Large Fields for Smaller Facility Sources

Compared to conventional particle accelerators, plasmas can sustain accelerating fields that are thousands of times higher. To exploit this ability, massively parallel SciDAC particle simulations provide physical insight into the development of next-generation accelerators that use laser-driven plasma waves. These plasma-based accelerators offer a path to more compact, ultra-fast particle and radiation sources for probing the subatomic world, for studying new materials and new technologies, and for medical applications.

Introduction

Particle accelerators are among the largest and most powerful instruments of scientific discovery. At the energy frontier, where accelerators continue to unravel the structure of matter and forces that shape our understanding of the Universe, a proposed tera (trillion)-electron volt (TeV) class electron–positron linear collider is expected to require 20 km long accelerators. High-energy electron accelerators have already driven a revolution in materials science and biology by powering intense radiation sources from X-rays to the terahertz (THz), and new machines such as Stanford's Linac Coherent Light Source will use 15 giga (billion)-electron volt (GeV) electrons from a kilometer-scale accelerator oscillating in magnetic undulators to generate unprecedented X-ray brightness. Other applications include medical radiotherapy and imaging, and gamma beams to probe cargos for concealed nuclear material.

Accelerators developed during the past half-century use metallic cavities that shape radio frequency electromagnetic waves to produce accelerating fields. The electrical breakdown of these cavities limits the maximum accelerating field, which dictates the machine's size for a given energy. In turn, size is a major factor for both cost and location. To scale beyond TeV energies and to provide brighter and smaller (laboratory- and hospital-scale) radiation sources, accelerator scientists are developing machines that greatly increase the accelerating fields, and hence the energy achieved in a given length. At the same time, new accelerators must be efficient and produce precise beams that are capable of being focused to very small spot sizes with energy spreads at the 0.1% level.

Laser-driven plasma waves are a promising path for smaller accelerators. The accelerating structure is formed when the radiation pressure of an intense laser pulse displaces the electrons in an ionized gas or plasma, leaving the heavier ions stationary. This charge separation creates an electric field. The field in turn pulls the electrons back after the laser passes, forming a plasma-density wave (or wake) (figure 1, p14), and the field of this wake can be used to accelerate particles. The effect is analogous to a moving boat; water displaced by the boat rushes back after its passage and forms a wake.

Self-consistent simulations are important for plasma accelerators. As the laser drives the wake, the laser pulse is shaped by its interaction with the plasma, and the formation of a stable accelerating structure relies on balancing this process. Moreover, for efficient accelerators, a large part of the laser energy is depleted into the plasma, and the laser and plasma evolve together as this occurs. The wake is typically somewhat nonlinear, and when a particle bunch is accelerated, it further acts back on the wake. Simulations provide new information—such as nonlinear plasma response, beam trapping, and self-consistent laser propagation and beam acceleration—to understand and improve wakefield accelerators.

Recent experiments at the Lasers, Optical Accelerator Systems Integrated Studies (LOASIS) program at Lawrence Berkeley National Laboratory (LBNL), led by Wim Leemans, have demonstrated for the first time high-quality electron

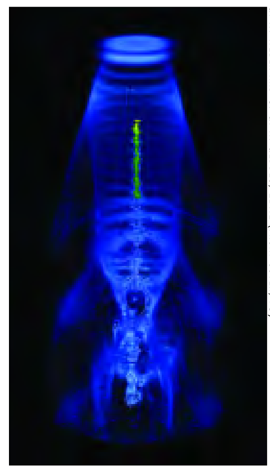

G. H. Weber and C. G. R. Geddes, LBNL (front cover; p3)

Laser-driven plasma waves are a promising path for smaller accelerators.

The wake travels with the laser pulse—close to the speed of light in low-density plasmas—allowing particles to be accelerated by the wave over long distances and to high energies.

C. G. R. Geddes, LBNL (left); Sean Fulton Photography (right)

Figure 1. *In laser plasma accelerators (left), the radiation pressure of a laser pulse (red, moving to the right) displaces plasma electrons creating a density (purple-blue) wave whose electric fields accelerate particles (green-yellow by energy). Similarly, oscillation of water behind a boat (right) creates a wake that can move a surfer.*

beams at 0.1–1.0 GeV using plasmas millimeters to centimeters long. Simulation is challenging for two reasons: (1) the disparity in scale between the micron-scale laser wavelength and the acceleration length, and (2) the large number of plasma particles (sidebar "Simulating Plasma Accelerators" p15). These factors can require million-CPU-processor-hour simulations that use tens of thousands of processors.

Physicists at LOASIS collaborate with colleagues at Tech-X Corporation, who develop the massively parallel VORPAL computational framework, to simulate the experiments. The team received a DOE INCITE grant in 2006 and collaborates with other groups, including a related effort centered at the University of California–Los Angeles (UCLA), through the SciDAC Community Petascale Project for Accelerator Science and Simulation (ComPASS).

Because the simulations produce terabytes of output, three-dimensional visualization is conducted through collaboration with researchers from the SciDAC Visualization and Analytics Center for Enabling Technologies (VACET) and the National Energy Research Scientific Computing (NERSC) Center Analytics Team, using the VizSchema data access plug-in from Tech-X.

The codes are validated via comparison with experimental data, verified against other simulation models, and compared to theory. This rigor provides a physical understanding of the accelerator and the design of new experiments, while shortcomings identify the need for new or improved algorithms and numerical techniques.

