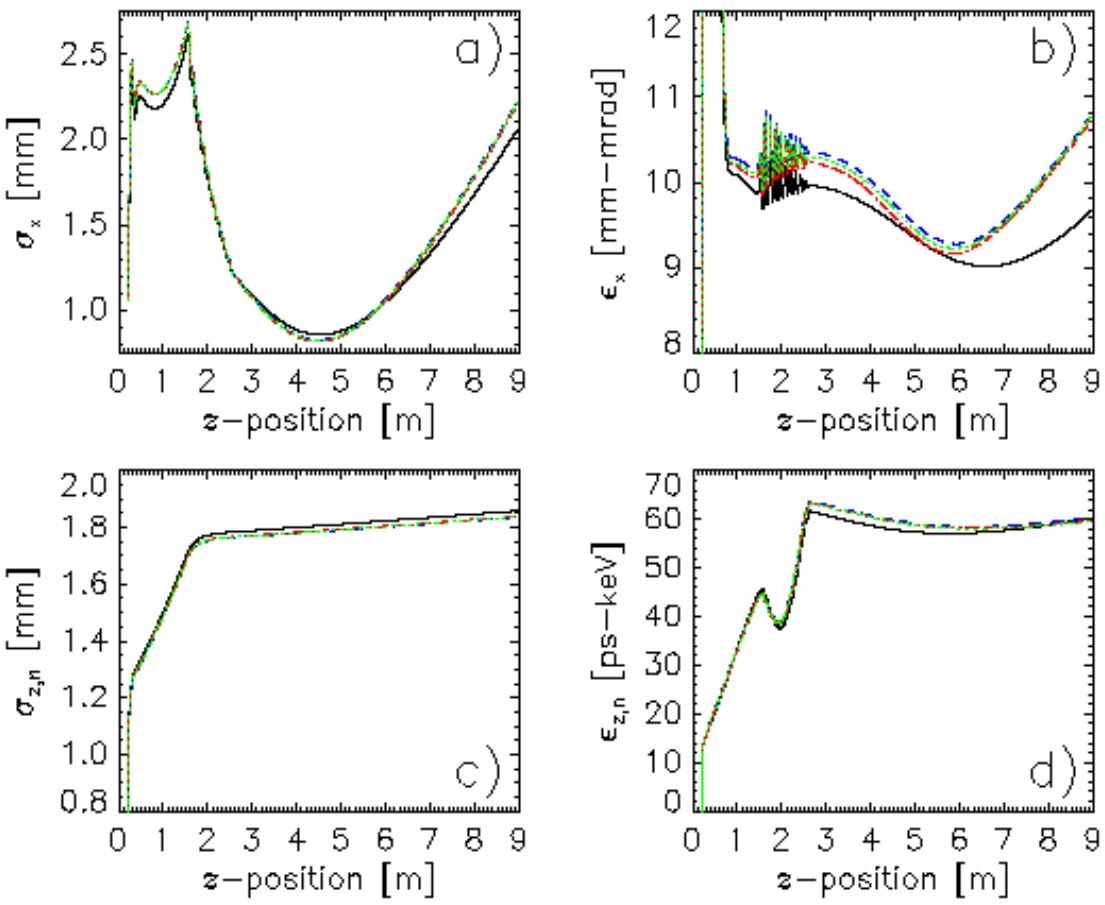
a)
b)
c)
d)
σ_x [mm]
ϵ_x [mm−mrad]
$\sigma_{z,n}$ [mm]
$\epsilon_{z,n}$ [ps−keV]
z−position [m]

FIGURE 4

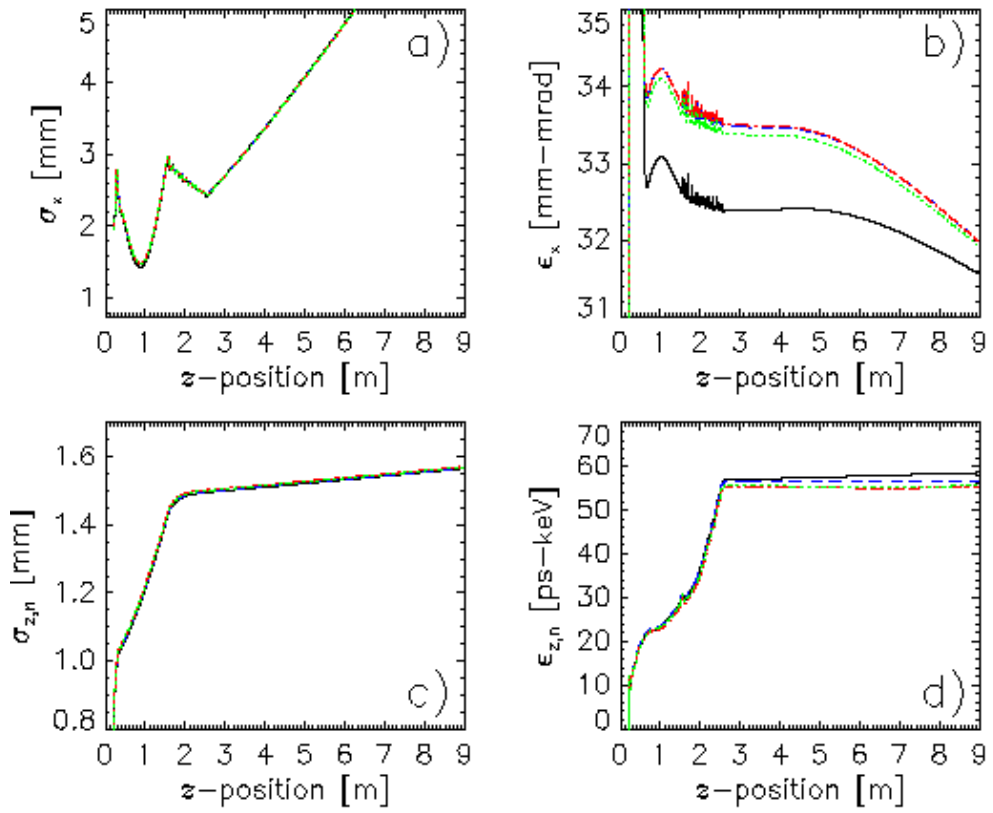
a)
b)
c)
d)
σx [mm]
εx [mm−mrad]
σz,n [mm]
εz,n [ps−keV]
z−position [m]

