

# OSIRIS.FRAMEWORK: AN INTEGRATED TOOL FOR MODELING ASTRYPHYSICAL AND LABORATORY PLASMAS.

Ricardo A. Fonseca<sup>1,2\*</sup>

1: Centro de Física dos Plasmas  
Instituto Superior Técnico  
Universidade Técnica de Lisboa  
Av. Rovisco Pais, 1049-001 Lisboa, Portugal  
e-mail: ricardo.fonseca@ist.utl.pt web: <http://cfp.ist.utl.pt/>

2: Departamento de Ciências e Tecnologias da Informação  
Instituto Superior de Ciências do Trabalho e da Empresa  
Av. Forças Armadas, 1649-026 Lisboa, Portugal  
e-mail: ricardo.fonseca@iscte.pt web: <http://dcti-server.iscte.pt/>

**Keywords:** Computational Physics, Plasmas, Particle-In-Cell, Parallel Computing, SEMNI

**Abstract.** *We describe the `osiris.framework` [1], a general purpose, three-dimensional, fully relativistic, massively parallel, object oriented particle-in-cell code for the numerical simulation of astrophysical and laboratory plasmas, complemented by a set of specially designed visualization tools [2]. Developed in Fortran 95, the code runs on multiple platforms (Cray T3E, IBM SP, Beowulf, Mac clusters) and can be easily ported to new ones. Details on the code's capabilities are given. We discuss the object-oriented design of the code, the encapsulation of system dependent code and the parallelization of the algorithms involved. We also discuss the implementation of communications as a boundary condition problem and also load balancing, as well as other key characteristics of the code, such as the moving window, open-space and thermal bath boundaries, arbitrary domain decomposition, 1D (cartesian), 2D (cartesian and cylindrical) and 3D geometry, ion sub-cycling, tunnel and impact ionization and diagnostics. Finally results from three-dimensional simulations are presented, in connection with the data analysis and visualization infrastructure developed to post-process the scalar and vector results from PIC simulations.*

## 1 INTRODUCTION

Based on the highly nonlinear and kinetic processes that occur during high-intensity particle and laser beam-plasma interactions, we use particle-in-cell (PIC) codes [3, 4], which are a subset of the particle-mesh techniques, for the modeling of these physical

problems. In these codes the full set of Maxwell's equations are solved on a grid using currents and charge densities calculated by weighting discrete particles onto the grid. Each particle is pushed to a new position and momentum via self-consistently calculated fields. Therefore, to the extent that quantum mechanical effects can be neglected, these codes make no physics approximations and are ideally suited for studying complex systems with many degrees of freedom.

Achieving the goal of one to one, two and three dimensional modeling of laboratory experiments and astrophysical scenarios, requires state-of-the-art computing systems. The rapid increase in computing power and memory of these systems that has resulted from parallel computing has been at the expense of having to use more complicated computer architectures. In order to take full advantage of these developments it has become necessary to use more complex simulation codes. The added complexity arises for two reasons. One reason is that the realistic simulation of a problem requires a larger number of more complex algorithms interacting with each other than in a simulation of a rather simple model system. For example, initializing an arbitrary number of lasers or particle beams in 3D on a parallel computer is a much more difficult problem than initializing one beam in 1D or 2D on a single processor. The other reason is that the computer systems, e.g., memory management, threads, operating systems, are more complex and as a result the performance obtained from them can dramatically differ depending on the code strategy. Parallelized codes that handle the problems of parallel communications and parallel IO are examples of this. The best way to deal with this increased complexity is through an object-oriented programming style that divides the code and data structures into independent classes of objects. This programming style maximizes code reusability, reliability, and portability.

The goal of this code development project was to create a code that breaks up the large problem of a simulation into a set of essentially independent smaller problems that can be solved separately from each other. This allows individuals in a code development team to work independently. Object oriented programming achieves this by handling different aspects of the problem in different modules (classes) that communicate through well-defined interfaces.

This effort resulted in a new framework called OSIRIS, which is a fully parallelized, fully implicit, fully relativistic, and fully object-oriented PIC code, for modeling intense beam plasma interactions. Details of the status of OSIRIS several years ago can be found in [10]. The evolution of OSIRIS is due to the combined efforts of many people. In this paper we describe some of this evolution and the current status of OSIRIS.

