ODYSSEUS: A DYNAMIC STRONG-STRONG BEAM-BEAM SIMULATION FOR STORAGE RINGS

E. B. Anderson, T. I. Banks, and J. T. Rogers,
Laboratory of Nuclear Studies, Cornell University, Ithaca, NY 14853

Abstract

We have developed a simulation of the beam-beam interaction in e+/e- storage ring colliders which is specifically intended to reveal the dynamic collective behavior of the colliding beams. This program is a true 6-dimensional strong-strong simulation in which the electromagnetic fields of longitudinal slices of the colliding beams are recalculated for each slice collision. Broadband wake fields are included and no constraints are placed on the distribution of particles in the beams. Information on tests of the code will be shown. Results will be presented including limiting beam-beam parameters for round and flat beams, deviations from the Gaussian distribution, effects of the beam-beam parameter on head-tail instability thresholds, and Landau damping rates. Possibilities for further improvements will be discussed.

1 INTRODUCTION

Our development of a beam-beam simulation program was motived, in part, by the observation in the Cornell Electron Storage Ring (CESR) of an m = -1 head-tail instability which occurs at lower beam current when beams are in collision. This change seems to indicate that the instability is interacting with the beam-beam effect. The head-tail instability involves particles moving forwards and backwards within the bunch, so any model for this interaction would have to describe the longitudinal dynamics of the bunch. In addition, collective effects arising from the beam-beam force alone can limit luminosity. The flip-flop instability is commonly observed in e^+e^- storage ring colliders. The DCI storage ring at LAL, Orsay, France, had four colliding beams in which the e^+ beam charge was compensated by the e^- charge, but the beam-beam limit was not significantly different from that for uncompensated beams. The beam-beam limit in DCI was attributed to a collective beam-beam instability [1]. This suggests that the beam-beam limit for two-beam collisions may also be due, in some cases, to a collective instability.

This paper presents a new beam-beam simulation program, ODYSSEUS (Optimized DYnamic Strong-Strong E-plus e-minUs Simulation). To the author's knowledge, ODYSSEUS is the first six-dimensional strong-strong beam-beam simulation in which no constraints are placed on the beams and is the first to include wake fields. These features make it possible to investigate any mode of oscillation of the colliding beams. ODYSSEUS is designed to serve as a flexible, efficient, and portable tool for investigating beam-beam effects.

2 BASIC IDEAS

ODYSSEUS uses macroparticles to model the six-dimensional motion of the particles in the beams. Typically, the number of macroparticles used is on the order of ten thousand in each beam. With this number of macroparticles the speed of the calculation is limited by the electromagnetic field calculations. ODYSSEUS adaptively chooses from a variety of different field computation methods. Different algorithms are used for the core and transverse tails of the beam and for longitudinal slices with large or with small charge. The parameters of the program can be changed to model flat or round beams. Further, inclusion of the longitudinal degrees of freedom and wake fields allows the investigation of previously inaccessible physics.

2.1 Particle Tracking

On each simulated turn through the storage ring, each macroparticle is propagated from the collision point and back again through the linear optics of the storage ring, including chromaticity, synchrotron radiation excitation and damping, RF phase focusing, and wake field deflections.

The magnetic optics of the ring are approximated with linear transport theory as described in many sources including a popular article by M. Sands [2]. The fields in the RF cavities are approximated as sinusoids, while the change in position is handled using a momentum compaction. Macroparticles which have migrated past a transverse aperture are no longer considered in the simulation. As they pass the aperture, the positions and velocities of these particles are recorded for later analysis. ODYSSEUS handles the longitudinal variation of the electromagnetic field of the beam by dividing the beam into slices that were typically of equal thickness. Because the particles are moving at ultra-relativistic speeds, it is approximated that the fields are entirely transverse.

Individual macroparticles undergoing longitudinal oscillations may migrate from slice to slice, so on each turn the macroparticles are sorted according to their longitudinal position and reassigned to slices. This is necessary for the calculation of both the electromagnetic wake field and the actual beam-beam force. The longitudinal motion in most accelerators is slow, so during collisions each macroparticle is assumed to remain within its slice. Because the motion of a macroparticle from slice to slice is slow, Heapsort is an effective sorting algorithm [3].

