

Plasma Accelerators at the Energy Frontier and on Tabletops

Charged particles surfing on electron density waves in plasmas can experience enormous accelerating gradients.

Chandrashekhar Joshi and Thomas Katsouleas

Experiments using charged-particle accelerators have led to remarkable discoveries about the nature of fundamental particles and the behavior of nuclear matter. These breakthroughs have been made possible by dramatic advances in our understanding of the physics and technology of particle acceleration.¹ Accelerator beam energies increased exponentially—by an order of magnitude every decade—for half a century after the pioneering days of John Cockcroft and Ernest Lawrence in the early 1930s, as established technologies were pushed to their limits and superseded by new ones (see figure 1).

The present state of the art for proton synchrotrons is the Large Hadron Collider (LHC), under construction at CERN on the French–Swiss border. Its 7-TeV proton beams will, in effect, provide experiments with constituent quark energies on the order of 1 TeV. Looking further ahead, particle physicists are already planning for an electron–positron collider with 250-GeV beams, a neutrino factory, and even a TeV muon collider. These energies represent the so-called energy frontier, where particle physicists confidently expect to discover fundamentally new phenomena.

Although the accomplishments of accelerator-based physics have been spectacular, there is much more to do. The accelerators now in operation or contemplated are expected to lead the way beyond the present manifestly incomplete standard model of particle physics. In particular, they should unearth new classes of particles and enhance our understanding of the asymmetry between matter and antimatter, the masses of the quarks and fundamental leptons, and the transition to the primordial quark–gluon plasma.

We must ask, however, whether the rapid pace of discovery can continue without further breakthroughs in accelerator technology. Irrespective of accelerator topology or the types of particles being accelerated, the fact remains that high-energy accelerators that rely on radio-frequency technology are simply getting too big and too expensive. (See the article by Maury Tigner in *PHYSICS TODAY*, January 2001, page 36.) Is there a different paradigm for building particle accelerators at the energy frontier while dramatically reducing their size—and hopefully their cost?

Chandrashekhar Joshi is a professor of electrical engineering and director of the Center for High Frequency Electronics at the University of California, Los Angeles. **Thomas Katsouleas** is a professor of electrical engineering at the University of Southern California in Los Angeles.

(See the articles in *PHYSICS TODAY* by Andrew Sessler, January 1988, page 26, and by Jonathan Wurtele, July 1994, page 33.)

The impact that much smaller low-energy particle accelerators have had on other branches of science and

technology has been equally impressive. Accelerators are being used, among other applications, for materials science, structural biology, nuclear medicine, fusion research, food sterilization, transmutation of nuclear waste, and cancer therapy. The requirements for such machines are very different from those of high-energy accelerators, and they vary from one application to another. Nonetheless, many low-energy applications would benefit from extremely compact “tabletop accelerators” that could provide beams of GeV electrons, protons, or ions with energies of a few hundred MeV per unit charge.


In this article, we survey new approaches to charged-particle acceleration by collective fields in plasmas. These approaches show considerable promise for realizing plasma accelerators at the energy frontier as well as tabletop electron and ion accelerators. The plasmas would not only provide unprecedentedly high acceleration gradients; they would also serve to focus the accelerated beams down to very small spot sizes.

Surfing on plasma waves

Instead of using RF waves to accelerate charged particles, as conventional accelerators do, plasma accelerators use plasma-oscillation waves excited by lasers or by “driver beams” of charged-particles. The accelerating gradients and focusing strengths that have been demonstrated in plasma experiments have been orders of magnitude greater than those achieved thus far by RF accelerators. The greater the accelerating gradient—if it can be maintained over sufficient distance—the shorter would be the accelerator required to reach a given energy. The impressive plasma-acceleration results already demonstrated in experimental structures raise hopes that this revolutionary technology may miniaturize future accelerators in the same way that semiconductor processors miniaturized electronics.

In addition to the customary “Livingston curve” that charts the progress of working particle physics accelerators over the decades, figure 1 also indicates the highest energies achieved in various plasma acceleration experiments over the past 10 years. Admittedly, plasma schemes have a long way to go before they can produce beams with sufficiently high intensity and low energy spread and “emittance” (that is, angular spread) to be useful for doing high-energy physics. Nevertheless, one can see from figure 1 that the peak energy gain in plasma experiments has been increasing by an order of magnitude every *five* years.

