FERMILAB-PUB-14-117-CD
ACCEPTED

High Performance Computing Modeling Advances Accelerator Science for High Energy Physics

James Amundson, Alexandru Macridin and Panagiotis Spentzouris, Fermilab

April 29, 2014

Abstract

The development and optimization of particle accelerators are essential for advancing our understanding of the properties of matter, energy, space and time. Particle accelerators are complex devices whose behavior involves many physical effects on multiple scales. Therefore, advanced computational tools utilizing high-performance computing (HPC) are essential for accurately modeling them. In the past decade, the DOE SciDAC program has produced such accelerator-modeling tools, which have been employed to tackle some of the most difficult accelerator science problems. In this article we discuss the Synergia beam-dynamics framework and its applications to high-intensity particle accelerator physics. Synergia is an accelerator simulation package capable of handling the entire spectrum of beam dynamics simulations. We present the design principles, key physical and numerical models in Synergia and its performance on HPC platforms. Finally, we present the results of Synergia applications for the Fermilab proton source upgrade, known as the Proton Improvement Plan (PIP).

1 Particle Accelerators for HEP

The U.S. Department of Energy (DOE)'s Office of High Energy Physics (HEP) promotes a broad, long-term particle physics program by supporting current operations and experiments and research as well as development for future facilities and experiments at three interrelated frontiers of particle physics [1].

- The Energy Frontier directly explores the fundamental constituents and architecture of the universe 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

Operated by Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the United States Department of Energy.

essential to meeting these grand scientific HEP challenges because it enables and catalyzes advancement in accelerator science.

1.1 Science Challenges and Research Needs

In the next decade, the HEP community will explore the intensity frontier by designing high intensity proton sources for neutrino physics and rare process searches, such as the 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, and 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 by 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 multi-scale, multi-physics simulations of the most challenging accelerator science projects and demonstrated the impact of large scale simulations in accelerator science [3]. In this article we present the Synergia beam dynamics framework and discuss its current applications, which focus on the design of high-intensity proton accelerators.

2 Intensity Frontier of HEP

The Intensity Frontier (IF) addresses central questions in particle physics that are not directly accessible with current or planned accelerators at the Energy Frontier. Experiments at the IF study rare 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 IF is the Proton Improvement Plan (PIP) at Fermilab (stages I and II) [4, 5], leading to a multi-megawatt proton accelerator that will produce intense beams of neutrinos and muons (and possibly kaons, as well as heavy nuclei, with further upgrades). The culmination of the PIP efforts will be the delivery of a new superconducting RF linac and major improvements to the Booster and Main Injector accelerators at Fermilab. High-fidelity simulations are an important component for the success and cost effectiveness of this campaign.

2.0.1 Fermilab Proton Improvement Plan

Every proton for the domestic United States high energy physics experimental program will be accelerated by the existing, now 40-year-old, Fermilab Linac and Booster 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 GeV and lower energy, and no sooner than well into the next decade to serve demand for beams at higher energy. The domestic high energy physics program for the next 15 years is dependent on the viability and vitality of the Fermilab Linac and Booster. Fermilab has established a charge for developing a plan to assure this viability and vitality. Specifically, the charge for PIP-I calls for "delivering 2.25×10^{17} protons/hour (at 15 Hz) by January 1, 2016" (more than 2 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 will achieve the target intensities for the Proton Improvement Plan it is 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 radiofrequency system, etc.). High-fidelity simulations which incorporate all the relevant physics effects are employed to guide the accelerator scientists as they undertake these tasks.

3 Synergia

3.1 Accelerator simulation problem definition

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 pre-programmed operating cycles and/or in response to active feedback. The beams in these accelerators are usually bunched longitudinally with $\mathcal{O}(10^{12})$ 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 typical longitudinal extent of meters.

Synergia[7] is an accelerator simulation package designed to take advantage of computational resources varying from desktop machines to leadership 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 of the core computations are done in C++ while end-user simulations are described using Python. The combination provides efficiency and great flexibility.

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,

$$\mathcal{H} = \mathcal{H}_i + \mathcal{H}_c, \tag{1}$$

The split-operator approximation for the time step evolution operator, $\mathcal{M}_{full}$ is given by

$$\mathcal{M}_{full}\left(t\right) = \mathcal{M}_i\left(\frac{t}{2}\right)\mathcal{M}_c\left(t\right)\mathcal{M}_i\left(\frac{t}{2}\right) + \mathcal{O}\left(t^3\right), \tag{2}$$

where $\mathcal{M}_c$ and $\mathcal{M}_i$ are the evolution operators corresponding to $\mathcal{H}_c$ and $\mathcal{H}_i$, 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 consists of thousands to a hundred thousand turns. Real accelerator cycles can be 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 (and/or user-defined) *diagnostics* can be performed on the bunch(es) to monitor the state of the beam 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. The end-user program may use only existing Synergia classes, or may include custom extensions to those classes. The entire state of the simulation, including end-user extensions, may be checkpointed and/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 which are exposed to Python via Boost.Python[8]. The independent-particle dynamics are handled by the CHEF[9] C++ libraries. Both Synergia and CHEF were developed at Fermilab.

