Final Scientific/Technical Report

Community Petascale Project for Accelerator Science and Simulation

Period of Performance: September 01, 2007 to August 31, 2012

For Grant DE-FC02-07ER41500

P.I.: Warren B. Mori
UCLA Departments of Physics and Astronomy and of Electrical Engineering
mori@physics.ucla.edu

Summary:

The UCLA Plasma Simulation Group is a major partner of the "Community Petascale Project for Accelerator Science and Simulation". This is the final technical report. We include an overall summary, a list of publications, and individual progress reports for each year.

During the past five years we have made tremendous progress in enhancing the capabilities of OSIRIS and QuickPIC, in developing new algorithms and data structures for PIC codes to run on GPUS and future many core architectures, and in using these codes to model experiments and in making new scientific discoveries. Here we summarize some highlights for which SciDAC was a major contributor:

Continued development of OSIRIS:
During the past five years we have significantly improved OSIRIS. We have added perfectly matched layers for transmitting electromagnetic waves, added options for higher order particle shapes, optimized the code so that it runs effectively on SIMD units such as the SSE (at 30% of peak speed for problems with good load balance), added dynamic load balancing in 1, 2, and 3D, added a hybrid OpenMP/MPI parallelization strategy, added options for both energy and momentum conserving force interpolation, and added a framework for using the ponderomotive guiding center option for pushing the particles and advancing the laser. OSIRIS was chosen to be part of ASCRs software effectiveness program in 2010/2011. Under this program OSIRIS was scaled to the full Jaguar machine and was shown to be able to run at >30% of the peak speed. In addition, we added the capability of using OSIRIS to model LWFA in a Lorentz boosted frame including having a moving antenna, adding additional field solvers, and current and field filters. We have also improved the beam initialization routine as well as incorporated the framework for implementing the ponderomotive guiding center approximation.

Continued development of QuickPIC:
During the past few years, we have also significantly improved QuickPIC. Key developments include a novel pipelining routine to improve the parallel scalability of QuickPIC. QuickPIC is based on the quasi-static approximation in which the drive beam

is assumed to remain static during the time it passes by a plasma particle. In the code, a drive beam passes by 2D slice of plasma. After the entire beam has passed through, the plasma fields from that slice are calculated and then used to advance the drive beam forward large distances (compared to a cell size). Since the 2D slice is not that large, then the routine can only scale to ~100 -1000 processors. However, once one part of the drive beam has passed by the slice this part can be advance forward before the entire beam has passed by. As a result multiple copies of the basic algorithm can be running simultaneously. This can increase the parallel scalability by roughly the number of longitudinal cells. Such a pipelining algorithm was implemented and QuickPIC now scales to at least 10,000 cores. Another key development was finding an field solver that requires less iterations in the predictor corrector loop in order to converge to known results from a full PIC codes such as OSIRIS. We experimented with different field solvers (different gauges) and reduced the number of iterations that are required from ~4-7 to only 1.

Development of PIC algorithms for GPUs and Many Core Platforms:
During the several years, we have developed key concepts and key codes for PIC to run on GPUs. This includes optimal data structures and algorithms. The optimal data structure has evolved as the hardware has evolved. For GPUs, the data is organized into tiles (several computational cells) that can fit into small shared memory. The data in each tile is computed by either a single thread or a block of threads depending on the architecture. The crucial component to enable efficient use of the GPU is an optimal particle reordering scheme. This scheme currently provides 940ps per particle (including push, deposit, and reordering) for an electrostatic code. In addition, we have integrated the tile decomposition with MPI decomposition for running on multiple GPUs.

Modeling LWFA using Lorentz boosted frames:
During the past five years, we have developed the capability to carry out LWFA simulations in a Lorentz boosted frame. Soon after this idea was published by Vay, we implemented a laser initialization and field smoothing routines. The field smoothing is necessary in order to mitigate a numerical instability that results when a plasma drifts across a grid. We used Lorentz boosted LWFA simulations to study in 3D what was possible with next generation lasers. We have recently developed a theory to predict the location of the instability in k-space which will help find mitigation strategies. We have recently implemented the framework for using a spectral field solve in a boosted frame simulation.

Modeling LWFA experiments:
We have used OSIRIS extensively to model experiments at UCLA and LLNL. These experiments have demonstrated self-guiding, self-guided acceleration, ionization induced self-injection, and two-stage acceleration using ionization induced injection. OSIRIS simulations have been critical to the success of these experiments and to the development of the theoretical threshold for ionization self-injection. The self-guiding experiments were motivated by previous OSIRIS simulations. We are currently modeling the two-stage acceleration concept for 500TW lasers in anticipation of experiments on such

systems in the next year. We also have been using OSIRIS to make discoveries regarding self-injection such as the use of magnetic fields to trigger the process.

Modeling PWFA experiments:

We have also used OSIRIS and QuickPIC extensively to study PWFA using electron, positron, and proton beams. We showed that PWFA of 25 GeV beams with 3x10^10 particles can accelerate 1x10^10 particles with ~50% efficiency while maintaining energy spreads less than 1%. We have also simulated planned and ongoing experiments at FACET. Such experiments include a two bunch experiment in which a second bunch will be accelerated in the wake of the first bunch while maintaining its emittance and energy spread. This experiment will be a milestone in plasma based accelerator research. We have also been studying wakefields generated by compressed or uncompressed proton bunches. The later are motivated by possible experiments at FNAL and CERN.

Publications:

1. Huang C., Lu W., Zhou M., Clayton C.E., Joshi C., Mori W.B., Muggli P., Deng S., Oz E., Katsouleas T., Hogan M.J., Blumenfeld I., Decker F.J., Ishebeck, Iverson R.H., Kirby N.A., Walz D., "Hosing instability in the blowout regime for plasma wakefield Acceleration," Physical Review Letters Vol. 99, No. 25 pp. 255001/1-4 December 2007.

2. Kniep S., Nagel S.R., Bellei C., Bourgeois, N., Dangor A.E., Gopal A., Heathcote R., Mangles S.P., Marques J.R., Maksimchuk A., Nilson P.M., Ta Phuoc K., reed S., Tzoufras M., Tsung F.S., Willingale L., Mori W.B., Rousse A., Krushelnick K., Najmudin A., "Observation of synchrotron radiation from electrons Acceleratorserated in a petawatt laser-generated plasma cavity," Physical Review Letters Vol. 100, No. 10 pp. 105006/1-4, March 2008.

3. Muggli P., Blue B.E., Clayton C.E., Decker F.J., Hogan M.J., Huang C., Joshi C., Katsouleas T., Lu W., Mori W.B., O'connell C.L., Siemann R.H., Walz D., "Halo formation and emittance growth of positron beams in plasmas", Physical Review Letters Vol. 101, No. 5, pp 055001/1-4, August 2008.

4. Vieira J., Fiuza F., Fonseca R.A., Silva L.O., Huang C., Lu W., Tzoufras M., Tsung F.S., Decyk V.K., Mori W.B., Cooley J., Antonsen T., "One-to-one full-scale simulations of laser wakefield Acceleration using QuickPIC," IEEE Transactions on plasma science Vol. 36, No. 4, pp. 1722-1727, August 2008.

5. Huang C.K., Tsung F.S., Mori W.B., "Quasi-static particle-in-cell simulation of the plasma wakefield afterburner," IEEE Transactions on plasma science Vol. 36, No. 4, pp. 1294-1295, August 2008.

6. Vieira J.F., Martins S.F., Fiuza F. Fonseca R.A., Silva L.O., Huang C., Lu W., Tzoufras M., Tsung F.S., Mori W.B., Cooley T., Antonsen T., "Three-dimensional structure of the laser wakefield Accelerator in the blowout regime," IEEE Transactions on plasma science Vol. 36, No. 4, pp. 1124-1125, August 2008.

7. Geddes C.G.R., Bruhwiler D.L., Cary J.R., Mori W.B., Vay J-L., Martins S.F., Katsouleas T., Cormier-Michel., E., Fawley W.M., Huang C., Wang X., Cowan B., Decyk V.K., Esarey E., Fonseca R.A., Lu W, Messmer P., Mullowney P., Nakamura K., Paul K., Plateau G.R., Schroeder C.B., Silva L.O., Toth Cs., Tsung F.S., Tzoufras M., Antonsen TJ. Vieira J., and Leemans W.P., "Computational studies and optimization of wakefield Accelerators," Journal of Physics: Conference Series Vol. 125 pp. 12002/1-11, August 2008

8. Wang S, Ischebeck R., Muggli P., Katsouleas T., Joshi C., Mori W.B., "Positron injection and Acceleration on the wake driven by an electron beam in a foil-and-gas plasma," Physical Review Letters, Vol. 101, No. 12, pp. 124801/1-4, September 2008.

9. Tzoufras M., Lu W., Tsung F.S., Huang C., Mori W.B., Katsouleas T., Vieira J., Fonseca R.A., Silva L.O. "Beam loading in the nonlinear regime of plasma-based Acceleration," Physical Review Letters, Vol. 101 No. 14 pp. 145002/1-4, September 2008.

10. Fonseca R.A., Martins S.F., Silva L.O., Tonge J., Tsung F.S., Mori W.B., "One-to-one direct modeling of experiments and astrophysical scenarios: pushing the envelope on kinetic plasma simulations," Plasma physics and controlled fusion, Vol. 50 No. 12 pp. 124034, December 2008.

11. Martins S.F. Vieira J., Fiúza F., Fonseca R.A., Huang C., Lu W., Trines R., Norreys P., Mori W.B., Silva L.O, "Numerical simulations of LWFA for the next generation of laser systems", AIP Conf. Proceedings Vol. 1086, pp. 285, January 2009.

12. Paul K., Huang C., Bruhwiler D. L., Mori W. B., Tsung F. S., Cormier-Michel E., Geddes C. G. R., Cowan B., Cary J.R., Esarey E., Fonseca R.A., Martins S.F., and Silva L.O., "Benchmarking the codes VORPAL, OSIRIS, and QuickPIC with Laser Wakefield Acceleration Simulations", AIP Conf. Proceedings Vol. 1086, pp. 315 January 2009.

13. Ralph J.E., Marsh, K.A., Pak, A.E. Lu W., Clayton C.E., Fang F., Tsung F.S., Mori W.B., and Joshi C., "Self-Guiding of Ultrashort Relativistically Intense Laser Pulses to the Limit of Nonlinear Pump Depletion," AIP Conf. Proceedings Vol. 1086, pp. 196 January 2009.

14. Feng B., Huang C., Decyk V., Mori W.B., Hoffstaetter G.H., Muggli P., and Katsouleas T., "Simulation Of Electron Cloud Effects On Electron Beam At ERL With Pipelined QuickPIC," AIP Conf. Proceedings Vol, 1086, pp. 340 January 2009.

15. Blumenfeld I., Clayton C.E., Decker F.J., Hogan, M.J., Huang C., Ischebeck R., Iverson R.H., Joshi C., Katsouleas T., Kirby N., Lu W., Marsh K.A., Mori W.B., Muggli P., Oz E., Siemann R.H., Walz D.R., and Zhou M., "Measurement of the Decelerating Wake in a Plasma Wakefield Accelerator," AIP Conf. Proceedings Vol. 1086, pp. 569 January 2009.

16. Gholizadeh Reza, Katsouleas Tom, Muggli Patric, and Mori Warren B., "Preservation of Ultra Low Emittances Using Adiabatic Matching in Future Plasma Wakefield-based Colliders," AIP Conf. Proceedings Vol. 1086, pp. 575 January 2009.

17. Wang X., Ischebeck, R., Muggli P., Katsouleas T., Joshi C., Mori W.B., and Hogan M.J., "Threshold for Trapping Positrons in the Wake Driven by a Ultra-relativistic Electron Bunch," AIP Conf. Proceedings Vol. 1086, pp. 586 January 2009.

18. Ralph J.E., Marsh K.A., Pak A.E., Lu W., Clayton C.E., Fang F., Mori W.B., Joshi C., "Self-guiding of ultra-short, relativistically intense laser pulses through underdense plasmas in the blowout regime," Physical Review Letters, Vol. 102 No. 17 pp. 175003/1-4 May 2009.

19. Wang X, Muggli P., Katsouleas T. Joshi C., Mori W.B., Ishebeck R., Hogan M.J., "Optimization of positron trapping and Acceleration in an electron beam driven plasma wakefield Acceleration," Physical Review Special Topics – Accelerators and Beams, Vol. 12 No. 5 pp. 051303/1-4 May 2009.

20. Kirby N., Blumenfeld I., Clayton C.E., Decker F.J., Hogan M.J., Huang C., Ischebeck R., Iverson R.H., Joshi C., Katsouleas T., Lu W., Marsh K.A., Martins S.F., Mori W.B., Muggli P., Oz E., Siemann R.H., Walz D.R., Zhou M., "Transverse emittance and current of multi-GeV trapped electrons in a plasma wakefield Accelerator," Physical Review Special Topics-Accelerators and Beams, Vol. 12 No. 5, pp. 051302 May 2009.

21. Tzoufras M., Lu W., Tsung F.S., Huang C., Mori W.B., Katsouleas T., Vieira J., Fonseca R.A., Silva L.O., "Beam loading by electrons in nonlinear wakes," Physics of Plasmas Vol. 16 No. 5, pp. 056705/1-12, May 2009.

22. Silva L.O., Fiuza F. Fonseca R.A., Martins J.L., Martins S.F., Vieira J.V., Huang C., Lu W., Tsung F.S., Tzoufras M., Mori W.B., "Laser electron Acceleration with 10PW lasers, "Comptes Rendus Physique Vol. 10 No. 2 pp. 167-175, April 2009.

23. Feng B., Huang C., Decyk V., Mori W.B., Muggli P., Katsouleas T., "Enhancing parallel quasi-static particle-in-cell simulations with a pipelining algorithm," Journal of Computational Physics Vol. 228 No. 15 pp. 5340-5348, August 2009.

24. Kniep S. Nagel S.R., Martins S.F., Mangles S.P.D., Bellei C., Chekhlov O., Clarke R.J., Delerue N., Divall E.J., Doucas G., Ertel K., Fiuza F., Fonseca R., Foster P., Hawkes S.J., Houker C.J., Krushelnick K., Mori W.B., Palmer C., Ta Phouoc K., Rajeev P., Schreiber J., Streeter M.J.V., Urner D., Vieira J., Silva L.O., Najmudin Z., "Near GeV Acceleration of electrons by a nonlinear plasma wave driven by a self-guided laser pulse," Physical Review Letters, Vol. 103 No. 3 pp. 035002/1-4, July 2009.

25. Huang C., An W., Decyk V.K., Lu W., Mori W.B., Tsung F.S., Tzoufras M., Morshed S., Antonsen T., Feng B., Katsouleas T., Fonseca R.A., Martins S.F., Vieira J., Silva L.O., Esarey E., Geddes C.G.R., Leemans W., Cormier-Michel E., Vay J-L., Bruhwiler D.L., Cowan B., Cary J.R., Paul K., "Recent results and future challenges for large scale particlein-cell simulations of plasma-based Accelerator concepts," Journal of Physics: Conference Series Vol. 180 pp. 012005, August 2009.

26. Vay J.L., Bruhwiler D.L., Geddes C.G.R., Fawley W.M., Martins S.F., Cary J.R., Cormier-Michel E., Cowan B., Fonseca R.A., Furman M.A., Lu W., Mori W.B., Silva L.O., "Simulating relativistic beam and plasma systems using an optimal boosted frame," Journal of Physics: Conference Series Vol. 180 pp. 012006, August 2009.

27. Froula D.H., Clayton C.E., Doeppner T., Fonseca R.A., Marsch K.A., Barty C.J., Divol L., Glenzer S.H., Joshi C., Lu W., Martins S.F., Michel P., Mori W.B., Palastro J.P., Pollock B.B., Pak A., Ralph J.E., Ross J.S., Siders C., Silva L.O., Wang T., "Measurements of the critical power for self-injection of electrons in a laser wakefield Accelerator," Physical Review Letters, Vol. 103, No. 21 pp. 215006/1-4 (Nov 2009).

28. An, W., Huang, C., Lu, W., Mori, W.B., Katsouleas, T.C., "Positron acceleration by using a particle beam-driven wake field in plasma," Proceedings of the Particle Accelerator Conference, Vancouver, B.C., May 4-8, 2009, pp. 3013-3015 (2009).

29. Martins, J., Fonseca, R.A., Martins, S., Silva, L.O., Joshi, C., Mori, W.B., "Emission of collimated x-ray radiation in laser-wakefield experiments using particle tracking in PIC simulations," Proceedings of the Particle Accelerator Conference, Vancouver, B.C., May 4-8, 2009, pp. 2958-2960 (2009).

30. Pak, A., Joshi, C., Marsh, K.A., Martins, S.F., Mori, W.B., "Investigation of Ionization Induced Trapping in a Laser Wakefield Accelerator," Proceedings of the Particle Accelerator Conference, Vancouver, B.C., May 4-8, 2009, pp. 3031-3033 (2009).

31. Lu, W., An, W., Huang, C., Joshi, C., Mori, W.B., Hogan, M. Raubenheimer, T., Seryi, A., "High transformer ratio PWFA for application on XFELs," Proceedings of the Particle Accelerator Conference, Vancouver, B.C., May 4-8, 2009, pp. 3028-3030 (2009).

32. Huang, C., An, W., Clayton, C., Joshi, C., Lu, W., Marsh, K., Mori, W.B., Tzoufras, M., Katsouleas, T., Blumenfeld, I., Hogan, M., Kirby, N., Raubenheimer, T., Seryi, A., Muggli, P., "Simulations of 25 GeV PWFA sections: path towards a PWFA linear collider," Proceedings of the Particle Accelerator Conference, Vancouver, B.C., May 4-8, 2009, pp. 3025-3027 (2009).

33. Seryi, A., Hogan, M., Pei, S., Raubenheimer, T., Tenenbaum, P., Katsouleas, T., Huang, C., Joshi, C., Mori, W.B., Muggli, P., "A concept of plasma wake field acceleration linear collider (PWFA-LC)," Proceedings of the Particle Accelerator Conference, Vancouver, B.C., May 4-8, 2009, pp. 2688-2690 (2009).

34. Martins, S.F., Fonseca, R.A., Silva, L.O., Lu, W., Mori, W.B., "Boosted Frame Pic Simulations of LWFA: Towards the Energy Frontier," Proceedings of the Particle Accelerator Conference, Vancouver, B.C., May 4-8, 2009, pp. 3160-3162 (2009).

