

High-Performance Computing Modeling Advances Accelerator Science for High-Energy Physics

James Amundson, Alexandru Macridin, and Panagiotis Spentzouris | Fermilab National Accelerator Laboratory

Particle accelerators are essential for advancing our understanding of matter, energy, space, and time. Because they exhibit many physical effects on multiple scales, advanced computational tools are essential for accurately modeling them. The authors focus here on Synergia, an accelerator simulation package capable of handling the entire spectrum of beam dynamics simulations.

The US Department of Energy (DOE)'s Office of High-Energy Physics (HEP) promotes a broad, long-term particle physics program by supporting current operations, experiments, and research as well as development for future operations at three interrelated frontiers of particle physics¹:

- The energy frontier directly explores the universe's fundamental constituents and architecture through the highest-energy particle beams.
- The intensity frontier enables a second, unique, investigation of fundamental interactions through a combination of intense particle beams and highly sensitive detectors.
- The cosmic frontier reveals the nature of dark matter and dark energy by using particles from space to explore new phenomena.

These scientific frontiers form an interlocking framework that addresses fundamental questions about the laws of nature and the cosmos. The development and deployment of high-performance computing (HPC) accelerator modeling capabilities is essential to meeting these grand scientific HEP challenges because it enables and catalyzes advancement in accelerator science. In the next decade, for example, the HEP community will explore the intensity frontier by designing high-intensity proton sources for neutrino physics and rare process searches, such as Proton Improvement Plans (PIP) I and II at Fermilab.

The design, cost optimization, and successful operation of modern accelerators require the optimization of many parameters, as well as the understanding and control of many physics processes. This can only be accomplished by employing high-fidelity computational accelerator models that efficiently utilize HPC resources. A comprehensive picture of current HPC accelerator modeling capabilities in the US can be obtained by reviewing the codes developed under the Scientific Discovery through Advanced Computing (SciDAC) program within the Community Petascale Project for Accelerator Science and Simulation (ComPASS).² These codes obtain good scalability and parallel efficiency on thousands to hundreds of thousands of processors and are routinely used to perform single-physics, single-scale simulations on a few thousand processors. Their applications have enabled large multiscale multiphysics simulations of the most challenging accelerator science projects and demonstrated the impact of large-scale simulations in accelerator science.³

Synergia is an accelerator simulation package that can handle the entire spectrum of beam dynamics simulations. Here, we present the Synergia beam dynamics framework and discuss its current applications in the design of high-intensity proton accelerators.

HEP's Intensity Frontier

The intensity frontier addresses central questions in particle physics that aren't directly accessible with current or planned accelerators at the energy

frontier. Experiments at the intensity frontier study are processes that indirectly probe higher-mass scales and exotic physics using intense beams of particles such as neutrinos, muons, kaons, and nuclei, providing powerful probes of new phenomena.

In the US, the key to long-term leadership at the intensity frontier is the PIP at Fermilab (stages I and II),^{4,5} which is building a multimegawatt proton accelerator that will produce intense beams of neutrinos and muons (and possibly kaons and heavy nuclei with further upgrades). The culmination of these PIP efforts will be the delivery of a new superconducting radio frequency (RF) Linac and major improvements to Fermilab's Booster and Main Injector accelerators. High-fidelity simulations are an important component for the campaign's success and cost-effectiveness.

Every proton in the domestic HEP experimental program will be accelerated by the existing (and now 40-year-old) Fermilab Linac and Booster accelerators until new machines become operational to replace them. The leading replacement candidate—the proposed superconducting linear accelerator—is anticipated for completion no sooner than 2020 to serve demands for beams at 3 gigaelectronvolts (GeV) and lower energy and no sooner than well into the next decade to serve demand for beams at higher energy. The domestic HEP program for the next 15 years therefore depends on the viability and vitality of the Linac and Booster, which Fermilab has established a charge to assure, promising “2.25 × 10¹⁷ protons/hour (at 15 Hz) by January 1, 2016” (more than two times the beam rate in current operations) while “ensuring a useful operating life of the proton source through 2025.”^{4,6}

Booster intensities and repetition rate are currently limited by radiation due to uncontrolled losses. These losses are a problem both because of prompt radiation levels and equipment activation. To ensure that we'll achieve the PIP's target intensities, it's necessary to understand and suppress beam instabilities as well as to develop and study techniques to control and minimize beam losses (examples include optimized operating parameters, an improved collimation system, a second harmonic radio frequency system, and so on). High-fidelity simulations that incorporate all the relevant physics effects will help guide accelerator scientists as they undertake these tasks.

Synergia

Realistic accelerator simulations require treatment of many devices and physical effects, and cover a

broad range of size and detail. Modern particle accelerators are complex devices, typically consisting of thousands of components. The settings of these components are often varied during a run either through preprogrammed operating cycles or in response to active feedback. The beams in these accelerators are usually bunched longitudinally with $O(10^{12})$ particles in a bunch (the longitudinal direction is defined along the direction of the beam propagation). These bunches have transverse dimensions on the order of millimeters and a typical longitudinal extent of meters.

Synergia (<https://web.fnal.gov/sites/Synergia/SitePages/Synergia%20Home.aspx>) is an accelerator simulation package designed to take advantage of computational resources ranging from desktop machines to computing facilities. It utilizes particle-in-cell (PIC) methods to combine advanced independent-particle dynamics with state-of-the-art collective effects. The current version (2.1), which we describe here, is a hybrid Python/C++ implementation: all core computations are done in C++, while end-user simulations are described using Python. The combination provides efficiency and great flexibility.