3.2 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, etc.) are much smaller than the forces due to interactions with other particles in the beam. The CHEF libraries[9], which we use in Synergia, provide a comprehensive set of tools for linear and non-linear independent-particle accelerator physics.

CHEF not only models the propagation of particles through accelerators, it also performs symbolic algebra calculations on the transfer maps for the propagation, allowing the user to perform analyses such as non-linear map analysis of resonant structures, etc. 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_i, are related to the initial coordinates, u_i by

$$u_f = \mathcal{M}(u_i), \tag{3}$$

where $\mathcal{M}$ 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.

3.3 Collective effects

Synergia includes a general treatment of collective accelerator physics effects via the split-operator approximation described above. 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 the like-charged particles in the beam. Wake fields 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. The design of Synergia allows for the implementation of additional collective effects and/or different approximations/implementations of the existing space charge and wake field models by the user.

3.3.1 Space charge

Calculating the effects of space charge requires solving the Poisson Equation,

$$\nabla^2\phi = -\frac{\rho}{\epsilon_0}, \tag{4}$$

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 of order meters) is usually much greater than the transverse dimensions of the beam (typically of order 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. Fig. 1 shows a space charge solve with macroparticles and the resulting scalar field.

3.3.2 Wake fields

The effects of induced wake fields in beam pipes with horizontal and vertical mirror symmetry can be summarized using the expressions [10]

$$\beta c\Delta p_z = -qQW^{||}(z), \tag{5}$$
$$\beta c\Delta p_x = -qQ\left(W_X(z)X + W_x(z)x\right) \tag{6}$$
$$\beta c\Delta p_y = -qQ\left(W_Y(z)Y + W_y(z)y\right), \tag{7}$$

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

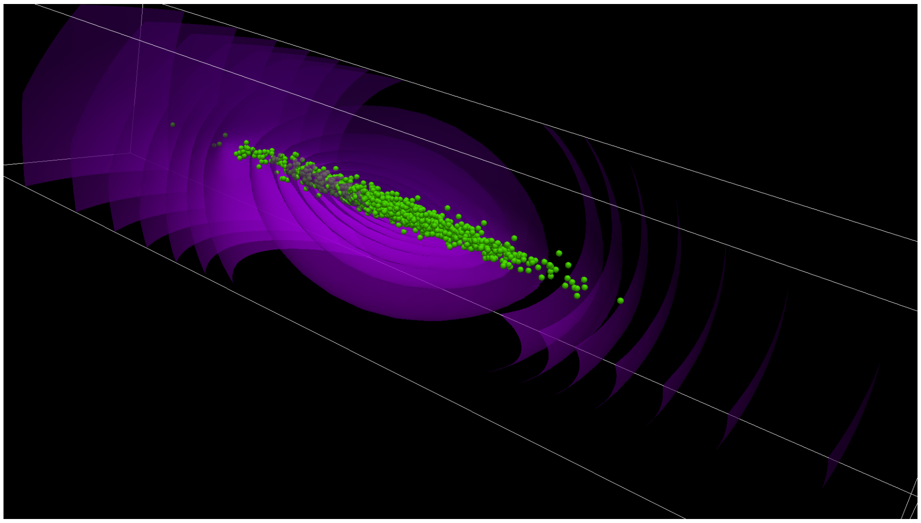

Figure 1: Macro particles with space charge (scalar electric) field. The number of macroparticles has been reduced by a factor of 1000 for clarity. An animated version of this figure is available at `http://compacc.fnal.gov/~amundson/animation-synergia.m4v`.

3.4 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 one can easily utilize a range of parallel resources ranging from multi-core desktops to 100,000+ core supercomputers. Fig. 2 demonstrates the scalability of the code. We find excellent strong scaling behavior over a range of a little less than a thousand on a Blue Gene/Q machine. We also obtain excellent weak scaling in the number of macroparticles used (not shown) and number of bunches, shown in Fig. 2 -b).

Scalability in Synergia is limited by the communication requirements of collective effects. The scaling we have achieved to date relies on two complementary techniques: communication avoidance and hybrid MPI/OpenMP optimizations. In communication avoidance we perform multiple, redundant field solves in order 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 is enabled, the number of threads per process can be chosen at run time.

In preparation for the next generation of HPC technologies, we are 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 implementations is the charge deposition calculation. In that case, multiple threads need to write to a single