35. Mori, W.B., Spentzouris, P., "Summary Report of Working Group 2: Computation," Proceedings of

the Advanced Accelerator Concepts Workshop, Annapolis, MD, June 13-19, 2010. AIP Conference Proceedings, Vol. 1299, pp. 88-91 (Nov 2010).

36. Martins, S.F., Fonseca, R.A, Mori, W.B., Silva, L.O., "Exploring the future of laser-plasma acceleration with massively parallel simulations in OSIRIS," Proceedings of the Light at Extreme Intensitie – Opportunities and Technological Issues of the Extreme Light Infrastructure (LEI 2009) Brasov, Romania, AIP Conference Proceedings, Vol. 1228, pp. 301-4 (Nov 2010).

37. An, W., Lu, W., Joshi, C., Mori, W.B., Huang, C., Hogan, M.J., Martins, S.F., Silva, L.O., "Simulations of Two-Bunch Plasma Wake Field Accelerator experiments at FACET," Proceedings of the Advanced Accelerator Concepts Workshop, Annapolis, MD, June 13-19, 2010. AIP Conference Proceedings, Vol. 1299, pp. 472-7 (Nov 2010).

38. Xia, G., Caldwell, A., Lotov, K., Pukhov, A., Kumar, N., An, W., Lu, W., Mori, W.B., Joshi, C., Huang, C., Muggli, P., Assman, R., Zimmerman, F., "Update on proton driven plasma wakefield acceleration," Proceedings of the Advanced Accelerator Concepts Workshop, Annapolis, MD, June 13-19, 2010. AIP Conference Proceedings, Vol. 1299, pp.510-515 (Nov 2010).

39. Pak A. Marsh K.A., Martins S.F., Lu W., Mori W.B., Joshi C., "Injection and trapping of tunnel ionized electrons into laser-produced wakes," Physical Review Letters, Vol. 104, No. 2, pp. 025003/1-4 (15 Jan 2010).

40. Martins S.F., Fonseca R.A., Lu W., Mori W.B., Silva L.O., "Exploring laser-wakefield-accelerator regimes for near-term lasers using particle-in-cell simulations in Lorentz-boosted frames," Nature Physics, Vol. 6, No. 4, pp. 311-316 (April 2010).

41. Gholizadeh, R., Katsouleas, T., Muggli, P., Huang, C., Mori, W., "Preservation of Beam Emittance in the Presence of Ion Motion in Future High-Energy Plasma-Wakefield-Based Colliders," Physical Review Letters, Vol. 104, No. 15, pp. 155001/1-4 (April 2010).

42. Vieira, J., Fiúza, F., Silva, L. O., Tzoufras, M., Mori, W. B., "Onset of self-steepening of intense laser pulses in plasmas," New Journal of Physics, Vol. 12, No. 4, pp. 045025/1-12 (30 April 2010).

43. Muggli, P., Blumenfeld, I., Clayton, C. E., Decker, F. J., Hogan, M. J., Huang, C., Ischebeck, R., Iverson, R. H., Joshi, C., Katsouleas, T., Kirby, N., Lu, W., Marsh, K. A., Mori, W. B., Oz, E., Siemann, R. H., Walz, D. R., Zhou, M., "Energy gain scaling with plasma length and density in the plasma wakefield accelerator," New Journal of Physics, Vol. 12, No. 4, pp. 045022/1-12 (30 April 2010).

44. Ralph, J. E., Clayton, C. E., Albert, F., Pollock, B. B., Martins, S. F., Pak, A. E., Marsh, K. A., Shaw, J. L.; Till, A., Palastro, J. P., Lu, W., Glenzer, S. H., Silva, L. O., Mori, W. B., Joshi, C., Froula, D. H., "Laser wakefield acceleration at reduced density in the self-guided regime," Physics of Plasmas, Vol. 17, No. 5, pp. 056709/1-6 (May 2010).

45. Martins, S. F., Fonseca, R. A., Vieira, J., Silva, L. O., Lu, W., Mori, W. B., "Modeling laser wakefield accelerator experiments with ultrafast particle-in-cell simulations in boosted frames," Physics of Plasmas, Vol. 17, No. 5, pp. 056705/1-6 (May 2010).

46. Martins S.F., Fonseca R.A., Silva L.O., Lu W., Mori W.B., "Numerical simulations of laser wakefield accelerators in optimal Lorentz frames," Physics Communications, Vol. 181, No. 5, pp. 869-875 (May 2010).

47. Hogan, M. J., Raubenheimer, T. O., Seryi, A., Muggli, P., Katsouleas, T., Huang, C., Lu, W., An, W., Marsh, K. A., Mori, W. B., Clayton, C. E., Joshi, C., "Plasma wakefield acceleration experiments at FACET," New Journal of Physics, Vol. 12, No. 5, pp. 055030/1-19 (May 2010).

48. Lu, W., An, W., Zhou, M., Joshi, C., Huang, C., Mori, W. B., "The optimum plasma density for plasma wakefield excitation in the blowout regime," New Journal of Physics, Vol. 12, No. 8, pp. 085002/1-8 (Aug 2010).

49. Clayton, C.E., Ralph, J.E., Albert, F., Fonseca, R.A., Glenzer, S.H., Joshi, C. Lu, W. Marsh, K.A., Martins, S.F., Mori, W.B., Pak, A. Tsung, F.S., Pollock, B.B., Ross, J.S., Silva, L.O., Froula, D.H., "Self-Guided Laser Wakefield Acceleration beyond 1 GeV Using Ionization-Induced Injection," Physical Review Letters, Vol. 105, No. 10, pp. 105003/1-4 (Sept 2010).

50. Kaluza, M.C., Schlenvoigt, H.-P., Mangles, S.P.D., Thomas, A.G.R., Dangor, A.E., Schwoerer, H., Mori, W.B., Najmudin, Z., Krushelnick, K.M., "Measurement of Magnetic-Field Structures in a Laser-Wakefield Accelerator," Physical Review Letters, Vol. 105, No. 11, pp. 115002/1-4 (Sept 2010).

51. Blumenfeld, I., Clayton, C.E., Decker, F. J., Hogan, M. J., Huang, C., Ischebeck, R., Iverson, R. H., Joshi, C., Katsouleas, T., Kirby, N., Lu, W., Marsh, K. A., Mori, W. B., Muggli, P., Oz, E. Oz, R.H. Siemann, D. R. Walz, and M. Zhou, "Scaling of the Longitudinal Electric Field and Transformer Ratio in a non-linear Plasma Wakefield Accelerator, Physical Review Letters ST Accel Beams, Vol. 13, No. 11, pp. 111301/1-5 (Nov 2010).

52. Gholizadeh R., Katsouleas T., Huang C., Mori W.B., Muggli P., "Effect of temperature on ion motion in future plasma wakefield accelerators," Phys. Rev. Special Topics – Accelerators & Beams, Vol. 14, No. 2, pp. 021303/1-5 (16 Feb 2011).

53. Kimura W.D., Milchberg H.M., Muggli P., Li X., Mori W.B., "Hollow plasma channel for positron plasma wakefield acceleration," Phys. Rev. Special Topics – Accelerators & Beams, Vol. 14, No. 4, pp. 041301/1-11 (18 April 2011).

54. Vieira, J., Martins, S.F., Pathak, V.B., Fonseca, R.A., Mori, W.B., Silva, L.O., "Magnetic Control of Particle Injection in Plasma Based Accelerators," Physical Review Letters, Vol. 106, No. 22, pp. 225001/1-4 (31 May 2011).

55. Pollack, B.B., Clayton, C.E., Ralph, J.E., Albert, F., Davidson, A., Divol, L., Filip, C., Glenzer, S.H., Herpoldt, K., Lu, W., Marsh, K.A., Meinecke, J., Mori, W.B., Pak, A., Rensink, T.C., Ross, J.S., Shaw, J., Tynan, G.R., Joshi, C., Froula, D.H., "Demonstration of a Narrow Energy Spread, □0.5 GeV Electron Beam from a Two-Stage Laser Wakefield Accelerator," Physical Review Letters, Vol. 107, No. 4, pp. 045001/1-4 (18 July 2011).

56. Sahai, A., Katsouleas, T.C., Tableman, A., Tonge, J., Tsung, F.S., Mori, W.B., "Proton Acceleration by Trapping in a Relativistic Laser Driven Uphill Plasma Snowplow," Proceedings of the 2011 Particle Accelerator Conference, New York, NY, March 28-April 1 (2011) 3pp.

57. England, R.J., Frederico, J., Hogan, M.J., Joshi, C., An, W., Lu, W., Mori, W.B., Muggli, P., "A High Transformer Ratio Plasma Wakefield Accelerator Scheme for FACET," Proceedings of 2011 Particle Accelerator Conference, New York, NY, March 28-April 1 (2011) 3pp.

58. Davidson, A.W., Lu, W., Joshi, C., Silva, L.O., Martins, J., Fonseca, R.A., Mori, W.B., "Numerical Study of Self and Controlled Injection in 3-Dimensional Laser-Driven Wakefields," Proceedings of 2011 Particle Accelerator Conference, New York, NY, March 28-April 1 (2011) 3pp.

59. Xia, G., Caldwell, A., Huang, C., Mori, W.B., "Simulation Study of Proton-Driven PWFA Based on CERN SPS Beam," Proceedings of the 2011 Particle Accelerator Conference, New York, NY, March 28-April 1 (2011) 3pp.

60. Tochitsky, S., Ya., Tsung, F., Haberberger, D.J., Mori, W.B., Joshi, C., "Enhanced Laser-Driven Ion Acceleration via Forward Raman Scattering in a Ramped Gas Target," Proceedings of 2011 Particle

Accelerator Conference, New York, NY, March 28-April 1 (2011) 3pp.

61. Shi, Y., Chang, O., Muggli, P., Huang, C., An, W., Mori, W., "Betatron Radiation from an Off-axis Electron Beam in the Plasma Wakefield Accelerator," Proceedings of the 2011 Particle Accelerator Conference, New York, NY, March 28-April 1 (2011) 3pp.

62. Viktor K. Decyk and Tajendra V. Singh ,"Adaptable Particle-in-Cell algorithms for graphical processsing units," Computer Physics Communications 182, 641

63. Marsh, K.A., Clayton, C.E., Joshi, C., Lu, W., Mori, W.B., Pak, A., Silva, L.O., Lemos, N., Fonseca, R.A., Albert, F., Doeppner, T., Filip, C., Glenzer, S.H., Price, D., Ralph, J., Pollock, B.B., "Laser Wakefield Acceleration Beyond 1 GeV Using Ionization Induced Injection," Proceedings of 2011 Particle Accelerator Conference, New York, NY, March 28-April 1 (2011) 5pp.

64. Xia, G., Caldwell, A., Lotov, K., Pukhov, A., Assmann, R., Zimmermann, F., Huang, C., Vieira, J., Lopes, N., Fonseca, R.A., Silva, L.O., An, W., Joshi, C., Mori, W., Lu, W., Muggli, P., "A Proposed Experimental Test of Proton-Driven Plasma Wakefield Acceleration Based on CERN SPS," Proceedings of 2011 Particle Accelerator Conference, New York, NY, March 28-April 1 (2011) 3pp.

65. Viera, J., Huang, C.K., Mori, W.B., Silva, L.O., "Polarized beam conditioning in plasma based acceleration," Physical Review Special Topics – Accelerators and Beams, Vol. 14, No. 7, pp.071303/1-11 (28 July 2011).

66. Fiuza, F., Fonseca, R.A., Silva, L.O., Tonge, J., May, J. Mori, W.B., "Three-Dimensional Simulations of Laser-Plasma Interactions at Ultrahigh Intensities," IEEE Transactions on Plasma Science, Vol. 39, No. 11, pp. 2618-19 (7 November 2011).

67. Haberberger, D., Tochitsky, S., Fiuza, F., Gong, C., Fonseca, R.A., Silva, L.O., Mori, W.B., Joshi, C., "Collisionless shocks in laser-produced plasma generate monoenergetic high-energy proton beams," Nature Physics, Vol. 8, No. 1, pp. 95-99 (January 2012).

YEAR 1

I. Pipelining and improvements to QuickPIC:

In the QuickPIC computation cycle, the plasma response to the driver is evaluated by sweeping a slice of plasma along the longitudinal direction and calculating the trajectories of the plasma particles as the slice moves from the front of the driver to the back. This setup allows a software pipelining technique to be used, which can dramatically decrease the turn-around time of a PWFA/LWFA or electron cloud simulation by adding more processor groups to the pipeline. We have implemented the pipelining algorithm into both the basic and full quasi-static version of QuickPIC for the particle beam driver. The HDF diagnostic routine has also been modified to allow merging data of each processor group on the fly without post-processing. This pipelining algorithm has been successfully tested on NERSC platforms with 2048 processors in which a maximum of 64 processor groups was used. The pipelining algorithm is verified to produce the same result as the non-pipelining version. Performance measurement shows that the speedup of the 2D plasma slice solver in the pipelining mode is nearly ideal, this is because the data transfer between two successive groups is relatively inexpensive compared to the time spent on the solver itself and the transfer also overlaps with the computation. When the 3D beam pusher is included, the overall efficiency of each processor group in the pipelining mode reaches 85% relative to the non-pipelining mode, thus using 64 processor groups leads to a 54 times reduction in turn-around time for a long simulation. A description of the pipelining routine will be submitted to J. Comp. Physics shortly. We also identified two key factors for improving the efficiency and processor number scaling of the pipelining algorithm, they are (1) choosing a 2D domain decomposition strategy in the 3D beam pusher and (2) implementing load-balancing in the same pusher. We have also begun to improve the predictor corrector routine for both preionized and self-ionized plasmas. A simple finite difference solver was recently developed for QuickPIC and used to subtitute the FFT-based solver for test purpose. The finite difference scheme allows a position-dependent diffusion coefficient for the field solver. Comparisons on the benchmark problem show that the finite difference solver improves the stability of the solver when the plasma density is hightly nonuniform.

II. Beam Loading:

A theory of how electrons can be loaded into the nonlinear wakefield generated by either a laser or particle drive beam has been developed. The maximum amount of charge that can be loaded and the optimum shape are derived and confirmed by OSIRIS simulations. In the ideal case, the efficiency of beam-loading is shown to be 100%. A preliminary plasma wakefield afterburner design based on this theory has been simulated in a QuickPIC simulation on NERSC. In the simulation, the drive electron beam has 4.4×10^{10} electrons and a length of 58 microns, the trailing electron beam has 1.7×10^{10} electrons and a length of 22 microns. The drive beam has a triangular longitudinal current profile while the trailing beam has a inverse triangular profile, this choice of the beam profiles

has lead to complete flattening of the decelerating and the accelerating fields in the plasma wake. Therefore, remarkably small energy spread is achieved for the accelerated trailing beam! A paper has been submitted to Phys. Rev. Lett.

***III. Benchmarking codes*:**

In collabortation with Tech-x and LBNL, we have begun benchmarking of the three major simulation codes used in advance accelerator modeling effort, OSIRIS, VORPAL and QuickPIC. A series of 3D benchmarks covering the linear and nonlinear regimes of LWFA problem were run with each code. In these benchmarks, the laser pulse has the following fixed parameters, laser wavelength λ=0.8 micron, pulse duration τ = 60 fs, spot size w_0 = 8.2 micron, and the magnitude of the normalized vector potential of the laser are a_0=0.5, 1, 2 and 4, respectively. The plasma density is 1.38×10^{19} cm^{-3}. The simulation box size is 80×80×20 microns with 512×512×512 grid points. The time step is 0.2075 ω_0^{-1}. The longitudinal wake electric field and the laser electric field in OSIRIS and VORPAL simulation at 800 and 1600 time steps were compared with each other. In addition, longitudinal electric field from QuickPIC simulation is also compared with the full PIC results. The benchmarks show that the OSIRIS and VORPAL results are nearly identical. QuickPIC results are also seen to agree extremely well with the full PIC results for a_0=0.5, 1, 2 cases, while for the a_0=4 case deviation is seen near the spike of the accelerating field. Details about the benchmark and the results are being prepared for publication in the SciDAC proceedings. We will next run benchmarks for which the laser evolves, as well as test ideas for how to improve QuickPIC for better agreement for the larger values of a_0. QuickPIC and OSIRIS are also being compared to the code WAKE. The quasi-static codes can be more than two orders of magnitude faster than the full PIC codes.

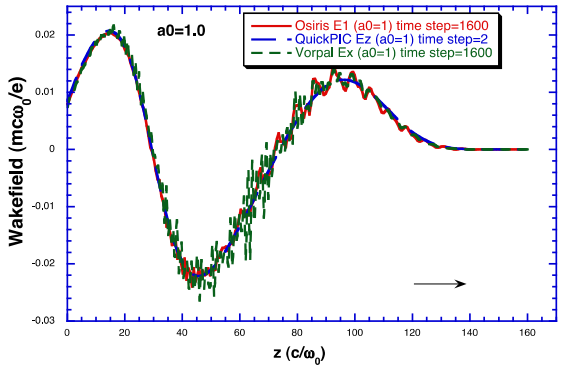

Figure 1. Three major simulation codes used in advanced accelerator modeling, OSIRIS, VORPAL and QuickPIC are benchmarked against each other.

***IV. Modeling experiments*:**

We have conducted QuickPIC simulations with relevant parameters for the future SLAC 2-beams PWFA experiments. The simulations use beam profiles derived from particle tracking data. Optimal charge distribution study and plasma parameter scan were carried

out for designing a 25GeV single-stage PWFA afterburner. Both pre-ionized and self-ionized plasma are simulated. Based on the beam-loading theory in the nonlinear blow-out regime, ideal parameters for the drive beam and the trailing beam was obtained and energy doubling of 25 GeV beams with energy spread within 1% was demonstrated in the simulation (Fig. 1). Several high resolution QuickPIC simulations are conducted to study the relevant physics in future plasma afterburners. The design of the beam parameters is based on the beam-loading theory described above. The spot size of the drive beam is 10 microns to avoid significant ion motion, while the trailing beam spot size has to be sub micron for collider application. In the simulation, we use 2048×2048×256 grids with box size of 600×600×270 microns. The transverse resolution is ~0.3 micron which allows us to simulate a 0.6 micron spot size trailing beam. Both the drive beam and the trailing beam have initial energy of 250 GeV. After propagating about 4 meters in the plasma, the trailing beam has gain around 100 GeV energy. We have conducted simulations with/without ion motion and synchrotron radiation loss to investigate their effects on beam acceleration and propagation. We have also assisted in carrying out OSIRIS simulations of laser guiding experiments conducted at UCLA.