Accelerator Simulation Problem Definition

The split-operator technique is the core mathematical concept for combining independent-particle and collective effects in Synergia. When the Hamiltonian for a system can be split into independent (i) and collective (c) components,

$$H = H_i + H_c. \quad (1)$$

The split-operator approximation for the time-step evolution operator, M_{full} is given by

$$M_{full}(\tau) = M_i\left(\frac{\tau}{2}\right)M_c(\tau)M_i\left(\frac{\tau}{2}\right) + O(\tau^3), \quad (2)$$

where M_c and M_i are the evolution operators corresponding to H_c and H_i , respectively. Synergia abstracts this technique by defining each simulation as a series of steps, each of which is defined by an ordered set of operators; a set of steps through an entire accelerator is called a *turn*. Typical simulations of circular accelerators may consist of thousands to 100,000 turns—real accelerator cycles can mean millions of turns. Accelerator operations often require changing the accelerator parameters such as altering magnet settings and RF phase shifts, possibly in response to beam conditions. Synergia can also handle linear accelerators, through which the

beam only passes once. Then, the (poorly named, in this context) number of turns is simply one.

A Synergia simulation consists of propagating a single *bunch* or train of bunches through a given number of turns. Along the way, various user-selected (or user-defined) diagnostics can be performed on the bunches to monitor beam state as it propagates through the machine. Operators act on a bunch of macroparticles representing the beam bunch, propagating them forward in time. The simulation parameters are defined by a brief Python or C++ program written by the end user. This program may use existing Synergia classes or include custom extensions to those classes. The entire state of the simulation, including end-user extensions, can be checkpointed or resumed at any point. This checkpointing mechanism allows both for recovery from hardware failures and the chaining of multiple jobs in time-limited queues to complete one long simulation.

Synergia consists of a core set of C++ classes that are exposed to Python via Boost.Python (www.boost.org/doc/libs/1_56_0/libs/python/doc). The independent-particle dynamics are handled by the CHEF C++ libraries.⁷ Both Synergia and CHEF were developed at Fermilab.

Independent-Particle Physics

Most modern particle accelerators are dominated by independent-particle physics, which means that the forces felt by beam particles due to externally applied fields (magnets, accelerating structures, and so on) are much larger than the forces due to interactions with other particles in the beam. The CHEF libraries, which we use in Synergia, provide a comprehensive set of tools for linear and nonlinear independent-particle accelerator physics.

CHEF not only models particle propagation through accelerators, but it also performs symbolic algebra calculations on the transfer maps for the propagation, allowing the user to perform analyses such as nonlinear map analysis of resonant structures. In Synergia, we take advantage of this ability by incorporating two different types of particle propagation: direct numerical integration and polynomial map application. In the latter, the six-dimensional final phase space coordinates of a particle u_f are related to the initial coordinates u_i by

$$u_f = M(u_i), \quad (3)$$

where M is a polynomial in the phase space coordinates u . This approach is useful both in

terms of efficiency for some calculations and for comparison with other accelerator simulation packages, many of which are limited to fixed-order polynomial maps for particle propagation.

Collective Effects

Synergia includes a general treatment of collective accelerator physics effects via the split-operator approximation described earlier. Synergia 2.1 includes implementations of the two most important effects for high-intensity proton accelerators—space charge and wake fields. Space charge effects arise from the electromagnetic repulsion between like-charged particles in the beam, and wake field effects occur when leading portions of the beam induce currents in the beam pipe or other structures that give rise to residual electromagnetic fields felt by the passing particles. Synergia's design lets the user implement additional collective effects and/or different approximations or implementations of the existing space charge and wake field models.

Space charge. Calculating the effects of space charge requires solving the Poisson equation,

$$\nabla^2 \phi = -\frac{\rho}{\epsilon_0}, \quad (4)$$

where ϕ is the scalar electric potential and ρ is the charge density due to the beam particles. In a bunched beam, the longitudinal separation between bunches (typically on the order of meters) is usually much greater than the transverse dimensions of the beam (typically on the order of millimeters), so bunch-to-bunch space charge effects can be neglected.

Synergia includes space charge solvers at different levels of approximation. The simplest solver uses an analytic approximation for the field due to a two-dimensional Gaussian charge distribution with open boundary conditions. All the other solvers obtain numerical solutions to the Poisson equation on a discrete grid. The solvers include two- and three-dimensional approaches and multiple boundary conditions. Figure 1 shows a space charge solve with macroparticles and the resulting scalar field.

Wake fields. The effects of induced wake fields in beam pipes with horizontal and vertical mirror symmetry can be summarized using the expressions⁸

$$\beta c \Delta p_z = -q Q W^{\parallel}(z) \quad (5)$$

$$\beta c \Delta p_x = -qQ(W_X(z)X + W_x(z)x) \quad (6)$$

$$\beta c \Delta p_y = -qQ(W_Y(z)Y + W_y(z)y), \quad (7)$$

where Q , X , Y (q , x , y) represent the charge and horizontal and vertical displacements of the leading (trailing) particle, and $W_{x,y,\parallel}$ represents the wake functions that need to be calculated for a given geometry. The momentum of a trailing particle going through a lattice element is kicked, that is, instantaneously changed, by a term proportional to a wake function that depends only on the distance z between the leading and the trailing particle. Wake field effects include both intra- and inter-bunch interactions.

Parallel Performance

The Synergia design includes parallelism at its core. Simple Synergia simulations can be run on a desktop machine with a single core, but users can easily utilize a range of parallel resources ranging from multicore desktops to 100,000+ core supercomputers. Figure 2 demonstrates code scalability. We find excellent strong scaling behavior over a range of a little less than a thousand on a Blue Gene/Q machine and weak scaling in the number of macroparticles (not shown) and bunches (Figure 2b).

