Simulation of the beam halo from the beam-beam interaction

T. Chen, J. Irwin, and R. Siemann
Stanford Linear Accelerator Center, Stanford University, Stanford, California 94309
(Received 25 October 1993)

A technique for simulating the beam halo in circular e^+e^- colliders is introduced, tested, and applied. Amplitude space is divided into core and halo regions, and only halo particles are tracked saving a factor of 100 or more in CPU time. The methods for determining the regions, selecting and tracking particles, and connecting the core and halo are described. Results agree with conventional simulations. The beam halo is strongly influenced by nonlinear beam-beam resonances indicating that resonance streaming and phase convection are the dominant mechanisms for particles reaching large amplitudes.

PACS number(s): 41.75.Ht, 29.27.Bd, 29.20.Dh, 41.85.Ew

I. INTRODUCTION

Since the first e^+e^- collider, the beam-beam interaction has been an important, but not well understood, issue that dominates performance. Generally, based on operational experience, there are two beam-beam limits [1]. The first is the saturation of the beam-beam strength parameter,

$$\xi_y = \frac{r_e N \beta_y^*}{2\pi\gamma\sigma_y(\sigma_x+\sigma_y)} , \tag{1}$$

where r_e is the classical electron radius, N is the number of particles in a bunch, β_y^* is the vertical amplitude function at the interaction point, γ is the beam energy in units of rest energy, and σ_y and σ_x are the rms vertical and horizontal beam sizes, respectively. Flat beams, $\sigma_x \gg \sigma_y$, are assumed for the rest of this paper. Above some N the vertical beam size blows up leading to a constant value of ξ_y and the luminosity,

$$L = \frac{N^2 f_c}{4\pi\sigma_x\sigma_y} \approx \frac{N f_c \gamma \xi_y}{2 r_e \beta_y^*} , \tag{2}$$

being proportional to N rather than N^2 (f_c is the collision frequency). This first beam-beam limit is due to blow-up of particles in the beam core, i.e., small amplitude particles at high population densities. Coherent effects and nonlinearities can both be involved.

The second beam-beam limit comes from a short lifetime as N keeps increasing, finally stopping the luminosity increase. The bad lifetime results from particles driven to such large amplitudes by the beam-beam interaction that they fall outside the accelerator aperture and are lost. It has been observed that the beam halo deviates substantially from an extrapolation of the approximately Gaussian core [2].

The two beam-beam limits could be due to different physics, or they could be different manifestations of the same physics. Many beam-beam simulations have been written to investigate the beam core and to make luminosity predictions [3,4]. However, little progress has been made in studies of the beam halo. This paper introduces a simulation algorithm for investigating beam halos in e^+e^- circular colliders.

II. METHOD

The major difficulty of simulating the beam tail is that the particles determining the lifetime are rare; their density is 10^{-5}–10^{-8} of the core. Tracking for a huge number of particle turns is necessary to simulate the beam halo at an interesting level, and this requires an unacceptable amount of CPU time even for modern computers. However, since we are only interested in the particles in the halo when determining the lifetime, we do not have to track all particles, and not following core particles saves a tremendous amount of computing time. This is the basic idea of this algorithm. The remainder of this section deals with essential details which follow from earlier work by Irwin [5].

The simulation starts with 1000 particles that are Gaussian distributed in six-dimensional (6D) phase space. Those particles are tracked through the ring and receive beam-beam kicks from a counter-rotating beam with fixed transverse dimensions and bunch length, i.e., the weak-strong picture is used. The weak-strong model is valid because the forces on tail particles are determined by the core and there are so few tail particles that they cannot cause coherent motion. We are assuming that parametric driving from possible coherent motion of the core can be neglected. The particles are tracked for a few damping times to determine the equilibrium distribution, and, then, a boundary is determined in amplitude space so that about 100 particles are outside the boundary and 900 particles inside. After that, more tracking is performed to save all coordinates of particles crossing from inside to outside the boundary. This information is important for connecting the regions separated by the boundary. A large number of crossing coordinates is important to provide this boundary condition. The program saves up to 200 000 sets of crossing coordinates, and typically, the number of crossings in 10 000 turns is under 200 000. These crossing coordinates are important

 ©1994 The American Physical Society

to make connections between regions. Their usage in the next step is discussed below. The maximum amplitudes of a particle before the crossing are also saved for the lifelike calculation. At the same time the crossing coordinates are being obtained, the density distribution inside the boundary is saved and 1000 particle coordinates outside the boundary are saved by randomly choosing one particle outside the boundary every 10 turns and saving its phase-space coordinates. These are used as the initial coordinates for the next step.