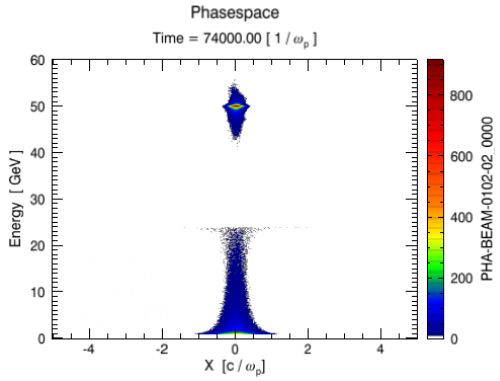

Figure 2. A trailing beam (initial energy 25GeV) with 0.34×10^{10} electrons gains 25 GeV energy in a PWFA simulation for the planned SLAC 2-beams experiment. The plasma is preionized.

***V. Scaling PIC codes to 10,000 processors*:**

We have done extensive benchmarks with the UPIC Framework on large systems to identify areas where current algorithms can be improved to enable scaling of PIC codes to 100,000 processors. This is directly applicable to the QuickPIC code, which depends on this Framework, but new successful algorithms can also be implemented in the other PIC codes. A billion particle benchmark simulation with a 512x256x512 grid, using an electromagnetic, relativistic plasma model has been run successfully on an Opteron-based cluster with Infiniband. This strong scaling study in which the number of processors was varied from 128 to 8192 while keeping the particle size fixed showed good scaling (92%) for the particle part of the calculation (which typically dominates PIC codes). The FFT part scaled well up to 4096 processors, then saturated at 8192. Top processing speeds of more than 10 billion particles per second (for an entire step including field solve etc.) were obtained. To enable scaling to 100,000 processors, two approaches are being tested. First, strategies to improve existing algorithms in the particle manager and the FFT were

identified. Second, a mixed MPI/threaded programming model has been implemented and is now being tested.

A nearly identical scaling study is done using the 3D electromagnetic PIC code OSIRIS. These two codes (OSIRIS and UPIC) have different strengths and we have initially our timing results to reflect these differences. For example, UPIC is more mature and is more optimized on a single CPU, but its spectral field-solver makes perfect scaling difficult for a large number (>5,000 CPU's on the ATLAS cluster) of CPU's. OSIRIS, on the other hand, is not as optimized on a single processor but its finite difference field solver which is scalable for a large number of processors. For these reasons, we anticipated that the behaviors of these two codes to be different over such a wide range of CPU numbers. Nonetheless, the two codes both show similarly excellent scaling up to ~4,000 CPU's. The scaling results are shown in the figure below.

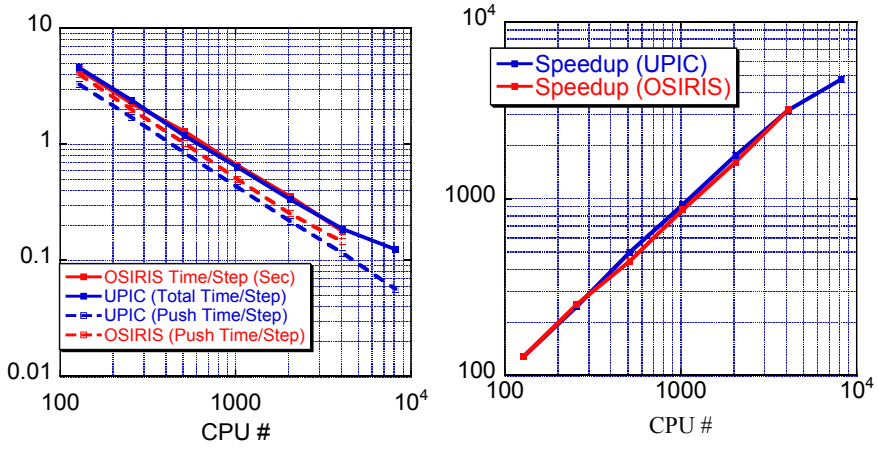

Figure 3. Timing and speedup of OSIRIS (in red) and UPIC (in blue) on the ATLAS supercomputer.

YEAR 2

I. Scale OSIRIS to 10,000+ processors:

During the previous year we have spent time demonstrating that OSIRIS scales well to over 10,000 processors. We have also spent time testing how OSIRIS performs on a single processor. Arguably the most important figure of merit for the performance of a PIC code is the time it takes to push one particle in one time step. A code which does this well is harder to scale to more processors. The effort on OSIRIS has also included developing a single precision version that works well on vector and SIMD processors without loss in accuracy for the particles. A beta version is working in 2D and one will soon be working in 3D. In addition, Viktor Decyk has been developing new data structures for PIC codes that should allow efficient use of GPUs and next generation multi-core devices. With these new data structures, single precision calculations also do not lose accuracy for the particles. Preliminary benchmarks for both the charge deposit and particle push have been very promising, showing at least an order of magnitude improvement for a GPU over a core.

As an example of the testing we conducted, we show in table 1 how OSIRIS performs in 3D on two systems. HYDRA is an i7 3.2 Ghz node with four cores. Hoffman 2 shared cluster housed at UCLA. We used one quad core AMD 2.2 Ghz node on the Hoffman 2 cluster. The loop time includes the field solve, deposit, and particle push. The simulation used 4 million particles on a 64x64x128 mesh (8 particles per cell). We expect even better efficiency on when the dual quad core i7 nodes are released later this year. Linear weighting was used. For quadratic weighting the numbers are 1.5 slower. In a 2D version of OSIRIS where pieces were rewritten to take advantage of vector processors in single precision a speed up of a factor of 2-3.5 was seen depending on the processor.

	Loop time/ particle/step [ns]	**Speedup from 1 core**	**Efficiency**
HYDRA			
1 core	176		
4 cores	50	**3.5**	89%
8 cores*	30	**5.8**	73%
HOFFMAN2			
1 core	346		
4 cores	100	**3.5**	86%
8 cores*	64	**5.4**	67%

* forced with *mpirun -np 8*

Table 1. Timings for the OSIRIS in 3D on quad core i7 3.2 Ghz node (HYDRA) and on a quad core 2.2 Ghz node (HOFFMAN 2).

Last year we reported on weak and strong scaling on the 9,000 processor ATLAS system at LLNL. This past year, we have run performance tests of OSIRIS at the Argonne BlueGene Supercomputer Intrepid (www.alcf.anl.gov). The parallel scalability tests were

performed with up to 32768 CPUs. Two scenarios are used to test the performance. In one scenario (Warm) the plasma is stable so the load will stay balanced and communication is moderate. In this case an electron plasma is simulated in 3D on a 512 x 512 x 512 grid with more than 1 billion particles (8 particles per cell). The second scenario involves the collision of 2 hot electron-positron clouds flowing in opposite directions on a 512 x 512 x 256 grid with more than 1 billion particles (16 particles per cell). This scenario causes severe turbulence and cavitation, greatly stressing communication between nodes and parallel scalability, and was used as a "worst case" scenario. We have dynamic load balancing in the code, but we have not had a chance to test it on this system.

Strong scaling for each scenario relative to a 256 CPU's are shown in Figure 1. The efficiency for strong scaling from 256 CPUs to 32768 CPUs is 86% in the warm case and 57% in the collision case. Even in our worst case scenario (the run with colliding plasmas) this strong scaling is excellent. This shows that OSIRIS scales to well over 10,000 CPU's.

We have also been working on scaling QuickPIC to 10,000+ processors. Last year, we reported on the successful implementation of a software pipelining technique into QuickPIC, which can dramatically decrease the turn-around time of a PWFA/LWFA or electron cloud simulation by adding more processor groups to the pipeline (currently pipeline mode is available for PWFA and electron cloud problems). We also identified two key factors for improving the efficiency and scaling on the number of processors for the pipelining algorithm. They are (1) using a 2D domain decomposition strategy in the 3D part of the code and (2) implementing load-balancing in the beam pusher and deposition routines. The first work will enable better allocation of the computation resources in a pipeline stage and eliminate the limit of memory per CPU on high resolution simulations. The second work will enable fine-tuning of the CPU allocation which will be beneficial for achieving higher parallel efficiency. As the first step, we are implementing the 2D domain decomposition in QuickPIC. The 2D domain decomposition is supported in the UPIC framework which is the building block of QuickPIC. We have imported the necessary changes from the UPIC framework to QuickPIC. In addition, special attentions need to be given to the domain decomposition in the pipeline mode. We have implemented the major upgrades to the pipeline routines to handle the improved domain decomposition. Currently the 2D domain decomposition algorithm is under testing and preliminary result shows scaling beyond the limit of the original 1D domain decomposition. This is work in progress.

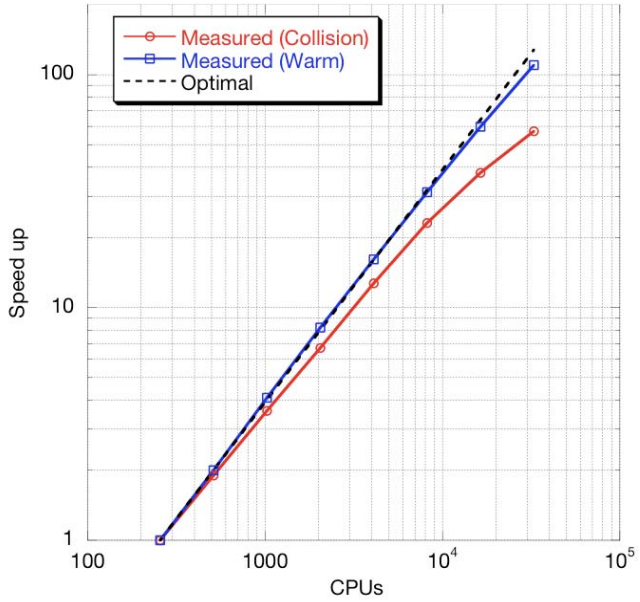

Figure 1: Speed up for the two test cases and comparison with optimal scaling.

II. Compare quasi-static PIC vs. PGC PIC:

In collaboration with Tech-X, we are conducting benchmarking of Quasi-Static PIC (QuickPIC) and the Ponderomotive Guiding Center PIC (in VORPAL) algorithms for LWFA simulations. Both algorithms assume the fast laser oscillation can be averaged out and the plasma response can be approximated by the responses to the ponderomotive potential (the envelope model). The Quasi-Static algorithm further assumes that the laser envelope changes slowly on the time it takes the laser to pass by a plasma particle (this would not be true for a trapped electron). In the benchmark, the laser wavelength is 0.8 micron, the pulse duration (FWHM) is 15fs, the pulse width is 8.2 microns and the plasma density is 1.38×10^{19} cm^{-3}. The normalized vector potential of the laser will be varied from 0.5, 1.0, 2.0 and 4.0. This work is in progress and Figure 2 shows the agreement of both algorithms for the case where the normalized laser vector potential is 0.5.

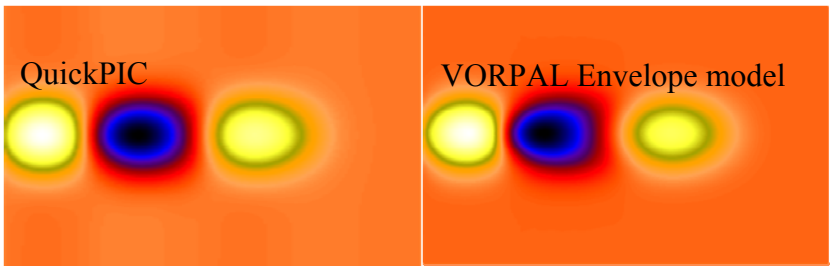

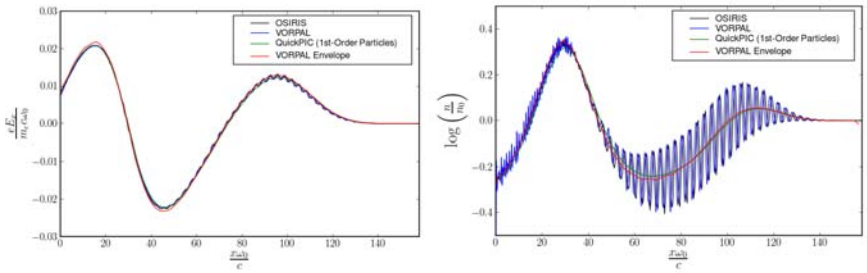

Figure 2 Top: Benchmark of the longitudinal acceleration field in LWFA problem with a0=0.5 using Quasi-Static PIC(QuickPIC) and PGC PIC (VORPAL – envelope model). Bottom: Benchmark of the on-axis longitudinal acceleration field (left) and plasma density (right) for a0=1.0 between OSIRIS, VORPAL (full PIC and envelope model) and QuickPIC.

III. Model FACET, L'OASIS, and other relevant experiments with QuickPIC and OSIRIS:

QuickPIC and OSIRIS have been used extensively to model recent and future LWFA/PWFA experiments at SLAC/LBL/RAL. One of the goals of the proposed FACET program at SLAC is to achieve acceleration of a trailing electron beam in the wakefield of a drive electron beam. The designed facility can deliver 3~5.6 nC electron beams in two bunches with 25 GeV nominal energy. Some of the compelling questions needed be addressed to design future experiments include determining if self-ionized induced beam head erosion limits the energy transfer efficiency and how the energy spread and loaded gradient depend on the spacing between the drive and trailing beams. A series of QuickPIC simulations using the actual profiles of the beams from the particle tracking code ELEGANT were carried out with total charge varying from 3~5.6 nC to test the possibility of doubling the energy of the trailing beam in the self-ionized plasma. In the simulation, the spacing between the bunches was verified to produce small (a few percent) energy spread of the trailing beam. And the final energy of the trailing beam can reach more than 50 GeV with sufficient charge in the drive beam. Such simulations will provide important guiding of the operation of the future experiments. We have also been using OSIRIS to understand how ionization induced particle trapping might affect the wake. The code OSIRIS is being used by the SLAC AARD group to understand data produced in the e-168 experiment and to better understand ionization induced particle trapping. We have also been using OSIRIS and QuickPIC to study various beam loading scenarios for positrons. We have primarily concentrated on loading positrons into weakly nonlinear wakes driven by electron beams.

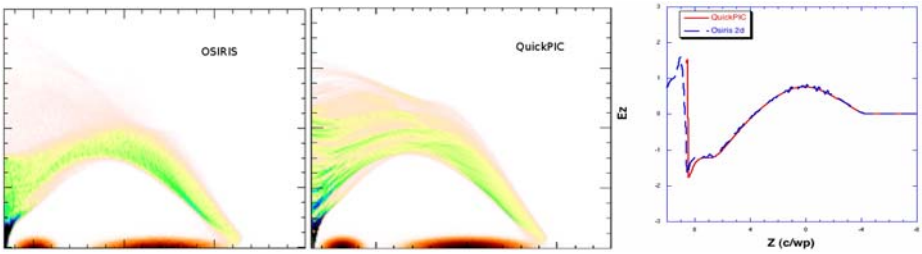

Figure 3. Left and middle: Drive and trailing beams and the plasma wake from the beam-ionized lithium gas in OSIRIS 2D cylindrical and QuickPIC simulations. Right: the on axis longitudinal wakes in the beam-ionized plasma from the OSIRIS and QuickPIC simulations.

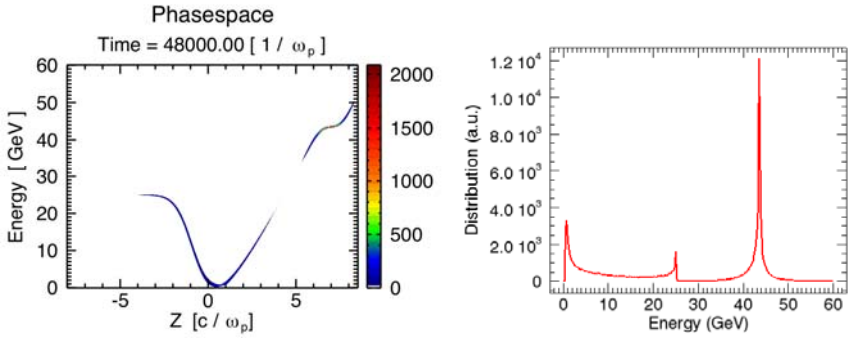

Figure 4. Left: Phasespace of the drive and trailing beams after 80 cm propagation in plasma. Right: the spectrum of the drive and trailing beams at the same distance.

We have also continued to carry out OSIRIS and QuickPIC simulations of possible next generation LFWA experiments. For example, as part of a collaboration with the plasma group in IST, Portugal, OSIRIS and QuickPIC has also been used to simulate the possible LWFA experiment at RAL 10 PW laser upgrade.

The development of new laser systems based on OPCPA will push Laser Wakefield Accelerators (LWFA) to a new energy range. The optimal configurations for particle acceleration with these upgraded laser sources are not yet completely known, both due to the inherent nonlinearities in the process, and to the extreme computational requirements of one-to-one numerical simulations. . Based on the prospective design parameters for the future Vulcan 10~PW OPCPA laser system, we have determined the optimal parameters for a single LWFA stage using theoretical scalings of Lu et al., for such system and a combination of QuickPIC and OSIRIS simulations. Based on this work, we believe such a laser system could produce self-injected electrons in excess of 10~GeV in a self-guided configuration and produce 50~GeV bunches with externally-injected electrons in a laser-guided configuration. We have also shown that it is possible to perform ultra-fast LWFA simulations in three dimensions with all the relevant physics in these regimes by using a Boosted frame. In figure 3, a simulation of a 10 PW laser operating in the bubble regime is presented. For this case the 10 PW laser has a pulse length of 30fs, it is focused to an a_0=43, and the plasma density if 1.5x10^19 cm^-3. For this case the maximum electron energy is 2.3 GeV. The simulation was done in the lab frame. In figure 4, results

are shown from the blowout regime where the laser now has a pulse length of 110 fs (and the same energy). The plasma density is now 2.7x10^17 and the focused a_0=5.8. For this case the maximum energy is 13 GeV and the simulation was done in a boosted frame. Based on this work, we believe that the work of Lu et al., is an excellent starting point to design single stages (or even multiple stages) that can reach energies beyond 10GeV. This work will be submitted for publication shortly. Furthermore, with sufficient support the boosted frame algorithm might allow a new generation of numerical simulations that can perform parameter testing and scanning for the setup and optimization of experiments into the next decade.