Figure 1. Exponential growth of accelerator beam energy, traditionally depicted by the Livingston curve (green), began tapering off around 1980 as conventional radio-frequency technology approached its limits. The extrapolation to 2010 ends with the 250-GeV electron and positron beams of the 30-km-long Next Linear Collider (NLC), which particle physicists are proposing to build with state-of-the-art RF technology. Progress toward a more radical solution that might eventually reach the high-energy frontier with much shorter linear accelerators is indicated by the red dots, which mark energies already obtained in some of the plasma-acceleration experiments in the US, Japan, and Europe. (Institutional abbreviations are spelled out in the text.) They all represent electron acceleration, except for E162 at SLAC, which is for positrons. The blue dot marks the record electron energy anticipated from a plasma experiment now in progress.



We describe here the ideas and developments that have led to these promising results, as well as the considerable challenges that lie between proof-of-principle experiments and the realization of useful plasma accelerators.

Generating the waves

John Dawson of UCLA has aptly been called the father of plasma-based accelerators. (See his obituary in *PHYSICS TODAY*, July 2002, page 78.) The four plasma acceleration schemes invented by Dawson and coworkers all share a common underlying principle: Charged particles are accelerated by surfing on the longitudinal electric field of a relativistically propagating electron plasma wave. Relativistic propagation means that the wave's phase velocity is very close to c , the speed of light, so that the particle cannot quickly outrun the accelerating electric field of the wave and therefore interacts with it over a long distance.

The electric field intensity in volts per centimeter is approximately given by the square root of the plasma's electron density ρ in cm^{-3} . But the characteristic wavelength of plasma oscillation decreases with increasing density like $\rho^{-1/2}$. Thus, a plasma with $\rho = 10^{16}/\text{cm}^3$ can support a wave that has a peak field of about 10^8 V/cm. That's orders of magnitude more than the 10^5 to 10^6 V/cm accelerating gradient of existing and proposed RF electron linacs. But such high gradients are achieved at the expense of ever shorter wavelengths. For $\rho = 10^{16}/\text{cm}^3$, the plasma wavelength is only $300 \mu\text{m}$.

Figure 2 illustrates the different plasma accelerator schemes for exciting a plasma wave:

► In a *laser wakefield accelerator* (LWFA), as shown in figure 2a, the radiation pressure of a short, intense beam of laser photons pushes plasma electrons forward and aside. As the laser pulse passes, these displaced electrons snap back and overshoot their original positions because of the restoring force exerted by the heavier, less mobile ions. The displacement sets up a plasma density oscillation behind the laser pulse, much like the wake of a ship. In a dilute plasma, the photon pulse propagates with a group velocity

very close to c , and therefore the phase velocity of the wake oscillation is also near c . The wake's electric field—the so-called wakefield—thus accelerates the relativistic particles over a long distance before they finally outrun it. So the particles' energy gain is substantial.²

► In a beam-driven *plasma wakefield accelerator* (PWFA), a high-gradient plasma wakefield is excited by a short, high-charge, relativistic beam of charged particles rather than by laser photons (see figure 2b). Now it is the Coulomb force of the beam's space charge that expels the plasma electrons (or, in the case of a positron driver beam, pulls them in). Again, they rush back and overshoot when the driving beam passes, setting up a plasma oscillation.³ In both the LWFA and PWFA schemes, the optimum length of the drive pulse turns out to be approximately half a wavelength of the plasma wave. For $\rho = 10^{16}/\text{cm}^3$, the optimum drive pulse length is about 500 femtoseconds.

► The *plasma beat-wave accelerator* (PBWA) scheme (figure 2c) was one of two approaches devised to address the difficulty that, until recently, subpicosecond photon and electron bunches powerful enough to drive a high-gradient plasma wake did not exist. The PBWA makes use of a beat wave produced by the interference of two higher-frequency laser beams whose frequency difference is exactly equal to the natural frequency of the plasma oscillation.⁴ The amplitude-modulated beat wave looks like a series of short pulses that can resonantly drive the plasma wave. One can then use lasers of much lower intensity than one needs in the LWFA case.