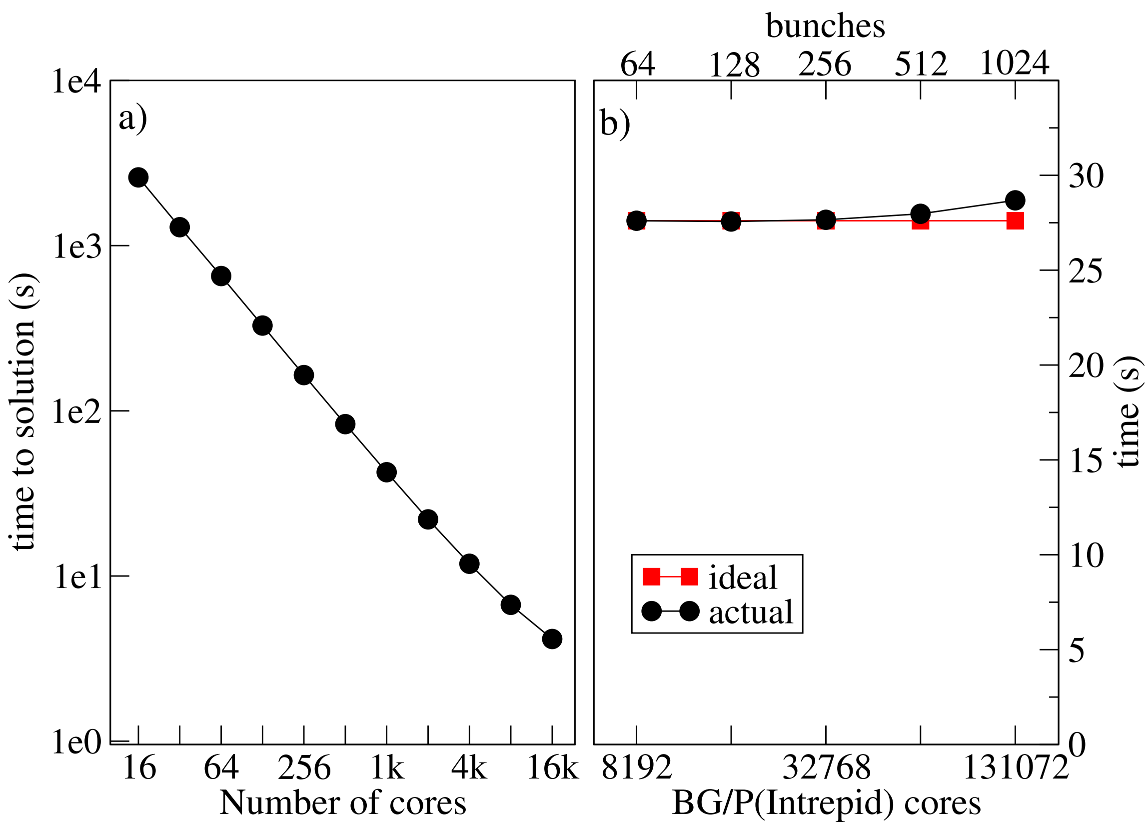

Figure 2: a) Strong scaling of a single-bunch space charge simulation on ALCF's Mira, a Blue Gene/Q machine. The space charge calculation uses a $32 \times 32 \times 1024$ grid and 100 million macroparticles. The red curve is the total time; other curves detail the scaling of the most computationally demanding portions of the calculation. 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 $32 \times 32 \times 1024$ grid and 100 million macroparticles per bunch; the largest simulation has a total of over 13 billion macroparticles.

grid in memory, resulting in locking issues. After extensive experimentation, we have settled on a red-black scheme for interleaved writes. Fig. 3 shows preliminary results we have obtained running a full Synergia benchmark on Tesla and Kepler GPUs in comparison with results from a single Intel Xeon processor and a cluster of Intel Xeons. The results are promising so far. We plan to release production versions of Synergia ready for GPUs and MICs in the next year.