Scalability in Synergia is limited by the communication requirements of collective effects. The scaling we've achieved to date relies on two complementary techniques: communication avoidance and hybrid message passing interface (MPI)/OpenMP optimizations. In communication avoidance, we perform multiple, redundant field solves to reduce the size of the necessary communication steps. The hybrid MPI/OpenMP optimizations involve reducing the number of MPI processes while using OpenMP threads to take advantage of the multiple cores available for each process. Here, communication is reduced because of the smaller number of MPI processes. The use of OpenMP is optional; once it's enabled, the number of threads per process can be chosen at runtime.

In preparation for the next generation of HPC technologies, we're actively working on adapting Synergia for GPUs and Intel's MIC architecture. The prototype GPU implementation is able to take advantage of multiple GPUs using communication avoidance—a single field solve fits naturally on a current GPU. The most difficult portion of the collective calculations in both the OpenMP and GPU

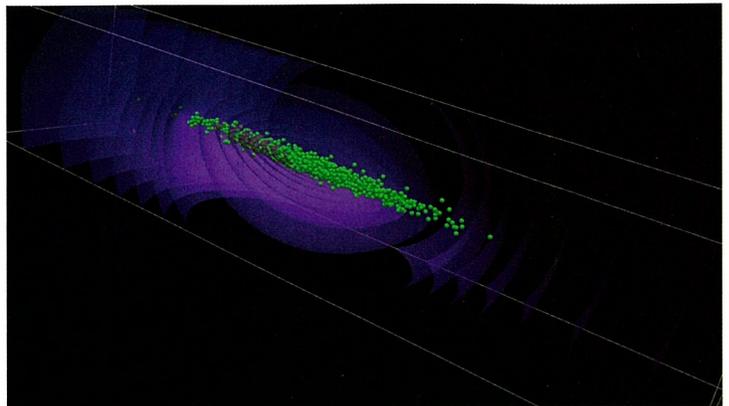


Figure 1. Macroparticles with space charge (scalar electric) field. The number of macroparticles has been reduced by a factor of 1,000 for clarity. An animated version of this figure is available at <http://compacc.fnal.gov/~amundson/animation-synergia.m4v>.

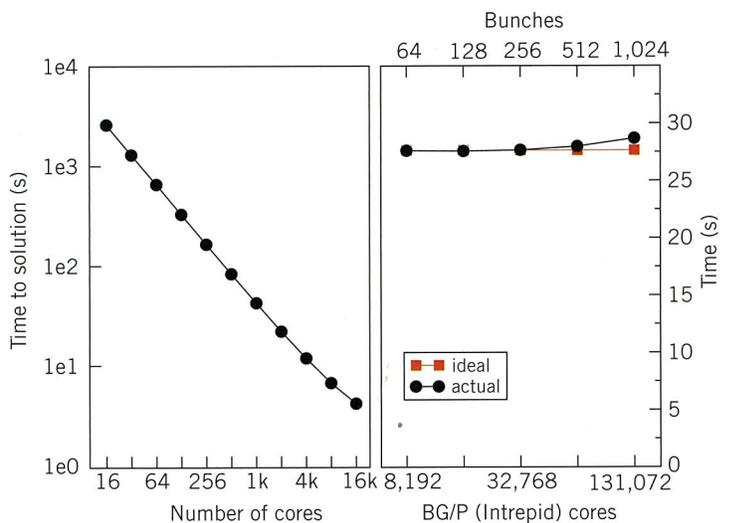


Figure 2. Code scalability. (a) Strong scaling of a single-bunch space charge simulation on the Argonne Leadership Computing Facility's (ALCF's) Mira, a Blue Gene/Q machine. The space charge calculation uses a $32 \times 32 \times 1,024$ grid and 100 million macroparticles. (b) Weak scaling of a multiple-bunch space charge simulation on ALCF's Intrepid, a Blue Gene/P machine. The space charge calculation uses a $32 \times 32 \times 1,024$ grid and 100 million macroparticles per bunch; the largest simulation has a total of over 13 billion macroparticles.

implementations is the charge deposition calculation. In this case, multiple threads need to write to a single grid in memory, resulting in locking issues. After extensive experimentation, we've settled on a red-black scheme for interleaved writes. Figure 3 shows our preliminary results after running a full Synergia benchmark on Tesla and Kepler GPUs in comparison with results from a single Intel Xeon

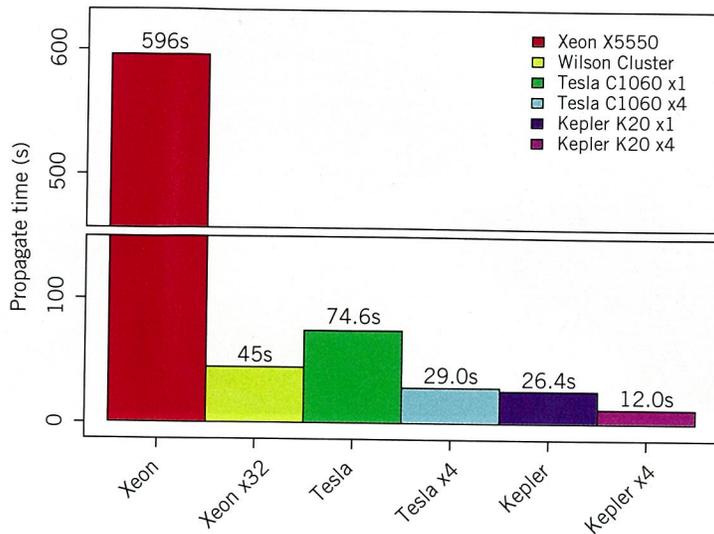


Figure 3. The overall simulation time for one turn of a particle beam on different platforms. Running a full Synergia benchmark on Tesla and Kepler GPUs offers promising results.

processor and a cluster of Intel Xeons. The results are promising so far. We plan to release production versions of Synergia for GPUs and MICs next year.

