



Introduction of a New longitudinal Electric Field Evaluation at Beam Position for HiPACE

John Moroney, Trinity College Dublin, Ireland

Supervisor: Alexander Aschikhin, FLA group, DESY

September 8, 2016

Abstract

The HiPACE particle-in-cell code for simulating plasma wakefield acceleration scenarios employs a zeroth-order scheme for interpolating field values to particle positions in the longitudinal direction. This gives rise to an unphysical structure in visualisations of longitudinal phase space of particle beams. This project aimed to implement a new longitudinal interpolation scheme that yields more physically accurate results. Two such alternative schemes were developed, one which mixes zeroth-order and first-order interpolation to obtain an improved field evaluation, and one which employs a degree of temporal extrapolation in search of a smooth first-order scheme throughout. The former was successfully implemented, while the latter yielded encouraging results, but requires some further examination

Contents

1	Introduction	3
2	Theory	3
2.1	Plasma Wakefield Acceleration	3
2.2	Particle-in-Cell simulations	5
2.3	HiPACE, the Quasi-Static Approximation and Parallelisation	5
2.4	Interpolation Schemes	6
2.5	Field Interpolation in HiPACE	6
3	Procedure	8
3.1	Initial Visualisations of Interpolation Scheme	8
3.2	Extending the Transverse Interpolation Scheme to the Longitudinal Direction	10
3.3	Improved Scheme in Parallel	11
3.4	An Alternative Solution	13
4	Summary and Conclusions	15
5	Acknowledgements	15

1 Introduction

The FLASHForward project is a plasma-wakefield particle-acceleration experiment that currently in development at DESY, Hamburg. High current density electron beams from the FLASH linear accelerator beamline will be used to drive a plasma wave through a cell of ionised Hydrogen a few centimetres long [1]. The wakefields produced by the driver beam can be used to accelerate particles to velocities near the speed of light, and the experiment aims to produce GeV-order beams that can be used in a free-electron laser. Plasma acceleration is a field that is still in its infancy, and it could bring great improvements in the ease of producing high intensity laser beams, which would benefit a wide variety of fields from biology to materials science. It was demonstrated in the Stanford Linear Accelerator Centre (SLAC) in 2006 that such plasma accelerators are capable of producing an energy gain in electrons of more than 40 GeV over a distance of just 85cm [2]. Compared with conventional linear accelerators such as the European XFEL, plasma-based accelerators have the potential to obtain similar particle energies, but on a length-scale four orders of magnitude smaller.

Given the difficulties associated with carrying out experiments with plasma in the laboratory (most notably the cost of building a device that can maintain the necessary conditions for the plasma), and also the difficulty in controlling the quality of the electron bunches injected into a wakefield accelerator, progress in this area is generally made through computer simulation [3]. HiPACE (Highly efficient Plasma Accelerator Emulation Code) is a particle-in-cell (PIC) code that was developed by Timon Mehrling at DESY to simulate beam-driven plasma-wakefield acceleration such as will occur in FLASHForward. The aim of my project this summer was to introduce a new scheme for evaluating the longitudinal electric field at the beam particle positions, that is more accurate than the zeroth order interpolation scheme currently used by HiPACE.

2 Theory

2.1 Plasma Wakefield Acceleration

Broadly speaking, there are two different types of plasma wakefield acceleration, but the concept is essentially the same. In Laser Wakefield Acceleration (LWFA), a plasma wave is driven by a short, intense laser pulse, whereas Plasma Wakefield Acceleration (PWFA) involves using a high current-density particle beam as a driver instead [1]. FLASHForward, and thus HiPACE, use PWFA, and specifically electron beams as drivers so this report will always refer to this method unless stated.