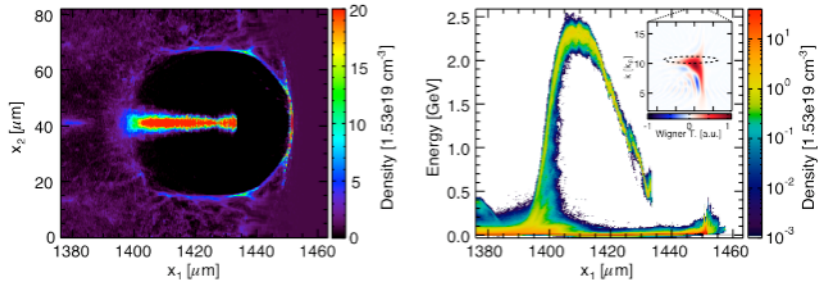

Figure 5 Simulation of a 300J laser with a pulse length of 30 fs. This simulation was done in the lab frame. The plot on the left is a slice of the plasma density and the plot on the right is of the energy vs. position of the self-trapped electrons.

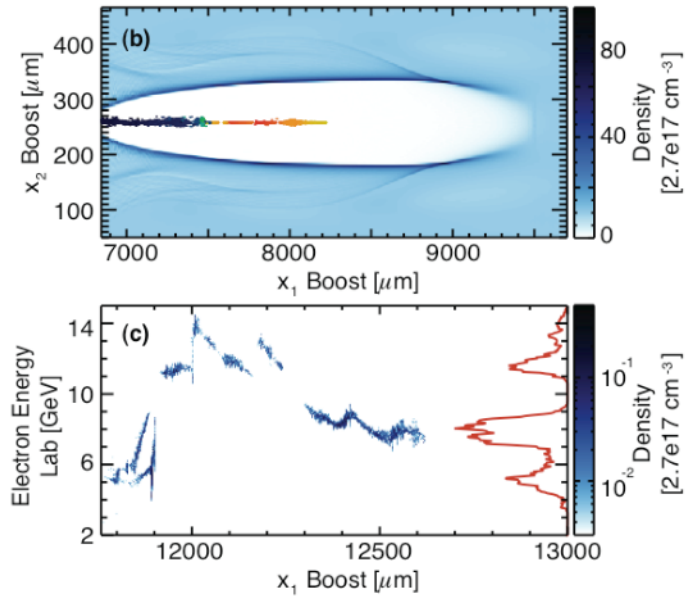

Figure 6. Simulation of a 300J laser with a pulse length of 110 fs. This simulation was done in a boosted frame. The plot on the top is a slice of the plasma density and the plot on the bottom is of the energy vs. position of the self-trapped electrons.

QuickPIC has also been used to simulate the possible LWFA experiment at RAL PW laser upgrade, which is expected to deliver 300 J laser pulses. In order to use such laser to

extract multi-GeV electron beams from centimeter to meter scale plasmas, it is required that the radiation losses are kept to a minimum, and that efficient guiding of the transverse laser profile is achieved. QuickPIC simulations are designed to explore these possibilities. Simulations were carried out to examine the acceleration of externally injected electron beams using PW-class lasers in plasmas with densities of 10^{16}-10^{17} cm^{-3} including the effects of radiation damping. The externally injected electron beam is placed at the end of the first bucket in the wake in the region of maximum accelerating gradient. In addition, low charge beams with 1-10 pC are used, so that the electron beams are accelerated in a test-particle regime. The laser evolution is analyzed in both uniform and in pre-formed parabolic plasma channels. In the first scenario, self-guided laser propagation regimes are examined, whereas the second allows stable propagation of lower power laser pulses in comparison to self-guided propagation regimes. With uniform plasmas, stable laser propagation is achieved for more than 10 cm. The resulting plasma wave accelerates the externally injected electrons to more than 15 GeV. Using a preformed plasma channel guarantees stable laser propagation for more than 5 meters, accelerating electrons to more than 30 GeV.

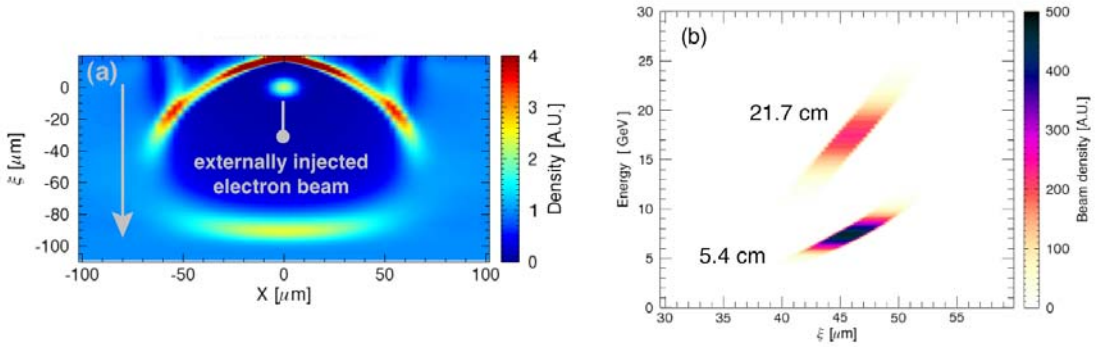

Figure 7 (a) Plasma response at the beginning of the simulation for the self-guided propagation regime. The arrow indicates the direction of the laser beam. (b) Externally injected beam phase-space for the self-guided propagation regime, for two different propagation distances, showing final energies in excess of 15 GeV.

IV. Model a 500 GeV PWFA Afterburner for ILC parameters with pipelined QuickPIC:

This work was modified to study the staging of several 25 GeV stages together. This is the ultimate goal of the FACET collaboration. The goal is to use 20 or more stages to reach 500 GeV or higher energies. Therefore 25 GeV PWFA stages are the key components of a PWFA-based TeV linear collider concept. In this exercise keeping the beam energy spread low and acceleration efficiency high are the major considerations. To achieve an energy spread less than 1%, will require the beam-loaded wake to be flat within the beam. This may require tailing the beam current profile. Furthermore, to achieve high efficiency of energy deposition into the wake, a customized current profile for the drive beam might also be needed. Using standard Gaussian shaped current profiles for both the drive and trailing beams we have found that the energy spread can be kept

below a few percent and the transfer efficiency can be near 50%. To improve on this we have used our recent beam-loading theory in the relativistic blow-out regime. Using this theory we find that the optimal profiles for the beams turn out to be trapezoidal for the drive beam and inverse-trapezoidal for the trailing beam. In the detailed design of the 25GeV PWFA stage, the drive beam has 3.65E10 electrons in its trapezoidal body and has 0.82E10 electrons in a triangular precursor. The length of the precursor and the body are 13.4 microns and 44.7 microns respectively. The trailing beam has 1.62E10 electrons and a length of 22.35 microns. The plasma density is 5.66E16 cm^{-3} and its length is 0.7 meter. The spot sizes of the beams are 3 microns, and the emittances are 62.9 mm mrad, which are matched to the plasma density, except for the precursor which has smaller emittance of 10 mm mrad to prevent head erosion. The transformer ratio is 1.2 with a loaded wake at 45GeV/m. The QuickPIC simulation user 1024x1024x256 grids with a resolution of 0.98x0.98x1.06 microns. Two simulations are conducted where the trailing beam energy is 25 GeV and 475 GeV, which correspond to the first and the last stage of a PWFA-based TeV linear collider concept design. Figure 8 shows the design of a PWFA stages and the simulation result for a plasma-based TeV collider.

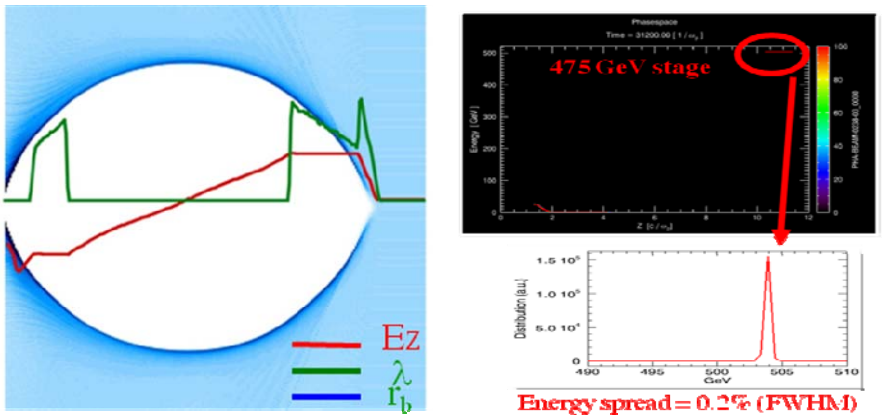

Figure 8. Left: The current profile of the drive and trailing beams(green), the longitudinal wakefield (red) and the blow-out plasma channel (blue) in a 25GeV PWFA stage for TeV linear collider. Right: The phasespace of the drive and trailing beam (top) of the last PWFA stage in a TeV collider and the spectrum of the trailing beam (bottom) showing a 0.2% energy spread.

YEAR 3

I. Scale QuickPIC to 10,000+ processors:

In the last report we described our efforts to scale QuickPIC to 10,000+ processors. As described we successfully implemented a pipelining algorithm together with new domain decomposition scheme which permitted scaling large simulations to more than 10,000+ processors. However, the QuickPIC routine has room for improvement in algorithmic efficiency. Specifically, during the past period of performance, we have been working on improving the algorithm including the field equations and the iteration loop.

In QuickPIC, a full Maxwell solver is reduced to a non-radiative EM solver using the quasi-static approximation. As a result, the three dimensional plasma response to the driver (a particle beam or a laser pulse) can be calculated slice by slice along the longitudinal direction. This allows one to convert the 3D problem into a 2D solver wher the time advance corresponds to pushing the 2D slice back through the driver in the axial direction. In the solution for one transverse 2D plasma slice, the fields are solved according to the current and charge density in the plasma slice and the particles are pushed in the fields to advance the current and charge density forward in time. Unfortunately, the discretization of the basic equations under the quasi-static approximation shows that the fields at one time step depend on the current and charge density at same time step. Therefore, the leap-frog algorithm, which is widely used in the full PIC codes, will not work in QuickPIC. Instead, an iteration (predictor-corrector) method is applied for solving those coupled equations. Generally speaking, in the 2D plasma calculation procedure, predicted fields are used to push the particles and generate predicted plasma current and charge densities. These current and charge densities are used to solve for corrected fields which are then used to calculate the corrected current and charge densities. This iteration can be carried out as many times as is needed.

In the original implementation of QuickPIC the scalar and vector potentials are solved for in the Lorentz gauge. In practice, we often need 5 or 6 iterations for very nonlinear problems. Recently, we have experimented a new field solver (and iteration loop) which uses the the transverse Coulomb gauge. Preliminary results indicate that only 2 iterations are required which will lead to a savings of more than a factor of 2. We are investigating whether additional improvements can reduce the number of iterations to 1 and we are working to improve the speed of each iteration (particle push and field solve).

II. Carry out high resolution PWFA simulations

Even with a highly efficient tool like QuickPIC, simulating beam-plasma interactions in the PWFA-LC concept is still very challenging due to the extreme beam dimensions and energy requirements of the next generation linear collider. And to develop PWFA-LC into a successful design requires detail understanding of the relevant physics such as hosing, ion motion, beam-loading, radiation damping etc, their interplay, and carefully choosing the operation point in the parameter space. Our simulation work aims at exploring the road map for building a linear collider utilizing the plasma wakefield

acceleration mechanism with our highly efficient and accurate simulation tools.

The set of physical design parameters for a 20-stages TeV class PWFA-LC is based on the study in [Huang-PAC2009]. In this preliminary design, a driving beam and a witness beam with 3E10 and 1E10 electrons are used respectively. The driving and witness beam, both having 25 GeV initial energy, propagate in a meter long plasma with 100 microns separation. The longitudinal electric field of the plasma wakefield excited by the drive beam and sampled by the witness beam exceeds 20 *GeV/m*.

While the drive beam that excites the plasma wakefield can be made with a relatively large spot size, the major computation requirements come from modeling the witness beam that needs to be accelerated. The smallest emittance of the witness beam in a TeV collider can be 0.093 mm·mrad (geometric mean for the two transverse planes), which implies that the matched spot size of the witness beam at 25 GeV in a plasma of 1×10^{17}cm^{-3} will be 100 nm, which is three orders of magnitudes smaller than the longitudinal spot size or the plasma wavelength. In our first simulation, the transverse box size is around 24 c/ω_p, which is ~400 microns. We use 8192 grids in both transverse directions . In the longitudinal direction, the plasma wavelength needs to be well resolved using O(1000) grids. Here we use 2048 grids. Such a 3D simulation of the would need $8192\times8192\times2048 = 1.3\times10^{11}$ grids and 5.2×10^{11} plasma particles (assuming 4 particles/cell). In QuickPIC each 2D slice will have $2.5\text{x}10^{8}$ paricles. The smallest time step and the corresponding number of time steps for meter-long propagation distance are ~ 1.3 ps and ~2600 respectively for a quasi-static simulation. Very recently, we started the initial QuickPIC simulation for studying the first stage of a TeV class plasma wakefield accelerator linear collider (PWFA-LC). A simulation was started on Jaguar utilizing 16,384 processor cores. Figure 1. shows a snapshot of the simulation with drive and witness beam and the plasma response after 100 3D timesteps

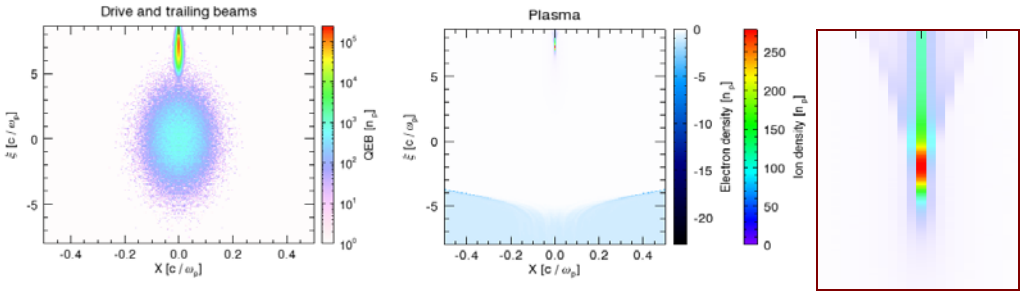

Figure 1. Left: Drive beam and witness beam in high resolution QuickPIC simulation for PWFA-LC. Middle: The plasma response in the simulation showing ion concentration due to tight witness beam at the top center. Right: A blowup of the ion compression region.

***III. Model previous FFTB and future FACET experiment using QuickPIC and OSIRIS*:**

Next generation experiments at the FACET facility which is under construction at SLAC will be aimed at showing that a trailing bunch of electrons can be accelerated with high transfer efficiency while maintaining its emittance and energy spread. The FACET facility will have electron and positron bunches at 23 GeV with 3nC of charge. In addition, the beams can have bunch lengths as short as 14mm which corresponds to peak currents of 22 kA. Designs of these experiments indicate that two bunches with ~1nC in the drive beam and ~.5nC in the trailing beam and a separation ~100 mm are possible. These conditions have been simulated with QuickPIC and OSIRIS to find the optimum density.

The most important experiment we are simulating is the two-bunch PWFA. The two-bunch configuration needed to demonstrate a narrow energy spread at FACET will be crafted by manipulating the electron beam phase space in the dispersion plane. In earlier PWFA experiments in the FFTB the plasma was created by self-ionization from the beam's own electric field. However, the process of generating the two bunches at FACET will lower the peak current of the drive beam thereby requiring a gas with a lower ionization threshold. This might be accomplished by choosing Cs which has an ionization threshold of 3.9 eV as opposed to previously used Li which has an ionization threshold of 5.4 eV. Simulations done using the code QUICKPIC and OSIRIS show that the drive beam is able to produce a meter long plasma with density of up to 5 x 10^16 cm^-3 before beam head erosion effects terminate the wake. The appropriately placed trailing beam can be accelerated by the wake to give energy gains on the order of 10 GeV with a narrow energy spread. See figure 2.

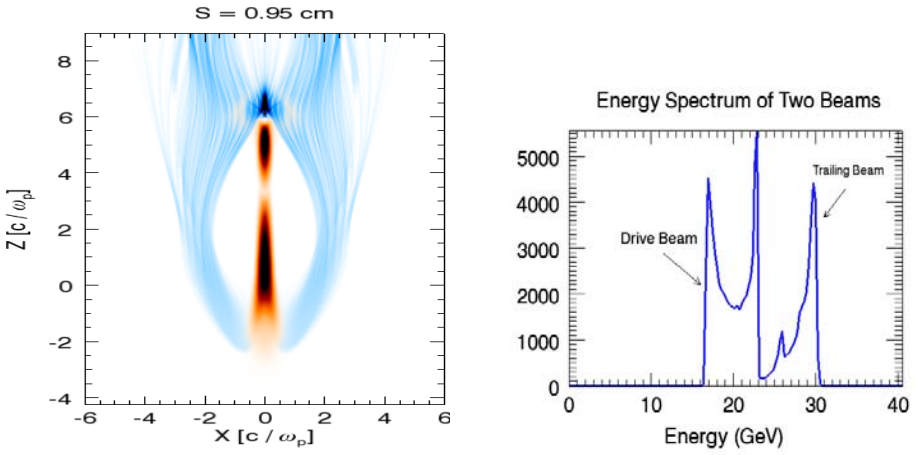

Figure 2. Simulation of the 2 beam FACET experiments in a field ionized plasma using a Cesium gas cell. On the left are the drive and trailing beams (moving downward) and the plasma density in the wakefield in a self-ionized plasmas. On the right is the energy spectra for the drive and trailing beams.

The energy gain can be increased by propagating the two beams in a pre-ionized plasma. In figure 3 we show that for a preformed plasma, that the trailing beam can be accelerated from 23 GeV to 50 GeV with a few % energy spread in less than 1 meter.