The halo extends differently in the longitudinal, horizontal, and vertical dimensions, so we have chosen the ellipsoid shown in Fig. 1 as the boundary. The axes in the figure are the longitudinal, horizontal, and vertical amplitudes, normalized to the nominal sizes, and the principal axes of the ellipsoid are along these amplitude axes with lengths A_{sb}, A_{xb}, and A_{yb}, respectively. The boundary parameters are found as follows: (1) find amplitudes a_{sb}, a_{xb}, and a_{yb}, such that 10 particles have larger amplitudes in each direction, i.e., $N(A_i > a_{ib}) = 10$, $i = s,x,y$; (2) in the longitudinal, take $A_{sb} = a_{sb}$; (3) define a factor α by $A_{xb} = \alpha a_{xb}$ and $A_{yb} = \alpha a_{yb}$ and adjust α so there are 100 particles outside the boundary and 900 particles inside.

The boundary could be determined in other ways. A two-dimensional (A_x and A_y only with no consideration of A_s) and a three-dimensional cylindrical boundary (a maximum value of A_s but a profile in A_x and A_y that is independent of A_s) were tested. The resultant distributions in these tests were not sensitive to boundary parameters or shapes. We prefer the ellipsoid boundary because it retains particles with large longitudinal amplitudes.

The second step starts with the 1000 sets of coordinates saved in previous step. These particles are outside the first boundary, and they represent 10% of the total population. We are tracking the equivalent of 10 000 particles for a factor of 10 gain. Core particles, particles inside the boundary, are not tracked. During this second step, many particles initially outside the boundary fall inside it. We replace each such particle with a new one randomly chosen from among those that crossed the boundary from inside to outside on the previous step. That step was 10 000 turns long, and so each set of the coordinates tends to be used 10 times per 10 000 turns because of the particle number gain. Even though the same starting coordinates are used many times, subsequent phase-space trajectories are different due to quantum fluctuations.

Similar to the first step. (1) particles are tracked for a few damping times to determine the distribution, (2) a second boundary with 100 particles outside and 900 inside is found, (3) boundary crossing information is saved, (4) 1000 sets of coordinates of particles outside this second boundary are saved, and (5) the distribution between the first and second boundaries is obtained. This is repeated several times, as shown in Fig. 2, with a gain of ten in the equivalent number of particles each time. Each time the distribution inside the latest boundary is updated by multiplying by the gain.

There are two keys to making this algorithm reproduce the beam distribution. One is the randomness involved. Quantum excitation produces a random component of motion for every particle on every turn; a particle "forgets" its microhistory after several turns. In addition, there is randomness in selecting the 1000 particles that are tracked and in picking the boundary-crossing coordinates when replacing a particle. This randomness is essential for modeling the huge number of beam particles with a limited number of test particles. The second key is that particles are tracked for long enough that the distribution and boundary-crossing flux reach equilibrium. Effectively, the flux across the inner boundary serves as a boundary condition. In this way, information about the distribution inside the inner boundary is not ignored.

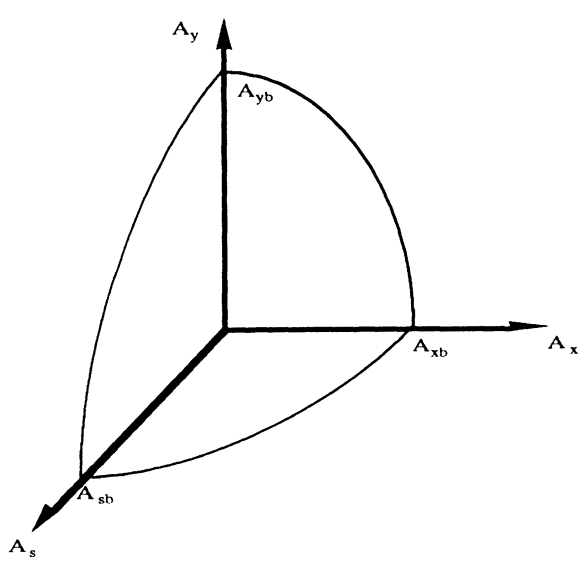

FIG. 1. The ellipsoid boundary.

III. THE PROGRAM

The simulation is written in standard FORTRAN. The tracking algorithm is like other strong-weak simulations. The strong beam is assumed Gaussian in all three dimensions, and the beam-beam interaction is calculated using the formula of Bassetti and Erskine [6]. The complex error function that appears in that result is evaluated using a Padé approximation [7]. The strong beam can be sliced into 3, 5, or 7 pieces longitudinally. This increases computing time significantly, but it is necessary for bunch lengths comparable to β_y^* [4]. The longitudinal positions of the beam-beam kicks, as well as the strong beam sizes at those positions, are modulated by synchrotron motion.

There is a self-consistent mode in which a strong-strong simulation is performed for about two damping times at the beginning. In this mode, the strong and weak beams have about the same core profile. The resulting core size is used as the strong beam size in the subsequent weak-strong beam halo simulation.

