IOP PUBLISHING PLASMA PHYSICS AND CONTROLLED FUSION

Plasma Phys. Control. Fusion **49** (2007) 373–384 doi:10.1088/0741-3335/49/4/002

Spectral control in proton acceleration with multiple laser pulses

A P L Robinson[1], D Neely[1], P McKenna[2] and R G Evans[1,3]

[1] Central Laser Facility, CCLRC Rutherford-Appleton Laboratory, Chilton, Didcot, Oxfordshire, OX11 0QX, UK
[2] SUPA, Department of Physics, University of Strathclyde, Glasgow, G4 0NG, UK
[3] Blackett Laboratory, Imperial College London, Prince Consort Road, London, SW7 2BZ, UK

Received 22 November 2006, in final form 1 February 2007
Published 22 February 2007
Online at stacks.iop.org/PPCF/49/373

Abstract

We address the question of whether multiple high intensity ($>10^{18}\,\mathrm{W\,cm^{-2}}$) laser pulses (each O(100 fs)) can produce proton beams with a modified energy spectrum on irradiating a foil target. This has been studied in one dimension with both Vlasov and particle-in-cell codes. A reduction in the maximum proton energy, and the generation of spectral peaks, is observed. This is the first theoretical demonstration of spectral peak generation by purely optical means. The mechanism, which has been termed multiple pulse sheath acceleration, that produces these spectral peaks is described, and the same mechanism occurs in both sets of simulations.

1. Introduction

The production of multi-MeV protons and ions in ultraintense ($>10^{18}\,\mathrm{W\,cm^{-2}}$) laser–solid interactions has been studied extensively both experimentally [1–15] and theoretically [16–28] in the past few years. This means of producing energetic ions has the potential to be applied to both medicine and fast ignition (FI) ICF, as well as other areas.

One of the main problems with this method, especially as regards applications, is that in most experiments the energy spectrum of the protons or ions is a broad quasi-Maxwellian distribution with a sharp cut-off. Proton cancer therapy demands a narrow energy spread (i.e. a monoenergetic beam) to ensure that the proton energy is deposited over a small distance in the patient's tissue. In the case of FI a very tight energy spread may not be required, but ignition calculations [29] have made it clear that it is desirable to be able to control the energy spectrum as this may lead to a significant reduction in the proton beam energy required for ignition.

Recent experiments have shown energy spectra with quasi-monoenergetic features, and this has mainly been achieved by careful target engineering [12–15]. At the moment the number of ions in such features is very small, but this may well be improved on in future

experiments. It should also be noted that rings of protons were found on detector plates (in CR-39 stacks) in some experiments, which were quasi-monoenergetic [1].

The physical mechanism of ion acceleration has been debated, however there is good evidence that the proton emission in a number of experiments is due to the target normal sheath acceleration (TNSA) mechanism [18]. This is sometimes also referred to as simply 'plasma expansion' [16]. This is electrostatic ion acceleration that occurs as the fast electrons generated by the laser interaction try to escape into the vacuum. When laser intensities reach $10^{23}\,\mathrm{W\,cm^{-2}}$ then the acceleration may be dominated by the 'Laser-Piston' mechanism, but observations at 10^{18}–$10^{20}\,\mathrm{W\,cm^{-2}}$ are often attributed to TNSA. Shock acceleration and other mechanisms may make a significant contribution under certain circumstances [30].

In this paper the question of the control of the proton energy spectrum by using multiple O(100 fs) laser pulses is addressed for the first time. The minimal problem to consider is whether two pulses of fast electrons that propagate to a plasma–vacuum interface produce a significantly different expansion than a single pulse of electrons. This was investigated in one dimension (1D) using a Vlasov solver with one spatial and one momentum dimension (1D1P). Since the outcome may be highly dependent on any assumptions made about the fast electrons, the next problem to consider is the 1D problem in which one uses a realistic source of fast electrons. This problem was studied using a 1D3P electromagnetic particle-in-cell (EM PIC) code. We show, for the first time, that two high intensity 40 fs pulses incident on a foil produce an energy spectrum with spectral peaks above the single pulse reference spectrum. This is the result of a novel two-stage process, that was previously not known to exist within the general TNSA mechanism. The new term we have given to this variant of TNSA is multiple pulse sheath acceleration (MPSA). The first pulse is at a quarter to a tenth of the intensity of the second pulse, so this increases the effective efficiency in certain proton energy ranges. This study is quite different from previous studies of laser pre-pulse. In the work of Grismayer and Mora [31], an exponential density gradient created by a low temperature expansion was considered in the case of a single ion species and a single electron temperature. Similarly in the work of Andreev *et al* [32], linear density ramps were considered in the case of a single ion species.