Simulations, experiments, and theory show the accelerating field created by the plasma wake is limited by the trapping of particles from the plasma. This limit can be thousands of times that of a conventional accelerator. The wake travels with the laser pulse—close to the speed of light in low-density plasmas—allowing particles to be accelerated by the wave over long distances and to high energies. The wake period, driven by the electron attraction back to the ions, is typically tens of microns. This compact structure naturally produces electron bunches only microns long, meaning femtoseconds (one quadrillionth of a second) in duration, with kiloampere peak currents. Such ultra-short beams, which otherwise require complicated beamline manipulation, are well-suited to drive radiation sources.

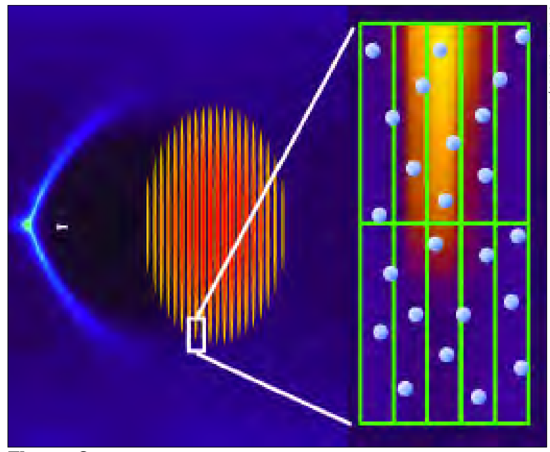

C. G. R. Geddes, LBNL

Figure 2. *Particle-in-cell simulations of a laser wakefield accelerator: the inset shows an expansion of the computational grid, which resolves the laser period (yellow) and plasma particles (blue).*

Meeting high-energy physics and radiation-source requirements for beam quality using laser plasma accelerators will require improved control. For example, one parameter is beam emittance, or roughly the product of beam divergence and size, which governs how well the beam can be focused. Beam momentum, momentum spread, as well as stability are also important. Recent simulations are increasing our under-

Simulating Plasma Accelerators

Modeling laser plasma accelerators challenges state-of-the-art computing by requiring continued parallel scaling, and also collaboration between accelerator and computational scientists to increase speed and accuracy by developing and verifying new models. Explicit particle-in-cell (PIC) simulations are the traditional tool, solving Maxwell's equations by finite difference, resolving the laser period in time step and on a grid in space (figure 2). They self-consistently include fields and interactions of the laser, plasma, and bunch. Simulations of present experiments use up to hundreds of millions of cells and particles and a million time steps, requiring hundreds of thousands of processor hours and effectively using more than 10,000 processors. Even this relatively direct approach is an approximation—the real plasma contains hundreds of trillions of particles and fields with wavelengths unresolved by the simulations (for example, X-ray radiation). The grid itself can add numerical dispersion affecting the laser pulse velocity and can cause momentum errors (sidebar "Algorithms Resolve High-Quality Bunches" p17).

Meter-scale 10 GeV experiments will be at least three to four orders of magnitude more costly to simulate than present experiments, because of increased accelerator length and width, stricter requirements to resolve laser pulse velocity at low density, and the resolution of beam quality for applications. Continued scaling, which appears realistic given recent petascale results on Roadrunner, and also new models will be required. In addition to scaled simulation, approaches to increasing speed include modeling the laser's envelope but not the fast optical oscillation (figure 3), and relativistically shifting the calculation frame to reduce the disparity between scales. In some cases these reduced models allow simulation with 100- to 10,000-fold savings in CPU time. Each also makes additional approximations, such as a relatively fixed laser frequency or neglect of backscatter, which require validation and are not always valid. Solvers that are dispersion-free in vacuum along the grid axes are also being explored to better model laser propagation, which can reduce required resolution and run time. Combinations of these techniques are used to model the accelerators in order to use the strengths of each, to check the models against one another for accuracy, and to derive physical understanding from the differences each approximation reveals.

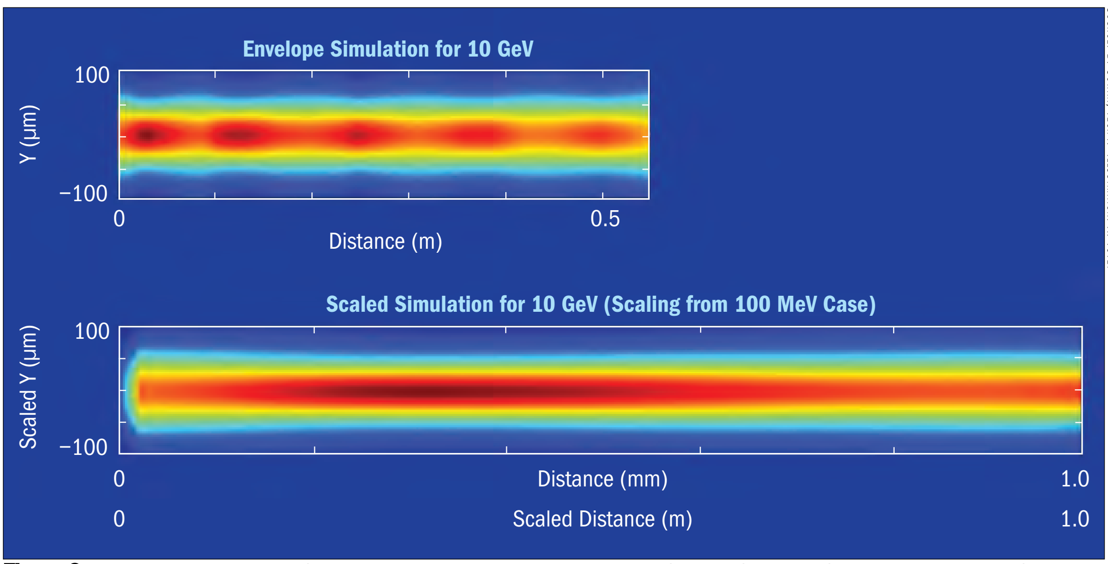

SOURCE: B. COWAN, TECH-X ILLUSTRATION: A. TOVEY

Figure 3. *Envelope model (top) of laser propagation showing laser spot profile as a function of distance along part of a meter-scale 10 GeV laser plasma accelerator, parameters inaccessible to standard PIC simulations. Scaling of smaller standard PIC simulations (bottom) can evaluate propagation further into laser depletion, where the laser frequency shifts and broadens, but the laser spot oscillation period is not correctly modelled, motivating use and development of multiple methods.*

