Simulation and analysis of using deflecting cavities to produce short x-ray pulses with the Advanced Photon Source

M. Borland

Argonne National Laboratory, Argonne, Illinois 60439, USA

(Received 9 March 2005; published 13 July 2005)

We present detailed simulations and analysis of Zholents's [A. Zholents, P. Heimann, M. Zolotorev, and J. Byrd, Nucl. Instrum. Methods Phys. Res., Sect. A **425**, 385 (1999).] concept for using deflecting cavities in a synchrotron light source storage ring for the purpose of producing short x-ray pulses. In particular, we look at the optimization and performance of such a system for the Advanced Photon Source. We find the concept is practical and that x-ray pulse durations of about 1.5 ps FWHM should be achievable with more than 15% of the original intensity retained. Issues covered include lattice design, emittance degradation, lifetime, photon beam modeling, errors, and optimum choice of rf parameters.

DOI: 10.1103/PhysRevSTAB.8.074001 PACS numbers: 29.20.−c, 29.27.−a, 41.60.Ap

I. INTRODUCTION

This paper presents detailed simulation studies and analyses of x-ray compression using transverse-deflecting rf cavities [1] in the Advanced Photon Source (APS), a 7-GeV synchrotron radiation source. As is well known, efforts to produce short x-ray pulses in storage rings are fraught with difficulty. In the APS, the beam is naturally long (about 20 ps rms) because of the relatively large energy spread (about 0.096% rms) and the limited rf voltage possible with continuous-duty systems. The desire for long Touscheck lifetimes also motivates against making the electron bunch too short. What is more, as the single-bunch current is increased, considerable bunch lengthening is seen due to potential well distortion [2]. As a result, the bunch length for the standard APS operating mode with 5 mA per bunch is about 40 ps rms. Similarly, use of an isochronous lattice gives short bunches, but at greatly reduced intensity [3]. Use of laser slicing [4] at 7 GeV would require an impractical amount of laser power [5], and the intensity is in any case quite small with this scheme. In contrast, many experimenters are eager for intense subpicosecond x-ray pulses. Zholents's concept of using transverse-deflecting cavities promises a way out of this dilemma.

We begin with a review of the concept and potential problems. We then discuss lattice choices and constraints, showing a lattice that will provide benefits to four undulator beam lines. Following this, an analytical treatment of x-ray compression is given, which provides an indication of how much deflecting voltage will be required. Limitations on the deflecting voltage and frequency related to beam lifetime requirements are covered next. After reviewing tracking methods, we explore through analysis and simulation a major limitation, namely, emittance-degrading effects, some of which are intrinsic to the machine and others of which are related to errors. Finally, we describe how to model the photon beam and its compression, both analytically and in simulation, and give predictions of ultimate performance.

II. REVIEW OF THE CONCEPT

In this concept, transverse-deflecting rf cavities ("crab cavities") are used to impose a correlation ("chirp") between the longitudinal position of a particle within the bunch and the vertical momentum. Two cavities are required in order to avoid extremely large vertical emittance growth and to provide the same chirp on each turn. The second cavity must be placed at a vertical phase advance $n\pi$ downstream of the first, where n is an integer, so as to cancel the chirp. The cavities must have a TM110 mode frequency that is a harmonic h of the ring rf frequency (which itself is the 1296th harmonic of the revolution frequency). We also characterize their deflection using an effective voltage V, which can impart a maximum slope change of V/E, where E is the beam energy.

If an undulator or bending magnet is placed between the cavities, then the emitted photons will have correlations among time (t), vertical position (y), and vertical slope (y'). Given the typical minimum length of beam lines at the APS (about 30 m), the time/slope correlation provides the most promising capability. At that distance, a strong time/position correlation develops, which can be used for either pulse slicing or pulse compression. In order to reduce the beam size in the undulator source and maximize the chirp of the slope, it is advantageous to place the undulators at locations that are approximately $m\pi$ (where $m < n$ is an integer), distant from the first cavity in vertical phase advance.

Once the photon beam has drifted a sufficient distance, the pulse can clearly be shortened simply by using vertical slits. This will, of course, throw away considerable intensity. Another option [1] is to use asymmetric-cut crystals to perform pulse compression. Such crystals allow imposing a time-of-flight variation that is proportional to vertical position. While this also results in some attenuation [6] (see Sec. X), it is advantageous nonetheless.

Ideally, the beam outside the region between the chirping cavities will see no effect from the chirping system. This is only true for a perfectly linear system without errors. In real systems, errors, nonlinearities, and other details destroy the perfect cancellation of the kicks, resulting in vertical emittance growth. Some emittance-degrading effects we have examined are listed here. These are detailed in subsequent sections.

(i) Nonzero momentum compaction and energy spread, which leads to time-of-flight variation between the cavities.

(ii) Chromaticity and energy spread, which leads to variation in the phase advance.

(iii) Quantum excitation in the longitudinal plane, which amplifies the effect of the momentum compaction and chromaticity.

(iv) Sextupoles between the cavities, resulting in x-y coupling, tune shift with amplitude, and phase-space distortion.

(v) Lattice errors resulting in phase-advance errors and beta-function differences at the cavities.

(vi) Roll, either relative to each other or relative to "true vertical," of the cavities about the beam axis.

(vii) Lattice coupling between cavities. This may be due, for example, to rolled elements between the cavities or due to skew quadrupoles that are powered to correct the global coupling.

(viii) rf phase and voltage errors.

III. LATTICE CHOICES AND CONSTRAINTS

A. Phase-advance requirement

As mentioned in the introduction, a basic requirement for the optics is to have vertical phase advance of $n\pi$ between the cavities. Without this, there is no cancellation of the deflecting kicks and the vertical emittance will grow very large. The original proposal for APS [5] also required $p\pi$ phase advance in the horizontal plane, where p is an integer, and where p and n must both be either odd or even. This was thought to be necessary to provide cancellation of horizontal kicks for off-axis particles.

We have found that this is not necessary. This can be seen from the expressions for the Cartesian components of the magnetic field [see Eq. (A7) in the appendix]. The term we are concerned about is B_x, which deflects in the "other" plane in this case. The ratio of B_x/B_y is

$$\frac{B_x}{B_y} \approx \frac{\omega^2 xy}{4c^2}. \tag{1}$$

Requisite power sources are available [7] for rf frequencies as high as 2.8 GHz, i.e., the $h = 8$ harmonic of the APS 352-MHz rf system. The rms beam sizes in the horizontal and vertical planes are, respectively, approximately 280 and 11 μm. At 10σ, we have $(B_x/B_y) \approx 0.0003$. Hence, we do not anticipate a problem from deflection in the horizontal plane.

To confirm this, we tracked with ELEGANT [8] using a lattice that has 2π phase advance between the cavities in the vertical plane, and 3.64π phase advance in the horizontal plane. The tracking used the RFTM110 element, which has the TM110 fields included via a sixth-order expansion. We tracked a 5σ Gaussian beam, using 5×10^6 simulation particles, from the entrance of the first cavity to the exit of the second using linear optics only. We used a harmonic number of 8 and deflecting voltages of up to 6 MV, which is the largest chirp we are considering (see Sec. V). The variation in horizontal emittance for these runs was less than 1 part in 10^6. Hence, we dropped the constraint on the horizontal phase advance.

In addition to the TM110 fields, there are TE11-like fields [9] due to the irises of the cavity. TE11-mode fields have the same transverse spatial dependence as the TM110 fields [10]. They also add to magnetic deflection [9] and hence have the same orientation. This implies that the TE11-like fields will, like the TM110 fields, not corrupt the emittance in the nondeflecting plane.

B. Lattice options

The APS lattice presently has a vertical tune of about 19.27, meaning that the phase advance in each of the 40 sectors is close to π. This is fortunate as it means that cavities in successive straight sections are already close to satisfying our requirement. The simplest possible arrangement, with minimal optics changes, involves putting the cavities 1 or 2 sectors apart in equivalent locations in the straight section. These options are shown in Fig. 1. One then makes a slight lattice adjustment to bring the phase advance to the required value. These changes are small enough that the beta functions are changed very little. These are shown in Fig. 2.

The disadvantage of this scheme is that space is taken up in one straight section without delivering chirped pulses to that straight section. Another scheme, which resolves this issue, would be to exchange the position of the first cavity and insertion device, as shown in Fig. 3. The advantage is that this arrangement would provide chirped pulses to more insertion devices. The difficulty is that one must now *reduce* the phase advance between the cavities, as in the existing lattice it would be too large (since the cavities are now on opposite sides of the beta-function waists). This requires increasing the vertical beta function in the insertion device straight section from the nominal value of about 3 m.

We have developed a two-sector lattice along these same lines, providing chirped pulses to four insertion devices. This offers similar benefits to more beam lines for the same cost. The lattice functions for this solution are shown in Fig. 4, where one sees that the vertical beta function in the

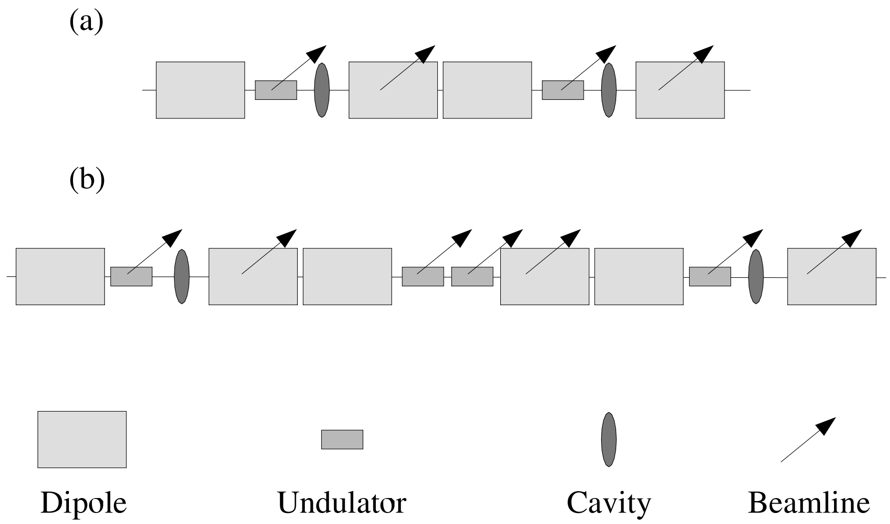

FIG. 1. Arrangement of components requiring minimal changes to the APS optics (after Sajaev [24]). In the first arrangement (a) chirped pulses are provided to a single undulator beam line and a single bending magnet beam line. In the second arrangement (b) chirped pulses are provided to three undulator beam lines (assuming canting of the middle pair [12]) and two bending magnet beam lines. The many quadrupoles and sextupoles that exist in the lattice are omitted from this drawing, for clarity.

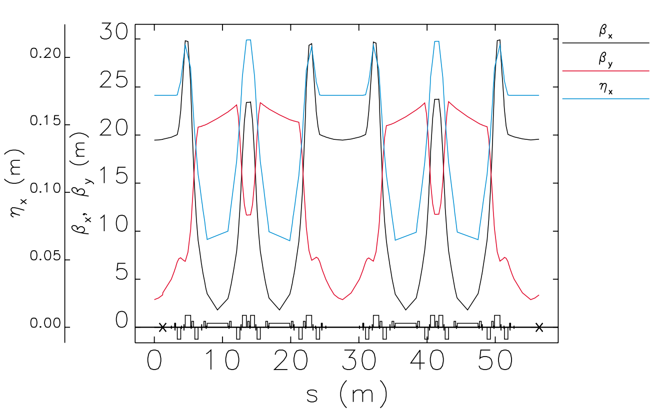

FIG. 2. (Color) Lattice functions corresponding to arrangement (b) from Fig. 1. The approximate locations of the cavities are marked with an “X”.

straight section has been raised from about 3 m to about 5 m. We have used this lattice in all subsequent work.