2. Vlasov method

The Vlasov code employed in this study is the same one as used in [33]. This is an Eulerian Vlasov solver based on upwind methods [34], which solves

$$\frac{\partial f_\alpha}{\partial t} + v_x \frac{\partial f_\alpha}{\partial x} + q_\alpha E_x \frac{\partial f_\alpha}{\partial p_x} = 0, \tag{1}$$

where α denotes the particle species. Three species are included—electrons, protons and a heavy ion species (C^{4+}). The code was initialized with 186 μm of plasma at an electron density of $1 \times 10^{29}\,\mathrm{m^{-3}}$. This choice of length is not related to the size of a real target, but this length is required so that all the fast electrons are contained in the mesh at the start of the simulation (i.e. so we do not have to inject any fast electrons artificially). From 0–140 μm a fast electron population was present with a distribution function given by

$$f = \frac{n_{\mathrm{f},1}}{N_1} \exp\left[\frac{-c\sqrt{p_x^2 + m_\mathrm{e}^2 c^2}}{k_\mathrm{B} T_{\mathrm{f},1}}\right], \tag{2}$$

where N_1 is a normalization constant. The fast electron density, $n_{\mathrm{f},1}$, was set at $1 \times 10^{27}\,\mathrm{m^{-3}}$, and the temperature, $T_{\mathrm{f},1}$ was set at 2 MeV. This population of fast electrons was the main

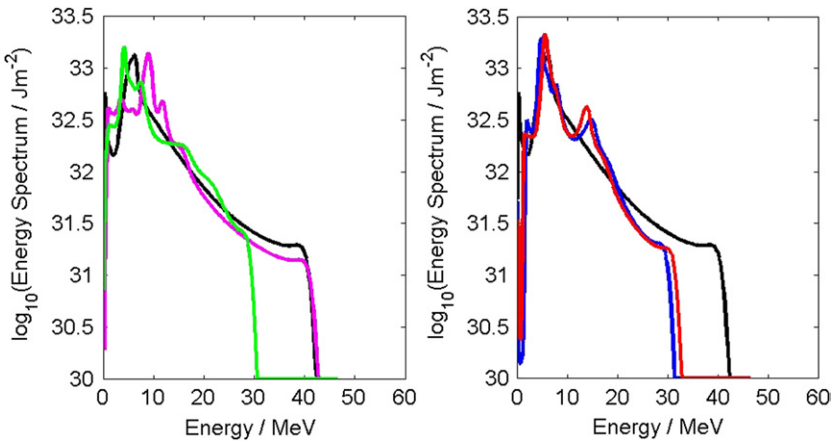

Figure 1. (Left) Proton energy spectra at 750 fs in run A (black), B (magenta) and E (green). (Right) Proton energy spectra at 750 fs in run A (black), C (red) and D (blue).

drive pulse. From 156–186 μm there was a 'set-up' pulse of fast electrons with a distribution function given by

$$f = \frac{n_{\mathrm{f},2}}{N_2} p_x^2 \exp\left[\frac{-c\sqrt{p_x^2 + m_\mathrm{e}^2 c^2}}{k_\mathrm{B} T_{\mathrm{f},1}}\right], \tag{3}$$

where N_2 is a normalization constant. The fast electron density, $n_{\mathrm{f},2}$, was set at $1 \times 10^{27}\,\mathrm{m}^{-3}$, and the temperature, $T_{\mathrm{f},2}$, was varied between runs. The cold electron temperature was set at 2.5 keV to reduce computation time. The majority of the target consisted of C^{4+} ions at $2.5 \times 10^{28}\,\mathrm{m}^{-3}$, except for a 40 nm surface layer that consisted of H^+ and C^{4+} in a 2 : 1 ratio, i.e. a typical contamination layer composition. The simulation was run up to 750 fs. A non-uniform momentum grid was employed for the electrons, but the momentum grids were uniform for the proton and ion species. A non-uniform spatial grid was used. Both spatial boundaries were reflective.

3. Results and discussion of Vlasov simulations

Initially five runs were carried out using the basic configuration (labelled A–E). In run A there was no set-up pulse of electrons, thus making this the 'comparison' runs. In runs B–E the temperature of the set-up pulse, $T_{\mathrm{f},2}$, was set at 1 MeV, 500 keV, 400 keV and 250 keV respectively. The proton energy spectra at 750 fs are shown in figure 1. This clearly shows that spectral changes have been induced by the set-up pulse.

In the single pulse run, A, a typical expansion from an inhomogeneous target is observed. The energy spectrum is predominantly quasi-exponential. The dip at very low energies is due to both the electrostatic shock due to the separation of fast and cold electron populations (rarefaction shock), and the electrostatic shock due to separation of protons and C^{4+} [3, 22]. The maximum proton energy reached is in good agreement with what is expected from the Mora expression [16] which for long expansion times becomes

$$E_{\max} \approx 2Zk_\mathrm{B}T_\mathrm{e}\left[\log\left(\frac{2\omega_{\mathrm{pi}}t}{\sqrt{2e^1}}\right)\right]^2, \tag{4}$$

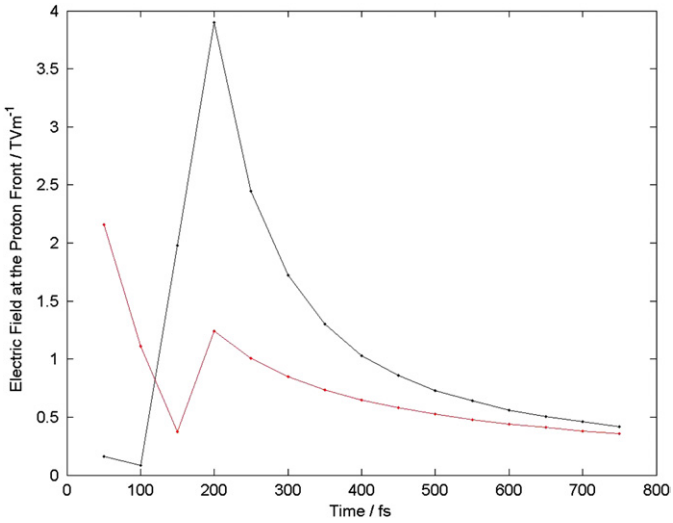

Figure 2. Electric field at the proton front against time for run A (black) and run C (red).