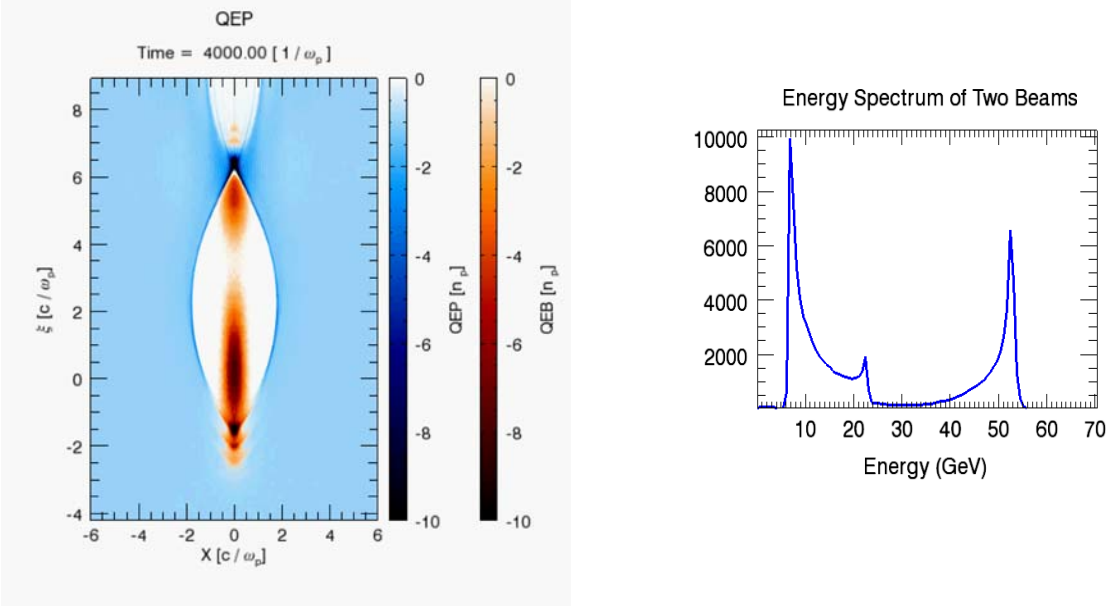

Figure 3 QuickPIC simulations of possible two bunch experiments at FACET. On the left is the beam and plasma density. The plasma density is chosen so that the trailing bunch loads the rear of the wake. On the right is the energy spectra of both bunches. The drive bunch is been decelerated from 23 GeV to 6 GeV while the trailing bunch has been accelerated to 50 GeV.

The initial QuickPIC simulations were done on Franklin at NERSC. On Jaguar we have also been carrying out the first 3D full-scale OSIRIS simulations of FACET to compare against QuickPIC and to study ionization induced trapping (self-trapping is not included in QuickPIC). Typically, to model this problem using OSIRIS, the size of the simulation box is 400*300*300 (in grid); the total particle number is ~10^8;the total time steps is ~10^6. We use higher order particle shapes. We use 3000~6000 processors in a typical simulation. Preliminary results are shown in figure 4.

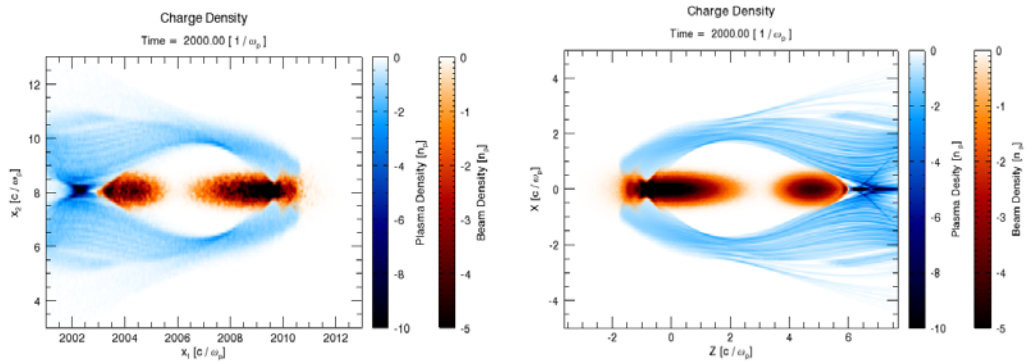

Figure 3. OSIRIS (left) and QuickPIC (right) simulations of the 2 beam experiment on FACET. In these figures, the plasma density is in blue while the electron beam density is shown in shades of red. The beams are moving in opposite directions in the figure.

IV. Modeling possible LWFA experiments using OSIRIS

Research in LWFA acceleration is now focused on novel mechanisms and configurations to control the acceleration and injection processes so that .1 to 1 nC beams with low energy spread and low emittance can be produced and controlled. This is crucial for the wide use of plasma based accelerators from high-energy-physics experiments, to light sources, and to medical applications.

Using the resources of Jaguar we have used OSIRIS to explore using a transverse static magnetic field to tailor the final self-injected electron bunch properties. The simulations revealed that depending on the gradient of the magnetic field, two distinct injection mechanisms can be distinguished: in uniform field regions injection may occur due to the conversion of perpendicular into longitudinal momentum. In the down ramp regions of the magnetic field, injection may occur due to the expansion of the magnetized plasma wave.

Jaguar simulations showed that the magnetic field can relax the trapping thresholds in laser or particle-beam driven wakefields, allowing for an effective electron self-injection for arbitrary wake velocities. With the additional steering of the longitudinal injection position, the magnetic field can also control the output charge and energy. According to our simulations, and for the next generation PBAs aiming at producing 10-100 GeV electron bunches in external injection scenarios, our scheme requires external magnetic fields as low as 5 T.

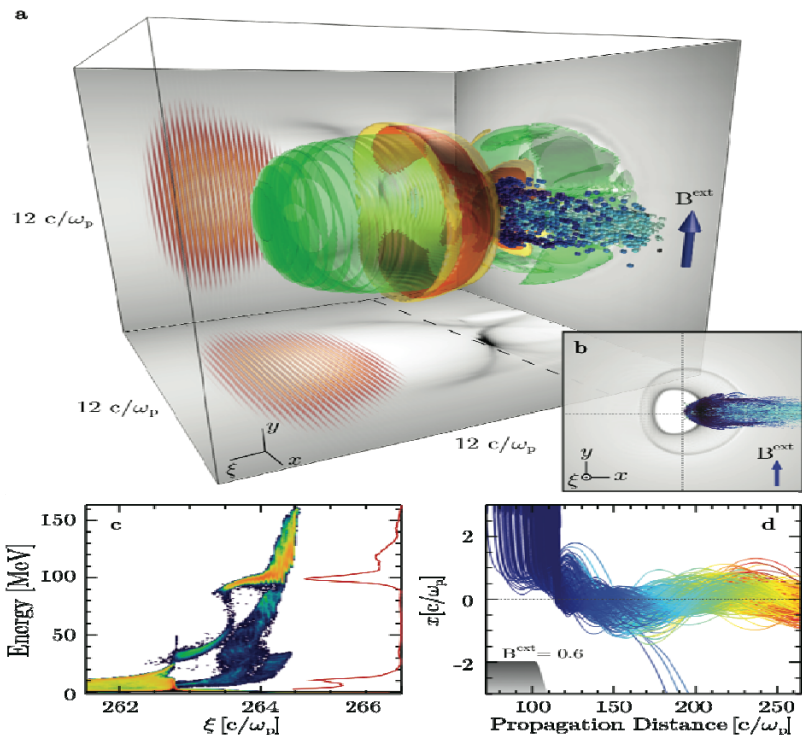

Figure 5. 3D OSIRIS simulation of a magnetized LWFA where the B field rises during 10 c/ω_p, is constant for 40 c/ω_p, and falls down to zero for 10 c/ω_p. **a**. Electron density iso-surfaces in the uniform B field region. Electrons (blue spheres) are self-injected closer to the axis, while the particles with larger radial positions escape from the trapping region. **b**. Transverse density profile (gray) at the region where self-injection occurs (dashed line in **a**) revealing the off-axis injection. **c**. Phase-space of the plasma electrons (blue-red) and spectrum (red line), showing quasi-monoenergetic ~6% FWHM spread) electron bunch. **d**. Trajectories of the self-injected beam particles which illustrate synchronized betatron oscillations (0 MeV-blue, and 125 MeV-red). The propagation axis is shown by the dashed line, indicating an average oscillation radius of $r_0 \sim 0.4\ c/\omega_p$. The magnetic field profile is represented in grayscale.

Numerous three-dimensional and two-dimensional PIC simulations were performed using OSIRIS on Jaguar. Figure 5 shows the results from one of these simulations, chosen in order to illustrate this controlled injection mechanism in the laser wakefield accelerator. A linearly polarized laser pulse with central frequency $\omega_0/\omega_p=20$ (being ω_p the plasma frequency and ω_0 the laser central frequency) was used, with a peak normalized vector potential $a_0=qA_0/m_ec=3$, a duration $\omega_p\,\tau_{FWHM}= 2a_0^{1/2}$, and a transverse spot size matched to the pulse duration such that $W_0=c\,\tau_{FWHM}$. A matched plasma channel was used to guide the laser, and a density ramp ensured a smooth vacuum-plasma transition. For this case, a static external magnetic field pointing in the positive y-direction was used. At the point where the plasma density reaches its maximum value, the external field rises with a $\sin^2$ profile (ensuring a smooth profile, preserving the relevant physics) from zero to $B_{y0}^{ext}=0.6$ (or $\omega_c/\omega_p= 0.6$). The simulation box was 24x24x12 $(c/\omega_p)^3$ divided into 480x480x1200 cells with 2 particles per cell, giving $5.5\text{x}10^8$ particles in total. This figure

clearly shows that the presence of the external magnetic field leads to injection in a well defined angular region, which then leads to synchronized betatron oscillations of the self-injected beam. In addition, the self-injected beam has a charge of 0.15 nC, a peak energy of 100 MeV and an energy spread of only 6 %.

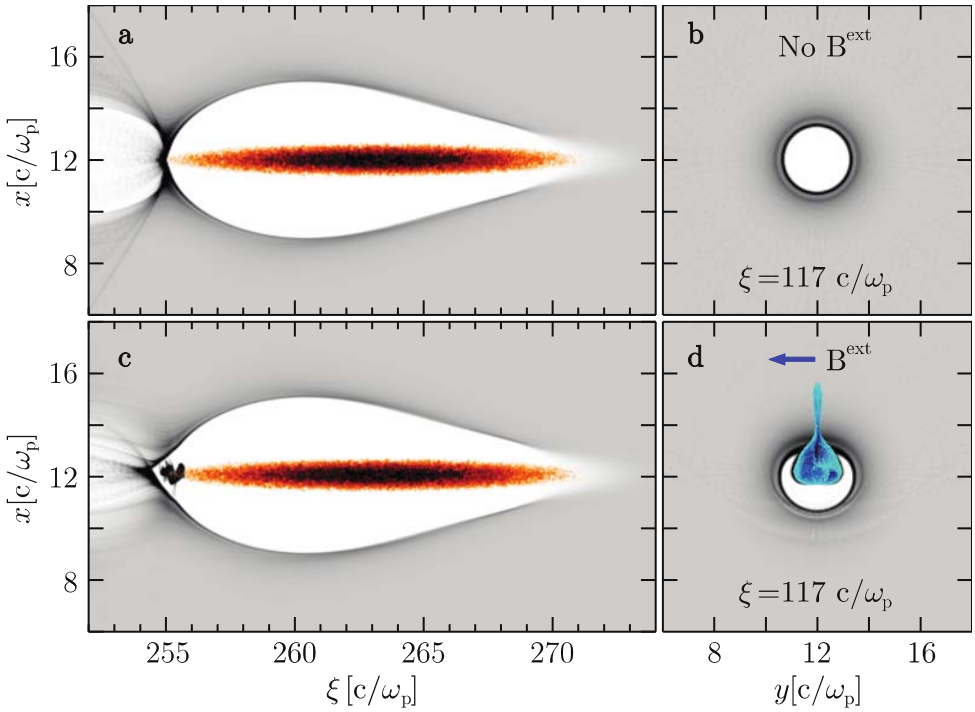

Figure 6. 3D OSIRIS simulation of B-injection in PWFA (2D slices represented). **a**. Electron density, and **b**. transverse density slice at the back of the bubble in the unmagnetized case. **c**. Electron density after the magnetic field down-ramp and **d**. transverse density slice at the back of the bubble, showing trapped particles colored in blue.

This scheme can also be applied to the PWFA, as shown in Fig. 6. A bi-gaussian 30~GeV beam with spot size $\sigma_\perp$=0.17 c/ω_p, and length σ_z=1.95 c/ω_p was initialized in an uniform plasma with a beam to plasma density ratio n_b/n_0=19, chosen to provide a similar r_b in comparison to the 3D LWFA simulation. The magnetic field profile is similar to the LWFA case, with a flat top at B_y^{ext}=0.6. The simulation box was 16x16x20 $(c/\omega_p)^3$, and it was divided into 800x800x800 cells with 2 particles per cell, giving a total number of particles of 1.02×10^9. Despite the strong blowout, injection is absent in the unmagnetized PWFA (Fig. 6a-b).

V. Optimization on multi-core and SIMD architectures

We have been designing new Particle-in-Cell algorithms for advanced multi-core architectures, expected to be important building blocks for future exascale computer

systems. Our goal is to develop algorithms which will be general and adaptable to many different architectures. The algorithms can be applied to advanced accelerators as well as in beam dynamics codes.

To establish a stable framework in this fast-changing environment, we defined a simple hardware abstraction of future architectures. In the abstraction, there are three layers. The lowest layer consists of an abstract SIMD (or vector) unit, which works in lockstep, with fast memory and synchronization. (This abstract SIMD unit can be just a single core.) Multiple SIMD units are then coupled via slow shared memory and synchronization. At the highest layer, multiple nodes of such SIMD units are coupled via MPI and do not share memory. It is assumed that memory access is substantially slower than computation. The architecture closest to this hardware abstraction is a cluster of graphical processing units (GPUs), so our current development is initially focused on that.

We believe that streaming algorithms are best suited to this kind of architecture. Streaming algorithms are those where the bulk of the data is read and written only once. To implement streaming algorithms for PIC, we keep particles constantly ordered by cell, where a cell consists of one or more neighboring grid points. Streaming algorithms minimize memory access and allow very fine grain parallelism. The challenge is to develop optimized particle re-ordering schemes.

As a first step, we have developed a simple 2D electrostatic PIC code on NVIDIA GPUs. The fundamental unit of parallelization (a single thread) is assigned to a sorting cell and the particles which are associated with that cell. The SIMD (vector) length (the number of threads running in lock step) is set to the block size, the number of cores which share fast memory, typically 32-128. To avoid memory hazards (two processors trying to update the same memory location), guard cells are used, just as in distributed memory computers. The guard cells are added up or filled in as part of the field solver.

The first version used only slow global memory and achieved a performance of about 2.3 nsec/particle/time step on the Tesla C1060, a speedup of about 13 over a 2.66 GHz Intel Nehalem (i7) processor. A later version which made use of fast shared memory achieved a performance of about 1.2 nsec/particle/time step, a speedup of about 22.

Maintaining particle order uses a stream compaction techniuque, with two different schemes. For particles moving within a SIMD unit, a thread-racing technique is used. For particles moving between two different SIMD units a message-passing technique is used.

As a second step, we have implemented a 2-1/2D, relativisitic electromagnetic code. This code has more computational intensity (more operations per memory access), and achieved better relative performance, about 2.2 nsec/particle on a Fermi C2050 GPU, for a speedup of about 55. We are continuing work to improve the algorithms.

YEAR 4

I. Optimize QuickPIC on a single processor:

We have continued to improve the performance of QuickPIC on a single processor and improved its accuracy. In QuickPIC, a full Maxwell solver is reduced to a non-radiative EM solver using the quasi-static approximation. As a result, the three dimensional plasma response to the driver (a particle beam or a laser pulse) can be calculated slice by slice along the longitudinal direction. This allows one to convert the 3D problem into a 2D solver where the time advance corresponds to pushing the 2D slice back through the driver in the axial direction. In the solution for one transverse 2D plasma slice, the fields are solved according to the current and charge density in the plasma slice and the particles are pushed in the fields to advance the current and charge density forward in time. Unfortunately, the discretization of the basic equations under the quasi-static approximation shows that the fields at one time step depend on the current and charge density at same time step. Therefore, the leap-frog algorithm, which is widely used in the full PIC codes, will not work in QuickPIC. Instead, an iteration (predictor-corrector) method is applied for solving those coupled equations. Generally speaking, in the 2D plasma calculation procedure, predicted fields are used to push the particles and generate predicted plasma current and charge densities. These current and charge densities are used to solve for corrected fields which are then used to calculate the corrected current and charge densities. This iteration can be carried out as many times as is needed.

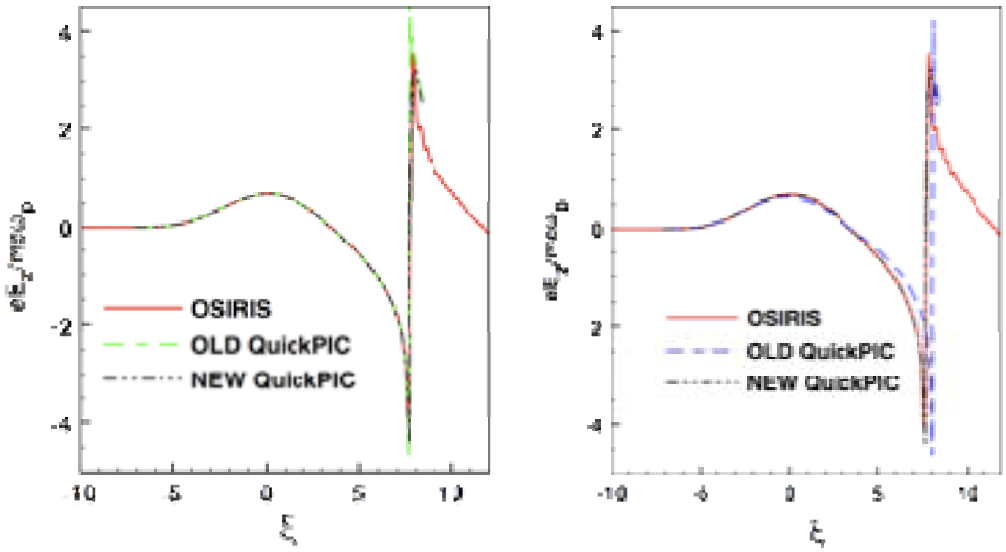

(a) With 8 Iterations (b) With 1 Iteration

Fig 1. Simulation results of longitudinal electric fields in an electron beam driven PWFA. The drive beam's spot size is σ_r = 1 μm, pulse length is σ_z = 30 μm and the particle number is N = 3.0 x 10^{10}. The plasma density is 1 x10^{17} cm^{-3}.