2.2 Radiation Effects

There are two major ways that random perturbations due to synchrotron radiation are handled in beam-beam simulations. The first is to use the description of the storage ring to find the cumulative effects of the radiative kicks and apply these to the beam. The second is to pick a radiative kick that gives the known beam size or emittance. In the horizontal direction the design of the magnetic optics closely determines the magnitude of the radiative perturbations. In the vertical direction errors in the optics and small deviations from a horizontal orbit dominate the radiative perturbations. For these reasons, ODYSSEUS uses information derived from the synchrotron radiation integrals to determine the size of the radiative kicks in the horizontal case and phenomenological perturbations are added in the vertical direction to agree with the actual beam size.

2.3 Wake Fields

Longitudinal and transverse, single bunch, short-range wake fields are included in the simulation. One of the program inputs is a list of longitudinal and transverse resonators with values for the resonant frequencies, shunt impedances, and quality factors of the resonators. The wake functions are therefore the sum of exponentially damped sinusoids. The wake fields are calculated by summing up the effect of each of the effective resonators.

2.4 Collisions

During its passage though the opposing bunch, the transverse position of each macroparticle may change appreciably because the vertical interaction point beta function, β_V^*, in CESR and many other colliders is comparable to the bunch length, σ_z. The simulation collides each pair of slices sequentially, updating the transverse momenta and positions of each macroparticle after each pairwise collision of slices. For each slice collision $\langle x \rangle$, $\langle y \rangle$, $\langle x^2 \rangle$, $\langle y^2 \rangle$, and the total charge for the slice distribution are found, and the slice electromagnetic field is calculated.

3 FIELD CALCULATION

For purposes of calculating the electromagnetic field from the beam, each beam is divided into longitudinal slices. Although the number of slices can be set arbitrarily, about fifteen are typically used. The field from each slice, integrated over the length of the slice, is calculated independently. The beams are assumed to be ultra-relativistic, so the field due to each slice is transverse and affects only the particles within the region of that slice.

The calculation of the electromagnetic field of each beam is adaptive in order to maximize the speed of the program. One method uses moments of the beam to calculate an approximate electromagnetic field, while others make calculations of the fields on a rectangular grid. Different methods are used depending on whether the field is calculated for the region of the beam core or for the beam tails, whether the number of macroparticles within a slice, N, is large or small, and whether the number of grid points, N_g, used in the field calculation is large or small.

3.1 Beam core

3.1.1 Small N

If the number of macroparticles, N, within a slice is very small, the integrated field at a probe beam macroparticle is calculated from the exact radius vector from each opposing source beam macroparticle. The field must be calculated at the position of each macroparticle in the probe beam, so the number of calculations goes as N^2, making this method efficient only for very small N. In practice this method is only used when N is less than fifty.

3.1.2 Large N, Small N_g

For larger values of N, the electromagnetic field is calculated on a rectangular grid using pre-calculated Green's functions for charges on the grid points. Since the Green's function describes the effect of a single unit charge particle, the Green's function on the grid is found by calculating the field at each grid point due to a unit charge at the origin:

$$G(\vec{r}) = \int_{+\infty}^{-\infty} \vec{E}(s, r) ds = \frac{1}{2\pi\epsilon_0} \frac{\hat{r}}{r} \qquad (1)$$

where $\vec{r}$ is the position vector for the grid point. Notice that the Green's function describes the integrated field strength and has been doubled to take into account the effects of the magnetic field. The macroparticle charge is assigned to the grid points using one of two area-weighted techniques (both techniques described are in Section 5). For small values of the number of grid points, the convolution of the charge density and Green's function is done as a summation in real space. The number of calculations required for this convolution goes as N_g^2. The portion of the code whose speed is dependent on the number of macroparticles is now only linear in N. This technique is used only when the number of grid points is quite small, generally under two hundred.

