Advances in large-scale computations of multi-particle beam dynamics

J. Amundson,1 J. Cary,2 P. Lebrun,1 A. Macridin,1 L. Michelotti,1 C. S. Park,1 P. Spentzouris,1 P. Stoltz,2 J. Qiang,3 J.-L. Vay,3 and S. Veitzer2

[1]Fermi National Accelerator Laboratory, Batavia, IL, USA
[2]Tech-X, Boulder, CO, USA
[3]Lawrence Berkeley National Laboratory, Berkeley, CA, USA

E-mail: amundson@fnal.gov

Abstract. The design and operation of modern particle accelerators is making ever-increasing demands on computer simulations of beam dynamics. We describe recent large-scale computations of collective effects utilizing high-performance computing to address the problems posed by the leading accelerators in operation today and upgrades to these accelerators and facilities planned for tomorrow. We discuss electron cloud, beam-beam and space charge simulations of the Fermilab Tevatron, Main Injector and Debuncher, as well as CERN's LHC, SPS and proposed PS2. In addition, because accelerator science serves the larger purpose of basic physics research, we also describe some of the physics topics that will become accessible due to the upgrades to accelerators we simulate.

1. Introduction

Under the SciDAC program, the Community Petascale Project for Accelerator Science and Simulation (ComPASS) [1] has developed a suite of beam dynamics simulation tools to address the modeling needs of accelerator physicists working to design tomorrow's machines as well as optimizing the machines in operation today. The beam dynamics topics that require high-performance computing are generally collective effects; single-particle effects are important to basic accelerator modeling, but do not on their own require high-performance computing to address. Collective effects, however, require the application of modern, massively parallel computing platforms. We will focus here on three of the most important collective effects in accelerator physics: electron cloud, beam-beam interactions and space charge.

In these proceedings we will describe several ComPASS-based modeling activities, focusing on current and future machines at Fermilab and CERN. We emphasize that these simulations are not done in isolation. Instead, they are each portions of larger design and optimization efforts. Information flows from the simulations to the accelerator physicists in the form of results and, in turn, flows back again in the form of experimental information and new designs. We will highlight four ComPASS codes: VORPAL, BeamBeam3d, IMPACT and Synergia. In addition, we will describe some of the particle physics topics supported by the accelerators targeted by these simulation efforts in order to provide context for the work.

2. Particle Physics at the Frontier(s)

The Particle Physics Project Prioritization Panel (P5) [2] has identified three frontiers in particle physics: the Energy Frontier, the Intensity Frontier and the Cosmic Frontier. Of these, only the Cosmic Frontier does not depend on advances in accelerator technology. The Energy Frontier is the area most commonly discussed in conjunction with particle physics. It is where the largest, most expensive machines are built and operated, and holds out the hope of the most dramatic new discoveries. The Intensity Frontier has historically received less coverage in, e.g., the popular press, but it is an equally important cornerstone of the particle physics program. The highest intensities enable, among other thing, the highest precision measurements and the possibility of discovering new physics at energy scales that will not be directly accessible to the Energy Frontier for at least another generation. In an effort to bring some of the excitement of the Intensity Frontier to a wider audience, we will emphasize it here, leaving a discussion of the Energy Frontier to the already extensive public literature.

Fermilab is proposing to dramatically extend the Intensity Frontier with Project-X [3], which includes a new superconducting 8 GeV linac as well as upgrades to existing Fermilab accelerators to handle the higher intensities demanded by the program. These upgrades require, of course, a better understanding of intensity-dependent effects, and as such are a bread-and-butter aspect of the ComPASS project's beam dynamics efforts. The physics opportunities afforded by Project-X include support for the long-baseline neutrino program, rare-decay searches and high-precision experiments. Ultimately, these experiments will help us address questions about the origins of mass, the observed family structure of particles and the matter-antimatter asymmetry of the universe.

We will highlight one proposed Intensity Frontier experiment at Fermilab, Mu2e. The Mu2e experiment will search for decays of muons into electrons. Such decays are forbidden at tree-level in the Standard Model. In fact, the only way such decays can proceed in the Standard Model is via the so-called "penguin" diagram of Figure 1. The rate for such decays is extraordinarily small. Although the figure contains a virtual photon, the photon can also be real, resulting in $\mu \to e\gamma$ decays. In this latter case, the branching ratio is given by

$$Br(\mu \to e\gamma) = \frac{3\alpha}{32\pi}\left|\sum_{i=2,3} U_{\mu i}^{*} U_{ei} \frac{\Delta m_{1i}^{2}}{m_{W}^{2}}\right|^{2} < 10^{-54}, \tag{1}$$

a rate so small that it is completely beyond the reach of any conceivable experiment. Any observation of $\mu \to e$ transitions, therefore, would be a gold-plated signal of physics beyond the Standard Model. The Mu2e experiment intends to improve the existing limit of $\sim 10^{-11}$ by more than four orders of magnitude before the Project-X era, and another two orders of magnitude during Project-X. In the end, Mu2e will require on the order of one muon for each grain of sand on earth.*

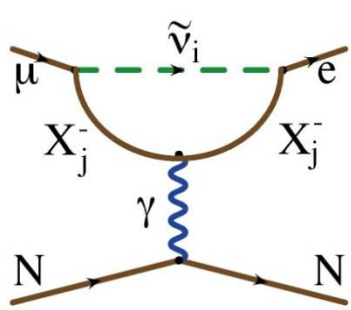

Figure 1. "Penguin" diagram for $\mu N \to eN$ decays in the Standard Model.

