Computational studies and optimization of wakefield accelerators

C G R Geddes[1], D L Bruhwiler[2], J R Cary[2,3], W B Mori[4], J-L Vay[1], S F Martins[7], T Katsouleas[6], E Cormier-Michel[1], W M Fawley[1], C Huang[4], X Wang[6], B Cowan[2], V K Decyk[4], E Esarey[1], R A Fonseca[7], W Lu[4], P Messmer[2], P Mullowney[2], K Nakamura[1], K Paul[2], G R Plateau[1], C B Schroeder[1], L O Silva[7], Cs Toth[1], F S Tsung[4], M Tzoufras[4], T Antonsen[5], J Vieira[7] and W P Leemans[1]

[1] Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
[2] Tech-X Corporation, Boulder, CO 80303, USA
[3] University of Colorado, Boulder, CO 80309, USA
[4] University of California, Los Angeles, CA 90095, USA
[5] University of Maryland, College Park, MD 20742, USA
[6] University of Southern California, Los Angeles, CA 90089, USA
[7] GoLP/IPFN Instituto Superior Tecnico, Lisbon 1049-001, Portugal

E-mail: cgrgeddes@lbl.gov

Abstract. Laser- and particle beam-driven plasma wakefield accelerators produce accelerating fields thousands of times higher than radio-frequency accelerators, offering compactness and ultrafast bunches to extend the frontiers of high energy physics and to enable laboratory-scale radiation sources. Large-scale kinetic simulations provide essential understanding of accelerator physics to advance beam performance and stability and show and predict the physics behind recent demonstration of narrow energy spread bunches. Benchmarking between codes is establishing validity of the models used and, by testing new reduced models, is extending the reach of simulations to cover upcoming meter-scale multi-GeV experiments. This includes new models that exploit Lorentz boosted simulation frames to speed calculations. Simulations of experiments showed that recently demonstrated plasma gradient injection of electrons can be used as an injector to increase beam quality by orders of magnitude. Simulations are now also modeling accelerator stages of tens of GeV, staging of modules, and new positron sources to design next-generation experiments and to use in applications in high energy physics and light sources.

1. Introduction

Particle accelerators are among the most powerful instruments of scientific discovery. A TeV-class linear collider to extend the energy frontier will require order of 20 km-long conventional radio-frequency (RF) accelerating linacs [1], while machines such as the LCLS will use kilometer-scale linacs to drive undulators for unprecedented x-ray brightness [1]. Greatly increased accelerating gradient is then needed to scale beyond TeV energies and to provide bright, laboratory-scale radiation sources.

Figure 1. Schematic of an LWFA, from a VORPAL 3D simulation on 4096 processors of a centimeter-scale 0.7 GeV stage (A): the radiation pressure of a laser pulse (red) displaces plasma electrons creating a space charge wave (purple-blue) which accelerates particles, creating a compact bunch (green). Simulations resolve phase space and internal dynamics to optimize WFAs, showing for instance (B) particles in px-x-y space above plasma density in a 100 MeV self-trapped stage (VORPAL 3D).

In laser or particle beam driven plasma wakefield accelerators (LWFAs or PWFAs) [2, 3], the radiation pressure of an intense laser pulse (or space charge of a particle beam) displaces electrons, and the subsequent oscillation of the plasma creates a plasma wave (wake) following the driver (figure 1). The field of the wake is limited by trapping of particles and can be thousands of times that in RF accelerators. The wake period is micrometer scale, producing femtosecond bunches well suited for light sources [4].

Laser wakefield accelerators have produced electron bunches with percent energy spread and low emittance from electrons self-trapped by the wake. Bunches at ~0.1 GeV were produced in a few millimeters using plasma channels [5] or large laser spots [6,7] to extend the interaction length. Using a plasma channel to extend the interaction to centimeter scale, GeV bunches and stable operation at 0.5 GeV were observed [8]. Control over injection has now produced bunches with an order of magnitude lower absolute momentum spread and with stability over several days using plasma ramps [9], and produced bunches with tunable energy [10] using the colliding pulse method [11]. A PWFA has doubled the energy of a fraction of the SLAC beam, achieving energy gain of 42 GeV [12]. Plasma accelerator applications now require development and staging of ≥10 GeV LWFAs, further control over trapping and acceleration to reduce and stabilize momentum spread, and narrow energy spread PWFAs. Plasma scale will increase from centimeter to meter scale, and detailed bunch kinetics will be essential.