► The *self-modulated laser wakefield accelerator* (SMLWFA) is the other scheme originally designed to address the absence of sufficiently powerful subpicosecond bunches. It was already known to plasma physicists in the 1970s, but its potential for acceleration was not appreciated until the 1980s. In this scheme, a long, intense pulse modulates itself into something like a beat-wave structure (figure 2d) via the Raman forward-scattering instability, so called because it resembles the Raman scattering of light from molecules. But in this case, it is a parametric instability in

which the scattered light and the plasma wave grow exponentially at the expense of the pumping radiation.⁵ The incident electromagnetic wave couples its energy into a relativistically propagating plasma wave and two electromagnetic daughter waves with frequencies shifted up and down by the plasma frequency.

The interaction between a short, intense laser pulse and a plasma can also be used to accelerate ions.⁶ For instance, at the Rutherford Appleton Laboratory (RAL) in England, a 50-terawatt laser focused on a solid target has generated a collimated beam of 10 to 400 MeV protons and other ions. The laser's electric field first ionizes the target atoms, and the electrons are then accelerated forward by the laser magnetic field's Lorentz force. The space charge of the electrons, in turn, accelerates the ions leaving the target's surface.

The ions can also be accelerated by the shock front that propagates through the solid material after it is hit by the laser pulse. The efficiency of converting laser energy into high-energy ions can exceed 10%. With petawatt lasers now becoming available, ion beams with energies up to a GeV can be expected from such a relatively simple device. Protons and ions in this energy range are well suited for a variety of applications: radiography for nuclear stockpile stewardship, drivers for fast-ignition fusion, generation of medical isotopes, and perhaps even as injectors for conventional accelerators.

Recent experiments with electron acceleration

The first excitation of relativistic plasma waves by the PBWA technique was accomplished in 1985 by the UCLA group led by one of us (Joshi).⁷ In the 1990s, electron acceleration by the PBWA scheme was demonstrated by the UCLA group and by groups at Chalk River Laboratories in Ontario, the Ecole Polytechnique near Paris, and the Institute for Laser Engineering (ILE) in Osaka, Japan. From the observed energy gains of the electrons in those experiments, one could infer acceleration gradients in excess of a gigavolt per meter, confirming the theoretical predictions of extremely high electric fields for relativistic plasma waves.⁸

In the early 1990s, the use of chirped-pulse amplification to produce high-power pulses with Nd:glass lasers led to laser systems that could generate 100-TW pulses. Christine Coverdale and coworkers at Lawrence Livermore National Laboratory (LLNL) showed not only that such a laser focused in a plasma produced Raman forward scattering, but also that the plasma wave accelerated some of the electrons from the plasma itself. And in the mid-1990s, pioneering experiments on the SMLWFA scheme were done at the RAL by a collaboration led by Bucker Dangor of Imperial College, London.⁹

Those experiments demonstrated that, in a plasma of electron density $10^{19}/\text{cm}^3$, one could produce a 5-nanocoulomb beam of electrons with energies up to 100 MeV. Amazingly, that energy gain occurred in less than a millimeter. That's the highest terrestrial acceleration gradient ever achieved—almost 200 GeV/m. Subsequently, groups from the US Naval Research Laboratory, the University of Michigan, Lawrence Berkeley National Laboratory (LBNL), and KEK, Japan's high-energy accelerator laboratory, demonstrated similarly impressive results.¹⁰

With the emergence of Ti:sapphire lasers in the late 1990s, it has become possible to generate 50-fs, 100-TW




Figure 2. Four different kinds of driver pulses (red, propagating from right to left) that can excite a relativistic electron density wave (blue) in a plasma accelerator. **(a)** In a laser wakefield accelerator, the radiation pressure of a short, intense laser pulse pushes plasma electrons away to initiate the accelerating plasma wakefield. **(b)** In a plasma wakefield accelerator, the plasma electrons are pushed away by the Coulomb force of a short, high-charge relativistic bunch of electrons from a linac. **(c)** In a laser-driven plasma beat-wave accelerator, the interference between two high-frequency lasers produces a laser beat wave whose frequency equals the plasma's natural oscillation frequency. **(d)** In a self-modulated laser wakefield accelerator, a long, intense laser pulse interacting with the plasma acquires a beat-wave-like amplitude modulation by the Raman forward-scattering instability.