standing of topics such as the formation of narrow energy spread bunches, improving bunch quality by controlling the injection of particles into the wake, and the preservation of these qualities by controlling the accelerator structure. Development and staging of 10 GeV accelerator modules to achieve high energies, efficiently transfer laser energy into the particle beam, and positron acceleration are also being studied.

Controlling Beam Trapping and Acceleration: The First High-Quality Beams

Early experiments in the 1990s demonstrated that laser wakefield accelerators achieved high accelerating gradients and could trap and accelerate electrons out of plasma. However, the electrons were spread out in energy, which made them unsuitable for applications. Simulations published in 2002 by Pukhov et al. suggested that

Experiments at LOASIS produced high-quality beams for the first time in laser accelerators.

G. H. Weber and C. G. R. Geddes, LBNL

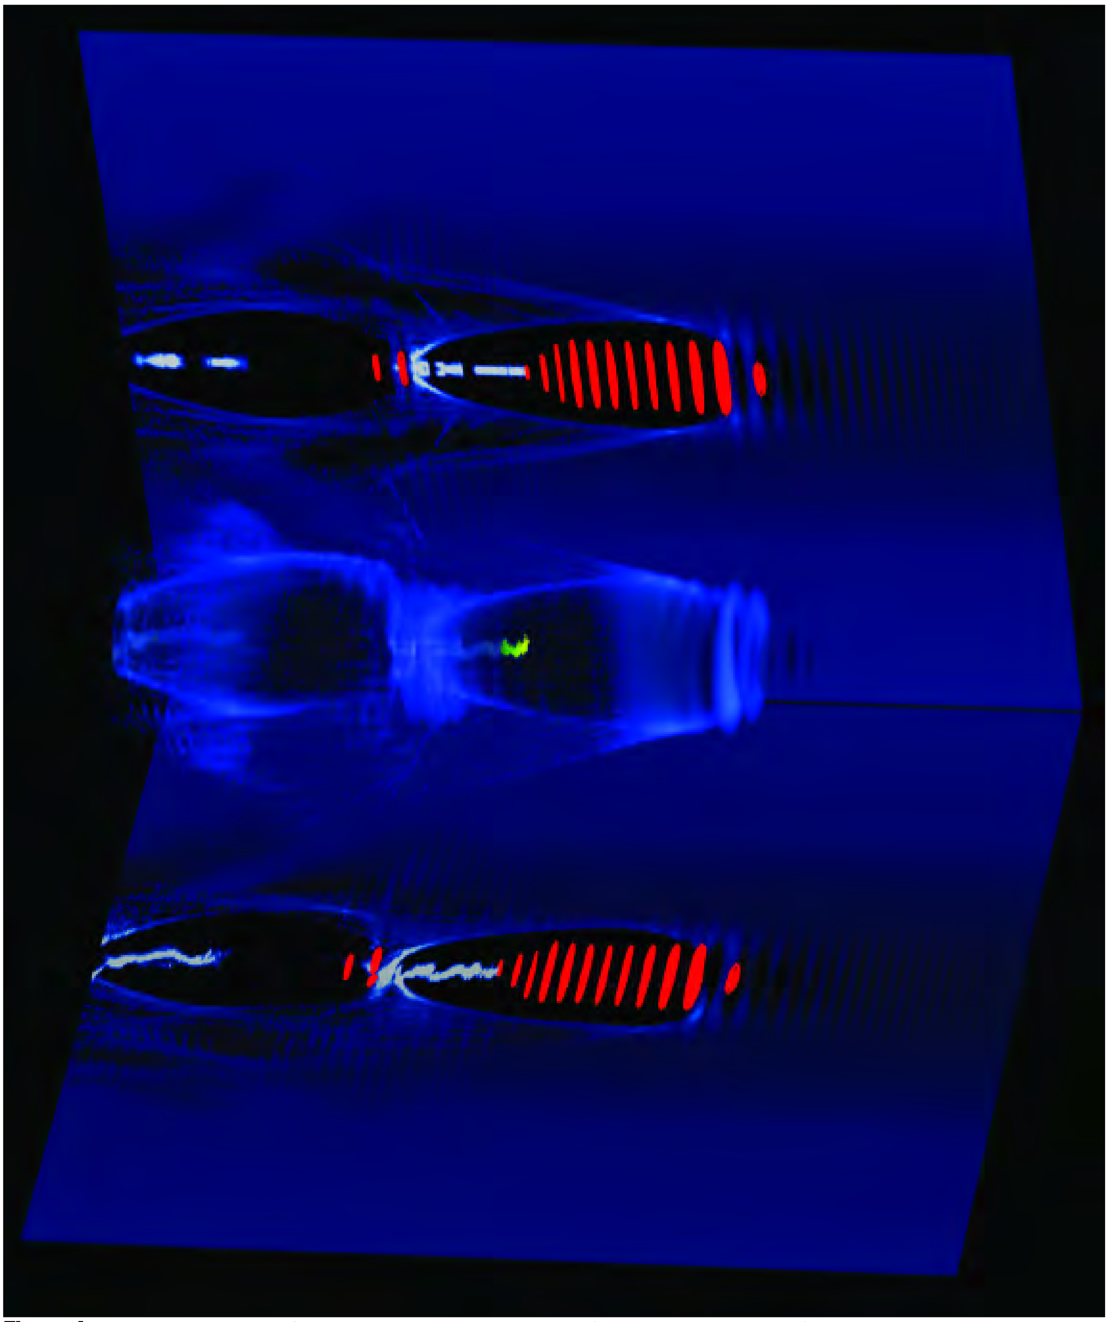

Figure 4. *Three-dimensional VORPAL simulations model the self-consistent evolution of laser pulse and wake, and acceleration of particles. Shown in volume rendering and cut-aways are the laser pulse (red, moving from left to right), wake (blue) and particle bunch (green low energy, yellow high energy).*

strongly nonlinear wakes could produce quasi-monoenergetic electrons by using lasers an order of magnitude beyond what was available.