3. Electron cloud

An electron cloud is generated in an accelerator when stray electrons are accelerated by a passing positive beam bunch, accelerating the electrons to high enough energy that they generate secondary

* This comparison is based on the estimate of the number of grains of sand on earth from http://www.hawaii.edu/suremath/jsand.htmlhttp://www.hawaii.edu/suremath/jsand.html.

electrons when they collide with the beam pipe. Under the right conditions, a cascade effect causes the bunch to grow exponentially, eventually creating a cloud of electrons with a large enough electric field to cause beam instabilities. Electron cloud is the conceptually most complex collective effect of those we address in this work. While space charge and beam-beam interactions are calculable from first principles, electron cloud models ultimately rely on phenomenological models for the production of secondary emission of electrons. Since the details of the secondary emission yield (SEY) depend on the details of the beam pipe surface, there are unavoidably large uncertainties in the phenomenological models. We describe here two electron cloud simulation efforts. The first, in the Fermilab Main Injector, addresses the buildup of the cloud itself. It also includes a simulation of a microwave electron cloud detection experiment. The second simulation, at the CERN SPS, focuses on the destabilizing effect of a pre-generated electron cloud on the beam. It includes a simulation of a proposed compensation mechanism.

3.1. Electron Cloud Simulations in the Fermilab Main Injector

The Fermilab Main Injector is currently able to handle its designed beam intensity with losses within acceptable limits. In the Project-X era, however, the Main Injector will be required to handle approximately a factor of three increase in beam power, which could increase instabilities and lead to unacceptable losses. In the conjunction with experimental studies currently begin pursued in the Main Injector, we have performed a series of studies of electron cloud formation [4] utilizing the fully 3D and self-consistent electromagnetic PIC code VORPAL [5].

The physical model consists of an elliptical stainless steel beam pipe (of main and major axes 2.34 cm and 5.88 cm, respectively) located in a static magnetic field. Two variations were studied in detail: a short section of 0.25 m and a long section of 16 m, taken from a representative MI arc consisting of a dipole followed by a quadrupole, another dipole and, finally, a field free region. The grid used was 48×48 transversely, with 384 and 6144 longitudinal divisions for the long and short sections, respectively. In both cases, a perfectly matching layer was included at both ends of the gridded area to avoid unphysical reflections from the artificial longitudinal boundaries. The simulation included a small seed of electrons uniformly distributed along the beam direction. The details of the seed distribution proved to be irrelevant, as expected. During the simulation, there is a phase where the number of electrons grows exponentially. It was necessary to modify the ratio of macro particles to real particles during the simulation by periodically randomly culling the electron distribution. Depending of the growth speed of the cloud, roughly from 3 to 10 culling phases were needed before reaching stable saturation.

The effects on the electrons of the passage of a beam bunch can be seen in Figure 2. On the left hand side of the graph, no beam is present and the average electron kinetic energy decreases as energy is lost in collisions with the wall. When the beam enters in the center of the graph, the kinetic energy quickly increases as the electrons feel the pull of the positively-charged beam. The bunch quickly passes and the average kinetic energy again begins to decline.

Figures 3 and 4 show representative slices of the electron cloud density once saturation has been reached. Figure 3 shows transverse distributions during the relaxation (no beam present) and pinch (beam present) phases. Figure 4 shows the longitudinal distribution in the long pipe simulation. The simulation is long enough to include multiple bunches. The beam density is given by the red line.

The beam-position monitor (BPM) has been simulated via the use of VORPAL's pseudo-voltages. In a dipole, a displaced beam in an electric cloud of sufficient density ought to result in an echo of the beam pulse detectable by the BPM. An FFT analysis of the simulated BPM response shows a weak resonance at the cutoff frequency of the beam pipe. The excited electron cloud carries electromagnetic waves that can propagate in the beam pipe.

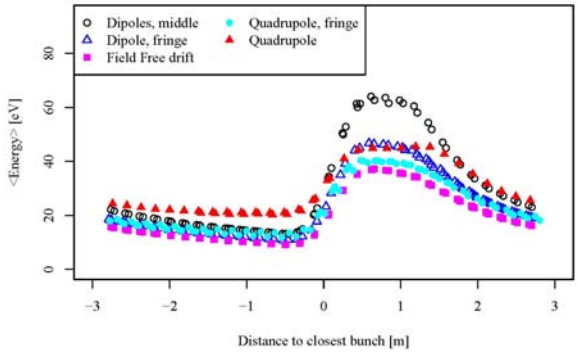

Figure 2. Variation of the kinetic energy distribution as function of the distance to the nearest bunch.