Simulations describe nonlinear plasma response, beam trapping, and self-consistent acceleration not accessible to analytic theory. They provide information on internal dynamics to optimize WFAs, and were, for example, essential in modeling the physics of narrow energy spread bunches from self-trapped electrons observed in recent years and evaluating the scaling of these experiments to higher energies as described in [5, 13-16] and references therein, among others.

The VORPAL [17] and OSIRIS [18, 19] codes provide explicit particle-in-cell simulations that resolve the laser wavelength (the shortest major scale) and particle kinetics and were used to model recent experiments. VORPAL now allows fluid description of the wake (reducing noise) with particles for the bunch. Resolving the laser wavelength over the propagation distance (~10^6 steps for centimeter-scale GeV simulations) and wake volume (~10^8 cells in 3D) drives the computational cost, which is order 10^6 hours for centimeter-scale GeV runs. Envelope simulations decrease cost by resolving the laser envelope but not its fast oscillation, allowing reduced resolution, and quasistatic codes further reduce cost by assuming slow evolution [20]. These assumptions, which require validation in the regime simulated, allow simulation of the same plasma and laser parameters with a 100 to 10000 times savings in CPU time. Envelope models are implemented in VORPAL and quasi-static in QuickPIC [21]. Performing calculations in a boosted frame [22] can also reduce cost, and this

method is now being used in WARP[23], VORPAL, and OSIRIS. Optimizing methods for kinetic accuracy are also vital [24].

Here, we describe recent large-scale kinetic simulation of wakefield accelerators by the ComPASS [25] SciDAC project. Section 2 describes a benchmarking and scaling program that confirmed agreement of the codes and is being used to develop the reduced models needed to simulate future large experiments. Section 3 describes modeling of controlled injection in LWFAs towards development of multi-GeV stages with increased stability and reduced momentum spread, and modeling of 10-100 GeV LWFAs. Section 4 describes modeling of PWFA's at 10's of GeV and new positron sources. Finally, in section 5, development of the boosted frame algorithm is described which will be important to simulate m-scale plasmas. PIC algorithm studies were also conducted to optimize kinetic accuracy.

2. Benchmarking and code comparison

Software verification is an important part of the SciDAC mission, and benchmarking exercises are an effective way to achieve this goal in regimes with no analytical results, which is the case for nonlinear 3D plasma wakes. We consider the important example of an intense $[I_{peak} \sim 10^{18}$ W cm$^{-2}]$, ultra-short $[\tau_{fwhm} = 30$ fs$]$ Ti-Sapphire $[\lambda_0 = 0.8$ μm$]$ laser pulse entering a uniform density plasma $[n_e = 1.38 \times 10^{19}$ cm$^{-3}]$, close to the parameters of recent experiments. We compare 3D simulation results from the time-explicit PIC codes OSIRIS and VORPAL with each other and with the quasi-static code QuickPIC, for several values of a_0, the normalized peak laser field. The rectangular mesh has 512x512x512 cells for the time-explicit case, with 8 macro-particles per cell, for a total of 1.34×10^8 cells and 1.07×10^9 particles. The longitudinal grid spacing is 0.04 μm (20 cells per λ_0), and transverse spacing is 0.16 μm.