In the original implementation of QuickPIC the scalar and vector potentials are solved for in the Lorentz gauge. In practice, we often needed 5 or 6 iterations for very nonlinear

problems. During the past year, we have experimented with a new field solver (and iteration loop) which uses the transverse Coulomb gauge. In the last report, we mentioned that preliminary results indicate that only 2 iterations are required which will lead to a savings of more than a factor of 2. We have now improved it so that only 1 iteration is needed for a wide class of problems. One example is shown in Fig. 1, where we plot the accelerating field with the old field equations and with the new field equations for a highly nonlinear case. The results show that with enough iterations (which is 8 in our example) both algorithms will lead to the same solution (and they also match with the OSIRIS result), while for the new algorithm a very accurate solution is found after only one iteration.

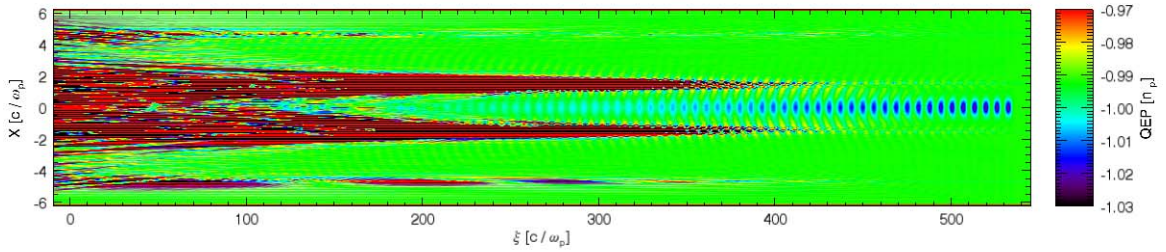

(a) Old QuickPIC result

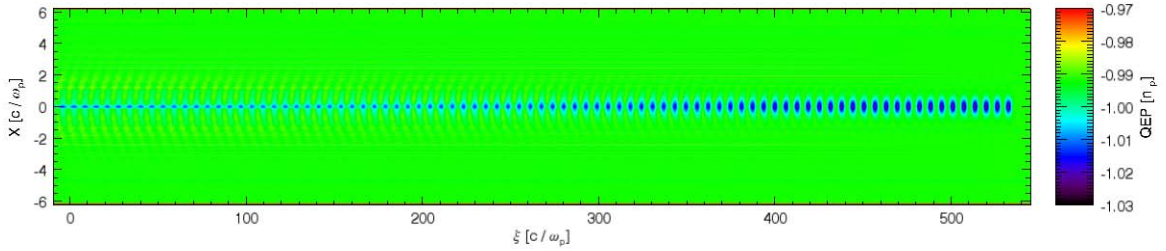

(b) QuickPIC result with full Fourier domain field solver

Fig 2. Plasma density of a long proton beam driven plasma wakefield. The proton beam is a half-cutoff Gaussian beam with higher density in the head and lower density in the tail. The beam's spot size is σ_r = 200 μm, pulse length is σ_z = 12 cm and the particle number is N = 5.75 x 10^{10}. The plasma density is 1 x10^{14} cm^{-3}.

Another improvement is achieved by implementing the whole field solver in Fourier domain avoiding frequent transforms to and back from Fourier domain. In this way, half of the Fourier transformations inside the solver are reduced, which leads to a significant time saving according to our tests. Another important benefit from this full Fourier domain field solver is the accuracy improvement. Fig. 2 shows an example of a long proton beam driven plasma wakefiled in the linear region (The beam is moving to the right). The length of the simulation box is larger than 500 plasma skin depths in this case. When using the old field solver, significant noise can be seen in the plasma density (beginning 100 plasma skin depths behind the head of the beam). But when using the full Fourier domain field solver, the noise is completely absent.

II. Provide QuickPIC and OSIRIS simulation support for the PWFA FACET experiments:

Next generation experiments at the FACET facility which is under construction at SLAC will be aimed at showing that a trailing bunch of electrons can be accelerated with high transfer efficiency while maintaining its emittance and energy spread. The FACET facility will have electron and positron bunches at 23 GeV with 3nC of charge. In addition, the beams can have bunch lengths as short as 14mm which corresponds to peak currents of 22 kA. Designs of these experiments indicate that two bunches with ~1nC in the drive beam and ~.5nC in the trailing beam and a separation ~100 mm are possible. These conditions were simulated with QuickPIC and OSIRIS to find the optimum density and the some results were reported in the previous report.

The use of a field-ionized plasma source greatly simplifies the experiment. The conversion of a single bunch into a drive and trailing bunch leads to a loss of charge and a loss of peak current. For the FACET parameters, initial experiments will operate with at peak current that is near the ionization threshold. Therefore, the ionization module becomes significantly important for modeling those experiments. We have rewritten the ionization module in QuickPIC to make it merge with the pipelining algorithm and the most updated 2D solver. This ensures that the QuickPIC with field-ionization can also scale to more processors, run faster and be more accurate. Fig. 3 and Fig. 4 show the comparison between QuickPIC and OSIRIS.

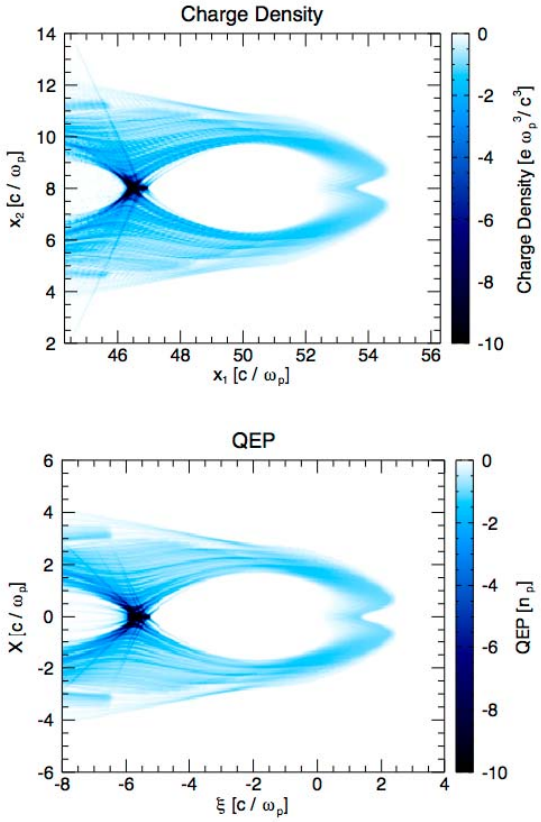

(a) OSIRIS result (b) QuickPIC result

Fig 3. Plasma density of an electron beam driven PWFA in a field ionized plasma. The drive beam's spot size is σ_r = 10 μm, pulse length is σ_z = 34.1 μm and the particle number is N = 9.6 x 10^9. The plasma density is 5 x10^{16} cm^{-3}.

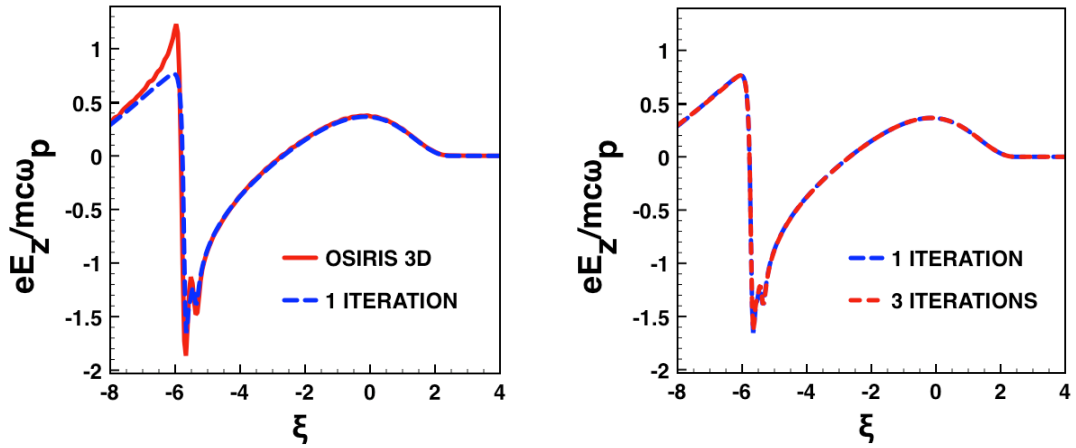

(a) (b)

Fig 4. Accelerating field of an electron beam driven PWFA in a field ionized plasma. (a) Comparison between OSIRIS and QuickPIC (with 1 iteration);(b) Comparison between 1 and 3 iterations using QuickPIC. The drive beam's spot size is σ_r = 10 μm, pulse length is σ_z = 34.1 μm and the particle number is N = 9.6 x 10^9. The plasma density is 5 x10^{16} cm^{-3}.

In addition, multiple ionization based on the Ammosov-Delone-Krainov (ADK) model has also been implemented into the QuickPIC ionization module instead of only single ionization. This is required to model Cs as a plasma source instead of Li (which is used before) because Cs has a much lower threshold of second ionization. The new ionization package will help to study the influence of multiple ionization.

III. *Carry out high resolution QuickPIC simulations of PWFA-LC stages on Jaguar*

In preparation for using the new QuickPIC algorithms for large simulations on Jaguar we have been merging them with the pipelining and advanced parallel decomposition routines that have been reported previously. We are now testing the merged code. We have implemented the 2D decomposition in 3D domain and pipelining algorithm to make QuickPIC scale to more than 10,000 processors. This ensures QuickPIC to simulate high resolution PWFA cases. But then we found that the diagnostic algorithm became an important issue, which can affect the code performance. In the old algorithm, we used a single processor to write output data file. Before doing this, the data managed by other processors will be gathered to this single processor. When the simulation system increases, this process of gathering all other processors' data to a single one will become very slow. In order to keep QuickPIC still in a good level when using more processors,

we have upgraded the diagnostic module by using HDF5 parallel library instead of HDF4. With this parallel algorithm of diagnostic, QuickPIC shows great performance when simulating high resolution PWFA problems with more than 10000 processors. In addition, this high performance version of QuickPIC has also been equipped with the new 2d solver introduced in the first section. Now we are able to run QuickPIC more efficiently on large scale simulations.
Once it is working, we will commence with high resolutions simulations described in past reports of a single stage of the current PWFA-LC concept. These simulations will uses a 8192x8192x2048 grid. The matched spot sizes for the witness beam in this design is ~10-100 nm while the radius of the accelerating structure is ~10um. This disparity in scales requires the use of many grid points. In addition, at these small spot sizes the density of the witness beam is $\sim 10^4$ to 10^5 times larges than the plasma density. As a result the ions will move which will degrade the focusing and acceleration fields. The improved version of QuickPIC together with Jaguar's capabilities will allow us to study the physics of beam quality preservation of narrow witness bunches for the first time.

IV. Optimize OSIRIS for SSE and vector units and scale to full Jaguar computer and continue the development of PIC algorithms for GPUs :

We have been designing new Particle-in-Cell algorithms for advanced multi-core architectures, expected to be important building blocks for future exascale computer systems. Our goal is to develop algorithms which will be general and adaptable to many different architectures. The algorithms should be able to be applied to advanced accelerators as well as in beam dynamics codes. We have also been working on scaling OSIRIS out to 200,000+ Jaguar cores for self-trapped LFWA simulations.

The code OSIRIS was chosen to be part of the OASCR Joule Metric program for this year as an accelerator physics code. The problem we targeted to study was generating 1-10 GeV, high beam quality, electron beams in the self-injected and self-guided regime of LWFA. This problem was chosen because studying self-injection requires full particle-in-cell simulations as well as large computational resources. While there are various self-injection scenarios including using intense lasers in a self-guided regime, e.g., using field-ionization by the laser, using density gradients, and using multiple lasers, from a high performance computing perspective their is much overlap in the simulation requirements to study these scenarios. Therefore, the high performance computing lessons learned from studying one injection method will apply to other injection methods.

A critical issue in performing these simulations on 10000 or more processing elements (PEs) is that of parallel load balance. The self-injection mechanism itself places particles into a local region of real space. This leads to a severe accumulation of particles in a

very narrow region at the back of each plasma wavelength. If a large number of PEs is used and the region where particles accumulate resides on a very small number of the total number of PEs then a severe load balance will arise. Furthermore, the position of these particles in the simulation window moving at the speed of light, will not remain constant so that any redistribution of particles across the PEs will have to be done regularly during the simulation.

To study these effects, we chose a set of 4 problems for an initial study. The first problem models a uniform warm plasma with the goal of measuring code performance under a perfectly load balanced simulation. In this scenario, particle diffusion across parallel nodes happens uniformly so the total number of particles per node remains approximately constant. This is also a good overall performance test, as these plasma conditions will resemble those on most of the simulation box for the laser wakefield runs.

For the laser wakefield scenarios we chose 3 problems. Two of them correspond to the interaction of a 200 TW (6 Joule) laser interacting with uniform plasma with a density of 1.5×10^{18} cm^{-3} plasma with an intensity sufficient to trigger self-injection, under different numerical and physical conditions. We chose different grid resolutions, different number of particles per cell, and mobile/immobile ions. The third LWFA run corresponds to a PW (30J) laser propagating in a 5×10^{18} cm^{-3} plasma.

We used quadratic shaped particles for the current deposition and field interpolation for all the simulations, as this gives the best tradeoff between reducing aliasing at the expense of higher floating point counts for these problems. Finally the simulations were stopped before the complete details of the observables could be measured in order to fit in a 1 million cpu hour limit. However, each simulation ran long enough for key physical processes (such as self-injection) to occur and to capture the relevant flow of the computation. These simulations will be completed later to obtain all the physics results.

The warm plasma test performed quite well yielding a global performance of 7.62×10^{10} particle pushes per second on 55,000 PEs, which gives an average of 1.38 million particles pushes per second per PE (for quadratic interpolation and current smoothing). The performance of the code did not fluctuate over the length of the simulation. The average load imbalance, measured as the ratio between the maximum number of particles in a PE over the average number of particles in a PE was 1.00011, effectively maintaining a perfect load balance over the full course of the simulation, as expected. The performance measurement therefore set a baseline for comparison with other scenarios where the load imbalance will hinder code performance.

The performance of the laser wakefield simulations was, as expected, lower than the reference warm plasma test. This was due mostly to significant load imbalance that occurs as a result of the physics of the underlying problem. The average load imbalance over the 3 runs ranged from 18.8 to 78.6 and the overall code performance was reduced by the same factors, when compared to the warm test. There is a significant accumulation of particles in a small number of nodes that correspond to the simulation region at the back of the accelerating buckets, with a maximum of 1.37 million particles

in this node, compared with an average of only 55.6 thousand particles in all nodes. Using different initial domain decompositions where is broken up into fewer partitions along the laser propagation direction reduced the load imbalance (for one case to below 10) and we note that if 50 times few PEs were used then the load imbalance could be ~ 2.

During the second half of this year, we will experiment with different techniques to improve the load imbalance including using a mixture of OpenMP and MPI and using the dynamic load balancing capability within OSIRIS. On Jaguar there are only 12 PEs per node so we will adopt a strategy in which we will have twelve copies of the current on each node. We will also experiment with the use of the SSE capability within OSIRIS to see if the single core performance can also be improved without loss of accuracy.

During the first half of this year, we have also continued our investigation of algorithms for PIC codes on many core architectures. We investigated the portability of our GPU algorithms (described in previous reports) to other languages and architectures. We began by translating the Cuda C program into OpenCL. OpenCL matches our hardware abstraction very well and the kernel routines in Cuda can be translated directly to OpenCL without conceptual difficulty. However, OpenCL is a very difficult environment to work with at present. The language is complex and has no debugging facilities. The OpenCL version of the code was able to run with three different OpenCL implementations: from NVIDIA, Apple, and AMD. On NVIDIA hardware, OpenCL is slower than Cuda C, typically 20-60%. We also translated the GPU code to PGI's Cuda Fortran. This language is higher level than Cuda C and is easier to use. Performance was about 10% slower than Cuda C.

We have also been translating traditional PIC algorithms into OpenMP as well as mixed MPI/OpenMP without duplicating the current or charge. These algorithms use domain decomposition as in MPI, and the procedures which do not communicate are similar to those in MPI. However, the procedures which communicate make use of the additional information available in a shared memory architecture. Results so far indicate that OpenMP can give results very similar to MPI, and that a mixed MPI/OpenMP program looks promising in improving performance of global communications and in maintaining local load balancing.

V. Provide OSIRIS simulation support of present and future LLNL/UCLA LWFA experiments:

During the last half of the year, we have been modeling LWFA experiments for the LLNL Callisto laser. These experiments have a two-stage design. The first stage, which we designate as the Injection regime, contain a mixture of 99.5% He and 0.5% N, whereas the second stage only consists purely of Helium. Electrons are ionization-injected in the first stage, which then propagates into the second, longer stage to form a mono-energetic beam.

Typical simulation parameters consists of a 60 fs, 40 TW laser with $w_0 = 15.0$ μm. The plasma itself has a peak density of $n_p = 3 \times 10^{18}$ cm^{-3}. The Injection Regime consists of a 1 mm upwards ramp, a 1.5 mm flat, and a 2.0 mm mixing region where the Nitrogen population drops until the gas is purely of He, with total plasma density remaining constant. The Accelerator Regime is simply a flat distribution of He plasma over the course of 6 mm. The laser is focused at the top of the up ramp of the injector regime.

We discovered that near the end of the injection regime (2.7 mm into plasma) a mono-energetic peak forms at 275 MeV of about 360 pC. After traversing a substantial distance into the accelerator regime (8.6 mm into plasma) there is a peak at 535 MeV of about 880 pC. These peaks are filtered for particles which satisfy $(p_2^2 + p_3^2)/p_1^2 < 0.005^2$, since the small-angle particles are what are physically observed by detectors. The increase in charge is indicative of a redistribution of particles in phase space, rather then the injection of additional electrons. The plots of the filtered spectrum are given below in figure 5.

