Parallel Particle-In-Cell Simulation of Colliding Beams in High Energy Accelerators*

Ji Qiang, Miguel A. Furman, and Robert D. Ryne
Lawrence Berkeley National Laboratory,
Berkeley, CA 94720

In this paper we present a self-consistent simulation model of colliding beams in high energy accelerators. The model, which is based on a particle-in-cell method, uses a new developed shifted-Green function algorithm for the efficient calculation of the beam-beam interaction. The model uses transfer maps to treat the external focusing elements and a stochastic map to treat radiation damping and quantum excitation of the beams. In the parallel implementation we studied various strategies to deal with the particular nature of the colliding beam system – a system in which there can be significant particle movement between beam-beam collisions. We chose a particle-field decomposition approach instead of the conventional domain decomposition or particle decomposition approach. The particle-field approach leads to good load balance, reduced communication cost, and shows the best scalability on an IBM SP3 among the three parallel implementations we studied. A performance test of the beam-beam model on a Cray T3E, IBM SP3, and a PC cluster is presented. As an application, we studied the effect of long-range collisions on antiproton lifetime in the Fermilab Tevatron.

I. INTRODUCTION

High energy accelerators are essential to study the inner structure of nuclear and elementary particles in modern physics. In a high energy collider, two counter-rotating charged particle beams moving at speeds close to the speed of light collide at one or more interaction points where detectors are located. The rate at which particle production and other events occur inside the detectors depends on a proportionality factor called the *luminosity*. Maximizing luminosity is therefore a key issue in high energy colliders. The electromagnetic interaction between two beams, i.e. the beam-beam interaction, places a strong limit on the luminosity. An accurate simulation of the beam-beam interaction is needed to help optimize the luminosity in high energy accelerators.

The beam-beam interaction has been studied for many years. However, the extreme computational cost required to accurately and self-consistently model the beam-beam interaction as the beams circulate for many (typically 10^4 to 10^6) turns has caused most previous studies to use simplified models. Examples include "weak-strong" models, in which only the "weak" beam is affected by the higher intensity "strong" beam [1–4], soft Gaussian models, where one beam is assumed *apriori* to have a Gaussian shape, or the pancake model, where each beam is assumed to be a single two-dimensional disk perpendicular to the direction of motion of the beam [5–11]. To study the beam-beam interaction fully self-consistently for both beams (i.e. a "strong-strong"

SC'03, November 15-21, 2003, Phoenix, Arizona, USA Copyright 2003 ACM 1-58113-695-1/03/0011... $5.00.

formulation), and to include all the physical processes of long range off-centroid interactions, finite beam bunch length effects, and crossing angle collisions, requires computation resources far beyond the capability of current serial computers. As far as we know, there is no previously developed code that can simultaneously handle all of these physical processes accurately. In this paper we present a parallel beam-beam simulation model, with weak-strong and strong-strong capabilities, that can simulate these physical processes accurately using high performance computers.

The organization of the paper is as follows: The physical model and computational methods are described in Section 2. The parallel implementation is given in Section 3. An application to the study of the effect of long-range collisions on antiproton lifetime in the Fermilab Tevatron is given in Section 4. We summarize our results in Section 5.

II. PHYSICAL MODEL AND COMPUTATIONAL METHODS

In our model of beam dynamics in an accelerator, each charged particle is characterized by its charge, mass, and phase space coordinates $(x, p_x, y, p_y, \Delta z, \Delta p_z/p_0)$. Here, the independent variable, s, is the arc length along a reference trajectory inside the accelerator, $p_{x,y}$ is the transverse momentum normalized by the total momentum of a reference particle ($p_0 = E_0/c$), $\Delta z = s - ct(s)$ with c the speed of light, $\Delta p_z = |p| - p_0$ with p_0 the absolute momentum value of the reference particle. The motion of particles will be determined by several factors, all of which must be included in the model. Externally applied magnetic fields guide the beam and provide transverse and longitudinal focusing. Particles will also lose its energy through synchrotron radiation, a process that involves radiation damping and quantum excitation. The Coulomb interaction among the charged particles within a bunch is negligible due to the cancellation of the electric and magnetic forces at relativistic speeds. However, in the collisions with the oppositely moving beams, the electric and magnetic forces add up. The resulting beam-beam force is a strongly nonlinear interaction that can significantly affect the motion of the charged particles. Fig. 1 gives a schematic plot of two colliding beams with a finite crossing angle. In the figure, α is the collision crossing angle, and IP is the interaction point.