Figure 2 presents 2D contour plots of the simulated accelerating electric field from each code, in a slice taken from the center of the 3D domain, with longitudinal position along the horizontal axis and transverse position along the vertical axis. The peak accelerating field for trapped electrons is near the middle of the center-left blue region. Acceleration and transverse focusing overlap throughout the right half of this region (roughly), which is where one typically finds trapped, accelerated beams. The lower-right plot in Figure 2 shows lineouts of the accelerating electric field for all three codes along the center of the simulated domain, normalized to the cold nonrelativistic wavebreaking field $E_{wb} = m_e c \omega_0 / e$. We note that E_0 is about 10x smaller than E_{wb}. Position is normalized to the central wavenumber of the laser pulse. Agreement between VORPAL and OSIRIS is good, lending confidence in the validity of both codes. The fact that QuickPIC also agreed well gives further confidence as it is based on a completely different reduced algorithm. QuickPIC does not resolve the space/time oscillations driven by the laser pulse but accurately models features on the scale of the
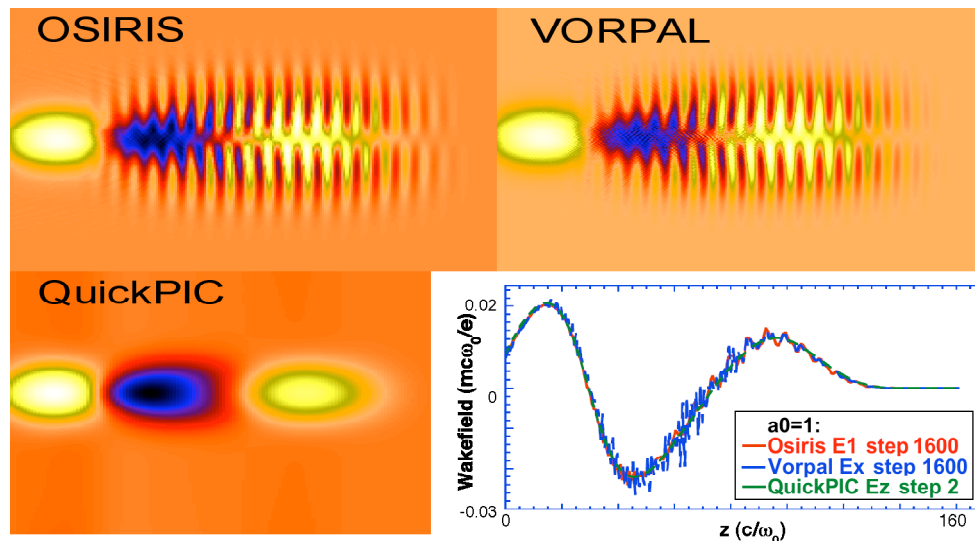
Figure 2. 3D simulation results from OSIRIS, VORPAL, and QuickPIC for an LWFA benchmark problem with $a_0=1$ show good agreement. The normalized accelerating wakefield is shown as a 2D slice (upper and left), and also overlapping lineouts down the center (lower right).

plasma wavelength so long as there are no self-injected electrons. The benchmarking effort is being extended to include VORPAL fluid and envelope models as well as the quasi-static model in WAKE.

These benchmarking exercises were repeated for $a_0=0.5$, 1, 2 and 4. Overlapping lineouts like that in figure 2, are shown in figure 3 for all values of a_0. For $a_0=1$, 2, and 4, these plots include data from OSIRIS and VORPAL simulations using second-order splines for the particle shapes. The higher-order shape, used both for charge-conserving current deposition and force interpolation, filters out high-wavenumber components of the particle currents and forces, which can greatly reduce noise and unphysical effects such as grid heating. Filtering must be done with care, however, as splines and higher-order particle shapes can modify physical effects, like the dispersion relation for plasma waves and the wavelength of the wake. As seen in figure 3, the spline-based particles in OSIRIS and VORPAL agree well with standard first-order shapes and with each other, for $a_0=1$ and 2, and show only modest differences for $a_0=4$. QuickPIC also supports higher-order particle shapes.

Extensive scaling of PIC has been performed on large computers to identify areas where current algorithms can be improved to enable scaling of PIC codes to 100,000 processors. Strong scaling studies with OSIRIS, VORPAL, and UPIC have recently been conducted on a billion-particle benchmark simulation with a 512x256x512 grid, using an electromagnetic, relativistic plasma model. The UPIC study is directly applicable to QuickPIC, which depends on this framework. Scaling was carried out on an Opteron-based cluster with Infiniband (ATLAS). The number of processors was varied from 128 to 8192 for UPIC, 4096 for OSIRIS, and 1024 for VORPAL keeping the simulation size fixed. UPIC showed good scaling (92%) for the particle part of the calculation (which typically dominates PIC codes). The fast Fourier transform (FFT) field solve in UPIC scaled well up to 4096 processors, then saturated at 8192. In order to enable this code (and other spectral PIC codes) to scale to 100,000 processors, strategies to improve existing algorithms in the particle manager and the FFT were identified, and a mixed MPI/threaded model has been implemented and is being tested. For OSIRIS both the particle push and field solve scaled well up to the maximum used (4096 processors). Top processing speeds of more than 10 billion particles per second (for an entire step including field solve, etc.) were obtained for UPIC. OSIRIS was found to be 30% slower. VORPAL demonstrated 90% efficiency on up to 1024 processors for the same case, and runs on 4096 processors are in progress.