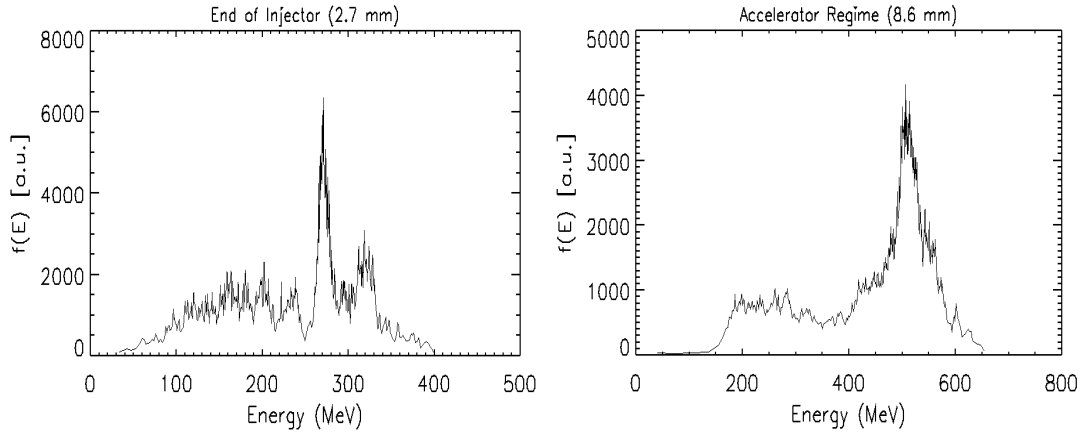

Figure 5. Electron spectra within a half angle of 5mr after the injector and after the full acceleration length.

Although this had been the intended effect, detailed analysis of the data show that the dynamics of the process is highly complicated. Most notably, at the end of the simulation the charge injected into the second bucket enter first bucket and begin to accelerate. In a 2D simulation with identical parameters, this population of electrons in fact accelerates to energies above those that were injected into the first bucket. In addition, it may be possible that by reducing the Nitrogen doping in the injector regime, one may trap fewer charge and reduce the distortion in the wake – allowing the particles in the second bucket to accelerate differently.

YEAR 5

I. Developing PIC algorithms for clusters of GPUs and next generation many cores :

During the past year, the research into developing Particle-in-Cell (PIC) algorithms for Graphical Processing Units (GPUs) has focussed on two main areas. The first area is enabling PIC codes to run on a cluster of GPUs. The second is the evolution of GPU PIC algorithms as the GPU hardware changes.

In prior years of this project, we developed optimal PIC algorithms for a single GPU [Decyk1], and achieved speedups compared to a single CPU of about 20 for a 2D electrostatic model and 55 for a 2-1/2D electromagnetic model. The optimal algorithm for the GPU divided space into small tiles (typically 2x3 or so) for each thread, and particles were reordered every time step. The partitioning was similar to that used by traditional distributed memory algorithms, except the domains were much smaller (micro-domains) and data between micro-domains could be read directly without sending messages. The reordering of particles on the GPU had some similarities to how the particle manager on clusters reordered particles using MPI. To implement such designs on a cluster of GPUs requires that two different domains had to be integrated, one micro-domain used by GPUs and a larger domain used by MPI.

For the first part of the project year, we developed a prototype system with a hierarchical domain decomposition, based on MPI and OpenMP, using a 2D electrostatic code from the UPIC Framework. With a traditional cluster, one assigns one MPI node to each core, so that a multi-core node with N cores would have N MPI nodes. In the new hierarchical scheme, we would assign one OpenMP thread to each domain on the core, and connect the domains at the edges with MPI. There could be one or more OpenMP domains in a single multi-core. The OpenMP domains could be later replaced by GPU domains. Unlike earlier work with OpenMP/MPI hybrid systems, the OpenMP domains also partitioned memory, but did not use message-passing to read or write the data controlled by adjacent threads. This worked very well, and had an important advantage for spectral codes because FFTs scaled substantially better. Some details were presented in [Decyk2].

A hierarchical domain decomposition for GPU/MPI is now under development. Currently, all the particle subroutines have been written and tested, and the GPU reordering and the MPI particle manager have been integrated in an optimal way. Currently a cluster of 3 GPUs gives a speedup compared to a single CPU of about 60 for a 2D electrostatic model on the Dawson2 cluster at UCLA, which uses NVIDIA M2090 GPUs. The field solver has not yet been implemented, and additional optimizations present in the single GPU code have not yet been ported to this hybrid GPU/MPI system, so we expect additional improvements.

The GPU architecture has also been evolving. The current NVIDIA Fermi architecture now has cache, a larger shared memory, and hardware atomic operations. As a result, simpler algorithms which were not optimal on the earlier architecture now work quite well. Our second research focus this year was to reexamine our earlier work in light of

the newer architecture. As a result, we found that a scheme with larger tiles with a block of threads per tile was more optimal on the new architecture. This scheme made use of faster floating point atomic operations. The original scheme used smaller tiles with one thread per tile, and did not require any atomic operations. Both schemes use a similar algorithm, but implemented differently. The new scheme was faster on the newer architecture, and original scheme was faster on the original architecture. A newer NVIDIA architecture is just coming out, called Kepler, and additional design changes may be needed for optimal operation.

Decyk1] Viktor K. Decyk and Tajendra V. Singh ,"Adaptable Particle-in-Cell algorithms for graphical processsing units," Computer Physics Communications 182, 641 (2011).

[Decyk2] Viktor K. Decyk and Tajendra V. Singh, "Hierarchical Domain Decompositions for Particle-in-Cell Codes, Proc. of the Plasma Conference 2011 (PLASMA 2011), Kanazawa, Japan, November, 2011.

[Decyk3] "Particle-in-Cell Simulations on Modern Computing Platforms", 22nd International Conference on Numerical Simulation of Plasmas, Long Branch, New Jersey, Sept. 8, 2011.

[Decyk4] "Particle Simulations on GPU supercomputing systems." 21st Intl. Toki Conference, Toki-City, Gifu, Japan, December 1, 2011.

[Decyk5] "Development of Numerical Algorithms for Next Generation Supercomputers," 2011 US-Japan Workshop on The Next Stage in the Progress of Simulation Science in Plasma Physics, Toki-City, Gifu, Japan, December 2, 2011.

[Decyk6] "Developing PIC Codes for the Next Generation Supercomputer using GPUs." Workshop on Computational Methods in High Energy Density Plasmas, UCLA Institute for Pure and Applied Mathematics, Los Angeles, CA, March 29, 2012.

[Decyk7] "GPU as a Protoytpe for Next Generation Supercomputer." Workshop on High Performance Computing in High Energy Density Plasmas, UCLA Institute for Pure and Applied Mathematics, Los Angeles, CA, May 14, 2012.

II. Modeling of LWFA in a Lorentz boosted frame using OSIRIS and UPIC :

During the past year, we have continued to study the use of a Lorentz boosted frame to study LWFA. We continue to work on a moving antenna for launching the laser, on developing a method for initializing a trailing particle beam for studying beam loading, and on understanding the noise that can be prevalent when the plasma drifts across the grid. The noise is the result of a numerical instability. We have developed a theory for this instability that predicts the location of the unstable mode numbers and their growth rates. The instability is due to the coupling on a plasma beam mode with electromagnetic modes with numerical phase velocities less than the speed of light and of aliased plasma beam modes with light waves. We have also made extensive studies of how the location

of the unstable modes depends on the field solver. It was found that spectral codes (schemes with light waves with numerical phase velocities greater than light) have advantages since only the aliased modes occur and these occur at high wavenumbers that can be filtered. As a result we have been developing the capability of modeling LWFA with the UPIC framework. This work was presented at the recent Advanced Accelerator Concepts Workshop.

III. Large Scale OSIRIS Simulations of Current and Future Plasma Accelerator Experiments:

We have been performing OSIRIS simulations of current and future experiments of both PWFA and LWFA.

IIIa. Protron driven PWFA

If currently available protron bunches could be compressed then they could be used to excite relativistic plasma waves leading to electron acceleration up to the ~ 1 TeV energy barrier. However, it is not currently possible go compress existing proton bunches so wakes are exciting through long proton bunches self-modulating. In this case, instead of exciting large amplitude plasma waves from the beginning of the propagation, plasma wakes are excited through a self-modulational instability. Full scale OSIRIS simulations were performed to explore the self-modulation of long particle bunches in conditions that are relevant for experiments. The linearized quations for self-modulation of particle beams are self-similar between using proton, positron, and electron beams. Therefore, aspects of the self-modulation of a proton beam can be studied in experiments on the self-modulation of electron and positron beams. Therefore, we have used OSIRIS simulations to design possible self-modulation experiments of electron beams at FACET. We have also been simulating possible experiments at CERN and FNAL on the self-modulation of proton beams.

Aii. Self-modulation of lepton bunches

The self-modulation of electron/positron bunches and proton/anti-proton bunches are the same if one scales the propagation distance to the respective betatron wavelengths. Therefore, we have performed simulations of the self-modulation of electron/positron beams to learn about the self-modulation of proton beams. The electron bunch parameters were chosen in order to be close to those available at SLAC FACET. We then modeled 20 GeV lepton bunches with $2x10^{10}$ leptons per bunch, and with σ_z = 500 mm long and σ_r = 10 mm wide. Plasmas with densities in the range of $n_0=10^{17}$ cm^{-3} were also used. The simulations showed that the self-modulation instability saturates in less than 10 cm for well controlled noise sources. Large accelerating gradients above 20 GV/m were observed. In addition striking differences between the self-modulation of electrons and positrons were found. This was attributed to the fact that the blowout regime was reached for which electrons are accelerated and focused and positron beams are defocused. Figure 1 shows a simulation result from a 2D cylindrically symmetric

OSIRIS simulation illustrating a self-modulated FACET electron bunch.

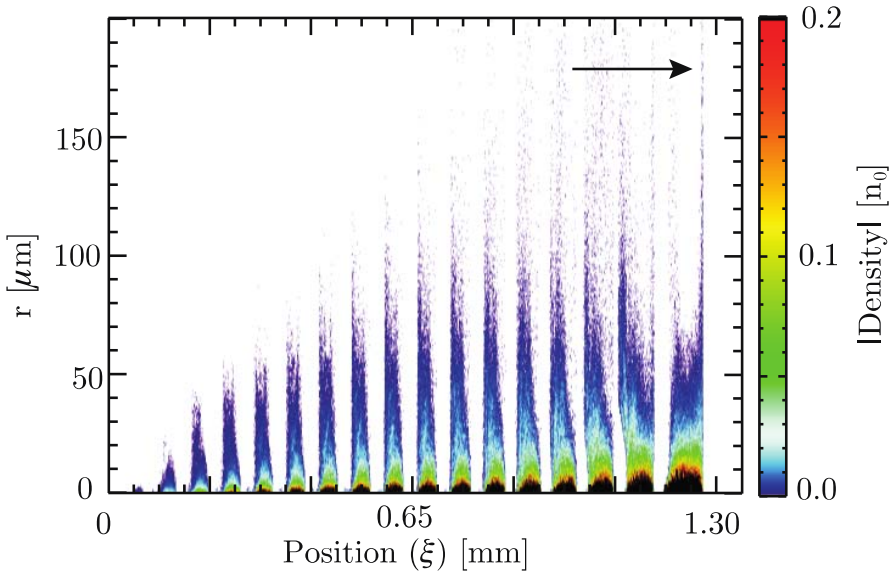

Figure 1: Cylindrically symmetric simulation results illustrating the density of a long SLAC electron bunch after propagating in an 1 meter plasma. The bunch travel from left to right. Several electron beamlets are visible at the plasma exit. This is a clear illustration of the self-modulation instability. The simulations also revealed generation of peak accelerating gradients above 20 GeV/m, and energy gain and loss by the driving electrons of several GeVs. These results suggest that the propagation of long electron bunches in plasmas provides an excellent surrogate for a proton driven plasma wakefield accelerator experiment.

Aiii) Proton Driven Plasma Wakefield Acceleration

We have also simulated the self-modulation of proton beams using parameters of relevance to both Fermilab in the US and CERN in Europe.. Due to the long pulse length (~2 x 10^3 larger than the electron collisionless skin depth, c/ω_p), the proton beam may be self modulated when propagating inside the plasma, which will result in a large amplitude of the wake field. Our current study is focusing on how the proton beam will evolve when the self-modulation happens.

OSIRIS in its 2D cylindrical coordinate mode was used for the simulation. We also updated the code to enable the beam initialization with Twiss parameters (which describing the transvers phase space distribution of the beam). This is important for the simulation because the proton beam's pulse length is comparable to the beam's β^*. Figure 2 shows snapshots of a proton beam's charge density at the time when the beam's center has propagated into the plasma for 62cm's. As shown in the figure, the proton beam has undergone a significant amount of self modulation, resulting in the micro bunching in the middle and rear part of the beam. The distance between each two micro bunches is equal to the plasma wavelength. Fig. 3 shows the energy modulation on the

proton beam at the same time as that of Fig. 2. The maximum energy change at this time is around 200 MeV.

A typical 2D simulation like those shown here uses ~10 million cells and ~200 million simulation particles. The simulation uses 18,000 cores on Jaguar and requires 100,000 cpu hours.

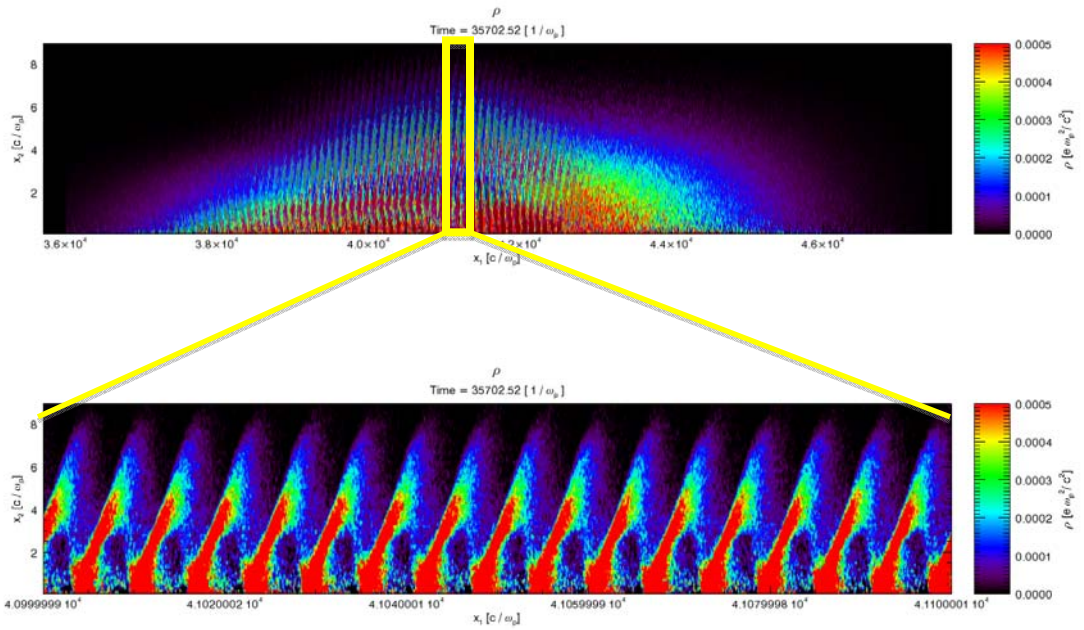

Fig. 2 Snapshot of a self modulated proton beam at the time when the beam center propagating in the plasma for 62 cm. The total beam particle number is N = 1 x 10^{11}; the initial beam focus spot size (r.m.s. value) is σ_r = 100 μm; the initial beam length (r.m.s. value) is σ_z = 10 cm; the initial beam emittance is ε_N = 3.33 mm mrad; the beam is focused at 60 cm inside the plasma. The initial plasma density is 1 x 10^{16} cm^{-3} . (The beam density is normalized to the initial plasma density.)

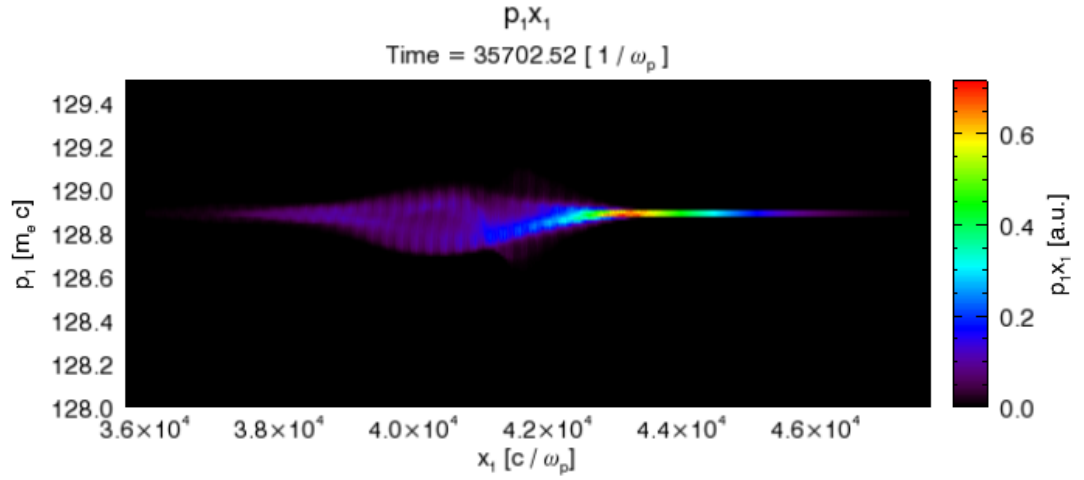

Fig. 3 The proton beam's longitudinal phase space distribution at the time when the beam center propagating in the plasma for 62 cm.

Seeding the self-modulation instability is critical to control experimental outputs and to achieve reproducible results. One idea is to use half-cut proton bunches. However, due to their large relativistic mass on the order of 0.5 TeV, the generation of hard-cut proton bunches is experimentally very demanding. Therefore, we have explored the possibility to use an ionizing laser pulse to seed the instability. Since full PIC simulations are very demanding because the laser wavelength (1 micron) needs to be resolved we employed a ponderomotive guiding center approximation (PGCA) for the laser solver [gordon:00]. The PGCA was implemented into OSIRIS this past year. Under the PGCA the laser evolution is obtained through an envelope equation for the laser vector potential, thus strongly relaxing the computational requirements.