Experiments at LOASIS produced high-quality beams for the first time in laser accelerators. A VORPAL simulation result appeared on the cover of *Nature* on September 30, 2004, and VORPAL simulations were used to understand the physics behind the experiments reported by Geddes et al. in that issue. Laser propagation was controlled by shaping the plasma density with a coaxial precursor laser that made a low-density "channel" along the laser axis. The 10 terawatt (TW) drive laser was guided like an optical fiber and extended propagation ten-fold. Bunches with 4% energy spread and 3 mrad divergence containing 300 pC of charge at 86 MeV were observed using 2 mm plasmas, and began to address the needs for applications. The same issue of *Nature* published similar efforts by Laboratoire d'Optique Appliquée in France and Imperial College London on the use of large laser spots to extend propagation, and quasi-monoenergetic bunches were predicted by UCLA simulations using the OSIRIS code.

Algorithms Resolve High-Quality Bunches

Accelerators must produce bunches with very low momentum spreads both longitudinally and transversely, making high kinetic accuracy vital for useful simulation. Particles oscillate many times in strong laser and wake focusing fields, and small errors can accumulate and contaminate results. Numerical studies by Cormier-Michel et al. showed errors arise from discretization (of grid and particles) and from interpolation of forces from the grid. High-order spline interpolation for forces and currents, and smoothing of the current, reduced momentum errors by approximately two orders of magnitude, allowing more accurate modeling of beam momentum spread and plasma temperature at orders of magnitude lower cost than using increased resolution. New experiments will require even higher accuracy and indicate the importance of continued model improvement. This development will include new particle pushers, fluid models of the wake to reduce noise, and alternative techniques for modeling the beam, possibly including mesh refinement near the particle bunch. Modeling of scattering and radiation processes will be important in some cases, which will challenge not only model accuracy but also particle statistics of the codes.

Two-dimensional simulations (which simulate a slab, or a plane of space) available at the time of the experiments showed the essential physics behind the new regime. However, the simulated energy, energy spread, and charge did not match the experiments. By driving the wake to amplitude sufficient to accelerate electrons to the wake's velocity in one wave period, electrons were trapped and accelerated from the plasma—so called "self-trapping." The observed isolated bunches were formed when these initial trapped electrons damped (or beam loaded) the wake, suppressing further trapping. The electrons then outran the accelerating part of the wake, which occurred because the laser, and thus wake speed, is less than the speed of light in plasma. This "dephasing" concentrated the electrons in energy, because the lead (highest energy) electrons started to decelerate while the tail continued accelerating. Channeling of the laser extended the accelerating distance to match this dephasing, producing the high-quality beams.

Full three-dimensional simulations on Seaborg at NERSC, made possible by an INCITE grant, and ongoing work on Franklin at NERSC and Atlas at Lawrence Livermore National Laboratory (and similar simulations by UCLA) showed important differences in wake and trapping structure. The wake is highly nonlinear, with a region almost evacuated of plasma and a thin high-density sheath (figure 4). The simulations came close to the experimental beam energy and charge but initially did not match observed energy spread or divergence. These simulations stretched 2006-era computing resources, requiring approximately 300,000 Seaborg processor hours for resolution of the laser wavelength over the wake volume (200 million cells and a billion particles) and over the experimental length of 2 mm (100,000 simulation steps). Many smaller two-dimensional simulations and fundamental numerical studies (sidebar "Algorithms Resolve High-Quality Bunches") were used to test numerical methods and parameters to increase simulation accuracy.

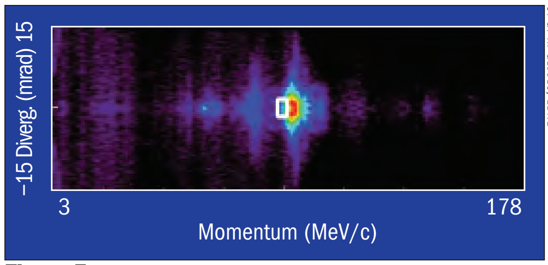

C. G. R. Geddes, LBNL

Figure 5. *A simulated electron spectrum from VORPAL shows a narrow divergence, narrow energy spread bunch, similar to the white overlain circle that shows the experimental beam envelope.*

Three-dimensional simulations that incorporate the developed methods are approaching experimental divergence and energy spread (figure 5), though still fractionally higher. Discrepancies are reduced with increased resolution, and results are best with increased resolution in all spatial dimensions and time simultaneously, which increases cost as resolution to the fourth power. Because the codes scale efficiently as the number of cells increases, design of laser accelerators will therefore continue to benefit from increased computing capabilities. The remaining disagreements are being resolved by adding diagnostics (agreement has already improved with the experimental transmitted laser spectrum and spot), more precisely measuring experimental parameters and using these improved inputs for simulations, and further algorithm development and scaling.