Figure 3. Overlapping lineouts down the center of the simulation domain show good agreement between 3D OSIRIS, VORPAL and QuickPIC simulations of the LWFA benchmark problem for a_0=0.5, 1, 2, and 4, including the use of second-order spline-based particle shapes.

The VORPAL FDTD Maxwell update scaled well to 8192 processors on the Franklin supercomputer at NERSC, demonstrating 95% efficiency with respect to a 4-processor run in weak scaling where the domain was increased to keep equal cells per processor [26]. Tests of I/O scaling on up to 8192 processors also showed that VORPAL's I/O, which uses parallel HDF5, can match that of other optimized applications if the domain sizes are identical on every processor. The restriction appears to come from the HDF5 implementation. For more than 2048 processors, I/O rates begin to decrease, and individual dump files can be used to work around this.

3. Controlled injection, staging, and acceleration

While the energies achieved by LWFAs [8] are sufficient for many radiation source applications [4] and could be staged in series to provide energies needed for high energy physics, improvement of bunch momentum spread and day-to-day accelerator stability are required. Toward this goal, experiments coupled with simulations recently produced stable bunches with longitudinal and transverse momentum spreads an order of magnitude lower than previously observed [9], and show that these bunches can function as injectors to reduce energy spread of high energy LWFAs. The experiments focused a 10 TW laser at the downstream edge of a thin gas jet where density is decreasing, producing bunches with 0.17 (0.02 MeV/c) longitudinal (transverse) momentum spread and with central momenta stable at 0.76±0.02 MeV/c over several days.

Simulation of the experiments [9] using VORPAL explicit PIC showed plasma density gradient control of trapping produced the low energy spread bunches, and indicated that use of such bunches as injectors can greatly improve LWFA bunch quality. In a decreasing gradient, the plasma wavelength λ_p increases, causing the wake fronts to slip behind the laser and decreasing the wake velocity v_ϕ and hence the threshold wake amplitude for trapping and accelerating plasma electrons [27,28]. The simulations showed this modulation of v_ϕ produced trapping without significant modulation of the laser pulse (which is unstable) and at low wake amplitude. It also caused the bunch to quickly outrun the wake structure, producing bunches at ~ 1.5MeV/c, similar to the experiments. The low wake amplitude at trapping also produced low bunch momentum spread, with longitudinal (transverse) momentum spread of 0.2 (0.05 MeV/c), reasonably consistent with the experiments. This low momentum spread, together with the compact bunch diameter of 5μm, indicated a normalized

emittance (focusability) of ~0.4 π mm-mrad, approximately an order of magnitude better than other LWFAs. Consistent with observed stability, simulations showed MeV-class bunches over ±10% variation in laser power, plasma density. and plasma length. Simulating focusing through the jet required a large domain, and hence simulations were performed in 2D (with 10^7 cells versus 10^{10} required in 3D), which may explain differences from the experimental energy. Results have now been reproduced by using the VORPAL envelope model, which will enable 3D simulations.
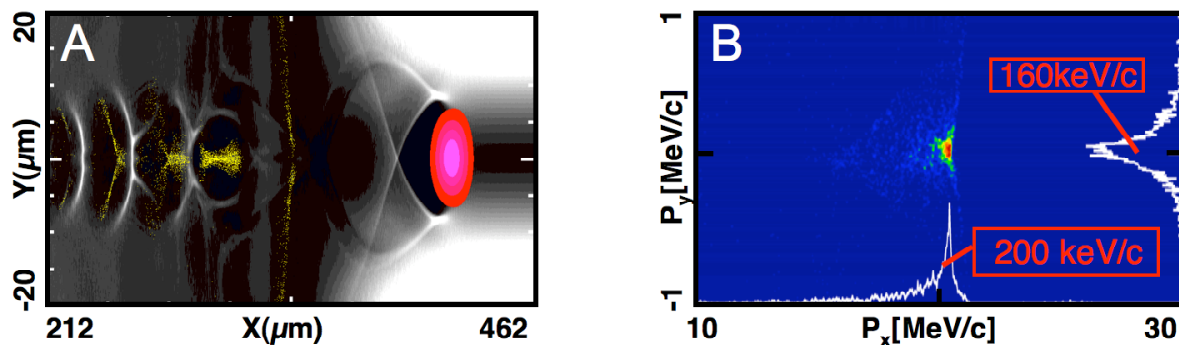

Figure 4. VORPAL simulations merging density gradient injection to an accelerating channel (A) show the laser (red), plasma wake density (grey) and accelerating particles (yellow). Acceleration in the channel (B) produces 20 MeV bunches and preserves the bunch's 0.2 MeV/c momentum spread.