Parasitic crossings near the interaction point have been included. The parasitic crossing points appear symmetrically on both sides of the interaction point, and each parasitic interaction is represented by a single beam-beam

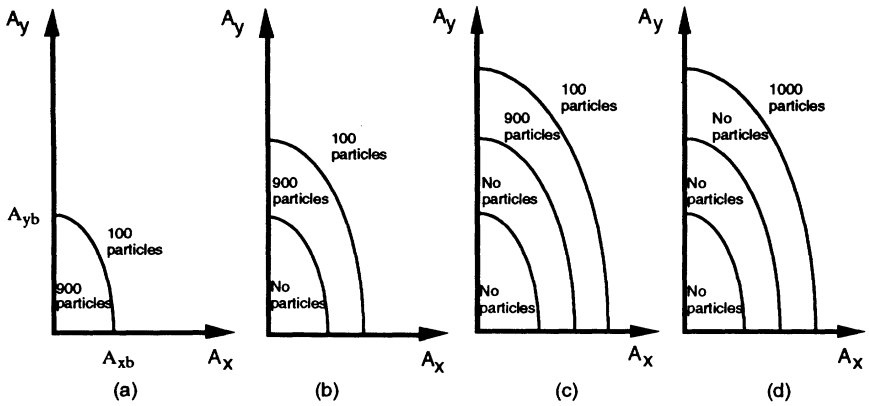

FIG. 2. An illustration of the simulation process. Two-dimensional boundaries are plotted for simplicity, and three boundaries are used in this example.

kick. The program can also handle two beams with different parameters. This allows studies of asymmetric colliders such as B factories.

Particles are transported through the rest of the ring with three linear transfer matrices, one for each of the three spatial dimensions. Synchrotron radiation damping and excitation are applied each turn. The dependence of β^* on energy, up to third-order chromaticity, and tune dependence on amplitudes can be included, but the effects of these nonlinearities on the beam halo has not been studied yet. It is also possible to replace the transfer matrices with a more accurate symplectic map. This is planned for the future.

The structure of the program is shown in Fig. 3. Each step illustrated in Fig. 2 is split into two loops. In the first, the *settling loop,* the distribution settles down to its equilibrium. The next boundary is determined at the end of this loop. The coordinates of particles crossing the boundary, the starting coordinates for the 1000 particles used in the next step, and the density distribution inside the boundary are then found in the *tracking loop.* Each loop in each step after the first includes insertion of new particles to replace those that fall inside the boundary.

The final step has two loops, but there is no need to find a boundary after the settling loop and to save crossing information in the tracking loop. The last tracking loop covers the large-amplitude region that particles rarely reach. It gives the tail distribution. The equivalent number of particle turns equals the length of the final tracking loop, multiplied by the total gain 10^B, where B is the number of boundaries. Therefore, the length of the final loop and the number of boundaries depend on the total particle turns desired.

Lifetimes are calculated as a function of aperture. To do this, we select a series of horizontal and vertical amplitudes as the presumed apertures. Once a particle's amplitude exceeds an aperture, this particle would be lost. The lost particles are counted at each aperture, and the lifetime in turns can be calculated from