where $\omega_{\rm pi}$ is the ion plasma frequency ($\sqrt{Zn_{\rm e}e^2/m_{\rm i}\varepsilon_0}$). Taking the fast electron density and temperature, and an expansion time of 600 fs (allowing for the time taken for the main pulse to reach the plasma edge) an energy of 37.6 MeV is obtained, which is close to the observed value of 39 MeV at 750 fs. Overall, run A represents a fairly well understood scenario. In runs B–E we observe two important differences in this spectrum. Firstly there is a reduction in the maximum proton energy from that of run A, and secondly there is the presence of spectral peaks above the same energy in run A. In order to explain these two aspects of the simulation set it is necessary to consider the details of run C in particular, as well as those of runs A and B.

In the case where $T_{\rm f,2} = T_{\rm f,1}$, it is reasonable to expect that the proton spectra will not differ much, since the two pulses are very similar. Therefore it is not entirely unexpected that the results for run A and run B are similar. However, another argument might say that the higher the set-up pulse temperature is the longer the scale length of the expansion due to the set-up pulse. Therefore, on the basis of the work of Grismayer and Mora, the maximum proton energy should decrease as $T_{\rm f,2}$ increases. When one tracks the electric field at the proton front in each run from 200–750 fs, one does indeed find that the electric field is lower in run B than in run C. However at 150 fs, the maximum proton energy in run B is 14.5 MeV whereas in run C it is 7 MeV. So despite the reduction in the electric field that the main drive pulse produces, the net result is that the maximum proton energy in run B is higher than runs C–E, and approximately equal to run A. In the work of Grismayer and Mora the protons in the pre-expanded plasma had zero velocity, so one always observes a decrease in the maximum proton energy with initial scale length. This is not necessarily true if the protons have a non-negligible initial velocity. However for $T_{\rm f,2} < T_{\rm f,1}$ a reduction in the maximum proton energy is generally observed. This is because the effect of reducing the electric field due to the expansion caused by the set-up pulse is more significant than acceleration caused by the set-up pulse. This is illustrated by figure 2 where we plot the electric field at the proton front against time for runs A and C. When the main drive pulse electrons arrive between 150 and 200 fs, the electric field rises to a much smaller value in run C due to the effect of the set-up pulse. If one progressively reduces $T_{\rm f,2}$ below 250 keV then one would expect the maximum energy to then rise again as the effect of the set-up pulse is reduced.

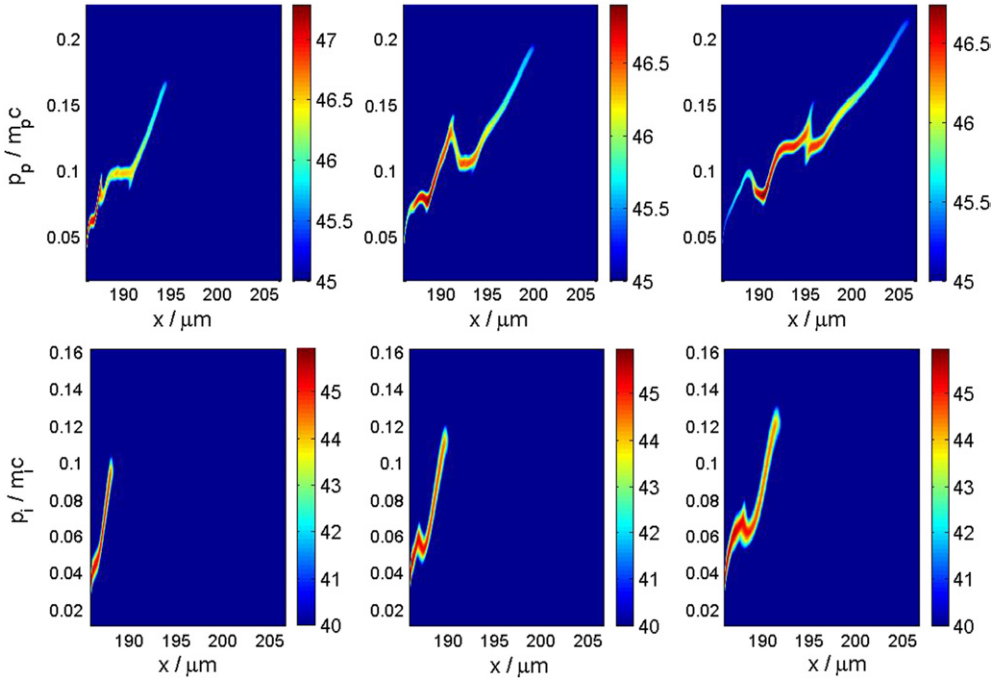

Figure 3. Top: Proton phase space in run C at 250 fs (left), 350 fs (centre) and 450 fs (right). Bottom: ion phase space at 250 fs (left), 350 fs (centre) and 450 fs (right).