Simulations further showed that merging the ramp into a plasma channel with constant axial density (figure 4) immediately after trapping allows use of these bunches as an injector to improve LWFA bunch quality. This is possible because experiments and simulations show the laser is transmitted through the jet without significant depletion or modulation of the spot, allowing it to drive a wake in the channel. The bunch must also be shorter than λ_p to allow efficient acceleration, and simulations show the trapped bunch meets this condition. The simulated bunch length at the location where the density is correct for emission of THz radiation was benchmarked to THz experimental measurements [29], showing that the simulated length is accurate. This work extended the experimental diagnostics used to compare to simulations to include THz, laser pulse transmission, and bunch transverse momentum in addition to the electron spectrum, increasing the detail and reliability of simulations evaluating injection and postacceleration.

Because the bunch is short compared to the plasma wavelength, it sees a nearly even accelerating field and its momentum spread is nearly preserved as it accelerates in the channel, producing 0.2 MeV/c class momentum spread at high energy. Bunches with 0.2 MeV/c energy spread at energies greater than 20 MeV have so far been demonstrated, limited by computational time with the large domain size. Longer and 3D simulations using the envelope model are in progress to optimize bunch quality, and related simulations [30] indicate this may enable bunches at GeV energies and beyond with < 0.1% energy spread.

Proposed next-generation experiments will use controlled injection coupled with meter-scale plasmas to produce high-quality bunches at ≥10 GeV from a PW class laser. Together with staging, this will be an important step toward high energy physics applications, which could stage multiple 10–100 GeV modules to reach TeV energies. Because meter-scale explicit simulations in multiple dimensions are beyond the capacity of current computers, a combination of Lorentz boosted, envelope, and scaled explicit simulations are being used to detail 10 GeV stage designs (two of which are illustrated in figure 7). Continued simulations of 0.1 to 1 GeV self-trapped stages [15, 31] are validating the codes against additional diagnostics and using this detail to further optimize stage performance. These topics are the subject of upcoming publications; related work was summarized in [15,16] and references therein.

Simulations using QuickPIC have also been used to design high energy, efficient LWFA stages. These designs start from a phenomenological theory [32] that includes the concepts of nonlinear multi-dimensional wake excitation [33], local pump depletion, dephasing, and laser guiding. The simulations show that in this nonlinear blowout regime, a laser can excite a stable wake over distances hundreds of Rayleigh lengths long, as long as its spot size and duration are properly matched: $k_p w_0 = \omega_p \tau_L = 2(a_0)^{1/2}$. In the simulations a_0 is held fixed at 2, and the plasma density is decreased while the spot size is kept matched. Under these conditions the laser power is equal to the critical power for self-focusing, P_c. A preformed channel is used to keep the leading edge of the laser guided. Stages that provide an average gradient 3.6 GV/m (7.2 GV/m) with a final energy of 100 GeV (25 GeV) were demonstrated (figure 5).


Figure 5. Final electron spectra for a series of LWFA simulations in the nonlinear blowout regime for P/Pc=1. In each case a trailing bunch of electrons was injected. The energy gain is shown, where each color corresponds to a different laser power. Energy gain agreed with scalings (QuickPIC).

4. Optimized PWFA design and positron sources

The theory of Lu et al. [33] can allow one to accurately determine the optimum bunch shapes for a particle driver or the load [34] for either the PWFA or LWFA. As an example, we have conducted QuickPIC simulations of idealized beams for a PWFA stage in which the energy of a 25 GeV beam is doubled to 50 GeV. One could imagine stringing 20 of these stages together to produce a 500 GeV electron beam. The parameters are also relevant to possible two-beam PWFA experiments at SLAC. Based on beam-loading theory [34] in the nonlinear blow-out regime, ideal parameters for the drive beam and the witness beam were obtained, and energy doubling of a 25 GeV beam with energy spread within 1% was demonstrated in the simulation (figure 6). In the simulation, we use 512×512×256 grids with box size of 600×600×270 microns. Both the drive beam and the trailing beam have initial energy of 25 GeV and spot size of 3 microns. After propagating about 0.7 meters in the plasma, the trailing beam has gained approximately 25 GeV in energy. We have also conducted higher resolution (with 2048×2048×256 grids) simulations with and without ion motion and synchrotron radiation loss to investigate their effects on beam acceleration and propagation.