IV. X-RAY COMPRESSION

In this section, we give analytical results for x-ray compression assuming Gaussian beams, a linear chirp, and linear compression. This will help establish the magnitude of the chirp required to obtain the desired compression. As in electron bunch compressors, the degree to which one can compress a photon pulse depends on the slice emittance of the photon beam. For photons from an undulator, the emittance results from a combination of the electron beam emittance and the intrinsic opening angle of the undulator radiation.

A commonly used approximation [11] is to assume the single-electron radiation distribution is a Gaussian function of the polar angle θ, with rms parameter

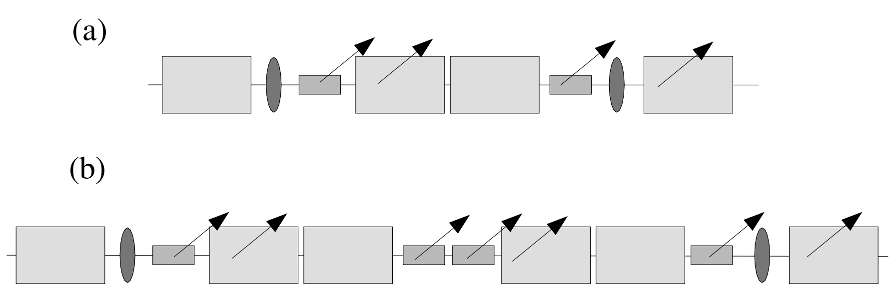

FIG. 3. Arrangements of components requiring significant but manageable changes to the APS optics (after Sajaev [24]). These arrangements provide chirped pulses to more beam lines for the same cost compared those in Fig. 1. In the first arrangement (a) chirped pulses are provided to two undulator beam lines and one bending magnet beam line. In the second arrangement (b) chirped pulses are provided to four undulator beam lines and two bending magnet beam lines. The many quadrupoles and sextupoles that exist in the lattice are omitted from this drawing, for clarity.

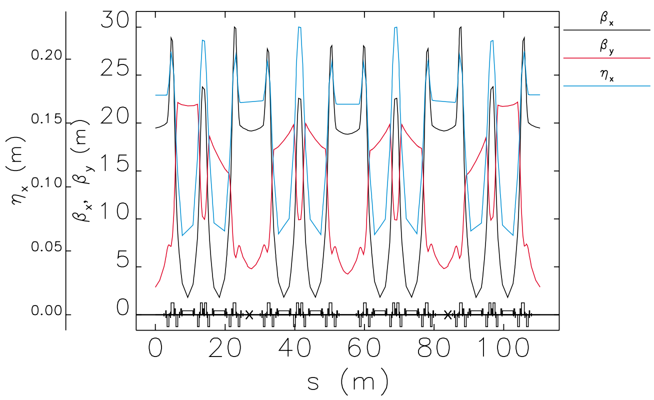

FIG. 4. (Color) Lattice functions corresponding to arrangement (b) from Fig. 3. The approximate locations of the cavities are marked with an X.

$$\sigma_\theta = \sqrt{\frac{\lambda}{2L_u}}, \tag{2}$$

where λ is the radiation wavelength and L_u the undulator length. We will see below (Sec. X) that this is a reasonable approximation to a more accurate distribution function that is used in the simulations.

To determine how much compression might be possible, we have adopted a simple model for the action of the asymmetric cut crystal. In particular, we assume that it is simply a linear process [1]

$$\delta t = K(y + Ly'), \tag{3}$$

where δt is the variation in time of flight for a photon and K is a constant determined by the parameters of the crystal. A typical APS beam line has the monochromator about 30 m from the source point, so we will take $L = 30$ m.

We can analytically compute the K value required for maximum compression if we assume that the chirp is linear. The arrival time offset (relative to pulse center) of a ray at the exit to the compression optics is

$$\Delta t_2 = \Delta t_1 + K y_1, \tag{4}$$

where the "1" subscripts refer to the entrance and "2" subscripts refer to the exit of the optics. Of course, $y_1 = y_0 + Ly'_0$, where "0" refers to the source point. The mean square of t_2 is

$$\langle \Delta t_2^2 \rangle = \langle \Delta t_1^2 \rangle + 2K\langle \Delta t_1 y_1 \rangle + K^2 \langle y_1^2 \rangle. \tag{5}$$

Minimizing this with respect to K gives

$$K = -\frac{\langle \Delta t_1 y_1 \rangle}{\langle y_1^2 \rangle}, \tag{6}$$

and thus the minimum achievable x-ray pulse duration is

$$\langle \Delta t_2^2 \rangle_{\min} = \langle \Delta t_1^2 \rangle - \frac{\langle \Delta t_1 y_1 \rangle^2}{\langle y_1^2 \rangle}. \tag{7}$$

There are three contributions to y'_0: the chirp, the angular spread due to the uncorrelated vertical beam emittance, and the angular spread due to the radiation. Taking the terms in order, we can write this as

$$y'_0 = C\Delta t_0 + \delta y'_0 + \sigma_\theta, \tag{8}$$

where $C = V\omega/E$ is the deflecting cavity parameter, which depends on the voltage V, the beam energy E, and the angular rf frequency ω. This rf frequency is a harmonic h of the APS rf frequency, $\omega_m/(2\pi) = 351.9$ MHz. Recognizing that $y'_1 = y'_0$ and $\Delta t_1 = \Delta t_0$, we can work out the terms in Eq. (7):

$$\langle \Delta t_1^2 \rangle = \langle \Delta t_0^2 \rangle, \tag{9}$$

$$\langle \Delta t_1 y_1 \rangle = CL\langle \Delta t_0^2 \rangle, \tag{10}$$

and

$$\langle y_1^2 \rangle = \langle y_0^2 \rangle + L^2\langle \delta y_0'^2 \rangle + L^2\sigma_\theta^2 + C^2L^2\langle \Delta t_0^2 \rangle, \tag{11}$$

where for simplicity we have assumed $\langle y_0 y'_0 \rangle = 0$. This occurs if $\alpha_y = 0$ at the center of the undulator and if there is exactly $n\pi$ phase advance from the cavity to the center of the undulator. The minimum pulse duration is thus

$$\langle \Delta t_2^2 \rangle_{\min}^{1/2} = \langle \Delta t_0^2 \rangle^{1/2} \times \sqrt{1 - \frac{C^2L^2\langle \Delta t_0^2 \rangle}{\langle y_0^2 \rangle + L^2\langle \delta y_0'^2 \rangle + L^2\sigma_\theta^2 + C^2L^2\langle \Delta t_0^2 \rangle}}. \tag{12}$$

It is informative to compare the magnitude of the individual terms. If we make a modest choice of $hV = 8$ MV, we get $C = 2.52$ MHz . Ignoring for the moment the degradation of the vertical emittance, for the APS with 1% coupling, the unchirped rms vertical beam size and divergence are 11 μm and 2.2 μrad, respectively. We also assume a 40-ps rms bunch duration. Hence, in order of decreasing magnitude, we have

$$C^2L^2\langle \Delta t_0^2 \rangle = 9.2 \times 10^{-6}\ \mathrm{m}^2, \tag{13}$$

$$L^2\sigma_\theta^2 = 1.8 \times 10^{-8}\ \mathrm{m}^2, \tag{14}$$

$$L^2\langle \delta y_0'^2 \rangle = 4.4 \times 10^{-9}\ \mathrm{m}^2, \tag{15}$$

$$\langle y_0^2 \rangle = 1.2 \times 10^{-11}\ \mathrm{m}^2. \tag{16}$$

The resulting minimum rms pulse duration is 2.0 ps.

One might think that increasing the distance to the monochromator is advantageous. However, because $\langle y_0^2 \rangle$ is much smaller than the other terms in the denominator in Eq. (12), the minimum pulse duration is very insensitive to further increases in L. The only way to make the pulse shorter is to increase the slope of the kick, i.e., to increase hV. If we desire less than 1 ps FWHM, then we would need $hV \approx 38$ MV. Of course, we must keep in mind that we have assumed the chirp is linear, which is not a good assumption even for $h = 4$. In addition, we have used the nominal vertical emittance, whereas we will see that the actual vertical emittance will increase as we increase the chirp. This implies the possibility of a point of diminishing returns as the chirp is increased.

Given the relative values of the terms, we can derive the following approximate form of Eq. (12) by expanding the denominator to first order in the limit $L \to \infty$

$$\langle \Delta t_2^2 \rangle_{\min}^{1/2} \approx \frac{E}{V\omega}\sqrt{\langle \delta y_0'^2 \rangle + \sigma_\theta^2}. \quad (17)$$

We will see below that the vertical emittance will increase by a factor of about 7.4 for 4 MV and $h = 8$. This will increase the slice divergence $\langle \delta y_0'^2 \rangle^{1/2}$ by a factor of about 2.7, making it slightly larger than the radiation opening angle and leading to a 50% increase in the minimum achievable pulse duration. The estimated minimum pulse duration for $h = 8$, $V = 4$ MV, $\lambda = 1.2$Å (10 keV), and a 2.4-m undulator is thus 0.77 ps rms or 1.8 ps FWHM. This should be achievable in spite of the sinusoidal nature of the rf deflection, but, as we will see, only at the expense of using vertical slits.

We noted above that the single-electron undulator radiation distribution is proportional to $\sqrt{\lambda/L_u}$. Hence, one might think that more compression could be achieved for shorter radiation wavelengths and longer undulators. However, the degree to which this is useful is reduced greatly by the emittance degradation. With the large emittance degradation we see when using a strong chirp, there is not a sufficient benefit compared to the option of using two 2.4-m devices in a canted configuration to supply two independent beam lines [12].

V. LIFETIME ISSUES

As just seen, in the simplest understanding of Zholents's concept, it is beneficial to increase the chirp either by increasing the rf voltage V or the rf harmonic h or a combination of both. This allows compression to shorter pulse duration. We also noted that emittance degradation may result in a point of diminishing returns. In addition to emittance degradation, we must consider the possible effect on the beam lifetime.

Tracking gives us the rms vertical beam size due to the combination of the chirp and the vertical emittance. This is shown in Fig. 5 using as an example $h = 1$ with $V = 6$ MV. Not surprisingly, we see that between the deflecting cavities, the projected beam size is much larger than elsewhere.

Using, for the moment, a linear approximation for the chirp, the chirp contribution to the beam size at any point can be written as

$$\Sigma_y = \frac{hV\omega_m\sigma_t}{E} F(s), \quad (18)$$

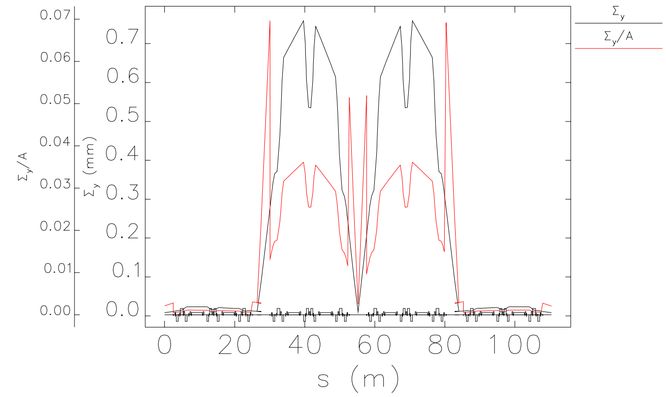

FIG. 5. (Color) Projected rms vertical beam size for the APS with $h = 1$ and $V = 6$ MV, along with the ratio of the beam size to the vertical half-aperture.

where $\omega_m/(2\pi) = 351.9$ MHz is the ring main rf frequency, $\sigma_t = 40$ ps is the rms bunch duration, and E is the beam energy. The shape function $F(s)$ is related to the beta function and phase advance. Figure 5 shows a plot of function $\Sigma_y/A(s)$, where the aperture is $\pm A(s)$, $h = 1$, and $V = 6$ MV. Like most light sources, APS has the smallest apertures in the insertion device straight sections.