$$\tau=\frac{N}{\Delta N}\ , \tag{3}$$

where N is the total number of particle turns, and ΔN is

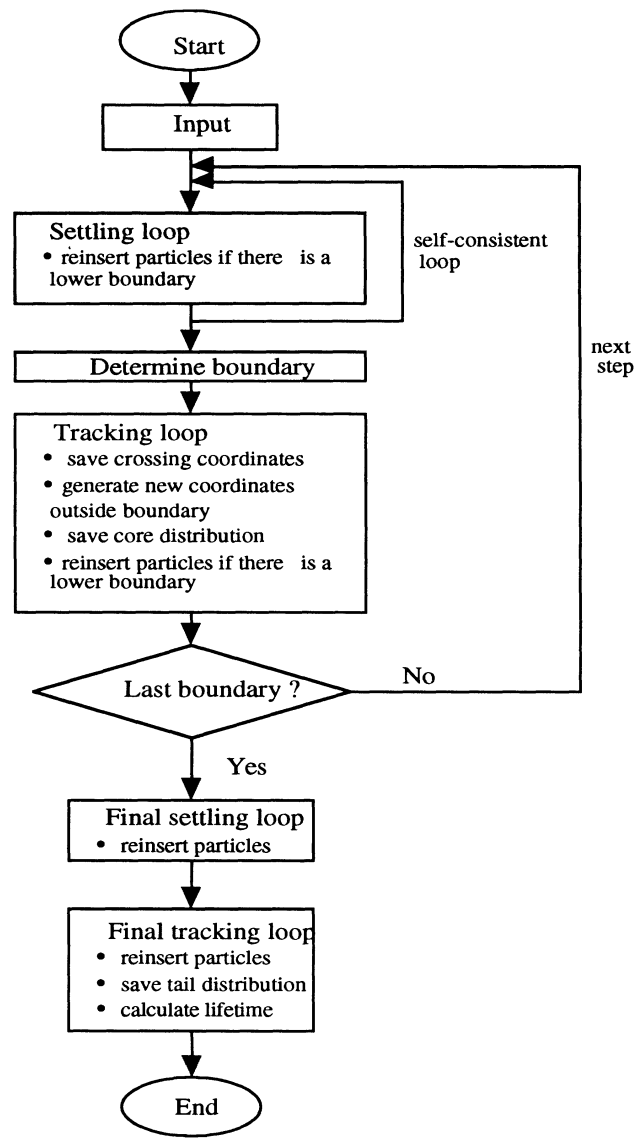

FIG. 3. Program flow chart.

the number of the particles that exceeded this aperture. The lifetime calculation is carried out only in the tracking loops. Each particle should be counted only once for exceeding a certain aperture. To avoid over counting particles that pass through an aperture more than once, we have to keep track of the maximum amplitudes of each particle and only count the first crossing. For a particle being reinserted at the boundary, its previous maximum amplitudes are also reinstated to keep its history. These maximum amplitudes were saved together with the crossing coordinates. The number of lost particles ΔN is counted in each step in the region between the previous boundary and the next boundary and needs to be updated by multiplying by the gain, just like we did in the distribution calculation.

IV. TESTS

Comparisons were made between this program and conventional tracking to verify the method. The first comparison was with a single beam-beam kick without synchrotron motion. The strong and weak beams had the same parameters, these are listed in Table I. The tunes are chosen such that two high-order coupling resonances, $4q_x+2q_y=4$ and $2q_x+6q_y=5$, were within the tune spread. Resonance streaming [8] and particles flowing from the sixth- to the eighth-order resonance were expected [9]. This is discussed more in Sec. V.

Figure 4 shows the distribution contours. The number of particles in the differential area $dA_x dA_y$ is parametrized as

$$n(A_x, A_y) \sim \exp[-\Phi(A_x, A_y)] \ . \tag{4}$$

The contours are lines of constant Φ that differ by $\Delta\Phi = 0.1|\Phi_{\max} - \Phi_{\min}|$. Figure 4(a) is the result of conventional tracking for 5.9×10^9 particle turns. Figure 4(b) is the result of this method for the equivalent number of tail particle turns. Two boundaries were used, and 1.29×10^8 particle turns were actually tracked. The agreement between the two distributions is good, and only about 2% of particle turns were used.

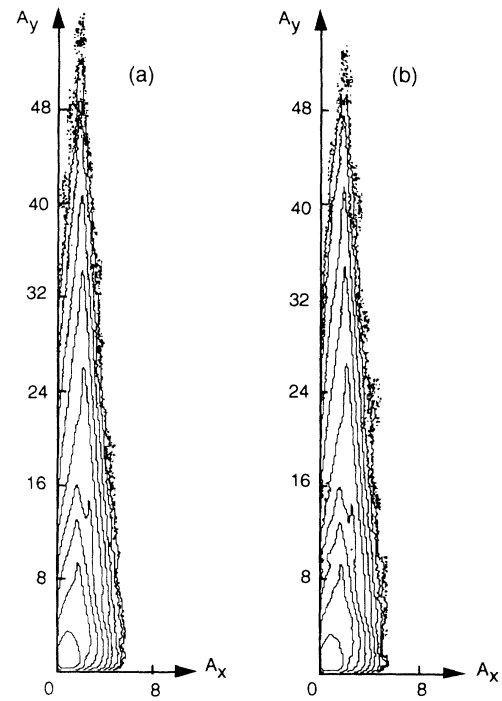

FIG. 4. Particle distribution contours in amplitude space. The parameters are listed in Table I. (a) was generated by conventional tracking, and (b) by the method proposed here. The particle density changes by a factor of 5.4 between contour lines.

A second comparison included synchrotron motion. The parameters are also listed in Table I. They are from the PEP-II Conceptual Design Report [10]. The comparison between conventional, brute-force tracking, and the method proposed here is shown in Fig. 5. The brute-force tracking was for 3×10^9 particle turns, while the method proposed here used 8×10^7 particle turns. Again,

TABLE I. Beam parameters. β_x^* and β_y^* are the horizontal and vertical amplitude functions at the interaction point. q_x, q_y, and q_s are horizontal, vertical, and synchrotron tunes. τ_B is the radiation damping time for betatron motion. ε_x and ε_y are the horizontal and vertical emittances. σ_L is the rms bunch length. ξ_x and ξ_y are the horizontal and vertical beam-beam strength parameters; ξ_y is given by Eq. (1) and ξ_x is given by Eq. (1) with subscripts x and y interchanged.

