

COMPASS, the COMMunity Petascale project for Accelerator Science and Simulation, *a broad computational accelerator physics initiative*

J R Cary^{1,9}, P Spentzouris², J Amundson², L McInnes³, M Borland³, B Mustapha, B Norris³, P Ostroumov³, Y Wang³, W Fischer⁴, A Fedotov⁴, I Ben-Zvi⁴, R Ryne⁵, E Esarey⁵, C Geddes⁵, J Qiang⁵, E Ng⁵, S Li⁵, C Ng⁶, R Lee⁶, L Merminga⁷, H Wang⁷, D L Bruhwiler¹, D Dechow¹, P Mullowney¹, P Messmer¹, C Nieter¹, S Ovtchinnikov¹, K Paul¹, P Stoltz¹, D Wade-Stein¹, W B Mori⁸, V Decyk⁸, C K Huang⁸, W Lu⁸, M Tzoufras⁸, F Tsung⁸, M Zhou⁸, G R Werner⁹, T Antonsen¹⁰, T Katsouleas¹¹

¹Tech-X Corporation, Boulder, CO 80303

²Fermi National Accelerator Laboratory, Batavia IL, 60510

³Argonne National Laboratory, Argonne, IL 60439

⁴Brookhaven National Laboratory, Upton, NY 11973

⁵Lawrence Berkeley National Laboratory, Berkeley, CA 94720

⁶Stanford Linear Accelerator Center, Menlo Park, CA 94025

⁷Thomas Jefferson National Accelerator Facility, Newport News, VA 23606

⁸University of California at Los Angeles, Los Angeles, CA 90095

⁹University of Colorado, Boulder, CO 80309

¹⁰University of Maryland, College Park, MD 20742

¹¹University of Southern California, Los Angeles, CA 90089

E-mail: cary@txcorp.com

Abstract. Accelerators are the largest and most costly scientific instruments of the Department of Energy, with uses across a broad range of science, including colliders for particle physics and nuclear science and light sources and neutron sources for materials studies. COMPASS, the Community Petascale Project for Accelerator Science and Simulation, is a broad, four-office (HEP, NP, BES, ASCR) effort to develop computational tools for the prediction and performance enhancement of accelerators. The tools being developed can be used to predict the dynamics of beams in the presence of optical elements and space charge forces, the calculation of electromagnetic modes and wake fields of cavities, the cooling induced by comoving beams, and the acceleration of beams by intense fields in plasmas generated by beams or lasers. In SciDAC-1, the computational tools had multiple successes in predicting the dynamics of beams and beam generation. In SciDAC-2 these tools will be petascale enabled to allow the inclusion of an unprecedented level of physics for detailed prediction.

1. Introduction

The just-funded Community Petascale Project for Accelerator Science and Simulation (COMPASS) is a computational initiative to provide predictive and analysis capability through

use of the large computing facilities extant and future to the accelerator physics community of the Department of Energy. Accelerators are the largest experimental instruments supported by the Department of Energy, with costs upwards of \$1B for construction, and they are fundamental to the research carried out by the Offices of High Energy Physics (HEP), Nuclear Physics (NP), and Basic Energy Sciences. Hence, COMPASS is a multi-office project to support the needs of all three of those offices.

The main thrust areas of this project are:

- Beam Dynamics (BD): The evolution of beams through beam optics systems, including self forces and other forces of interaction.
- Electromagnetics (EM): The design and optimization of electromagnetic structures to maximize acceleration while minimizing deleterious effects of wakefields, other modes, surface heating, and multipactoring.
- Electron Cooling (EC): Simulate the dynamical friction of ions in relativistic copropagating electron beams to assess the potential luminosity increase for future heavy ion and ion-electron colliders.
- Advanced Acceleration (AA): Creation of large accelerating fields in plasma to accelerate beams to large energies in very short distances.