To calculate the electromagnetic force from the beam-beam interaction, we have used a multiple slice model. In this model, each beam bunch is divided into a number of slices along the longitudinal direction in the moving frame as shown in Fig. 1. Each slice contains nearly the same number of particles at different longitudinal locations Δz. The collision point between two opposite slices i and j is determined by

$$s_c = \frac{1}{2}(\Delta z_i^+ - \Delta z_j^-) \tag{1}$$

The transverse coordinates of the particles at the collision point are given by

$$x^c = x + s_c\ p_x \tag{2}$$
$$y^c = y + s_c\ p_y \tag{3}$$

The slopes of the particles are updated using the beam-beam electromagnetic forces at the collision point following

$$p_{x_{new}} = p_x + \Delta p_x \tag{4}$$
$$p_{y_{new}} = p_y + \Delta p_y, \tag{5}$$

where

$$\Delta p_{x_2} = \frac{2q_1q_2N_1}{\gamma_2 4\pi\epsilon_0 m_2 c^2} E_{x_1} \tag{6}$$

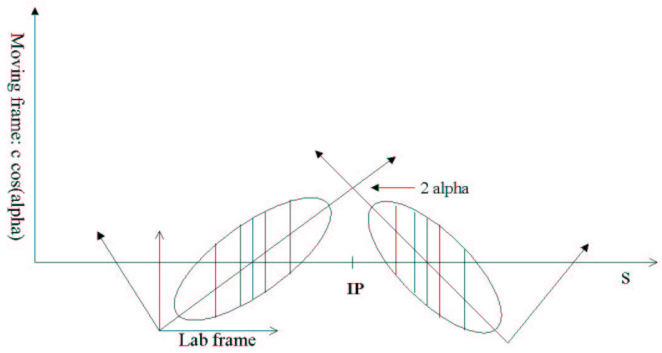

FIG. 1: A schematic plot of the two colliding beams with finite crossing angle.

$$\Delta p_{y_2} \;=\; \frac{2q_1q_2N_1}{\gamma_2 4\pi\epsilon_0 m_2 c^2} E_{y_1}. \tag{7}$$

In the above equations, the subscripts 1 and 2 pertain to each of the two beams, the corresponding equations for the other beam are obtained from the above by the exchange $1 \leftrightarrow 2$, $\gamma = 1/\sqrt{1-\beta^2}$, $\beta_i = v_i/c$, $i = x, y, z$, c is the speed of light, ϵ_0 is the vacuum permittivity, q is the charge of the particle, m is the rest mass of particle, N is the number of particles in a bunch, and E_x and E_y are the transverse electric fields generated by the opposite moving beam. After the collision, the particles of each slice drift back to their original locations according to

$$x \;=\; x^c - s_c\; p_{x_{new}} \tag{8}$$

$$y \;=\; y^c - s_c\; p_{y_{new}} \tag{9}$$

The electric fields generated by the opposite moving beam can be obtained from the solution of Poisson's equation. The solution of Poisson's equation can be written as

$$\phi(x,y) \;=\; \int G(x,\bar{x},y,\bar{y})\rho(\bar{x},\bar{y})\; d\bar{x}d\bar{y} \tag{10}$$

where G is the Green's function, ρ is the charge density, $(\bar{x},\bar{y})$ is the spatial location of the charged particles. For the case of transverse open boundary conditions, the Green's function is given by:

$$G(x,\bar{x},y,\bar{y}) \;=\; -\frac{1}{2}\ln((x-\bar{x})^2+(y-\bar{y})^2) \tag{11}$$

Now consider a simulation of an open system where the computational domain containing the particles has a range of $(0, L_x)$ and $(0, L_y)$, and where each dimension has been discretized using N_x and N_y points. From Eq. 16, the electric potentials on the grid can be approximated as

$$\phi(x_i,y_j) \;=\; h_x h_y \sum_{i'=1}^{N_x}\sum_{j'=1}^{N_y} G(x_i-x_{i'},y_j-y_{j'})\rho(x_{i'},y_{j'}) \tag{12}$$

where $x_i = (i-1)h_x$ and $y_j = (j-1)h_y$. This convolution can be replaced by a cyclic convolution expression in a double-gridded computational domain [12]:

$$\phi_c(x_i, y_j) \;=\; h_x h_y \sum_{i=1}^{2N_x} \sum_{j=1}^{2N_y} G_c(x_i - x_{i'}, y_j - y_{j'}) \rho_c(x_{i'}, y_{j'}) \tag{13}$$

where $i = 1, \cdots, 2N_x$, $j = 1, \cdots, 2N_y$, and

$$\rho_c(x_i, y_j) \;=\; \begin{cases} \rho(x_i, y_j) & : \;\; 1 \le i \le N_x;\; 1 \le j \le N_y \\ 0 & : \;\; N_x < i \le 2N_x \;\; or \;\; N_y < j \le 2N_y \end{cases} \tag{14}$$

$$G_c(x_i, y_j) \;=\; \begin{cases} G(x_i, y_j) & : \;\; 1 \le i \le N_x + 1;\; 1 \le j \le N_y + 1 \\ G(x_{2N_x - i + 2}, y_j) & : \;\; N_x + 1 < i \le 2N_x;\; 1 \le j \le N_y + 1 \\ G(x_i, y_{2N_y - j + 2}) & : \;\; 1 \le i \le N_x + 1;\; N_y + 1 < j \le 2N_y \\ G(x_{2N_x - i + 2}, y_{2N_y - j + 2}) & : \;\; N_x + 1 < i \le 2N_x;\; N_y + 1 < j \le 2N_y \end{cases} \tag{15}$$