Because the vertical beam size is dominated by the chirp, which is related to the bunch duration, we can analyze this as a quantum lifetime limit for the *longitudinal* plane. First, we must find the aperture limiting location, which is the point s_l at which $\Sigma_y/A(s)$ is maximized. We see that this quantity spikes at the end of the straight section downstream of the first cavity and at the start of the straight section upstream of the second cavity, with a value of 0.0696. This is related to the asymmetric placement of the cavities, which are at the centers of the upstream and downstream halves of the straight sections.

Given that the limit is at the end of the undulator straight section opposite the cavity, we can derive the limit analytically. The distance from the cavity center to the far end of the extrusion is $D = 3.682$ m. Since the beam size there is dominated by the chirp, it is given by

$$\sigma_y = \frac{hVD\omega_m\sigma_t}{E}. \quad (19)$$

This gives 279 μm. Given that the half-aperture is $A = 4$ mm, we estimate $(\sigma_y/A)_{\max} = 0.0698$, essentially identical to the result from ELEGANT.

The limit for the time deviation for a particle that makes it through the chamber is related to A and D by

$$\frac{A}{D} = \frac{hV\omega_m\Delta t_{\text{limit}}}{E}, \quad (20)$$

giving

$$\frac{\Delta t_{\text{limit}}}{\sigma_t} = \frac{EA}{hVD\omega_m\sigma_t}. \quad (21)$$

The quantum lifetime is [13]

$$\tau_q = \frac{1}{2}\tau_\delta \frac{e^\xi}{\xi}, \tag{22}$$

where $\tau_\delta = 4.75$ ms is the longitudinal damping time and, in this case,

$$\xi = \frac{\Delta t_{\rm limit}^2}{2\sigma_t^2}. \tag{23}$$

The APS lifetime is dominated by Touschek scattering and is normally seven hours in 24-bunch mode. If we want to keep the lifetime above six hours, we need $\tau_q \geq 42$ h. This requires $\xi \geq 21$, or

$$\Delta t_{\rm limit} \geq 6.5\sigma_t. \tag{24}$$

We would like to estimate the maximum value of hV we can tolerate. Using Eqs. (20) and (24), we find

$$\frac{hV}{A} \leq 3.33 \text{ MV/mm}. \tag{25}$$

For the actual value of $A = 4$ mm, we get $hV \leq 13.3$ MV, which is well short of the 38 MV we would need to reach below 1 ps FWHM.

In this analysis, we have assumed a linear chirp. This is actually a very pessimistic assumption insofar as lifetime is concerned. In reality, the sinusoidal nature of the chirp limits the number of particles that get large deflections. Because of this, there is a lifetime advantage to achieving higher hV by using primarily higher h, because electrons with large time deviations get smaller kicks than our linearized analysis assumes. Thus, one can use a much higher voltage slope than might be assumed, without hurting the lifetime.

We found above that we needed to preserve beam out to $\Delta t_{\rm lim} = \pm 6.5\sigma_t$ in order to get adequate lifetime with $hV = 3.33$ MV and linearized rf. The kick for the extremal particles in this case is 7.65 MV. Hence, we can preserve or improve on the previous lifetime requirement if

$$V \max[\sin(2\pi h f_m \Delta t)] \leq 7.65 \text{ MV}, \tag{26}$$

where $|\Delta t| \leq 6.5\sigma_t$ and $f_m = 351.9$ MHz is the ring rf frequency. As h increases, we reach a point at which the maximum value of the sine function is 1 for $|\Delta t| \leq 6.5\sigma_t$. This occurs when $h > 3$. Hence, $V = 7.65$ MV is acceptable for all $h > 3$. This means that once $h > 3$ the voltage slope can be increased arbitrarily by increasing h without any negative impact on lifetime. Of course, higher h means a less linear chirp in the photon beam, the implications of which are discussed in Sec. X.

Simulations have shown that using 7 MV actually results in small beam losses, a result of the blowup of the vertical emittance. Hence, we have limited ourselves to 6 MV in our calculations, since no chirp-related beam losses are observed in the simulations at this level. In fact, as we will see below, we will most likely not operate above 4 MV. Hence, we expect no issues with beam lifetime. (In simulation, some losses are observed due to scattering of particles outside the rf bucket by quantum excitation. This is a result of using an artificially low rf voltage in the simulations, the reasons for which are discussed in the next section.)

VI. TRACKING METHODS

While we have provided analytical estimates of many effects, tracking is necessary both to confirm the analytical estimates and go beyond simplifying assumptions. All tracking used the program ELEGANT, which performs tracking of 6D phase space, with simulation of basic optical elements such as bending magnets, quadrupoles, and sextupoles available via matrices or symplectic integration. Given that the bending radius of the APS dipoles is quite large (38 m), we have used first-order matrices for the dipoles, which makes tracking faster. Quadrupoles and sextupoles were modeled with a fourth-order integrator [14] using the exact Hamiltonian, except in cases where we explicitly needed to turn off higher-order effects.

Accelerating rf cavities were modeled using ELEGANT's RFCA element, which has exact time dependence. APS has four sets of four cavities in four different straight sections. For simplicity in longitudinal matching, we have modeled all the cavities using a single zero-length element at the start of the ring.

As mentioned, the rms bunch duration in the APS is approximately 40 ps. In normal operation with 100 mA in 24 bunches, it is lengthened from the nominal 20 ps by potential well distortion [2]. This presents a difficulty for the tracking, since it seems to force us to track with impedances. While ELEGANT supports this (see for example [15]), using this feature would make time-intensive simulations even more computationally expensive.

The nominal total rf voltage for the APS is 9 MV. In order to get the desired bunch duration without the computational overhead of simulating the longitudinal impedance, we could use an artificially low rf voltage that provides the desired bunch duration. Of course, using this work-around means the coherent synchrotron tune is artificially low in the simulations. The nominal calculated (and measured) coherent synchrotron tune is 0.0077, or 2.1 kHz. The coherent synchrotron tune in these simulations, 0.0039, is about half what it should be. However, the single-particle synchrotron tune (what each particle sees as it oscillates inside a bunch with a stable distribution) should be about right. We are essentially using the lower rf voltage to model the effect of the potential well distortion, which results from the whole bunch, on the motion of individual particles.

Originally, we assumed that it was not necessary to model synchrotron radiation (i.e., damping and quantum excitation) for these simulations. The reason is that, if one tracks without synchrotron radiation, one sees that the emittance growth due to the chirping cavities is quite rapid and stabilizes in about 100 turns. In contrast, the transverse and longitudinal damping times are, respectively, 2600 and 1300 turns. Therefore it was thought that synchrotron radiation effects could be ignored. We will show that this assumption is incorrect and that modeling synchrotron radiation is in fact very important.

VII. EMITTANCE DEGRADATION IN A PERFECT MACHINE

In this section, we examine various causes of emittance degradation. In particular, we look at those effects that are present even in a perfect machine. The next section examines emittance degradation due to errors.

For many calculations in this section, we use the extreme case of $h = 8$ and $V = 6$ MV. Of course, the effects will be diminished if the voltage is reduced. Also, calculations do not include synchrotron radiation effects unless stated otherwise.

A. Momentum compaction and energy spread

This effect is present even if there are no errors and no nonlinearities. Short of a very significant change to the APS optics, the momentum compaction factor is not subject to change from the present value of 2.8×10^{-4}. Similarly, the energy spread is essentially fixed at 0.096%. We can get a sense of why this might matter by computing the path length variation between the two cavities for an energy deviation of 0.096%. The value, 15μm, corresponds to a mere 0.05° of phase for $h = 8$. However, assuming $V = 6$ MV, a particle that was at exactly zero phase in the first cavity will get a kick of 0.8 μrad from the second cavity if its phase there is 0.05°. This is significant compared to the rms vertical beam divergence, which is typically 2 to 3 μrad.

It is straightforward to estimate the single-pass emittance increase due to this effect. Let the phase of the ith particle in the first cavity be $\phi_{1,i}$. The particle gets a vertical kick

$$\Delta y'_{1,i} = \frac{V}{E}\sin\phi_{1,i}. \tag{27}$$

If the time of flight to the second cavity is the same for all particles, the phase in the second cavity is $\phi_{2,i} = \phi_{1,i} + n\pi$. However, the time of flight depends on the fractional momentum offset δ_i via the momentum compaction factor α_c. Since the cavities are two sectors apart, the time-of-flight differential is

$$\Delta t_i = \frac{2}{40}\alpha_c\delta_i T_0, \tag{28}$$

where $T_0 = 3.68$ μs is the revolution time and 2/40 is the ratio of the number of lattice cells between the cavities to the total number of cells. At the second cavity, the ith particle gets a canceling vertical kick of

$$\Delta y'_{2,i} = -\frac{V}{E}\sin(\phi_{1,i} + \omega\Delta t_i), \tag{29}$$

where ω is the angular frequency of the deflecting cavity. Assuming $\phi_{1,i} \ll 1$ and $\omega\Delta t_i \ll 1$ the net vertical kick is simply

$$\Delta y'_i = -\frac{V\omega\Delta t_i}{E}. \tag{30}$$

The rms slope error is

$$\sigma_{\Delta y'} = \frac{V\omega\sigma_{\Delta t_i}}{E}. \tag{31}$$

For $V = 6$ MV and $h = 8$, this works out to 0.75 μrad. For comparison, the rms vertical angular spread for the lattice we are considering is 2.2 μrad. Since the slope error is unrelated to the unperturbed slope, these add in quadrature and so the emittance increases according to

$$\frac{\Delta\epsilon_y}{\epsilon_y} = \frac{\sqrt{\sigma_{y'}^2 + \sigma_{\Delta y'}^2}}{\sigma_{y'}} - 1. \tag{32}$$

This works out to 5.7%, which agrees reasonably well with tracking results for linear transport, the latter giving 6.3%. The difference is presumably due to using the linear approximation for the chirp in the analysis. This increase does not strictly build up turn after turn, as we see in Sec. VII C.

B. Chromaticity and energy spread

We will see in Sec. VII E that it is somewhat advantageous to turn off the sextupoles between the cavities in order to avoid large emittance growth. In this case, there is no chromatic correction between the cavities. The natural chromaticity of the APS in the vertical plane is $\partial\nu/\partial\delta = -43$. Hence, the betatron phase error at the second cavity of an oscillation that starts at the first cavity is

$$\Delta\nu_i = -43\delta_i\frac{2}{40} = -2.15\delta_i. \tag{33}$$

If the kick from the first cavity was y'_1, then the residual after the second cavity would be

$$y_2 = \beta y'_1\sin 2\pi\Delta\nu_i, \tag{34}$$

$$y'_2 = y'_1(1 - \cos 2\pi\Delta\nu_i), \tag{35}$$

where $\beta \approx 5$ m is the common beta function at the two locations. The error in y'_2 is second order in δ, so we ignore it. The rms value (over all particles) of the residual oscillation amplitude y_2 is

$$\langle y_2^2\rangle^{1/2} = \frac{4.3\pi\beta V\omega\sigma_\delta\sigma_t}{E}, \tag{36}$$

which works out to 41μm for $h = 8$, $V = 6$, $\sigma_t = 40$ ps, and $\sigma_\delta = 0.096\%$. This is about 3.7 times the nominal vertical beam size of 11 μm, from which we estimate that the emittance will increase to 3.9 times the initial value.

In passing, note that the effect is worse when the beta function at the cavities is higher. For fixed emittance blowup, a higher beta function helps achieve more compression by reducing the slice divergence of the electron beam. These effects essentially cancel each other, and we do not expect a strong effect from different beta functions. As noted above, a higher beta function was one effect of the lattice option we chose, which allowed more undulators to take advantage of the chirped pulses.