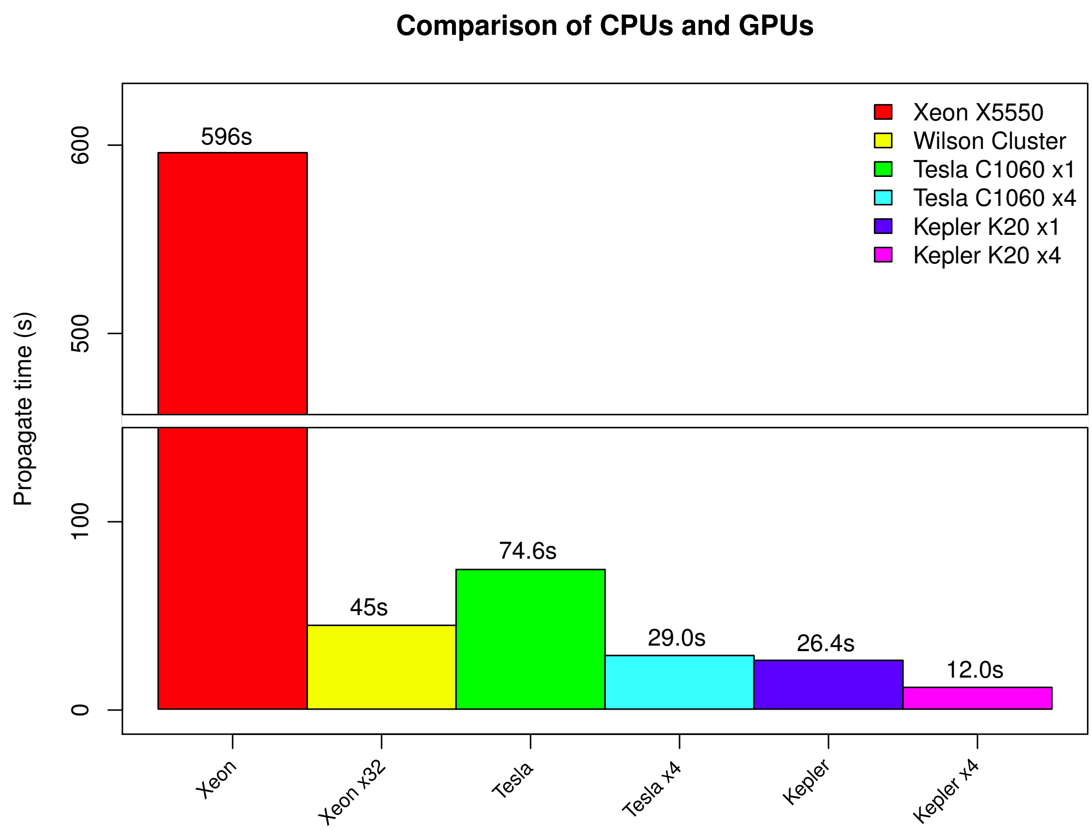

Figure 3: The overall simulation time for one turn of particle beam on different platforms.

4 Modeling the Fermilab Booster

The FNAL Booster accelerator is an approximately 150-meter diameter proton synchrotron with an injection energy of 400 MeV and an extraction energy of 8 GeV. The Booster is made up of 96 dipole-quadrupole combined function (focusing (F) and defocusing (D)) magnets 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% 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/final $E = 0.4GeV/8GeV$), the space charge is also strong. The Booster 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, $\sim 6 \times 10^{12}$ protons per pulse at 15Hz.

Two intensity-dependent effects are important in the Fermilab Booster: space charge and wake fields in the laminated magnets. We have developed a detailed model of wake fields in laminated structures in Synergia and validated it with experimental results [11]. 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 affect the propagation of the particles through 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 the beam loss, but it is necessary that all the relevant physics, including accurate single-particle dynamics (determined by the accuracy of the accelerator description) and accurate multi-particle dynamics (space-charge, single and multi-bunch wake fields effects), should be properly considered.

There is 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 Run II. Our models enabled the first-ever simulation of linac microbunch capture, debunching, and acceleration, including beam position feedback, three-dimensional space-charge [12], and multi-bunch impedance effects [11]. These simulations provided guidance to machine operators to reduce losses and maximize intensity and to commission the Booster collimators, which were essential to the success of the neutrino program during the Tevatron Run II. Our current work aims to support the PIP program, which has much more demanding intensity requirements, by running simulations of a Booster model that incorporates all the essential single- and multi-particle dynamic effects. In order to validate the accuracy of our model we compare our simulation results with beam experiments. The results of the validations and some preliminary conclusions regarding the behavior of the Booster are presented below. To demonstrate the scale of the computation, the results presented below used roughly 20 million core-hours, primarily running 16,384-core jobs with $\sim 10^{12}$ particle-steps per job.

4.1 Model construction

The single-particle dynamics of the accelerator are determined by the positions and strengths of the different type magnets and accelerating cavities that comprise the accelerator *"lattice."* To ensure agreement between the lattice model and the real lattice, the parameters of the dipole and quadrupole correctors in the model lattice were determined using Orbit Response Measurement fitting of the measured data [13]. For a realistic simulation it is important to have an accurate estimate of the wake functions. Calculation of the wakes requires solving the electromagnetic problem for the chamber that the beam propagates in. The solution is dependent on the chamber geometry and on the boundary conditions for the electromagnetic field at the vacuum chamber walls. Due to the exposure of the laminations, the impedances (which is 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 calculate the impedances