Parameter	Figs. 4 and 8 both beams	Figs. 5–7 strong beam	Figs. 5–7 weak beam
β_x^*	1 m	0.5 m	0.375 m
β_y^*	2 cm	2 cm	1.5 cm
q_x	0.682	0.57	0.57
q_v	0.6012	0.64	0.64
q_s			0.05
τ_B	5000 turns		5014
ε_x	2.4×10^{-7} m	4.6×10^{-8} m	6.1×10^{-8} m
ε_v	1.2×10^{-9} m	1.8×10^{-9} m	2.4×10^{-9} m
σ_L		1 cm	1 cm
ξ_x	0.03		0.03
ξ_y	0.06		0.03

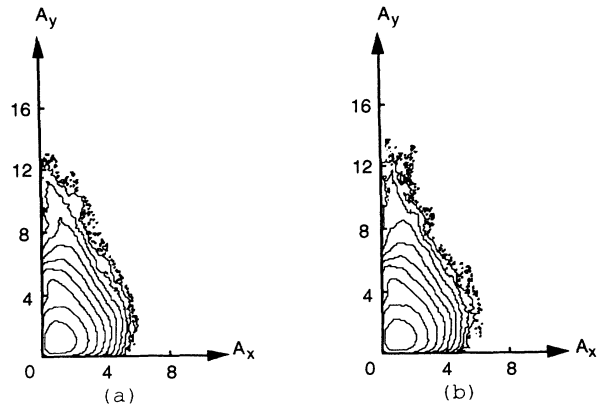

FIG. 5. Particle distribution contours in amplitude space. Simulation includes synchrotron motion. (a) was generated by conventional tracking, and (b) by the method proposed here. The particle density changes by a factor of 5.4 between contour lines.

the agreement is good.

We have studied the sensitivity to parameters of the simulation itself. The most important issue is the length of the settling loop. Figure 6 gives the vertical tail of settling for 2, 4, and 6 damping times. The tail is defined as

$$V_T(A_y)=\frac{1}{N}\int_0^\infty dA_x \int_{A_y}^\infty dA_y' n(A_x, A_y') , \tag{5}$$

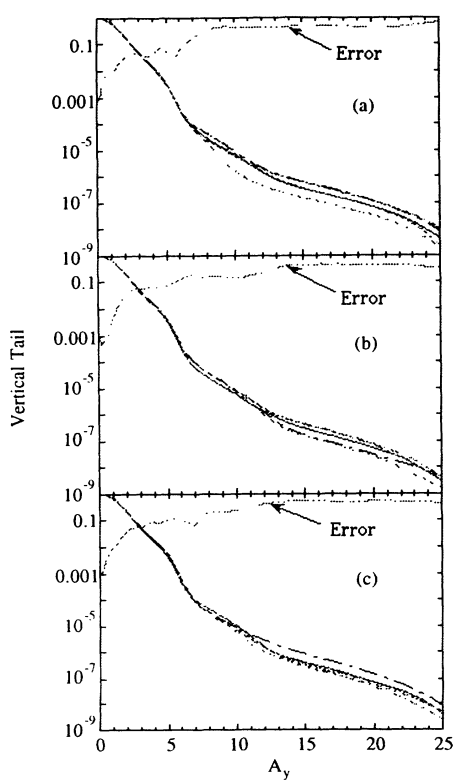

FIG. 6. Vertical tail distribution with different settling times: (a) 2 damping times; (b) 4 damping times; (c) 6 damping times. The vertical axis is V_T, the fraction of particles beyond a certain amplitude, a dimensionless number defined in Eq. (5).

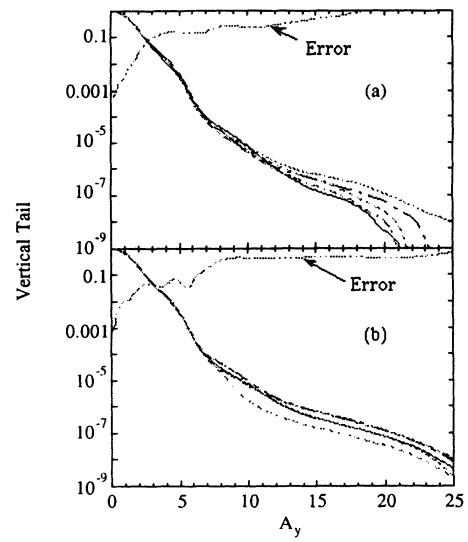

FIG. 7. Vertical distributions with different number of boundaries: (a) 3 boundaries and (b) 4 boundaries. The final tracking loop was 50 000 turns long in both cases, and the settling loop is 2 damping times.

where N is the total number of particle turns. In each case, five different random seeds are used to estimate the relative error; that is plotted also. That error is about 0.1% in the core and rises to the 10–40 % level for most of the tail. It reflects the statistics of the sampling process used in initializing the core distribution and compiling boundary-crossing information.