The spectral peaks that are observed in runs C and D are generated by a two-stage process. The first stage occurs after the main drive pulse of fast electrons has reached the rear surface around 150 fs. The increase in the fast electron temperature greatly increases the strength of the electrostatic shock associated with the C^{4+} front, driving a surge of protons ahead of it. This, however, is only transient and the surge is only driven for ≈200 fs Some protons are even accelerated to higher energies than protons ahead of them. This leads to a 'wave-breaking' type of process [31] as these protons overtake the ones ahead of them (around 450 fs). Note that the term 'wave-breaking' is only used in the sense of particle over-taking. This results in a spike in the proton density at this point. This is shown in figure 3, which shows both the proton and the C^{4+} phase space at 250, 350 and 450 fs in run C. This shows the surge of protons at the C^{4+} front (250 and 350 fs), and the 'wave-breaking' event which happens away from the C^{4+} front.

From 500 fs onwards the peak in the spectrum actually emerges. The spike in the proton density creates a corresponding spike and dip in the electric field. Under quasi-neutral conditions one would expect this from,

$$E_x = -\frac{k_B T_e}{e n_p} \frac{\partial n_p}{\partial x}, \tag{5}$$

where n_p is the proton density. In the section where we present the PIC results this is illustrated by figure 8. The peak of the electric field lies ahead of the peak in the proton density, which one also expects from equation (5). The peak in the field changes its position, and decays as the proton density is smoothed out, however it does cause a sharp increase in proton velocity. The rate of increase of proton density beyond the peak is reduced. This is therefore accompanied by an accumulation of protons in momentum space which is the peak in the energy spectrum.

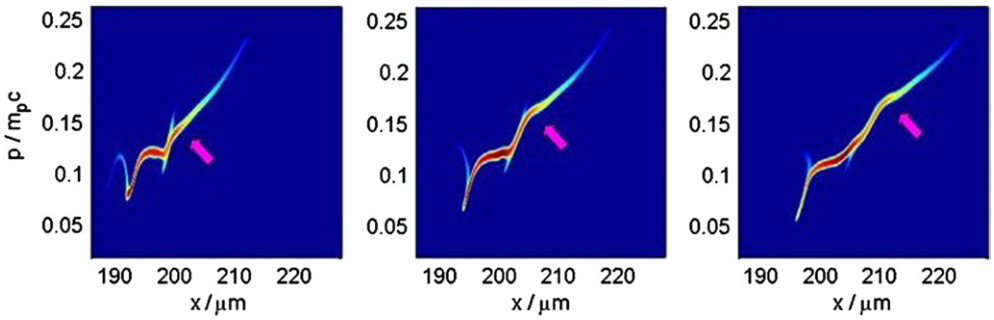

Figure 4. Proton phase space in run C at 550 fs (left), 650 fs (centre) and 750 fs (right).

This is shown in figure 4, which shows the proton phase space of run C during the second stage, at 550, 650 and 750 fs. The sequence in figure 4 clearly shows that the feature responsible for the spectral peak originates from the 'wave-breaking' region and is accelerated to higher energy. In figure 4 this feature is indicated by the purple arrow. It is important to note that this accumulation occurs some distance from the C^{4+} front, and it is not directly associated with it, as this distinguishes this mechanism from the direct generation of spectral peaks by heavy ion fronts [3, 22, 23, 33].

Given that this mechanism involves a degree of complexity beyond that involved in previous studies of spectral peak generation [3, 22, 23, 33], further investigation of the mechanism's sensitivity to the initial conditions is appropriate. This leads to the PIC simulations which we discuss later, and to further Vlasov simulations. Run B was repeated with the density of the set-up pulse set to $5 \times 10^{26}\,\mathrm{m}^{-3}$ (run F) and also with the density of the set-up pulse set to $2.5 \times 10^{26}\,\mathrm{m}^{-3}$ (run G). Since it is unreasonable to expect the fast electron density in a laser-solid scenario to be significantly higher than

$$n_\mathrm{e} \approx \frac{4\pi^2 a_0 m_\mathrm{e} c^2 \varepsilon_0}{e^2 \lambda^2}, \tag{6}$$

which is the relativistically corrected critical density, so only fast electron densities below $1 \times 10^{27}\,\mathrm{m}^{-3}$ were considered.

Figure 5 shows that the spectral peak is also generated in these runs, as well as the reduction in the maximum proton energy. Later runs were also carried out in which the length of the set-up pulse region was doubled (i.e. to 126–186 μm, and the main drive pulse was confined to 0–100 μm) and the spectral peak and the maximum energy reduction were still observed. It therefore seems to be the case that despite the mechanism being somewhat more complex than conventional TNSA, the existence of the spectral peak does not have a 'chaotic' sensitivity to initial conditions. The mechanism may actually be quite robust in this regard. On this evidence it is reasonable to conclude that two different electron pulses, with parameters relevant to ultraintense laser interactions, could induce significant changes in the proton energy spectrum. To verify the robustness of this conclusion PIC calculations in which the electron distribution function would be self-consistently calculated, and the effects of a finite-sized target would be accounted for, were carried out. It is this two-stage mechanism that underpins the MPSA mechanism and thus differentiates it from 'normal' TNSA.