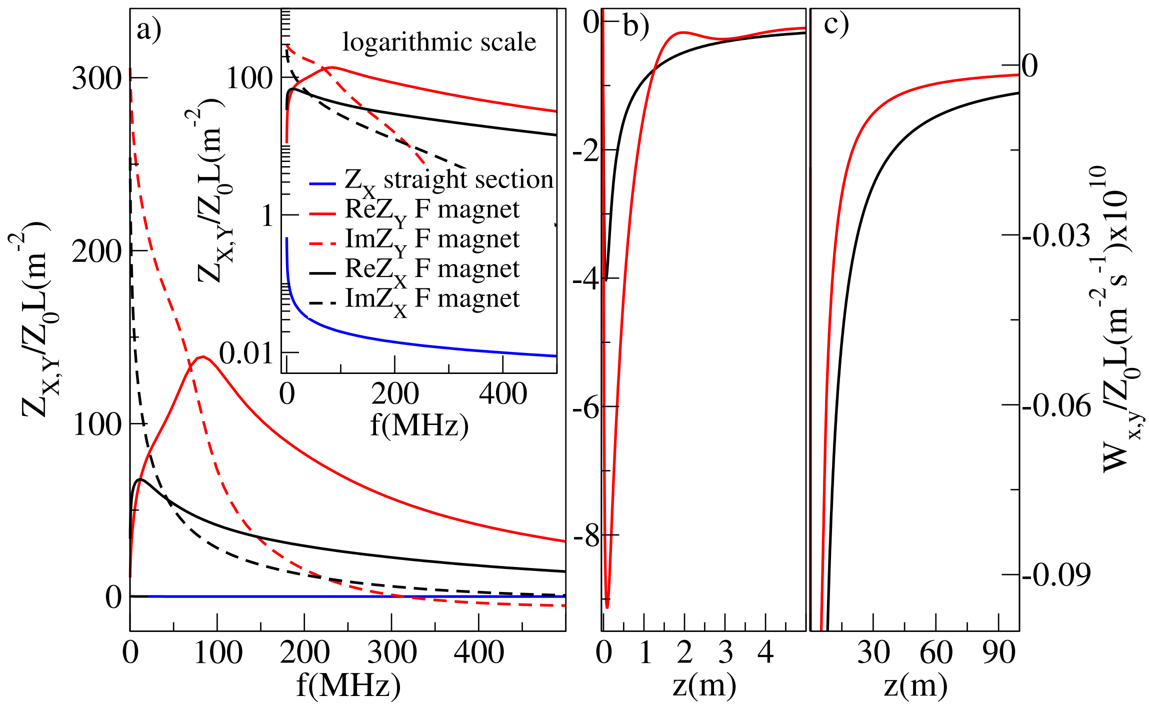

Figure 4: a) Transverse impedances Z_X and Z_Y for the laminated F magnet and the metallic straight section in Fermilab Booster. The impedance in the laminated magnets (see inset, logarithmic scale) is three-four orders of magnitude larger than in the metallic pipe and has a very different frequency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

and the wake fields of the Booster laminated magnets at injection energy. The details of impedance calculations for flat laminated chambers in the ultrarelativistic limit are presented in [11], while we discuss the nonultrarelativistic effects in [14]. In Figure 4 -a) we show 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 very different.

4.2 Model validation and simulation analysis

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

For example, the vertical coherent tune (*i.e.*, the frequency of the oscillations of the

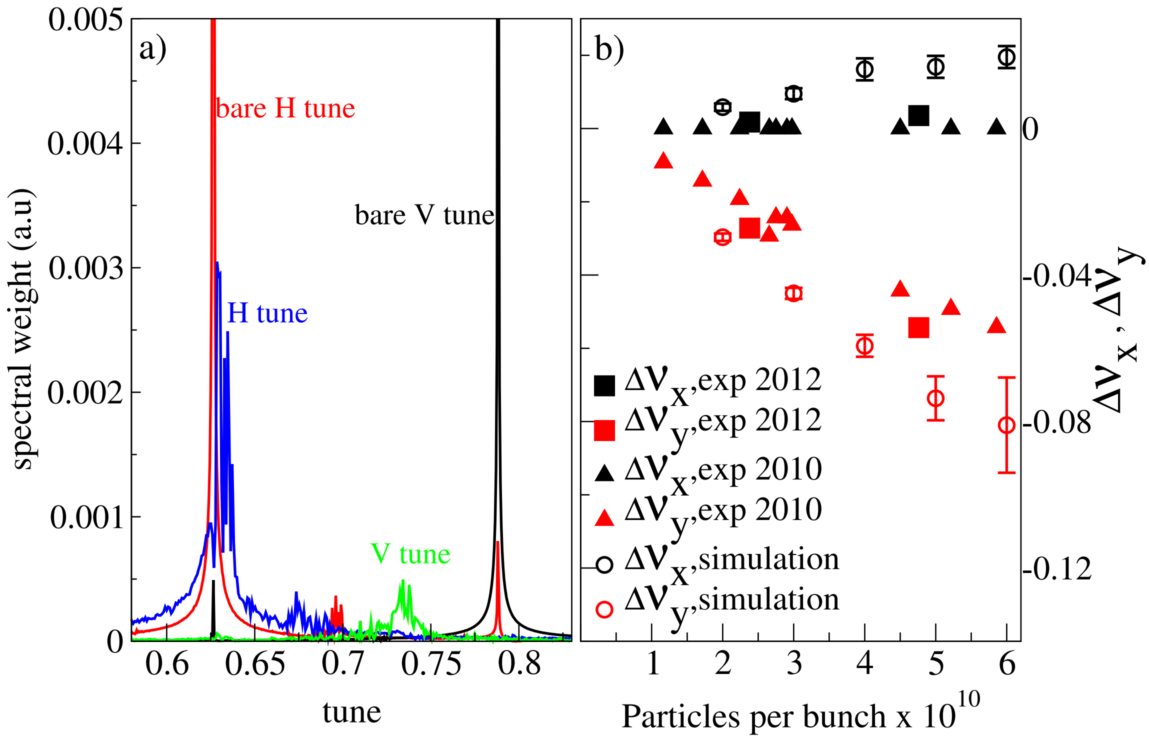

Figure 5: a) Fourier transform of the beam centroid horizontal and vertical displacements at intensity 4×10^{10} p per bunch for the full machine (84 bunches). When the collective effects are neglected (red and black), the spectral weight exhibits sharp peaks for frequencies corresponding to the bare 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.

