Plasma Wakefield Experiments

N. Barov*, M.C. Thompson†, K. Bishofberger†, J.B. Rosenzweig†,
H. Edwards** and J. Santucci**

*Dept. of Physics, Northern Illinois University, DeKalb, IL 60115
†UCLA Dept. of Physics and Astronomy, 405 Hilgard Ave, Los Angeles, CA 90095
**Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL 60510

Abstract. We discuss the recent experiments in the field of Plasma Wakefield Acceleration (PWFA), with emphasis on the FNAL experiment. After completion of this work, the next round of experiments will need to push the envelope on accelerating gradient, interaction length, stability, and accelerated charge. We present theoretical results dealing with plasma accelerator performance in the limit of high driver charge, and short wavelengths. We comment on the impact of such results on both the afterburner (single module) and staged (multiple modules) approaches to plasma-based accelerators.

INTRODUCTION

There has been much recent activity in the field of plasma wakefields. In the Plasma Wakefield Accelerator (PWFA), a dense high-energy beam displaces electrons in the plasma, resulting in strong longitudinal and transverse wakes. In the blowout regime, where the drive beam has enough charge to radially eject all plasma electrons from its path, the plasma wakes have the desirable property of being linear close to the axis. The focusing force in the blowout region is linear in radius and does not depend on longitude, and the accelerating field has no radial dependence. In this regard, the field uniformity of the PWFA scheme is able to match our expectations of conventional accelerator devices.

The PWFA can be powered by a variety of drive beam parameters, spanning the range from at least the 100 nC, 0.3 cm σ_z bunches found at the AWA facility[4][1], to the 4 nC, 20 micron bunches for the proposed E-164 experiment[2] at SLAC. The drive bunch length is generally matched to the wavelength of the plasma λ_p, with the optimum wakes developing when $\lambda_p \simeq \pi\sigma_z$. Given this flexibility, the PWFA can bridge the gap between the high end of copper-based structures such as the W-band (92 GHz) devices or even CLIC (30 GHz), to the operating wavelengths of the laser-driven plasma accelerators, generally around a λ_p of 30 microns. Given that one of the main difficulties with staging (accelerating a beam in a collection of separately powered plasma stages) comes from the tight spatial tolerances at these short wavelengths, the option of using a longer wavelength, along with a reduced acceleration gradient, can be very useful in this regard. In this paper, we explore some scaling law arguments that might help in identifying the relative merits of a given choice in operating frequency.

In order to propagate a drive beam in a plasma over interesting lengths, one has to select parameters which avoid the electron-hose instability, a deflection mode instability

CP647, *Advanced Accelerator Concepts: Tenth Workshop,* edited by C. E. Clayton and P. Muggli

that's analogous to head-tail effects in conventional linacs. Work in understanding how this instability affects long beams has been going on for many years, and has yielded convenient asymptotic growth formulas. The short beam case ($\sigma_z < \lambda_p$) doesn't obey this formula, and a fully 3-D computational approach must be used[3]. Due to the expense and time needed to perform these simulations, a complete characterization of the instability over a variety of input conditions remains a challenge in the field.

Experimental work in the field of plasma wakefields has been ongoing for some time, with results on PWFA and plasma lenses having been reported by groups including ANL[5], UCLA[6], Univ. of Tokyo, LANL[7], and SLAC[8]. Each new set of experiments has pushed the envelope toward higher acceleration gradients, as well as toward a greater understanding of the underlying mechanisms. The present round of experiments are producing gradients in the region of 100 MeV/m (with peak fields possibly reaching several times this number). While at first glance this number seems to be much less interesting than the 100 GeV/m shown in the laser wakefield experments, this mostly has to do with the added expense and complexity of generating the required drive pulses. Because of this, most PWFA experiments have relied on existing linac facilities for their drive beam. A purpose-built facility could make a considerable improvement in the peak attainable wakefield E_z, which scales as $E_z \propto Q/\sigma_z^2$, either by increasing the drive beam charge Q, or using compression to attain shorter bunches. In the absence of electron-hose, the plasma focusing force acting on the drive beam allows for a near-equilibrium beam propagation where the overall plasma length is set by the energy depletion length. Using, for example, a 10 GeV drive beam would result in a single-pass energy gain of the witness beam of nearly 20 GeV, assuming a transformer ratio of two.