In addition, the project includes enabling activities in the areas of computer science (component technologies) and applied math (solvers, optimization, algorithms, and visualization). Of course, these areas are not in isolation; beam dynamics is affected by wakefields computed by electromagnetics, particle-in-cell methods are common to beam propagation in cavities and advanced acceleration, electron cooling is needed for heavy ion beams because of emittance growth due to intra-beam scattering and beam-beam interactions. Hence, the project is integrated to ensure that the results of all areas are available to all.

In order to bring together the expertise needed to address the many areas of computational accelerator physics, COMPASS has participation from 11 institutions, including one research corporation, six national laboratories, and four universities. These 11 institutions bring a wide range of expertise to this project including experience in beam optics, space-charge forces, electromagnetics computations and analysis, cavity design, two-body forces and different approaches to modeling them, self-consistent interactions with electromagnetics in plasmas, modern computer science technologies, and applied mathematics.

The COMPASS approach to computation is to avoid duplication but keep complementarity and duality for verification. Thus, the electromagnetics effort has both an unstructured and structured mesh/embedded boundary approaches to obtaining accurate answers. In electron cooling, there is one approach using two-body force summation and another using a velocity-space collision operator. In advanced acceleration, both full-physics, particle-in-cell methods and reduced models are being used. This not only provides the ability to cross-compare computations for verification, but it also allows multiple types of computations providing results with varying amounts of detail.

This paper is organized to provide a description of each of the above application areas with some discussion of past accomplishments and future directions. This is followed by a discussion of enabling technologies and then some discussion of physics integration. We conclude with a summary and discussion of future directions.



Figure 1. Aerial view of the Tevatron.

2. Beam Dynamics (BD)

Beam Dynamics is the study of the evolution of beams through a variety of beam optics systems including the effects of external forces from beamline elements, space-charge forces, and other the forces of interaction. Questions addressed include: How does a beam spread apart after repeated interactions with its opposite in a collider? How does the buildup of beam intensity, with the concomitant space charge forces of the beam, affect the dynamics?

Figure 2 illustrates a type of calculation in this area. In this case, the linac injector to the Fermilab booster operates at a frequency that is five times higher than that of the Booster (a circular machine that accelerates the beam to a mid range). As a consequence, one must merge linac bunches into a single booster bunch, ultimately bringing them together, such that space-charge forces come into play. SciDAC-1 led to the development of the capability of modeling space charge by computing the beam-frame electrostatic potential on a $33 \times 33 \times 257$ grid. The beam was modeled by 5M particles, and the full multiturn injection process, including beam optics, was computed. Problems like these tend to scale to an intermediate range of processors, just around 100.

As another example, the BeamBeam3D computational application is a parallel application able to model colliding beams including the effects of finite crossing angle and long-range forces. It was used in the first million-particle, million-turn simulation of the LHC, and has been used for large-scale modeling of colliding beams at the Tevatron, LHC, and RHIC.

An important part of this project is finding problems with designs, ideally before they are built. For example, modeling with PARMELA and ELEGANT [Borland 2001] predicted a CSR-driven microbunching instability in the LCLS design [Borland 2007]. By virtue of integration with the GENESIS FEL model, it was possible to assess the impact on FEL output, which was so severe that the LCLS was redesigned. Next-generation light sources, such as LCLS and proposed energy recovery linacs, need enhanced, high-resolution modeling of the gun and early injector, as well as efficient, seamless integration of modeling of the accelerator and radiation producing devices. Both ELEGANT and IMPACT have been used for LCLS-related modeling. IMPACT was recently used for large-scale studies of the microbunching instability with simulations performed using up to 1 billion particles [Qiang 2007]. In regard to ELEGANT, parallelization and improvement of CSR and space charge algorithms in ELEGANT would be beneficial to a wide community of researchers who use this code already for such designs. We also need reliable, on-demand, high-fidelity simulations that are integrated into the control room, providing tools to rapidly understand and tune complex new accelerators.