pulses at a high repetition rate. So now one can, at last, realize the promise of the LWFA scheme as it was originally conceived. Recently, Victor Malka's group at the Ecole Polytechnique's Laboratoire d'Optique Appliquée (LOA) has used wakefield acceleration to obtain extremely promising results: a maximum electron energy of 200 MeV in just a millimeter (see figure 3).¹¹ Having measured the angular divergence of the accelerated electrons at various energies, the group concluded that the highest-energy electrons also have the smallest emittance.

Tabletop prospects

How might one further increase the energy of electrons in laser-driven schemes to produce a GeV tabletop plasma accelerator? To go from 200 MeV in a couple of millimeters to a GeV and beyond, one needs to increase the length of the plasma wave by guiding the intense laser pulse. A particularly promising approach is the use of a preformed plasma channel. The channel would have an electron density minimum along its axis. The refractive index is higher on the axis of the channel than at its edges and therefore the channel acts like an optical fiber to guide light beams.

How to produce such plasma channels for guiding ultra-intense laser pulses while exciting a wakefield is an active area of research that is meeting with some success.¹² For instance, Michael Downer's group at the University of Texas at Austin has propagated 80-fs laser pulses with peak intensities of $2 \times 10^{17} \text{ W/cm}^2$ over 1.5 cm. That distance is 60 times the characteristic diffraction length in such a plasma channel. At such intensities, however, the wake amplitude is not large enough to accelerate electrons from the plasma itself. In any case, such plasma electrons would have too large an energy spread. Most applications require a nearly monoenergetic electron beam.

For monoenergetic acceleration, an extremely short electron bunch must be phase-locked to the accelerating




Figure 3. Tabletop laser wakefield acceleration of electrons to a few hundred MeV. (a) A typical tabletop setup would fire focused ultrashort laser pulses into a laminar-flow gas jet to produce a plasma and excite a wakefield. The wakefield traps plasma electrons and accelerates them to about 100 MeV in just a millimeter. (b) Recent results with 50-fs, 100-TW laser pulses at the Laboratoire d'Optique Appliquée (LOA) in France demonstrate acceleration of electrons to 200 MeV in 1 mm of plasma.¹¹ (c) The angular distribution of the electrons in the LOA experiment, for three energy ranges, shows that the spread of exit angles, and therefore the emittance of the accelerated beam, decreases with increasing electron energy. For each energy range, the blue outer circle indicates the full width at half maximum of the angular distribution for that range. (Courtesy of V. Malka.)

plasma waves, whose wavelengths may be only tens of microns at the requisite ultrahigh accelerating gradients. Fortunately, several all-optical injection schemes have been proposed that use two or more laser beams: one to excite the plasma wake and the other(s) to inject electrons into the wake with femtosecond precision.¹³

In one such scheme, Donald Umstadter and coworkers at the University of Michigan have suggested that a second ultra-intense laser pulse, propagating at right angles to the first, could perturb the orbits of the electrons

in the wakefield excited by the first pulse and thus cause them to be trapped and accelerated while maintaining a relatively small energy spread. Combining plasma guiding with all-optical injection is thought by many to be the most promising route toward GeV electron beams in centimeter-scale plasmas.

Scaling up to the energy frontier

How hard is it to extend these centimeter-scale acceleration techniques to the meter scales needed for approaching the energy frontier? One approach to higher energies would be to stage hundreds of smaller laser-plasma acceleration modules, each providing a gain of a few GeV. The problems associated with staging, however, are enormously complicated and beyond the scope of this article. But if one could make a single, ultrahigh-gradient plasma-acceleration stage several meters long, staging could be greatly simplified—or even rendered unnecessary.