3.1.3 Large N, Large N_g

For larger values of N_g, the convolution of the Green's functions and charge density is done as a simple multiplication in wavenumber space. The speed of this method is limited by the speeds of the necessary Fourier transform to wavenumber space and the inverse transform back to real space. The number of calculations goes as $N_g \log_2 N_g$. To suppress edge effect problems in the Fourier transforms, the size of the wavenumber space is doubled in both directions and padded with zeros [4].

4 BEAM TAILS

The tails of the beam, typically taken to be particles with a displacement of more than $(10/3)\sigma$ in the horizontal, vertical, or longitudinal directions, are treated differently than

the core particles. The tail particles have very little effect on the beam-beam force. They do, however, respond to the beam-beam force and must be tracked to determine the beam lifetime. Performing a strong-strong calculation for the beam tails with the grid method is computationally inefficient and unnecessary, so a weak-strong calculation is used.

4.1 Longitudinal Tails

Longitudinal tail particles are subject to forces from the core of the opposing beam. This is a weak-strong calculation. A full calculation of the field from the opposing beam slice is performed, as described above for the beam core. The tails are assumed to have no effect on the other beam. It should be noted that the user chooses the number of slices that will be treated in a weak-strong manner and that all slices can be treated as strong-strong slices if the user chooses to do so.

4.2 Transverse Tails

The transverse tail particles are subject to a beam-beam force of similar magnitude to that experienced by the core particles. The fine structure of the charge distribution of the core has little influence on the field in the transverse tails, so the field there is calculated from a two-dimensional Gaussian charge distribution with the same $\langle x \rangle$, $\langle y \rangle$, $\langle x^2 \rangle$, $\langle y^2 \rangle$, and total charge as the charge distribution of the slice. The field from this Gaussian charge distribution is calculated from the rational approximation of Talman and Okamoto [5] for the complex error function solution of Bassetti and Erskine [6].

5 INTERPOLATION TECHNIQUES

Whenever a grid-based technique is used, it is necessary to interpolate. The charges of the macroparticles must be distributed on a grid for the field calculations, and the fields calculated on the grid must be applied to particles at arbitrary locations. In order to conserve momentum, the same interpolation scheme must be used in these two situations[4].

The lowest order interpolation scheme used in ODYSSEUS is the Cloud-In-Cell (CIC) technique. In this scheme the macroparticle is treated as a uniform cloud the size of a grid rectangle. The portion of this cloud closest to a grid point is assigned to that point. Since the interpolation in ODYSSEUS is done on a two-dimensional grid, this involves the four nearest grid points.

The second-order techniques that are most useful for this type of calculation are the symmetrical five- and nine-point interpolation schemes. A nine-point interpolation scheme has been coded as an option in ODYSSEUS and is typically used instead of the CIC. The nine-point scheme that is used is a natural extension of a one-dimensional technique called Triangular-Shaped Cloud (TSC). In the one-dimensional case, TSC represents a macroparticle with a triangular cloud two grid spaces wide. The fraction of the area of the cloud that is closest to each of the grid points is assigned to that point. ODYSSEUS uses the two-dimensional extension of the technique. A set of three fractions is found for each dimension, and these are multiplied to find the weights at all nine nearby points.

When higher order interpolation schemes are used, the charge of the macroparticle can be spread out so that it conceals some of the structure of the charge distribution. This is corrected by using a "sharpening function". In ODYSSEUS sharpening is done during the convolution of the charge density and Green's function in Fourier space. To determine the sharpening function on the grid, a unit charge is placed exactly on a grid point. The interpolation scheme is then used, and some fraction of the charge will be deposited on grid points other than the one where the unit charge is actually located. The Fourier transform of this grid is then found. If G represents the Green's function, $\rho^\ddagger$ represents the spread-out charge distribution, S represents the spreading function, F represents the appropriate electromagnetic potential or field, and $F^\ddagger$ the spread-out potential or field, then:

$$S(\vec{r}) \star F(\vec{r}) = F^\ddagger(\vec{r}), \tag{2}$$

$$F^\ddagger(\vec{r}) = G(\vec{r}) \star \rho^\ddagger(\vec{r}), \tag{3}$$

and

$$S(\vec{r}) \star F(\vec{r}) = G(\vec{r}) \star \rho^\ddagger(\vec{r}). \tag{4}$$