Developing a wakefield-based collider requires schemes for accelerating positrons as well as electrons. One solution is to use a nearly linear wake; in this regime dynamics are similar for electrons and positrons. The nonlinear regime offers higher fields, but there are then significant differences between electrons and positrons. It is found that the positron beam wake is weaker than the electron beam wake motivating schemes to accelerate a positron beam in an electron beam or laser driven wake. Generating a positron beam behind an electron beam is, however, challenging. Simulations of a novel idea to do this were carried out with the code OSIRIS [35]. In order to create a properly phased positron beam, a double-pulse electron beam is collided with a thin foil target embedded within the plasma. This creates a double-pulse positron beam of lesser charge overlapping the electron beam through bremsstrahlung. The first electron pulse drives a weakly nonlinear wake (bubble), and the positrons overlapping this pulse are expelled because they reside in a wake region that focuses electrons but defocuses positrons. Nonlinear wakes have restricted positron focusing phase, but in the weakly nonlinear regime there is a phase of the wake just behind the bubble that has both accelerating

and focusing fields for positrons (defocusing for electrons). Placing the second bunch here, the trailing group of electrons will be defocused while the newly created positrons will remain focused and accelerated. Simulations of a possible two-bunch experiment at SLAC show that a large number of positrons (1.2×10^7) are injected and accelerated to 6.2 GeV and with relatively narrow energy spread (15%) after 99 cm of plasma.
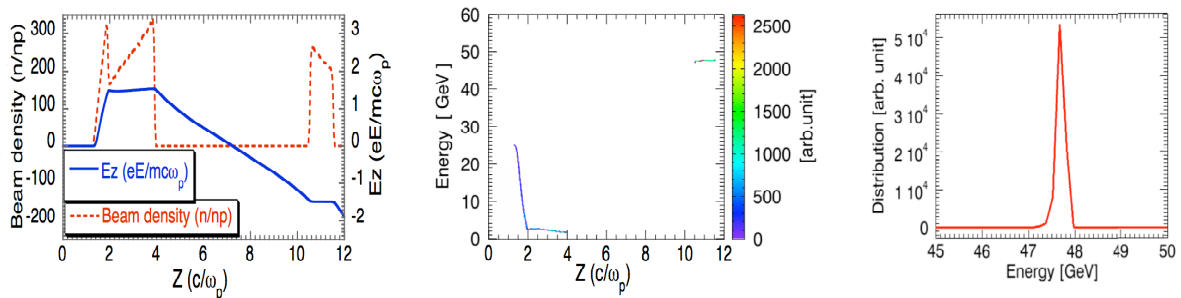
Figure 6. A witness beam with initial energy of 25 GeV and 1.7×10^{10} electrons gains ~25 GeV energy with 0.46% energy spread in a PWFA simulation. The plasma is preionized. Left: the density profile of the drive and witness beam and the longitudinal wake they generated. Middle: Phase space plot of the drive and witness beam, showing energy doubling of the witness beam and depletion of the energy of the drive beam. Right: Energy distribution of the accelerated witness beam (QuickPIC).

5. Boosted frame and PIC kinetic model development

While it is common to model WFAs in the laboratory frame, under certain conditions orders of magnitude speedup can be achieved in modeling relativistic systems by working in a boosted frame [22]. Because of relativistic length contraction and time dilation, the separation of scales between the laser wavelength and the plasma length (which drives the number of timesteps required) is not invariant, and the total number of computer operations can be greatly reduced by choosing a frame for the calculation that is moving at an optimal relativistic velocity.

The boosted frame method is being benchmarked against simulations in the laboratory frame. For example, two-dimensional simulations of coherent spontaneous emission from prebunched beams in linearly polarized undulators agree with standard theory and showed speedups of ~5 orders of magnitude (for $\gamma \sim 250$) over standard EM PIC calculations [36]. Standard PIC or other algorithms can be applied in the boosted frame (with corrections as below), so that the fast oscillation is still resolved unlike other reduced models. Compared to standard eikonal and wiggler-period averaging approximations, the boosted frame undulator calculation recovers the "backward wave" emission and also sideband development in a harmonic cascade.