Although much of the work on plasma wakefields is driven by the promise of plasma-based collider modules and final focus systems, other applications of the technology have also been proposed. The LANL experiment was started in an effort to build a source of VUV radiation. Another scheme which is nearing the proof-of principle experiment stage[9] consists of trapping plasma electrons by a sharp downward transition in plasma density, with the eventual application being a high-brightness source of electrons.

At the moment, the two places where the PWFA scheme is actively being researched are SLAC and FNAL. As the SLAC experiment is covered elsewhere in these proceedings, this paper will concern itself mainly with the FNAL experiment. A very interesting result coming from the SLAC experiment is the work done on beam steering by a transverse gradient in plasma density. As the drive beam comes close to the edge of the plasma region, the ion column remaining after the blowout is asymmetric. The electrostatic forces arising from this tend to pull the beam toward the center of the plasma. This result can help establish the amount of spatial inhomogeneity that can be tolerated in order to keep beam steering under control in a collider application.

WAKEFIELD SATURATION AT HIGH CHARGE

Although this section discusses a theoretical result which does not yet manifest itself in the current set of experiments, we believe that parameter sets developed for future experiments will need to consider these effects.

In an earlier paper[10], we developed an analytical treatment of the energy loss of a drive beam in a plasma. Energy coupling between the drive beam and the plasma is crucial to the functioning of the PWFA, since if the energy is not coupled into the plasma wave, it can not be used to accelerate a trailing beam. In order to make any headway with a purely analytical approach, we assume a drive beam of infinitesimal length (beams with a finite σ_z will be treated later with simulations). In the linear regime, the plasma response due to a disk-like beam of radius a is given by solving the differential equation for the azimuthal magnetic field H in the plane of the disk,

$$\frac{\partial H}{\partial r^2} + \frac{1}{r}\frac{\partial H}{\partial r^2} - \frac{H}{r^2} - H = \frac{Q}{\pi a^2}\delta(r-a)\delta(z-ct),$$

where Q is the drive beam charge expressed as a dimensionless quantity, $Q = 4\pi k_p r_e N_b$, with N_b the number of drive beam electrons. The quantity H is likewise expressed as a unitless quantity, defined in terms of the wavebreaking field, while a is expressed in terms of k_p^{-1}. The drive beam is assumed to be ultrarelativistic, with its velocity given by $v \simeq c$. The solution of this equation is,

$$H(r) = E_r(r) = \begin{cases} K_1(a)I_1(r)\delta(z-ct) & (r < a) \\ K_1(r)I_1(a)\delta(z-ct) & (r > a) \end{cases}$$

where K_1 and I_1 are modified Bessel functions. The longitudinal wake directly behind the bunch, with $a \ll 1$, is given by,

$$E_z(r) \simeq \frac{Q}{2\pi}\left[\ln\left(\frac{2}{a}\right) - 0.577\right].$$

In physical units, this formula becomes.

$$eE_z \simeq 2e^2 k_p^2 N_b \ln\left(\frac{1.123}{k_p a}\right).$$

Although derived using a slightly different method, this result is identical to that given in J.D. Jackson.

So far, the derivation assumes that $Q \ll 1$, so that relativistic effects in the plasma electron motion are not important. As Q is raised we must consider the plasma electrons' relativistic response. The differential equation above evaluates the unit kick from the electromagnetic fields to be $p_r = nv_r = \int H dt$. The relativistic case differs from this in that there is also a longitudinal momentum impulse, given by $p_z = \frac{1}{2}p_r^2$, which follows from the electrons' motion in crossed electric and magnetic fields where $E_r = H$. This momentum is clearly small when $p_r \ll 1$, but dominates the overall momentum in the opposite limit. The expression for v_r then becomes,

$$v_r = \frac{p_r}{\gamma} = \frac{p_r}{\sqrt{1 + p_r^2 + \frac{1}{4}p_r^4}} = \frac{p_r}{1 + \frac{1}{2}p_r^2}.$$

As the plasma electrons are pushed in the direction of beam travel, the plasma electron density is modified according to,