With results quantitatively approaching experiments, simulations allow the evaluation of optimization by allowing controlled parameter variation and provide detailed diagnostics of internal dynamics not accessible in experiments. Tracking particles through the simulations shows they were injected from the sides through the strong fields of the nonlinear wake, which increases their transverse momentum and

Full three-dimensional simulations on Seaborg at NERSC, made possible by an INCITE grant, and ongoing work on Franklin at NERSC and Atlas at Lawrence Livermore National Laboratory (and similar simulations by UCLA) showed important differences in wake and trapping structure.

Visualizing Particle Trajectories

Terabyte output from three-dimensional simulations motivates researchers to use parallel tools to extract physics information. Accelerator scientists collaborate with developers of VisIt parallel visualization tools to provide three-dimensional images for analysis of the wake and laser (figure 4, p16). Particle bunches must be identified and the particles must be tracked through the simulation to understand where they are injected and where accelerator properties can be modified to increase beam quality (figure 6). Clustering and machine learning techniques are being used to automatically identify beam formation and quality in the accelerator. The particles are then tracked using the FastBit indexing tool, which is under development to allow real-time exploration of particle trajectories and associated fields. These projects are in collaboration with the SciDAC VACET center, the NERSC and LBNL visualization and analytics groups, and the Tech-X VizSchema data access team.

D. Ushizima (left) and O. Rübel (right), LBNL

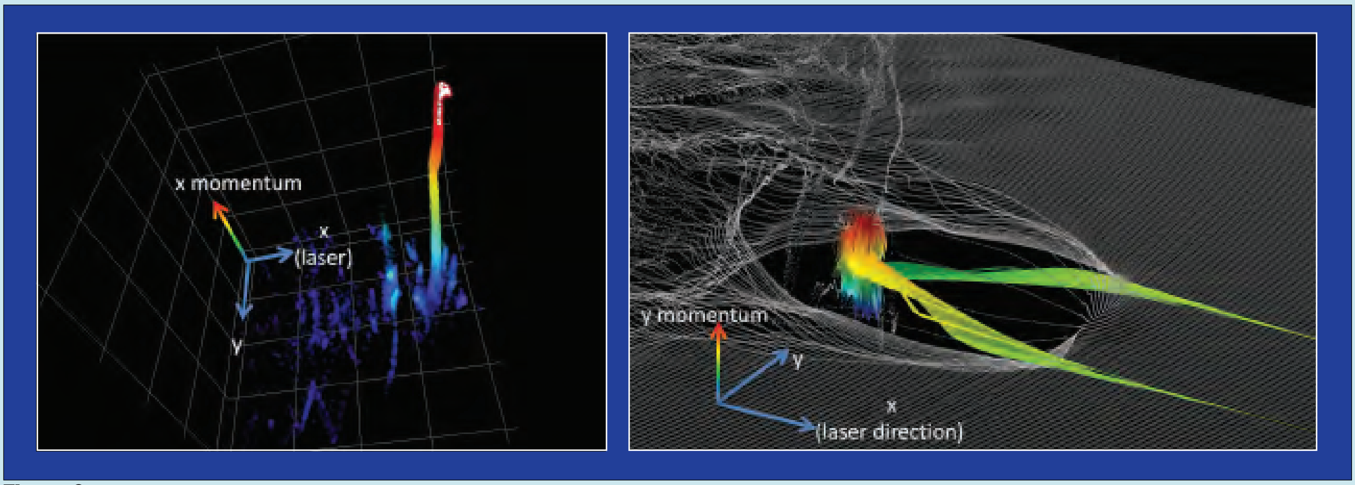

Figure 6. *Formation of high-quality particle beams is automatically identified (white dots, left) by clustering techniques . Tracks through the simulation of accelerated particles, obtained by FastBit, are overlain on a snapshot of the wake particles (gray) showing injection from the sides of the wake and oscillation.*

also, undesirably, the beam divergence (sidebar "Visualizing Particle Trajectories"). Varying laser amplitude in the simulations showed that increasing laser amplitude (or plasma density) increases charge as well as accelerator stability to small fluctuations in laser power, because the trapping process is no longer at threshold. However, the energy spread increases, because increased beam loading is required to turn off trapping, which results in a bunch that covers a larger part of the wake. Similar behavior has now been observed in experiments, and such results are motivating next-generation experiments to control injection of electrons to further stabilize and improve beam quality.

Energies in Centimeter-Scale Plasmas

Electron beams at GeV energies are used to drive X-ray light sources, and are a stepping-stone to yet higher energies. By reducing density, the energy limit is raised when the electrons outrun the wake; both the laser (and hence wake) velocity increases, and also the wake period lengthens. Beams at 1 GeV were produced at LOASIS, in experiments published by Leemans et al. in 2006, using a 40 TW laser to drive a longer (few centimeters), lower-density plasma channel created by an electrical discharge inside a capillary tube.

Simulations with parameters close to the GeV experiment show that the internal dynamics of laser-pulse evolution, trapping, and dephasing-controlled beam formation are similar to the 2004 experiments. The simulated electron beam (figure 7) can closely reproduce the experimental result, with bunch energy 1 GeV with 4% rms energy spread, 2.4 mrad divergence, and approximately 25–60 pC charge. Experimental energy spread was 2.4% and beam divergence 1.6 mrad with 30 pC charge. Use of additional diagnostics such as laser mode, spectrum, and depletion are in progress to further constrain the simulations.