Three-dimensional simulations of electron cloud effects in high-energy physics accelerators reproduced lab frame results with a speedup of 10^3 but showed the need for new numerical techniques. A new particle pusher was required [37] because errors appeared in the standard Boris push for particles in the boosted frame. The new push may benefit modeling of WFAs also if the improvement is not within EM solve accuracy. The high grid resolution required by standard electromagnetic solvers to accurately model laser group velocity may be aggravated when performing the calculation in the boosted frame due to numerical Cerenkov [38] acting on the plasma traveling near the speed of light. Solvers that are dispersion-free in vacuum [39] along the grid axes for a "magic" time step are being explored [40, 41], but this approach requires mitigation of an even-odd oscillation and potential instability that can develop [40]. Back-scattered radiation is also unresolved in the boosted frame, which is important for many applications.

The Lorentz method has been implemented in all dimensions (1D, 2D, and 3D) in OSIRIS 2.0 [19] and in 1D and 2D in VORPAL [42], dramatically reducing computational resources required to model LWFAs (typically by~γ^2, with γ being that of the boosted frame). While the algorithm needs no

change, running in this frame poses particular difficulties. In addition to backscatter and electromagnetic wave considerations, it is important to eliminate round-off errors that can accumulate from the current of the relativistically moving background electrons and ions.

Simulation of laser plasma acceleration setups were reported in one, two, and three dimensions with speedups of respectively x1,500 [42], x150, and x75 [43]. A 1.5 GeV self-injected beam, already fully studied in the laboratory frame [32], was simulated in OSIRIS, and the final bunch energy (1.5 GeV) and injected charge (~0.5 nC) were in reasonable agreement with the laboratory results and theoretical predictions (figure 7A). The simulation was performed with a gamma of 5, giving a total speedup of ~20 times relative to the standard laboratory simulation. A 10 GeV LWFA stage (0.4 m long at $a_0=1.6$ and $n_e=2\times10^{17}$ cm^{-3}) was simulated in VORPAL using lab and boosted frames. Lab frame simulations required ~3×10^7 steps and 5,000 processor hours, and required special techniques (current smoothing, second-order particle shapes) to prevent artificial injection of plasma electrons (i.e. dark current). Boosted frame simulations used $\gamma=27$ and obtained a speedup factor of ~1,500 with no artificial dark current problems. Figure 7B shows the boosted-frame longitudinal phase space of test electrons, externally injected with a wide range of phases to sample the wakefields. Energy gain in the lab frame was ~8 GeV, in good agreement with lab frame simulations.

While the above methods extend simulation of high energy stages, improved kinetic accuracy is equally important to accurate accelerator design. Studies of the PIC algorithm in modeling of LWFAs [24] using a local charge conserving deposition code [44] showed that use of high-order spline interpolation for force and current deposition, coupled with smoothing of the current on the grid,

Figure 7. Longitudinal phase space from LWFA simulation in the boosted frame of (A) self-injected 1.5 GeV electrons, with energy transformed to the laboratory frame (OSIRIS) and (B) test electrons accelerated to 8 GeV in 1D (VORPAL).

reduced momentum errors by approximately two orders of magnitude. This greatly improved kinetic accuracy at a cost of order 25%. Improvement was slow with resolution, at a cost O[resolution^4] for the 2D studies conducted for all spatial dimensions, time, and particles (O[resolution^5] in 3D), emphasizing the importance of these techniques. The errors arise from discretization (of grid and macroparticles) and from interpolation of forces from the grid and can lead to momentum and orbit displacements that result in unphysical trapping, especially important in modeling of dark current free structures for future experiments. Suppression of unphysical trapping for the first few periods of the wake (which are typically modeled) was demonstrated. These methods are implemented in the codes VORPAL and OSIRIS.

6. Conclusions

Large-scale one-to-one simulations, benchmarked to experiments, revealed the physics of formation of narrow energy spread bunches from self-trapping at 0.1 GeV and 1 GeV. They are now providing quantitative understanding for the design of new accelerators. Simulations are being used to develop controlled injection for increased beam quality, accelerator stages for ≥ 10 GeV energies, positron

sources, and staging of accelerators to reach high energies. Code development will continue to improve and verify reduced and boosted frame models and to improve the kinetic accuracy of the codes. Inclusion of additional experimental diagnostics will further constrain the codes to increase the detail of information available for accelerator optimization. The new models developed will be used together with scaling to tens of thousands of processors to simulate controlled injection, staging, and meter-scale 10 GeV class experiments with high accuracy and to design next-generation accelerators to push the energy frontier in high energy physics and to develop new light sources.

