Advanced accelerator simulation research: miniaturizing accelerators from kilometers to meters

W.B. Mori, F.S. Tsung, C. Huang, M. Zhou, W. Lu, M. Tzoufras, V.K. Decyk

Departments of Physics and Astronomy and of Electrical Engineering,
University of California, Los Angeles, Los Angeles, CA 90095

D.L. Bruhwiler, J.R. Cary,* P. Messmer, D.A. Dimitrov, C. Nieter

Tech-X Corp, Boulder, CO 80303
*and U. of Colorado, Boulder, CO 80309

T. Katsouleas, S. Deng, A. Ghalam

Department of Electrical Engineering ,
University of Southern California, Los Angeles, CA 90089

E.H. Esarey, C.G.R. Geddes

Lawrence Berkeley National Laboratory, Berkeley, CA 94720

J.H. Cooley, T.M. Antonsen Jr.

Departments of Physics and of Electrical Engineering ,
University of Maryland, College Park , MD 20742

Mori@physics.ucla.edu

Abstract: Advanced accelerator research is aimed at finding new technologies that can dramatically reduce the size and cost of future high-energy accelerators. Supercomputing is already playing a dramatic and critical role in this quest. One of the goals of the SciDAC Accelerator Modeling Project is to develop code and software that can ultimately be used to discover the underlying science of new accelerator technology and then be used to design future high-energy accelerators with a minimum amount of capital expenditure on large-scale experiments. We describe the existing hierarchy of software tools for modelling advanced accelerators, how these models have been validated against experiment, how the models are benchmarked against each other, and how these tools are being successfully used to elucidate the underlying science.

1. Introduction

The long-term future of experimental high-energy physics research using accelerators depends on the successful development of novel ultra high-gradient acceleration methods. New acceleration techniques using lasers and plasmas have already been shown to exhibit gradients and focusing forces more than 1000 times greater than conventional technology. The challenge is to control these high-gradient systems and then to string them together. Such technologies would enable the development of

ultra-compact accelerators. The potential impact on science, industry, and medicine of placing such compact accelerators in research organizations, high-tech businesses, and hospitals is staggering.

Under the Accelerator Modeling SciDAC Project, the Advanced Accelerator effort has emphasized developing a suite of parallelized particle-in-cell (PIC) codes, ensuring that all code be reusable and easily extendable, benchmarking these codes and their underlying algorithms against each other and against experiments, adding more realism into the models, and applying them to advanced accelerators as well as more mainstream problems in accelerator physics, such as the electron cloud instability. Furthermore, the effort has included running these codes to plan and interpret experiments and to study the key physics that needs to be understood before a 100+ GeV collider based on plasma techniques can be designed and tested. The Advanced Accelerator effort at the Universities has supported PhD students at both UCLA and USC.

The application codes and Frameworks used or developed in this effort are OSIRIS [1], VORPAL [2], QuickPIC [3], UPIC [4], and Chombo/EB. OSIRIS is a fully explicit three-dimensional PIC code written in Fortran95; VORPAL is also a fully explicit PIC code, but is written in C++; QuickPIC is a quasi-static (the idea will be described shortly) but fully 3D PIC code based on the UPIC Framework; UPIC is a highly optimized Framework for quickly constructing new parallelized PIC codes; and Chombo/EB is an embedded boundary capability for modeling complex geometries. UPIC has allowed us to quickly write parallel electrostatic and/or fully explicit electromagnetic PIC codes as well as reduced description codes such as QuickPIC. Furthermore, it will be used as a test bed for optimization techniques which can eventually incorporated into OSIRIS and VORPAL. Chombo/EB – developed by the APDEC ISIC – provides the geometrical capability needed to model the complex geometries encountered in the gas jet problem.

The following describes our major accomplishments. They are organized into three areas: code and algorithm development, benchmarking and validation, and applications of the production codes to accelerator projects.

2. Codes and algorithms

In some advanced accelerator concepts a drive beam, either an intense particle beam or laser pulse, is sent through a uniform plasma. The space charge or radiation pressure creates a space-charge wake on which a trailing beam of particles can surf. To model such devices accurately usually requires following the trajectories of individual plasma particles. Therefore, the software tools developed fully or partially under this project, OSIRIS [1], VORPAL [2], OOPIC [5], QuickPIC [3], and UPIC [4] rely on the particle-in-cell (PIC) techniques.