FIGURE 5

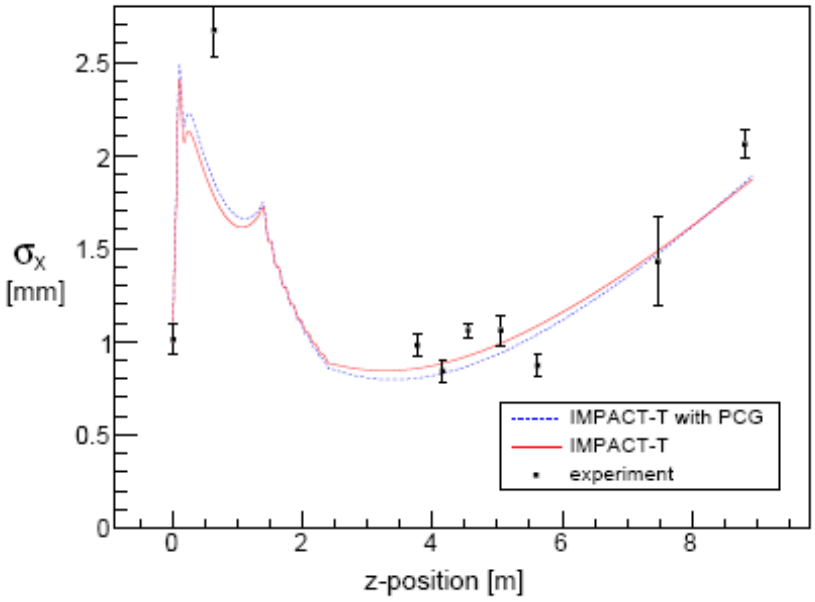
2.5
2
1.5
1
0.5
0
σ_x
[mm]
0
2
4
6
8
z-position [m]
IMPACT-T with PCG
IMPACT-T
experiment

FIGURE 6

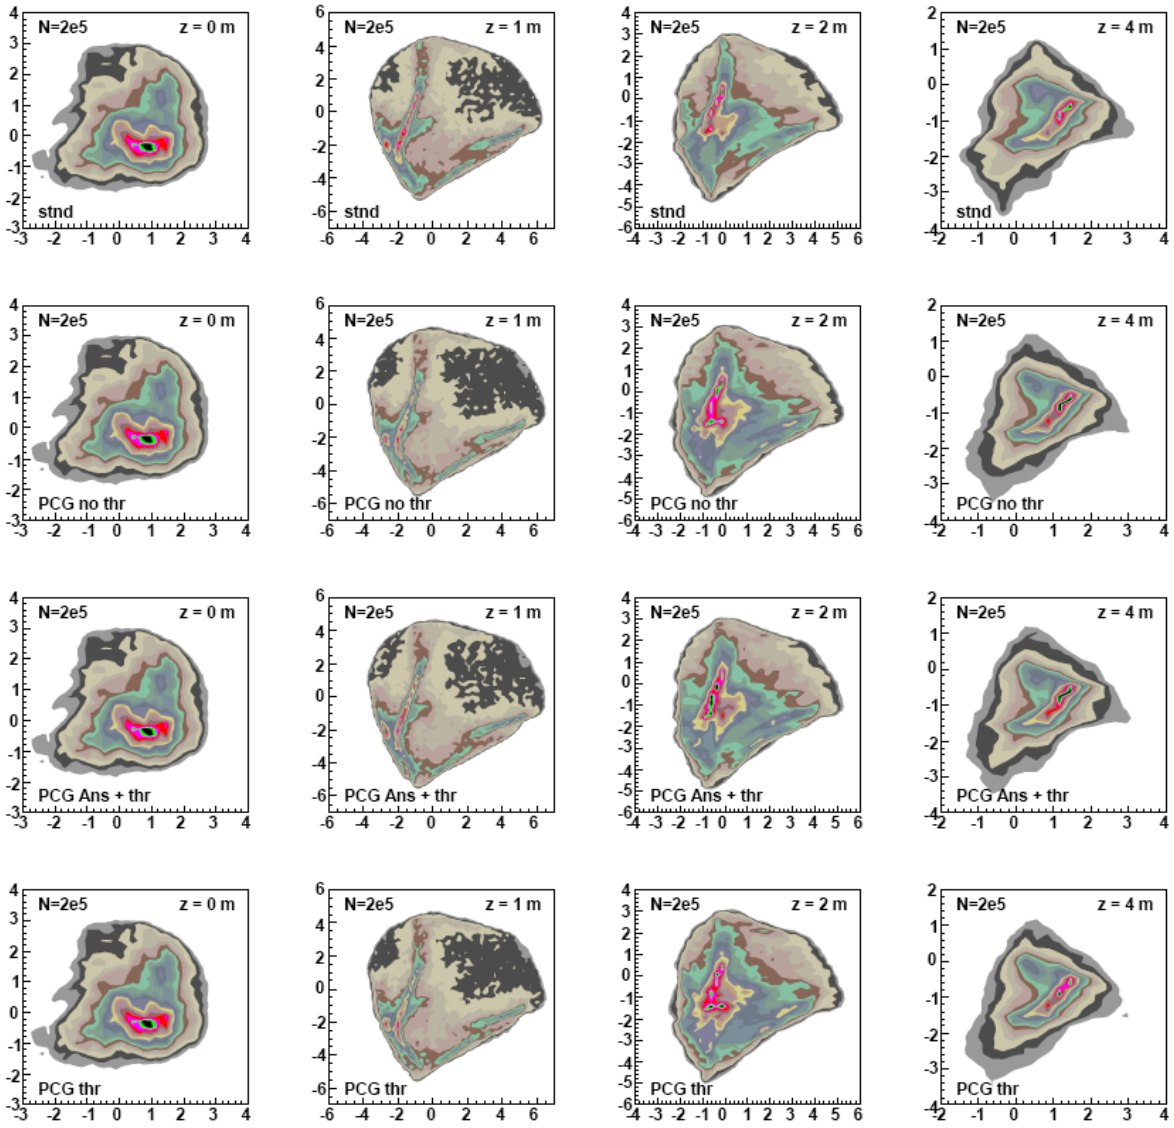
N=2e5
z = 0 m
stnd
N=2e5
z = 1 m
stnd
N=2e5
z = 2 m
stnd
N=2e5
z = 4 m
stnd
N=2e5
z = 0 m
PCG no thr
N=2e5
z = 1 m
PCG no thr
N=2e5
z = 2 m
PCG no thr
N=2e5
z = 4 m
PCG no thr
N=2e5
z = 0 m
PCG Ans + thr
N=2e5
z = 1 m
PCG Ans + thr
N=2e5
z = 2 m
PCG Ans + thr
N=2e5
z = 4 m
PCG Ans + thr
N=2e5
z = 0 m
PCG thr
N=2e5
z = 1 m
PCG thr
N=2e5
z = 2 m
PCG thr
N=2e5
z = 4 m
PCG thr