## 2 DEVELOPMENT

The programming language chosen for this purpose was Fortran 90, mainly because it allows us to more easily integrate already available Fortran algorithms into this new framework that we call OSIRIS. We have also developed techniques where the Fortran 90 modules can interface to C and C++ libraries, allowing for the inclusion of other libraries

that do not supply a Fortran interface. Although Fortran 90 is not an object-oriented language *per se*, object-oriented concepts can be easily implemented [5, 6, 7] by the use of polymorphic structures and function overloading.

In developing OSIRIS we followed a number of general principles in order to assure that we were building a framework that would achieve the goals stated above. In this sense all real physical quantities have a corresponding object in the code making the physics being modeled clear and therefore easier to maintain, modify and extend. Also, the code is written in a way such that it is largely independent from the dimensionality or the coordinate system used, with much of the code reused in all simulation modes.

Regarding the parallelization issues, the overall structure allows for an arbitrary domain decomposition in any of the spatial coordinates of the simulation, with an effective load balancing of the problems in study. The input file defines only the global physical problem to be simulated and the domain decomposition desired, so that the user can focus on the actual physical problem and does not need to worry about parallelization details. Furthermore, all classes and objects refer to a single node (with the obvious exception of the object responsible for maintaining the global parallel information), which can be realized by treating all communication between physical objects as a boundary value problem, as described below. This allows for new algorithms to be incorporated into the code, without a deep understanding of the underlying communication structure.

### 3 DESIGN

#### 3.1 OBJECT-ORIENTED HIERARCHY

Figure 1 shows the class hierarchy of OSIRIS. The main physical objects used are particle objects, electromagnetic field objects, and current field objects. The particle object is an aggregate of an arbitrary number of particle species objects. The most important support classes are the variable-dimensionality-field class, which is used by the electromagnetic and current field classes and encapsulates many aspects of the dimensionality of a simulation, and the domain-decomposition class, which handles all communication between nodes.

Benchmarking of the code has indicated that the additional overhead from using an object oriented framework in Fortran 90 leads to only a 8% slowdown in speed.

#### 3.2 PARALLELIZATION

The parallelization of the code is done for distributed memory systems, and it is based on the MPI message-passing interface [13]. We parallelize our algorithms by decomposing the simulation space evenly across the available computational nodes. This decomposition is done by dividing each spatial direction of the simulation into a fixed number of segments ( $N_1, N_2, N_3$ ). The total number of nodes being used is therefore the product of these three quantities (or two quantities for the 2D simulations).

The communication pattern follows the usual procedure for a particle-mesh code [14].