2.1. Algorithms

2.11 *Full PIC*

The fully explicit PIC algorithm is straightforward. Basically, a chosen number of particles are loaded onto a grid. The charge and current densities can then be calculated by "depositing" the particles onto the grid; $\rho = \sum_{particles} q$ and $\vec{j} = \sum_{particles} q\vec{v}$. These current and charge densities are used to advance the fields via Maxwell's equations,

$$-\nabla \times \vec{E} = \frac{1}{c}\frac{\partial \vec{B}}{\partial t} \quad \text{and} \quad \nabla \times \vec{B} = \frac{4\pi}{c}\mathbf{j} + \frac{1}{c}\frac{\partial \vec{E}}{\partial t}$$

The updated fields are used to advance the particles to new positions and velocities via the relativistic equation of motion,

$$\frac{d\vec{P}}{dt} = q\left(\vec{E} + \frac{\vec{v} \times \vec{B}}{c}\right)$$

Although simple in concept, there are many subtle issues that arise when solving these equations on a computer, e.g., the deposited ρ and j may not satisfy the continuity equation and the use of smoothers and splines to reduce the noise from aliasing. In addition, there are issues related to how to efficiently parallelize this basic algorithm. All codes described here use domain decomposition and MPI.

2.1.2. Ponderomotive guiding center

When using a fully explicit PIC code the time step must be smaller than the smallest cell size. Therefore, when modeling a laser-plasma accelerator scheme the smallest cell size must be $\sim.3c/\omega_o$ in order to resolve the laser wavelength ($\sim$20 cells per λ). In the ponderomotive guiding center scheme one separates out the fields into plasma fields and laser fields. These equations were clearly derived by Mora and Antonsen [6].

The plasma fields are solved explicitly and the time averaged laser field is solved using an envelope, i.e, a paraxial wave type, equation,

$$\left(2i\omega\partial_\tau - 2\partial_{\tau\xi} - \nabla_\perp^2\right)a = \chi\, a$$

The plasma particles are pushed using the plasma fields and the ponderomotive force from the laser's envelope as follows

$$\frac{d\vec{P}}{dt} = q\left(\vec{E} + \frac{\vec{v} \times \vec{B}}{c}\right) - \frac{1}{4}\frac{q}{\gamma}\nabla|a|^2 \text{ where } \gamma = \sqrt{1 + P^2 + \frac{|a|^2}{2}}$$

The particles are weighted onto the grid in the standard way to evolve the plasma fields and are used to calculate a susceptibility term for the laser's envelope equation via, $\chi = -\sum\frac{q}{\gamma}$. This idea was successfully implemented in a code called turboWAVE [7] and more recently in VORPAL [8].

When using the ponderomotive guiding center approximation the smallest spatial scalelength is now the wavelength of the wake. In addition, satisfying the Courant condition requires that the time step resolve the plasma frequency. Therefore, this approximation could lead to a savings of $(\omega_o/\omega_p)^2$; however, the savings is not quite this much it is found that one still needs to resolve spatial scales of the harmonics of the wake since it typically gets nonlinear.

2.1.3 Quasi-static PIC

Another level of approximation is to use the quasi-static or frozen field approximation. There are various ways of attempting to implement such an approximation, e.g., the choice of gauge. One such implementation is that in QuickPIC, which starts from the Maxwell equations in the Lorentz gauge,

$$(\frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \nabla^2)\phi = 4\pi\rho \tag{1}$$

$$(\frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \nabla^2)\mathbf{A} = \frac{4\pi}{c}\mathbf{j} \tag{2}$$