Such issues have been investigated in recent experiments on beam-driven plasma-wakefield acceleration at Argonne National Laboratory, Fermilab, and at SLAC.¹⁴ The work at SLAC was done by a collaboration of scientists from SLAC, UCLA, the University of Southern California, and LBNL. A series of experiments at SLAC (designated E157 to E162) has demonstrated a prototype, 1.4 meters long, of a positron and electron PWF stage.¹⁵ This is a very significant accomplishment, because the particle physics community has assigned its highest priority to a linear electron-positron collider at the energy frontier (see PHYSICS TODAY, September 2001, page 22).

The drivers of the plasma wakefield in the SLAC experiments were 4-ps-long beam bunches of 28.5-GeV positrons and electrons from the laboratory's 3-km linac, with peak currents of nearly a kiloamp. The passage of the intense electron beam through a 1.4-m-long lithium plasma not only excited the plasma wave but also accelerated a significant amount of charge from the back end of the same drive bunch.¹⁵

These experiments also showed that beam propagation over 1.4 meters of plasma is controllable, despite the strong focusing forces that arise in the plasma. (Those forces are comparable to what one would get in a 6000-T/m quadrupole electromagnet.) Aside from energy loss and gain, the experiments have quantitatively verified many of the other predicted phenomena, such as periodic focusing of the beam and x-ray emission due to the wiggling motion of electrons about the axis of the wakefield.

For plasma waves that are not too large, a plasma wakefield accelerates positrons in much the same way as it accelerates electrons—except, of course, that the positrons are carried along on the opposite phase of the plasma wave. For the strongly driven plasma waves of interest for the highest accelerating gradients, however, the physical mechanism for wakefield excitation and acceleration turns out to be different for electrons and positrons. In plasma-wakefield acceleration with an intense electron-beam driver, the Coulomb force of the beam's space charge completely expels the plasma electrons. But the space charge of a positron beam pulls in plasma electrons laterally, thus creating an electron density on axis that far exceeds the initial plasma density. In a uniform plasma, the wakes produced by an otherwise identical positron beam tend to be somewhat smaller than those produced by an electron beam.

Figure 4 compares the energy change of the positron beam observed in the SLAC experiments with the prediction from the self-consistent three-dimensional computer code OSIRIS.¹⁵ The agreement between experiment and simulation is excellent: The bulk of the beam loses about




Figure 4. Observation and simulation of positron energy changes resulting from traversal of a 1.4-m-long plasma. (a) In a plasma wakefield experiment at SLAC,¹⁵ energy changes were recorded for different temporal slices of the 10-ps relativistic positron beam pulse. Whereas the bulk of the pulse (near $t = 0$) showed energy loss of about 55 MeV, its tail (t later than +1 ps) demonstrated energy gain of as much as 80 MeV. This was the first demonstration of high-gradient plasma acceleration of positrons. (b) A computer simulation of the SLAC experiment group predicted much the same pattern of deceleration of the pulse's head followed by acceleration of its tail. (Adapted from B. Blue, et al., *Phys. Rev. Lett.*, in press.)

55 MeV, whereas particles in its tail gain a maximum energy of 80 MeV. These results represent the first high-gradient acceleration of positrons using collective plasma fields. Such meter-scale prototype stages for accelerating both electrons and positrons are something of a milestone in the development of plasma accelerators for the energy frontier.

Plasma focusing of particle beams

If one wants a miniature high-energy collider, one must also miniaturize the final-focusing optics that make the colliding-beam cross sections small enough to provide a useful collision rate. Fortunately, plasmas can produce lenses of extremely short focal length for beams of ultra-relativistic charged particles.¹⁶

The physical mechanism of a plasma lens is easily understood. In vacuum, an ultrarelativistic electron beam




Figure 5. Plasma focusing of a 30-GeV positron beam from the Stanford Linear Accelerator. At a fixed distance beyond the plasma lens, the demagnification of the beam diameter, which was initially $50 \mu\text{m}$, is plotted against the electron density of the focusing plasma. As the plasma density is increased toward its optimum ($2 \times 10^{12}/\text{cm}^3$), the focusing strength increases and the beam diameter is pinched to half its initial value. But at still higher plasma densities, the beam is overfocused and the spot size grows again. The red error bar indicates the diameter uncertainty of the beam without plasma focusing. (Adapted from M. Hogan et al., *Phys. Rev. Lett.*, in press.)