bunch centroid) is strongly suppressed at large intensity due to the strong wakes and to 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 [10]. Using our simulation, we calculate the coherent tune by doing a Fourier transform of the beam centroid position as a function of the propagation length. The spectral features are present in Fig. 5 -a). 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 the

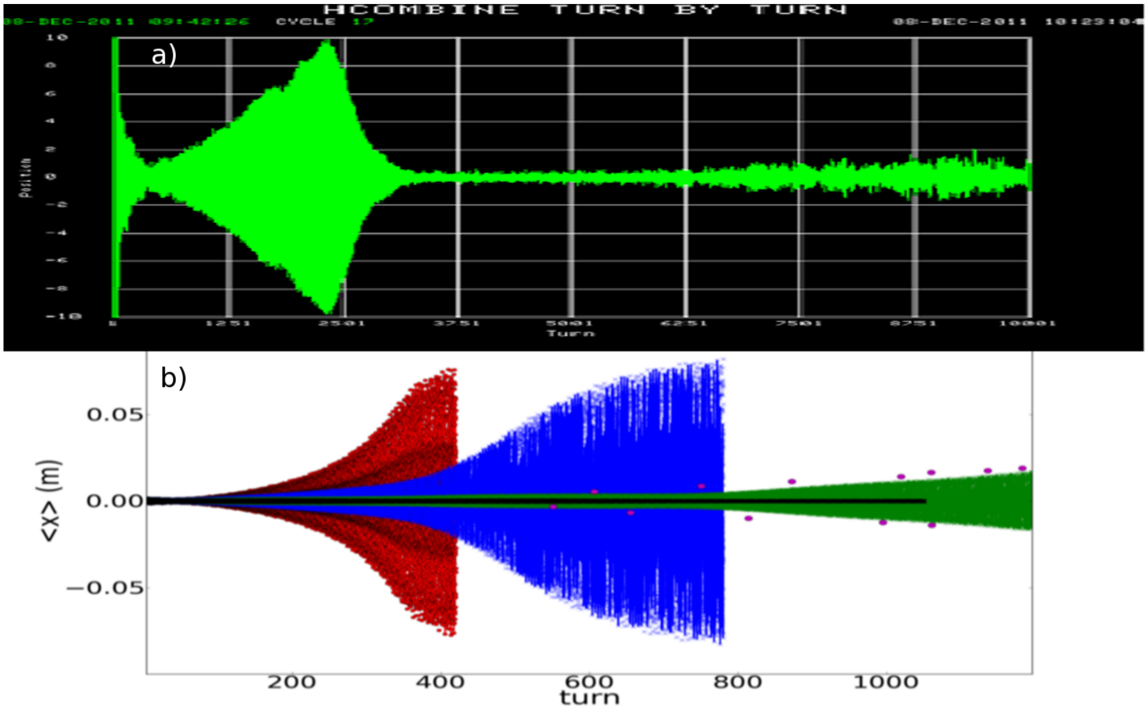

Figure 6: a) Beam horizontal instability measured at $(\frac{\omega_{\xi x}}{\beta c}, \frac{\omega_{\xi y}}{\beta c}) = 2\pi \times (0.06, 0.025)m^{-1}$ for the intensity 4×10^{12} p per batch [15]. b) Beam centroid horizontal displacement at BPMs location versus turn number for different horizontal 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, $\frac{\omega_{\xi y}}{\beta c} = 2\pi \times 0.023m^{-1}$. 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 onset of the instability. 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 [15], is considered.

evidence of an interaction between multiple modes. The calculated tune shifts are compared with the experimental data in Fig. 5 -b).

A rather puzzling effect observed in the Booster is the presence of a 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 are able to reproduce this behavior. A comparison of experimental data to our simulation is presented in Fig. 6. The beam envelope at the onset of the instability is extracted from the experimental data shown in Fig. 6 a) and plotted with magenta on top of the simulation results in Fig. 6 b). A large horizontal chromaticity[1] is required to stabilize

[1]The chromaticity measure the tune dependence on the momentum/energy spread. We define the