To demonstrate the feasibility of the scheme we have modeled the propagation of an ionizing laser pulse located at the center of the proton bunch. Several tests were performed varying the ionizing laser parameters. Fig. 4 illustrates simulation results using an ionizing laser with normalized vector potential is a_0=0.1. The transverse spot size is W_0=1 mm and was chosen to obtain a 5 meter Rayleigh length. The laser duration is τ_l = 5.9 ps. The laser propagates in a plasma with $n_0 = 10^{14}$ cm^{-3}. A 10 cm long, 0.5 TeV SPS-like proton bunch with σr = 200 microns and 10^{11} particles was considered. The simulations demonstrated that this scheme can be used to ionize the plasma and to effectively seed the instability. The ionization front acts as if the front of the proton bunch is sharply cut therefore creating an initial wake amplitude similar to that provided by hard-cut bunches. A typical simulation uses a 27,200x 320 grid with 4 particles per cell and is run for ~10^6 time steps.

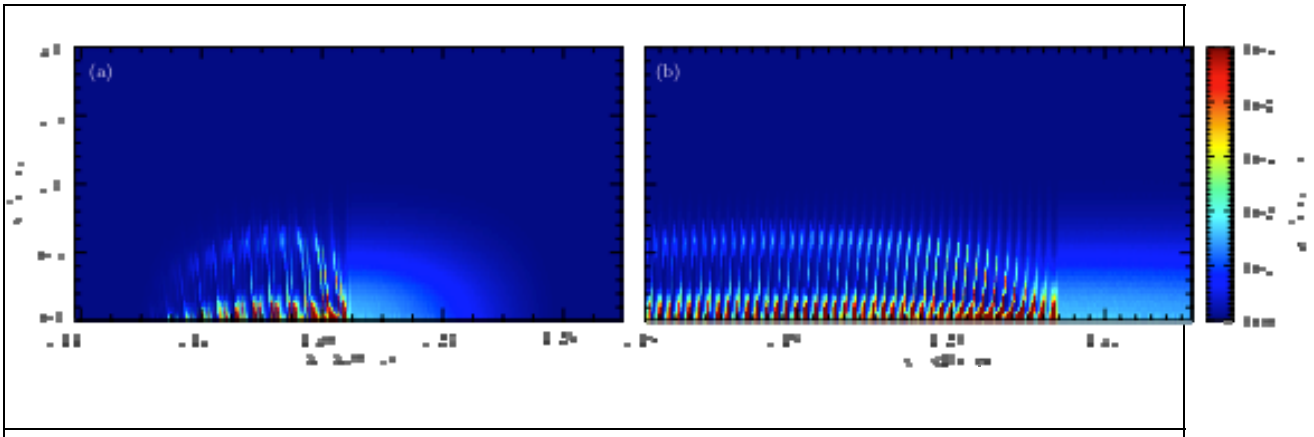

Figure 4: Cylindrically symmetric simulation results illustrating ionization seeding of the proton driven plasma wakefield accelerator using a ponderomotive guiding center solver for the laser fields. The pictures show the proton bunch density. The laser is placed at the middle of the proton bunch, and propagates to the right. (a) full proton bunch density profile. (b) zoom of the proton bunch density near the location where the laser propagates. The SMI only grows from the region of the laser to the back of the proton bunch clearly demonstrating that an ionizing laser pulse can be used to seed the SMI in

the conditions of the proton driven plasma wakefield accelerator.

Simulations of self-modulation of proton beams are very challenging. Currently full PIC simulations in 3D of these processes are not feasible. We have also been using QuickPIC to study 3D effects of self-modulation of proton beams such as those at FNAL. QuickPIC is a 3D PIC code using quasi-static approximation, which is usually 1000 times faster than the ordinary PIC code for special problems like PWFA. In figure 5 show a comparison between Osiris 2D (in cylindrical coordinates) and QuickPIC simulation results of a proton beam self modulation. Above we described OSIRIS simulations of possible proton beam self-modulation for parameters at CERN in which an ionizing laser was used to trigger the self-modulation. In these sets of simulations we model parameters for a possible experiment at FNAL.

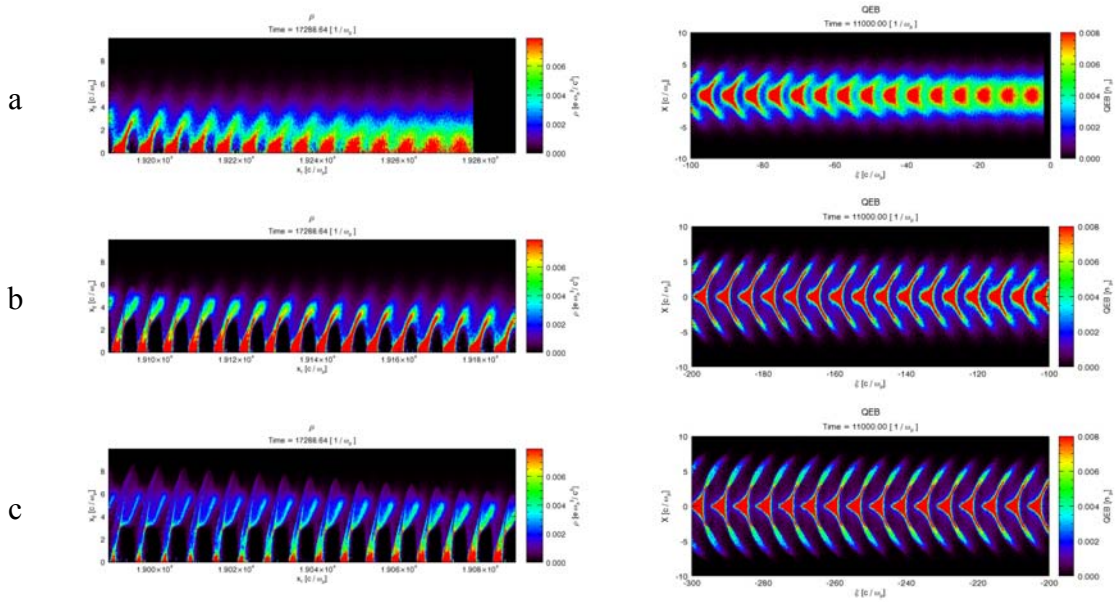

Figure PWFA:5: Proton beam density snap shots from Osiris 2D and QuickPIC simulations. Part a, b and c are three successive parts of the proton beam, in which the part a is the beam head.

From the beam density plots, we can find that two different simulation codes give similar results. But we can see the difference especially in the part c. The reason for that is two codes are using different coordinates for describing the longitudinal position and the time. OSIRIS uses (z,t) but in QuickPIC it becomes (ξ=ct-z, s=z). So that in the QuickPIC simulations, all the beam particles have the same "z" but not the same "t". If the beam is very short and the transverse velocity of the beam particle is much smaller than c, the difference will be negligible without changing the coordinates. But in the proton beam self modulation case, the beam pulse length is very long and the results from two codes cannot be compared directly. More detailed comparison between two codes with the same coordinates will be done in the future. A typical 2D OSIRIS simulation uses a 120000x100 cell grid and follows 2x10^8 particles for ~10^6 time steps. A typical

QuickPIC simulation uses a 262144x128x128 cell grid with 4 particles per cell and ~10^4 3D time steps.

III b. OSIRIS Simulations of two stage LWFA's using ionization induced self-injection

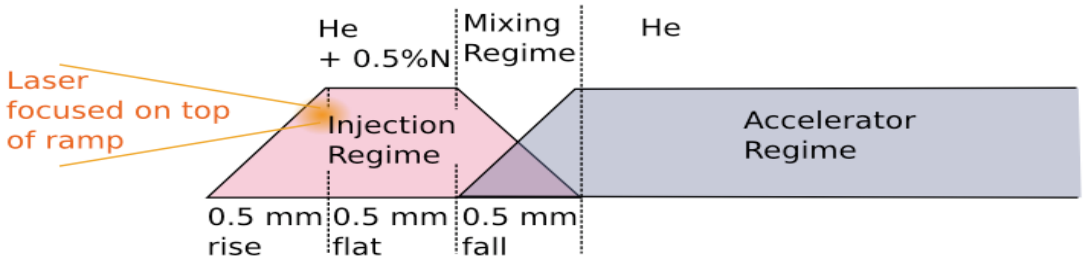

Figure 6: Here the general layout of the two-stage ionization injected laser wakefield accelerator is presented. The Injection Regime is 1.5mm long total. We have run the laser 6.5 mm into the Accelerator Regime (8 mm total into the plasma).

In order to explore methods of injecting a high quality, mono-energetic beam of electrons into an accelerating wakefield, we performed large scale 2D simulations of a two-stage laser wakefield accelerator (LWFA). This concept is shown in figure 6. The two stages are the Injector Regime (1.5 mm long) and the Accelerator Regime (21.5 mm long). The injector Regime consists of a gas (total density $1x10^{7}$ cm^{-3}) composed of 99.5% He and 0.5% N. The 6th and 7th ionized electrons of the Nitrogen atoms inject into the wakefield produced by the He electrons and the rest of the N electrons.

We simulated this scheme with lasers of 90 TW, 100 TW, and 500 TW. The 90 TW and 100 TW have slightly mismatched spot sizes, chosen to correspond to the experimental parameters of the Callisto laser. Results from the 90TW case are shown in figure 7. Both having a beam width of 15 µm, the 90 TW laser is 20% mismatched and the 100 TW laser is 12% mismatched. In contrast the width of the 500 TW laser was chosen to be 26.2 µm, which is perfectly matched. We can simultaneously study the energy gain under increasing laser energies, as well as the changes that are due to only slight variations of spot size mismatch unexplained by mere changes in energy. After 8 mm into the plasma, the 90 TW run reached a beam energy of 0.68 GeV with a 32 % energy spread, whereas the 100 TW only devolved to a 28% energy spread after having traversed 22.8mm, having an energy of 1.4 GeV. In order to retain a good energy spread across a larger region of plasma, and thereby achieve a higher usable energy beam, it is especially important to match the spot size.

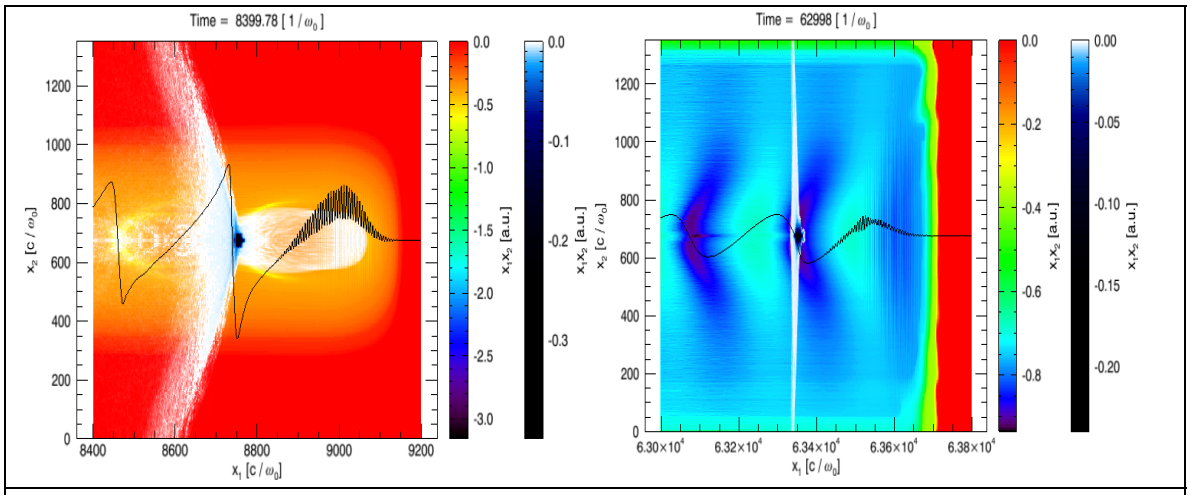

Figure 7. This figure on the left shows the trapping process 1mm into the plasma (in the Injection Regime, of the 90 TW simulation). The rainbow scale showing the density of the He electrons (forming the wake), the blues showing the Nitrogen electrons of the 6th and 7th energy levels (forming the trapped particles), and the E_1 field lineout. The Figure on the right shows the simulation after traversing 8mm. Due to spot size mismatch, the shape of the wake bubble is significantly devolved.

All of these simulations ran 6 particles per cell. The 90 and 100 TW simulations ran 4000 x 300 x 300 cells over 7200 processors. The 500 TW simulation is running 4000 x 360 x 360 cells, over 16000 processors. The simulations are run for ~10^6 time steps.

IV. OSIRIS and QuickPIC simulations of current and future PWFA experiments at FACET

Many simulations have been run for the PWFA experiments at FACET. One important improvement is that we can import the beam particle data (positions and velocities) from the accelerator simulations done at SLAC. Instead of simply assuming a Gaussian profile of the beam, this method will make the QuickPIC simulation much closer to the reality. The following plots show an electron beam imported into the QuickPIC.

The beam has 45.1 □m and 9.8 □m r.m.s. spot sizes in two transverse directions. There is a little bump at the beam tail when integrating the beam profile in the transverse directions (figure 8). But this bump in the tail is very diffuse as shown in the 3D beam density plot (the blue plots on the wall are projections of the beam density) and cannot flatten the accelerating wake field in the plasma. The peak energy of the beam is around 23 GeV.

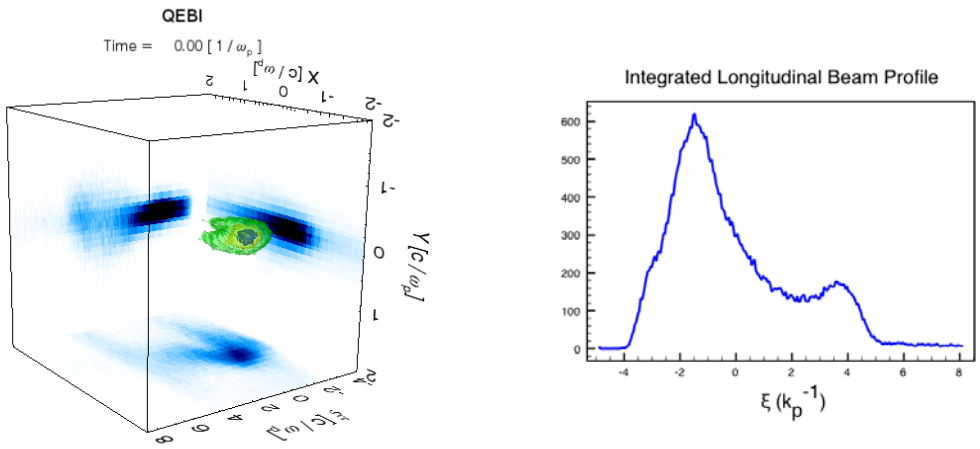

(a) (b)

Figure 8: External imported beam density plots from QuickPIC. (a) Initial 3D beam density plot; (b) Transversely integrated beam profile.

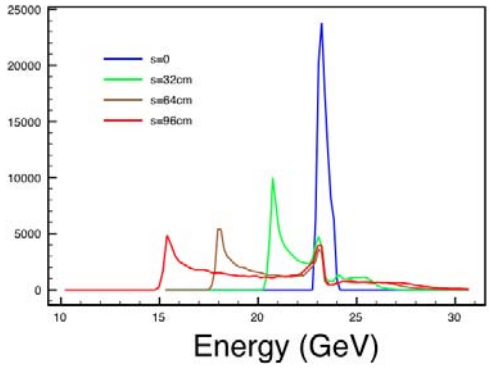

Figure 9: The beam energy spectra at different propagation distances in a field-ionized Rb plasma. The initial Rb gas density is 3.0 x 10^{16} cm^{-3}.

The beam spectra (in **Figure 9**) shows that the energy spread of the accelerated particles is large since the accelerating field is not uniform locally. And the peak values of the decelerated particles in the spectra almost have equal differences, which means the decelerating field does not change very much in the whole process.

The coherent transverse motion of the beam also appears in the simulation. From the beam phase space plots (**Figure 10**), we can find that the beam centroid moves to the right in the plots when the beam propagating inside the plasma. This phenomenon becomes significant when the beam has a big transverse emittance (e.g. 500 mm mrad) as well as a big initial tilt (e.g. 5%) of the beam centroid along the longitudinal direction. More investigation on this problem will be done in the future.

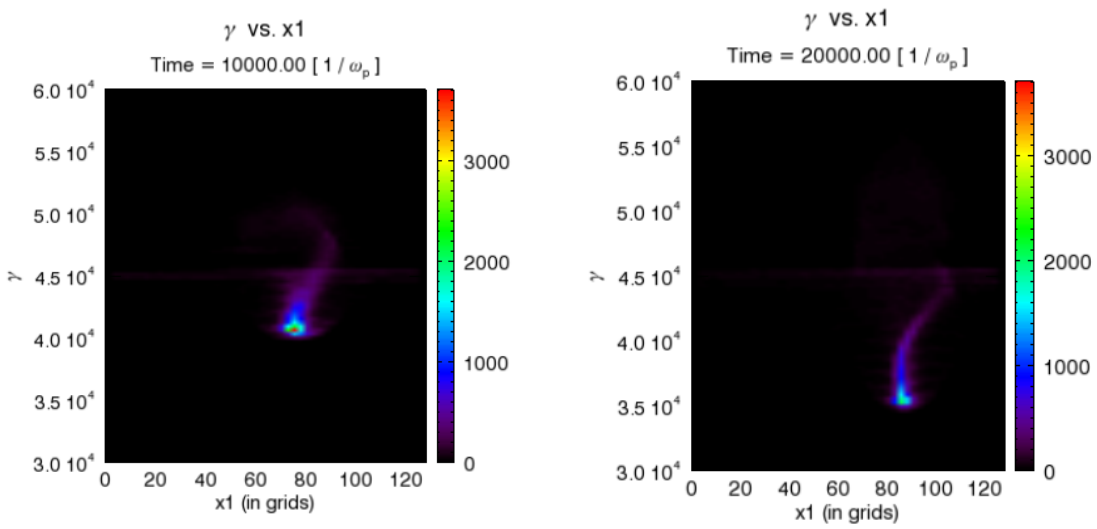

(a) s = 32 cm (b) s = 64 cm

Figure 10: Beam phase space (pz - x) plots at different propagation distances.