Modeling the Fermilab Booster

The Fermi National Accelerator Laboratory (FNAL) Booster accelerator is an approximately 150-meter-diameter proton synchrotron with an injection energy of 400 megaelectron volts (MeV) and an extraction energy of 8 GeV. The Booster accelerator is made up of 96 dipole-quadrupole combined function magnets (focusing [F] and defocusing [D]) in a series of 24 repeating periods. The magnets' vacuum chamber has a quasi-flat geometry consisting of two parallel planes along the horizontal direction. The combined function magnets cover about 60 percent of the machine length—the rest of the machine is made up of straight metallic beam pipe sections. A consequence of the presence of bare laminations is the formation of very large wake fields. Since the machine runs at low energy (injection $E/\text{final } E = 0.4 \text{ GeV}/8 \text{ GeV}$), the space charge is also strong. The Booster accelerator runs with an average repetition rate of 9 Hz at an intensity of 4.5×10^{12} protons per batch, which is about two times larger than the originally designed intensity. A batch consists of 84 bunches. The next generation of Fermilab neutrino experiments require even higher proton output, approximately 6×10^{12} protons per pulse at 15 Hz.

Two intensity-dependent effects are important in the Fermilab Booster: space charge and wake fields in the laminated magnets. We've developed a detailed model of wake fields in laminated structures in Synergia and validated it with experimental results.⁹ The wake field and space charge calculations are inherently coupled through the boundary conditions in the space charge field, so our model has to include space charge with compatible boundary conditions. These collective effects in turn affect particle propagation through the accelerators, often creating instabilities that result in particle losses and degradation of beam quality. Simulations can provide insight into the instabilities' mechanism and the necessary guidance to improve beam loss, but all the relevant physics, including accurate single-particle dynamics (determined by accelerator description accuracy) and accurate multiparticle dynamics (space charge, single and multibunch wake field effects), should be properly considered.

There's a long history of successful Synergia applications in support of the high-intensity FNAL physics program. Synergia simulations were used to study emittance growth and beam halo generation in the Fermilab Booster during the second collider mode run of the FNAL program, which began in 2001. Our models also enabled the first-ever simulation of Linac microbunch capture, debunching, and acceleration, including beam position feedback, three-dimensional space charge,¹⁰ and multibunch impedance effects.⁹ These simulations provided guidance to machine operators to reduce losses, maximize intensity, and commission the Booster collimators, which were essential to the neutrino program's success during the second Tevatron run. Our current work aims to support the PIP program, which has much more demanding intensity requirements, by running simulations of a Booster accelerator model that incorporates all the essential single- and multiparticle dynamic effects. To validate our model's accuracy, we compared our simulation results with beam experiments. To demonstrate the computation's scale, the results presented next used roughly 20 million core-hours, primarily running 16,384-core jobs with 10^{12} particle steps per job.

Model Construction

The accelerator's single-particle dynamics are determined by the positions and strengths of the different magnets and accelerating cavities that comprise the accelerator *lattice*. To ensure agreement between the lattice model and the real lattice,

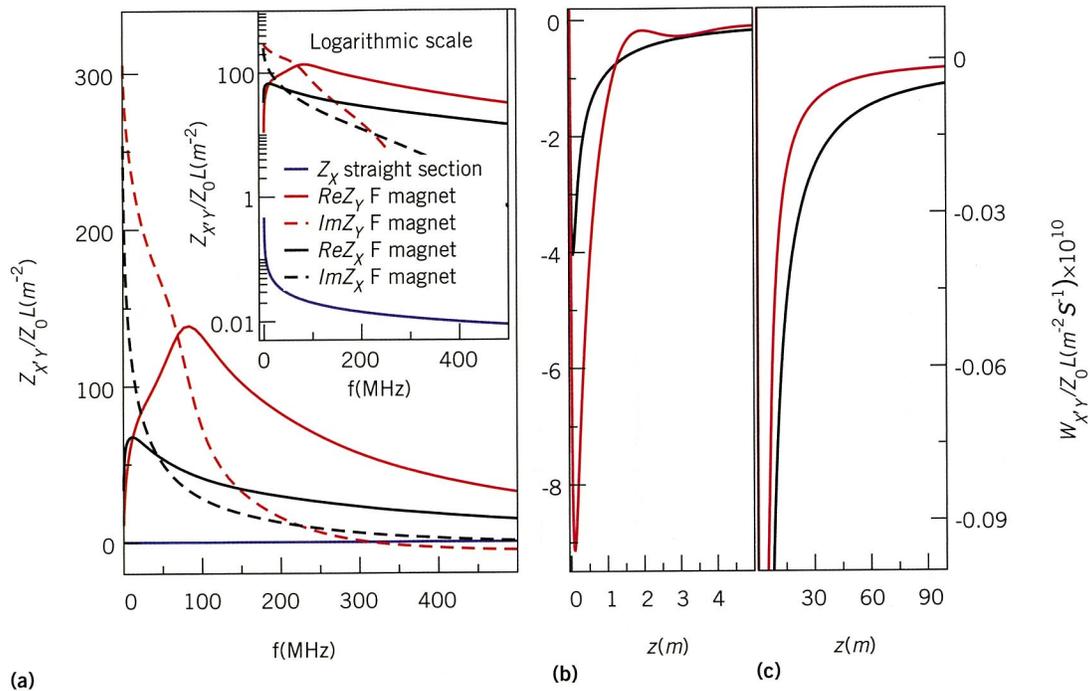


Figure 4. Impedances and wakes. (a) Transverse impedances Z_X and Z_Y for the laminated F magnet and the metallic straight section in the Fermilab Booster. The impedance in the laminated magnets (see inset, logarithmic scale) is three to four orders of magnitude larger than in the metallic pipe and has a very different frequency behavior. At small frequencies, the horizontal impedance is larger than the vertical one. (b) Horizontal and (c) vertical wakes. At a short distance, the vertical wake is about two times larger than the horizontal one, but the situation changes at larger distances on the order of the few bunch lengths (1 bunch length = 5.654 m) relevant for instabilities.