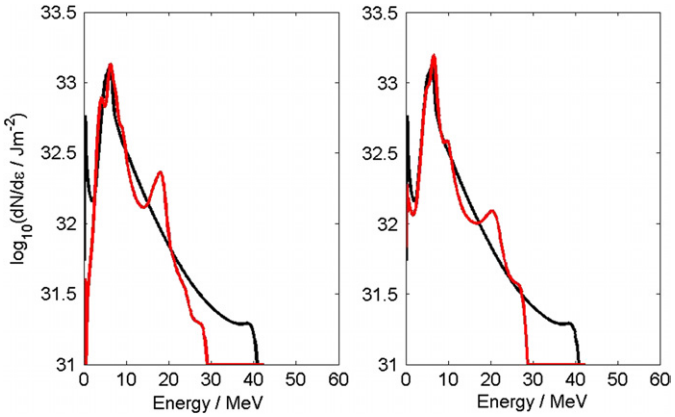

Figure 5. Left: proton energy spectrum at 750 fs in run F (red line). Right: proton energy spectrum at 750fs in run G (red line). Black is the comparison run A at 750 fs.

4. Particle-in-cell method

The 1D3P EM PIC code uses standard explicit PIC methods [35]. Linear interpolation is used for particle and grid weighting and the electric field obeys Gauss' law. In the simulations a grid of 50000 spatial points was used with a cell size of 2.5×10^{-9} m. A foil target with a thickness of 0.4 μm and an electron density of 2×10^{28} m^{-3} was placed at 60 μm. The electron temperature was set to 10 keV, and the ions consisted of protons and a heavy ion ($Z = 1$, $m_i = 3m_p$) in a 1 : 1 ratio. In all runs the 'laser' pulses had $\sin^4$ envelope functions, with a total length of $2ct_{pulse}$. Both pulses had $t_{pulse} = 40$ fs, and a wavelength of 1 μm. Quasi-particles of 1.6×10^6 were used for the electrons, and 8×10^5 quasi-particles for each ion species, so that initially there were 10000 quasi-electrons per cell, and 5000 quasi-ions per cell. The simulations were carried out up to 350 fs. The authors draw the readers attention to the fact that this initialization differs considerably from the way that the Vlasov simulations were initialized, and that the two sets of simulations are not intended to correspond to one another in a one-to-one fashion. In choosing two different scenarios, the robustness of the spectral peak generation mechanism to changes in the initial conditions is demonstrated, as well as the fact that the mechanism also works in the case of a homogeneous target (i.e uniform composition).

5. Results and discussion of PIC simulations

For comparison, a run (run I) was carried out in which a single (main drive) pulse was incident on the foil. The main drive pulse had a peak intensity of 4.4×10^{19} W cm^{-2} (i.e. $a_0 = 4$). The peak of the pulse was located at 48 μm at the start of the simulation. In the two pulse PIC simulations the set-up pulse had its peak located at 48 μm and the main drive pulse had its peak located at 15 μm. The main drive pulse was kept fixed with $a_0 = 4$. Three further runs were carried out. The set-up pulse had a peak normalized vector potential of $a_0 = 2$, 1.5, 1, 0.5 in runs II–V respectively. Note that when comparing runs II–V to run I one must account for the the fact that in run I the main drive pulse starts closer to the foil. This means that the time at which we should compare is 100 fs earlier in run I than in the other runs, i.e. the propagation time to the foil. We therefore take our single pulse reference for the spectra at 250 fs in run I.

In figure 6 the proton energy spectra of runs II–V at 350 fs is shown, along with the reference case of run I. There is a reduction in the maximum proton energy that is greatest in

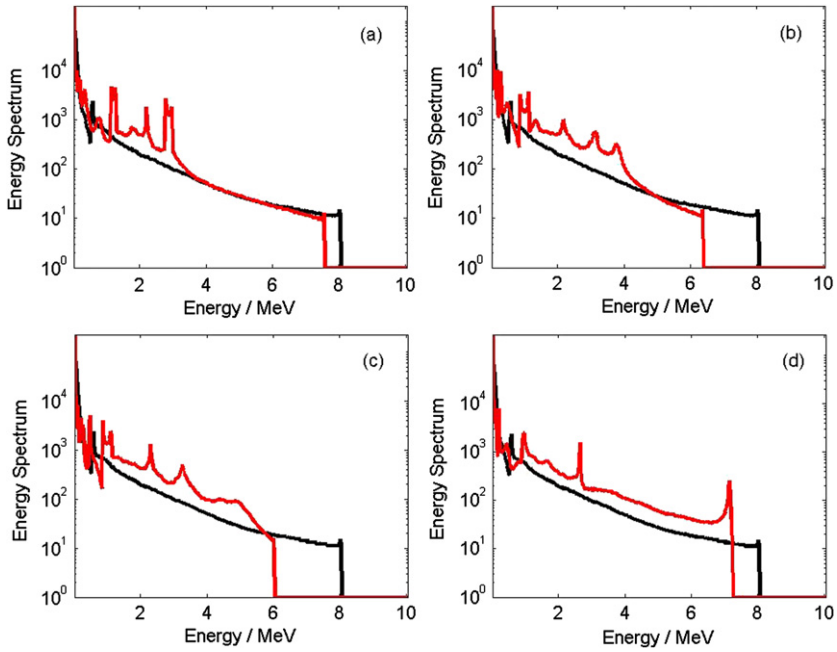

Figure 6. Proton energy spectra at 350 fs in PIC runs (red) II (*a*), III (*b*), IV (*c*) and V (*d*), against the single pulse reference case (black).

run IV. It is also clear that, as in the Vlasov simulations, there is the generation of peaks in the energy spectrum. Is this the result of the same two-stage process, i.e. MPSA? Let us consider run II, in which two similar stages have been identified, the first taking place between 140 and 220 fs, and the second taking place between 220 and 350 fs.