We tested Eq. (36) by using ELEGANT to track with sextupoles turned off but chromatic effects from quadrupoles turned on. We tracked a single pass through the system with and without energy spread. Figure 6 shows the results. As expected, the emittance grows very significantly when there is energy spread. For 6 MV, the ratio is very close to a factor of 4, in reasonable agreement with the estimate of 3.9 made above. We see that with no energy spread, there is no emittance growth, as expected.

Also included is data with interior sextupoles on. This is worse at the highest voltage, but better at low voltage. This seems to indicate that in some cases leaving the sextupoles on is better, but tracking a single pass does not tell the whole story.

C. Choice of the vertical tune

The vertical tune for the APS is normally $\nu_y = 19.27$. Certain fractional tunes, such as $1/4$ and $1/3$ have an advantage in that the single-turn emittance increase largely cancels after, respectively, 4 or 3 turns. To see this, we ignore the small variation in momentum offset for a particle over the course of making a few turns in the ring. With this assumption, the slope and position errors at the exit of the second cavity due, respectively, to chromaticity and momentum compaction, receive identical increments on each turn.

Consider a pure incremental position error of Δy_0 on each turn and a fractional tune of $1/4$. After four turns, the total position error is

$$\Delta y = \Delta y_0(1 + \cos 2\pi\nu_y + \cos 4\pi\nu_y + \cos 6\pi\nu_y) \quad (37)$$

$$= \Delta y_0(1 + 0 - 1 + 0) \quad (38)$$

$$= 0. \quad (39)$$

A similar result holds for pure slope errors and by superposition for arbitrary combinations of the two. Hence, it appears that having a fractional vertical tune of 19.25 would be advantageous for the APS.

Of course, exact cancellation does not in fact occur because the momentum deviation varies turn by turn and the vertical-plane motion is subject to alteration by the chromaticity and sextupole nonlinearities of the rest of the ring. Figure 7 shows the emittance as a function of turn for several different tunes, from tracking with ELEGANT. Tracking shows that a fractional tune of $1/3$ is better than $1/4$. The reason is that in the former case, cancellation occurs after fewer turns. Also, during the time before cancellation occurs, the particle amplitudes are smaller. Unfortunately, a fractional tune near $1/3$ is not practical in APS due to lifetime issues. Surprisingly, a tune of 19.25 is worse than a tune of 19.27, again presumably because of the larger amplitudes experienced in the former case. We have used the nominal 19.27 tune in subsequent simulations.

One of the assumptions of the above analysis is that the momentum deviation varies little over three or four turns. We see in the figure that the emittance is gradually increasing and that cancellation is not perfect. One reason for this is variation in the momentum due to synchrotron oscillations.

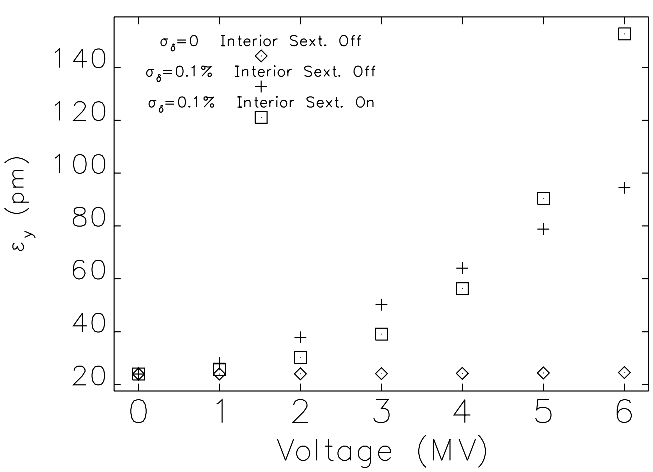

FIG. 6. Vertical emittance after a single pass through the system as a function of deflecting voltage for $h = 8$. With interior sextupoles off, tracking with and without momentum spread illustrates the impact of the natural chromaticity.

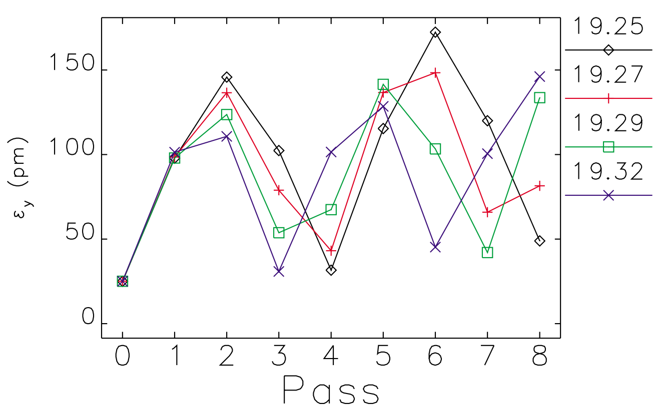

FIG. 7. (Color) Vertical emittance on successive passes in the ring for various vertical tunes, with 6 MV and $h = 8$. Interior sextupoles are off.

D. Synchrotron oscillations and quantum excitation

Clearly if the momentum offset of a particle varies significantly on the scale of a few turns, then the cancellation of the emittance increase will be incomplete. However, for a linear machine, if synchrotron oscillations are regular and symmetric in momentum deviation, we should see complete cancellation of the emittance increase after one synchrotron period. Figure 8 shows an example of tracking that illustrates what actually happens. We see that the emittance has a distinctive pattern that is related to synchrotron oscillations, but that the transport is not sufficiently linear for the beam to completely recohere.

Since synchrotron motion is important, one might wonder if quantum excitation of the energy deviation of a particle is also important, and indeed it is. Here we are concerned not about changes in the rms energy spread, which is of course constant, but rather about noise injected into the synchrotron motion of an individual particle by quantum excitation.

The dipoles in the APS [16] have a bending radius of $\rho = 38.1$ m and a bending angle of $\theta = \pi/40$. The rms change in the momentum offset of individual electrons in making a single turn in the APS is [17]

$$\langle(\Delta\delta)^2\rangle = 1.44 \times 10^{-27}\frac{\gamma^5\theta N_d}{\rho^2}, \tag{40}$$

where $\gamma = 1.37 \times 10^4$ is the relativistic factor and $N_d = 80$ is the number of dipoles. This works out to 3.0×10^{-9}, or an rms change of 5.5×10^{-5} in a turn. This is about $1/17$th of the rms momentum spread, so we estimate that in about $17^2 \approx 300$ turns, the typical electron's momentum deviation will have been largely randomized. Recalling that the single-particle synchrotron period is about 250 turns, we have reason to expect that quantum excitation will significantly reduce the extent to which the emittance degradation cancels after one synchrotron oscillation. We will show shortly that this is indeed the case.

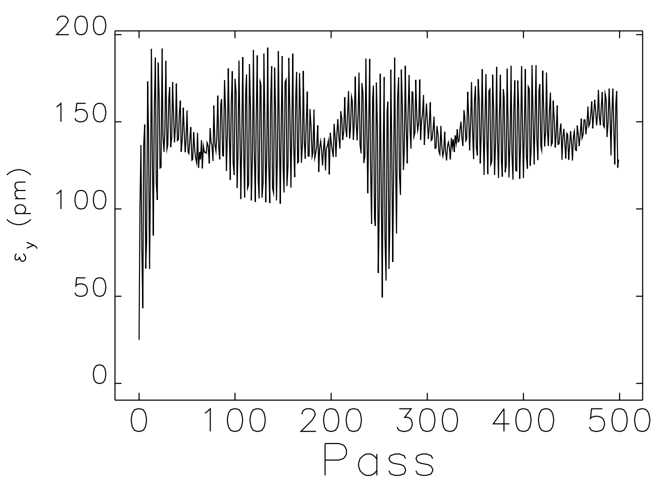

FIG. 8. Vertical emittance on successive passes in the ring, with 6 MV and $h = 8$, showing the partial recovery of the emittance after a synchrotron oscillation (about 250 turns). Interior sextupoles and synchrotron radiation effects are off.

E. Sextupole nonlinearities

We noted above that single-turn emittance growth from sextupole nonlinearities was sometimes better and sometimes worse than emittance growth from uncorrected chromaticity (i.e., with interior sextupoles off). We have just seen that quantum excitation may be expected to have a serious impact on the emittance growth when the interior sextupoles are off. At this point, it is unclear what to do with the interior sextupoles.

To further explore these trade-offs, we need to track many turns with interior sextupoles on and off. In both cases, the chromaticity of the ring was adjusted to the standard values of 5 in both planes. The results, when quantum excitation is excluded, are shown in Fig. 9. We see that over many turns, having interior sextupoles on is much worse. Even when the single-turn growth is worse for the interior-sextupoles-off case, it is better in the long run. This is in part a result of greater partial cancellation of the emittance growth on successive turns when interior sextupoles are off. Although cancellation will occur when interior sextupoles are on, as Fig. 10 shows, it is not as effective because the motion is distorted by large amplitudes inside sextupoles.

As seen in Fig. 11, having interior sextupoles on also has dramatic deleterious effects on the horizontal plane. The sextupoles cause large horizontal kicks due to the large vertical amplitudes that occur between the cavities. It is quite clear that, in the absence of synchrotron radiation effects, leaving interior sextupoles on would restrict us to relatively small deflecting voltages.

Including synchrotron radiation effects reduces the differences between these two cases dramatically. First, quan-

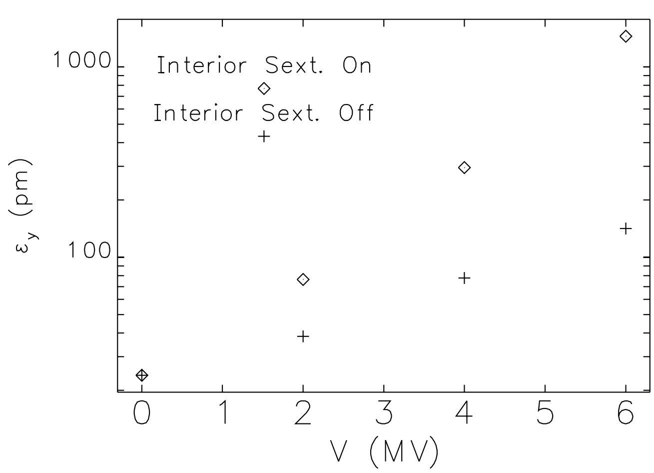

FIG. 9. Eventual vertical emittance as a function of deflecting voltage for $h = 8$, with interior sextupoles on and off, without synchrotron radiation effects.

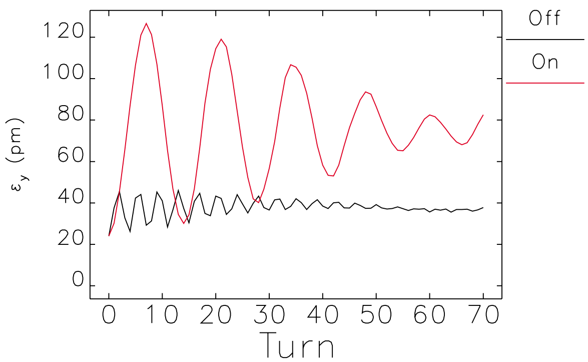

FIG. 10. (Color) Vertical emittance on subsequent turns for $V = 2$ MV and $h = 8$, with interior sextupoles on and off, without synchrotron radiation effects.