The scaling of beam energy from 0.1 GeV to 1.0 GeV in experiments and simulations verified the understanding of scaling in laser accelerators with plasma density and laser power, and indicated even higher energies are achievable. Simulations also showed higher energies and reduced divergence could be obtained by using controlled injection. Scaling to very high energies is also possible. Related experiments used the high energy available in the SLAC electron beam in place of a laser to drive the wake and demonstrated energy gains of 40 GeV, though with broad energy spread. Laser development is progressing rapidly: commercial systems are now available at petawatt (PW) powers and shrinking in size and electrical-power requirements, which is laying the groundwork for future laser-driven accelerators.

Injection and Staging for Improved Beams

Both light sources and high-energy physics require bunch-momentum spread well below present experiments in both the longitudinal and transverse (divergence) directions. They also require day-to-day accelerator stability. Recent experiments and simulations show that using gradients in the plasma density to control trapping of electrons can help address these goals. Experiments focused 10 TW laser pulses at the edge of a thin gas jet, so the laser moves through a decreasing density "ramp." Bunches were produced with 0.17 MeV/c longitudinal and 0.02 MeV/c transverse momentum spread, an order of magnitude below previous experiments, and with only 3% momentum fluctuation over several days of operation. Although the beam energy was less than 1 MeV, simulations showed these beams can be post-accelerated to reduce momentum spread at high energies. This arrangement parallels the approach in conventional accelerators, where a low-momentum spread/low-energy injector produces beams that are then accelerated to high energies in many accelerator stages.

As the laser travels in the decreasing density plasma, the plasma wake wavelength increases. Simulations showed this makes the wake peaks slip further behind the laser, decreasing wake velocity and thus the threshold amplitude for trapping and accelerating plasma electrons. This low-wake velocity was responsible for the observed low energies, because it caused the bunch to quickly outrun the wake structure. The low-wake amplitude at trapping also reduced bunch-momentum spread, which is consistent with experimental observations. This reduction gave the bunches emittance and momentum spread an order of magnitude better than self-trapped wakefield accelerators. Consistent with observed stability, simulations showed stability over variation in laser power, plasma density, and plasma length. Simulations that focused through the jet required a large simulation domain, and therefore simulations to date were performed in two dimensions. This limitation may explain differences from the experimental energy. Three-dimensional simulations that use the envelope model in VORPAL are now in progress (sidebar "Simulating Plasma Accelerators" p15).

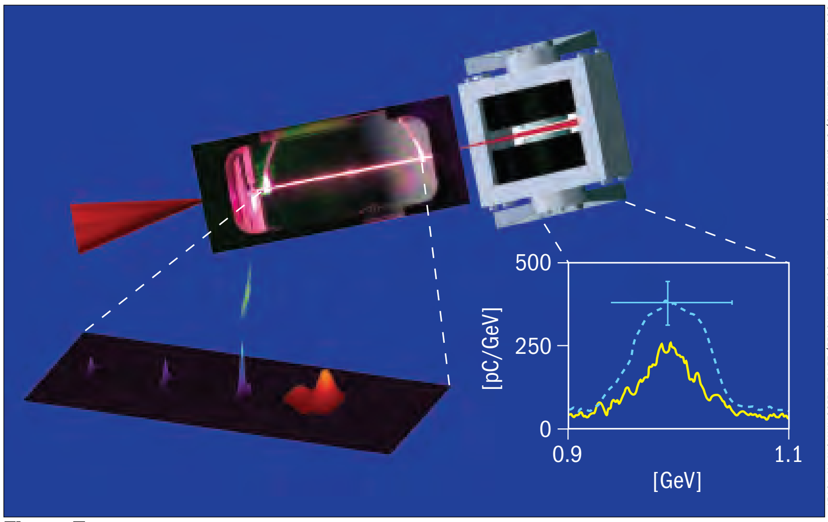

Source: W.P. Leemans, C. G. R. Geddes, and E. Cormier-Michel, LBNL Illustration: A. Tovey

Figure 7. *GeV laser accelerator experiments (top) use a 40 TW laser driving a 3 cm capillary plasma and observe electrons using a magnetic spectrometer. The simulation shows the plasma response to the laser (left). The simulated spectrum (right, solid, A.U.) can closely reproduce the experimental electron spectrum (dotted, with error bars). Charge was 30 pC in experiments, and 25–60 in simulations, also in reasonable agreement.*

Stable performance allowed extended comparison between simulations and experiments. To allow post-acceleration in a second wakefield accelerator, the bunches must be shorter than the wave period. The simulated bunch length was therefore benchmarked to THz experimental measurements, showing that the simulated length was accurate, which in turn showed the bunches were short enough for post-acceleration. The laser pulse was also compared, and showed transmission without severe modulation or depletion. These added diagnostics increased the detail and reliability of simulations evaluating injection and post-acceleration.

Simulations (figure 8, p20) indicate the down-ramp-produced bunches can be used as an injector to improve high-energy accelerators by post-accelerating them in a plasma-channel-guided wakefield accelerator such as those in the 0.1–1.0 GeV experiments. The channel density would be set to avoid additional trapping in the channel. Because the bunch is short compared to the plasma wave, it sees a nearly even accelerating field. The momentum spread is nearly preserved as it accelerates in the channel, producing 0.2 MeV/c class momentum spread at high energy. Energies over 20 MeV have so far been demonstrated but are limited by the computational time with the large domain required. Related simulations with fluid codes by Dr. Bradley Shadwick and co-workers at the University of Nebraska and LBNL indicate that further acceleration should enable bunches at GeV energies and beyond, with a less than 0.1% energy spread. Alternative-controlled injection techniques include using colliding laser pulses to inject particles, which has also

Consistent with observed stability, simulations showed stability over variation in laser power, plasma density, and plasma length.