One transforms to the (x, y, s, ξ) coordinates, where $s = z$ (z is the direction in which the beam is moving), $\xi = t - z/c$; then the quasi-static approximation amounts to assuming $\partial_s \approx 0$. Then a set of full quasi-static equations can be written as,

$$\nabla_\perp^2 \phi = -4\pi\rho \tag{3}$$

$$\nabla_\perp^2 \mathbf{A}_\perp = -4\pi\mathbf{j}_\perp/c \tag{4}$$

$$\nabla \cdot \mathbf{A}_\perp = -\frac{\partial\Psi}{\partial\xi} \tag{5}$$

and the equations of motion are,

$$\frac{d\mathbf{P}_{b\perp}}{ds} = -q_b\nabla_\perp\Psi, \tag{6}$$

$$\frac{d\mathbf{P}_{p\perp}}{d\xi} = \frac{q_p}{1 - V_{e//}/c}[\mathbf{E}_\perp + \frac{(\mathbf{V}_p \times \mathbf{B})_\perp}{c}], \tag{7}$$

In Eq. (5) and (6), $\Psi = \phi - A_{//}$, where $A_{//}$ is the longitudinal component of vector potential, q is the charge of particle, and ρ is the charge density. $\mathbf{V}$, $\mathbf{P}$ are velocity and momentum respectively. The subscripts b and p denote beam and plasma respectively.

The equations above only involve two dimensions, which are perpendicular to the beam propagation direction, so Eqns. (3) and (4) can be solved in 2D space. Once the potentials are calculated, the velocity and position of particles can be updated using (6) and (7). The axial velocity can then be calculated using the constant of the motion,

$$P_{pz}/(m_e c) = \frac{1 + P_{p\perp}^2/(m_e c)^2 - [1 - q_p\psi/(m_e c^2)]^2}{2[1 - q_p\psi/(m_e c^2)]}$$

The charge and current density depend on the axial velocity via

$$\rho_e = \frac{1}{V}\sum_i \frac{q_{ei}}{1 - V_{e//i}/c}, \tag{8}$$

and,

$$\mathbf{j}_{e//} = \frac{1}{V}\sum_i \frac{q_{ei}V_{e//i}}{1 - V_{e//i}/c}\hat{z}. \tag{9}$$

For plasma electrons or ions, the above loop is done every time step $\Delta\xi$, which needs to resolve the plasma frequency. On the other hand for beam electrons, positrons, or for a laser, the time step is Δs, which only needs resolve the betatron frequency or the laser Rayleigh length.

2.2 Codes

Under SciDAC three basic production codes (OSIRIS, VORPAL, and QuickPIC) and one parallel PIC Framework (UPIC) are continually being developed, benchmarked against each other, validated against experiment, and used to unravel complex and nonlinear physics.

2.2.1 OSIRIS

OSIRIS: OSIRIS [1] is a fully explicit, multi-dimensional, fully relativistic, parallelized PIC code. It has a moving window for modeling short lasers or particle beams moving through long regions of plasma. It is written in Fortran95 and takes advantage of advanced object oriented programming techniques that allow for modifications to the code while maintaining full parallelization. The parallelization is done using 1D, 2D, and 3D domain decomposition using MPI. There are 1D, 2D, and 3D versions that can be selected at compile time by changing one line of code. It also has a sophisticated array of diagnostic and visualization packages for rapidly processing large 1D, 2D, and 3D data sets. Visualization scripts exist for IDL, OpenDX, and MatLAB.

It was already very mature before the SciDAC project began. During the SciDAC project the code has been thoroughly benchmarked. This benchmarking includes comparisons against QuickPIC which is based on a completely different model. The fact that they agree gives much confidence in both models and codes. An important new feature is sub-cycling for beam particles where beam particles are only pushed every N time steps [9]. This can lead to a substantial savings if the value of N greatly exceeds the number of cells in the beam propagation direction or if the number of beam particles is considerably more than the background plasma particles. Other relatively new features include, electrostatic and non-relativistic particle push routines that can be used to more efficiently push the ions; a diagnostic which keeps track of the accumulated $p \cdot E$ for a pre-selected group of particles to give a detailed account of particles' energy gain; load balancing for "four corner" partitions; smoothing routines for the fields (as well as the current); and a function parser that permits complicated functions to be used to describe the density profile or external fields. It has also been highly optimized on a single processor while maintaining high parallel efficiency, i.e., >90% efficiency on over 1000 processors. OSIRIS has also been ported to a variety of platforms including clusters of Intel processors (Xeon's) running the Linux operating system and Apple G5 x-server processors running MaxOSX.