FIGURE 7

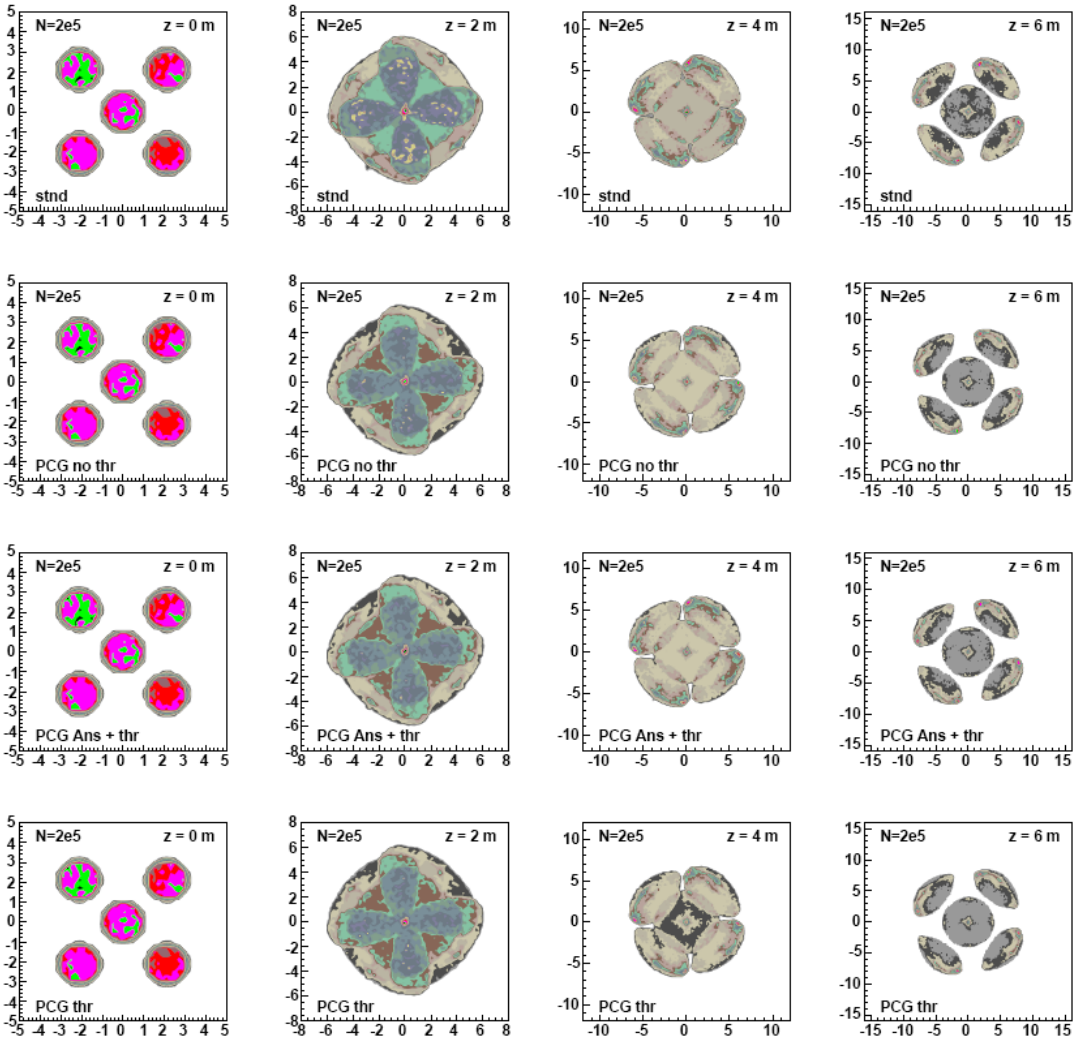
N=2e5
z = 0 m
stnd
N=2e5
z = 2 m
stnd
N=2e5
z = 4 m
stnd
N=2e5
z = 6 m
stnd
N=2e5
z = 0 m
PCG no thr
N=2e5
z = 2 m
PCG no thr
N=2e5
z = 4 m
PCG no thr
N=2e5
z = 6 m
PCG no thr
N=2e5
z = 0 m
PCG Ans + thr
N=2e5
z = 2 m
PCG Ans + thr
N=2e5
z = 4 m
PCG Ans + thr
N=2e5
z = 6 m
PCG Ans + thr
N=2e5
z = 0 m
PCG thr
N=2e5
z = 2 m
PCG thr
N=2e5
z = 4 m
PCG thr
N=2e5
z = 6 m
PCG thr

FIGURE 8

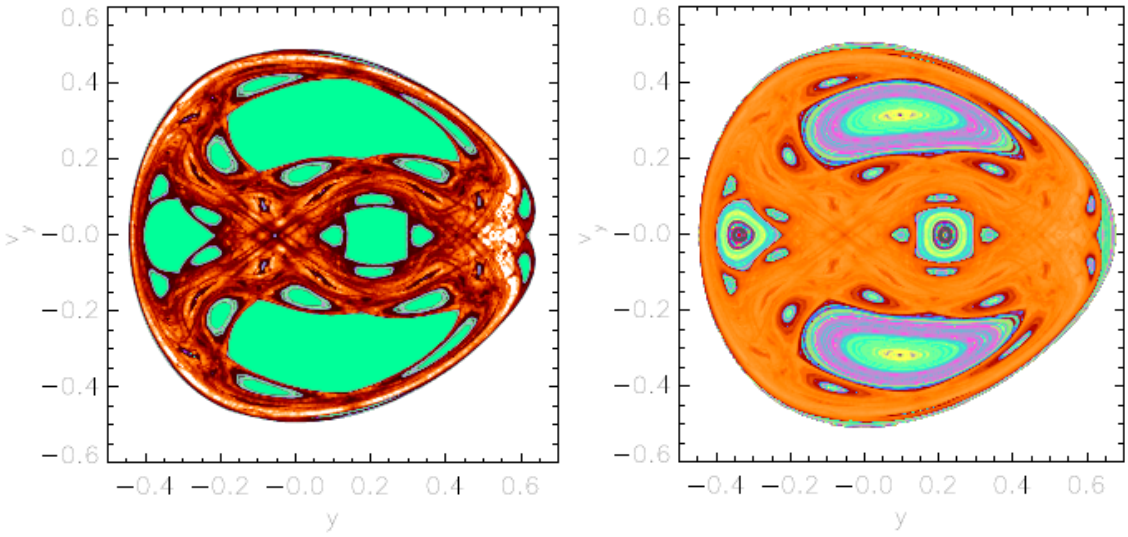
0.6
0.4
0.2
−0.0
−0.2
−0.4
−0.6
v y
−0.4 −0.2 −0.0 0.2 0.4 0.6
y
0.6
0.4
0.2
−0.0
−0.2
−0.4
−0.6
v y
−0.4 −0.2 −0.0 0.2 0.4 0.6
y