Then in Fourier space it can simply be written that:

$$\tilde{F}(\vec{k}) = \frac{\tilde{G}(\vec{k})\tilde{\rho}^\ddagger(\vec{k})}{\tilde{S}(\vec{k})}. \tag{5}$$

Because the fields are actually spread out twice, once when interpolating to the grid and once when interpolating back to the macroparticles, the actual expression used is:

$$\tilde{F}(\vec{k}) = \frac{\tilde{G}(\vec{k})\tilde{\rho}^\ddagger(\vec{k})}{\tilde{S}(\vec{k})^2}. \tag{6}$$

6 PRE- AND POST-PROCESSING

The preprocessing program Penelope was written to maintain consistency between the input variables and provide a easy to use, portable interface. ODYSSEUS was designed to investigate the coherent oscillations of the bunch, so post-processing to analyze the bunch spectrum is necessary. Post-processing and spectral analysis is done in a *Mathematica* [7] notebook.

7 TESTING

7.1 Field Errors

A small program was written that generates the fields by all the methods used in ODYSSEUS and compares them. There is no difference between the real and Fourier space PIC calculations.

There are significant differences between the time required for field calculations for flat and round beams. With

round beams a field calculation grid can be constructed out of nearly square cells with equal numbers of cells in each dimension. For example, a 32 by 32 grid has only 1024 grid cells, which allows it to run in a reasonable amount of time. In contrast, the often extreme aspect ratios of flat beams force the use of either large numbers of cells or individual cells with poor aspect ratios. For instance, a minimal 8 by 512 grid, appropriate for CESR, has 4096 cells and takes about five times as long as the round beam calculation above.

When trying to accommodate flat beams with a reasonable number of grid cells, one option is to give the individual grid cells aspect ratios other than unity, but there are dangers in this method. Unless the grid cells are perfectly square, there is no choice for the value of the potential at the origin which provides symmetry between the x and y components of the gradient of the potential. Using an electric field-based calculation and a separate Green's function for each component removes this problem, but many interpolation techniques will break down as the aspect ratio of the cells increases. ODYSSEUS uses a distinct Green's function for each component of the electric field. The maximum aspect ratio usually allowed was 1.4, but this could be relaxed to significantly larger values for less demanding calculations.

7.2 The Number of Slices

The experience of previous investigators [8] had indicated that low numbers of slices, five or less, are necessary. This estimate was not reasonable for ODYSSEUS for two reasons. One reason is that most previous simulations included a significant natural vertical emittance that made many slicing errors insignificant. The second is that the low β_y^* in CESR is comparable to the bunch length, making the hourglass-like effect at the interaction point more important. With round beams the β_y^* tends to be higher, which decreases the hourglass effect and thus the number of slices.

In the case of flat beams, simulations of beams with low, non-colliding, vertical emittances need very large numbers of slices. As discussed below, uniform slicing is partially to blame. Other methods may exist which require fewer slices, which is important since the speed of ODYSSEUS scales as the square of the number of strong slices in each beam.

An analytic estimate was made of the maximum possible tune shift that could occur *solely* from the uniform slicing method. In a particle-tracking simulation using longitudinal slices of uniform width, particles at the front and back of each slice receive a different deflection than would actual beam particles. The deflection error, to first order in the derivatives of $\rho(z)$, is

$$\delta y'_m \approx -\Delta y'_m \frac{1}{\rho}\frac{d\rho}{dz}\delta z \tag{7}$$

where $\Delta y'_m$ is the deflection due to slice m in the absence of the error, and $\delta_z = z - (z_m + z_{m+1})/2$. Slice m, where $m = -(M-1)/2, \ldots, -1, 0, 1, \ldots, (M-1)/2$, spans the interval (z_m, z_{m+1}). For a Gaussian $\rho(z)$,

$$\delta y'_m = \sqrt{2\pi} w(mw) e^{-m^2 w^2/2} \frac{\xi_y}{\beta_y^*}\left(y_0 + y'\frac{\sigma_z}{2}mw\right)\frac{\delta z}{\sigma_z} \tag{8}$$