Wakefield acceleration works as follows. A charged particle in an inhomogeneous oscillating electromagnetic field experiences a ponderomotive force, which is given by $F_p = -\frac{e^2}{4m\omega^2} \nabla(E^2)$, where e and m are the charge and mass of the particle, and ω and E are the angular frequency of oscillation and amplitude of the field respectively. This force acts to push the plasma electrons out of the path of the driver beam, thus

separating them from the ions, and creating a travelling longitudinal electric field in the wake of the driver. This 'wakefield' has a phase velocity roughly equal to the mean velocity of the driver, so for relativistic electron beams, one obtains a phase velocity close to the speed of light [5]. By injecting electrons with sufficient energy into a period of the wakefield, they are trapped and accelerated by the field, and it is this ultrashort 'witness beam' that undergoes a significant step-up in energy [4]. Figure 1 shows the fields for such a configuration, generated by HiPACE.




Figure 1: An illustration of the electric field produced by a driving electron beam in a plasma accelerator

When dealing with plasma systems, it is convenient to use so-called plasma parameters. The time, t , is normalised to the inverse plasma frequency ω_p^{-1} , where $\omega_p = \sqrt{\frac{4\pi n_0 e^2}{m}}$ (n_0 is the plasma electron density), while distances are normalised to $k_p^{-1} = c/\omega_p$, the plasma skin depth. When dealing with relativistic particles, we work in the speed-of-light frame, with ζ defined as $z - t$. The dynamics of the particles in the witness beam are given by the Lorentz force

$$\frac{du}{dt} = \frac{q}{e} \frac{m}{M} (E + \beta \times B) \quad (1)$$

where $u = \frac{p}{Mc}$ is the normalised momentum, and q and M are the charge and mass respectively, of the particle under consideration, while e and m are the charge and mass of the electron. E and B are the electric and magnetic fields, which are both normalised to $E_0 = \omega_p mc/e$ - the cold-non-relativistic wave breaking field [6].

2.2 Particle-in-Cell simulations

The Particle-in-Cell (PIC) technique is a numerical method commonly used in simulating plasmas. In order to make computations possible on a realistic timescale, particles (e.g. plasma particles or electrons) are not considered individually, but instead the entire system is treated as a number of macroparticles, each of which represent a large number of real particles of a certain type. This is necessary because medium-sized plasma simulations contain in the order of 10^7 particles, which, if all treated individually, would give rise to a simulation that would take hundreds of years to complete on modern machines [3]. This rescaling is possible, because Equation (1) depends only on the charge-to-mass ratio of the particle, which is left unchanged. A spatial grid is also introduced, and the particles' charge and current densities are interpolated to this grid. Maxwell's equations for the electric and magnetic fields are then solved on the grid, and the fields at the gridpoints are then interpolated back to the particles to find the forces acting on them. The equations of motion are then computed for each particle, and the beam is thus 'pushed' along.

2.3 HiPACE, the Quasi-Static Approximation and Parallelisation

Even though a fully explicit PIC code that works as is outlined in the preceding section saves computing time by using macroparticles, 'pushing' each particle individually is still hugely time-consuming, and such codes today can only be run on supercomputers. The high computational cost is due to the limitation on the time-step size used in the partial differential equation solving routine, which is constrained by the Courant-Friedrichs-Lewy condition [6], [3]. This condition effectively couples both the spatial and temporal resolutions. However, HiPACE employs a Quasi-Static Approximation (QSA) which allows for the decoupling of the timescales of the plasma and the particle beams, making possible a much larger time step size than in full PIC codes. The QSA is applicable if the condition that the particle beam and plasma evolve on vastly different timescales holds - it assumes that the plasma can be treated as stationary while the electrons (in a Lorentz-boosted frame) are pushed.

Another way in which HiPACE saves computational time is by employing parallelisation in all three spatial directions. The Message Passing Interface (MPI) is used to distribute the simulation over a number of processes running on a several cores, which can communicate with one another. MPI enables communication between processes using so-called 'halo cells' around the edge of each processor's domain. Each halo cell has a corresponding cell in the neighbouring processor, so when information is updated in a halo cell in one process, MPI informs the next process to update the information in its corresponding halo cell. Because of the QSA, HiPACE reduces a 3D problem over a finite area to a set of 2D problems on transverse slices of this area. This allows plasma and field information to be advanced in the negative ζ -direction.