$$\rho_c(x_i, y_j) \;=\; \rho_c(x_i + 2(L_x + h_x), y_j + 2(L_y + h_y)) \tag{16}$$

$$G_c(x_i, y_j) \;=\; G_c(x_i + 2(L_x + h_x), y_j + 2(L_y + h_y)) \;. \tag{17}$$

These equations make use of the symmetry of the Green function in Eq. 17. From the above definition, one can show that the cyclic convolution will give the same electric potential as the convolution Eq. 18 within the original domain, i.e.

$$\phi(x_i, y_j) \;=\; \phi_c(x_i, y_j) \qquad for \;\; i = 1, N_x;\; j = 1, N_y \;. \tag{18}$$

The potential outside the original domain is incorrect but is irrelevant to the physical domain. Since now both G_c and ρ_c are periodic functions, the convolution for ϕ_c in Eq. 19 can be computed efficiently using an FFT as described by Hockney and Eastwood. [12].

In the above FFT-based algorithm, the particle domain and the electric field domain are contained in the same computational domain. Here, the particle domain is the configuration space containing the charged particles, and the field domain is the space where the electric field is generated by the charged particles. In the beam-beam interaction, the two opposite moving beams might not overlap with each other. For example, in the long-range interaction, the two colliding beams could be separated by more than several σ, where σ is the rms size of the beam. Thus the field domain where the electric field is generated by one beam can be different from the particle domain containing the beam. Fig. 2 gives a schematic plot of the two separated domains. In this figure, the particle domain has a range from $-R$ to R for x and y, and the field domain has a range from 0 to $2R$ for x and y, where R is maximum extent of the beam. The origin of the field domain in this figure is $x_c = R$, $y_c = R$, where the origin is chosen to be at the beam centroid. In the beam-beam simulation, the origin of the field domain can be at an arbitrary location and varies from turn to turn. To apply Hockney's algorithm directly will require the computational domain to contain both the particle domain and the field domain, i.e. both beams. Since there is a large empty space between two beams, containing both beams in one computational domain will result in a poor spatial resolution of the beams. This is also computationally inefficient because the electric fields in the empty space between two beams are not used.

To avoid this problem, we have defined a *shifted Green function* as

$$G_s(x, \bar{x}, y, \bar{y}) \;=\; -\frac{1}{2} \ln((x_c + x - \bar{x})^2 + (y_c + y - \bar{y})^2) \tag{19}$$

where x_c and y_c are the center coordinates of the field domain. The electric potential in the field domain is written as

$$\phi(x + x_c, y + y_c) \;=\; \int G_s(x, \bar{x}, y, \bar{y}) \rho(\bar{x}, \bar{y}) \; d\bar{x} d\bar{y} \;. \tag{20}$$

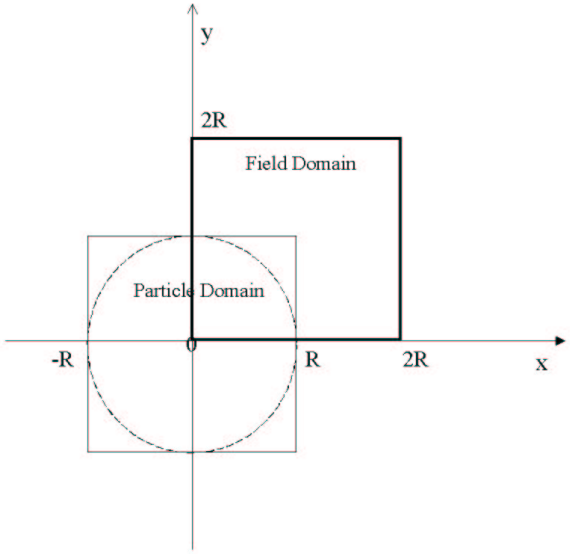

FIG. 2: A schematic plot of the particle domain and the field domain.