FIGURE 9

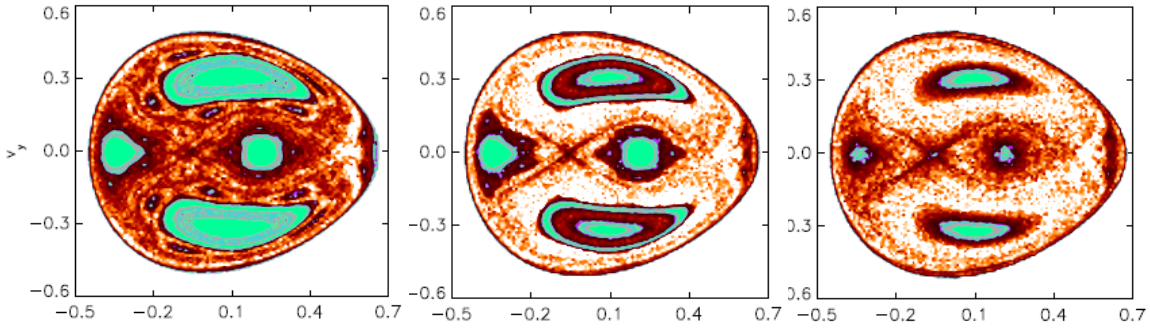
0.6
0.3
0.0
−0.3
−0.6
v_y
−0.5
−0.2
0.1
0.4
0.7

FIGURE 10

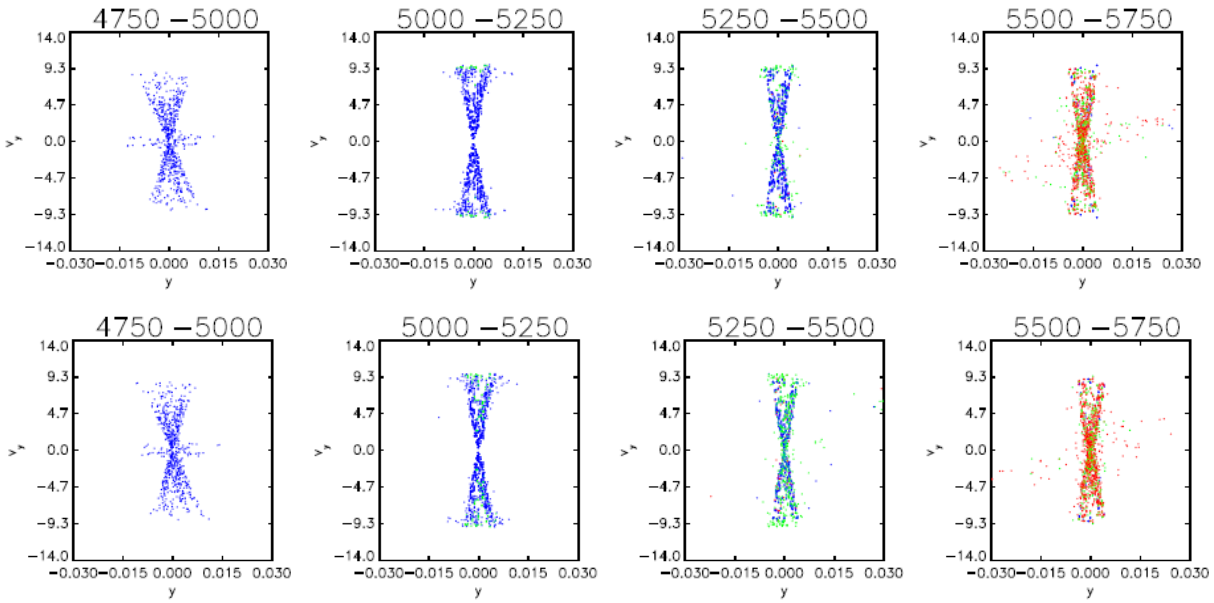
4750 −5000
5000 −5250
5250 −5500
5500 −5750
4750 −5000
5000 −5250
5250 −5500
5500 −5750
14.0
9.3
4.7
0.0
−4.7
−9.3
−14.0
−0.030 −0.015 0.000 0.015 0.030
y
v_y