we determined the parameters of the dipole and quadrupole correctors in the model lattice by using Orbit Response Measurement to fit the measured data.¹¹ For a realistic simulation, it's important to have an accurate estimate of wake functions. Wake calculation requires solving the electromagnetic problem for the chamber in which the beam propagates—specifically, the solution depends on chamber geometry and boundary conditions for the electromagnetic field at the vacuum chamber walls.

Due to the exposure of the laminations, the impedances (which are related to the wake function via a Fourier transform) and wakes are orders of magnitude larger in the combined function magnets than in the metallic pipes. We calculated the impedances and wake fields of the Booster's laminated magnets at injection energy. The details of impedance calculations for flat laminated chambers in the ultrarelativistic limit appear elsewhere,⁹ as do the nonultrarelativistic effects.⁸ Figure 4a shows the horizontal and vertical impedances for the laminated F magnet and the transverse impedance in the straight metallic pipe. Besides the

fact that the impedance in the gradient magnets is three or four orders of magnitude larger than the one in the metallic pipe, the frequency behavior is also quite different.

Model Validation and Analysis

The measurements in the Booster accelerator exhibit many unconventional effects. Our simulations were able to capture very well the experimentally observed behavior.

For example, the vertical coherent tune—that is, the frequency of the bunch centroid's oscillations—is strongly suppressed at large intensity due to the strong wakes and the large space charge. However, the horizontal tune increases slightly with beam intensity. This difference between the horizontal and the vertical tune behaviors is a result of the flat geometry of the combined magnet's vacuum chamber.¹²

Using our simulation, we calculated the coherent tune by doing a Fourier transform of the beam centroid position as a function of propagation length. Figure 5a shows the spectral

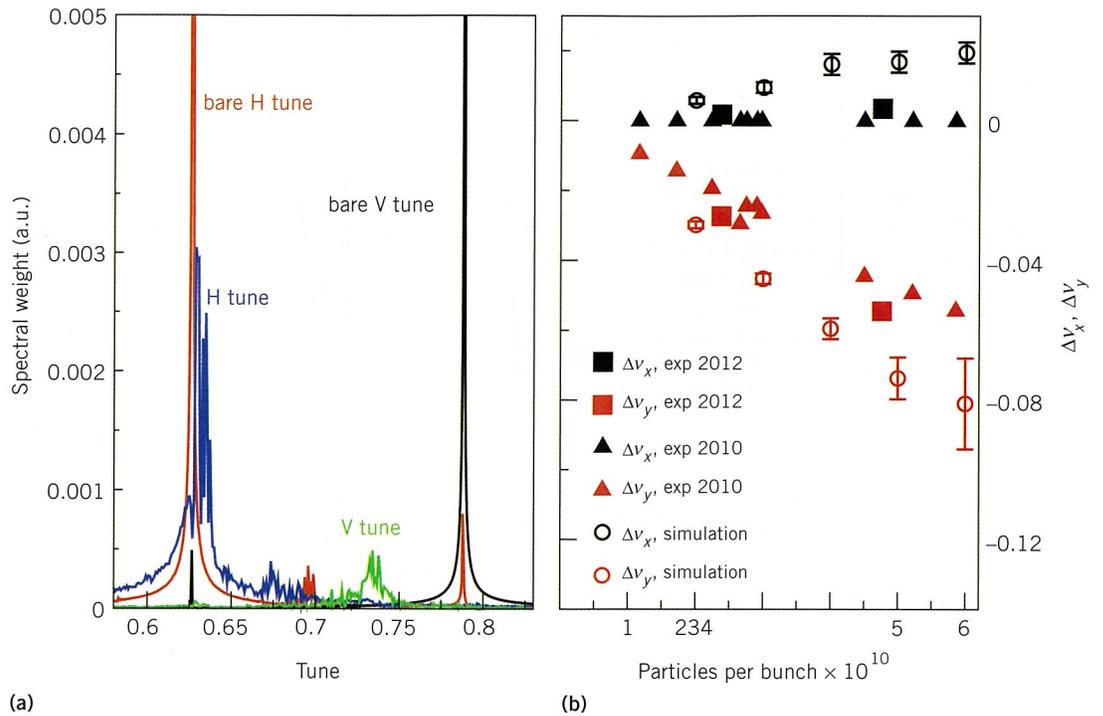


Figure 5. Simulations and comparisons with data. (a) Fourier transform of the beam centroid horizontal and vertical displacements at intensity 4×10^{10} p per bunch for the full machine (84 bunches). When the collective effects are neglected (red and black), the spectral weight (given in arbitrary units, or a.u.) exhibits sharp peaks for frequencies corresponding to the bare tunes. The spectral weight shows small positive horizontal (blue) and large negative vertical (green) tune shifts when the collective effects are present. Note the wide spectral features when the collective effects are included. (b) Coherent tune shift versus intensity in the Fermilab Booster at injection. The simulation results are compared with two experimental measurements, one in 2010 and the other in 2012. The horizontal tune increases slightly with intensity, while the vertical tune shows a strong decrease.

features. When no collective effects are included, the spectrum shows sharp peaks at frequencies corresponding to the bare tunes. With the collective effects taken into account, the spectral weight shows small positive horizontal and large negative vertical tune shifts. Aside from that, the spectral features are broad, indicating an interaction between multiple modes. Figure 5b compares calculated tune shifts with experimental data.