Using the shifted Green function, the center of the field domain is shifted to the center of the particle domain. The range of x and y cover both the particle domain and the field domain in one computational domain. The FFT can be used to calculate the cyclic convolution in Eq. 19 using the new Green function. Here, on the doubled grids, the Green function is given as

$$G_c(x_i, y_j) \;=\; -\frac{1}{2}\begin{cases} \ln((x_c + x_i)^2 + (y_c + y_j)^2) & : \;\; 1 \le i \le N_x;\; 1 \le j \le N_y \\ \ln((x_c - x_{2N_x - i + 2})^2 + (y_c + y_j)^2) & : \;\; N_x < i \le 2N_x;\; 1 \le j \le N_y \\ \ln((x_c + x_i)^2 + (y_c - y_{2N_y - j + 2})^2) & : \;\; 1 \le i \le N_x;\; N_y < j \le 2N_y \\ \ln((x_c - x_{2N_x - i + 2})^2 + (y_c - y_{2N_y - j + 2})^2) & : \;\; N_x < i \le 2N_x;\; N_y < j \le 2N_y \; . \end{cases} \tag{21}$$

To summarize, using the shifted Green function:

- avoids the requirement that the particle domain and the field domain be contained in one big computational domain,
- leads to better numerical resolution for the charge densities and the resulting electric fields than the conventional method, because the empty space between the beams is not included in the calculation,
- is far more efficient, in terms of computational effort and storage, than the traditional approach of gridding the entire problem domain.

As an example of the above FFT-based algorithms, we have computed the radial electric field distribution generated by a round beam with a Gaussian density distribution using the particle domain and the field domain shown in Fig. 2. Fig. 3 shows the radial electric field E_r as a function of distance along the diagonal line of the field domain using the the shifted Green function, and 128x128 grid. The electric field from the analytical calculation is also given in the same figure for comparison. It is seen that the agreement between the numerical solution and the analytical calculation is excellent.

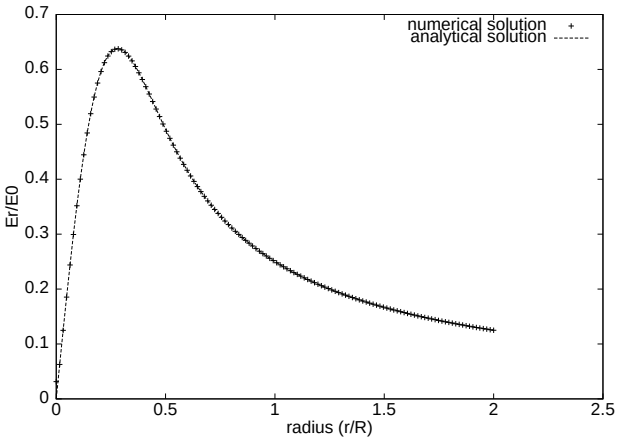

FIG. 3: Radial electric field as a function of distance along the diagonal line of the particle domain.

The effects of external fields can be represented, in the small-amplitude approximation, by a one-turn linear map, i.e.

$$x_{n+1} = (\cos(2\pi\nu_{0x}) + \alpha_x \sin(2\pi\nu_{0x}))\ x_n + \beta_x \sin(2\pi\nu_{0x})\ p_{x_n} \tag{22}$$

$$p_{x_{n+1}} = -\gamma_x \sin(2\pi\nu_{0x})\ x_n + (\cos(2\pi\nu_{0x}) - \alpha_x \sin(2\pi\nu_{0x}))\ p_{x_n} \tag{23}$$

where α_x, β_x and γ_x are lattice functions at the interaction point, and ν_{0x} is horizontal lattice tune. A similar map applies to the vertical phase space y and y' by replacing $x \to y$ in above equations. For the longitudinal phase space, the one-turn map is defined by

$$\begin{pmatrix} \Delta z/\sigma_z \\ \Delta p_z/\sigma_{p_z} \end{pmatrix}_{n+1} = \begin{pmatrix} \cos(2\pi\nu_s) & \sin(2\pi\nu_s) \\ -\sin(2\pi\nu_s) & \cos(2\pi\nu_x) \end{pmatrix} \begin{pmatrix} \Delta z/\sigma_z \\ \Delta p_z/\sigma_{p_z} \end{pmatrix}_n \tag{24}$$

where σ_z and σ_{p_z} are the rms beam sizes in z and p_z, and ν_s is the synchrotron tune.

The effects of radiation damping and quantum excitation can be represented using a localized stochastic map. For each particle, with lattice function $\alpha_x = \alpha_y = 0$, the map consists of the following transformations [5]:

$$x_{n+1} = \lambda_x x_n + r_1 \sigma_x \sqrt{1-\lambda_x^2} \tag{25}$$

$$p_{x_{n+1}} = \lambda_{p_x} p_{x_n} + r_2 \sigma_{p_x} \sqrt{1-\lambda_{p_x}^2} \tag{26}$$

$$y_{n+1} = \lambda_y y_n + r_3 \sigma_y \sqrt{1-\lambda_y^2} \tag{27}$$

$$p_{y_{n+1}} = \lambda_{p_y} p_{y_n} + r_4 \sigma_{p_y} \sqrt{1-\lambda_{p_y}^2} \tag{28}$$

$$\Delta z_{n+1} = \lambda_z \Delta z_n + r_5 \sigma_z \sqrt{1-\lambda_z^2} \tag{29}$$

$$\Delta p_{z_{n+1}} = \lambda_{p_z} \Delta p_{z_n} + r_6 \sigma_{p_z} \sqrt{1-\lambda_{p_z}^2} \tag{30}$$

where the σ's are the nominal rms equilibrium beam sizes in each dimension, the λ's are given in terms of the damping time τ (measured in units of turns) by $\lambda_i = \exp(-1/\tau_i)$ where i denotes x, y, or z, and the r's are independent random numbers satisfying

$$\langle r_i \rangle = 0 \tag{31}$$

$$\langle r_i r_j \rangle = \delta_{ij} \tag{32}$$

The first term in above transformation represents the radiation damping, and the second term represents the quantum excitation.