FIGURE 11

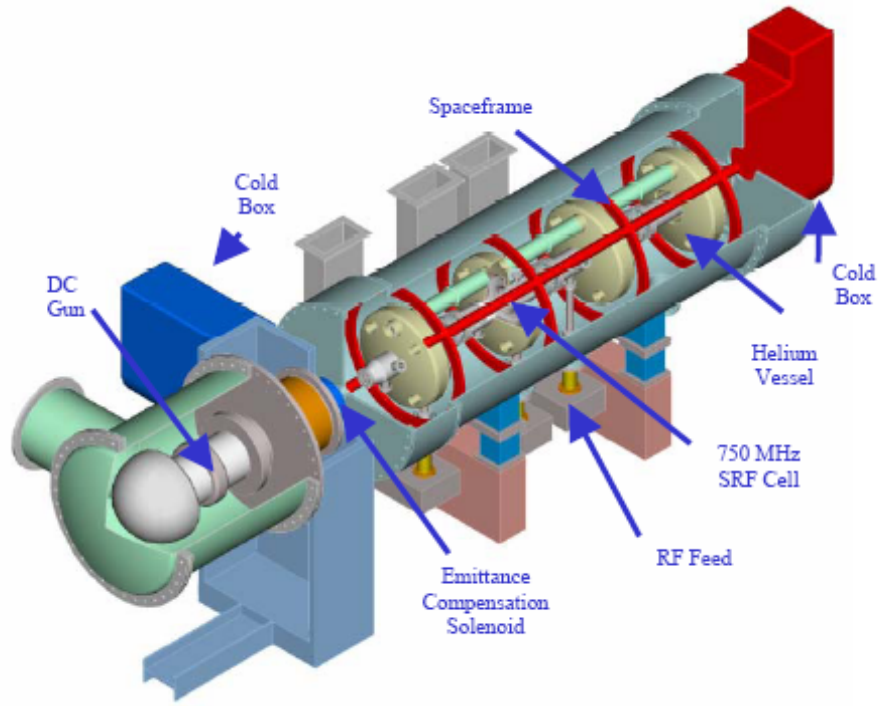
Spaceframe
Cold
Box
DC
Gun
Cold
Box
Helium
Vessel
750 MHz
SRF Cell
RF Feed
Emittance
Compensation
Solenoid

FIGURE 12

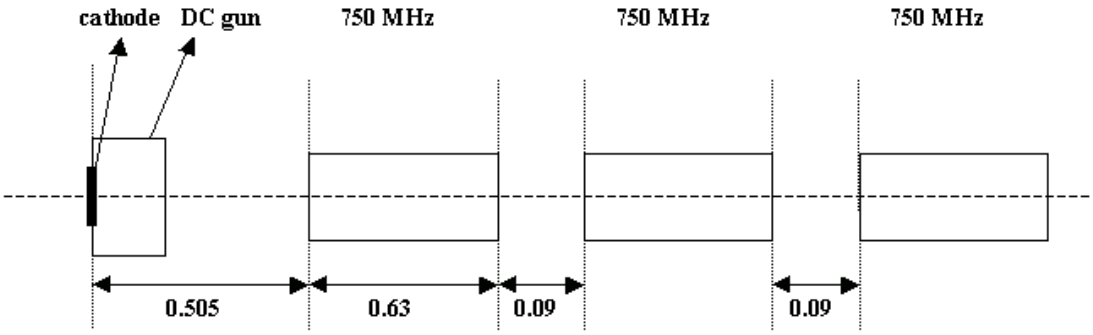
cathode
DC gun
750 MHz
750 MHz
750 MHz
0.505
0.63
0.09
0.09

FIGURE 13

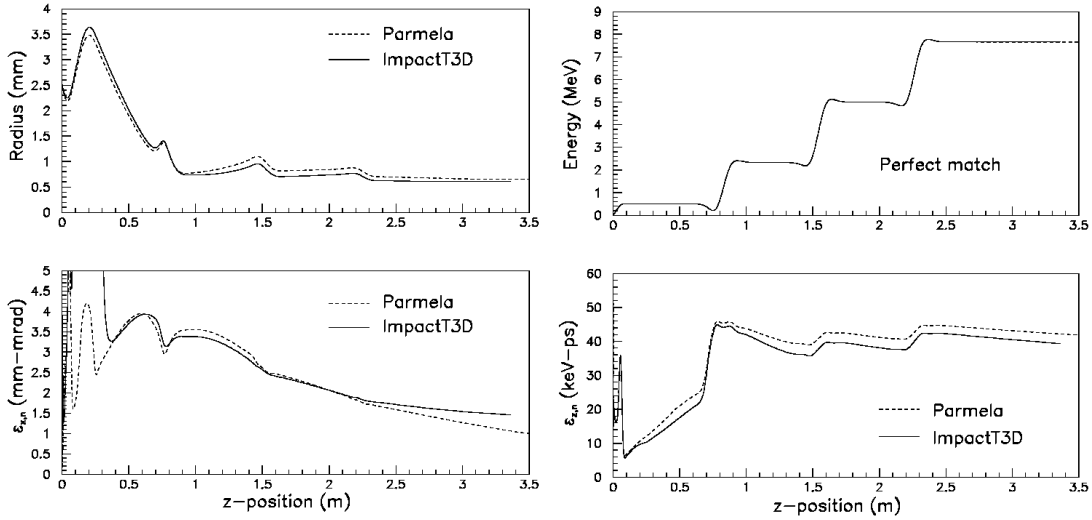
Radius (mm)
Parmela
ImpactT3D
Energy (MeV)
Perfect match
$\varepsilon_{x,n}$ (mm-mrad)
Parmela
ImpactT3D
z-position (m)
$\varepsilon_{z,n}$ (keV-ps)
Parmela
ImpactT3D
z-position (m)

FIGURE 14

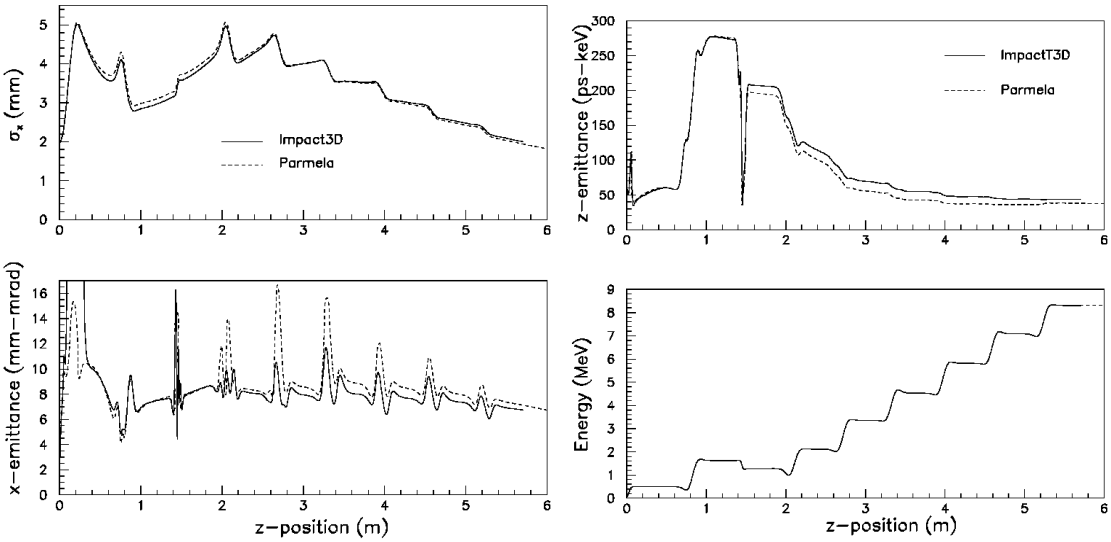
σx (mm)
Impact3D
Parmela
z−emittance (ps−keV)
ImpactT3D
Parmela
x−emittance (mm−mrad)
z−position (m)
Energy (MeV)
z−position (m)