A rather puzzling effect observed in the Booster accelerator is the presence of horizontal instability. Whereas the vertical wake is much larger than the horizontal wake, the bunch propagation is subject to instability in the horizontal plane. Our simulations were able to reproduce this behavior; Figure 6 compares experimental data to our simulation. The beam envelope at the instability's onset is extracted from the experimental data shown in Figure 6a and plotted with magenta on top of the simulation results in Figure 6b. A large horizontal

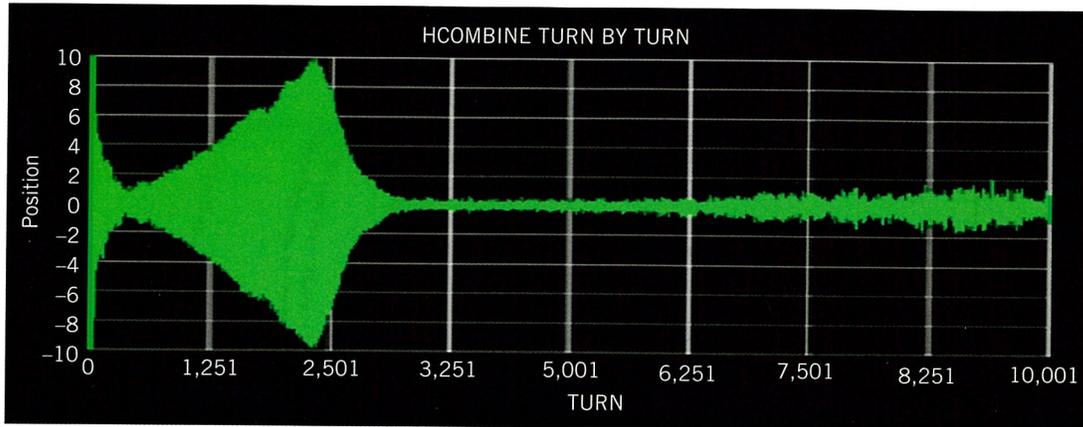
chromaticity is required to stabilize the beam.¹³ (Chromaticity measures tune dependence on momentum/energy spread. We define chromaticity

$$\text{as } \omega_\xi = \omega_0 \frac{\xi}{\eta}, \text{ where } \omega_0 \text{ is the revolution angular frequency of the synchronous particle, } \xi = \frac{p\delta\nu}{\delta p},$$

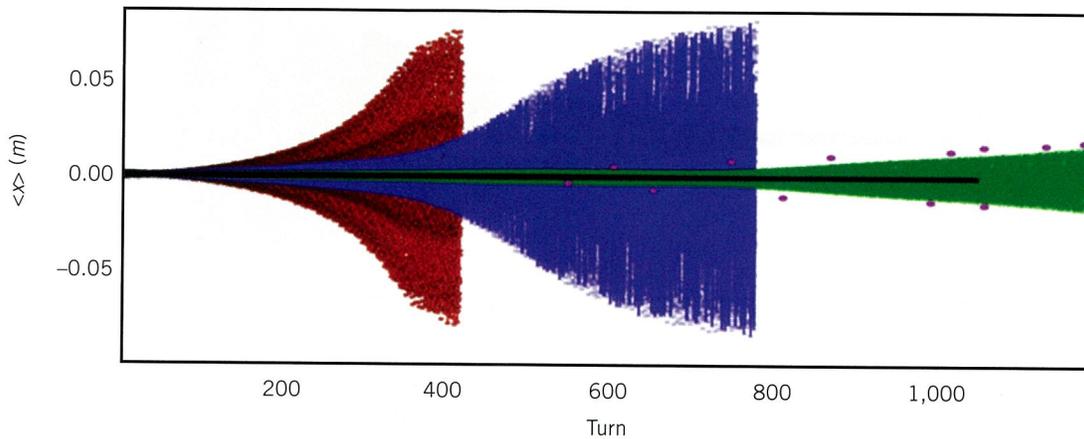
with $\delta\nu$ and $\frac{\delta p}{p}$ being the tune and momentum relative

spread, respectively, and η the slippage factor that measures the revolution frequency's dependency on the particle momentum spread.¹⁴)

To understand the instability present in the system, we performed simulations with modified interaction terms. We neglected the space charge term and focused on the different wake terms in Equations 5, 6, and 7. Contrary to previous speculations,¹³ we found that the vertical wake wasn't responsible for the instability—it was still present when all wake terms except W_X were set to zero, as



(a)



(b)

Figure 6. Experimental instability data versus our simulation. (a) Beam horizontal instability measured at $\left(\frac{\omega_{\xi x}}{\beta c}, \frac{\omega_{\xi y}}{\beta c}\right) = 2\pi \times (0.06, 0.025)m^{-1}$ for the intensity 4×10^{12} p per batch.¹³ (b) Beam centroid horizontal displacement at beam position monitor (BPM) location versus the turn number for different horizontal chromaticities, $\frac{\omega_{\xi x}}{\beta c} = 2\pi \times 0.023m^{-1}$ (red), $\frac{\omega_{\xi x}}{\beta c} = 2\pi \times 0.046m^{-1}$ (blue), $\frac{\omega_{\xi x}}{\beta c} = 2\pi \times 0.069m^{-1}$ (green), and $\frac{\omega_{\xi x}}{\beta c} = 2\pi \times 0.091m^{-1}$ (black). The vertical chromaticity is kept constant, $\frac{\omega_{\xi y}}{\beta c} = 2\pi \times 0.023m^{-1}$, and the intensity is 4.2×10^{12} p per batch. The magenta points represent the beam envelope extracted from the measurement in (a) after the instability's onset. The beam shows horizontal instability unless a large horizontal chromaticity $\frac{\omega_{\xi x}}{\beta c} = 2\pi \times 0.091m^{-1}$, similar to the experiment,¹³ is considered.

the green curve in Figure 7 illustrates. Therefore, the instability was caused by the horizontal wake term, which couples with the leading particle's displacement. One reason why instability wasn't present in the vertical plane, despite the vertical wake being larger than the horizontal one, is that instability growth rate is proportional not only to wake strength but also to the square root of the particles' oscillation amplitude. In the Booster accelerator, the

amplitude of the transverse oscillations in the gradient magnets is much larger in the horizontal plane.