$$n = (1 - v_z)^{-1} = 1 + \frac{1}{2}p_r^2,$$

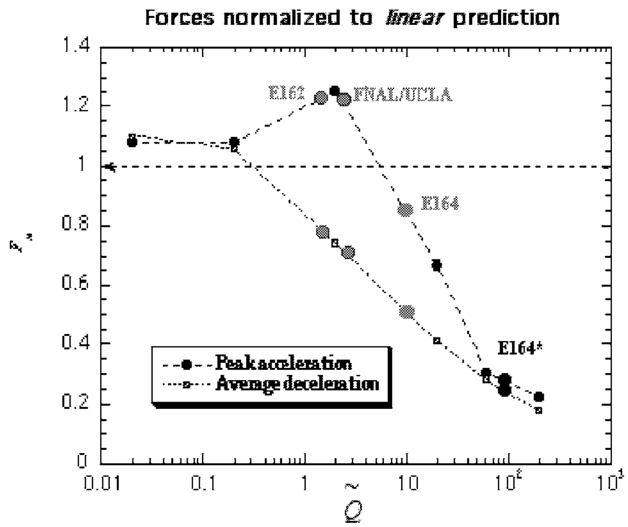

FIGURE 1. Accelerating and decelerating gradients as a function of normalized charge in present and future experiments. Unity on the y-axis corresponds to linear scaling.

with the result that the product of the two quantities $nv_r = p_r = \int H dt$ is unchanged from the linear response case $(Q \ll 1)$. This leads to the remarkable result that the above differential equation, and hence the system response, remains the unchanged in the large Q limit. This prediction was confirmed with the code NOVO, which showed that for very short bunches $\left(\sigma_z = 0.01 k_p^{-1}\right)$, the energy loss per charge stayed constant within a few percent in a range starting from $Q = 0.002$ to $Q = 2$, after which the method used in this code can no longer be applied.

The analytic results developed above are no longer valid for finite in time bunches, and a PIC code must be used. In contrast to the results in the infinitely short bunch limit, in the finite bunch case, we notice a saturation of energy coupling at high charge, which shows up in both the energy loss and peak energy gain in the plasma. Although the finite σ_z beam case is undoubtedly more complicated than the mechanics of the analytical result, we believe the PIC results are suggestive of the relativistic mass effect dominating over the snowplow effect in this regime. Figure 1 summarizes the results from several runs using parameters appropriate for the present round of experiments, as well as the proposed SLAC E-164 and E-164* set of experiments. We note that the E-164* result is a factor of three lower than predicted by an extrapolation of linear theory.

SCALING OF SYNCHROTRON LOSSES AND MULTIPLE SCATTERING

In the previous section, the use of dimensionless quantities is most helpful in deriving a result that can be applied over a broad range of plasma conditions. Conceptually, this is the same as applying frequency scaling to a problem. Unlike RF resonant cavity based accelerators, the scaling laws are exact for a plasma, due to the absence of a medium such as copper which is not scaled along with the rest of the problem. Incidentally, it is the deviation of the skin depth from ideal scaling that ultimately leads to shock heating limits for cavities at high frequencies. In this spirit, it becomes a worthwhile pursuit to find plasma phenomena which don't obey the usual scaling laws, as these might form the basis for new limits of plasma accelerators.

In frequency scaling, the field quantities E and H scale as ω, while lengths x, and time t, scale as $1/\omega$. Beam quantities like the beam emittance ε, and the charge Q also scale as $1/\omega$. The scaling preserves relativistic dynamics quantities such as the Lorentz factor γ, normalized velocity $\beta = v/c$.

In addition to the synchrotron losses and multiple scattering effects discussed below, the work done by Bruhwiler indicates that ionization phenomena should also be added to this list, since the effect turns on rapidly above 4 GeV/m for lithium and 20 GeV/m for hydrogen.

Synchrotron radiation

Synchrotron radiation is one phenomenon which does not obey frequency scaling. A particle traveling in a curved trajectory will radiate power at a rate $P \propto \beta^4\gamma^4/\rho^2$, where ρ is the radius of curvature. Holding γ and β constant, the amount of energy radiated in one plasma period will scale as $\omega^{-1}\rho^{-2} \propto \omega$, which increases at higher plasma frequencies, while scaling demands that it be invariant. Therefore, synchrotron radiation effects will tend to become more important at higher plasma densities.

The influence of synchrotron radiation on a misaligned beam is discussed in a paper by Montague and Schnell[13]. In their example of a 1 TeV beam transversely misaligned by 100 microns in a 1.1×10^{16} plasma, the betatron-averaged synchrotron loss is 766 GeV/m, almost a factor of 100 larger than the wavebreaking limit for these parameters. To compare this to a case with a 1.1×10^{14} plasma, with the misalignment scaled to 1 mm, the same beam loses energy at a rate of 7.6 GeV/m, or nearly a factor of 10 times greater than the wavebreaking limit. This difference of a factor of 10 between dimensionless energy loss for the two cases corresponds to the factor of ω predicted above from basic properties of synchrotron radiation. Although the examples above assume very extreme misalignments, to the point that the beam is likely to fall outside of the pre-formed ion channel, the large deceleration numbers should alert one to the seriousness of the problem.