The lengths of the tracking loops in all but the last step are determined by the need for sufficient boundary-crossing information. We keep the length of the loop the same order as the settling loop to provide enough crossing coordinates. For 10 000 turns, there are 100 000–200 000 crossings typically, and this is adequate.

The dependence on the number of boundaries is shown in Fig. 7 where the results of using three and four boundaries are compared. Four boundaries are necessary for the distribution at $A_y=20$ to have roughly 30% accuracy. In this particular case for a 10^9 turn lifetime, 5×10^{11} equivalent particle turns give a good lifetime estimate.

The percentage of particles that determines the boundaries is also a free parameter in this method. This parameter is usually chosen as 10%. A larger value will reduce the efficiency of the method, and a smaller value can result in larger statistical error. It has been tested that the variation of this parameter around 10% does not change the result.

V. UNDERLYING PHYSICS

There are different phenomena that could lead to beam halos. The traditional interpretation is based on the Chirikov criterion [11,12]. When high-order resonances are wide enough, or close enough, they can overlap leading to chaotic motion with particles moving from one resonance to another and reaching large amplitudes [13].

Resonance overlap is expected to be more important for large energy oscillation amplitudes because the number of important synchro-betatron resonances increases with energy amplitude.

A second proposed phenomenon is diffusion. Particles starting at locations throughout the core slowly diffuse to large amplitudes where they move as oscillators driven by noise from the beam-beam kick which is treated as a random process because of radiation damping and quantum excitations [14]. This could explain a non-Gaussian tail and is not dependent on nonlinear resonances.

A third possible phenomenon is resonance streaming [8,9]. Quantum fluctuations drive particles into nonlinear resonances. These particles oscillate around the resonance center which is a locus in the A_x-A_y plane satisfying the resonance condition

$$pq_x(A_x, A_y)+rq_y(A_x, A_y)+mq_s=n \ , \tag{6}$$

where p, r, m, and n are integers. The effect of damping is to move the oscillation center along this locus, and, for sum resonances where both p and r are positive, a decrease of A_x can result in an increase of A_y. Radiation damping actually increases the vertical amplitude and can transport a particle quickly to large amplitudes. In this case, the halo depends on the resonance structure of the beam-beam interaction.

Our program is designed to rapidly simulate the beam halo without introducing bias towards any one mechanism. The beam-beam interaction is treated as a kick; therefore, it includes all overlapping or isolated resonances. Diffusion would be accommodated naturally by the global expansion of the boundary separating core and halo, and the 3D ellipsoidal boundary guarantees the inclusion of particles with large energy oscillation amplitudes. The program was tested for conditions where brute-force tracking showed that resonance streaming is dominant [15].

Figure 4 shows the results of one such test. The distribution in Fig. 4 is plotted again in Fig. 8 with two sum resonances overlaid. The resonance lines are given by Eq. (6) with the amplitude dependence of the tunes, q_x and q_y, determined to first order from the average value of the beam-beam potential [1]. One might question whether averages of the potential could be expected to give accurate estimates of the resonance centers. The second-order terms are proportional to the derivatives of their first-order terms. At large amplitudes, since the first-order terms are nearly flat, the second-order perturbations can be neglected. For small amplitudes, the second-order perturbation is weak by definition. Therefore, the first-order perturbation theory is a good approach for determining resonance locations in the beam-beam problem, if the working point is not near a low-order resonance [16]. Particles follow the resonance $4q_x+2q_y=4$ from the core to $A_y \sim 15$, and then they follow a second resonance, $2q_x+6q_y=5$, to $A_y > 50$. This simulation of a beam distribution known to be determined by resonance streaming shows that the method is not biased against it.

The PEP-II parameters were chosen by optimizing a complete collider design. Required performance, site geometry, single particle dynamic aperture, RF, etc. entered into the optimization. The mechanism for producing the beam-beam halo was not considered, and, therefore, the PEP-II parameters offer an opportunity for distinguishing between the possible mechanisms in a realistic, unbiased example. Figures 9 and 10 show the results of a PEP-II simulation. The parameters are the same as

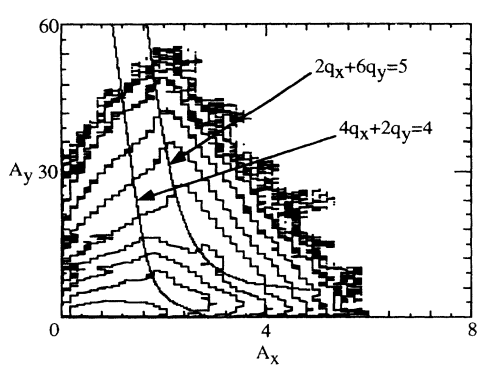

FIG. 8. The tail distribution as a result of resonance streaming. The contour plot is the same as Fig. 4(b).