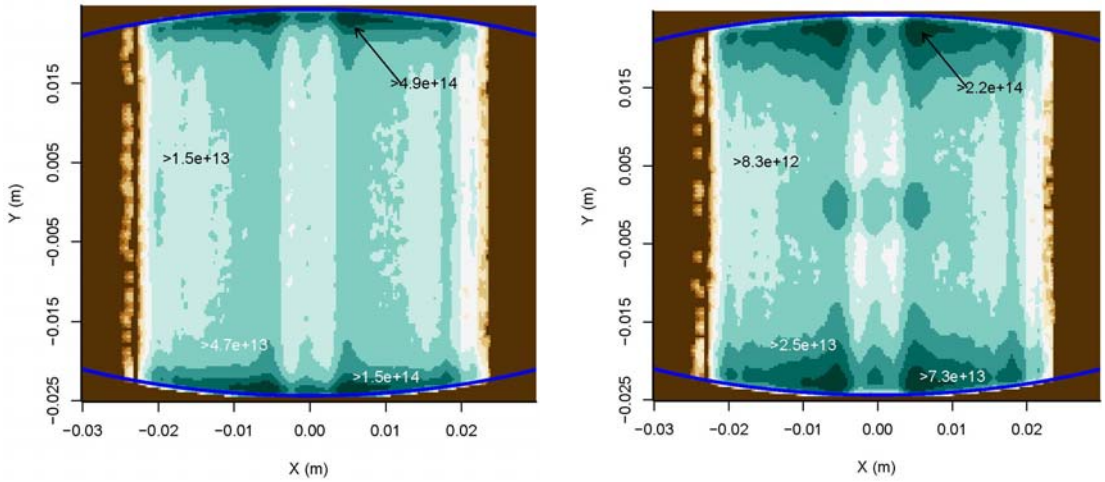

(a) (b)

Figure 3. Transverse density of the generated electron cloud in the Main Injector during the (a) relaxation and (b) pinch phases.

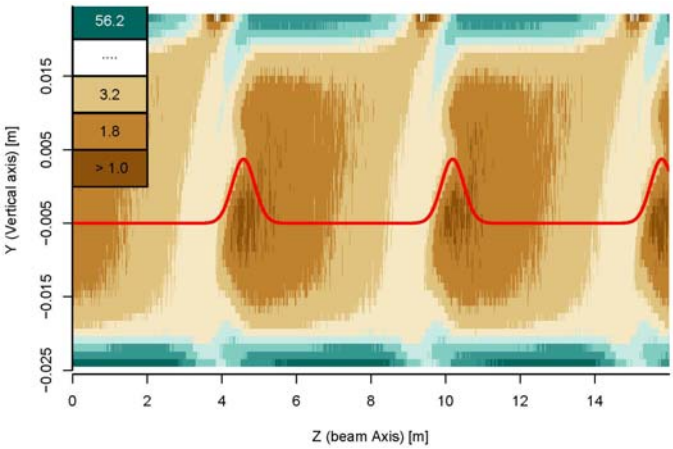

Figure 4. Y-Z projection of EC. The red line depicts the proton beam bunch profile. Here, a continuous dipole field of 0.234T is assumed.

The time scale of the electron cloud generation depends on the SEY parameters, the bunch intensity and, to a lesser extent, the magnetic field configuration. Under current conditions, in the MI dipoles, assuming a relatively high value for the maximum SEY (SEY_m) of about 2.0, the e-folding time for the generation of the cloud is about 50 ns. If SEY_m is less than 1.05, the EC dies away and saturation does not occur. Because of the uncertainties in the SEY model and the surface of the beam pipe itself, simulation alone can not *a priori* predict the EC density and its impact on the beam. However simulations provide valuable information used to guide and interpret the ongoing experimental program.

3.2. Electron Cloud in the CERN SPS

Electron clouds have been shown to trigger fast-growing instabilities on proton beams circulating in the CERN SPS [6]. A feedback system to control the instabilities is under active development [7]. To study this system, we have done a series of studies [9] of beam-electron cloud interactions and the feedback system in the SPS using WARP/Posinst [8]. Whereas the simulations in the previous section focused on the development of the electron cloud itself, this work focuses on beam-electron cloud interactions in a situation where a substantial electron cloud has already been shown to exist.

The model consists of a succession of N_s discrete interactions around the ring ("ecloud stations"). Two consecutive bunches of 1.0×10^{11} protons, representing bunches 35 and 36 of the bunch train, were modeled with an RMS bunch length of 0.23 m and an inter-bunch spacing of 25 ns. The initial electron distribution was extracted from a prior full-build-up simulation using Posinst alone to describe bunches 1 through 34, with parameters chosen to match the experimentally measured tune shift. Figure 5 shows the electron density around the two bunches projected onto the vertical and horizontal planes.