Figure 8 shows simulations where the horizontal dipole wake tail (that is, the wake at distances larger than a certain cutoff distance) is set to zero. We found the simulations with a full range wake to be very close to the one where the wake extended up to only five bunch lengths. By making the wake range shorter (on the order of two bunch lengths),

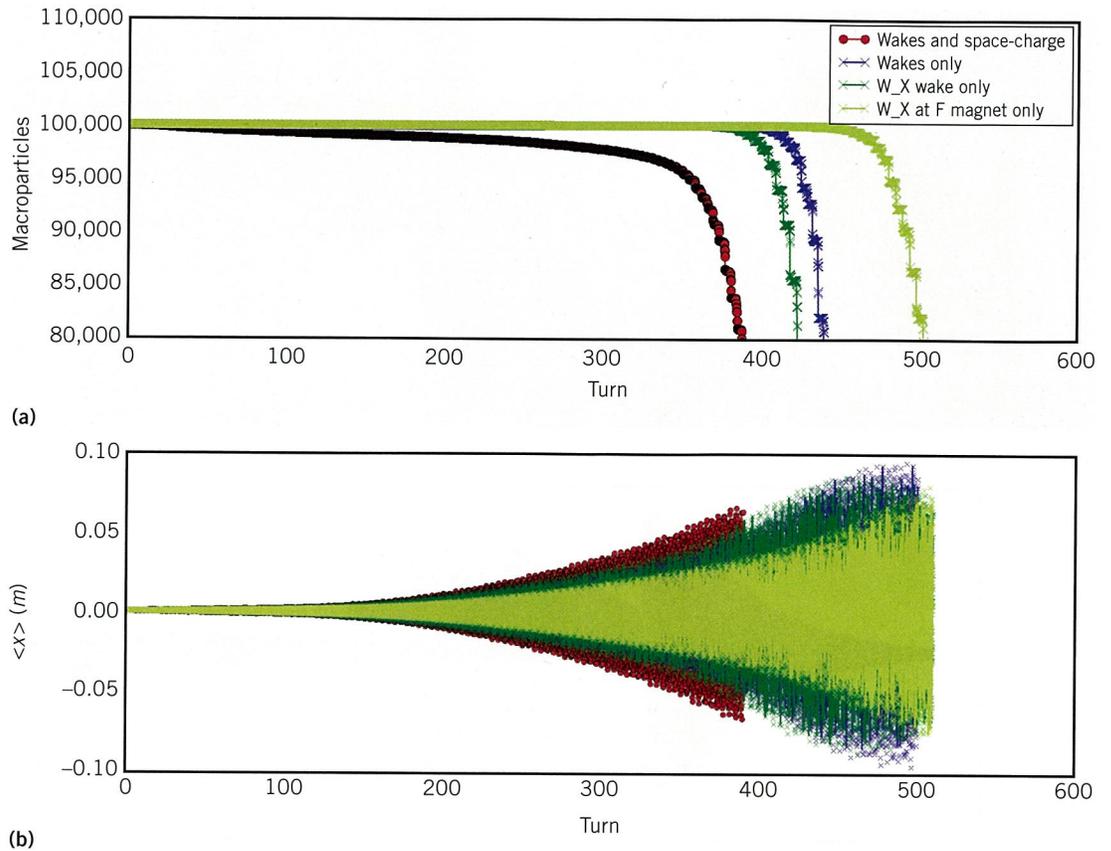


Figure 7. Simulations with modified wakes. (a) Number of macroparticles. (b) Beam horizontal centroid. The horizontal instability is present when the direct space charge and all wake terms except the horizontal one at the location of the focusing magnet are set to zero.

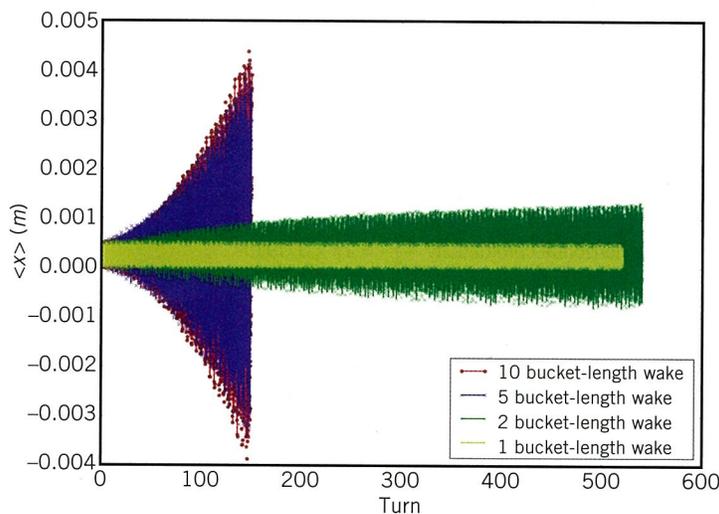


Figure 8. Simulations with short wakes. All wake terms and space charge except the dipole term W_x are set to zero. In the red (blue, green yellow) plot, W_x at a distance larger than 10 (5, 2, 1) bunch lengths is set to zero. The relevant wake range for instability is on the order of a few bunch lengths.

we found a strong suppression of the instability. One conclusion is that the relevant wake range for the instability is between one and five bunch lengths.