2.4 Interpolation Schemes

The Interpolation scheme currently used in HiPACE for electromagnetic fields in the transverse slices of each process is simple bilinear interpolation. Only the field values at the four grid points nearest to the particle are considered, as in figure 2 below. If Q_{11} , Q_{12} , Q_{22} and Q_{21} are the values of the fields at (x_1, y_1) , (x_1, y_2) , (x_2, y_2) and (x_2, y_1) respectively, then the field P at (x, y) is given by

$$P = \frac{1}{y_2 - y_1} Q_{11}(x_2 - x)(y_2 - y) + Q_{12}(x_2 - x)(y - y_1) + Q_{22}(x - x_1)(y - y_1) + Q_{21}(x - x_1)(y_2 - y) \quad (2)$$




Figure 2: A bilinear interpolation scheme

In three dimensions, trilinear interpolation simply involves performing bilinear interpolation onto a particle's projection on slices on either side of it, and performing linear interpolation to the particle from the two slices.

2.5 Field Interpolation in HiPACE

The HiPACE main loop consists of three steps. First, the beam current and charge density is deposited onto the grid. Then, the particle positions are unchanged while the plasma and field routines are performed - plasma and field information is 'back-streamed' in the negative ζ -direction. Finally, the new fields are interpolated onto the positions the beam macroparticles, which are then pushed. At the very beginning of the main loop, each processor inherits all the necessary information from its neighbouring processors using MPI, while at the very end of the loop, each processor passes all relevant information to its neighbours in the same way. After each loop, the system advances in time. While this structure provides for very efficient computations, one of its drawbacks is evident in the field interpolation step. When the code is parallelised in the longitudinal direction, each processor accesses a number of transverse slices. While

there is no problem with transverse interpolation on these slices, a particle on a given slice may not have access to the field values at points on the next slice to it, because this slice may be part of a different process. Furthermore, the back-streaming of the plasma and the fields means that a process running on a given core can not begin its plasma routine until the process running on the next core (in the ζ -direction) has finished its plasma routine. To minimise idle time that might result from this set-up, HiPACE is designed so that processes in consecutive domains run asynchronously such that the plasma routine in a domain finishes just as the plasma routine in the domain before it is called. The effect of this on interpolation is that if field values from a neighbouring process are required for interpolation to the position of a particle in a slice at the edge of one domain, the staggered time evolution means that these field values are not necessarily defined at the same time as the interpolation is occurring. For this reason, HiPACE interpolates fields only to zeroth order in the longitudinal direction. However, as MPI enables communication between adjacent processes, it should indeed be possible to implement a higher-order longitudinal interpolation scheme that involves some form of temporal extrapolation near process boundaries, and such an implementation was the goal of this project.

3 Procedure

3.1 Initial Visualisations of Interpolation Scheme

While HiPACE as it is now does not produce unphysical results in terms of how the macroparticles 'see' the electric and magnetic fields, the limitations of a zeroth-order longitudinal interpolation scheme become quite apparant at low resolutions, and for certain initial configurations. Below is a plot of the longitudinal phase space of a low-energy witness beam, with configuration that emphasises the effects of the longitudinal interpolation. The particle beam is initialised near the minimum of the longitudinal electric wakefield (E_z). As is evident from Figure 1, at this point, the field experienced by the particles at the back of the beam has a more negative value than the field for the frontmost particles. As q for electrons is also negative, the nature of the Lorentz force implies that the particles at the back are accelerated more (and therefore gain more energy/momentum) than those at the front. This explains the general shape of the graph below, and by initialising the beam with low energy ($\gamma = 100$ as opposed to $\gamma = 2000$ for the driver beam), this step-up in energy is more evident. The simulation is split evenly across four cores in the longitudinal direction. All further plots in this report will be of a witness beam with a similar configuration, in order to make a fair comparison.




Figure 3: Longitudinal Phase Space plot of witness beam at $t = 100$, simulation run on 4 cores for 10000 particles

A rather unphysical 'stepwise' structure is immediately evident on this plot. This is a direct result of particles having the field values of only their nearest 'slice' interpolated directly onto them. While at high resolutions this approximates a smooth curve, this structure could clearly be problematic for some configurations, so the aim of this project is to remove this structure altogether. It is also worth noting that this structure does not change over time. To illustrate this, Figure 4 shows the beam from the same simulation, but after 1000 timesteps.