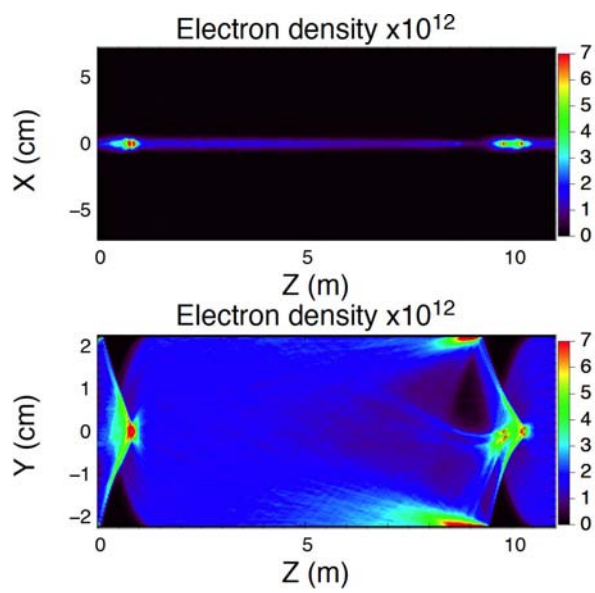

Figure 5. Electron density around two bunches in the SPS. The horizontal plane is shown on the top, the vertical on the bottom.

The feedback model takes the measured beam position at point y_{i-1} and uses the smooth focusing approximation to the full lattice to predict the beam position at time y_i ξ turns later than y_{i-1}. A gain factor g is added and a velocity kick proportional to the predicted beam position is applied. For these simulations, $\xi = 1$.

Figure 6 shows that the feedback system is successful in suppressing emittance growth. Initial comparisons with experiment show good qualitative and some quantitative agreement on key aspects of the observed instability [7]. More realistic simulations of the feedback system are planned for the future.

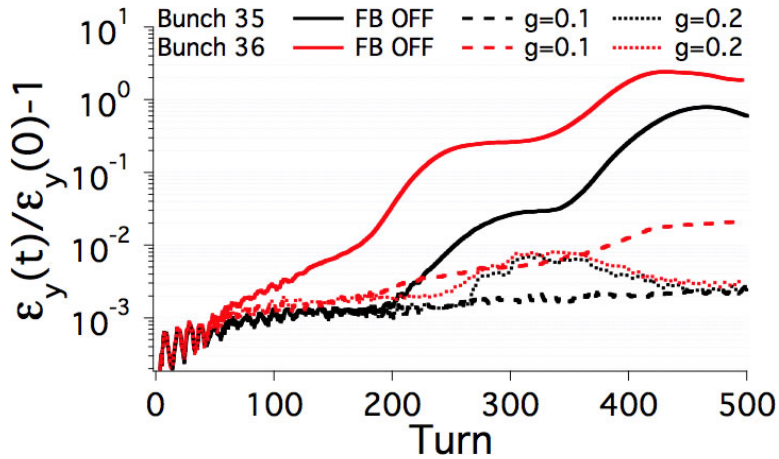

Figure 6. Emittance growth with and without compensation. g is the gain factor in the feedback system; "FB OFF" corresponds to no feedback. Note the logarithmic vertical scale.

4. Beam-beam

Parasitic interactions between colliding beams can lead to emittance growth and instabilities. They are another collective effect worthy of simulation. Beam-beam interactions may be modeled from essentially first principles, especially in comparison with electron clouds. In a typical collider beams are brought into collision, or at least near proximity, several times over the course of a turn. The basic problem is to calculate the field due to both beams during their period of close proximity and apply the effects of fields to both beams. The problem is somewhat complicated by the fact that trains of bunches are used, resulting in a combinatorial sequence of bunch i colliding first with j, then with k, etc.

The ComPASS code BeamBeam3d [10] has been developed to enable realistic simulations of beam-beam interactions. It includes an integrated, shifted Green function method for efficient field solves with spatially separated beams, a multiple-slice model for finite bunch length effects and a parallel particle-based decomposition model for ideal load balance. It can handle arbitrary closed-orbit separation, multiple bunches and multiple collision points. The optics calculations (for propagation between bunch-bunch interactions) include arbitrary transfer matrices, chromaticity effects, and thin lens multipole kicks. During the course of the ComPASS project it has been enhanced to include resistive wall wakefields, a conducting wire model, crab cavities and electron lens compensation. Figure 7 shows the model for longitudinal slices in the simulation.