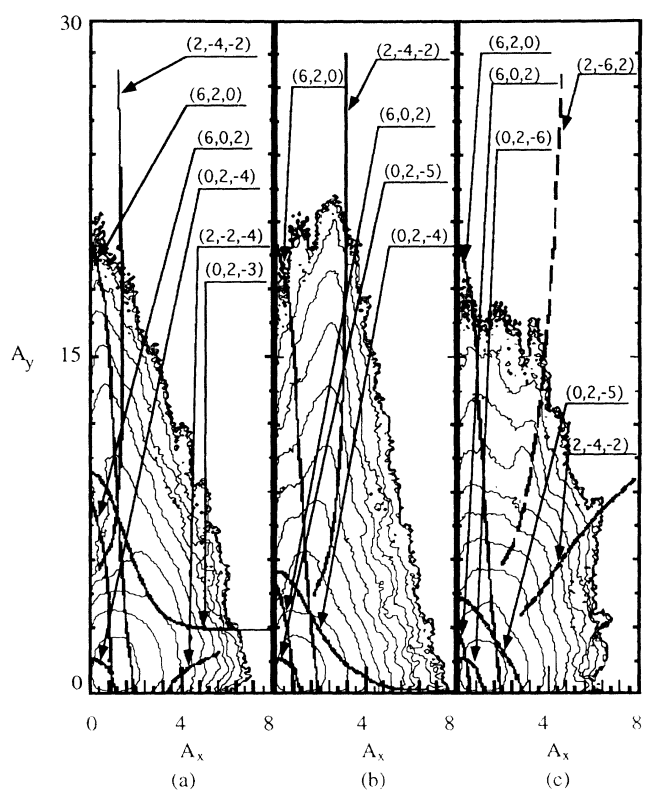

FIG. 9. Resonance lines and beam distributions for $q_x=0.63$, $q_y=0.552$, and different synchrotron tunes. (a) $q_s=0.03716$, $\sigma_L=1$ cm, (b) $q_s=0.029\,73$, $\sigma_L=1.25$ cm, and (c) $q_s=0.024\,77$, $\sigma_L=1.5$ cm. All resonances up to eighth order are drawn as solid lines. The dashed line in (c), is a tenth-order resonance selected because of its apparent role in the tail formation. Resonances $pq_x+rq_y+mq_s=n$ are labeled as (p,r,m) in the figure.

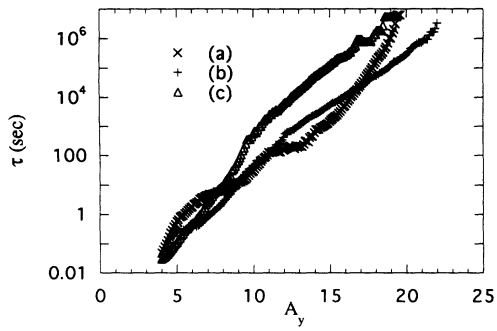

FIG. 10. The lifetime as a function of vertical aperture. The parameters in case (a), (b), and (c) are the same as those in Fig. 9. For (a) and (b), two hour lifetime requires $A_y > 16$, while $A_y > 12$ would be needed for (c).

those for Figs. 5–7, except for the tunes. The transverse tunes are $q_x = 0.63$, $q_y = 0.552$. The synchrotron tune is different for the three parts of the figure, and the bunch length changes accordingly.

The halo is strongly affected by resonances. The transverse resonance $6q_x + 2q_y = 5$ dominates it in Fig. 9(a), where $q_s = 0.037\,16$ and $\sigma_L = 1$ cm. Along with other changes this resonance is less prominent and the synchro-betatron resonance $2q_y - 3q_s = 1$ has moved out of the beam footprint in Fig. 9(b) ($q_s = 0.029\,73$, $\sigma_L = 1.25$ cm). Comparison of the two figures suggests that $2q_y - 3q_s = 1$ has enhanced the halo in Fig. 9(a) by providing a route for particles from the core to reach $6q_x + 2q_y = 5$. Gerasimov and Dikansky have developed the concept of a "most likely path" to large amplitudes; that path could involve one or several resonances [9]. (Figure 8 is another illustration of a path involving two resonances.)

The most prominent resonance in Fig. 9(b) is the synchro-betatron resonance $2q_x - 4q_y - 2q_s = -1$. That resonance moves to larger A_x and becomes less important in Fig. 9(c) where $q_s = 0.024\,77$ and $\sigma_L = 1.5$ cm. The resonance $6q_x + 2q_y = 5$ dominates again. However, it is not as strong as in Fig. 9(a), and a tenth-order resonance, $2q_x - 6q_y + 2q_s = -2$, contributes to the tail.