There are several sources of misalignments in a plasma accelerator. The drive and witness beams can enter the plasma at different transverse positions, or have slightly diverging momentum vectors. Additionally, effects such as the electron-hose instability can shift the electrical center of the transverse wake away from center.

Misalignments can be a serious problem for a staged plasma accelerator. Since the drive and witness beams travel different paths, they may accumulate different errors in beam position and angle. One cure for this is to provide a focusing force at the plasma that's tied to laboratory coordinates, such as a solenoidal field. Since this approach might prove too costly, a single drive beam could propagate through a series of plasma sections separated by quadrupole magnets. Feedback techniques can also be used to improve the drive beam pointing stability. One common design for a staged plasma accelerator passes each drive beam through a 180 degree bend prior to entering a plasma stage. The signal from a beam pickup located before the bend can be amplified and sent to a fast kicker after the bend.

A single-pass device such as the afterburner would seem to have an advantage in terms of misalignments, since the drive and witness beams are co-propagated over the entire beamline length, and thus see the same external field errors. One source of beam errors not necessarily present in the staged approach come from the drive beam's transverse wakes acting on the witness beam. Also, the afterburner argument is strongest in the context of doubling the energy of an existing facility such as the SLC. Space limitations arising from working within an existing infrastructure will favor a short plasma with a high plasma density, leading to tighter tolerances due to synchrotron effects.

Multiple scattering

Although a much less serious concern, emittance dilution due to multiple scattering effects also doesn't obey frequency scaling. Multiple scattering effects on the witness beam were discussed in the same paper that dealt with synchrotron effects[13]. In terms of scaling, the scattering angle, θ_{scatt} experienced by a beam particle after propagating through a scaled plasma distance of $k_p^{-1}$is:

$$\left\langle \theta_{scatt}(\Delta z = k_p^{-1}) \right\rangle \propto \sqrt{\omega}.$$

This number should be compared with the matched beam RMS angular divergence $\sigma_{x'}$, which remains constant as a function of scaling. Since the above expression increases with ω, multiple scattering is a bigger problem at higher plasma densities. This is a direct consequence of the stochastic nature of this phenomenon. Although a higher density plasma accelerator is shorter in length, the overall number of collisions is larger. The multiple scattering limit might be reached in a plasma accelerator that needs to deliver an ultra-low emittance beam.

RESULTS FROM THE FERMILAB EXPERIMENT

The Fermilab-NICADD Photoinjector Laboratory (FNPL)[14] is a 16 MeV electron accelerator built as a prototype source for the TESLA Test Facility. It features a 1.6 cell, 1300 MHz, normal-conducting (copper) gun, followed by a superconducting 9-cell standing-wave cavity. A high-efficiency photocathode based on cesium telluride creates

8 nC (design) charge per bunch. Next to the relatively high beam charge, the principal advantage of this facility for wakefield use is the ability to temporally compress the bunches using a dipole chicane. The best compressed case data for an 8 nC beam has yielded an RMS bunch length of 0.6 mm, as measured by a Hamamatsu C5680 streak camera looking at OTR light from a front-surface aluminized mirror. Although this bunch length is short enough to match to the wavelength of a $2 \times 10^{14}\mathrm{cm}^{-3}$ plasma, it is difficult to obtain on a typical run day. Given the beamline optimizations that are performed for the plasma experiment, a bunch length of $\sigma_z = 1\mathrm{mm}$, and a $1 \times 10^{14}\mathrm{cm}^{-3}$ nominal plasma density are the default set of run conditions for this experiment. The effective longitudinal plasma region is 8 cm.

A hollow cathode plasma source was developed for this experiment. Unlike a traditional hollow cathode arc, where the cathode is heated by the plasma itself, the tantalum cathode is directly heated by an external DC power supply. The relatively large cross-sectional area of the tantalum tube results in low resistance in the electrical circuit, which requires a heating supply current of > 1000 A to be delivered to the load. This allows for pulsed operation of the working gas (argon) and the arc current (110 A for 0.5 ms). Recent temperature measurements indicate a peak cathode temperature in the neighborhood of 2400 K, with a relatively good background pressure level of 4×10^{-8} Torr at this temperature, and no argon pulsed into the system.