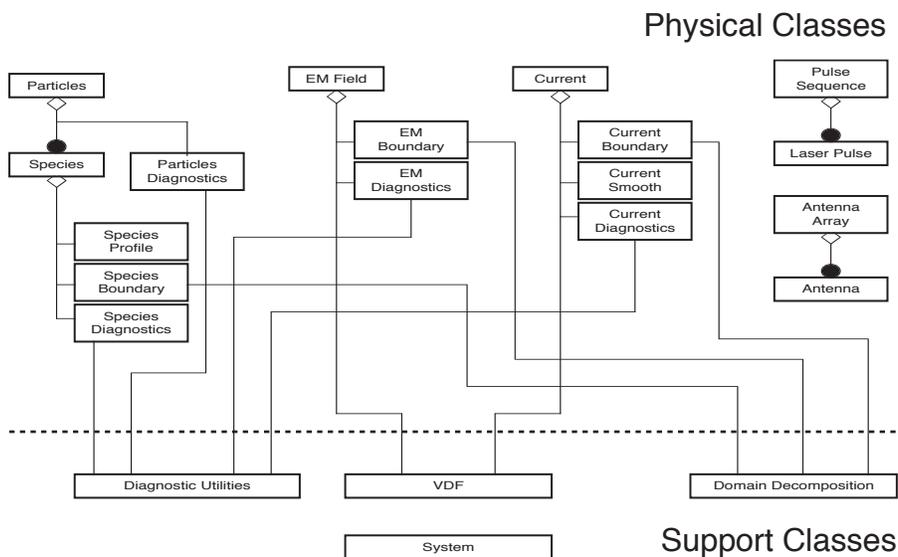


Figure 1: Osiris main class hierarchy

The grid quantities are updated by exchanging (electric and magnetic fields) or adding (currents) the ghost cells between neighboring nodes. As for the particles, those crossing the node boundary are counted and copied to a temporary buffer. Two messages are then sent, the first with the number of particles, and the second with the actual particle data. This strategy allows for not setting an *a priori* limit on the number of particles being sent to another node, while maintaining a reduced number of messages. Because most of the message are small, we are generally limited by the latency of the network being used. To overcome this whenever possible the messages being sent are packed into a single one, achieving in many cases twice the performance.

We also took great care in encapsulating all parallelization as boundary value routines. In this sense, the boundary conditions that each physical object has can either be some numerical implementation of the usual boundary conditions in these problems or simply a boundary to another node. The base classes that define grid and particle quantities already include the necessary routines to handle the latter case, greatly simplifying the implementation of new quantities and algorithms.

### 3.3 ENCAPSULATION OF SYSTEM DEPENDENT CODE

For ease in porting the code to different architectures, all code that is machine dependent is encapsulated in the system module. At present we have different versions of this module for running on the Cray T3E, the IBM SP, and for Macintosh clusters, running on both MacOS 9 (MacMPI [11]) and MacOS X (LAM/MPI [12]) clusters. The latter is actually a fortran module that interfaces with a POSIX compliant C module and should

therefore compile on most UNIX systems, allowing the code to run on PC-based (Beowulf) clusters. The MPI library has also been implemented on all these systems requiring no additional effort.

### 3.4 CODE FLOW

Figure 2 shows the flow of a single time step on a typical OSIRIS run. It closely follows the typical PIC cycle [4]. The loop begins by executing the diagnostic routines selected (diagnostics). It follows by pushing the particles using the updated values for the fields and depositing the current (advance deposit). After this step, the code updates the boundaries for particles and currents, communicating with neighboring nodes if necessary. A smoothing of the deposited currents, according to the specified input file, follows this step. Finally, the new values of the electric and magnetic field are calculated using the smoothed current values, and its boundaries are updated, again communicating with neighboring nodes, if necessary.

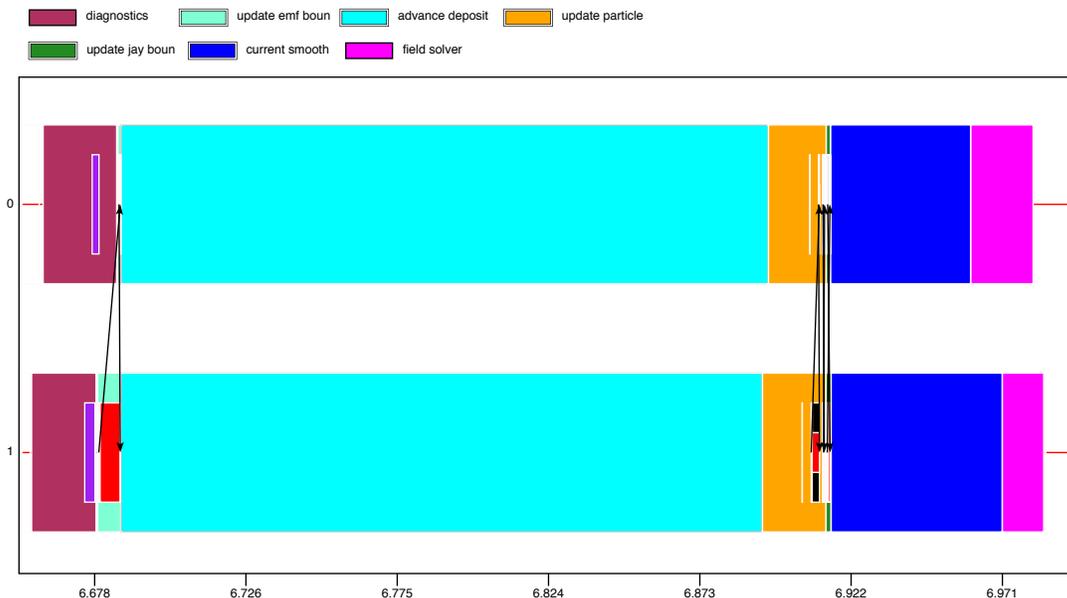


Figure 2: A typical cycle, one time step, in an OSIRIS 2 node run. The arrows show the direction of communication between nodes.