At a small distance—shorter than one bunch length—the vertical wake is roughly two times larger than the horizontal one. However, at larger distances in the relevant distance region for instability, the horizontal wake is larger than the vertical one. Figures 4b and 4c show this wake behavior, which can also be deduced by noticing that the horizontal impedance at a small frequency (Figure 4a) is larger than the vertical one.

Our analysis showed that instability is caused by short-range bunch-bunch interactions via dipole horizontal wake. The reason for instability in the horizontal plane is twofold: the large amplitude oscillations in the horizontal plane at the location of the focusing magnets, and at the relevant

distance (between one and five bucket lengths), the horizontal wake is larger than the vertical one. We are continuing both the experimental and simulation campaigns to better understand and characterize this effect. ■

Acknowledgments

This work was performed at Fermilab, operated by Fermi Research Alliance, LLC, under contract De-AC02-07CH11359 with the US Department of Energy. It was also supported by the COMPASS project, funded via the Scientific Discovery through Advanced Computing program in the DoE Office of High Energy Physics. We also used resources from the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the DoE under contract DE-AC02-06CH11357.

References

1. Particle Physics Project Prioritization Panel, "US Particle Physics: Scientific Opportunities, a Strategic Plan for the Next Ten Years," US Dept. of Energy, 2008; http://science.energy.gov/-/media/hep/pdf/files/pdfs/p5_report_06022008.pdf.
2. P. Spentzouris et al., "Community Petascale Project for Accelerator Science and Simulation: Advancing Computational Science for Future Accelerators and Accelerator Technologies," *J. Physics: Conf. Series*, vol. 125, 2008, p. 012005; www.mcs.anl.gov/research/LANS/publications/index.php?p=pub_detail&id=178.
3. Fermilab, *The Compass Project: High-Performance Computing for Accelerator Design and Optimization*, 2008; <https://web.fnal.gov/collaboration/compass/SitePages/Home.aspx>.
4. Fermilab, *Fermilab Proton Improvement Plan*, 2013; http://www-ad.fnal.gov/proton/PIP/PIP_index.html.
5. Fermilab, *Project-X, aka Fermilab Proton Improvement Plan II*, Fermilab, 2013; <http://projectx.fnal.gov>.
6. V. Lebedev et al., "Proton Source Improvement Plan," tech. memo FNAL/Beams-Doc 3781, Fermilab, 2011; <http://beamdocs.fnal.gov/AD-public/DocDB/ShowDocument?docid=3781>.
7. L. Michelotti, "C++ Objects for Beam Physics," *Proc. 14th IEEE Particle Accelerator Conf.*, 1991, pp. 1561–1563.
8. A. Macridin, P. Spentzouris, and J. Amundson, "Nonperturbative Algorithm for the Resistive Wall Impedance of General Cross-Section Beam Pipes," *Physical Rev. ST Accelerator Beams*, vol. 16, Dec. 2013, p. 121001.
9. A. Macridin et al., "Coupling Impedance and Wake Functions for Laminated Structures with an Application to the Fermilab Booster," *Physical Rev. ST Accelerator Beams*, vol. 14, 2011, p. 061003.
10. J. Amundson et al., "An Experimentally Robust Technique for Halo Measurement Using the IPM at the Fermilab Booster," *Nuclear Instruments and Methods*, vol. A570, 2007, pp. 1–9.
11. V. Lebedev et al., "Model Calibration and Optics Correction Using Orbit Response Matrix in the Fermilab Booster," *Proc. Int'l Particle Accelerator Conf.*, 2012, pp. 1251–1253; <http://accelconf.web.cern.ch/accelconf/ipac2012/papers/tuppc040.pdf>.
12. A. Chao, S. Heifets, and B. Zotter, "Tune Shifts of Bunch Trains Due to Resistive Vacuum Chambers without Circular Symmetry," *Physical Rev. ST Accelerator Beams*, vol. 5, Nov. 2002, p. 111001.
13. Y. Alexahin et al., "Observation of Instabilities of Coherent Transverse Oscillations in the Fermilab Booster," *Fermilab Conf. Proc.*, 2012, pp. 3129–3131.
14. M. Conte and W.W. MacKay, *An Introduction to the Physics of Particle Accelerators*, World Scientific, 1994.

James Amundson is the Synergia Project technical lead and head of the Scientific Software Infrastructure Department at *Fermilab National Accelerator Laboratory*. His research focuses on high-performance parallel computing and its intersection with computational accelerator physics. Amundson has a PhD in theoretical particle physics from the University of Chicago. He's a member of the American Physical Society. Contact him at amundson@fnal.gov.

Alexandru Macridin is a computational physicist in the Scientific Computing Division at *Fermilab National Accelerator Laboratory*. His research interests include accelerator physics, beam dynamics, collective effects in accelerators, numerical methods, parallel computing, and many-body physics. Macridin has a PhD in physics from the University of Groningen. Contact him at macridin@fnal.gov.

Panagiotis Spentzouris is head of the Scientific Computing Division at the Fermi National Accelerator Laboratory. His research interests include computational accelerator physics and neutrino physics. Spentzouris has a PhD in physics from Northwestern University. He's a member of the American Physical Society. Contact him at spentz@fnal.gov.



Selected articles and columns from IEEE Computer Society publications are also available for free at <http://ComputingNow.computer.org>.