Source: C. G. R. Geddes, LBNL Illustration: A. Tovey

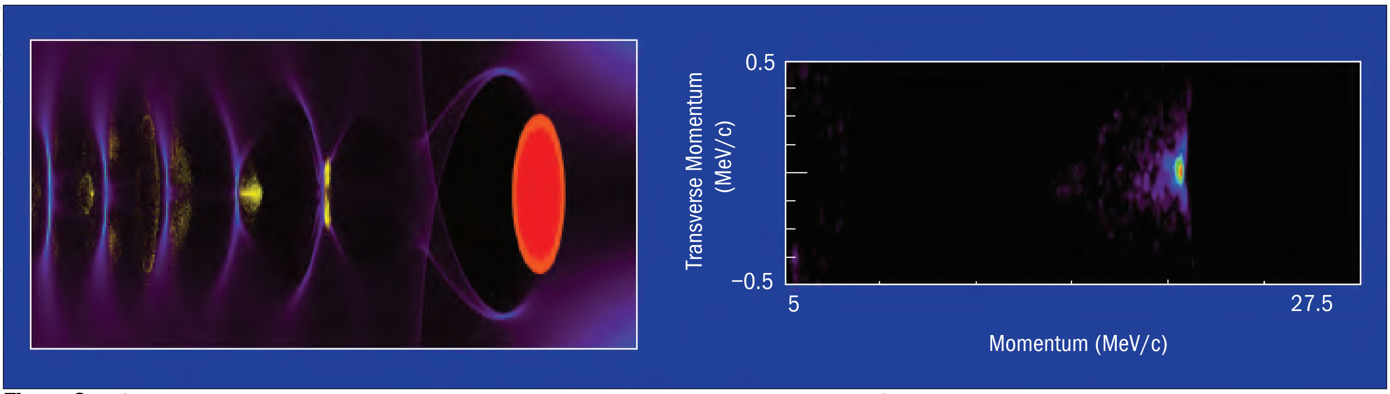

Figure 8. *VORPAL simulations merged density gradient injection to an accelerating channel (left) showing the laser (red), plasma wake density (purple-blue) and accelerating particles (yellow). Acceleration in the channel increases the bunch energy to 20 MeV/c while preserving its low momentum spread (right).*

Source: W. P. Leemans and C. G. R. Geddes, LBNL Illustration: A. Tovey

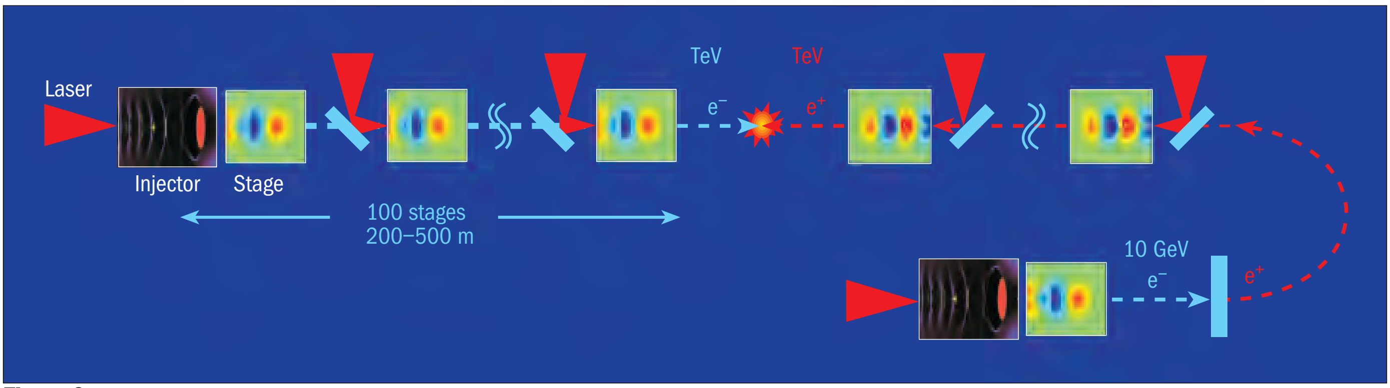

Figure 9. *Conceptual design for a laser plasma linear collider. Drive lasers can be turned at right angles using plasma mirrors, facilitating close coupling between stages to keep overall length short.*

been successful, and work continues to improve the quality of the particle beam by refining injection control using each of these methods.

Designing Next-Generation Efficient 10 GeV Stages

Proposed next-generation experiments will use controlled injection coupled with meter-scale plasmas to increase bunch quality and energy to 10 GeV. A collider could use many modules in series, each increasing the beam energy to reach TeV energies (figure 9). Tradeoffs of accelerating field, which decreases with decreasing density/increasing stage energy with the distance required to couple the laser to drive each stage, indicate that 10 GeV modules may be appropriate for collider applications. For both electrons and positrons, the stages must preserve low-energy spread and good emittance, which motivates the study of less strongly driven, quasilinear wakes wherein the dynamics are similar for electrons and positrons and the shape of the wake is controlled by the laser profile. This concept differs from the nonlinear regime of present experiments where the cavitated region offers high fields and good focusing for electrons but a small positron focusing and acceleration region.

Simulations together with theory show that meter-scale plasmas at a tenth of the densities of present experiments will be required to achieve 10 GeV in either regime. The low densities also lengthen the plasma wave period, making the accelerator structure longer and wider at the same time accuracy requirements increase to model high-quality beams. With simulation of present experiments pushing state-of-the-art computing, new techniques are required to model these stages.

Scaling the laser-spot size and pulse length with the plasma wake wavelength allows a shorter simulation at high density using traditional methods to deduce the properties of 10 GeV stages. Theory predicts scaling with density but does not predict how the plasma will focus the laser or the shape of the wake. A series of VORPAL simulations characterized wake structure and evolution and, by scanning plasma density in a series of runs, showed that they scale predictably with density allowing scaled design. Thus, theoretical scalings are used and verified, and as the gaps are filled, understanding of the accelerator develops.