If requested, at the end of each loop, the code will write restart information, allowing the simulation to be restarted later on at this time step.

## 4 OSIRIS FRAMEWORK

The code is fully relativistic and it presently uses either the charge-conserving current deposition schemes from ISIS [8] or TRISTAN [9]. We have primarily adopted the charge-

conserving current deposition algorithms because they allow the field solve to be done locally, i.e., there is no need for a Poisson solve. The code uses the Boris scheme to push the particles, and the field solve is done locally using a finite difference solver for the electric and magnetic fields in both space and time.

In its present state the code contains algorithms for 2D and 3D simulations in Cartesian coordinates and for 2D simulations in azimuthally symmetric cylindrical coordinates, all of which with 3 components in velocity (i.e. both 2D modes are indeed  $2\frac{1}{2}$ D or 2D3V algorithms). The loading of particles is done by distributing the particles evenly on the cell, and varying the individual charge of each particle according to the density profile stipulated. Below a given threshold no particles are loaded. The required profile can be specified by a set of multiplying piecewise linear functions and/or by specifying Gaussian profiles. The initial velocities of the particles are set according to the specified thermal distribution and fluid velocity. The code also allows for the definition of constant external electric and magnetic fields.

The boundary conditions we have implemented in OSIRIS are: conducting, and Lindmann open-space boundaries for the fields [19], and absorbing, reflective, and thermal bath boundaries for the particles (the later consists of reinjecting any particle leaving the box with a velocity taken from a thermal distribution). Furthermore, periodic boundary conditions for fields and particles are also implemented.

This code also has a moving window, which makes it ideal for modeling high-intensity beam plasma interactions where the beam is typically much shorter than the interaction length. In this situation the simulation is done in the laboratory reference frame. Simulation data is shifted in the direction opposite to the motion of the window whenever  $nc\Delta t > \Delta x$  where  $n$  is the first integer for which this inequality is satisfied. Since this window moves at the speed of light in vacuum no other operations are required. The shifting of data is done locally on each node, and boundaries are updated using the standard routines developed for handling boundaries, thus taking care of moving data between adjacent nodes. The particles leaving the box from the back are removed from the simulation and the new clean cells in the front of the box are initialized as described above.

OSIRIS also incorporates the ability to launch EM waves into the simulation, either by initializing the EM field of the simulation box accordingly, or by injecting them from the simulation boundaries (e.g. antennas). Moreover, a subcycling scheme [20] for heavier particles has been implemented, where the heavier species are only pushed after a number of time steps using the averaged fields over these time steps, thus significantly decreasing the total loop time. We have also implemented particle sorting routines that allow for better use of CPU cache memory, resulting in performance gains of up to 40%.

A great deal of effort was also devoted to the development of diagnostics for this code that goes beyond the simple dumps of simulation quantities. For all the grid quantities envelope and boxcar averaged diagnostics are implemented; for the EM fields we implemented energy diagnostics, both spatially integrated and resolved; and for the particles

phase space diagnostics, total energy and energy distribution function, and accelerated particle selection are available. The output data uses the HDF [15] file format. This is a standard, platform independent, self-contained file format, which gives us the possibility of adding extra information to the file, like data units and iteration number, greatly simplifying the data analysis process.

## 5 VISUALIZATION AND DATA-ANALYSIS INFRASTRUCTURE

It is not an exaggeration to say that visualization is a major part of a parallel computing lab. The data sets from current simulations are both large and complex. These sets can have up to five free parameters for field data: three spatial dimensions, time and the different components (i.e.,  $E_x$ ,  $E_y$ , and  $E_z$ ). For particles, phase space has seven dimensions: three for space, three for momentum and one for time. Plots of y versus x are simply not enough. Sophisticated graphics are needed to present so much data in a manner that is easily accessible and understandable.