In order to separate the high argon pressure needed to operate the arc from the upstream accelerator equipment, a 10 micron thick aluminum window was placed on the upstream side of the plasma chamber. The aluminum surface is also useful as a beam OTR profile monitor, although surface roughness in the present configuration limits its effectiveness as a quantitative diagnostic. The foil also serves as the upstream plasma boundary, with the high-density plasma region directly behind it. This directly exposes the foil to the intense wakefields created in the plasma, a fact which may have contributed to our recent foil failure.

After the plasma chamber, the beam energy is analyzed in an imaging, broadband spectrometer. This system includes a vertical focusing quadrupole, a 145 degree bend magnet, and a phosphor screen. The spectrometer can only give a time-integrated measurement of the particle distribution.

For a compressed beam with the bunch charge ranging form 4 to 8 nC per bunch, the mean beam energy with the plasma turned off ranges from 13.6 to 14.8 MeV, with the peak intensity at 13.8 MeV. With the plasma on, the energy distribution becomes much broader. Although the phosphor screen at the spectrometer output window covers almost a factor of two in energy, the beam energy range is broader than these limits, and it is not possible to capture the entire energy range in a single shot. Figure 2 shows the low end of this range, with a strong signal at 4 MeV, and a lower end-point of 3 MeV, which is not visible in this frame.

We have used the fluid code NOVO to simulate our experimental conditions. This code has been modified to include a super-particle representation of the beam electrons. The beam distribution has been modeled with a temporal asymmetry, to match our experimental conditions. The streak camera measurements of the current profile show a steep rising edge followed by a slower falling edge, as expected from the process of dipole chicane compression. This type of beam shows more deceleration in the plasma than a corresponding Gaussian distribution. The code was run with a split-Gaussian

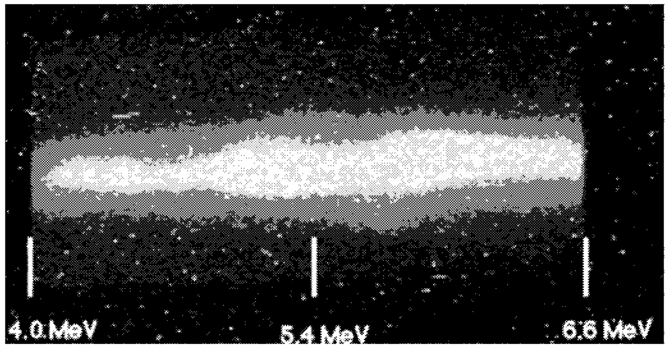

FIGURE 2. Beam phosphor screen image at spectrometer output window showing decelerated electrons.

representation for this beam, with a $\sigma_{rise} = 0.5k_p$, and $\sigma_{fall} = 2.0k_p$. The simulation presented here also assumes a beam charge of 7 nC, an initial beam spot of $\sigma_x = 200$ microns, and a normalized emittance of $\varepsilon_n = 45$ mm-mrad. The low-energy end-point of $E = \gamma mc^2 = 3.2$ MeV for this simulation is in good agreement with the observed result. However, in the experiment the compressed bunch length can have significant fluctuations which we are not able to record because the beam cannot be simultaneously sent to both the plasma and the streak camera OTR screen. Varying the bunch length in the simulations by +/- 25% can yield an end-point as low as 1.2 MeV.

The longitudinal phase space from this simulation is shown in Figure 3. In this simulation, particles in the tail of the beam are accelerated all the way up to 31 MeV, with an average acceleration gradient of 210 MeV/m. The actual number of particles observed at the highest energies will depend on a number of factors, including the precise shape of the beam tail, transverse effects such as beam head erosion, and 3-D effects like electron-hose. Additionally, some portions of the accelerated tail might be too dilute to definitively stand out above the noise.

When the spectrometer is tuned to transmit the high-energy end of the spectrum, we observe accelerated electrons up to 20.3 MeV, as shown in Figure 4. In order to properly display the high-energy tail on the printed page, the grayscale levels of the picture have been manipulated in a way that makes a large part of the image appear saturated. In order to compute the acceleration gradient, we use a conservative estimate for the starting energy of the electrons of 14.5 MeV. This number corresponds to the upper half-maximum of the energy distribution with no plasma. The average acceleration gradient is, therefore 72 MeV/m. However, the accelerated tail does not appear to stop at 20.3 MeV, which is the maximum range on the spectrometer. In the future, the system will be upgraded to allow the measurement of even higher energies and gradients.