III. PARALLEL IMPLEMENTATION

Following the above physical model, we have used a particle-in-cell method to calculate the electromagnetic fields at the beam-beam interaction point. Outside the interaction point, the particles are transported through the accelerator using the one-turn lattice map and the radiation damping/quantum excitation map. The fact that the lattice map can cause significant particle movement has important ramifications for the parallelization strategy. During the development of the code we have studied the performance of three methods: domain decomposition, particle decomposition, and particle-field decomposition.

In the domain decomposition approach the spatial domain is divided into a number of subdomains, with each subdomain mapped onto a single processor [13, 14]. The particles with their spatial positions inside the subdomain are assigned to that processor. When particles move out of their spatial subdomain, they are sent to the processor containing the corresponding spatial subdomain. After all particles are local to each processor, the Poisson equation is solved on the grid and the particles are advanced using the electromagnetic fields. To improve the efficiency, a load balance scheme can be used to ensure that each processor contains about the same number of particles. The domain decomposition works well when the particles do not move too far from their positions during each time step. This means that only neighboring processor communication is required. However, in the simulation of colliding beams, after each turn the particles can move a long distance due to the action of external maps. (The colliding beam code is a quasi-static particle-in-cell code, not a fully electromagnetic code, there is no Courant condition. The large movement of particles is not an artifact of the numerical implementation, but instead is a physical effect associated with the fact that particles undergo many oscillations as they are transported around the collider between beam-beam collision points.) A lot of communication is required to move these particles to their local processors. Meanwhile, even though the domain decomposition approach can achieve a load balance of particles, the solution of the Poisson equation is not balanced since each processor has a different number of computational grids, i.e. a different size of subdomain.

Perfect load balance can be achieved, and particle movement avoided, by using a particle decomposition approach [15]. In this approach, the particles are uniformly distributed among processors. Each processor contains the whole spatial domain. To solve the Poisson equation, the particles are deposited onto the global computational grid, collected and broadcast to all processors. Each processor now owns the charge density distribution of the whole domain, and the Poisson equation is solved within this domain. Unfortunately this implementation does not take advantage of the parallelism in the solution of the Poisson equation. To overcome this drawback, in this paper we have proposed a particle-field decomposition approach as the strategy that is best suited to the parallel implementation of the particle-in-cell method for modeling colliding beams.

In the particle-field decomposition approach, each processor possesses the same number of particles and the same number of computational grid points, i.e, the same size of spatial subdomain. Fig. 4 shows a schematic plot of the particle-field decomposition among three processors. We see that the global computational mesh has been uniformly distributed among three processors. Each processor also has the same number of particles. The spatial coordinates of the particles on each processor may not stay within the spatial mesh domain of that processor. In the process of solving the Poisson equation, the particles are deposited onto the computational grid to obtain the charge density distribution. For the particles with spatial positions outside the local subdomain, an auxiliary computational grid is used to store the charge density. After the deposition, the charge density stored on the auxiliary grid will be sent to the processor containing that subdomain. With charge density local to each processor, the Poisson equation is solved in parallel on a local subdomain using the shifted Green function method. Since each processor contains the same number of com-

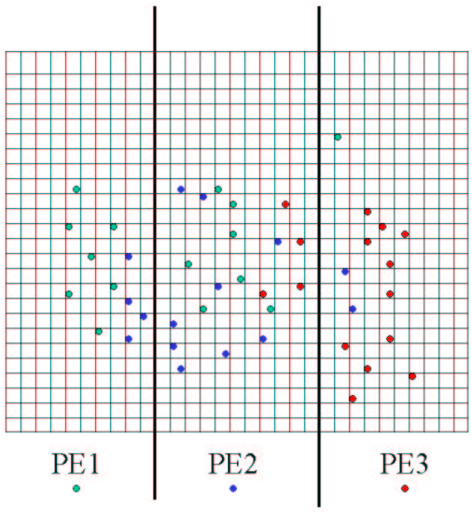

FIG. 4: A schematic plot of the particle-field decomposition among three processors.