FIGURE 15

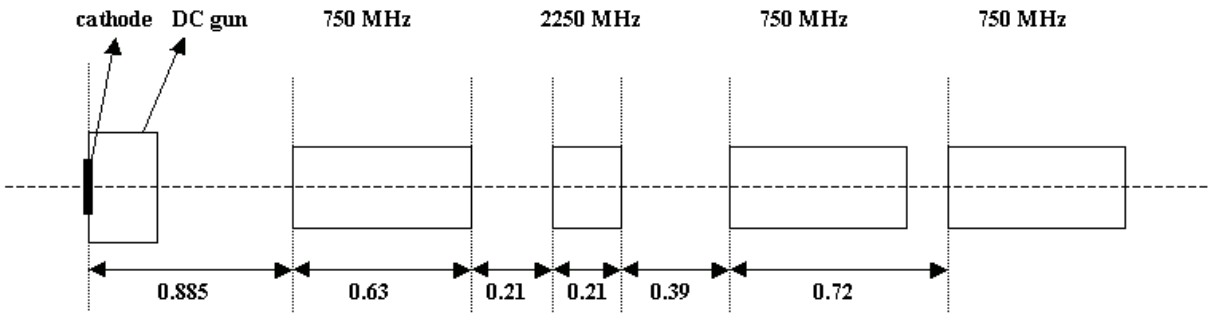
cathode
DC gun
750 MHz
2250 MHz
750 MHz
750 MHz
0.885
0.63
0.21
0.21
0.39
0.72

FIGURE 16

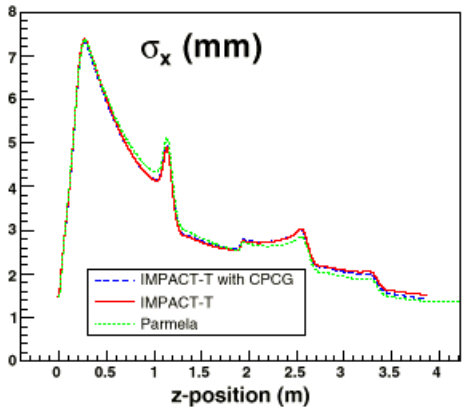
σx (mm)
IMPACT-T with CPCG
IMPACT-T
Parmela
z-position (m)

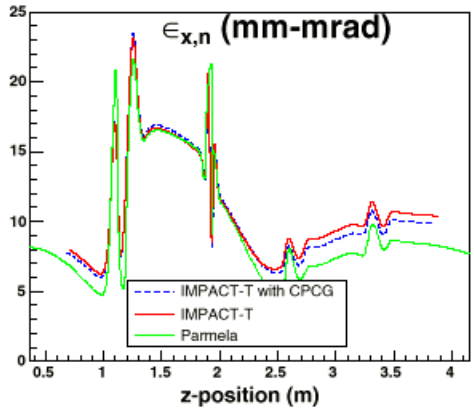
∈x,n (mm-mrad)
IMPACT-T with CPCG
IMPACT-T
Parmela
z-position (m)

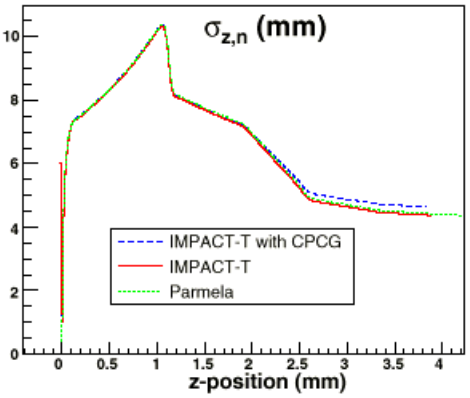
σz,n (mm)
IMPACT-T with CPCG
IMPACT-T
Parmela
z-position (mm)

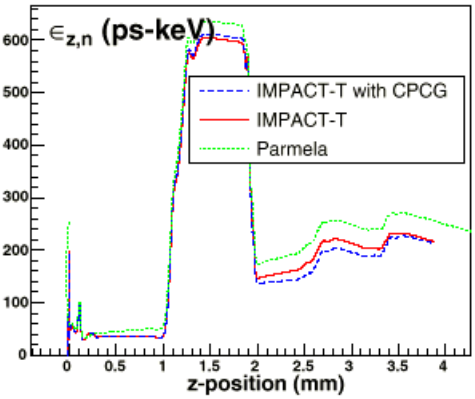
∈z,n (ps-keV)
IMPACT-T with CPCG
IMPACT-T
Parmela
z-position (mm)

FIGURE 17

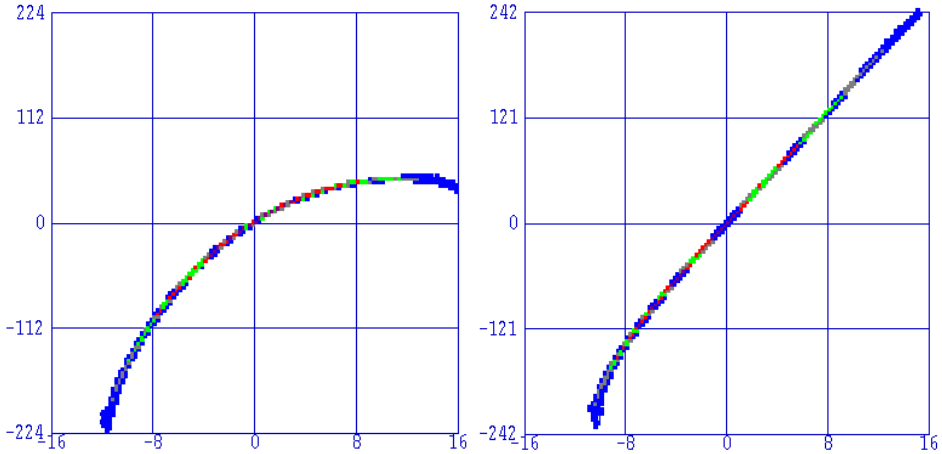
224
112
0
-112
-224
-16
-8
0
8
16
242
121
0
-121
-242
-16
-8
0
8
16