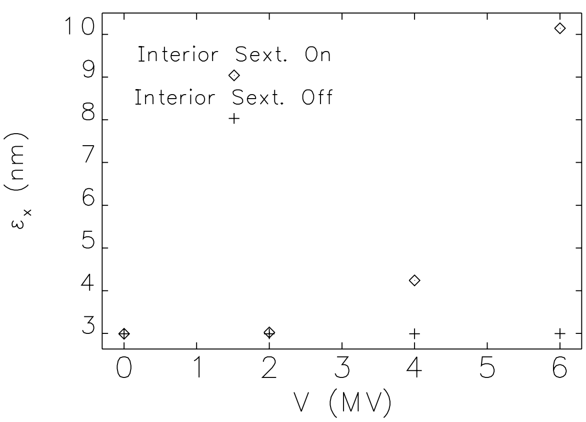

FIG. 11. Eventual horizontal emittance as a function of deflecting voltage for $h = 8$, with interior sextupoles on and off, without synchrotron radiation effects.

tum excitation results in long-term growth in the vertical emittance if the interior sextupoles are off. The reason for this was discussed at the end of the previous section. Second, radiation damping reduces the severe blowup due to nonlinearities in the case where interior sextupoles are on. This is shown in Figs. 12 and 13. The case with interior sextupoles off is still advantageous, particularly for 6 MV, where the sextupoles-on case suffers from large horizontal emittance growth. Even for 4 MV, the sextupoles-off case has a slight advantage.

Since the emittance degradation is quite severe for 6 MV and $h = 8$, we will find (in Sec. X) that we get little or no improvement in compression from 6 MV compared to 4 MV. Hence, most of the remaining analyses we used 4 MV and $h = 8$. We have chosen the case with interior sextupoles off because it has a marginal advantage and because we also expect that it will show less sensitivity to errors. Tracking and experiments show that turning off two sectors' worth of sextupoles still gives acceptable dynamic aperture. Figure 14 shows the expected vertical emittance as a function of rf voltage for several possible harmonic numbers.

F. Optimization of sextupoles

In light of these results, it may seem plausible that appropriate adjustment of the interior sextupoles will provide an advantage over the sextupoles-off case. We took a direct approach to this and used the optimization capability of ELEGANT to minimize the emittance blowup in both planes for a single pass through the system. This was done by adjusting the interior sextupoles, which were grouped into seven families in a symmetric fashion.

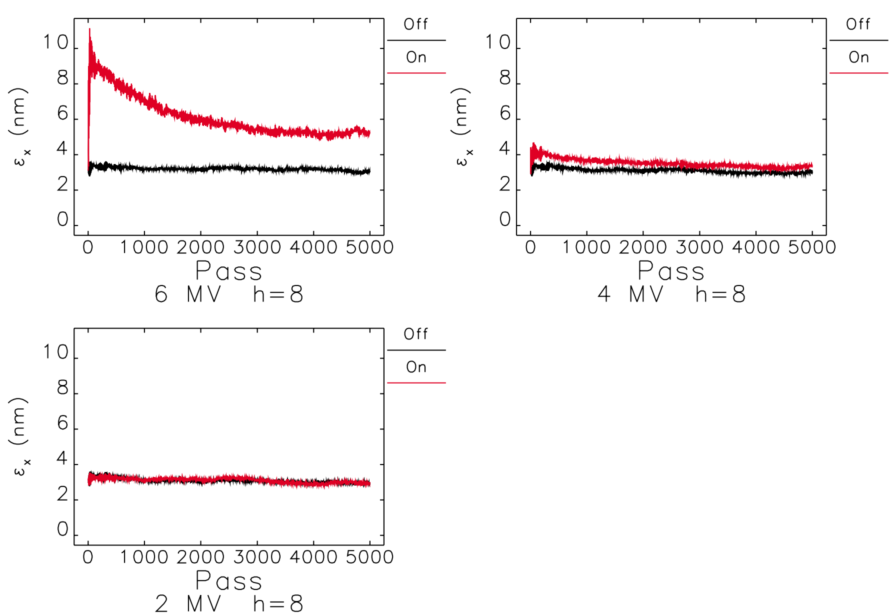

FIG. 12. (Color) Horizontal emittance vs pass including synchrotron radiation effects, with interior sextupoles on and off.

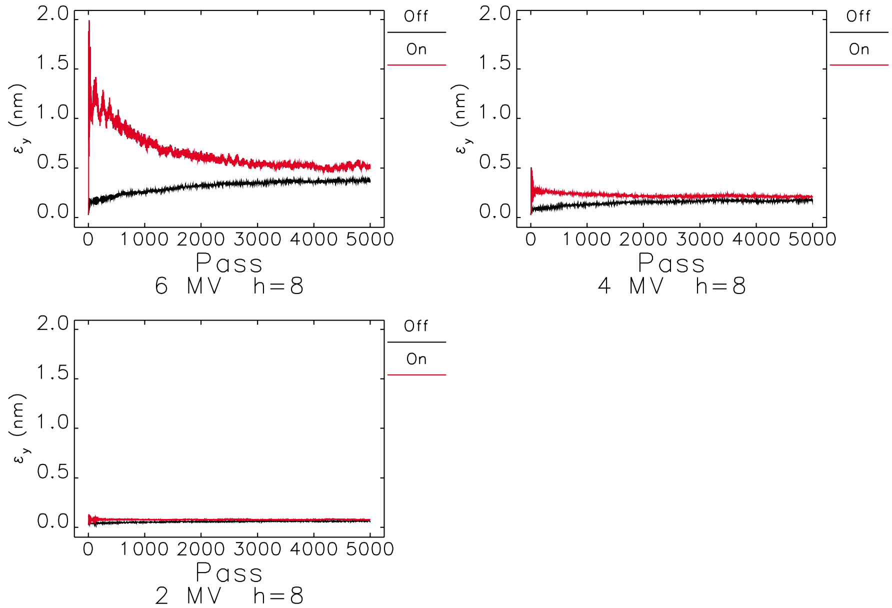

FIG. 13. (Color) Vertical emittance vs pass including synchrotron radiation effects, with interior sextupoles on and off.

Since the APS has independent power supplies for all magnets, this can be easily done in practice.

The minimization was very effective and gave a configuration that showed essentially no emittance growth when inserted into the full lattice, even for 6 MV and $h = 8$. Unfortunately, the changes to the interior sextupoles were quite large—in some cases changes of sign were needed—and as a result the dynamic aperture was dramatically smaller than normal and completely unworkable for operation of the APS. While it is still possible that an intermediate solution exists with reduced emittance growth and workable dynamic aperture, we elected to defer this investigation and have used the interior-sextupoles-off case in the remainder of our studies.

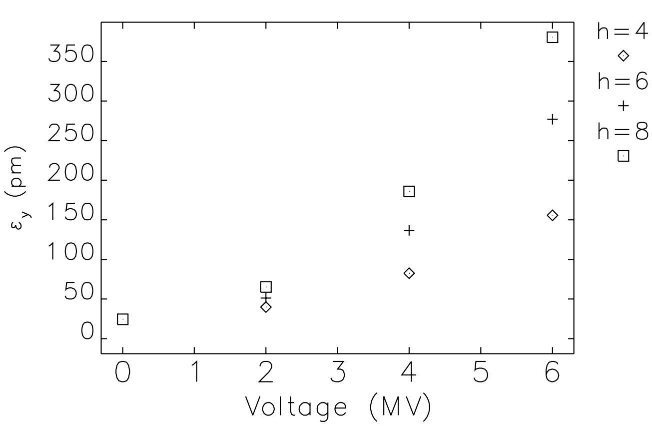

FIG. 14. Vertical emittance as a function of deflecting voltage for various harmonic numbers, with interior sextupoles off and including synchrotron radiation effects.

VIII. EFFECTS OF ERRORS

In the previous section, we reviewed various emittance-degrading effects that occur even in the perfect machine. In this section and the next, we look at effects that arise in a real machine with errors. Simulations included synchrotron radiation effects and started from the equilibrium results of tracking in the error-free machine.

To start these investigations, we ran ELEGANT for 10 000 turns with 1000 particles to find the equilibrium distributions for integer values of h for $4 \leq h \leq 8$, and voltages of 2, 4, and 6 MV. We took the start-of-ring distributions for the final turns of each simulation as the starting point for runs with errors. We tracked for 5000 turns in order to reach equilibrium, which is necessary since in the context of the simulation the error is instantaneously turned on. If this is not included, some sensitivities will be seriously overestimated.

Because of the extreme emittance growth seen for 6 MV in the last section, such a high voltage may be considered impractical. Also, we will see in Sec. X that it there is little benefit from 6 MV. Hence, we used only the 4-MV case for evaluating errors.

A. Lattice errors

Lattice errors should be readily controlled by the lattice correction methods [18] used at the APS. However, it is helpful to know how well the correction needs to perform. One concern is phase-advance error between the cavities. Another is differences in the beta functions at the cavities. If either type of error is present, we do not have an I or $-I$ transfer matrix between the cavities, and thus the cancellation is spoiled. We will look at each type of error separately.

Pure beta-function difference errors are readily simulated in ELEGANT using a simple trick. Specifically, we can apply a linear transformation to the vertical phase space at the entrance of the second cavity and apply the inverse transformation at the exit of the cavity. The first transformation has the form

$$T = \begin{pmatrix} M & 0 \\ 0 & 1/M \end{pmatrix}. \tag{41}$$

Applying the matrix T increases the beam size by the factor M while decreasing the divergence by the factor $1/M$. This is equivalent to multiplying the beta function by the factor M^2. Applying T^{-1} obviously undoes this.

Figure 15 shows the results for $V = 4$ MV and $h = 8$. We conclude that the beta-function error should be kept under about 1%. This is just on the edge of present capabilities [18]. One method [19] of dealing with residual beta-function errors is to adjust the relative voltage levels of the two cavities.

Phase-advance errors were simulated by another expedient method: we simply moved the cavities closer together or farther apart. If this is done symmetrically, by moving both cavities, the beta functions are still identical at the two locations and only the phase-advance changes. Figure 16 shows the results for $V = 4$ MV and $h = 8$. We see that errors must be larger than about 0.001 to be significant. It is possible to reproduce the tune of the entire APS to below this level [19], so this appears not to be a concern.

In passing, we note that the minimum of the curve is slightly offset from phase advance of 2π. This presumably results from our use of canonically integrated quadrupoles, which do not give exactly the phase advance expected from the lattice functions.

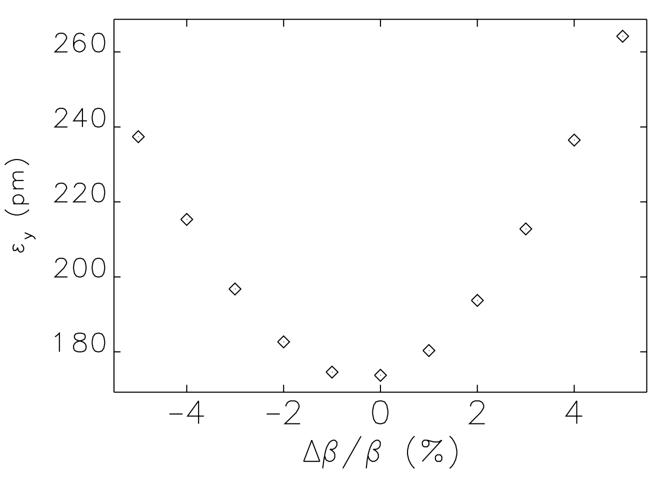

FIG. 15. Eventual vertical emittance as a function of beta-function difference between the two crab cavities, for $h = 8$ and $V = 4$ MV.

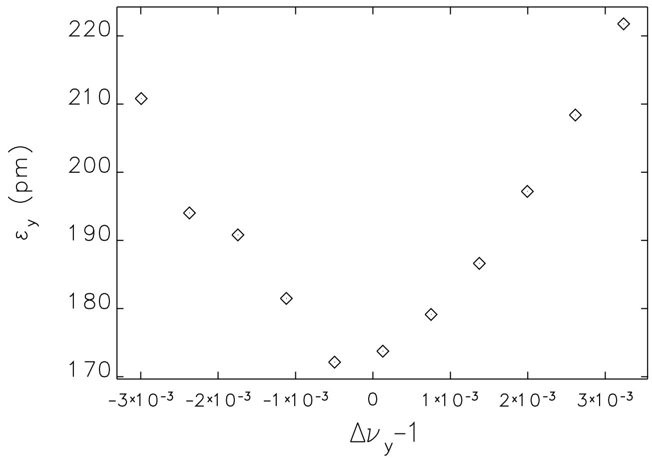

FIG. 16. Eventual vertical emittance as a function of betatron phase advance between the two crab cavities, for $h = 8$ and $V = 4$ MV.