BeamBeam3d has been used to model the machines at the Energy Frontier, the Tevatron and the LHC. In the case of Tevatron modeling, the combination of beam-beam interactions and chromaticity were studied [11]. The results of one of these simulations, combined with a data showing the same bunch-to-bunch behavior, are shown in Figure 8. The insight gained from these simulations was used to maximize luminosity in Fermilab's Run II. This result was described in a SciDAC Breakthrough presentation at this conference (see Reference [12]). BeamBeam3d has also been used to study beam-beam interactions at the LHC. In order consider the effects of machine optics on the beam-beam interaction, BeamBeam3d's built-in ability to do parallel parameter scans was used to survey 32 parameter combinations in parallel, effectively utilizing the availability of massively parallel computing to complete in 2.5 hours what would have taken 78 hours if jobs had to be launched individually. The results of the simulations have been fed back to the LARP team.

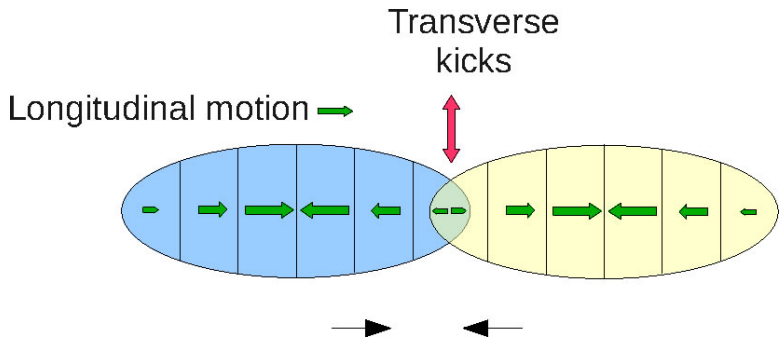

Figure 7. Longitudinal slices in the BeamBeam3d interaction calculation.

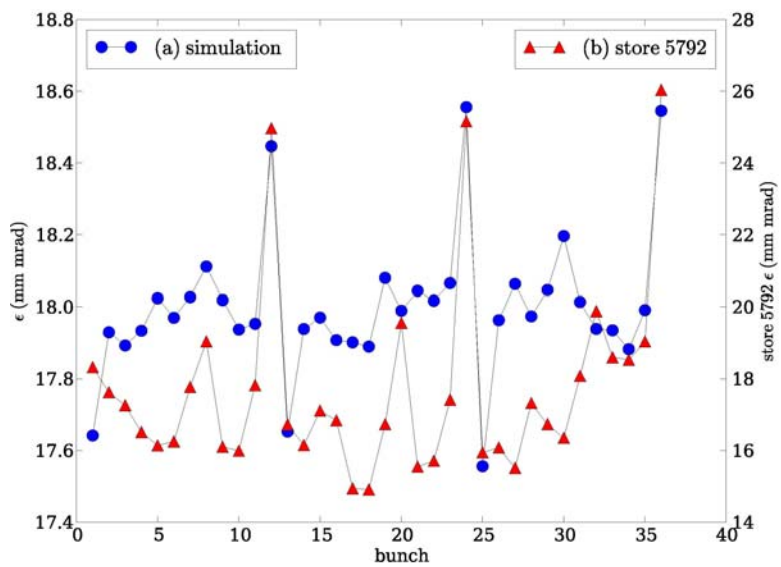

Figure 8. Vertical emittance data from the Tevatron overlaid with simulation showing similar patterns in bunch-to-bunch variation.

5. Space charge

Space charge is the canonical collective effect. It is simply the interaction of the beam with itself, and can be calculated (numerically) from very close to first principles. Space charge effects are most important for low-energy, high intensity machines. One of the ComPASS codes for studying space charge is IMPACT [13]. IMPACT contains multiple 3D Poisson solvers for a variety of boundary conditions, a detailed RF accelerating and focusing model, wakes, thin-lens nonlinear elements and more. IMPACT has recently been applied to studies for CERN's proposed PS2, the upgrade to the current CERN proton source. The results for emittance growth and beam characteristics have been used as input to the LARP design effort.

5.1. Simulations of Space Charge and Nonlinear Optics in the Fermilab Debuncher

The proposed Mu2e experiment, described above, will require running the existing Fermilab Debuncher and Accumulator rings with approximately 10^5 times the intensity seen in current operating conditions. Furthermore, the experiment plan includes resonant extraction, an inherently nonlinear process for extracting the beam from the Debuncher to the target which will create the muons for the experiment.

The combination of high-intensity beams and nonlinear optics make the Debuncher for Mu2e a prime candidate for simulation.

The ComPASS code Synergia [14] is ideal for the task. Synergia includes, among its features, fully 3D Poisson solvers for various boundary conditions and single-particle optics from CHEF [15], a highly-advanced single-particle code capable of handling all the necessary nonlinear effects at practically arbitrary levels of detail. A full simulation of the Mu2e resonant extraction process will include not only nonlinear optics and space charge, but also dynamical ramping of the lattice elements. Synergia's implementation as a fully programmable, Python-based simulation framework is perfectly suited to this problem.