where $w = \Delta z/\sigma_z$ is the slice length Δz in units of σ_z. Summing over all M slices,

$$\delta y'_{tot} = \sum_{m=-\left(\frac{M-1}{2}\right)}^{\left(\frac{M-1}{2}\right)} \delta y'_m \tag{9}$$

$$\approx \frac{\sqrt{2\pi}\xi_y y' \delta z}{2\beta_y^*}\int (mw)^2 e^{-(mw)^2/2} d\,(mw) = \frac{\pi \xi_y y' \delta z}{\beta_y^*} \tag{10}$$

Because δz is approximately uniformly distributed on the interval (z_m, z_{m+1}), $\left\langle (\delta z)^2 \right\rangle = (\Delta z)^2/12$, and

$$\left\langle \delta y'^2_{tot} \right\rangle = \left(\frac{\pi\xi_y}{\beta_y^*}\right)^2 \left\langle (\delta z)^2 \right\rangle \left\langle y'^2 \right\rangle = \left(\frac{\pi\xi_y}{\beta_y^*}\right)^2 \frac{(\Delta z)^2}{12}\sigma^2_{y'} \tag{11}$$

Because the δz of a particle changes in a non-periodic way, we will consider it to be random from turn to turn. When the beam is in equilibrium between random excitation and damping,

$$\left\langle \delta y'^2_{tot} \right\rangle = 4\delta\sigma'^2_y \tag{12}$$

where δ is the vertical damping decrement. From equations (11) and (12) we find that

$$\xi_{y,max} = \frac{4\sqrt{3}}{\pi}\frac{\beta_y^*}{\Delta z}\sqrt{\delta} \tag{13}$$

This is the maximum value of ξ_y which will be produced by the simulation due to the finite slice length. Other physical or numerical effects may further reduce ξ_y.

A series of calculations were made in the Gaussian approximation to investigate the effect of slicing. If increasing the number of slices used does not affect the beam's size, tune-shift, and luminosity, then the original number of slices is probably sufficient. At forty-five slices and more the effect of additional slices drops off rapidly. The noise introduced by slicing is typically unimportant compared to the natural vertical emittance for forty-five slices. For some applications of ODYSSEUS, forty-five slices are prohibitive, and a better slicing technique will need to be implemented.

7.3 Head-Tail Modes

Simulated head-tail damping modes for non-colliding beams were in accordance with expectations. With beams in collision, the $m = +1$ and $m = -1$ modes were both shifted upwards in frequency by one-half of the σ-mode / π-mode tune shift. This is in accordance with the predictions of Cornelis and Lamont [9].

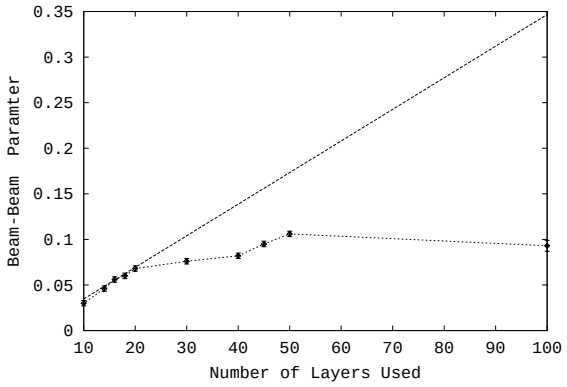

Figure 1: The figure above shows the vertical beam-beam tune shift, $\Delta\nu_y$, as a function of the number of slices for a series of runs. The straight line is an estimate of the maximum beam-beam parameter, ξ_y, from slicing errors alone.

7.4 Speed

ODYSSEUS is currently being run on a 500 MHz α-chip personal computer running LINUX. On this machine the Gaussian approximation runs in this paper were done in one to four hours, and the PIC methods took from less than a day to five days. Runs lasted for a few radiation damping times, from twenty to eighty thousand turns.