Figure 4: Longitudinal Phase Space plot of witness beam at $t = 1000$, simulation run on 4 cores for 10000 particles

Although the beam shape has evolved, one can observe that the longitudinal interpolation (or lack thereof) still has the same shape. This makes sense, as only spatial interpolation is being carried out.

3.2 Extending the Transverse Interpolation Scheme to the Longitudinal Direction

The most obvious approach to take was to simply attempt to apply the transverse interpolation scheme to the longitudinal direction to achieve trilinear interpolation. To achieve this, the interpolation routine was modified so that the indices of the slices immediately before and after of the particle's position were identified, bilinear interpolation was performed in the transverse direction, before interpolating linearly between these two slices (thus achieving a trilinear interpolation scheme). When this new code was run on a single processor, the following phase space plot was produced:




Figure 5: Longitudinal Phase Space plot of witness beam with simple longitudinal interpolation at $t = 100$, run on a single core for 10000 particles

In Figure 5 it is immediately evident that the stepwise structure has been replaced by a smooth outline that one would physically expect to see. While this graph fulfills our objectives by eliminating unphysical structures, it is important to note that this simulation was run in series on a single core, and the problems with this approach only become clear when one runs the same simulation in parallel.




Figure 6: Longitudinal Phase Space plot, as above, run in parallel across 4 cores

The anomalous region of the graph in Figure 6 occurs in the region of the (longitudinally speaking) uppermost slice of a processor. Inspection of the process using the debugger GDB reveals that particles in this upper slab have access to the fields in the slice to before it, but all of the field values stored in the next longitudinal index of this processor are zero. This produces the sharp 'zig-zag' in the graph. As the first slice of the next process contains the correct field values (otherwise the curve would not be smooth directly after $\zeta = 5.0$), the question to be addressed was whether the processor directly before could access these values in order to perform full interpolation. As the fields evolve in the negative ζ -direction, in theory, this is possible using MPI. In the 'Inherit all information' step at the beginning of the HiPACE main loop.

3.3 Improved Scheme in Parallel

In order to eradicate the unphysical jagged section of the graph that occurs before the boundary between two processors, it was needed to obtain the field values from the lowest slice of the upper processor for use in the lower processor's interpolation scheme. To do this, a number of MPI routines were added to the 'Inherit All Information' step at the start of the main loop in each processor, and to the 'Pass All Information' step at the end of the main loop to pass the various field values from the upper processor to the lower one. The 'MPI_Irecv' and 'MPI_Isend' functions work by specifying a buffer into which to read information in the former, and a buffer from which information is passed in the latter. These buffers store information in a one-dimensional array. For the purposes of other routines in HiPACE, a buffer containing the field values in each processor had

already been allocated, so all that was needed was to allocate a new buffer into which the values of each of the fields could be recieved in the 'MPI_Irecv' function, and then write this one-dimensional buffer onto a two-dimensional halo slice using nested for-loops. With this in place, each processor has access to a halo slice beyond the uppermost slice of its own domain which can be used to interpolate fields to the positions of particles that lie beyond the previous uppermost slice. Having made this change, the following plot was produced, having run HiPACE in a parallelised configuration across four cores in the longitudinal direction:




Figure 7: Longitudinal Phase Space plot of witness beam at $t = 100$, run on 4 cores for 1000 particles, simulation using new improved interpolation scheme

This is a much more encouraging plot. Gone is the jagged area of the graph corresponding to particles in a processor's uppermost slab. Unfortunately, due to time constraints, this simulation was run with a low number of particles. (When HiPACE is typically run on a cluster of cores, the witness beam contains over 100,000 particles.) Because of this, it is difficult to make out the structure of the graph clearly, so we cannot say for certain if this new interpolations scheme is completely successful.

As this interpolation scheme involves some temporal extrapolation, it is important to monitor the scheme as time evolves. Below is a plot of the same simulation after 1000 timesteps.




