VORPAL: A COMPUTATIONAL TOOL FOR THE STUDY OF ADVANCED ACCELERATOR CONCEPTS

C. Nieter* and J.R. Cary, University of Colorado at Boulder, Boulder, CO 80309, USA

Abstract

VORPAL, a fully object-oriented, dimension-free plasma simulation code, now has matured and is being used to study advanced accelerator concepts, particularly Laser Wake Field Acceleration. Both fluid and PIC models exist in the code and they can be used independently or in conjunction to perform hybrid simulations. VORPAL has a moving window and the ability to launch laser pulses from boundaries or initialize pulses in the plasma. VORPAL can accommodate a variety of different issues in LWFA, including chirped pulses, propagation of pulse in plasma channels, and optical injection of particle beams.

THE VORPAL FRAMEWORK

The VORPAL code framework make strong use of object oriented design to provide a variety of features included the ability to set the dimension of the simulation at run time, a general domain decomposition, and the the availability of multiple models for both the electromagnetic fields and the plasma. A full description of the code and the methods used in its development can be found in reference [1].

Multi-Dimensional

By templating most of the classes in the VORPAL libraries over dimension and floattype we are able to support simulations of one, two, and three dimensions and of either float or double precision with one code base. One of the principal challenges in writing a multi-dimensional code is the field updates. Normally the fields would be stored as multi-dimensional arrays and the updates would be done with nested loops. This can not be done for an arbitrary dimensional code since the dimension and hence the number of nested loops is not known at the outset.

To solve the problem of field updates, a combination of recursion and template meta-programming is used. The field data is stored internally in a one dimensional array. A generalization of an iterator is used which understands how the one dimensional index of the field data relates to the corresponding position in the multi-dimensional grid. These iterators have methods to bump their position on the grid in any direction. Updater classes are then created which consist of a collection of iterators, usually representing a stencil for a finite difference approximation of the relevant differential equation. These updater classes have a *updateCell()* method which updates the field at that point. These updater classes are then "walked" across the grid by what are referred to as walker classes. The walker classes are templated over dimension and they call the next lowest dimension recursively. The walker class for the lowest dimension is then specialized to actually call the *updateCell()* method of the updater class.

Domain Decomposition

VORPAL is designed to run as both a serial code on stand alone workstations and as a parallel code on systems that run MPI. As part of our general philosophy of flexibility, we have developed a general three dimensional domain decomposition for VORPAL. Any decomposition that can be constructed from a collection of bricks is possible. With this general domain decomposition static load balancing can be and dynamic load balancing is in the final stages of development. To minimize overhead, VORPAL overlaps communication with computation as much as possible.

To achieve this general domain decomposition, we use the idea of a slab. A slab is an object whose sides are straight lines and whose corners are all right angles. So in one dimension a slab is a straight line, in two dimensions a slab is a rectangle and in three dimensions a slab is a brick. The intersection of any two slabs of the same dimension is another slab. Each domain is described by two slabs, one which is referred to as the physical region is all the cells for which the processor is responsible for updating. The other, which is referred to as the extended region is the physical region plus one layer of guard cells in each direction. The cells that a domain needs to send to a neighboring processor is simply the slab that is the intersection of the sending processors physical region with the receiving processors extended region.

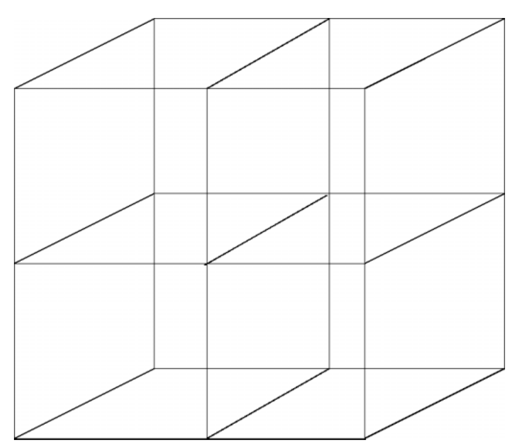

Figure 1: Standard 2D decomposition of a 3D cubic region.

This general domain decomposition allows for full load balancing. In Fig. 1 we see a standard 2D decomposition of a 3D cubic region. This decomposition is determined by

* nieter@colorado.edu

the two planes that cut across the region. To balance the load between the three different domains, the relative times between the first domain and the remaining three must be balanced. This requires three degrees of freedom but the position of the two planes only provide two. In Fig. 2 a more general decomposition that is possible using VORPAL has three movable planes providing the required number of degrees of freedom for load balancing to be done.