The structures in the beam distributions can be associated with resonances that have locations determined from first-order perturbation theory. Diffusion and chaotic motion by itself cannot explain this. Chaotic motion is expected at the resonance boundaries, but theoretical diffusion rates are in general too small to explain the particle motions observed. However, our observations are consistent with the resonance streaming and phase convection model. We conclude that this mechanism would determine the beam lifetime for a collider with the PEP-II parameters and conjecture that it is the most important one for other e^+e^- storage rings.

VI. CONCLUSION

This simulation method can generate tail distributions with a substantial reduction in computing time. The method has been verified by comparison with conventional tracking, and has been found to be insensitive to details such as choices of boundary shape. The method opens the door to the exploration of the second beam-beam limit.

A first step in this is the study of the mechanism leading to large-amplitude particles. The results for the PEP-II simulation show that resonance streaming is the dominant effect and lead to the conjecture that it is dominant in other colliders. Future work will include studies of parametric dependence, control of beam halo by modifying the resonance structure, and the simulation of operating storage rings.

This method could be useful in other computationally intensive branches of physics where a "self-generated boundary condition" would make it possible to simulate a particularly interesting regime while saving a significant amount of CPU time.

ACKNOWLEDGMENTS

This work was supported by the Department of Energy, Contract No. DE-AC03-76SF00515.

[1] R. Siemann, in *Frontiers of Particle Beams with e^+e^- Rings,* edited by M. Dienes, M. Month, B. Strasser, and S. Turner (Springer-Verlag, Berlin, 1994), p. 327.

[2] J. Seeman, in *Proceedings of the 12th International Conference on High-Energy Acccelerators,* edited by Francis T. Colea and Rene Donaldson (Fermilab, Batavia, 1983), p. 212.

[3] S. Myers, in *Nonlinear Dynamics Aspects of Particle Accelerators* edited by J. M. Jowett, S. Turner, and M. Month (Springer-Verlag, Berlin, 1986), p. 176; S. V. Milton and R. M. Littauer, in *Proceedings of the 1989 IEEE Particle Accelerator Conference,* edited by Floyd Bennett and Joyce Kopta (IEEE, New York, 1989), p. 824; S. Peggs and R. Talman, Phys. Rev. D **24**, 2379 (1981); A. Piwinski, IEEE Trans. Nucl. Sci. **NS-30**, 2378 (1983); G. Jackson and R. Siemann, Nucl. Instrum. Methods A **286**, 17 (1990).

[4] S. Krishnagopal and R. Siemann, Phys. Rev. D **41**, 2312 (1990).

[5] J. Irwin, in *Third Advanced ICFA Beam Dynamics Workshop on Beam-Beam Effects in Circular Colliders,* edited by I. Koop and G. Tumaikin (Institute of Nuclear Physics, Novosibirsk, 1989), p. 123; J. Irwin, SLAC Publication No. SLAC-PUB-5743, 1992 (unpublished).

[6] M. Bassetti and G. Erskine CERN Report No. CERN-ISR-TH/80-06 (unpublished).

[7] Y. Okamoto and R. Talman, CBN Report No. 80-13, 1980 (unpublished).

[8] J. Tennyson, Physica **5D**, 123 (1982).

[9] A. L. Gerasimov and N. Dikansky, Nucl. Instrum. Methods A **292**, 209 (1990); **292**, 221 (1990); **292**, 233

(1990).

[10] PEP-II Conceptual Design Report, Report Nos. LBL-PUB-5379, SLAC-418, CALT-68-1869, UCRL-ID-114055, UC-IIRPA-93-01 (1993) (unpublished).

[11] B. V. Chirikov, Phys. Rep. **52**, 263 (1979).

[12] A. L. Gerasimov, F. M. Izrailev, and J. L. Tennyson, in *Physics of Particle Accelerators,* Proceedings of the Fifth Annual U.S. Particle Accelerator School, edited by M. Month, AIP Conf. Proc. No. 153 (AIP, New York, 1987).

[13] A. W. Chao, in *Physics of High Energy Particle Accelerators,* Proceedings of the Third Annual U.S. Summer School on High Energy Particle Accelerators, edited by M. Month, P. F. Dahl, and M. Dienes AIP Conf. Proc. No. 127 (AIP, New York, 1985), p. 201.

[14] K. Cornelis, in *Tail Production Due to Beam-Beam,* Proceedings of the Third LEP Performance Workshop, edited by J. Poole (CERN, Geneva, 1993), p. 123.

[15] A. L. Gerasimov and R. Siemann, in *Nonlinear Problems in Future Particle Accelerators,* edited by W. Scandale and G. Turchetti (World Scientific, Singapore, 1991), p. 197.

[16] E. Forest (private communication).