In addition, more realism has been added including electron impact and field ionization (the ADK and a barrier suppression model), and a relativistically correct Coulomb collision operator. This work benefited greatly from the ability to benchmark the results against an existing ionization model in OOPIC [5]. OSIRIS has been used extensively and successfully to model laser wakefield acceleration, plasma wakefield acceleration, ion acceleration, and all optical injection. This includes full-scale 3D modeling of these phenomena.

2.2.2 VORPAL

VORPAL: VORPAL [2] is a plasma and particle simulation code that takes maximal advantage of object-oriented programming techniques in C++ to provide a greater level of flexibility. Template meta-programming techniques are used to enable simulations in 1, 2 or 3 physical dimensions, with dimensionality specified at run time. VORPAL runs on serial (Linux, Unix, MacOSX, Windows) and parallel (IBM SP, Linux cluster) platforms. Parallel execution, based on the message passing interface (MPI) includes a flexible domain decomposition that allows for any 1-D, 2-D or 3-D decomposition that can be represented as a collection of slabs. Output data is stored in the HDF5 hierarchal data format, and various visualization scripts and utilities have been developed using IDL, OpenDX and Python/GnuPlot.

The plasma or particle beam and the electromagnetic fields can be represented by a variety of different models. Both PIC and fluid representations for the plasma are available, with several different choices for the particle dynamics and both Euler and cold (no pressure term) models for the fluids. The cold

fluid model for electrons can correctly handle vacuum/fluid interfaces and passage of high-power laser pulses through density ramps. Similarly, the electromagnetic fields can be modeled in a variety of ways, including fully electromagnetic models based on the Yee mesh (2^{nd}-order explicit and implicit; 4^{th}-order explicit), a ponderomotive guiding center model for laser pulses [7,8], and an electrostatic model using the AZTEC libraries [10]. Perfectly matched layers (PML's) have been added to efficiently absorb electromagnetic radiation at the boundaries. External fields that vary as arbitrary space-time functions are also available. These models can be used in conjunction with each other to create hybrid simulations, such as modeling the electrons as a fluid and the ions via PIC.

VORPAL supports PIC simulations with tunneling and impact ionization processes using the ionization physics models implemented in the IONPACK library developed at Tech-X Corp. These models were successfully benchmarked against OOPIC [5], and hence against OSIRIS as well.

QuickPIC: QuickPIC [3] is a newly developed, highly efficient, fully parallelized, fully relativistic, three-dimensional particle-in-cell model for modeling plasma and/or laser wakefield acceleration. When QUICKPIC can be used instead of a full PIC code, *this algorithm reduces the computational time by 2 to 3 orders of magnitude without loss of accuracy.* In figure 1 we present benchmarking results which show that the remarkable agreement between QuickPIC and the conventional fully explicit models (OSIRIS) for highly nonlinear PWFA and LWFA cases. The model is based on the quasi-static or frozen field approximation described above, which reduces a fully three-dimensional electromagnetic field solve and particle push into 2d field solves and both 2D and 3D particle pushes. This is done by calculating the plasma wake assuming that the drive beam and/or laser does not evolve during the time it takes for it to pass through a region of plasma. The complete electromagnetic fields of the plasma wake and its associated index of refraction are then used to evolve the drive beam and/or laser using very large time steps.

The development of QuickPIC is a success story for the rapid construction of a new code using well trusted reusable parallel code. The basic equations and algorithms were developed from a deep understanding of the underlying physics involved in plasma and/or laser wakefield acceleration as well as ideas used in previous developed azimuthally symmetric 2D code [6]. The code embeds a two-dimensional (x,y) PIC code which advances the plasma particles in the ξ variable into a three-dimensional PIC code (ξ,x,y) which advances the beam particles in the s variable. All of the key pieces for QUICKPIC were taken from the UPIC Framework described below. As a result the code was constructed rapidly and both the 2D and 3D pieces were parallelized from the start and key routines in the code which use the most CPU time are highly optimized. It is also worth noting that UPIC has also been under construction during the development of QuickPIC, but because of sophisticated software engineering techniques, the simultaneous development of both has gone very smoothly.