A wide range of laser and plasma parameters were simulated, and they established conditions for a stage that efficiently transfers laser energy into

Simulations together with theory show that meter-scale plasmas at a tenth of the densities of present experiments will be required to achieve 10 GeV in either regime.

the particle bunch while maintaining high beam quality. These simulations included variation of laser pulse amplitude, width, and length to establish the parameters that best deplete the laser energy into the accelerating field of the wake while maintaining quasilinear structure. Next, electron beam charge, length, and width were adjusted to obtain efficient acceleration—that is, high transfer of laser energy deposited in the wake to the particles—while maintaining low momentum spread (figure 10). These simulations, published by Cormier-Michel et al. in 2008, showed that low-energy spread 10 GeV bunches of both electrons and positrons can be obtained using a petawatt laser. In this regime, the transverse fields of the wake are also shaped by the laser spot. Simulations are now using this feature, and further tailoring of the laser and plasma, to control beam propagation and further increase achievable performance. Simulations also show that nonlinear 10 GeV stages for electrons are accessible with a similar laser, providing further options (scalings in this regime were published by Lu et al. at UCLA).

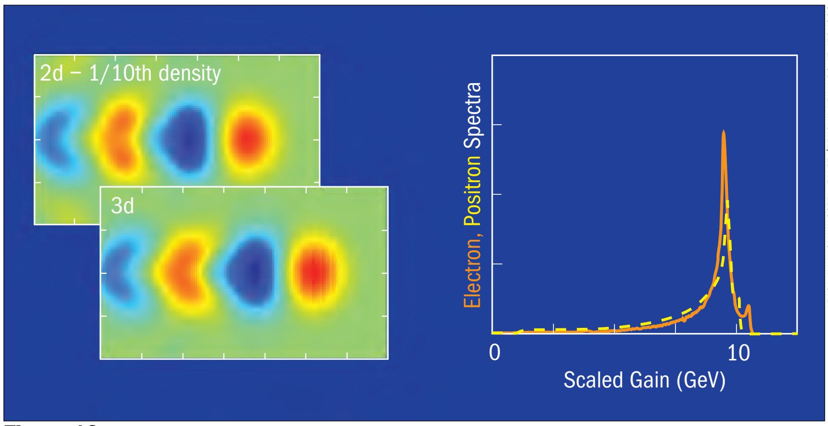

Source: E. Cormier-Michel, LBNL Illustration: A. Tovey

Figure 10. *An efficiently loaded 10 GeV quasilinear stage for the proposed BELLA laser was designed by establishing scaling of the fields (left) between two and three dimensions and versus plasma density using VORPAL simulations. Spectra (right) show electrons (orange) and positrons (yellow) can be accelerated nearly symmetrically achieving high energies with narrow energy spread.*

An advantage of scaled simulation is that laser evolution, up to depletion, can be modeled to evaluate efficiency because the laser period is resolved. However, some parameters, such as the electron bunch oscillation in the focusing field, do not scale. For this reason, and to verify accuracy, scaled simulations are used with envelope, Lorentz shifted, and other techniques (sidebar "Simulating Plasma Accelerators" p15). For collider design and many light source applications, advances are required in all methods to accurately model the very tight beam specifications required. In particular, collider applications require extremely low emittance, very high momentum resolution, and, in some cases, local grid refinement to resolve the small electron bunch.

Conclusions

Large-scale particle simulations provide essential understanding of accelerator physics to advance beam performance and stability of high-gradient, laser plasma particle accelerators. Such simulations demand both massive parallelism and careful model development, because a plasma with many particles must be modeled at fine spatial scale (to resolve the laser and electron beam) and with very high momentum resolution (in order to resolve the required particle beam quality). Computational, visualization, and plasma and accelerator scientists have developed simulations that now provide quantitative understanding of certain parameters and continue to advance the codes' reach. The presented work used the massively parallel VORPAL computational framework, and simulations of the same type are being conducted worldwide using a variety of codes. These simulations revealed important physics behind recent experimental results that demonstrated high beam quality. They are being used to understand and develop controlled injection for increased beam quality, 10 GeV accelerator stages, positron sources and accelerators, and staging of accelerators to reach high energies. These applications raise new code development challenges including improvement and verification of fast reduced models, continued improvement of the kinetic accuracy of the codes, and ability to resolve small-particle bunches. Benchmarking of codes to experiments continues to add new diagnostics, further constraining the simulation results and input parameters, and in turn increasing the physics detail and design utility the simulations offer. New models so developed will be used together with scaling to (and beyond) tens of thousands of processors for design of next-generation accelerators to push the energy frontier of high-energy physics and develop new light sources. ●

Large-scale particle simulations provide essential understanding of accelerator physics to advance beam performance and stability of high-gradient, laser plasma particle accelerators.

Contributors Cameron G. R. Geddes, Estelle Cormier-Michel, Eric H. Esarey, Carl B. Schroeder, Jean-Luc Vay, Wim P. Leemans, and the LOASIS team, all from LOASIS at LBNL; David L. Bruhwiler, John R. Cary, Ben Cowan, Marc Durant, Paul Hamill, Peter Messmer, Paul Mullowney, Chet Nieter, Kevin Paul, Svetlana Shasharina, Seth Veitzer, and the VORPAL development team, all at Tech-X; Gunther Weber, Oliver Rubel, Daniela Ushizima, Prabhat, and Wes Bethel, all at VACET; and John Wu, at the SciDAC Scientific Data Management Center.

Further Reading

VisIT
https://wci.llnl.gov/codes/visit/

VizSchema
https://ice.txcorp.com/trac/vizschema