B. Cavity roll

Cavity roll is of course a very easy effect to evaluate, since we can simply roll the cavities in ELEGANT. Two types of roll are relevant. First, we may have common roll of the field in both cavities, perhaps due to common construction errors, common coupling loop geometry, and so forth. The problem here would be that the phase advance in the horizontal plane is not a multiple of π, and hence we will get no cancellation. Second, we may have relative roll of the field in the two cavities, resulting not only in kicks in the horizontal plane, but perhaps more importantly, imperfect cancellation in the vertical plane.

To assess these issues, we used $h = 8$ and $V = 4$ MV for two sets of tracking runs. In the first, we rolled both cavities by equal amounts of up to 4 mrad. In the second, we rolled only the second cavity, again, by up to 4 mrad. Figure 17 shows the impact on the horizontal emittance, which is very small. The impact on the vertical emittance

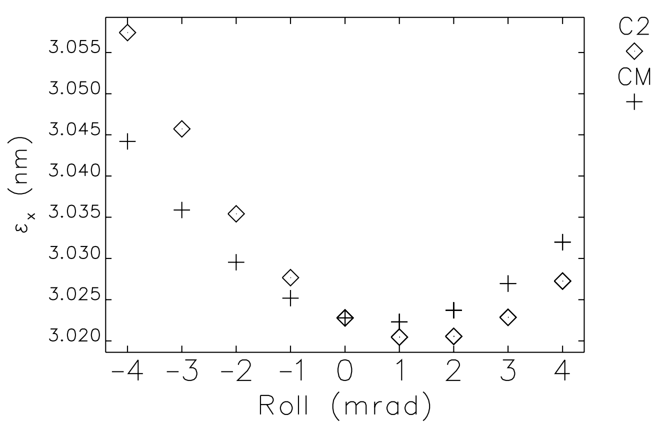

FIG. 17. Eventual horizontal emittance as a function of cavity roll for $h = 8$ and $V = 4$ MV. The "CM" data is for common-mode roll while the "DM" data is for differential-mode roll.

for roll of up to 4 mrad is negligible, and we do not bother to show it. From this we conclude that roll is not an issue, since it is easy to keep it under a few milliradians with present-day alignment techniques.

C. Lattice coupling

Lattice coupling is less straightforward to evaluate in that there is no single parameter that we can vary to parametrize it. We are concerned with rolled quadrupoles and sextupoles between the two cavities. (Rolled dipoles are not an issue as they only affect the orbit and do not couple betatron oscillations.) In the arrangement shown in Fig. 3(b) we would have only one skew quadrupole between the cavities. This is not sufficient to correct for the many coupling sources that might exist between the cavities. Hence, we must assess the effect of coupling by simulating the effect of random rolls of elements between the cavities, assuming that no correction is possible. We used the 0.25-mrad rms roll tolerance from the APS construction [16]. Of course, these rolls result in global coupling which will increase the vertical emittance irrespective of the presence of the crab cavities. Hence, we looked at the increase in vertical and horizontal emittances between the cavity-on and cavity-off cases.

We simulated 50 seeds with the crab cavities on and off. Histograms of the increase in the vertical and horizontal emittances due to having the cavities on are shown in Fig. 18. We see that the increase in the vertical plane is, perhaps not surprisingly, about the same as without errors. This reflects the existing, large increase due to other effects. The distribution of the increase in the horizontal emittance shows that the increase should be less than 10%, which is manageable. While it may be prudent to add skew quadrupoles between the cavities in order to control coupling from the vertical into the horizontal plane, it does not seem to be a serious issue.

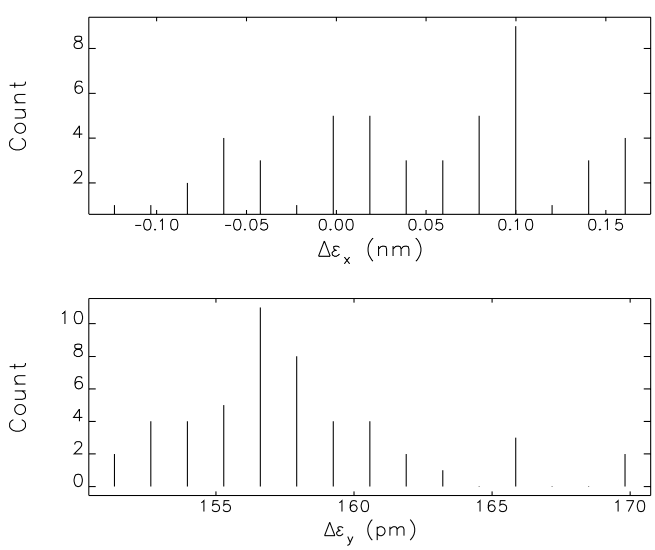

FIG. 18. Histograms of the increase in the eventual horizontal and vertical emittances for $h = 8$ and $V = 4$ MV with 0.25 mrad rms rolls of quadrupoles and sextupoles between the crab cavities.

IX. RF TOLERANCES

In this section, we investigate the relationship between rf errors and emittance growth, which allows us to specify some tolerances. Although vertical emittance growth is the main concern, there are other possibilities. Among these are variation in the length, arrival time, and pointing of the x-ray pulse. One might also see reduced lifetime in extreme cases.

We begin by enumerating some of the possible sources of error.

(i) Voltage difference between the crab cavities. This can result from variations in amplifiers and cavity tuning.

(ii) Beam phase error between the two crab cavities, which may be due to the following: (a) Actual phase error between the cavities, which has many possible sources, including variations in amplifiers, cavity tuning, source frequencies, cable temperature, etc., (b) Time-of-flight error between the crab cavities due to changes in the energy offset of the beam. These occur as the ring stretches under tidal forces, as ID gaps are varied, and also due to phase and voltage ripple in the main rf system.

(iii) Phase error between crab cavities and the main rf system. As in the previous point, there are two possible causes: (a) Actual phase error. (b) Time-of-flight variation due to energy offset variation.

Frequency error is not called out separately from phase error. It is assumed that the main and crab rf frequencies are derived from a common source. Thus, there should be no persistent frequency error and hence no phase slewing. Small frequency deviations may exist but will average to zero, manifesting themselves as phase errors.

For some of these errors, one can estimate the emittance growth for a single pass through the system. However, as we saw above, the accumulated effect over many turns may be greater or less than the single-pass effect. Hence, we use tracking studies exclusively in this section, with $h = 8$ and $V = 4$ MV.

A. Crab-cavity voltage errors

We start with voltage errors in the crab cavities as these are the simplest. The first study consists of scanning the voltage of the second cavity while holding the voltage of the first cavity fixed at 4 MV. This simulates a voltage error of varying magnitude that is essentially fixed on the scale of several damping times. The results, in Fig. 19, show that the impact on the vertical emittance for relative voltage errors of a fraction of a percent is modest. Requiring a voltage error of under 0.5% seems prudent.

Another type of error is voltage modulation, in which one cavity may vary in a time-dependent fashion relative to

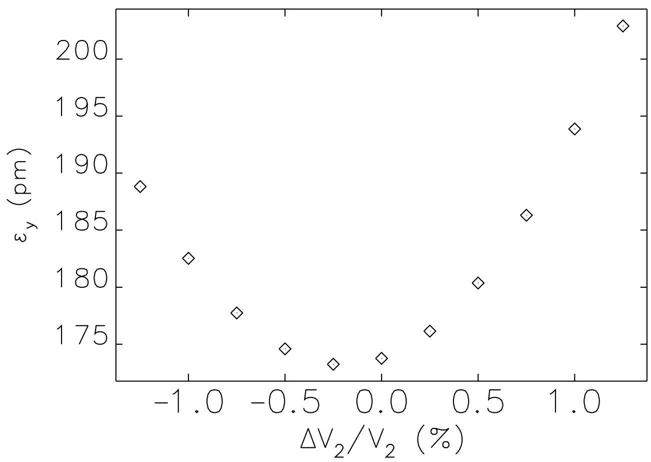

FIG. 19. Variation of the vertical emittance for $h = 8$ and $V = 4$ MV with voltage error in the second cavity.

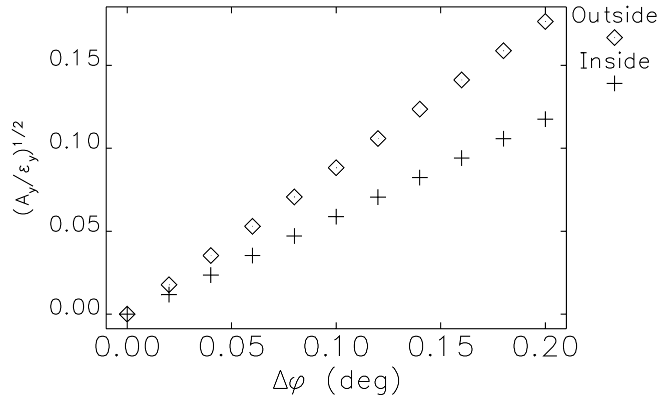

FIG. 20. Beam centroid normalized to beam size as a function of differential phase error between the crab cavities, for $h = 8$ and $V = 4$ MV.

the other, on a time scale that is comparable to or shorter than a damping time (4.8 ms for the longitudinal plane). The most worrisome modulation frequencies are presumably those that are multiples of the synchrotron frequency. Simulating this will require an upgrade to ELEGANT to support modulation of transverse rf cavity voltage, and has not been undertaken yet.

B. Inter-crab-cavity phase errors

In the previous section, we did a straightforward sweep of the voltage of one cavity with the voltage of the other cavity held fixed. This is useful because it is only the difference between the two cavity voltages that matters to the cancellation. For phase, we need to worry about both differential and common-mode phase errors. Differential-mode phase errors, where the opposite phase error is added to the cavities, affect the cancellation of the rf kicks, whereas common-mode phase errors do not. Both types of errors can change the beam position, although in the case of common-mode errors, the changes are restricted to the space between the cavities. If the phase error is relatively constant on the scale of the damping time, the closed orbit will change since the beam centroid is kicked by the same amount each turn. This can presumably be taken care of by standard orbit feedback techniques [20], at least for changes within the system bandwidth.

We find that the effect on the vertical emittance is very small and do not show it. Instead, we compare the vertical orbit amplitude to the beam size, since the main effect of phase errors is a displacement of the beam centroids. The fractional change in the beam centroid relative to beam size is characterized by $\sqrt{A_y/\epsilon_y}$, where $A_y = y^2(1 + \alpha_y^2)/\beta_y + 2\alpha_y y y' + \beta_y y'^2$ is the invariant corresponding to the orbit amplitude. A_y was computed using data from both inside and outside the region between the cavities. However, for ϵ_y we use only the emittance outside the cavity region, so that the increase in the apparent emittance between the cavities due to the chirp is not included. Doing otherwise would make the effects between the cavities appear much smaller.

Figure 20 shows the results of the simulations for differential-mode errors. We see that the beam centroid shows significant changes relative to the beam size. The vertical beam stability in the APS is presently 10% of the rms beam size and divergence. To avoid making this worse, differential-mode phase errors would need to be controlled below about 0.05°.

Figure 21 shows the results for common-mode phase errors. The effects here are somewhat smaller than those for differential-mode errors. As expected, the effect is largely confined to the region between the two cavities. As with voltage errors, we must recognize that phase modulation may be an issue. Again, simulating this requires upgrading ELEGANT.