The main drive pulse just reaches the foil at around 120 fs, and the interaction takes place between 120 and 200 fs. A 'surge' in protons across the heavy ion front between about 140 and 220 fs is observed. This is shown in figure 7, in which both the proton and heavy ion phase space (for the forward going particles only) at 140, 180 and 220 fs are plotted. This first stage is therefore similar to that observed in the Vlasov simulations . The second stage, however, is somewhat different, in that there is no 'wave-breaking' event. There seems to be no over-taking of quasi-particles, and thus the protons can still be described hydrodynamically. There is however a strong peaking in the proton density, and this causes the expected change in the electric field that drives a second stage that is again very similar to that observed in the Vlasov simulations.

The second stage is illustrated in figure 8, in which the proton density, heavy ion density and electric field at the rear of the foil are shown at comparable times in both runs I and II. This shows that the peak in the proton density in run II generates a very strong change in the electric field, specifically the negative and positive spikes in the field. On the other hand, in run I the ion front only perturbs the proton density slightly, and the proton density falls smoothly and exponentially with position. The electric field is therefore approximately constant between the point where the heavy ion density terminates and the proton density terminates in run I. The proton and ion phase space at 350 fs is shown in figure 9. It is clear from figure 9 that we have created an accumulation of protons away from the heavy ion front. Both figures 8 and 9 depict a second stage that is very similar to that observed in the Vlasov simulations. On this

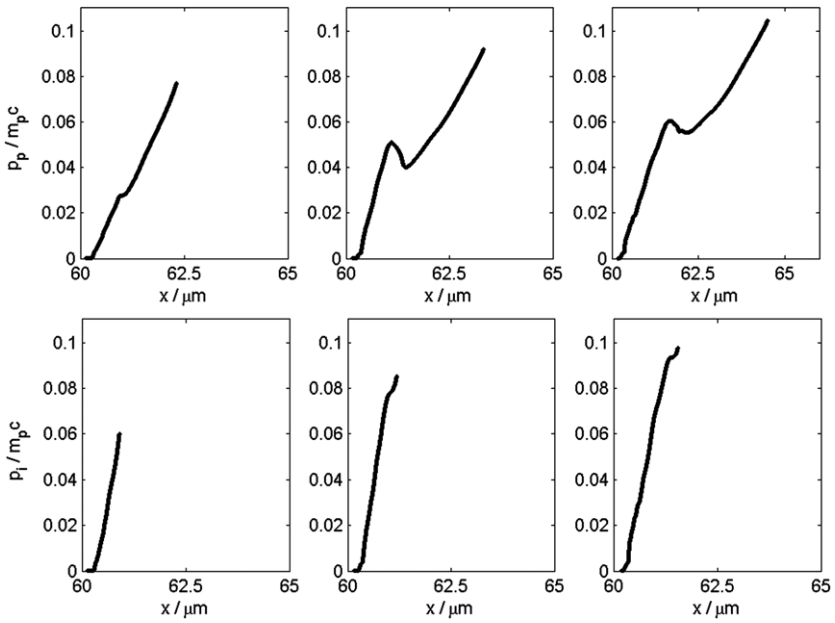

Figure 7. Top, left-right: Proton phase space in PIC run II at 140, 180 and 220 fs. Bottom, left-right: Heavy ion phase space in PIC run II at 140, 180 and 220 fs. Only forward going protons/ions are shown in both.

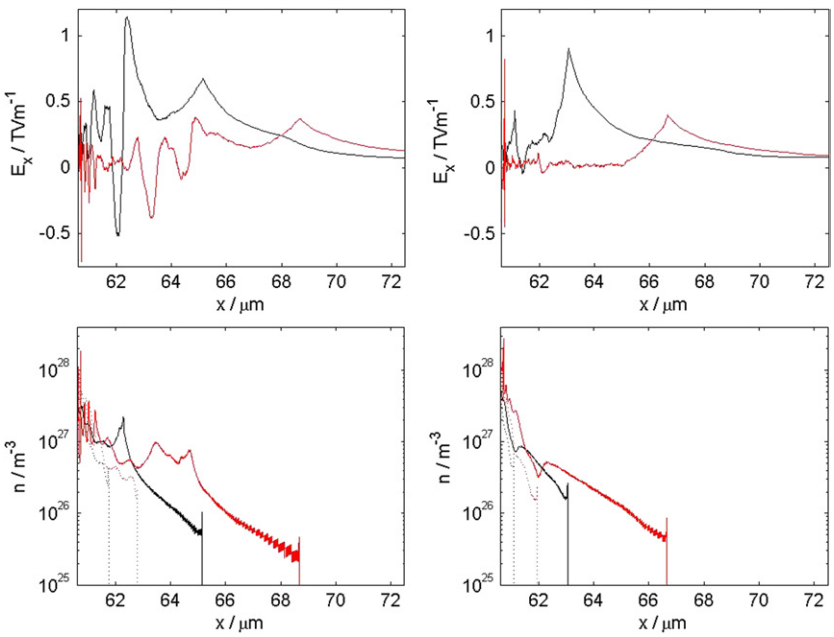

Figure 8. Electric field at rear of target at 240 fs (black) and 340 fs (red) in run II (top left) and at 140 fs (black) and 240 fs (red) in run I (top right). Proton/Heavy Ion density at rear of target at 240 fs (black/black dotted) and 340 fs (red/red dotted) in run II (bottom left) and at 140 fs (black/black dotted) and 240 fs (red/red dotted) in run I (bottom right).