3. Electromagnetics (EM)

The costs associated with the construction and cooling of accelerating cavities will be a dominant part of the cost of the International Linear Collider. Hence, optimization of cavity designs to minimize losses (e.g., due to surface resistance or multipactoring) while maximizing accelerating field will have a large impact on the cost of the ILC. Of course, improvements in cavities have a general impact on the field of accelerator physics, as electromagnetic cavities are central to any accelerator.




Figure 2. Five microbunches coalescing into one as the beam is injected into the FNAL Booster from the LINAC.




Figure 3. The simple Tesla cavity without end groups.

Petascale computing will allow electromagnetics computations of unprecedented size and fidelity. Without high-performance computation, it would have been typical to carry out serial or moderately parallel computations of single cavities in the simplest form – without including the end groups: the higher-order mode (HOM) couplers and the lower-order mode (LOM) couplers. An example of such a simple cavity is shown in Fig. 3. However, HOM couplers (for accelerating cavities) and LOM couplers [for, e.g., crab (steering) cavities] are critical parts of the devices, as without them, unwanted electromagnetic modes can build up in the cavities, ultimately causing the beams to break up or be lost through other mechanisms.

The presence of these end groups results in a need for HPC resources if the modeling is carried out in the most direct way. For example, the 9-cell Tesla superconducting RF (SRF) cavities are 1.3 m long, and the diameter is roughly 0.2 m, while typical gaps within the HOM coupler are of order 0.5 mm. Thus, a direct, constant mesh would have to have $1.3 \times 0.2^2 / 0.0005^3 = 4 \times 10^9$ cells. With 3 field components per cell, that amounts to of the order of 10^9 degrees of freedom - taking one well into the domain of HPC.

COMPASS is using two different approaches to electromagnetics. The approach primarily funded in SciDAC-1 was to represent the fields using finite elements on unstructured meshes. This is one natural approach, as one can choose the finite elements to conform to the surface of the cavity, and adaptive mesh techniques allow one to pack modes in regions where higher resolution is needed. The second approach, started at the end of SciDAC-1 and now a part of SciDAC-2, is to use regular grids with embedded boundaries. This *cut-cell* approach allows one to propagate particles through the system easily, as one need not use connectivity information to obtain the fields, which are instead obtained by modular arithmetic.

In SciDAC-1, the unstructured mesh approach was focused on obtaining the eigenmodes and frequencies for an 8-cavity cryomodule of the ILC. A typical calculation was to model this with 20 M degrees of freedom. A typical computation of one mode required one hour on 1024 (POWER-3) CPUs. A goal for SciDAC-2 is to enlarge these runs to compute an entire RF unit. This is estimated to require over 200 M degrees of freedom.

The structured mesh with embedded boundaries is based on the scheme of Dey and Mittra [Dey 1997]. Unlike stairstepped boundaries, which include or exclude entire mesh cells, embedded boundaries can simulate partial cells. Compared to stairstep boundaries, the Dey-




Figure 4. Tesla cavity showing the end groups for powering the cavity and also for HOM couplers.

Mittra scheme reduces the error by one order, so that the overall error (in, e.g., mode frequency) is second-order in the mesh cell size. For example, finding a cavity frequency to an accuracy of 1 part in 10^4 using stairstepped boundaries would require on the order of 10^4 mesh cells in each dimension, or 10^{12} cells total, while using embedded boundaries would require on the order of 10^2 cells in each dimension, or 10^6 cells total. Reducing the number of mesh cells needed also increases the maximum time step with which the electromagnetic fields can be advanced, further reducing computation time. Figure 5 shows the error in frequency of a spherical cavity mode, using stairstepped boundaries and embedded boundaries as implemented in VORPAL.

The development of a new algorithm [Werner 2007] to extract modes and frequencies was helpful for comparing stairstepped boundaries with embedded boundaries. To extract mode patterns and frequencies, the algorithm relies on the power of (massively parallel) time-domain algorithms along with standard small-scale linear algebra techniques, combining the mode-finding ability and accuracy of frequency-domain solvers with the scalability of time-domain codes.