spreads rather slowly as it propagates. That's because the repulsive space-charge force is nearly balanced at relativistic velocities by the attractive Lorentz force of the charged particles moving in parallel. However, as an electron beam enters a plasma and expels the plasma electrons, the ions left behind exert a focusing force. If the beam density exceeds the plasma density, all the plasma electrons leave. Then the focusing force of the remaining ion column is constant along the beam, and it grows linearly with transverse distance from the axis. In addition to its very short focal length, such a lens has, in principle, no spherical aberrations. That greatly simplifies the final-focusing optics before the collision point, which would otherwise require dozens of focusing elements.

The focusing of a positron beam is, alas, less ideal. The positrons pull in plasma electrons, which tends to neutralize the beam's charge. Plasma electrons from different distances converge on the positron beam at different times. The resultant focusing force is then neither linear in the radial direction nor constant along the beam. Nevertheless, the focusing gradients are orders of magnitude greater than those achievable by traditional magnetic focusing.

Plasma lenses have been shown to focus GeV electron and positron beams. In a set of experiments done by Johnny Ng and collaborators at SLAC,¹⁶ the $7\text{-}\mu\text{m}$ -diameter SLAC positron beam was further focused down to about $5 \mu\text{m}$ as it passed through a short column of plasma with an electron density of about $10^{18}/\text{cm}^3$. The focusing gradients in these experiments were enormous: 10^6 T/m . Figure 5 shows the focusing of positrons by a plasma lens in a recent experiment at SLAC by Mark Hogan and coworkers. The positron beam diameter is clearly seen to decrease, at a fixed distance beyond the lens, as the plasma density, and therefore the focusing strength, is increased to an optimum value. Beyond this optimum density, the beam is overfocused and the spot size grows again.

Plasma focusing has demonstrated spot-size reduction by a factor of two to four, for both positron and electron




Figure 6. To test the feasibility of plasma-wakefield afterburners added to a high-energy radio-frequency linear collider as an energy doubler, we propose the installation of two 10-m-long plasma afterburners at the ends of the 50-GeV electron and positron beams of the Stanford Linear Collider (SLC), just before the point where the beams are focused to collide with each other (see inset). That would double the collider's e^+e^- collision energy to 200 GeV. Unlike the plasma filling the electron afterburner, the positron afterburner's plasma would have a hollow axial channel. The SLC shares the downstream end of the 3-km-long Stanford Linear Accelerator with the PEP II asymmetric B-meson factory and other facilities.

beams. The logical next step is to demonstrate that plasmas can focus beams to submicron-diameter spots. For the eventual use of plasma lenses in a high-energy electron-positron linear collider, one would have to focus the colliding beams down to a few tens of nanometers.

Toward a plasma afterburner

The demonstration of plasma-wakefield acceleration of electrons and positrons over a meter, and the plasma focusing of multi-GeV e^+ and e^- beams down to spots only a few microns wide has fueled speculation about the scalability of the PWFA scheme to the energy frontier. The theoretical acceleration gradient in a PWFA scales as the electric charge of the driving bunch divided by the square of bunch length. Thus, if the 4-ps pulse length of the drive bunch in present experiments can be contracted by an order of magnitude, we should be able to increase the acceleration gradient from the present 100 MeV/m to something like 10 GeV/m. Using the ultrashort bunch facility at SLAC, we will soon be testing this crucial scaling law in an experiment labeled E164 (see figure 1). Should it prove possible to obtain such extraordinarily high gradients over a few meters without incurring any beam-plasma instabilities, one could contemplate building a single-stage PWFA energy doubler at the downstream end of a high-energy RF linear accelerator.¹⁷

How would such an “afterburner” work? As an illustrative example, figure 6 shows a pair of PWFA energy doublers, installed on either side of the collision point of the Stanford Linear Collider (SLC), that serve to double the energies of the electron and positron beams from the present 50 GeV to 100 GeV. That's the scale of demonstration experiment one would need to prove that the plasma wakefield accelerator is a realistic candidate for a future e^+e^- collider at the energy frontier. Each plasma afterburner in such an experiment would be about 10 m long. These

plasma sections play the role of a voltage transformer, increasing beam energy at the expense of beam current.

