1.1.1. Electromagnetic Modeling of Accelerator Structures

Prepared by: David Bruhwiler, TechX Corp.
Contributors: J.R. Cary, K. Amyx, T. Austin, B. Cowan, P. Messmer, P. Mullowney, K. Paul, P. Stoltz, D. Smithe and S. Veitzer, Tech-X Corp.; G. Werner, U. of Colorado
Supports the following HEP funded NERSC repositories:
"Community Petascale Project for Accelerator Science and Simulation (m778)," Principal Investigator: P. Spentzouris
"Particle simulation of laser wakefield particle acceleration," Principal Investigator: C.G.R. Geddes
"Simulation of photonic crystal structures for laser driven particle acceleration," Principal Investigator: B. Cowan

1.1.1.1. Scientific Objectives

The overarching scientific objective for the next three to five years is to enable rapid high-fidelity simulation and design of a wide range of accelerating structures of relevance to the Office of High Energy Physics.

For example, modeling of superconducting radio frequency (SRF) accelerating cavities requires rapid and accurate calculation of frequencies for all resonant modes (fundamental and high-order), the associated Q (quality factor) for each mode, and also the spatial structure of the modes. Surface heating and multipacting are key physical processes that limit the gradient of SRF cavities and must be modeled, with emphasis on moving from analysis to designs that reduce risk and cost while improving performance for new accelerator facilities. Relevant DOE/HEP applications include the Large Hadron Collider (LHC), Project X and the International Linear Collider (ILC).

Normal conducting (warm) RF cavities and waveguides are also critical technologies for present and future facilities. There is a worldwide effort to accurately simulate RF breakdown in warm structures. This will require major advances in modeling surface physics under extreme conditions, as well as the ability to couple very small scale surface simulations to large-scale electromagnetic simulations. A particular example explores the concept of "magnetic insulation" of novel RF cavities for muon acceleration. Relevant DOE/HEP applications include muon collider and neutrino factory concepts, RF power transport, and CLIC-like alternatives to the baseline SRF design of the ILC.

High-gradient RF cavities based on dielectric structures are key to the "advanced concepts" portfolio of research and development within the DOE Office of High Energy Physics. Two examples include: a) laser-driven photonic band gap (PBG) accelerating cavities; and b) novel, larger-scale RF structures with ultra-high Q and ultra-low wakefields. Field emission and secondary emission of electrons from dielectric surfaces, and resulting surface damage, are important issues that can limit the achievable gradients and must be addressed in future simulations.

1.1.1.2. Methods of Solution

The main project code, VORPAL, is a parallel framework for finite-difference time domain (FDTD) simulations of fields and particles of various types, employing a variety of algorithms. VORPAL uses an explicit stencil to update E and B fields with 2nd-order accuracy on structured 1D, 2D and 3D Cartesian meshes. The operators are all local, which enables efficient communication via MPI and excellent scaling up to 16,000 cores.

The use of cut-cell techniques makes it possible to accurately treat complicated metal geometries, and a recently developed 2nd-order FDTD algorithm for dielectrics with arbitrary geometry makes it possible to do new large-scale simulations of complicated RF cavities made from novel dielectric structures. Using the broadly filtered diagonalization technique, frequency related information can be obtained efficiently from a time-domain code. Multi-physics capabilities, like secondary emission of electrons, are made available in VORPAL through the freely available TxPhysics library. Algorithms for coupling implicit heat advection in metal surfaces to the explicit update of Maxwell's equations in vacuum are present in VORPAL and are under active development.

1.1.1.3. HPC Requirements

Recent simulations have been modest in size, typically using 10^6 to 10^7 cells and running for fewer than $1x10^5$ time steps on approximately1,000 cores. A 50-fold increase in resources would allow modeling of 50X larger structures by increasing the size of the mesh, while holding resolution constant. This would enable VORPAL to simulate full cryomodule assemblies containing multiple SRF cavities. Also, a 50X increase in resources would enable VORPAL to directly address multi-physics problems, coupling EM and heat transport solvers over order-of-magnitude longer time scales, and also coupling surface physics models to much larger scale EM simulations in order to understand and eventually mitigate RF breakdown physics.

Developers anticipate that VORPAL simulations will scale well from ~1,000 cores at present to ~100,000 cores in the near future, as the problem size is expanded to meet future challenges over the next three to five years. In addition, developers envision a new mode of routine operation, in which ~100 VORPAL simulations each using ~1,000 cores are run in a task-parallel fashion on ~10,000 cores, under the control of a parallelized nonlinear optimization algorithm. This approach to parallel computing at NERSC will enable a much-needed shift from the present workflow — doing a few large simulations in order to obtain physical insight — towards a new and more powerful workflow of optimizing existing accelerator designs and also creating completely novel designs.

1.1.1.4. Computational and Storage Requirements Summary

	Current (2009 at NERSC)	In 3-5 Years
Computational Hours	< 1 M	100 M
Parallel Concurrency	~1 K	~100 K
Wall Hours per Run	10	10
Aggregate Memory	10 GB	1 TB
Memory per Core	0.01 GB	0.01 GB
I/O per Run	10 GB	1 TB
On-Line Storage Needed	< 1 TB	~5 TB
Data Transfer	1 TB	100 TB

1.1.1.5. Support Services and Software

Parallel file I/O using HDF5 must be scaled to hundreds of thousands of processors.

Error checking and job-relaunch services that detect if a job has terminated partway through and automatically restart it will become more important as jobs take up increasing numbers of nodes, with corresponding increases in the possibility of failure.

To allow simulations to predictively guide experiments, scans of parameter space are needed (as are conducted in experiments), which will require the above job monitoring services together with automation to generate and run sequentially large numbers of jobs, and to extract the data from them.

Visualization services for visualizing and analyzing large datasets, and in extracting physics data from them, are important. VisIt and IDL are important for visualization. Parallel visualization and analytics tools must be further developed, to provide similar functionality to well-known serial tools such as IDL/MATLAB while providing access to petascale datasets. This is a serious challenge, as even simple operations such as smoothing require communication or guard cells. The tools need to be robust and script-callable so as to be integrated into a batch workflow providing automatic analysis after batch compute jobs complete.

1.1.1.6. Emerging HPC Architectures and Programming Models

The FDTD electromagnetic algorithms of VORPAL have already been ported to the NVIDIA GPU hardware, showing more than an order of magnitude speedup over modern CPU performance, updating $>4 \times 10^8$ cells per second. The new implementation works on a simple heterogeneous system – a small Opteron cluster with one GPU per node. Effective use of multiple GPUs per Opteron core is being considered and no major technical difficulties are anticipated. Hence, the FDTD algorithm has been shown to be well situated to take advantage of these new architectures, and we expect to obtain comparable benefits for simulations that include particles.