Before looking at the results from Synergia simulations of Mu2e, it is necessary to introduce some basic concepts from accelerator physics, especially particle *tune*. In the presence of linear beamline elements, the motion of a particle in an accelerator is described by Hill's Equation,

$$\frac{dx^2}{d^2s} + K(s)x = 0, \tag{2}$$

where s is the distance traveled along the closed orbit and $K(s)$ is the varying quadrupole magnet strength. The solution to this equation is

$$x = \sqrt{\varepsilon\beta(s)}\cos(\phi(s) + \phi_0), \tag{3}$$

where ε and ϕ_0 are constants of the motion, $\beta(s)$ is the amplitude variation due to the varying magnet strength and $\phi(s)$ is the phase advance, which also depends on the magnet strength. The *betatron tune*, ν is defined in terms of the total phase advance in a single turn, $\Delta\phi$,

$$\nu = \frac{\Delta\phi}{2\pi} = \int_0^{s_{turn}} \frac{ds}{\beta(s)}. \tag{4}$$

Parameters used for the Synergia simulations are tabulated in Table 1. The planned cycle for the Debuncher in the Mu2e experiment is to start with a mostly linear lattice, ramp up the (nonlinear) sextupoles to create a resonance condition, then move the beam on resonance, where the it will naturally be kicked out into extraction elements. In order to tackle this problem in detail, we first consider the nonlinear problem without space charge. Figure 9 shows the horizontal phase space distribution of the initial beam in the linear lattice and the same distribution after the sextupoles have been ramped to full power. The effect is quite dramatic—it clearly shows the threefold symmetry of the sextupole field.

Table 1. Beam parameters used for Synergia simulations, taken from the original Mu2e proposal.

Beam energy	8 GeV
Intensity	1.2×10^{13} protons/bunch
Transverse emittance	20π mm-mr
Longitudinal width	40 nsec RMS
Space charge tune shift (Laslett)	0.08

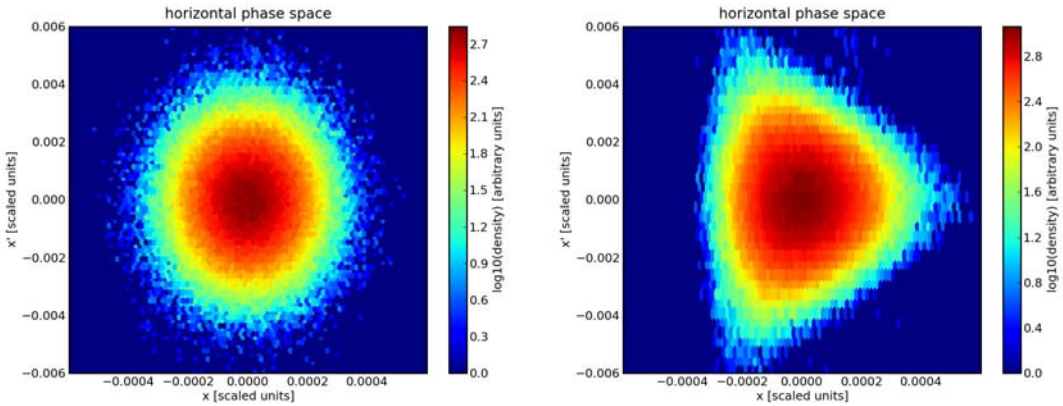

(a) (b)

Figure 9. Horizontal phase space in the Debuncher (a) before the sextupole ramp and (b) at the completion of the sextupole ramp.

The next step is to consider the effects of space charge in the purely linear lattice. In the absence of space charge, all of the particles have the same tune, known as *bare tune*. Figure 10 shows the distribution of tunes in the beam when the space charge effect is included. The bare tune is the yellow point in the upper right. The characteristic shape of the tune footprint is known as the space-charge "necktie." Figure 10 also shows the effect of combining the nonlinear lattice with the space charge effect. The effect is striking. The system displays resonant behavior, as is evidenced by the structures that appear in the tune footprint. The system also generates enough beam loss as to be unacceptable for actual operations. The uncontrolled losses before the beam can be extracted would lead to too much radiation in the beam tunnel.

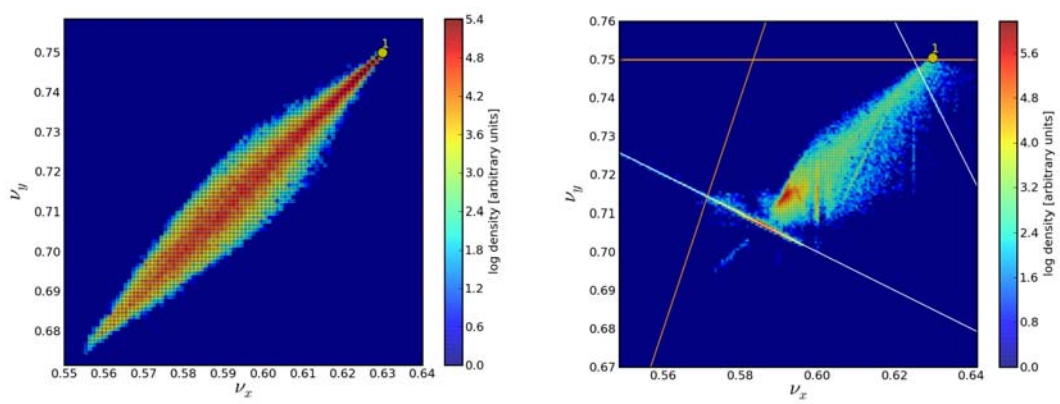

(a) (b)