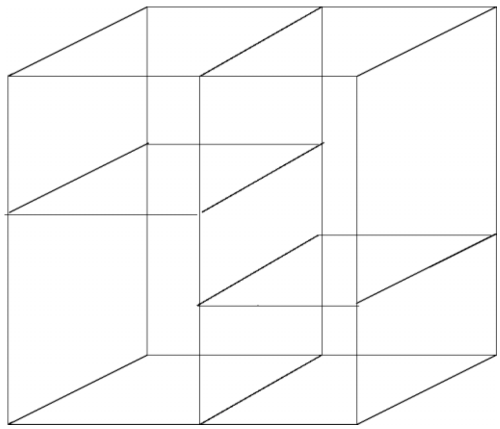

Figure 2: General 2D decomposition of a 3D cubic region.

Available Models

VORPAL flexible object oriented framework makes it possible to have multiple models for both the electromagnetic fields and the plasma in the same code. By interacting with the rest of the code through standard interfaces, different models can be used depending on the relevant physics being studied. Hybrid simulations can be done by using multiple models to represent different regions of the plasma or different species.

VORPAL has a finite difference finite time domain Maxwell solver for the electromagnetic fields based off the Yee mesh. Externally applied fields can be modeled using classes that produced profiles that varying in time or space according to some functional description. The Yee field can be combined with any number of externally applied fields. An electrostatic solver is currently under development.

The plasma particles can currently be modeled with either a cold fluid model or a particle-in-cell (PIC) model. The cold fluid model directly advects the fluid velocity rather than using flux transport to find the fluid momentum. This allows for regions of zero plasma density. The PIC model uses policy classes to provide for a variety of different particle dynamics, including a full relativistic Boris push, a non-relativistic electrostatic push, and free streaming particles.

APPLICATIONS TO LWFA

Several different features have been incorporated into VORPAL to make a useful tool in the study of Laser Wake Field Acceleration (LWFA) [4]. The laser wake field acceleration concept involve sending a high intensity short laser pulse into a plasma. The pondermotive force from the laser pulse blows out the plasma electrons creating a charge separation in the plasma. The electrons are pulled back by the ions and a plasma oscillation is created. This produces a wake field behind the laser pulse. These wake fields can generate field gradients on the of 100 GeV/m.

To study the propagation of the laser pulse and the generated wake field, the simulation region would have to be long enough to follow the pulse as it propagates through the plasma. This can become very expense computationally depending on long the pulse must be followed. However, the interesting physics occurs only in the region around the laser pulse itself. Therefore by including a moving window [2] [3] into VORPAL we can achieve considerable saving in compute time.

Since the various optical injection schemes that have been proposed required one or more pulses besides the pump pulse that generates the wake field, VORPAL has pulse launching boundaries that can launch multiple pulses of various polarizations and chirp from multiple boundaries. Plasma channels are one method of keeping the laser pulse collimated beyond the Rayleigh length in LWFA. VORPAL has the capacity to initialize a variety of plasma channels to study there effects on LWFA.

HYBRID SIMULATIONS

The LWFA is an example of a situation where a hybrid PIC/fluid simulation would be useful, since the beam and the plasma oscillations that generate the accelerating wake field are in some sense separate entities. The beam itself is best modeled by a collection of macro-particles, but by modeling the wake field with a fluid one can avoid the noise that is associated with modeling the bulk plasma with PIC and reduce the computational requirements of the simulation.

In the beat-wave (or colliding-pulse) injection scheme [5] for Laser Wake Field Acceleration (LWFA), three laser pulses are fired into the plasma. The first pulse, referred to as the pump pulse, is responsible for generating the wake field. The remaining two pulses are used to kick particles up to an energy that that puts them traveling in phase with the wake field for subsequent acceleration. This injection works by firing two counter propagating pulses with polarization perpendicular to the pump pulse. One pulse trails the pump pulse at a distance that puts it in the accelerating region of the wake field. The second injection pulse is launched from the other side of the plasma. Since injection pulses have a polarization perpendicular to the pump pulse, the second injection pulse passes through the pump pulse with minimal nonlinear interaction due to the plasma. When it reaches the other injection pulse the two pulses beat. This generates a short lived large electric field and a beat potential that inject the beam particles into the wake field.

To perform a hybrid simulation of this situation we model the front of the plasma with a collection of particles long enough that injection occurs with this region. Past this region the plasma is then modeled by a fluid. After injection the beam particles are traveling near the speed of light so they will remain in the simulation after the moving window is active. The remaining particles will leave the simulation as the window shifts the plasma, moving the wake field into the region where it will be modeled by the fluid. Although initially the computational cost of such a hybrid simulation is greater than a simple PIC simulation, once the moving window moves the bulk of the particles out of the simulation, only the beam particles are being updated. So the hybrid simulation reduces computational cost for long runs and reduces the noise in the simulation since the bulk of the plasma is being modeled by a fluid.