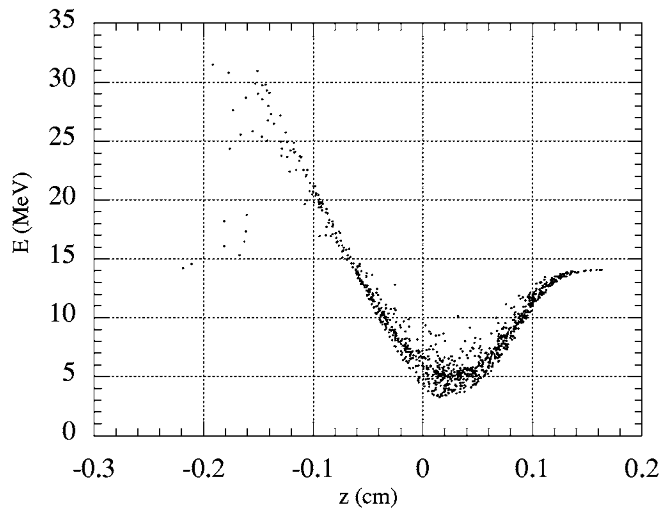

FIGURE 3. Simulated longitudinal phase space.

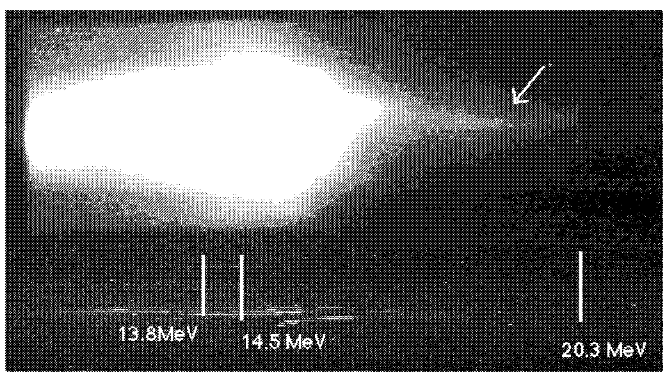

FIGURE 4. Spectrometer image with plasma on. The accelerated electrons appear as a whisker-like feature near the arrow.

CONCLUSIONS AND FUTURE WORK

Plasma wakefield experiments will continue to demonstrate higher gradients in future investigations. While this paper only touched on a few of the aspects of an overall plasma accelerator design, it is important that future experiments keep striving for improved results for all the parameters required of a well-engineered accelerator subsystem. Some of these advances will come from a better understanding of topics like the electron-hoseing, while others will come from improvements in plasma source design. Regardless of which accelerator layout is being considered, the role of drive beam stability cannot be overestimated. Some future advances in this field will undoubtedly come from improved drive beam handling techniques, combined with a better understanding of how to tie the capabilities of drive linac technologies to the requirements of plasma accelerators.

For the next stage of the Fermi/NIU/UCLA experiment, we have proposed to study beam loading of the plasma wave by a compressed witness beam. For a high enough witness beam current, the beam loading can cancel the natural energy slew which results in a higher energy gain for the beam tail. We will look at the narrowing of the witness beam energy spread as a signature of beam loading.

A beam test of the transition trapping experiment[9] of M.C. Thompson will also be scheduled at FNPL in the near future.

REFERENCES

1. M.E. Conde, et al., Phys. Rev. Special Topics -Accel. & Beams **1**, 04302 (1998).
2. W. Lu, these proceedings.
3. E.S. Dodd, et al. Phys. Rev. Lett. **88**, (2002).
4. J. B. Rosenzweig, et al., Phys. Rev. A **39**, 1586 (1989).
5. N. Barov, et al., Phys Rev. Special Topics - Accel. & Beams, 3, 011301 (2000).
6. G. Hairapetian, et al., Phys. Rev. Lett. **72**, 2404 (1994).
7. S. Russell, et al., Proceedings of the 2001 IEEE Particle Accelerator Conference.
8. M. Hogan, et al. these proceedings.
9. M.C. Thompson, these proceedings.
10. N. Barov, J.B. Rosenzweig, and M.C. Thompson, submitted to Phys. Rev. Lett.
11. J.D. Jackson, *Classical Electrodynamics*, 2nd Ed, (Wiley, New York, 1975).
12. D. Bruhwiler, these proceedings.
13. B.W. Montague and W. Schnell, *Laser Acceleration of Particles*, C. Joshi, ed. (AIP, New York, 1985).
14. J.P. Carneiro, et al., Proceedings of the 1999 IEEE Particle Accelerator Conference.