Figure 10. Tune footprint including space charge and (a) linear lattice elements only, (b) linear lattice plus full-strength sextupoles. The yellow point is the bare tune.

These results can be understood in terms of resonance theory. Resonances occur when

$$m + l\nu_x + q\nu_y = \nu_x \tag{5}$$

for integer m,l,q. The smaller the integers, the lower the order of the resonance. Equation (5) implies that resonance occur along straight lines in tune space, $(\nu_x,\ \nu_y)$. Figure 11 shows the same tune footprint, but with the 3rd-and 4th-order resonances highlighted. As one might expect, the most

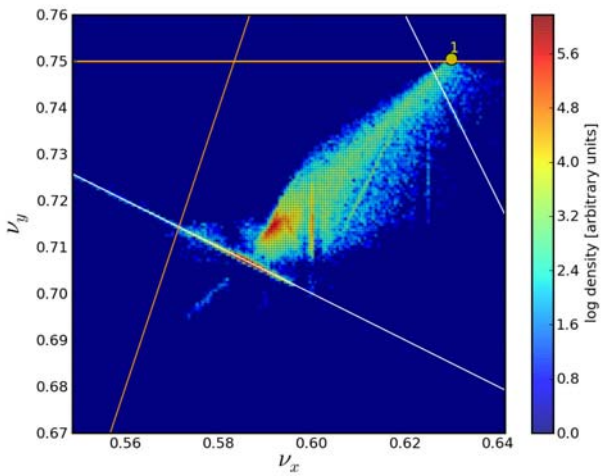

Figure 11. Tune footprint showing 3rd- and 4th-order resonance lines (white and orange, respectively). Higher order resonance structures are also present.

prominent resonance is the 3rd-order resonance, i.e., the one which is directly caused by the sextupoles. Other structures in the plot indicate the presence of higher-order resonances.

Because Synergia is capable of tracking individual particles, we can extract the trajectories of each individual particle that is lost and calculate its tune. Figure 12 dramatically demonstrates the correlation between resonance behavior and particle loss: nearly every lost particle falls neatly on a resonance line. Further structure is evident—particle loss seem to congregate around certain intersections of resonance lines. This phenomena was predicted over twenty years ago under the name "resonance streaming" [16]. It makes a stunning appearance in this analysis.

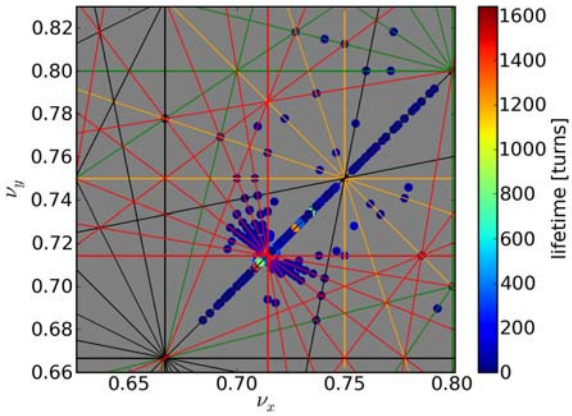

Figure 12. Lost particles in tune space and their relation to resonance lines up to 8th order: white 3rd-order, orange 4th-order, green 5th-order, black 6th-order, red 7th-order, and 8th-order. The color of the points themselves represents the lifetime of each particle before it is lost.

Because the structure of particle loss is tightly correlated with resonant behavior, it was hoped that an operational point in phase space could be found that would reduce losses to an acceptable level. After an extensive scan, a few points with somewhat better behavior were identified (along with many exhibiting much worse behavior). In the end, Synergia simulations indicate that it would be extremely difficult to run the Debuncher at this intensity with acceptable losses. This information has been fed back to the experiment, which is now considering lower-intensity operating scenarios.

6. Moving toward the Exascale

The entire HPC community is beginning to prepare for the next frontier in computing, the Exascale. ComPASS as a whole, and the beam dynamics portion of ComPASS in particular, is no exception. The first serious observation in considering the leap to the Exascale is that most, if not all, of our current parallel models are closely tied to the computing architectures available today. Indications are that

exascale computers will be very different architecturally from petascale computers. The next step is to start moving away from old assumptions by exploring the sorts of new architectures that are appearing today. Work on abstracting and parametrizing architectural assumptions in the context of generic PIC simulations have already been undertaken and shown to extract good performance from GPUs [17]. At the same time, we are exploring new decomposition strategies utilizing a full decomposition of phase space instead of just physical space with dynamical grids that track the particles, minimizing inter-processor particle movement.