In each plasma section, the particle bunch actually consists of two microbunches just 100 fs apart. The first microbunch excites the wake, while the second, containing only a third as many charged particles as its predecessor, is phased to accelerate on the wake of the first. To avoid sacrificing event rate at the SLC's collision point, one would have to offset the reduced number of particles in the accelerated bunch by having correspondingly smaller beam spot sizes. That would be done by placing a high-density plasma lens after each wakefield afterburner section, just before the collision point.

The positron beam's afterburner would be different from the one that boosts the electrons. As shown in figure 6, it would be a hollow channel on the axis, surrounded by plasma. The wakefield excited by a positron beam in such a hollow channel would be more coherent than one could get in a radially uniform plasma column. That's essentially because the positrons would be pulling in electrons from the cylindrically symmetric boundary of the channel.

Preliminary simulation studies of such plasma afterburners have shown that electrons in the second bunch are accelerated at a rate of 8 GeV/m. The 20% energy spread coming out of these simulations is too big for some classes of physics experiments that require a narrowly defined collision energy. But it might be tolerable for experiments that can rely on a robust production signal for the particle being sought—for example, the Higgs boson.

Important accelerator physics issues remain to be resolved. How stable, for example, would the plasma wakefield be over distances much longer than a meter? How efficiently could the accelerating beam extract energy from the wake? Then there are demanding technological issues like the alignment of the colliding beams. Such questions are common to all advanced accelerator schemes. They will

be addressed in the next few years by new 3D computer simulation tools and new, dedicated experimental facilities such as Orion, under construction at SLAC, and the Accelerator Test Facility, already in operation at Brookhaven National Laboratory.¹⁸

The quest for plasma acceleration is making rapid progress (see figure 1). It took accelerator builders 60 years to reach lepton (and quark) center-of-mass collision energies in the range of 200 GeV. If plasma-acceleration technology is to have an impact at the high-energy frontier, it had better make faster progress than that. We must address serious issues such as beam quality, event rates, and overall efficiency.

On the other hand, combining laser and plasma technologies could lead rather quickly to GeV tabletop electron and ion accelerators for a rich variety of applications. The future is, of course, full of challenges and uncertainties, but it is also full of exciting chances to make a difference.

We dedicate this article to the memory of John Dawson. We also thank our many collaborators and colleagues who have contributed to the field that Dawson pioneered.