4. Electron Cooling (EC)

Novel electron-hadron collider concepts, requiring orders of magnitude higher ion beam luminosities than present heavy ion colliders, are a long-term priority for the international nuclear physics community. Higher luminosity can only be achieved by some dissipative mechanism that reduces the effective phase space volume occupied by the ion beam. A promising technique for this purpose, known as electron cooling, propagates an overlapping electron beam with the same velocity as the ions, for a small fraction of the collider circumference, allowing the ions to give up some of their thermal kinetic energy via Coulomb collisions. This very brief and subtle interaction provides weak damping of the effective ion phase space volume (i.e. cooling), which accumulates turn by turn to successfully combat mechanisms that increase this volume, especially intra-beam scattering and perhaps beam-beam collisions. Direct simulation of electron cooling physics essentially involves the accurate beam-frame calculation of dynamical velocity drag and diffusion on a non-relativistic ion moving for a short time in a low-density electron distribution, in the presence of relevant external fields.

Simulations using a molecular dynamics (MD) approach, pushing particles with a 4th-order predictor-corrector algorithm that aggressively varies the time step to resolve close Coulomb collisions, have recently resolved a long-standing discrepancy between alternate analytical descriptions of the dynamical velocity drag for solenoid-based cooling systems [Fedotov 2006 PRSTAB, Fedotov 2006 NJP]. A modified MD algorithm, which accurately captures binary collisions with a semi-analytic model and uses 2nd-order operator-splitting techniques to include external fields [Bell 2006 AIP], has been used to verify a conjecture that the strong, coherent electron oscillations in a helical undulator magnet would only weakly decrease the velocity drag. The difficult problem of characterizing the negative effects of arbitrary magnetic field errors is a present concern [Bruhwiler 2005 PAC].

These algorithms were implemented in VORPAL and have to date been applied to small periodic 3D domains. The latter algorithm of Bell et al has been shown to scale with 60%




Figure 5. Dey-Mittra embedded boundaries greatly improve simulation accuracy.

efficiency on 96 opteron processors. Scaling to petascale hardware with much larger domains will be necessary to answer key questions, including: a) accurately modeling solenoid and undulator magnetic fields without the artificial assumption of transverse periodicity, b) using realistic ion densities to assess the potentially negative consequences of nearby ion trajectories, and c) including the full transverse extent of the overlapping electron and ion beams to observe the effect of transverse space charge fields on the electron wake and, hence, the dynamical friction. Moving to the petascale will require use of the electrostatic PIC algorithm, probably in a hybrid approach with MD techniques. VORPAL's electrostatic PIC algorithm has been shown to scale well up to several hundred processors [Messmer 2004 CompPhysComm] using the Aztec00 solver of the Trilinos library.

We have simultaneously pursued a Langevin approach for simulation of binary collisions via the Fokker-Planck equation, where the dynamic friction coefficient and diffusion tensor are calculated self-consistently for arbitrary phase space distributions, using the Rosenbluth potential (or Landau integral). These coefficients involve convolutions in 6D phase space, so the calculation is extremely challenging.

This approach was shown to work for a single species [Qiang 2000 Supercomputing]. Recent extensions to systems of two species of modest mass ratio has been demonstrated, showing correct thermal exchange. Near-term plans are to extend the technique to electrons and ions, with parameters relevant to the RHIC-II cooler, to include the effects of magnetic fields, to use a fast multipole method to calculate the Rosenbluth potential for arbitrary velocity distributions, to explore an efficient FMM-based technique to compute the F-P collisional operator, to implement an efficient FMM-based technique to deal with electron-ion mass disparity, and to integrate the F-P code with our other modules for multi-physics modeling.