A further observation about scaling beam dynamics simulations is that strong scaling is unlikely to be possible over several orders of magnitude. The types of calculations in most beam dynamics simulations simply are not large enough to scale to millions, much less billions, of processors, which also means that there are limits to straightforward weak scaling of our simulations. We are currently exploring new ways to exploit weak scaling, such as intelligent parallel optimization and new algorithms for taking longer time steps using more complex intermediate substeps.

At least one straightforward weak scaling topic is of potential interest. We currently perform macroparticle simulations of beam bunches containing $10^{10} - 10^{13}$ particles. While macroparticle calculations are usually sufficient, processes sensitive to real statistical fluctuations are difficult to model without the physical number of particles. We can easily anticipate utilizing exascale computers to simulate bunches of 10^{12} real particles.

Our next generation scaling efforts are just getting underway. We are forming collaborations between application people, computer science and applied math people and hardware vendors in order to come up with effective plans to exploit high-performance computing resources we expect to find in the upcoming era.

7. Conclusions

ComPASS beam dynamics simulations are actively involved in advancing the Energy and Intensity Frontiers in high energy physics. We have discussed how four of our codes, VORPAL, BeamBeam3d, IMPACT and Synergia have been used to perform large-scale simulations of electron cloud, beam-beam and space charge effects at CERN and Fermilab. These simulations have both benefited from and contributed to experimental and design programs in accelerator physics. We plan to extend our simulation capabilities to the next generation of high-performance computing in order to meet the needs of tomorrow's accelerators.

Acknowledgments

This work was supported by the United States Department of Energy under contract DE-AC02-07CH11359 and the ComPASS project funded through the Scientific Discovery through Advanced Computing program in the DOE Office of High Energy Physics. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This research also used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357.

References

[1] P. Spentzouris et al. https://compass.fnal.gov/.

[2] "US Particle Physics: Scientific Opportunities A Strategic Plan for the Next Ten Years," *Report of the Particle Physics Project Prioritization Panel (P5)*, http://www.er.doe.gov/hep/files/pdfs/P5_Report.

[3] Stephen D. Holmes, "Project X at fermilab," PAC'09, Vancouver, May 2009, ID-1167, http://www.JACoW.org.

[4] P. Lebrun, P. Spentzouris, J. Cary, P. Stoltz, and S. Veitzer, "Accurate Simulation of the Electron Cloud in the Fermilab Main Injector with VORPAL," IPAC'10, Japan, May 2010.

[5] C. Nieter and J. R. Cary, “VORPAL: a versatile plasma simulation code,” *J. Comp. Phys.* **196**, 448–472 (2004).
[6] G. Arduini et al., CERN-2005-001, 31–47 (2005).
[7] J. D. Fox et al., SPS Ecloud Instabilities—Analysis of machine studies and implications for Ecloud Feedback, IPAC10.
[8] J-L Vay et al., “Update on e-cloud simulations using the package Warp/Posinst,” PAC09, Proceedings, (2009).
[9] J-L Vay, et al., “Simulation of E-Cloud Driven Instability and its Attenuation Using a Feedback System in the CERN SPS,” IPAC10.
[10] J. Qiang, M.A. Furman, and R.D. Ryne, *J. Comp.Phys.*, **198** (2004), pp. 278–294; J. Qiang, M.A. Furman, and R.D. Ryne, *Phys. Rev. ST Accel. Beams*, 104402 (2002).
[11] E.G. Stern, A.A. Valishev, J.F. Amundson, and P. Spentzouris, “Fully 3D Multiple Beam Dynamics Processes Simulations for the Tevatron,” *Phys. Rev. ST AB*, 2009.
[12] E. Stern, et al., “Fully 3D multiple beam dynamics processes simulation for the Fermilab Tevatron,” these proceedings.
[13] J. Qiang, R. Ryne, S. Habib, and V. Decyk, *J. Comp. Phys.* vol. 163, 434, (2000).
[14] J. Amundson, P. Spentzouris, J. Qiang, and R. Ryne, 2006, *J. Comp. Phys.* **211** 229–248.
[15] L. Michelotti, “C++ objects for beam physics,” FERMILAB-CONF-91-159, *Presented at 14th IEEE Particle Accelerator Conf.*, San Francisco,CA, May 6–9, 1991.
[16] J. Tennyson, *Physica 5D* (1982) 123.
[17] V. Decyk, T. Singh, and S. Friedman, “Graphical Processing Unit-Based Simulations,” *Proc. Intl. Computational Accelerator Physics Conf.* (ICAP2009), San Francisco, CA, 2009.