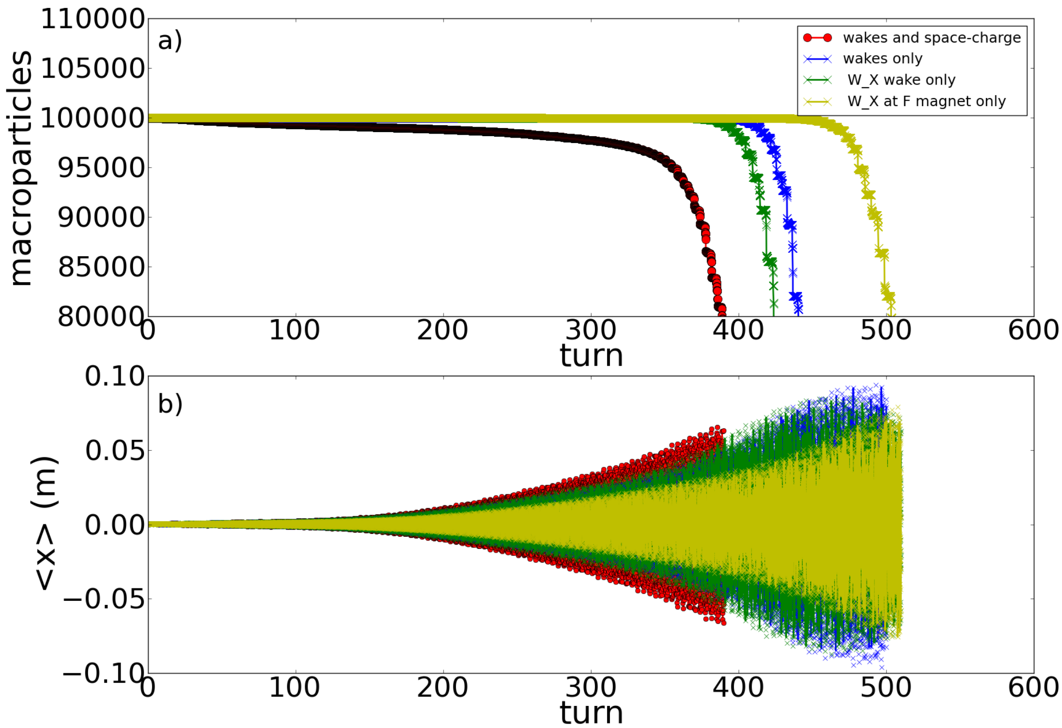

Figure 7: Simulations with modified wakes. a) Number of macroparticles. b) Beam horizontal centroid. The horizontal instability is present when the direct space charge and all wake terms expect the horizontal one at the location of the focusing magnet are set to zero.

the beam [15].

In order to understand the instability present in the system we perform simulations with modified interaction terms. We neglect the space-charge term and focus on the different wake terms in Equations 5, 6 and 7. Contrary to previous speculations [15] we find that the vertical wake is not responsible for the instability. The instability is still present when all wake terms except W_X are set to zero, as the green curve in Fig. 7 illustrates. Therefore the instability is caused by the horizontal wake term which couples with the displacement of the leading particle. One of the reasons why the instability is not present in the vertical plane, despite the fact that the vertical wake is larger then the horizontal one, is that the instability growth rate is proportional

chromaticity as $\omega_\xi = \omega_0 \frac{\xi}{\eta}$, where ω_0 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 is the slippage factor which measures the dependency of the revolution frequency on the particle momentum spread [16].

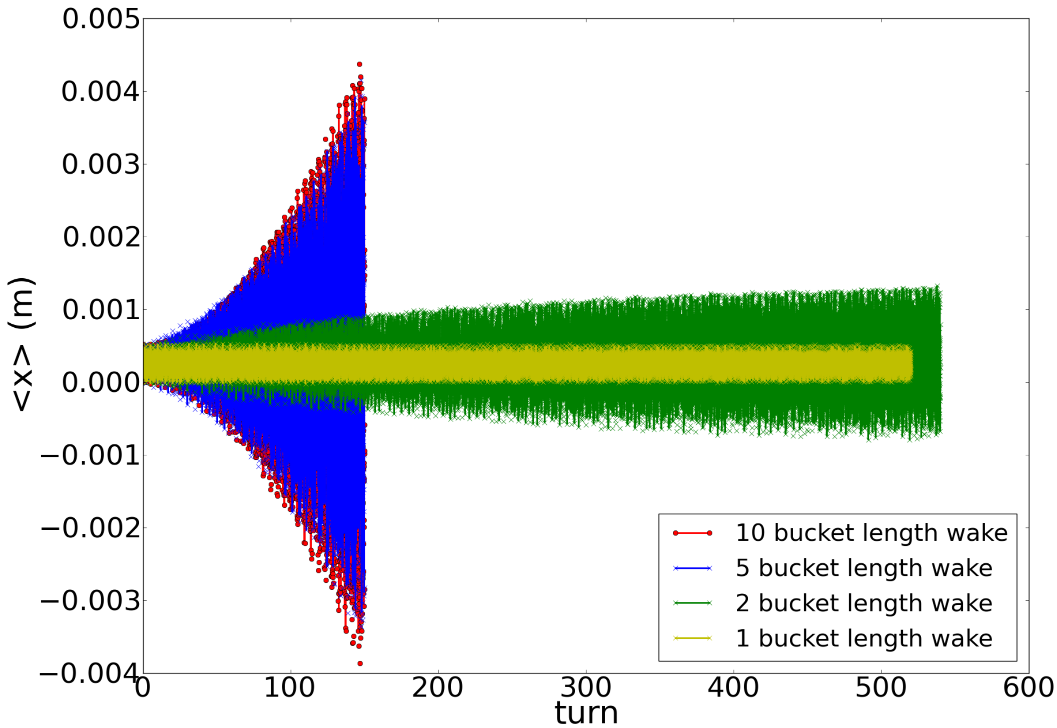

Figure 8: Simulations with short wakes. All wake terms and space charge except the dipole term W_X are set to zero. In the red (blue, green yellow) plot W_X at distance larger than 10 (5, 2, 1) bunch lengths is set to zero. The relevant wake range for instability is of order of a few bunch lengths.

not only to the wake strength but also to the square root of the particles' oscillation amplitude. In the Booster the amplitude of the transverse oscillations in the gradient magnets is much larger in the horizontal plane.