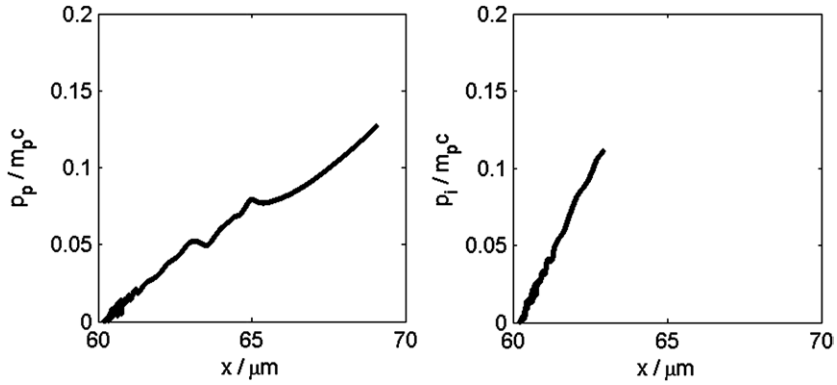

Figure 9. Left:Proton phase space at 350 fs in PIC run II. Right: Heavy ion phase space at 350 fs in PIC run II.

basis, the conclusion that was reached was that the MPSA mechanism was occurring in the PIC simulations as well. This process can have additional complications, however. In runs III and IV, the second stage repeats itself creating an additional spectral peak at high energy, and in run V the spectral peak is formed at the maximum energy (whereas it is formed at lower energy in run II).

It should also be noted that this explanation for the occurrence of the second stage is based on the proton density profile. No dependence on the detailed kinetics of the protons is required. Therefore the 'wave-breaking' event may be an incidental and non-essential phenomenon that occurs only under certain conditions. This matter will be the subject of future investigation.

Initially the PIC simulations were carried out at a reduced density ($n_e = 2 \times 10^{28}$ m^{-3}; $20n_{crit}$) and much stronger spectral peaks were generated. Simulations at reduced density are often carried out to reduce the computational resources required, especially in 2D and 3D. However these calculations suggest that this may lead to unrealistically optimistic predictions, and this implies that caution is required. Furthermore, it has also been found that the heavy ion phenomena are not critical to the generation of the highest energy spectral peak. This was found by repeating run II with a pure proton target (at the same electron density, i.e. $n_p = 8 \times 10^{28}$ m^{-3}), the results of which are shown in figure 10.

Finally a comment should be made on the difference in overall efficiency between the single and dual pulse runs. In the case of runs I and II it was found that run II was 68% more efficient than run I (the laser-proton efficiency in run I being about 4%), even accounting for the extra 25% expenditure of laser energy in run II. The total laser-proton efficiency is affected by two factors. Firstly the set-up pulse will change the absorption of the main drive pulse. Secondly there is the question of how efficient a normal plasma expansion is compared with one in which the two-stage process identified here occurs. One would expect the set-up pulse to produce a density gradient on the front surface that would enhance the absorption of the main drive pulse [36]. The second question is more difficult to answer, and this will be addressed in future work.

6. Summary and conclusions

1D1P Vlasov simulations of two fast electron pulses driving a plasma expansion and 1D3P EM PIC calculations of two short laser pulses driving ion acceleration from a thin foil have been performed. This was done to evaluate whether the use of two laser pulses produced

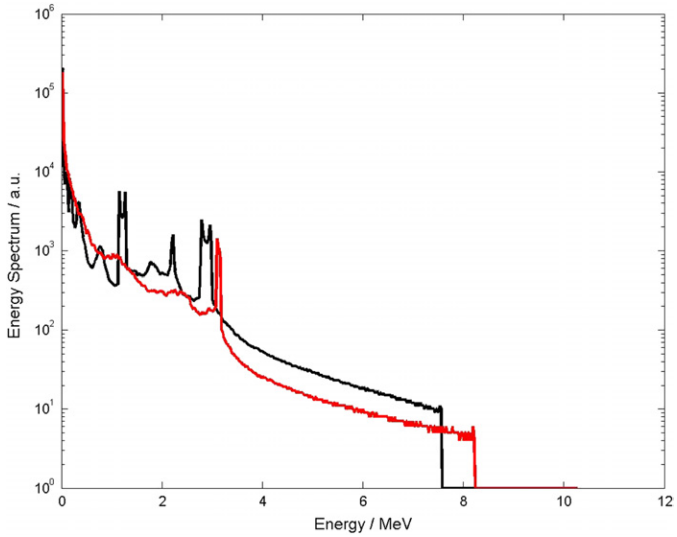

Figure 10. Proton energy spectrum in proton only run (red) and run II (black) at 350 fs.

any significant changes to the proton energy spectrum. It has been found that, in both sets of calculations, there are two important changes. Firstly, there is a reduction in the maximum proton energy and the number of protons at high energy. Secondly, there is the generation of spectral peaks. This was the result of a previously unknown two-stage process that can occur within the general TNSA mechanism. This gives rise to a TNSA variant that we have termed MPSA. The fact that two different algorithms provide similar results suggests that this is a genuine physical effect and is not a numerical artefact.