We developed a visualization and analysis infrastructure [2] based on IDL (Interactive Data Language). IDL is a 4GL language, with sophisticated graphics capabilities, and it is widely used in areas such as Atmospheric Sciences and Astronomy. It is also available on several platforms and supported in a number of systems, ranging from Solaris to the MacOS.

While developing this infrastructure we tried simplifying the visualization and data analysis as much as possible, making it user-friendly, automating as much of the process as possible, developing routines to batch process large sets of data and minimizing the effort of creating presentation quality graphics. We implemented a full set of visualization routines for one, two and three-dimensional scalar data and for two and three dimensional vector data. These include automatic scaling, dynamic zooming and axis scaling, integration of analysis tools, animation tools, and can be used either in batch mode or in interactive mode. We have also developed a comprehensive set of analysis routines that include scalar and vector algebra for single or multiple datasets, boxcar averaging, spectral analysis and spectral filtering, k-space distribution function, envelope analysis, mass centroid analysis and local peak tools.

## 6 RESULTS

The code has been successfully used in the modeling of several problems in the field of plasma based accelerators, and has been run on a number of architectures. Table 1 shows the typical push times on three machines, two supercomputers and one computer cluster.

We have also established the energy conservation of the code to be better than 1 part in  $10^5$ . This test was done in a simulation where we simply let a warm plasma evolve in time; in conditions where we inject high energy fluxes into the simulation (laser or beam plasma interaction runs) the results are better. Regarding the parallelization of the code, extensive testing was done on the old EPP cluster and also on the new ExPP cluster [21]

Table 1: Typical push time for three machines, in two and three dimensions. Values are in  $\mu\text{s}/\text{particle} \times \text{node}$

Machine	2D push time	3D push time
IBM SP	3.82	7.32
Cray T3E-900	5.77	11.2
EP <sup>2</sup> Cluster	4.96	9.82
ExPP Cluster	0.97	2.02

at IST in Lisbon, Portugal. We get very high efficiency, (above 91% in any condition), proving that the parallelization strategy is appropriate. It should be noted that the new cluster (20 dual G5 2.0 GHz Apple XServe G5, Gigabit Ethernet) is substantially faster than the IBM SP or the Cray T3E for the class of problems that it can run. However, for larger problems requiring up to 2048 cpus traditional supercomputers or large scale/fast interconnect clusters still provide the best results. One example of a three-dimensional modeling of a plasma accelerator is presented on figure 3. This is a one-to-one modeling of the E-157 Experiment [16] done at the Stanford Linear Accelerator Center, where a 30 GeV beam is accelerated by 1 GeV. The figure shows the Lorentz forces acting on the laser beam e.g.  $\mathbf{E} + \mathbf{z} \times \mathbf{B}$ , where  $\mathbf{z}$  is the beam propagation direction, and we can clearly identify the focusing /defocusing and accelerating/decelerating regions

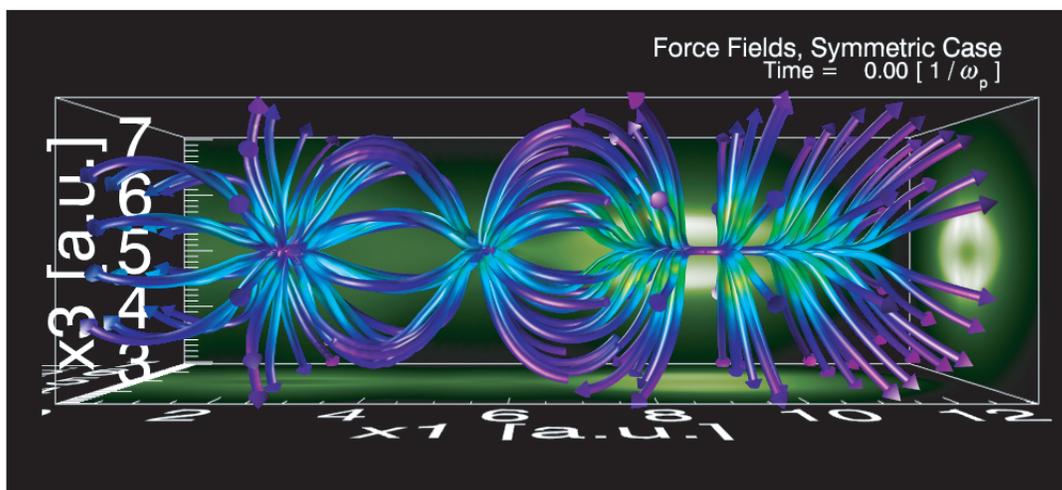


Figure 3: Force field acting on the 30 GeV SLAC beam inside a plasma column.