5. Advanced Acceleration (AA)

In the second half of SciDAC-1, two significant breakthroughs occurred in advanced accelerator research and both were published in Nature. In the first, the Dream Beam issue (*cf* Fig. 6) of Nature was published. For the first time, a beam with a narrow energy spectrum was produced by laser-plasma interaction, with the beam being accelerated to or order 100 MV in 0.5 cm, corresponding to an acceleration gradient of 20 GeV/m, nearly three orders of magnitude greater than that achievable by conventional accelerators. In the second, results from the e-167 collaboration were published where it was reported that the energy of some electrons in a SLAC bunch were doubled from 42 to 84 GeV by a plasma wave wake driven by the electron beam itself in less than one meter.

These publications in Nature illustrated the importance of the US program in computational science. In the Dream Beam issue three papers [Mangles 2004, Geddes 2004, Faure 2004] were published simultaneously, one each from experimental groups in England, France, and the Lawrence Berkeley National Laboratory. Significantly, all three papers contained computational analyses of those experiments, with two out of those three being from US computational groups. Moreover, those two groups (the OSIRIS group centered at UCLA and the VORPAL group centered at Tech-X Corp and the University of Colorado) are the two groups funded in Advanced Acceleration by this SciDAC project. Hence, there has been significant scientific payoff by support of this area. In the recent publication [Blumenfeld 2007], computation also played an



Figure 6. Cover of the Dream Beam issue of Nature illustrating the density contours from a computation of self-injection in the LWFA acceleration scheme.

important role. Simulations using the code QuickPIC [Huang 2006] (developed by UCLA, U. Maryland, and USC) identified the physical process that saturates the energy gain.

This area of computation has made extremely effective use of High Performance Computation. Even though the parallelization is far from trivial, it has been accomplished very effectively, with both the OSIRIS and VORPAL computational applications scaling well to thousands of processors. With these two applications scaling so well, it would not be surprising to see some of the first petascale applications coming from this area.

One example of a simulation underway is that to model acceleration to 3 GeV in a 3 cm long capillary waveguide. Our methodology is to run these simulations first in 2D, for which the computational requirements are orders of magnitude smaller. Once we believe that the 2D simulation is properly set up to match experiments, we move to 3D simulations. In this case, the 3D simulations require 2-3 Mhours of time on the IBM SP3. However, the good scaling helps to ensure that we can make full use of the provided computational resources. In this case we are able to run on 4096 processors effectively.

In SciDAC-2 we intend to apply additional effort to reduced models - models that capture the essential physics through use of approximations. One example of this is the QuickPIC computational application. This model makes effective use of the fact that a laser pulse changes little in the time it takes to propagate its own length. Another model is Ponderomotive PIC, which uses averaging only over the short oscillation frequency of the laser field, and which is implemented within VORPAL. In both of these cases, new computational challenges are arising, as the equations are no longer advanced by purely explicit algorithms. These newer, implicit methods require global matrix methods to obtain good solutions. Nevertheless, with a spectrum of computational capabilities coming on line, COMPASS is poised to contribute significantly to scientific discovery in SciDAC-2.

6. Enabling Technologies (ET)

Mathematical techniques, advanced scalable numerical algorithms, and computational tools are important ingredients in COMPASS. While many of the mathematical and computational tools we will employ are relatively mature, we must advance their capabilities and develop new approaches to meet the petascale computational challenges facing the project. Building on the collaborations and successes in SciDAC-1, our team is working closely with the APDEC [Colella 2007], CSCAPES [Pothen 2007], ITAPS [Diachin 2007], TOPS [Keyes 2007], TASCs [Bernholdt 2007], and VACET [Bethel 2007] Centers for Enabling Technologies as well as with the IUUV [Ma 2007] and PERI [Lucas 2007] Institutes to deploy state-of-the-art developments and to advance application-specific modeling capabilities in accelerator physics.