putational grid points, the work load is well balanced among all processors. The solution of the electric potential on the local subdomain is sent to all processors. With the electric potential on each processor, the electric field is calculated on the grid and interpolated onto individual particles of the opposite beam. The particles are advanced using the electromagnetic field and the external maps. Since each processor contains the same number of particles, the work of this process is also well balanced among processors. The volume of communication in the particle-field decomposition approach is proportional to the number of computational grid points instead of the number of moving particles in the domain decomposition approach. Since, in the study of beam-beam interactions, the number of particles is much larger than the number of computational grid points, e.g. 10^6 versus 10^4, the particle-field decomposition approach can significantly reduce the communication cost in the simulation. Fig. 5 shows a comparison of the speedup as a function of number of processors on an IBM SP3 computer using above three parallel implementation approaches for a single slice beam-beam model with one million particles and 128×128 grid points. We see that the particle-field decomposition method has the best scalability among three implementations. The speedup of the domain-decomposition approach saturates at 16 processors due to the large amount of time spent moving the particles among the processors. The particle decomposition approach does not take advantage of the parallelism in the field calculation and has a poorer performance than the particle-field decomposition. For example, on 32 processors, the domain-decomposition approach spends 86 seconds within the particle manager function moving the particles around the processors, while the total beam-beam simulation time is 114 seconds. In the calculation of electromagnetic fields, it takes about 16 seconds on 32 processors using the particle decomposition approach. It takes about 6 seconds using the particle-field decomposition approach. Even though the scalability of the field calculation in the particle-field decomposition approach is not perfect due to the significant communication cost associated with data collecting and transpose, it does show better performance than the particle decomposition approach, where the field calculation is done in serial.

Having adopted the particle-field decomposition approach, we next divide the total number of

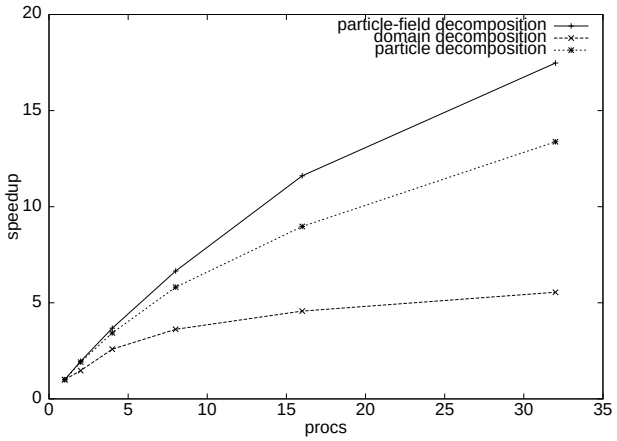

FIG. 5: Speedup as a function of number of processors on IBM SP3 with particle-field decomposition, particle decomposition and domain decomposition.

processors into two groups, with each group responsible for one beam, and each processor in a group containing the same number of particles. We furthermore divide each beam longitudinally into a specified number of slices. The processors in each group are arranged logically into a two-dimensional array with each column of the array containing a number of slices which are assigned to this column of processors cyclically along the row direction. This gives a good load balance of slices among different column processors. Within each column, the computational grid associated with each slice is decomposed uniformly among all the column processors. This provides the parallelization in the solution of the Poisson equation.

As a test of the parallel performance, we have measured the speedup as a function of number of processors on a Cray T3E, IBM SP, and a PC cluster at Lawrence Berkeley National Laboratory. The results are given in Fig. 6. Here, we have used five slices for each beam with two million particles and a computational grid of size 64×64. We see that the program scales up to 128 processors with an efficiency of $65\% - 80\%$ on all three machines. The speedup on the less expensive PC cluster appears to be better than that on the high performance, more expensive parallel machines, which is against general belief. To clarify this issue, we also show in Fig. 7 the total computing time as a function of number of processors on all three machines. We see that the dedicated parallel computers IBM SP3 and Cray T3E are still a factor of 2 to 3 faster than the PC cluster. A check of the time spent on the communication and computation indicates that in this example, the relative fraction of the communication time in the total computing time on the PC cluster is less than that on the other two parallel computers. For example, on 128 processors, the communication time is 3.4 seconds out of 19.0 seconds of total computing time on the PC cluster and 2.9 seconds out 10.5 seconds on the Cray T3E. While the communication time is slightly less on the T3E than that on the PC cluster, the computation speed on the T3E is about a factor of 2 faster than the PC cluster. (Here, the single processor double precision peak performance on the 450 MHz T3E is 900 Mflops while the 866 MHz Pentium III on the PC cluster takes a few cycles to do one double precision operation.) This results in a relatively larger fraction of communication cost on the T3E than that on the PC cluster, which causes the lesser scalability of the T3E for this test example. The speedup on IBM SP3 shows initial superlinear behavior below 64 processors. This superlinear behavior might be due to the finite cache size effects on a small number of processors. The IBM SP3 processor has about 200 KB L1 cache and about 8 MB L2 cache. This corresponds to about 166K particles in data storage, where each particle occupies 48 bytes. In a case using small number of processors, each processor will hold more particle and field data than the case using a larger

number of processors. This leads to more memory access time in the case involving a small number of processors. For example, on 4 processors, for 2 million particles, each processor holds $500,000$ particles which cannot be stored in the L2 cache. On 32 processors, each processor holds $62,500$ particles which can fit in the L2 cache. However, with continuing increase of processor number, the communciation time among the processors increases and the speedup falls below linear scalability eventually. For example, on 32 processors of SP3, the communication time is 2.8 seconds out of the 19.6 seconds of total computing time. On 128 processors, the communication time is 3.9 seconds out of 8.1 seconds of total computing time. The increased communication cost on large number of processors is due to the fact that each node of IBM SP3 has 16 processors but with only 2 switches to communicate among the nodes.