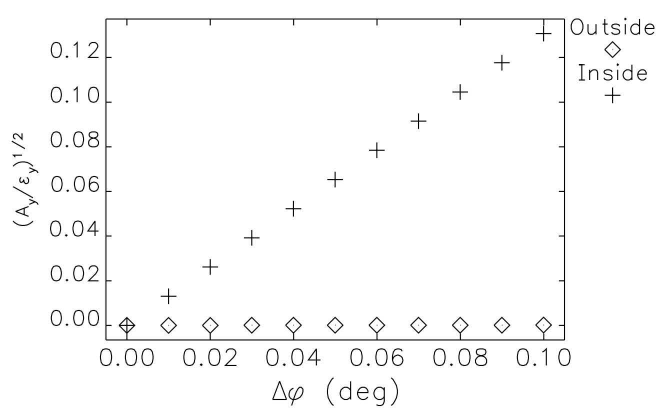

FIG. 21. Beam centroid normalized to beam size as a function of common phase error of the crab cavities, for $h = 8$ and $V = 4$ MV.

C. Time-of-flight errors

As indicated above, time-of-flight errors can result in problems because they change the relative phase of the beam in the two crab cavities. They also change the relative phase of the crab cavities and the main rf cavities.

The dominant source of time-of-flight errors is beam energy offset, which in turn is produced primarily by diurnal tidal stretching of the ring. Fortunately, at the APS this is controlled by the rf frequency feedback system. With this system, the rf frequency is adjusted based on the average orbit from about 300 BPMs that have nonzero dispersion. Assuming, conservatively, a resolution of 10 μm for each BPM and using the average dispersion function of 0.14 m, the resolution of the fractional momentum offset measurement is 4×10^{-6}. The frequency adjustment step is less than 0.1 Hz, corresponding to a change of 1×10^{-6} in δ, so we do not anticipate an issue correcting to levels much smaller than the momentum spread. Because the resolution and level of control are much smaller than the rms momentum spread, we conclude that time-of-flight errors are not a concern.

It should be pointed out that as the main ring rf frequency changes, so should the crab-cavity rf frequency in order to avoid phase slewing. The typical full range of variation is just under 30 Hz for a period of several weeks. This would correspond to 240 Hz for $h = 8$ crab cavities, or a relative change of less than 1 part in 10^7. Superconducting cavity bandwidth is less than $\Delta\omega/\omega = 1/Q \approx 1/(2 \times 10^9)$ [21], so we will probably need active tuning on the crab cavities unless there is significant external loading. Fortunately, the frequency variation is quite slow and the required range is small.

X. SIMULATION OF COMPRESSION

In Sec. IV we used a Gaussian approximation for the photon distribution and a linear approximation for the chirp, which permitted estimating the potential of this method when applied to APS. In this section, we show results of simulations that include the sinusoidal character of the chirp as well as a more accurate form of the photon distribution.

Near the axis and for exactly on-harmonic radiation, the distribution function is [11]

$$S(\theta) \approx \mathrm{sinc}^2\left(\frac{nN\pi\gamma^2\theta^2}{1 + K^2/2}\right), \qquad (42)$$

where n is the harmonic number and N the number of undulator periods. (Note that the reference omitted the γ^2 factor.) Figure 22 shows the angular distribution for $N = 73$ and $K = 0.524$, which corresponds to 1 Å radiation from APS's 2.4-m-long, 3.3-cm-period undulator A. Also shown is a Gaussian fit, which has $\sigma = 4.42$ μrad, in reasonable agreement with the value expected from Eq. (2), which gives $\sigma = 4.56$ μrad. The Gaussian clearly

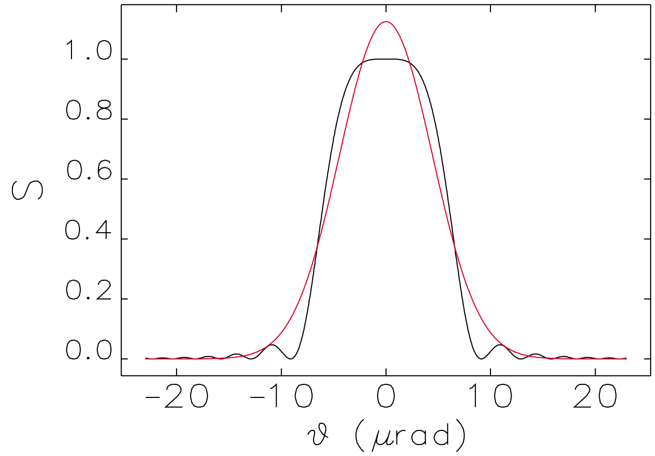

FIG. 22. (Color) The sinc-function angular distribution compared to a Gaussian fit, for 1 Å radiation from a 2.4-m-long, 3.3-cm-period undulator in the APS.

has intensity at larger angles than Eq. (42) exhibits, and we find that using a Gaussian approximation results in a slightly pessimistic analysis of the compression.

From tracking simulations, we obtain the full set of 6D coordinates for simulated electrons at the center of an undulator. In this case, we have chosen to work with the first undulator in the middle straight section (which contains two undulators, as shown in Fig. 3). As in the simulations with errors, we ran ELEGANT for 10 000 turns with 1000 particles to find the equilibrium distributions for integer values of h for $4 \leq h \leq 8$, and voltages of 2, 4, and 6 MV. We combined the start-of-ring distributions for the final 100 turns of each simulation to make a 100 000 particle distribution for each parameter set. This technique of combining multiturn data overcomes the formidable run times we would face if we tried to track more than 1000 particles to equilibrium. The large number of particles was needed in post-tracking analyses in order to reduce noise, particularly when the effect of small slits is investigated.

The electron distribution must be convolved with the photon distribution given by Eq. (42). Since the particle output from ELEGANT is in a self-describing SDDS file [22], this is readily done using postprocessing capabilities of the SDDS Toolkit [23]. In particular, we used the program SDDSSAMPLEDIST, which permits sampling general distributions. We first sampled the suitably normalized distribution $\theta S(\theta)$ to get a value of θ for each electron. We also sampled the uniform distribution on $[0, 2\pi]$ to get a value of the azimuthal angle, ϕ. The horizontal and vertical slopes for the photon relative to the emitting electron are then respectively $x'_r = \theta\cos\phi$ and $y'_r = \theta\sin\phi$. We create one (x'_r, y'_r) pair for each electron, using the transformations $x' \rightarrow x' + x'_r$ and $y' \rightarrow y' + y'_r$ to turn the electron distribution into a photon distribution, which is also stored as an SDDS file. (We could also have sampled the undulator radiation distribution more than once for each electron. This would simply have improved the statistics, which was unnecessary due to use of a large number of macroparticles.)

Because ELEGANT both reads and writes SDDS files, this photon distribution can be read back into ELEGANT in order to optimize the compression. The program does not specifically recognize these particles as photons, but if we use only drift spaces, slits, and a linear matrix representing the compression optics, it does not matter. The advantage of using ELEGANT is that it can optimize the results of tracking. These results include statistical measures of pulse duration that are more robust against outliers than the rms duration. We chose to optimize the 70-percentile duration, which is the time interval containing the central 70% of the particles. For Gaussian beams, this corresponds approximately to $2\sigma_t$ and is also comparable to the FWHM.

Our goal is to find the harmonic number, voltage, slit spacing, and compression parameter (K) that give the shortest pulse duration with the greatest transmission of photons through the slits. We are subject to two primary constraints: the harmonic number must be $4 \leq h \leq 8$ and the voltage must be less than 6 MV.

For each parameter set, we tracked 100 000 particles for ten turns, extracting coordinates at the undulator center on each turn. We then scanned the slits and computed the optimum compression for that slit spacing, using data from all ten turns together for the optimization. Having found the optimum compression parameter, we then compressed the distribution from each turn separately using that parameter to assess variation in the compression from turn to turn (it was very small). This complex and computationally intensive process was made manageable by use of the SDDS Toolkit, the Tcl/Tk scripting language, and a Linux cluster with approximately 90 processors.

Compression with asymmetric cut crystals is not loss free [6]. For the configuration under development at APS [6], the throughput from the exit of the slits to the sample is expected to be no worse that 30% over the range from 8 to

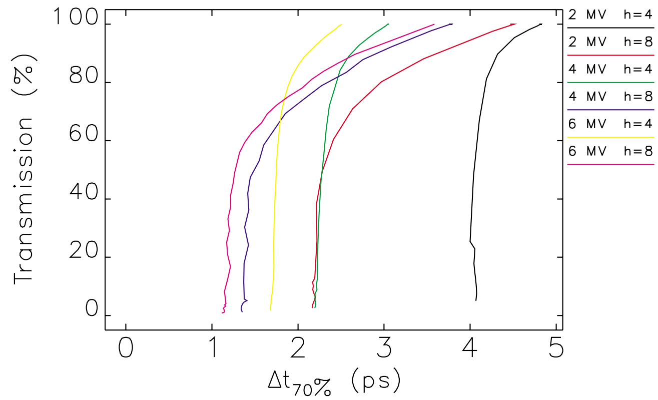

FIG. 24. (Color) Transmission through the slits as a function of the obtained duration of the central 70% of the x-ray pulse, for various deflecting voltages and harmonic numbers of 4 and 8.

18 keV. This number varies with photon energy, the choice of crystal planes, and the compression factor in a way that is related to the details of the x-ray optics design, which is beyond the scope of this paper.

Figures 23 and 24 show the results of the computations, assuming 10 keV photons (i.e., 1.2 Å radiation). The variation of the achievable pulse length with slit spacing is a result of the sinusoidal character of the deflecting waveform. For about 50% transmission through the slits, the compressed pulse duration has reached its ultimate value. The transmission for the ultimate pulse duration is slightly higher for lower harmonic values, a result of using the more linear part of the rf waveform. However, the minimum achieved pulse duration is also longer, since the chirp is less.

The final choice of operating point depends not only on the achievable pulse duration, but also on the impact on the rest of the APS users, which primarily means the impact on the vertical emittance. Figure 25 shows the vertical emittance as a function of the pulse duration achieved when 50% of the photons are allowed through the slits. We see that there is little to be gained from the case with 6 MV and

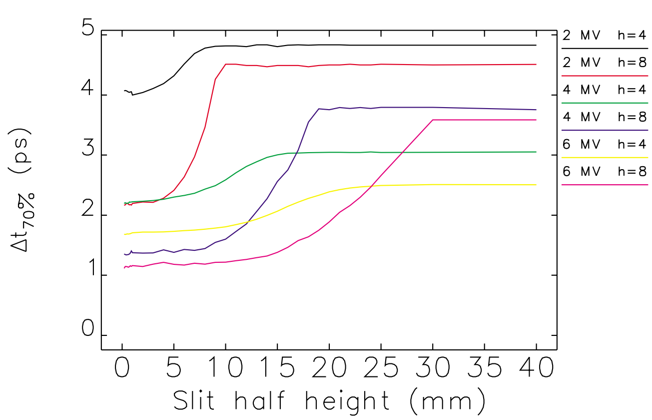

FIG. 23. (Color) Duration of the central 70% of the x-ray pulse as a function of the half-height of the slits, for various deflecting voltages and harmonic numbers of 4 and 8.

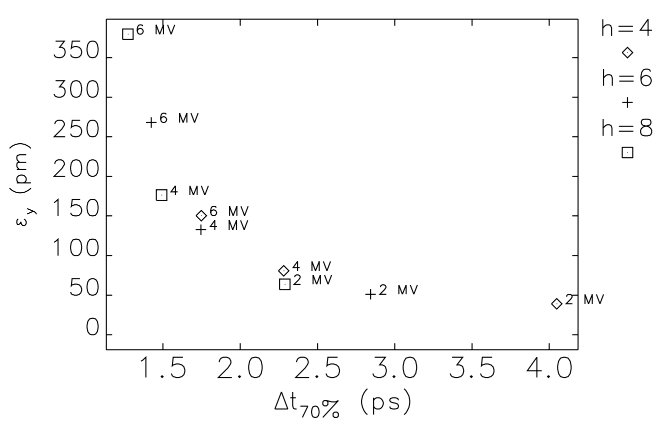

FIG. 25. Emittance as a function of achievable pulse length for constant 50% transmission through the slits.