A basic version of QuickPIC is also being used to study the electron-cloud problem [11]. This is also a major success story for rapid development of new accelerator code. Although the interaction of a beam with an electron cloud in a circular accelerator is similar to the beam-plasma interaction in a plasma wakefield accelerator, there are obviously certain differences that must be accommodated in the code. First of course is the effect of a circular machine instead of a straight plasma section. We have added the lattice equations to QuickPIC to include the effect of circulating in the accelerator, including chromaticity and betatron motion in external magnets. We showed that the image charges from the beam pipe play an important role and these are included via QuickPIC's conducting boundary conditions. In addition, subtle effects are important in the e-cloud problem due to the extremely long interaction (over 100,000 km has been modeled) that are not important in plasma wakefield accelerators
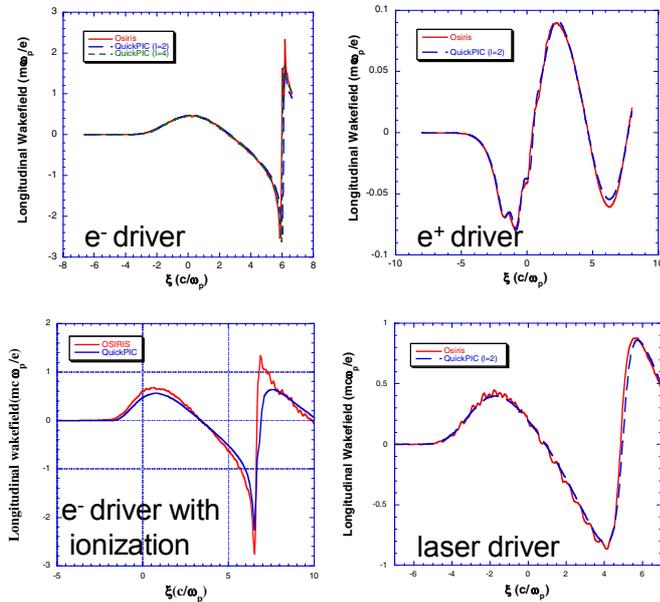
Fig. [1] Comparison between full PIC (OSIRIS) and quasi-static PIC (QuickPIC). Each plot has the accelerating fields with the electron, positron, or laser beam propagating to the left.

UPIC:

The UCLA Parallel PIC Framework (UPIC) [4] is being developed to provide trusted components for the rapid construction of new, parallel Particle-in-Cell (PIC) codes. The Framework uses object-based ideas in Fortran95, and is designed to provide support for various kinds of PIC codes on various kinds of hardware. It is carefully designed to hide the complexity of parallel processing.

The Framework is designed with layers. The lowest layer consists of highly optimized routines. The middle layer primarily provides a much safer and simpler interface to the complex legacy subroutines by encapsulating many details, such as data layouts on parallel machines. The upper layer consists of powerful high level classes that can easily be reused for those parts of the code which the new code developers do not intend to modify. The most CPU time-consuming parts of a PIC code are the particle push and charge deposit. These subroutines have been carefully written to provide the highest performance possible. The layer concept and optimization methods are language independent.

The UPIC Framework currently supports electrostatic, quasi-static and Darwin, and fully electromagnetic forces, with relativity as an option in 2D and 3D. Particle boundary conditions can be periodic, reflecting and mixed periodic/reflecting.

It is important for research codes to support multiple models and numerical schemes. For example, linear interpolation is generally used for PIC codes, but the easiest way to verify that linear interpolation was sufficiently accurate is to run a few cases with quadratic interpolation and see if the results changed. Different algorithms are used with different hardware. For example, a different scheme is used to deposit charge on a vector machine than on a RISC processor. The legacy layer has available subroutines to support linear and quadratic interpolation, both message-passing (MPI) and shared memory parallel programming (OpenMP or pthreads), and can support both RISC and vector architectures.

3. Scientific Discovery

There have been numerous scientific discoveries that have resulted from the above set of codes.