A simple hybrid simulation shows that VORPAL's ability to use multiple models to represent the plasma gives it the capacity to perform these hybrid simulations. The pump pulse and the right traveling colliding pulse are launched into vacuum from the left boundary in what we refer to here after as the x-direction and the right traveling pulse is created adiabatically within the plasma. All the pulses have a half cosine profile in the direction of propagation and are Gaussian in the transverse direction.

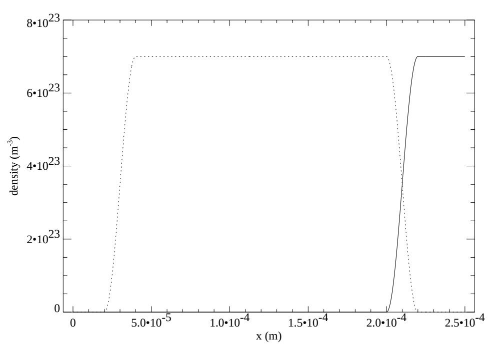

Figure 3: The initial distribution of particle and fluid densities used for a hybrid LWFA simulation. The dashed line is the density profile of the particles and the solid line is the density distribution for the fluid.

We take the rms length of the pump pulse, L_p to be half the plasma wavelength and the lengths of the two colliding pulses to be half a plasma wavelength. The wavelength of the two colliding pulses are slight detuned so they will beat. The forward moving pulse trails the pump by 55 μm measured center to center. This ensures that injection will occur at an accelerating and focusing region in the wake field. Fig. 3 shows the initial plasma density distribution. The plasma starts at zero and rises as a half cosine to the bulk density over 80 μm. The plasma is represented by particles for the next 100 μm so the injection pulses collide in within particles.. At this we have transition region of 20 μm where the particle density drops as the fluid density rises. After we have a constant fluid density.

In Fig. 4 we see the x-component of the relativistic velocity (γv) of the particles plotted against their x positions. A beam has clearly been formed and accelerated to an energy of approximately 14 MeV in less than half a millimeter. At this point in the simulation the initial particle region has been moved out of the simulation by the moving window, and only the particles traveling near the speed of light are present.

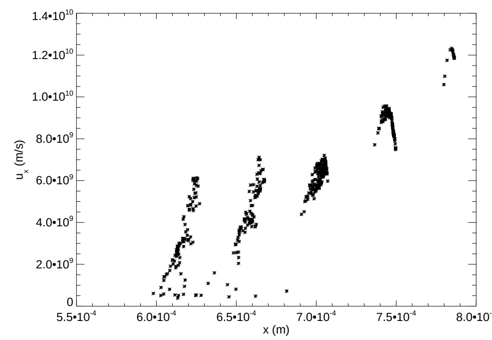

Figure 4: The relativistic velocity in the x-direction of the particles for the hybrid colliding pulse simulation plotted in the x-direction.

While such hybrid simulations are possible, we have found that they do not yet work well for highly asymmetric cells, as are required for laser pulses containing a large number of wavelengths. Asymmetric cells are known to be problematic in fluid numerics. Thus, the use of hybrid simulations for extreme cases appears to require algorithm development.

CONCLUSIONS

VORPAL is well suited as a computational code for studying LWFA. It has several features specificly developed with this problem in mind and due to its highly object oriented design it can change to accomodate new directions in research. A hybrid simulation of the colliding pulse injection scheme highlights some of ways that VORPAL can used to study LWFA.

REFERENCES

[1] C. Nieter, J. R. Cary, "VORPAL: a versatile plasma simulation code", to be published in J. Comp. Phys.

[2] Y. Pemper and B. Fidel and E. Heymana and R. Kastner and R. W. Ziolkowski, "Study of absorbing boundary conditions in the context of the hybrid ray-FDTD moving window solution" in The Thirteen Annual Review of Progress in Applied Computational Electrodynamics, ACES Symposium, 1997, pg. 17.

[3] C. D. Decker and W. B. Mori, *Phys. Rev. Lett.* **72**, 490 (1994).

[4] T. Tajima, J. M. Dawson, *Phys. Rev. Lett.* **43**, 267 (1979).

[5] E. Esarey, R. F. Hubbard, W. P. Leemans, A. Ting, P. Sprangle, *Phys. Rev. Lett.* **79**, 2682 (1997)