A key topic under investigation is the development of novel and scalable algorithms for large-scale nonlinear eigenvalue problems arising in trapped-mode analysis of waveguide-loaded SRF cavities in an RF unit consisting of three cryomodules. Another area of work is exploring novel, robust, and scalable direct and iterative solvers for extremely large sparse linear systems arising in driven frequency computations, time-domain simulations, shape determination and shape optimization, and eigenvalue calculations. We are also collaborating with the centers listed above on meshing, optimization, particle methods, visualization, data analysis, and integration.

7. Integration

Software infrastructure for multi-physics accelerator modeling is a central part of COMPASS. This effort demands a high degree of team collaboration to manage the integration of a wide range of accelerator physics applications and numerical libraries. Our approach builds upon ongoing work on Synergia2 [Amundson 2006] and MaryLie/IMPACT [Ryne 2006] and will extend to other subdomains of the project. Using the Common Component Architecture [Armstrong 2007, Bernholdt 2006], we are devising common interfaces that encapsulate preexisting physics

modules, thereby facilitating code reuse, the incorporation of new capabilities, and the exploration of performance tradeoffs among various algorithms and implementations for different simulation scenarios and target petascale machines. We also are collaborating with TASCs, PERI, and TOPS to provide support for computational quality of service [Norris 2004] to help select, parameterize, and reconfigure COMPASS codes dynamically during runtime, with initial emphasis on collective effects calculations for Synergia.

8. Summary and conclusions

The COMPASS project for High-Performance Computations of accelerators has just started, with an organization that accounts for the natural domain areas (Beam Dynamics, Electromagnetics, Electron Cooling, and Advanced Acceleration) as well as having contributions from Enabling Technologies, such as solvers and component technology. Multi-physics modeling will be important as the multiple separate domain areas are brought together in an integrated simulation.

These computational efforts will have a large impact on the DOE portfolio of facilities, the vast majority of which involve particle accelerators. With costs only rising (the ILC will cost multiple \$B), the use of computation to optimize designs and find improved operating regimes will lead to significant cost savings with modest investment.

Acknowledgments

Support is acknowledged from the US Department of Energy under grants DE-FG02-04ER41317, DE-FG02-06ER84484, DE-AC02-05CH11231, DE-FG02-05ER84173. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098. It also used resources of the Livermore Computing Services at Lawrence Livermore National Laboratory.

References

- [1] Amundson J et al 2006 Synergia: An accelerator modeling tool with 3-D space charge *Journal of Computational Physics* **211** 229–248
- [2] Armstrong R et al 2007 *Common Component Architecture Forum* <http://www.cca-forum.org>
- [3] Bell G I and Bruhwiler D L and Fedotov A V and Litvinenko V and Busby R and Abell D T and Messmer P and Veitzer S and Cary J R 2006 *AIP Conf. Proc.* **821** 319
- [4] Bernholdt D et al 2006 A Component Architecture for High-Performance Scientific Computing *Int J High Performance Computing Applications* **20** 163–202
- [5] Bernholdt 2007 *Center for Technology for Advanced Scientific Component Software (TASCs)* <http://www.scidac.gov/compsci/TASCs.html>
- [6] Bethel E 2007 *Visualization and Analytics Center for Enabling Technologies (VACET)* www.vacet.org
- [7] Blumenfeld I, Clayton C E, Decker F-J, Hogan M J, Huang C-K, Ischebeck R, Iverson R, Joshi C, Katsouleas T, Kirby N, Lu W, March K A, Mori W B, Muggli P, Oz E, Siemann R H, Walz D, Zhou M 2007 *Nature* **445** 741
- [8] Borland M 2000 elegant: A flexible SDDS-Compliant Code for Accelerator Simulation, APS Light Source Note LS-287, September 2000.
- [9] Borland M et al 2002 Start-to-End Simulation of Self-Amplified Spontaneous Emission Free Electron Lasers from the Gun through the Undulator, *NIM A* **483**, 268-272 (2002).
- [10] Bruhwiler D L and Busby R and Abell D T and Veitzer S and Fedotov A V and Litvinenko V N 2005 *Proc. Part. Accel. Conf.* 2583
- [11] Colella P 2007 *Applied Partial Differential Equations Center (APDEC)* <http://www.scidac.gov/math/APDEC.html>
- [12] Dey S and Mittra R *IEEE Microw. Guided Wave Lett.* **7** 273 (1997)
- [13] Diachin L 2007 *Interoperable Technologies for Advanced Petascale Simulations (ITAPS)* www.itaps-scidac.org
- [14] Dimitrov D A, Bruhwiler D L, Cary J R, Geddes C G R, Giacone R E, Esarey E and Leemans W P 2007 *Phys. Plasmas* **14** 043105