Modeling plasma wakefield acceleration experiments at SLAC: An important aspect of SciDAC is to provide full scale, 3D, high fidelity modeling of experiments. The current PWFA experiments (E-164x and E-167) present a great challenge. The plasma physics is highly nonlinear and relativistic, and to accurately model the experiments ionization physics as well as well as the correct neutral gas and beam profiles needs to be included. Full PIC (OSIRIS and OOPIC) as well as quasi-static PIC (QuickPIC) simulations have been used. Experiments closely coupled with simulations have shown clear evidence for plasma focusing, positron acceleration, electron acceleration, and x-ray generation from betatron motion [12]. Using QuickPIC with its factor of 100 speedup, we can now carry out full scale simulations in a few days and we make direct comparison with the experimental diagnostics. QuickPIC with ionization was validated by comparing its results to those from the recent E164X experiment [13] at SLAC. The maximum energy gain from QuickPIC was near 4.5GeV while that observed in the experiment was near 4GeV. Moreover, QuickPIC as helped to design the future E-167 experiment, by determining the optimum neutral density for wake excitation and the possible influence from long term beam dynamics (e.g. hosing, head erosion).

Modeling laser wakefield acceleration experiments at L'OASIS and Rutherford Appleton Laboratory: In recent issue of Nature [Sep. 30, 2004], three independent experimental groups reported observing mono-energetic beams at ~100 MeV, and with ~100nC of charge when a ~10TW laser propagates through ~mm of plasma at a density ~10^{19}cm^{-3}. Full PIC simulations are required in order to capture the self-injection of the electrons. Each of these articles had supporting simulation results. The L'OASIS results were modeled using VORPAL [14] and a VORPAL visualization made the cover of this "dream beam". The Rutherford/Imperial College results were modeled using OSIRIS [15].

In addition, prior to these experiments 3D OSIRIS simulations [16] had predicted that a modest 13TW laser could indeed self-inject electrons after the laser evolved due to a combination of frequency red shifting and group velocity dispersion and that a mono-energetic beam is produced as the electrons rotated in phase space as they dephase with the wake. These simulations did not use identical parameters to those of the experiments but they revealed the essential physics.

Modeling future LWFA experiments:
Within the next three years laser power will increase from ~10 TW to 200TW and perhaps even a PW. Two milestones will be the controlled (i.e. resonant) injection of electrons into the plasma wake and the acceleration of these electrons while maintaining good beam quality and the generation of 1+GeV mono-energetic beams. VORPAL has been used successfully to model a number of theoretical concepts related to controlled electron injection [17]. Full-scale 3D OSIRIS simulations have been carried out that predict that a 200TW laser will produce a mono-energetic 1.5GeV beam with 1nC of charge. This simulation followed .5x10^9 particles on a 4000x256x256 grid for 300,000 time steps.

Modeling 1TeV energy doubler stages: The PWFA experiments at SLAC described earlier indicate that it might be possible to double the energy of an existing beam using a short plasma cell. This energy doubler concept is called the "afterburner" [18]. Using QuickPIC one can now model a 1TeV PWFA stage in a 5,000 node hours as compared with 5,000,000 node hours using a full PIC algorithm. TeV afterburner simulations are being carried out on the DAWSON 512 G5 x-serve cluster. In one simulation, a drive beam with 3E10 electrons generates a plasma wake and a witness beam with 1E10 electrons samples the large acceleration gradient. Both beams start with 500 GeV energy and at the end of the simulation the witness beam is accelerated to 1 TeV with 5% energy spread, Fig. [2]. This is the first time a full-scale simulation for a 1 TeV afterburner has been done. The result shows that it is possible to double the beam energy in a PWFA. Many issues, such as hosing instability and wake evolution can now be studied in detail using QuickPIC.
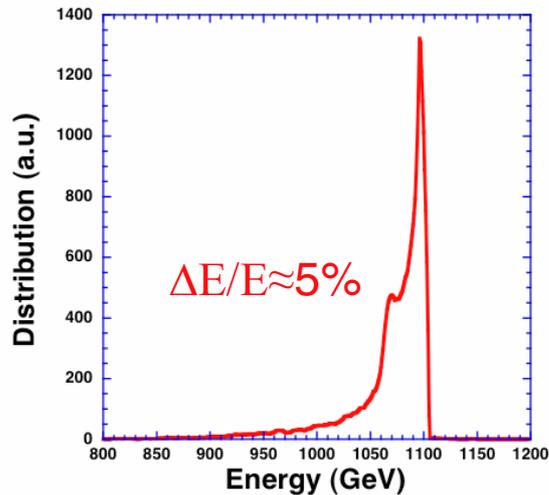
Fig. [2] Energy spectrum of the trailing beam in a 1 TeV afterburner simulation.