IV. APPLICATIONS

Our parallel beam-beam code has been applied to study several accelerators including the Tevatron at Fermi National Accelerator Laboratory, the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory, and the Large Hadron Collider (LHC) under construction at CERN. Most recently, we have performed extensive simulations to understand the effect of long-range collisions on the antiproton lifetime, and hence the integrated luminosity, in the Tevatron, as a function of several operating parameters. We have shown in Fig. 8, the antiproton lifetime as a function of proton intensity (number of proton per bunch) at injection energy. This is in reasonable agreement with experimental observation [17].

In regard to the LHC, our code has been used to study the efficacy of a beam sweeping procedure that will be used to monitor luminosity when the LHC is in operation [16]. We have furthermore recently used our code to perform the first-ever strong-strong beam-beam simulations involving one million particles propagating for one million turns.

V. SUMMARY

In this paper, we have presented a parallel simulation model to study beam-beam effects in high energy colliders. The electromagnetic fields between the two colliding beams are calculated using a parallel particle-in-cell approach with a new developed shifted-Green function algorithm. In the parallel implementation, a particle-field decomposition approach has been proposed. This approach shows better scalability than the domain decomposition approach and the particle decomposition approach due to the nature of the particle movement in the beam-beam code. A preliminary performance test of the multi-slice beam-beam model shows reasonable scalability up to 128 processors on a Cray T3E, IBM SP3 and a PC cluster. As an application, we have studied the effect of long-range collisions on antiproton lifetime in the Fermilab Tevatron. The computational results and the experimental results are in good agreement. In conclusion, the development of a parallel code that simultaneously treat the effects of multiple beam-beam phenomena (head-on collisions, long-range collisions, crossing-angle effects, finite bunch length effects, *etc.*), represents a new and powerful capability that will be useful for understanding and improving the operational characteristics of present and future colliders.

ACKNOWLEDGMENTS

We would like to thank Drs. V. Decyk and R. Gerber for helpful discussions about the parallel implementations and parallel performance. We also thank Drs. T. Sen and M. Xiao for the

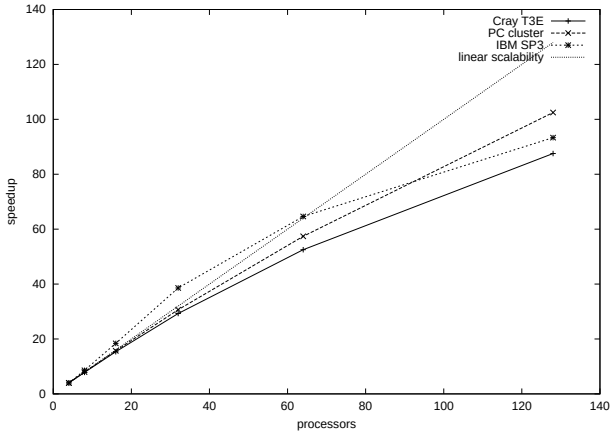

FIG. 6: Speedup as a function of number of processors on Cray T3E, IBM SP3, and PC PIII cluster.

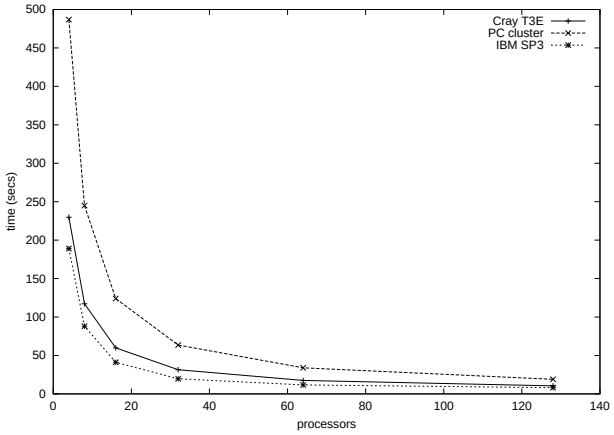

FIG. 7: Total computing time as a function of number of processors on Cray T3E, IBM SP3, and PC PIII cluster.