Since the spectral peak is a striking qualitative effect, an experimental study of two pulse interactions may further our general understanding of TNSA. At present it has not been demonstrated that the peaks have a sufficiently low energy spread or a high enough energy for a medical application. On the other hand, this scheme may be very useful for proton driven FI ICF. For additional laser energy input on the order of 10–25%, one gains an enhanced region where the number of protons increases by 2–10 times (taking the numbers from the PIC calculations). If one can design a specific laser-target configuration where this peak or enhanced region coincides with the 'useful' range of proton energies [37] for ignition then one can effectively enhance the efficiency. There is the additional benefit to proton driven FI of the reduction of the maximum proton energy, as this would reduce the undesirable 'pre-heat' of the compressed fuel. However more work needs to be done, including a study that incorporated an ignition calculation such as [29], before one can conclusively say that this would reduce the energy required for ignition. If the spectral peak generation can be improved then this may find application in proton radiography [38]. Importantly, this work points towards a method of spectral modification that is primarily optical rather than being based on target engineering. The optical route to spectral control is probably better for producing a high repetition rate laser-proton source, and it may well prove to be generally more practical. Another important aspect of this new optical approach is that it is not 'energy slicing', i.e. energy selection is not obtained by either eliminating protons at the post-acceleration stage or by using a limited proton source (either in thickness [13] or area [15]). Although this work does suggest that two pulse interactions can cause spectral modification, it is clear that multi-dimensional effects

need to be studied since the field changes induced by a peaking in proton density may differ significantly in 2D and 3D.

Acknowledgments

The authors are grateful for the use of computing resources provided by the CCLRC's e-Science facility.

References

[1] Clark E L 2001 Measurements of energetic particles from ultraintense laser-plasma interactions *PhD Thesis* Imperial College, London, UK
[2] Hegelich M *et al* 2005 *Phys. Plasmas* **12** 056314
[3] Allen M *et al* 2003 *Phys. Plasmas* **10** 3283
[4] Cowan T *et al* 2004 *Phys. Rev. Lett.* **92** 204801–1
[5] Roth M *et al* 2002 *Plasma Phys. Control. Fusion* **44** B99
[6] Roth M *et al* 2002 *Phys. Rev. ST-AB* **5** 061301
[7] Borghesi M *et al* 2004 *Phys. Rev. Lett.* **92** 05503-1–05503-4
[8] Zepf M *et al* 2001 *Phys. Plasmas* **8** 2323–30
[9] Zepf M *et al* 2003 *Phys. Rev. Lett.* **90** 064801–1
[10] Neely D *et al* 2006 *Appl. Phys. Lett.* **89** 1
[11] Beg F *et al* 2004 *Appl. Phys. Lett.* **84** 2766
[12] Toncian T *et al* 2006 *Science* **312** 410
[13] Hegelich M *et al* 2006 *Nature* **439** 441
[14] Ter-Avetisyan S *et al* 2006 *Phys. Rev. Lett.* **96** 145006
[15] Schwoerer H *et al* 2006 *Nature* **439** 445
[16] Mora P 2003 *Phys. Rev. Lett.* **90** 185002-1–185002-4
[17] Mora P 2005 *Phys. Rev.* E **72** 056401
[18] Wilks S C *et al* 2001 *Phys. Plasmas* **8** 542
[19] Sentoku Y *et al* 2003 *Phys. Plasmas* **10** 2009
[20] Esirkepov T Z *et al* 2002 *Phys. Rev. Lett.* **89** 175003
[21] Gibbon P 2005 *Phys. Rev.* E **72** 026411
[22] Tikhonchuk V T *et al* 2005 *Plasma Phys. Control. Fusion* **47** B869
[23] Bychenkov V Y *et al* 2004 *Phys. Plasmas* **11** 3242
[24] Murakami M and Basko M 2006 *Phys. Plasmas* **13** 012105.
[25] Dorozhkina D and Semenov V 1998 *Phys. Rev. Lett.* **81** 2691
[26] Pukhov A 2001 *Phys. Rev. Lett.* **86** 3562–5
[27] Ruhl H, Cowan T and Pegoraro F 2006 *Laser Part. Beams* **24** 181
[28] d'Humieres *et al* 2005 *Phys. Plasmas* **12** 062704
[29] Temporal M *et al* 2002 *Phys.Plasmas* **9** 3098
[30] Silva L *et al* 2004 *Phys. Rev. Lett.* **92** 015002–1
[31] Grismayer T and Mora P 2006 *Phys. Plasmas* **13** 032103
[32] Andreev A *et al* 2006 *Plasma Phys. Control. Fusion* **48** 1605
[33] Robinson A, Bell A and Kingham R 2006 *Phys. Rev. Lett.* **96** 035005
[34] Arber T D and Vann R G L 2002 *J. Comp. Phys.* **180** 339–57
[35] Birdsall C K and Langdon A B 1991 *Plasma Physics via Computer Simulation* (New York: Adam Hilger)
[36] Gibbon P 1994 *Phys. Rev. Lett.* **73** 664
[37] Roth M *et al* 2001 *Phys. Rev. Lett.* **86** 436
[38] Li C *et al* 2006 *Phys. Rev. Lett.* **97** 135003