<u>Applying plasma codes to e-clouds:</u> Electron clouds have been shown to be associated with limitations in particle accelerator performance in several of the world's largest circular proton and positron machines. The electron-cloud effect will be important to LHC, SNS and Fermilab upgrades, and is already important for RHIC, PEP-II, BNL booster, AGS, LANSCE PSR and KEK. Electrons accumulate in the vacuum chamber where a positively charged bunched particle beam propagates because of a multipacting process which involves primary electron generation (e.g., from residual gas ionization or from photo emission at the inner pipe wall due to synchrotron light) and their multiplication through secondary emission at the wall. The presence of an electron cloud inside the beam chamber can make the beam unstable via bunch-to-bunch or head-tail coupling. These instabilities are basically relativistic beam-plasma phenomena, albeit ones that take place over 100,000 km of beam propagation. The reusable parallel codes developed by the SciDAC team are ideal for modelling the complex beam dynamics with high-fidelity. QuickPIC has been modified to treat the cloud as a non-neutral plasma and to include the circular machine physics. It has successfully modelled LHC designs for thousands of turns without making the coarse multi-kick approximations used in all other models of the interaction (clouds at only a few points around the ring rather than the real continuous distribution). Still this represents only milliseconds in the 30 minute lifetime for the circulating beam, and further work is underway to enhance the speed of the model using a pipelining algorithm. An example of the cloud compression toward the beam in the LHC is shown in one of the magnetized dipole sections of the ring in Fig. 3.
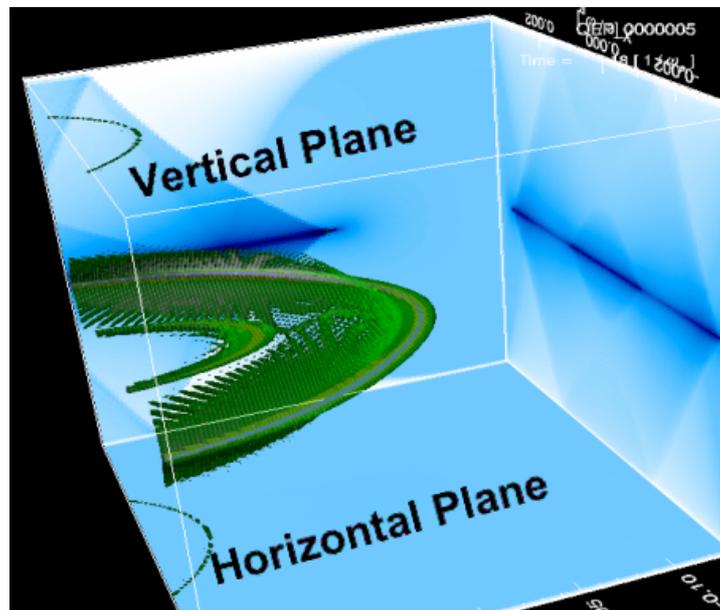
Fig. 3 3D image of perturbed cloud density caused by a proton beam moving to the right in the dipole magnet region of the LHC. Electrons drawn in the beam's potential compress primarily in the vertical plane (along the magnetic field).

3. Vision for the future

We have described the impressive capability and progress that has been achieved for modeling ultra-high gradient plasma acceleration. This capability includes full PIC, as well as quasi-static PIC and fluid models. Improvements that could dramatically improve the utility of full PIC are improved numerical dispersion and mesh refinement. For quasi-static PIC critical research areas include determining a procedure for determining when a plasma particle is self-trapped and then "promoting" it to a beam particle, (this would also be useful for ponderomotive guiding center PIC), determining how to extend ponderomotive guiding center PIC to higher laser intensities, mesh refinement and pipelining. In the current version of QuickPIC the beam particles are not pushed until the end of a 2D time step when all of the plasma fields have been calculated. In pipelining, beam particles are pushed as soon as the plasma fields at their axial location have been calculated. In addition, as the architectures of future massively parallel computers become known we will need to determine how to best partition the field and particle data as well as determine the optimum balance between MPI, OpenMP, and pthreads. In summary, both the science of and simulation capability for plasma acceleration are progressing rapidly and based on the lessons learned the road map towards continued progress is clear.