- [15] Faure J, Glinec Y, Pukhov A, Kiselev S, Gordienko S, Lefebvre E, Rousseau J-P, Burgy F and Malka V 2004 *Nature* **431** 541
- [16] Fedotov A V and Bruhwiler D L and Sidorin A O and Abell D T and Ben-Zvi I and Busby R and Cary J R and Litvinenko V N 2006 *Phys. Rev. ST Accel. Beams* **9** 074401
- [17] Fedotov A V and Ben-Zvi I and Bruhwiler D L and Litvinenko V N and Sidorin A O 2006 *New J. Phys.* **8** 283
- [18] Geddes C G R, Toth Cs, van Tilborg J, Esarey E, Schroeder C B, Cary J R and Leemans W P 2005 *Phys. Rev. Lett.* **95** 145002
- [19] Geddes C G R, Toth Cs, van Tilborg J, Esarey E, Schroeder C B, Bruhwiler D, Nieter C, Cary J R and Leemans W P 2004 *Nature* **431** 538
- [20] Huang C-K, Decyk V K, Ren C, Zhou M, Lu W, Mori W B, Cooley J H, Antonsen T M Jr and Katsouleas T 2006 *J. Comp. Phys.* **217** 658
- [21] Keyes D 2007 *Towards Optimal Petascale Simulations (TOPS)* <http://www.scidac.gov/math/TOPS.html>
- [22] Li R 2007 these proceedings
- [23] Lucas R 2007 *Performance Engineering Research Institute* <http://www.peri-scidac.org>
- [24] Ma K-L 2007 *Institute for Ultrascale Visualization* <http://www.scidac.gov/viz/ultraviz.html>
- [25] Mangles S P D, Murphy C D, Najmudin Z, Thomas A G R, Collier J L, Dangor A E, Dival E J, Foster P S, Gallacher J G, Hooker C J, Jaroszynski D A, Langley A J, Mori W B, Norreys P A, Tsung F S, Viskup R, Walton B R, Krushelnick K 2004 *Nature* **431** 535
- [26] Messmer P and Bruhwiler D L 2004 *Comp. Phys. Comm.* **164** 118
- [27] Norris B et al 2004 Computational Quality of Service for Scientific Components *Proc Int Symposium on Component-Based Software Engineering (CBSE7)* Edinburgh, Scotland
- [28] Pothen A 2007 *Combinatorial Scientific Computing and Petascale Simulations (CSCAPES)* <http://www.scidac.gov/math/CSCAPES.html>
- [29] Qiang J and Ryne R and Habib S 2000 *Proc. Supercomputing 2000 Conf.*
- [30] Qiang J and Pogorelov I and Ryne R 2007 *Proc. 2007 Particle Accelerator Conference*
- [31] Ryne R 2006 MaryLie/IMPACT: A Parallel Beam Dynamics Code with Space Charge, <http://scidac.nersc.gov/accelerator/marylielayer.html>
- [32] Shen B, Li Y, Nemeth K, Shang H, Soliday R, Crowell R, Frank E, Gropp W and Cary J 2007 *Phys. Plasmas* **14** 053115
- [33] Werner G R and Cary J R 2007 *J. Comp. Phys.* in review