Next we address the typical range of the wake function relevant for the instability. In Fig. 8 we show simulations where the horizontal dipole wake tail (i.e the wake at distances larger than a certain cutoff distance) is set to zero. We find that the simulations with a full-range wake is very close to the one where the wake extends up to only five bunch lengths. By making the wake range shorter, of order of two bunch lengths, we see a strong suppression of the instability. We may conclude that the relevant wake range for the instability is between one and five bunch lengths.

At small distance, shorter than one bunch length, the vertical wake is about two times larger than the horizontal one. However, at larger distance in the relevant distance region for instability, the horizontal wake is larger the the vertical one. This is shown in Fig. 4 -b) and -c). This wake behavior can also be deduced by noticing that the horizontal impedance at small frequency (see Fig. 4 -a)) is larger than the vertical one.

In conclusion, our analysis shows that the instability is caused by short range (order of a few bunch lengths) bunch-bunch interactions via dipole horizontal wake. The reason for instability to be in the horizontal plane is twofold. One cause is the large amplitude oscillations in the horizontal plane at the location of the focusing magnets. The second cause is that at the relevant distance (between one and five bucket lengths) the horizontal wake is larger than the vertical one.

Acknowledgments

This work was performed at Fermilab, operated by Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the United States Department of Energy. It was also supported by the ComPASS project, funded through the Scientific Discovery through Advanced Computing program in the DOE Office of High Energy Physics. We 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] Particle Physics Project Prioritization Panel, "US particle physics: Scientific opportunities, a strategic plan for the next ten years," text available at http://science.energy.gov/~/media/hep/pdf/files/pdfs/p5_report_06022008.pdf, 2008.

[2] P. Spentzouris *et al.*, "Community petascale project for accelerator science and simulation: Advancing computational science for future accelerators and accelerator technologies," *J. Phys. Conf. Ser.*, vol. 125, p. 012005, 2008.

[3] ComPASS Collaboration, "Community petascale project for accelerator science and simulation," http://compass.fnal.gov/scidac-breakthroughs, http://compass.fnal.gov/talks-and-publications, 2008.

[4] "Fermilab Proton Improvement Plan," see http://www-ad.fnal.gov/proton/PIP/PIP_index.html.

[5] "Project-X, aka Fermilab Proton Improvement Plan II," see http://projectx.fnal.gov.

[6] V. Lebedev, D. McGinnis, W. Pellico, and B. Webber, "Proton source improvement plan," Fermi National Accelerator Laboratory, Batavia, Illinois, Technical Memorandum FNAL/Beams-Doc 3781, 2011. [Online]. Available: http://beamdocs.fnal.gov/AD-public/DocDB/ShowDocument?docid=3781

[7] J. Amundson, "Synergia2.1: A framework for accelerator simulations." [Online]. Available: https://cdcvs.fnal.gov/redmine/projects/synergia2

[8] "The boost.python c++ library." [Online]. Available: http://www.boost.org/doc/libs/1_54_0/libs/python/doc/index.html

[9] L. Michelotti, "C++ objects for beam physics," in *Proceedings of the 14th IEEE Particle Accelerator Conference*, 1991.

[10] A. Chao, S. Heifets, and B. Zotter, "Tune shifts of bunch trains due to resistive vacuum chambers without circular symmetry," *Phys. Rev. ST Accel. Beams*, vol. 5, p. 111001, Nov 2002.

[11] A. Macridin, P. Spentzouris, J. Amundson, L. Spentzouris, and D. McCarron, "Coupling impedance and wake functions for laminated structures with an application to the Fermilab Booster," *Phys.Rev.ST Accel.Beams*, vol. 14, p. 061003, 2011.

[12] J. Amundson, W. Pellico, P. Spentzouris, T. Sullivan, and L. Spentzouris, "An experimentally robust technique for halo measurement using the IPM at the Fermilab Booster," *Nucl. Instrum. Meth.*, vol. A570, pp. 1–9, 2007.

[13] V. Lebedev, E. Prebys, A. Petrenko, S. Kopp, and M. McAteer, "Model Calibration and Optics Correction Using Orbit Response Matrix in the Fermilab Booster," *Conf.Proc.*, vol. C1205201, pp. 1251–1253, 2012.

[14] A. Macridin, P. Spentzouris, and J. Amundson, "Nonperturbative algorithm for the resistive wall impedance of general cross-section beam pipes," *Phys. Rev. ST Accel. Beams*, vol. 16, p. 121001, Dec 2013.

[15] Y. Alexahin, N. Eddy, E. Gianfelice-Wendt, V. Lebedev, W. Marsh *et al.*, "Observation of Instabilities of Coherent Transverse Ocillations in the Fermilab Booster," *Conf.Proc.*, vol. C1205201, pp. 3129–3131, 2012.

[16] M. Conte and W. W. MacKay, *An introduction to the physics of particle accelerators.* World Scientific, 1994.