References

1. See, for example, A. Chao et al. 2001 *Snowmass Accelerator R&D Report*, available at <http://www.hep.anl.gov/pvs/dpb/Snowmass.pdf>.
2. T. Tajima, J. M. Dawson, *Phys. Rev. Lett.* **43**, 267 (1979); P. Sprangle et al., *Phys. Rev. Lett.* **72**, 2887 (1994).
3. P. Chen et al., *Phys. Rev. Lett.* **54**, 693 (1985); J. Rosenzweig et al., *Phys. Rev. A* **44**, R6189 (1993).
4. C. Joshi et al., *Nature* **311**, 525 (1984).
5. D. Forslund et al., *Phys. Fluids* **18**, 1002 (1975); K. Estabrook, W. Kruer, *Phys. Fluids* **26**, 1892 (1983).
6. K. Krushelnick et al., *IEEE Trans. Plasma Sci.* **28**, 1185 (2000).
7. C. E. Clayton et al., *Phys. Rev. Lett.* **54**, 2343 (1985); M. Everett et al., *Nature* **368**, 527 (1994).
8. Y. Kitagawa et al., *Phys. Rev. Lett.* **68**, 48 (1992); C. Clayton et al., *Phys. Rev. Lett.* **70**, 37 (1993); N. Ebrahim, *J. Appl. Phys.* **76**, 7645 (1994); F. Amiranoff et al., *Phys. Rev. Lett.* **74**, 5220 (1995).
9. C. Coverdale et al., *Phys. Rev. Lett.* **74**, 4659 (1995); A. Modena et al., *Nature* **377**, 606 (1995); D. Gordon et al., *Phys. Rev. Lett.* **80**, 2133 (1998).
10. K. Nakajima et al., *Phys. Rev. Lett.* **74**, 4428 (1995); D. Umstadter et al., *Science* **273**, 472 (1996); A. Ting et al., *Phys. Rev. Lett.* **77**, 5377 (1996); W. Leemans et al., *Phys. Rev. Lett.* **89**, 174802 (2002).
11. V. Malka et al., *Science* **298**, 1596 (2002).
12. G. Durfee, J. Lynch, H. Milchberg, *Phys. Rev. E* **51**, 2368 (1995); Y. Ehrlich et al., *Phys. Rev. Lett.* **77**, 4186 (1996); E. Gaul, S. LeBlanc, A. Rundquist, R. Zgadzaj, H. Langhoff, M. Downer, *Appl. Phys. Lett.* **77**, 4112 (2000); P. Volfbeyn, E. Esarey, W. Leemans, *Phys. Plasmas* **6**, 2289 (1999).
13. D. Umstadter et al., *Phys. Rev. Lett.* **76**, 2073 (1996); C. Moore et al., *Phys. Rev. Lett.* **82**, 1688 (1999); E. Esarey et al., *Phys. Rev. Lett.* **79**, 2682 (1997).
14. N. Barov et al., *Phys. Rev. Lett.* **80**, 81 (1998); J. B. Rosenzweig et al., *Phys. Rev. Lett.* **61**, 98 (1988); see also <http://www.slac.stanford.edu/grp/arb/e162/>.
15. M. Hogan et al., in *Advanced Accelerator Concepts*, 10th Workshop, C. Clayton and P. Muggli, eds., AIP conf. proc. 647, AIP, Melville, N.Y. (2002); C. Joshi et al., *Phys. Plasmas* **9**, 1845 (2002); B. Blue et al. *Phys. Rev. Lett.* (in press); C. Clayton et al., *Phys. Rev. Lett.* **88**, 154801 (2002); S. Wang et al., *Phys. Rev. Lett.* **88**, 135004 (2002).
16. G. Hairapetian et al., *Phys. Rev. Lett.* **72**, 2403 (1995); R. Govil et al., *Phys. Rev. Lett.* **83**, 3202 (1999); P. Chen et al., *Phys. Rev. Lett.* **64**, 1231 (1990); J. Ng et al., *Phys. Rev. Lett.* **87**, 244801 (2001); M. Hogan et al., *Phys. Rev. Lett.* (in press).
17. S. Lee et al., *Phys. Rev. Special Topics B* **5**, 011001 (2002).
18. See, for example, <http://www-project.slac.stanford.edu/orion> and <http://www.bnl.gov/atf>.

MAGNETS

COMPLETE DESIGN FACILITIES • NEODYMIUM IRON BORON
SAMARIUM COBALT • CERAMIC • ALNICO • MOLDED/BONDED RARE EARTHS
ELECTROMAGNETS • MAGNETIC ASSEMBLIES



With a state of the art manufacturing facility which is certified to ISO 9001:2000 we can deliver a quality magnet, assembly or sub assembly **fast**. MCE can also fully engineer and design a solution for your magnet requirement. Call or FAX us with your requirement for an **immediate** quotation.



MAGNETIC COMPONENT ENGINEERING, INC.

2830 Lomita Blvd. • Torrance, CA 90505

Toll Free: (800) 989-5656

Main: (310) 784-3100 • Fax: (310) 784-3192

Email: sales@mceproducts.com • Website: www.mceproducts.com

Circle number 29 on Reader Service Card

www.janis.com

JANIS

3He - 4He DILUTION REFRIGERATORS

Janis Research offers the JDR Dilution Refrigerator line with:

- Base temperature of 7 mK - 15 mK
- Cooling power of 100 μ W - 400 μ W at 100 mK
- Sample in vacuum
- Integrated safety features
- Resistance and CMN calibrated thermometers.



Janis Research Company

2 Jewel Drive • Wilmington, MA 01887 USA

TEL 1 (978) 657-8750 FAX 1 (978) 658-0349 sales@janis.com

Circle number 30 on Reader Service Card