References

[1]   R.G.Hemker, "Particle-In-Cell Modeling of Plasma-Based Accelerators in Two and Three Dimensions," Ph.D thesis, University of California, Los Angeles, 2000; R.A.Fonseca, et al., " OSIRIS: A Three-Dimensional, Fully Relativistic Particle in Cell Code for Modeling Plasma Based Accelerators, " P.M.A. Sloot et al., editors, ICCS 2002, Lecture Notes Comput. Sci. 2331, pp. 342-351, 2002.

[2]   C. Nieter and J.R.Cary, "VORPAL: a versatile plasma simulation code", J. Comput. Phys. 196, 448 (2004).

[3]   C. Huang et al., "QuickPIC: A highly efficient PIC code for modeling wakefield acceleration in plasmas," submitted to J. Comp. Phys.

[4]   V. K. Decyk and C. D. Norton, "UCLA Parallel PIC Framework", Comp. Phys. Comm. 164, 80 (2004).

[5]   D.L.Bruhwiler et al., "Particle-in-cell simulations of plasma accelerators and electron-neutral collisions", Phys. Rev. STAB, 101302 (2001); D.L.Bruhwiler et al., "Particle-in-cell simulations of tunnelling ionization effects in plasma based accelerators", Phys.Plasmas 10, 2022 (2003).

[6]   P. Mora and T.M. Antonsen, "Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas," Phys. Plasmas 4, 217 (1997).

[7]   Gordon, D.F., et al., "IEEE Transactions on Plasma Science", **28**, No. 4, 1135-1143 (2000).

[8]   P.Messmer et al., "Particle-in-cell simulations of low density cm-scale capillary channels", Proc. Particle Accel. Conf. (2005) in press.

[9]   S.Deng et al., "Developing a multi-timescale PIC code for plasma accelerators," Proc. Particle Accel. Conf. (2005) in press.

[10]   P. Messmer and D.L. Bruhwiler, "A parallel electrostatic solver for the VORPAL code," Comp. Phys. Comm. **164**, 118 (2004).

[11]   A. Ghalam et al., "3-D Parallel Simulation Model of Continuous Beam-Cloud Interactions," Proc. ECLOUD04: 31st Advanced ICFA Beam Dynamics Workshop on Electron Cloud Effects, Napa, CA April 19-23, 2004; M. Furman, S. Henderson and F. Zimmerman, Eds., p. 325 CERN Report 2005-001 (CERN, Geneva, SW 2005); G. Rumolo et al., "Electron cloud effects on beam evolution in a circular accelerator," Phys.Rev.STAB, 6, 081002:1-9 (2003).

[12]   P. Muggli et al., "A meter-scale plasma wakefield accelerator driven by a matched electron beam," Phys. Rev. Lett., 93 014802:1-4 (2004); Hogan et al., "Ultra-relativistic positron beam transport through meter-scale plasmas," Phys.Rev.Lett., 90, 205002:1-4 (2004); B.E.Blue et al., "Plasma wakefield acceleration of an intense positron beam," Phys.Rev.Lett., 90, 214801:1-4 (2003); C.E.Clayton et al., "Transverse envelope dynamics of a 28.5 GeV electron beam in a long scale length plasmas," Phys.Rev.Lett. 88, 154801:1-4 (2002).

[13]   M.Hogan et al., "Multi-GeV energy gain in a plasma-wakefield accelerator", accepted in Phys. Rev. Lett.

[14]   C.G.R.Geddes et al., "High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding," Nature 431, 538 (2004).

[15]   S. P. D. Mangles, et al., "Monoenergetic beams of relativistic electrons from intense laser-plasma interactions, "Nature 431, 535 (2004).

[16]   F.S.Tsung et al, "Near GeV energy laser wakefield acceleration of self-injected electrons in a cm scale plasma channel," Phys. Rev. Lett. 93, 185004:1-4 (2004).

[17]   J.R. Cary *et al.*, "Clean beams from optical injection with a cleanup pulse," Phys. Plasmas (2005), accepted.

[18]   S. Lee et al, "Energy doubler for a linear collider," Phys.Rev. STAB, 5, 011001:1-4 (2002).