Figure 8: Longitudinal Phase Space plot of witness beam at $t = 1000$, run on 4 cores for 1000 particles, simulation using new improved interpolation scheme

Comparing this plot to Figure 4, one can see that the stepwise structure has been smoothed out. However, if one examines the area of the graph just before $\zeta = 5$, there appear to be some particles 'missing'. The area with the missing particles corresponds to the halo slice at the uppermost edge of a processor. This feature can possibly be explained by the fact that when the MPI routine inherits the field information for these halo cells from the upper processor in the new routine, this information was possibly calculated for the *previous* timestep. This is due to the asynchronous nature in which processes in HiPACE run. Therefore, this new scheme will be accurate as long as the fields at the grid points do not vary much from timestep to timestep. When this code was run for 1000 timesteps, It would appear that these errors accumulated to result in this empty region of the graph. It is also possible that this region can be explained by a fault within the MPI routine that passes field information back to the halo cells.

3.4 An Alternative Solution

While the above method appears to provide the ideal solution to the problem of longitudinal interpolation, the inaccuracies that appear over longer time periods suggest the possible need for an alternative working solution. A simple way to achieve this was to combine the success of the trilinear interpolation scheme when run in series, with the original zeroth-order scheme. This was done by implementing the trilinear interpolation scheme across a parallelised configuration, but with an if-loop enforcing the condition

that if the particle in question is in the right-most slab of a processor, the program reverts to zeroth-order interpolation for that particle. Figure 9 below shows the same graph plotted with such a scheme in place.




Figure 9: Longitudinal Phase Space plot of witness beam at $t = 100$, run on a single core for 10000 particles with longitudinal interpolation scheme 'turned off' in the final cell of each processor

While it is obvious that the above plot is still unphysical at the boundaries between processors where one can see the block-like structure, it is no less accurate than in the plots using the original HiPACE interpolation scheme, and throughout the rest of the processes, interpolation occurs with a higher degree of accuracy than before.

Note: the axes in the above plot should carry the same labels as all of the previous plots, and the different thickness of the beam is due to larger deviations in the particles' initial positions and energies than in previous plots.

4 Summary and Conclusions

The field interpolation scheme currently in use by HiPACE was examined, and two possible solutions to the lack of interpolation in the longitudinal direction were proposed and trialled. The first, which involved use of MPI to pass field values between adjacent processors for use in a trilinear interpolation scheme yielded promising results (Figure 7, Figure 8), but there are still some issues with the interpolation to particles in the halo cells which become apparent when the simulation is let run for a large number of time steps. More tests need to be run with this scheme before it is clear what is causing these issues, and whether a scheme of this sort is a viable replacement for the current one or not.

The second alternative scheme certainly is a viable replacement for the current zeroth-order interpolation scheme. It successfully eliminates the unphysical stepwise structure in all but one slab of each process. While it is not an ideal model of what would physically be observed, it is a better approximation, and in the field of numerical methods, Such 'better approximations' are what we constantly strive for.

5 Acknowledgements

I would like to thank the organisers of the summer student program for giving the opportunity to undertake this project. I would also like to thank the FLA group for welcoming me into their number, but most of all, I would like to extend immense gratitude to my supervisor, Alexander Aschikhin, who over the course of 8 weeks tutored and guided me through the many obstacles I encountered, and who put up with my constant questions, most of which could be answered with 'Did you remember to compile before running?'.

References

- [1] The FLASHForward Facility at DESY *A. Aschikhin et al* 2007
- [2] Energy Doubling of 42 GeV Electrons in a Metre-scale Plasma Wakefield Accelerator *I. Blumenfeld et al* 2007
- [3] Particle-in-Cell Simulation of Plasmas - A Tutorial *P. L. Pritchett* 2003
- [4] Theoretical and Numerical Studies on the Transport of Transverse Beam Quality in Plasma-Based Accelerators *T. J. Mehrling* 2014
- [5] Principles and Applications of Compact Laser-PLasma Accelerators *V. Malka et al* 2008
- [6] HiPACE: A Quasi-Static Particle-in-Cell Code *T. Mehrling et al* 2014