8 THE BEAM-BEAM RATIO

A number of predictions have been made for the beam-beam ratio $(\Delta\nu/\xi)$ [10, 11, 12, 13, 14, 15]. An interesting initial test for the program was to compare the results of ODYSSEUS with these previous results. The beam-beam ratio was calculated using ODYSSEUS. Flat beams at 8 mA under CESR conditions were used to measure $\Delta\nu/\xi$. This ratio was done for beams with a longitudinal extent and also for pancake-like beams. For three-dimensional beams the vertical beam-beam ratio was found to be 1.39, while the ratio was 1.18 for two-dimensional beams. The horizontal beam-beam ratio was 1.0 in both cases.

9 CONCLUSIONS

There are a number of important computational advances represented in ODYSSEUS. The most important of these is its adaptive nature. The code dynamically chooses to sum the forces over each particle individually, use a Gaussian charge density approximation, or use a PIC method in real or Fourier space. The grid used in the PIC calculations is pre-generated, then readjusted dynamically as the beam changes. ODYSSEUS also handles wake fields, something that has not been included in similar simulations. Approximations are used in the transverse and longitudinal tails in order to save time on calculations. Runs can be done in reasonable amounts of time, ranging from an hour to a few days depending on the approximations used and beam shape. ODYSSEUS is a significant advance in the simulation of beam-beam effects, and there are possibilities to improve it further.

There were some surprises in the writing and benchmarking of ODYSSEUS. One surprise was the importance of the individual cells' aspect ratios when deciding what sort of a grid to use for PIC calculations. Another was the importance of slicing algorithms and, consequently, the number of slices required. Improvements in the slicing portion of the code have the possibility of making ODYSSEUS significantly faster.

10 ACKNOWLEDGMENTS

This work was supported by the National Science Foundation. The authors extend their thanks to Richard Talman, John Irwin, Bob Siemann, Miguel Furman, Srinivas Krishnagopal, Rob Ryne, Ray Ng, Ray Helmke, Jerry Codner, and the other participants in the Strong-Strong Beam-Beam Workshop at SLAC in the summer of 1998.

11 REFERENCES

[1] J. LeDuff, M. Level, P. Marin, E. Sommer, and H. Zyngier, *Proceedings of the 11th International Conference on High Energy Accelerators*, 707 (1980).

[2] M. Sands, Stanford Linear Accelerator Center Report, SLAC-121, UC-28, ACC (1970).

[3] W.H. Press, *et al.*, *Numerical Recipes in FORTRAN: The Art of Scientific Computing* (Cambridge University Press, New York, 1992)

[4] R.W. Hockney and J.W. Eastwood, *Computer Simulation Using Particles* (Hilger, Bristol, U.K., 1988).

[5] Y. Okamoto and R. Talman, Cornell Electron Storage Ring Report, CBN 80-13 (1980).

[6] M. Bassetti and G.A. Erskine, European Laboratory for Particle Physics Report, CERN-ISR-TH/80-06 (1980).

[7] *Mathematica*, Wolfram Research, Inc., Champaign, Illinois.

[8] M. Furman, A. Zholenta, T. Chen, and D. Shatilov, PEP-II/AP Note 95.39/LBL-37680/CBP Note-152 (1995).

[9] K. Cornelis and M. Lamont, *Proceedings of the 4th European Particle Accelerator Conference, London, 1994*, 1150 (1994).

[10] A. Piwinski, *I.E.E.E. Transactions in Nuclear Science*, **NS-26**, 4268 (1979).

[11] A. Chao, *AIP Conference Proceedings, 127: Physics of High Energy Particle Accelerators*, (AIP Press, New York, 1983).

[12] R. Talman, Cornell Electron Storage Ring Report, CLNS-84/610 (1980).

[13] K. Hirata, *Nuclear Instruments and Methods in Physics Research*, **A269**, 7 (1988).

[14] R. E, Meller and R. H. Siemann, *I.E.E.E. Transactions in Nuclear Science*, **NS-28**, 2431 (1981).

[15] K. Yokoya, Y. Funakoshi, E. Kikutani, H. Koise, and J. Urakawa, High Energy Accelerator Research Organization (KEK) preprint 89-14 (1989).