FIGURE 18

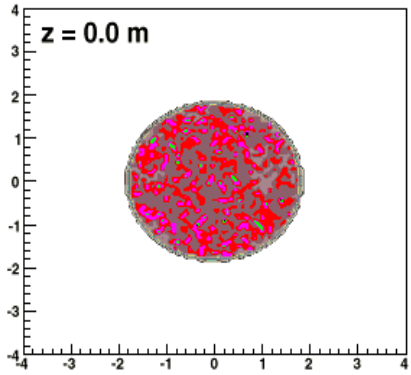
z = 0.0 m

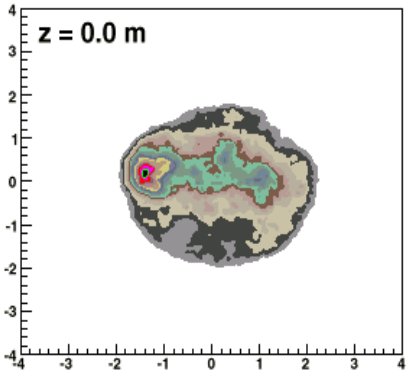
z = 0.0 m

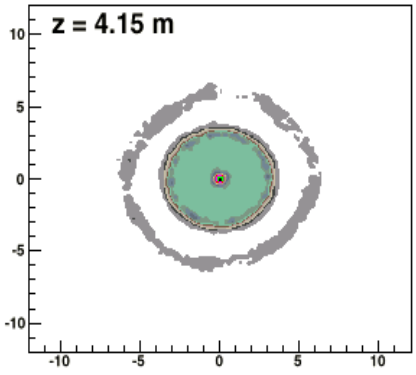
z = 4.15 m

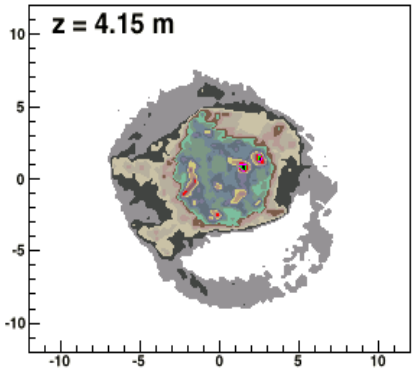
z = 4.15 m

FIGURE 19

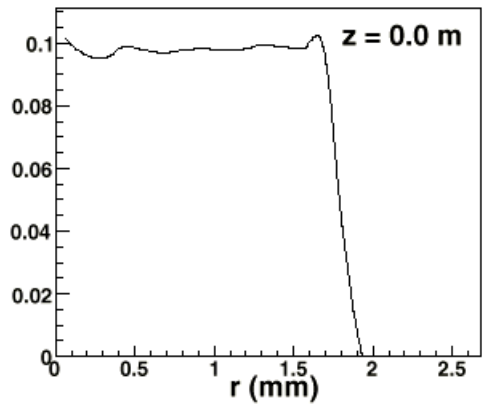
z = 0.0 m
r (mm)

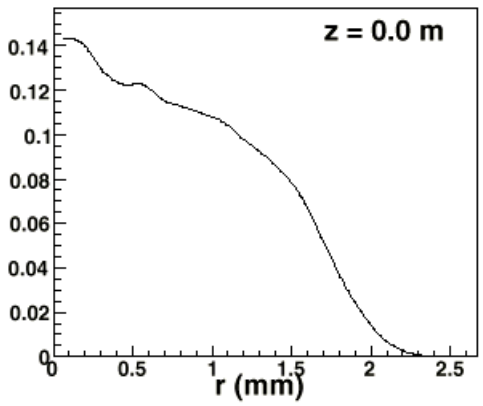
z = 0.0 m
r (mm)

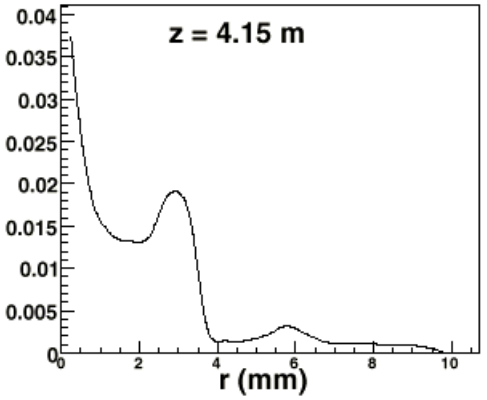
z = 4.15 m
r (mm)

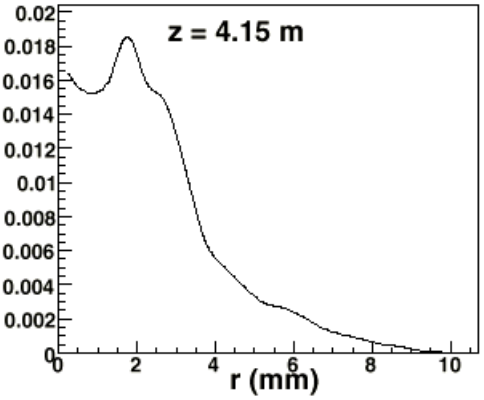
z = 4.15 m
r (mm)

FIGURE 20

FIGURE CAPTIONS

Figure 1. Examples of one-dimensional wavelets (top left) and a two-dimensional wavelet (top right), and decomposition of a signal into constituent wavelets of different scales using wavelets (bottom).

Figure 2. Flow chart of the wavelet-based Poisson solver using the preconditioned conjugate gradient (PCG) method. White boxes pertain to physical (configuration) space, and blue boxes pertain to wavelet space. Here, Δ is the Laplacian, u represents the potential, f is proportional to charge density, and L is the matrix representation of the Laplacian.

Figure 3. Schematic of the Fermilab/NICADD photoinjector. ''X'' refers to diagnostics stations (beam viewers, and/or slit location), ''L'' to emittance-compensating solenoidal lenses, ''Q'' to quadrupole magnets, and ''S'' to skew quadrupoles (skew-oriented when flat beam is wanted; not skew for our simulations). A bunch compressor (not shown) is located between the booster cavity and X3. All distances are in mm, with D =800 mm.