Notes
E Esarey is also at the University of Nevada, Reno; G R Plateau is also at the University of California, Berkeley; W P Leemans is also at Ecole Polytechnique, France.

References

[1]   ILC 2008 www.linearcollider.org/cms ; LCLS- www-ssrl.slac.stanford.edu/lcls/
[2]   Tajima T and Dawson J M 1979 *Phys. Rev. Lett.* **43:**267
[3]   Esarey E, *et al.* 1996 *IEEE Trans. Plasma Sci.* **24:**252
[4]   Leemans W P, *et al.* 2005 *IEEE Trans. Plasma Sci.* **33:**8
[5]   Geddes C G R, *et al.* 2004 *Nature* **431:**538
[6]   Mangles S P D, *et al.* 2004 *Nature* **431:**535
[7]   Faure J, *et al.* 2004 *Nature* **431**:541
[8]   Leemans W P, *et al.* 2006 *Nature Physics* **2:**696
[9]   Geddes C G R, *et al.* *Phys. Rev. Lett.* **100:**215004.
[10] Faure J, *et al.* 2006 *Nature* **444:**737
[11] Esarey E, *et al.* 1997 *Phys. Rev. Lett.* **79:**2682
[12] Blumenfeld I, *et al.* 2007 *Nature* **445:**741
[13] Pukhov A and Meyer-ter-Vehn J 2002 *Appl. Phys. B* **74:**355
[14] Tsung F S, *et al.* 2004 *Phys. Rev. Lett.* **93:**185002
[15] Geddes C G R, *et al.* 2007 *J. Phys. Conf. Ser.* **78:**12021
[16] Tsung F S, *et al.* 2007 *J. Phys. Conf. Ser.* **78:**12009
[17] Nieter C and Cary J 2004 *J. Comp. Phys.* **196:**448
[18] Hemker R G, *et al.* 1999 *Proc. 1999 PAC Conference*
[19] Fonscea R A, *et al.* 2002 *Lect. Notes Comp. Science* **2329** III-34 (Springer-Verlag)
[20] Mora P, *et al.* 1997 *Phys. Plasmas* **4:**217
[21] Huang C, *et al.* 2006 *J. Comp. Phys.* **217:**658-679
[22] Vay J-L 2007 *Phys. Rev. Lett.* **98:**130405
[23] Grote D P, Friedman A, Vay J-L and Haber I 2005 *AIP Conf. Proc.* **749:**55
[24] Cormier-Michel E, *et al.* 2008 *PRE* (accepted)

[25] Spentzouris P, *et al*. 2008 *Proc. SciDAC 2008* (also compass.fnal.gov)
[26] Cary J, *et al*. 2008 *Proc. SciDAC 2008*
[27 Katsouleas T 1986 *Phys. Rev. A.*, **33**:2056
[28] Bulanov S, *et al*. 1998 *Phys. Rev. E* **58:**R5257
[29] Plateau G, *et al*. 2007 *Proc. Particle Accel. Conf.* (IEEE: Piscataway, NJ)
[30] Andreev N E, *et al*. 2000 *Phys. Rev. ST-AB* **3:**021301
[31] Gonsalves A J, *et al*. 2007 Proc. Particle Accel. Conf. (IEEE: Piscataway, NJ)
[32] Lu W, *et al*. 2007 *Phys. Rev. ST Accel. Beams* **10:**061301 (2007)
[33] Lu W, *et al*. 2006 *Phys. Rev. Lett*. **96:**0165002; Lu W, *et al*. 2006 *Phys. Plasmas*. **13:**056709
[34] Tzoufras M, *et al*. 2008 (submitted for publication); Mtzoufras, Ph.D. UCLA 2008.
[35] Wang X, *et al*. 2008 (submitted for publication)
[36] Fawley W M, *et al*. 2008 (in preparation)
[37] J.-L. Vay, 2008 *Phys. Plasma* 15, 56701 (2008).
[38] review: Greenwood *et al*. *J. Comp. Phys*. 201, 665 (2004).
[39] Karkkainen M, *et al*. 2006 *Proc. ICAP* Chamonix, France (2006)
[40] Vay J-L 2008 (unpublished)
[41] Cowan B 2008 (unpublished)
[42] Bruhwiler D L 2007 (unpublished)
[43] Martins S F, *et al*. 2007 Mtg. Div. of Plasma Phys. (Orlando, FL)
[44] Ruhl H 2000 Habilitationsschrift (Technische Universitt Darmstadt)