Another example of the code capabilities is the modeling of the Laser Wakefield Accelerator (LWFA). In the LWFA a short ultrahigh intensity laser pulse drives a relativistic

electron plasma wave. The wakefield driven most efficiently when the laser pulse length  $L = c\tau$  is approximately the plasma wavelength  $\lambda_p = 2\pi c/\omega_p$  - Tajima-Dawson mechanism [17]. Figure 4 shows the plasma wave produced by a 800 nm laser pulse with a normalized vector potential of 2.16, corresponding to an intensity of  $10^{19}\text{W}/\text{cm}^2$  on focus, and a duration of 30 fs, propagating in an underdense plasma.

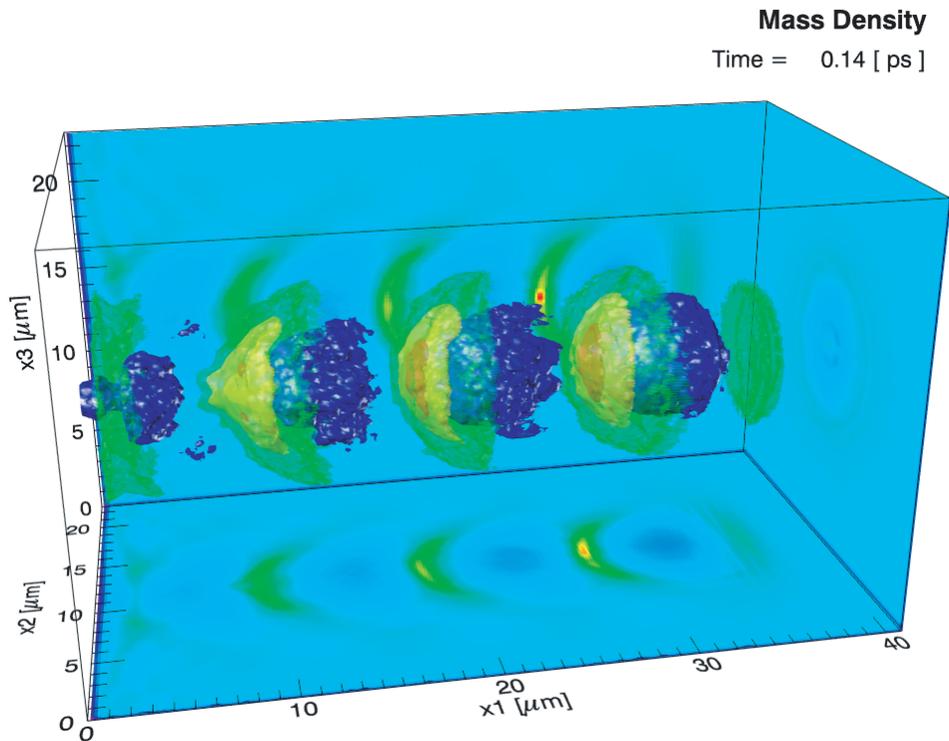


Figure 4: Plasma Wave produced in the LWFA. Isosurfaces shown for values of 0.5, 1.2, 2.0 and 5.0 normalized to the plasma background density.

## 7 FUTURE WORK

In summary, we have presented the OSIRIS framework for modeling plasma based accelerators. This is an ongoing effort; future developments will concentrate on the implementation of true open-space boundaries [18], ionization routines and dynamic load balancing. Regarding the visualization and data analysis infrastructure, a Web-Driven visualization portal will be implemented on the near future, allowing for efficient remote data analysis on clusters.