Figure 4. For *Distribution 1*: Simulation results for the Fermilab/NICADD photoinjector using 32^3 grid points and 200,000 macroparticles. Curves correspond to: conventional version of IMPACT-T (black), IMPACT-T with PCG without denoising (green), IMPACT-T with PCG with thresholding and Anscombe transformation (blue), and IMPACT-T with PCG with thresholding but without Anscombe transformation (red). Computed quantities are: (a) rms beam radius, (b) rms normalized transverse emittance, (c) rms bunch length, (d) rms normalized longitudinal emittance. For IMPACT-T with PCG, we use Daubechies wavelets of order 2.

Figure 5. Same as in Fig.5, but for *Distribution 2*.

Figure 6. For *Distribution 1*: measured rms beam radii for the Fermilab/NICADD photoinjector (black dots) versus those computed with conventional IMPACT-T (red) and IMPACT-T with PCG (blue). Numerical simulations were done on a 32^3 grid with 200,000 macroparticles using Daubechies wavelets of order 2 and no thresholding.

Figure 7. For *Distribution 1*: Integrated transverse cross-section of the beam at different positions along the Fermilab/NICADD photoinjector computed with 32^3 grid points, 200,000 macroparticles, and Daubechies wavelets of order 2: conventional IMPACT-T (first row), IMPACT-T with PCG and no thresholding (second row), IMPACT-T with

PCG and Anscombe transform and thresholding (third row), IMPACT-T with PCG and thresholding but no Anscombe transformation (fourth row). First column shows the beam leaving the cathode, second at z = 1m, third at z = 2m, and fourth at z = 4m.

Figure 8. For *Distribution 2*: Same as in Fig. 7, except z = 2m (second column), z = 4m (third column), and z = 6m (fourth column).

Figure 9. Contour plots of the Henon-Heiles (y,v_y) phase space (E = 0.125). The left panel emphasizes the chaotic sea. Here, the regular regions are easily distinguishable as green. The right panel emphasizes regular regions. The degree of regularity is different in different parts of these regular areas. These plots were made by integrating 20,000 orbits for about 300 orbital periods.

Figure 10. Contour plots of (y,v_y) phase space of the Henon-Heiles potential (E = 0.125) for short evolution times: ~30 periods (left), ~15 periods (center), ~10 periods (right).

Figure 11. Phase-space projections of the evolution of the same initial electron distribution computed with the conventional version (top) and wavelet version (bottom) of IMPACT-T. Blue denotes particles on regular orbits, green denotes weakly chaotic orbits, and red denotes strongly chaotic orbits.

Figure 12. Initial design of AES/JLab 100 mA photoinjector.

Figure 13. Schematic of AES/JLab photoinjector design of Fig. 12; dimensions in meters.

Figure 14. Beam moments for original AES/JLab 100 mA photoinjector design.

Figure 15. Beam parameters for 1100 mA version of the AES/JLab photoinjector.

Figure 16. Schematic of 1000 mA AES/JLab photoinjector; dimensions in meters.

Figure 17. Simulations of modified AES/JLab photoinjector. Bunch charge is 1nC. A 3rd harmonic cavity is used to minimize longitudinal emittance, but it does so at the expense of the transverse emittance.

Figure 18. Energy vs. longitudinal position along the bunch before (left) and after (right) the 3rd-harmonic cavity. The vertical axis represents the difference between the energy of a given electron and the energy of the reference particle in keV. The horizontal axis represents the difference between the longitudinal position of a given electron and the z-

position of the reference particle in degrees. Blue: Impact-T with CPCG, Red: Impact-T, and Green: PARMELA.

Figure 19. Transverse distributions at cathode (upper plots) and, correspondingly, at the exit of the modified AES/JLab photoinjector (bottom plots); dimensions in millimeters.

Figure 20. Radial charge-density distributions, in arbitrary units, corresponding to the transverse distributions shown in Fig. 19.

BIOGRAPHICAL SKETCHES

Courtlandt L. Bohn holds a Ph.D. in astrophysics from the University of Chicago and is Presidential Research Professor at Northern Illinois University. He spent 14 years at national laboratories and 10 years as an Air Force officer. Bohn led development of the energy-recovering accelerator for the IR Demo FEL at Jefferson Lab, as well as commissioning that FEL. As an officer, he served at the Pentagon managing most of the AF's directed-energy activities, and taught at the AF Academy where he won the AF R&D Award and the Seiler Research Award. He has published over 100 papers on accelerator physics.

Daniel Mihalcea received his B.Sc. degree in physics from Bucharest University, Romania, in 1981 and Ph.D. degree from Kansas State University in 1998. He is a Research Associate at Northern Illinois University. His research areas include beam physics, photoinjector design, computer simulations for beam and accelerator physics, and electron-beam diagnostics. Currently he is actively involved in experimental beam-physics research at the Fermilab/NICADD photoinjector and in theoretical projects aimed at understanding beam-halo formation and developing more sophisticated simulation tools for accelerator design.

Ilya Pogorelov received his B.Sc. degree in applied mathematics and physics from Moscow Institute of Physics and Technology in 1993, and his Ph.D. in physics from the University of Florida in 2001. He is a Computer Scientist at Lawrence Berkeley National Laboratory. His research interests include computational physics, beam dynamics, applied mathematics, and nonlinear dynamics.

Ioannis V. Sideris received his B.Sc. from the Aristotelion University of Thessaloniki, Greece, and his Ph.D. from the University of Florida in 2002. He is a Research Associate at Northern Illinois University. His main interests are chaos theory, the *N*-body problem, and evolution of stochastic systems. Sideris has worked on a variety of projects on galactic and beam dynamics, focusing on the investigation of chaotic phenomena and the effects of noise on macro- and micro-scales. Recently, he developed a new measure of chaos that can characterize signals and orbits locally in time.

Balša Terzić received his Ph.D. in applied mathematics from Florida State University in 2002. He is a Research Scientist at Northern Illinois University. He was previously a Research Associate in the Department of Astronomy at the University of Florida where he worked on problems relating to chaos and nonlinearity in multiparticle systems. His

research interests include development of multiscale algorithms for beam dynamics simulations, modeling of coherent synchrotron radiation, and a variety of astrophysical problems relating to morphology of elliptical galaxies and manifestation of chaos in self-gravitating systems.