$h = 8$, compared to 4 MV and $h = 8$. The practical minimum 70% pulse duration is about 1.5 ps. Using 2 MV is only 50% worse, and exhibits about a third the vertical emittance.

The 4-MV result is slightly better than our estimate of 1.8 ps FWHM above, for two reasons. First, for a Gaussian profile, the 70% pulse duration is expected to be about 15% less than the FWHM. Second, the actual undulator distribution is not as broad as implied by the Gaussian parameter in Eq. (2), as seen in Fig. 22.

As mentioned, Shastri's optics design shows optics throughput of no less than 30% between 8 and 18 keV. Hence, we expect our results with 50% slitting to correspond to total throughput of no less than 15%.

XI. CONCLUSION

We have investigated the feasibility from an accelerator physics standpoint of using transverse-deflecting cavities in the APS in order to allow production of short x-ray pulses. We find the concept is feasible and that pulse durations of approximately 1.5 ps may be achieved with 4-MV cavities at the eighth harmonic (2816 MHz) of the main ring rf frequency. Additional losses due to x-ray optics, not covered here, will reduce the total throughput, but total transmission of at least 15% is expected.

Several areas of concern were discovered and addressed:

(i) There will be a large increase in the vertical emittance for voltages over 2 MV. For 4 MV, the vertical emittance increases to about 7 times the normal value. The vertical emittance is still less than 6% of the horizontal emittance. This is thought to be compatible with storage ring operations and would not impact most APS users. Reducing the voltage to 2 MV virtually eliminates the emittance growth.

(ii) For 4 MV operation, we will probably need to turn off interior sextupoles to reduce emittance effects, which will adversely impact the lifetime. Experiments show that lifetime would suffer a 50% reduction, which is manageable. The vertical emittance increase will mitigate this impact. In addition, we can increase the rate at which top-up is performed to make up for the shorter lifetime. For 2 MV operation, we could leave all sextupoles on at the normal values, so that lifetime would be unaffected.

(iii) We found a tight tolerance on the intercavity rf phase error, of the order of 0.05° for 4 MV operation. Since a phase error results in an orbit distortion, it can be mitigated with a combination of orbit correction and fast feedback. This tolerance is relaxed by a factor of 2 for operation at 2 MV.

(iv) In order to achieve a rapid chirp with modest voltage, we need to use an 8th-harmonic cavity. This requires the use of slits, resulting in an approximately twofold reduction of the intensity (not including the reduction due to x-ray optics losses).

As indicated, the first three concerns can be mitigated by reducing the voltage to 2 MV. Indeed, we find that 2 MV operation would have no significant impact on normal APS operation. The trade-off is that we increase the achievable x-ray pulse length by about 50%. One option is to use 4 MV operation as a special mode for limited time periods and use 2 MV operation for all other times.

We briefly explored the option of optimizing the strength of interior sextupoles to eliminate the emittance growth and permit higher deflecting voltage. While we found that the dynamic aperture was unacceptable, it is possible that a compromise solution will be found that preserves dynamic aperture and reduces emittance growth.

Use of the chirping cavities is anticipated to require few changes to the accelerator. For example, it is compatible with existing magnets, power supplies, and vacuum chambers. Depending on the location chosen, existing insertion devices may need to be moved from one end of the straight section to the other, which is very simple. (In most APS sectors, insertion devices occupy only half of the straight section.) Most of the work involved in implementing the scheme will be related to design of the rf cavities, the rf system, the cryogenic plant, and the x-ray optics.

ACKNOWLEDGMENTS

The author wishes to acknowledge helpful discussions with and information from Y.-C. Chae, R. Dejus, L. Emery, R. Gerig, K. Harkay, D. Horan, R. Kustom, D. Mills, S. Milton, V. Sajaev, G. Waldschmidt, and M. White. I thank S. Shastri for information on intensity loss due to the asymmetric crystal compression optics. I thank R. Soliday for optimizing the performance of the Linux cluster to improve turnaround for this work. I thank H. Shang for assistance with graphics. Work supported by the U.S. Department of Energy, Office of Basic Energy Sciences, under Contract No. W-31-109-ENG-38.

APPENDIX: TM110 CAVITY MODEL

In support of this work, ELEGANT was upgraded to include the RFTM110 element, which provides a simulation of a TM110 mode. The cavity is modeled as a zero-length element.

To derive the field expressions, we start with some results from Jackson [10], Sec. 8.7. The longitudinal electric field for a TM mode is just

$$E_z = -2iE_0\Psi(\rho, \phi)\cos\left(\frac{p\pi z}{d}\right)e^{-i\omega t}, \tag{A1}$$

where p is an integer, d is the length of the cavity, and we use cylindrical coordinates (ρ, ϕ, z). The factor of $-2i$ represents a choice of sign, magnitude, and phase convention. We are interested in the TM110 mode, so we set $m = 1$ and $p = 0$. For a cylindrical cavity, we have

$$\Psi(\rho, \phi) = J_1(k\rho)\cos\phi, \tag{A2}$$

By choosing $\cos\phi$ for the aximuthal dependence, we get a magnetic field primarily in the vertical direction, giving a default horizontal deflection. (In ELEGANT, this deflection can be rotated to any desired angle using the TILT parameter.)

In MKS units, the magnetic field is

$$\vec{B} = \frac{2E_0}{kc} e^{-i\omega t} \left(\hat{\rho} \frac{J_1(k\rho)}{\rho} \sin\phi + \hat{\phi} \cos\phi \frac{\partial J_1(k\rho)}{\partial\rho} \right). \quad \text{(A3)}$$

We expanded these expressions to sixth order in $k * \rho$ for use in ELEGANT. Here, we present only the expressions to second order. Taking the real parts only, we now have

$$E_z \approx E_0 k\rho \cos\phi \sin\omega t, \quad \text{(A4)}$$

$$cB_\rho \approx E_0 \left(1 - \frac{k^2\rho^2}{8} \right) \sin\phi \cos\omega t, \quad \text{(A5)}$$

$$cB_\phi \approx E_0 \left(1 - \frac{3k^2\rho^2}{8} \right) \cos\phi \cos\omega t. \quad \text{(A6)}$$

The Cartesian components of $\vec{B}$ can be computed easily

$$cB_x = cB_\rho \cos\phi - cB_\phi \sin\phi \quad \text{(A7)}$$

$$= \frac{E_0}{4} \rho^2 k^2 \cos\phi \sin\phi \cos\omega t, \quad \text{(A8)}$$

$$cB_y = cB_\rho \sin\phi + cB_\phi \cos\phi \quad \text{(A9)}$$

$$= E_0 \left[1 - \frac{k^2\rho^2(2\cos^2\phi + 1)}{8} \right] \cos\omega t. \quad \text{(A10)}$$

It is helpful to get a feeling for the scale by picking a frequency for the cavity. A possible choice for APS is 1400 MHz (the fourth harmonic of our main rf frequency), which gives $k = 29.3\ \mathrm{m}^{-1}$. The APS dynamic aperture is less than 15 mm, giving $kx < 0.44$. The term proportional to $k^2\rho^2$ in the expression for B_y is thus about 7% of the leading term. The next term, proportional to $k^4\rho^4$, is less than 0.5% of the leading term.

The Lorentz force on an electron is $F = -eE_z\hat{z} - ec\vec{\beta} \times \vec{B}$, giving

$$F_x/e = \beta_z cB_y, \quad \text{(A11)}$$

$$F_y/e = -\beta_z cB_x, \quad \text{(A12)}$$

$$F_z/e = -E_z - \beta_x cB_y + \beta_y cB_x. \quad \text{(A13)}$$

We see that for $\rho \to 0$, we have $E_z = 0$, $B_x = 0$, and

$$cB_y = E_0 \cos\omega t. \quad \text{(A14)}$$

Hence, for $\omega t = 0$ and $E_0 > 0$ we have $F_x > 0$. This explains our choice of sign and phase convention above. Indeed, owing to the factor of 2, we have a peak deflection of eE_0L/E, where L is the cavity length and E the beam energy. Thus, if $V = E_0L$ is specified in volts, and the beam energy is expressed in electron volts, the deflection is simply the ratio of the two. As a result, we have chosen to parametrize the deflection strength simply by referring to the "deflecting voltage" V.

[1] A. Zholents, P. Heimann, M. Zolotorev, and J. Byrd, Nucl. Instrum. Methods Phys. Res., Sect. A **425**, 385 (1999).
[2] Y.-C. Chae, L. Emery, A. H. Lumpkin, J. Song, and B. X. Yang, in *Proceedings of the Particle Accelerator Conference, Chicago, Illinois, 2001* (IEEE, Piscataway, New Jersey, 2001), p. 1491.
[3] J. Feikes, K. Holldack, P. Kuske, and G. Wustefeld, in *Proceedings of the European Particle Accelerator Conference, Lucerne, Switzerland, 2004* (EPS-AG, Lucerne, 2004), p. 1954.
[4] A. Zholents *et al.*, Phys. Rev. Lett. **76**, 916 (1996).
[5] A. Zholents (private communication).
[6] S. Shastri (private communication).
[7] D. Horan (private communication).
[8] M. Borland, Advanced Photon Source LS-287, 2000.
[9] D. Li and J. N. Corlett, in *Proceedings of the 8th European Particle Accelerator Conference, Paris, 2002* (EPS-IGA and CERN, Geneva, 2002), p. 2253.
[10] J. D. Jackson, *Classical Electrodynamics* (J. Wiley, New York, 1975), 2nd ed., Sec. 8.7.
[11] K.-J. Kim, in *Physics of Particle Accelerators*, edited by Melvin Month and Margaret Dienes, AIP Conf. Proc. No. 184 (AIP, New York, 1989), p. 565.
[12] P. Den Hartog, G. A. Decker, and L. Emery, in *Proceedings of the Particle Accelerator Conference, Portland, Oregon, 2003* (IEEE, Piscataway, NJ, 2003), p. 833.
[13] See, for example, *Handbook of Accelerator Physics and Engineering*, edited by A. Chao and M. Tigner (World Scientific, Singapore, 1998), Sec. 3.1.4.6.
[14] E. Forest, in Physics of Particle Accelerators, AIP Conf. Proc. No. 184 (Ref. [11]), pp. 1106–1136.
[15] Y.-C. Chae, in Proceedings of the Particle Accelerator Conference, Portland, Oregon, 2003 (Ref. [12]), p. 3017.
[16] H. Bizek, ANL/APS/TB-26, 1996.
[17] G. Leleux, P. Nghiem, and A. Tkatchenko, in *Proceedings of the Particle Accelerator Conference, San Francisco, California, 1991* (IEEE, Piscataway, NJ, 1991), p. 571.
[18] V. Sajaev and L. Emery, in Proceedings of the 8th European Particle Accelerator Conference, Paris, 2002 (Ref. [9]), p. 742.
[19] L. Emery (private communication).
[20] F. Lenkszus, L. Emery, R. Soliday, H. Shang, O. Singh, and M. Borland, in Proceedings of the Particle Accelerator Conference, Portland, Oregon, 2003 (Ref. [12]), p. 283.
[21] G. Waldschmidt (private communication).
[22] M. Borland, in *Proceedings of the Particle Accelerator Conference, Dallas, Texas, 1995* (IEEE, Piscataway, New Jersey, 1995), p. 2184.
[23] R. Soliday, M. Borland, L. Emery, and H. Shang, in Proceedings of the Particle Accelerator Conference, Portland, Oregon, 2003 (Ref. [12]), p. 3473.
[24] V. Sajaev (private communication).