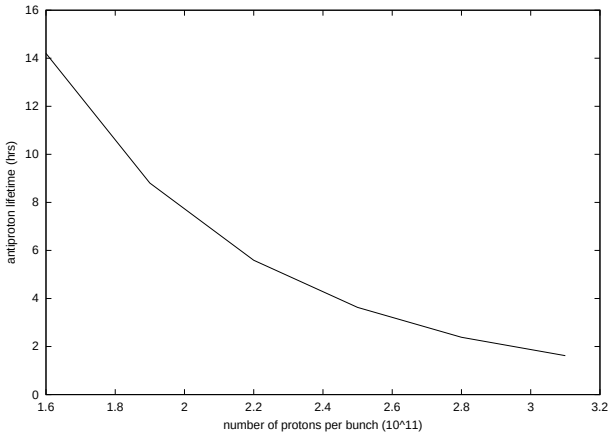

FIG. 8: Antiproton lifetime as a function of proton intensity (number of proton per bunch) at 150 GeV Tevatron.

collaboration on the beam-beam studies in the Tevatron. 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 (US DOE/SC) under Contract No. DE-AC03-76SF00098, and the resources of the Center for Computational Sciences at Oak Ridge National Laboratory. Some of the computational work for this project was done on the LBL/NERSC Alvarez Cluster, an 80 node Pentium III Myrinet cluster. This work was performed under the auspices of a Scientific Discovery through Advanced Computing project, "Advanced Computing for 21st Century Accelerator Science and Technology," which is supported by the US DOE/SC Office of High Energy and Nuclear Physics and the Office of Advanced Scientific Computing Research.

REFERENCES

[1] K. Hirata, H. Moshammer and F. Ruggiero, "A Symplectic Beam-Beam Interaction with Energy Change," Particle Accelerators, vol. 40, 1993, pp. 205-228.

[2] K. Hirata, "Analysis of Beam-Beam Interactions with a Large Crossing Angle," Phys. Rev. Lett., vol. 74, no. 12, March 1995, pp. 2228-2231.

[3] Y. Papaphilippou and F. Zimmermann, "Weak-Strong Beam-Beam Simulations for the Large Hadron Collider," Phys. Rev. Special Topics - Accel. Beams, vol. 2, Oct. 1999, 104001.

[4] L. H. A. Leunissen, F. Schmidt, G. Ripken, "Six-Dimensional Beam-Beam Kick Including Coupled Motion," Phys. Rev. Special Topics - Accel. Beams, vol. 3, Dec. 2000, 124002.

[5] M. A. Furman, A. Zholents, T. Chen, and D. Shatilov, "Comparisons of Beam-Beam Code Simulations," CBP Tech Note-59, 1996.

[6] S. Krishnagopal and M. A. Furman and W. C. Turner, "Studies of the Beam-Beam Interaction for the LHC," LBNL-43061, CBP Note 308, 1999.

[7] T. Koyama, "Single Kick Approximations for Beam-Beam Deflections," Phys. Rev. Special Topics - Accel. Beams, vol. 2, Feb. 1999, 024001.

[8] M. A. Furman, "Beam-Beam Simulations with the Gaussian Code TRS", LBNL-42669, CBP Note 272, 1999.

[9] M. P. Zorzano and F. Zimmermann, "Simulations of Coherent Beam-Beam Modes at the Large hadron Collider," Phys. Rev. Special Topics - Accel. Beams, vol. 3, April 2000, 044401.

[10] M. A. Furman and W. C. Turner, "Beam-Beam Simulations for Separated Beams in the LHC", LBNL-46223, CBP Note 350 (2000).

[11] W. Herr, M. P. Zorzano, and F. Jones, "Hybrid Fast Multipole Method Applied to Beam-Beam Collisions in the Strong-Strong Regime," Phys. Rev. Special Topics - Accel. Beams, vol. 4, May 2001, 054402.

[12] R. W. Hockney and J. E. Eastwood, *Computer Simulation Using Particles*, McGraw-Hill Book Company, New York, 1985.

[13] P. C. Liewer and V. K. Decyk, "A General Concurrent Algorithm for Plasma Particle-in-Cell Codes," J. Comput. Phys., vol. 85, 1989, 302.

[14] J. Qiang, R. D. Ryne, S. Habib, V. Decyk, "An Object-Oriented Parallel Particle-in-Cell Code for Beam Dynamics Simulation in Linear Accelerators," J. Comput. Phys., vol. 163, 2001, pp. 434-451.

[15] J. M. Dawson, V. K. Decyk, R. D. Sydora, and P. C. Liewer, "High Performance Computing and Plasma Physics," Physics Today, vol. 46, no. 3, 1993, 64.

[16] J. Qiang, M. A. Furman, and R. D. Ryne, "Strong-Strong Beam-Beam Simulation Using a Green Function Approach," Phys. Rev. Special Topics - Accel. Beams, vol. 5, Oct. 2002, 104402.

[17] T. Sen, "New Aspects of Beam-Beam Interactions in Hadron Colliders," in 2003 Particle Accelerator Conference proceedings, May 12-16, Portland, Oregon, 2003 (to be published).