## 8 ACKNOWLEDGEMENTS

We are grateful to Prof. J. M. Dawson for his constant support, deep insights and inspiration. The OSIRIS framework is the result of the collaboration of a large number of groups. The author would like to acknowledge the collaboration of J.C.Adam, V. K. Decyk, S.Deng, R. G. Hemker, T.Katsouleas, W. Lu, S. Lee, M. Marti, S. F. Martins, W. B. Mori, C.Ren, L. O. Silva, J. Tonge, F. S. Tsung, and M. Tzoufras. This work was supported by DOE(USA), under grants number DE-FC02-01ER41179, DE-FG03-98DP00211, and DE-FG03-92ER40727, NSF (USA), under grant number Phy-0078508, FLAD, GULBENKIAN, and by FCT (Portugal) under grants PESO/P/PRO/40144/2000, PESO/P/INF/40146/2000, CERN/P/FIS/40132/2000, POCTI/33605/FIS/2000, ESO/FAT/43741, and CERN/FAT/43743.

## REFERENCES

- [1] R. A. Fonseca *et al.*, OSIRIS: A Three-Dimensional, Fully Relativistic Particle in Cell Code for Modeling Plasma Based Accelerators”, *Lect. Not. Comp. Sci.* Vol. **2331**, pp. 342-351, (2002).
- [2] R. A. Fonseca, Visualization and Data Analysis of Computer Simulations”, *Proceedings of ISSS-7*, pp. 67-70 (2005).
- [3] J. M. Dawson, Particle simulation of plasmas. *Rev. Mod. Phys.*, Vol. **55**, pp. 403-447, (1983).
- [4] C. K. Birdsall and A. B. Langdon, *Plasma physics via computer simulation*. Bristol, UK: Adam Hilger, xxvi+479 pp, (1991).
- [5] V. K. Decyk, C. D. Norton, and B. K. Szymanski, How to express C++ concepts in Fortran 90. *Sci. Prog.*, Vol. **6**, pp. 363-, (1998).
- [6] V. K. Decyk, C. D. Norton, and B. K. Szymanski, How to support inheritance and run-time polymorphism in Fortran 90. *Comp. Phys. Com.*, No. **115**, pp. 9-17, (1988).
- [7] M. G. Gray and R. M. Roberts, Object-Based Programming in Fortran 90. *Comp. Phys.*, Vol. **11** pp. 355-361, (1997).
- [8] R. L. Morse and C. W. Nielson, Numerical simulation of the Weibel instability in one and two dimensions. *Phys. Fluids*, Vol. **14**, pp.830-40, (1971).
- [9] J. Villasenor and O. Buneman, Rigorous charge conservation for local electromagnetic field solvers. *Comp. Phys. Comm.*, Vol. **69**, (1992).
- [10] R. G. Hemker, Particle-In-Cell Modeling of Plasma-Based Accelerators in two and three dimensions. Phd. Thesis, UCLA, (2000).

- [11] V. K. Decyk and D. E. Dauger, How to Build an AppleSeed: A Parallel Macintosh Cluster for Numerically Intensive Computing. *Proceedings of ISSS-6* (2001); also at <http://exodus.physics.ucla.edu/appleseed/appleseed.html>
- [12] <http://www.lam-mpi.org/>
- [13] Message Passing Interface Forum.: MPI: A message-passing interface standard. *Int. J. Supercomp. App.*, Vol. **8**, (1994).
- [14] W. Gropp, E. Lusk, and A. Skjellum, *Using MPI*. MIT Press, xxii+371 pp. (1999).
- [15] <http://hdf.ncsa.uiuc.edu/>
- [16] P. Muggli, *et al.*, *Nature*, Vol. **411**, p. 43, (2001).
- [17] T. Tajima and J. M. Dawson, Laser Electron Accelerator, *Phys. Rev. Lett.*, Vol. **43**, pp. 267-270, (1979).
- [18] J. L. Vay, A new Absorbing Layer Boundary Condition for the Wave Equation. *J. Comp. Phys.*, No. **165**, pp. 511-521, (2000).
- [19] Lindmann, E. L.: Free-space boundary conditions for the time dependent wave equation. *J. Comp. Phys.*, No. **18**, pp. 66-78, (1975).
- [20] J. C. Adam, A. G. Sereniere, and A. B. Langdon: Electron sub-cycling in particle simulation of Plasmas *J. Comp. Phys.*, No. **47**, pp. 229-244, (1982).
- [21] <http://cfp.ist.utl.pt/golp/epp/>