

Low-emittance tuning of storage rings using normal mode beam position monitor calibration

A. Wolski*

University of Liverpool, Liverpool, United Kingdom

D. Rubin, D. Sagan, and J. Shanks

CLASSE, Cornell University, Ithaca, New York 14853, USA

(Received 14 April 2011; published 26 July 2011)

We describe a new technique for low-emittance tuning of electron and positron storage rings. This technique is based on calibration of the beam position monitors (BPMs) using excitation of the normal modes of the beam motion, and has benefits over conventional methods. It is relatively fast and straightforward to apply, it can be as easily applied to a large ring as to a small ring, and the tuning for low emittance becomes completely insensitive to BPM gain and alignment errors that can be difficult to determine accurately. We discuss the theory behind the technique, present some simulation results illustrating that it is highly effective and robust for low-emittance tuning, and describe the results of some initial experimental tests on the CesrTA storage ring.

DOI: [10.1103/PhysRevSTAB.14.072804](https://doi.org/10.1103/PhysRevSTAB.14.072804)

PACS numbers: 29.20.-c, 29.27.-a

I. INTRODUCTION

Lepton storage rings often require tuning to achieve very low vertical emittance: in light sources, this increases the beam brightness; in colliders, it improves luminosity. A number of facilities in recent years have reported vertical emittances of a few picometers [1–4]. Achieving emittances in this range requires precise correction of machine errors that generate vertical dispersion and betatron coupling. A common technique is to perform an analysis of the measured orbit response matrix (ORM) to determine the sources of error [5]. This includes diagnostic errors such as tilts of the beam position monitors (BPMs). However, data collection for ORM analysis is often rather slow, as it involves measuring the change in the closed orbit in response to each of the horizontal and vertical orbit correction (steering) magnets in the ring. In a large ring, there may be dozens of such orbit correctors. Fitting a model to the measured data (to determine the errors present in the ring) involves numerically intensive computation, and can also be time consuming. For large rings, such as those proposed, for example, for the damping rings of the International Linear Collider [6], it may become impractical to perform ORM analysis on a routine basis to achieve and maintain the specified vertical emittance.

In this paper, we describe an alternative technique for low-emittance tuning, based on excitation of the normal modes of the beam motion. Measurement of the turn-by-turn signals on the BPM buttons during resonant excitation

allows the BPMs to be calibrated in terms of the local normal mode axes. Measurement and correction of the “vertical” mode dispersion then provides a route to minimizing the vertical emittance. Since the BPMs are calibrated directly from the beam motion, the technique is intrinsically insensitive to BPM gain and alignment errors. Furthermore, given the appropriate hardware for resonant excitation and turn-by-turn beam position measurements, data collection and analysis are very fast and thus the technique can be applied as easily to a large ring as to a small ring.

In Sec. II we describe the technique in detail, and discuss the underlying theory. We illustrate its application by presenting simulation results in Sec. III. Some tests of the technique have already been carried out at Cornell’s CesrTA storage ring [7]: we present some relevant results from these tests in Sec. IV. Finally, we discuss some conclusions and outline possible further work in Sec. V.

II. THEORY

A. Normal mode dispersion and emittance

In a lepton storage ring, the beam is prevented from damping to zero emittance by quantum excitation. Consider a particle following the closed orbit. In a magnetic field (e.g. in a dipole magnet), the particle will randomly emit photons; if the dispersion is nonzero, then the change in energy of the particle resulting from a photon emission will mean that the particle is no longer on the closed orbit. Instead, it has some betatron amplitude around the closed orbit. The average of the betatron amplitude (more precisely, of the betatron action) over all particles in the beam gives the beam emittance.

Particle motion in a storage ring can be conveniently described in terms of action-angle variables. In the present

*a.wolski@liverpool.ac.uk

Published by the American Physical Society under the terms of the Creative Commons Attribution 3.0 License. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

work, we consider only linear motion. The relationship between action-angle and Cartesian variables is

$$\vec{x} \equiv \begin{pmatrix} x \\ p_x \\ y \\ p_y \\ z \\ \delta \end{pmatrix} = N \begin{pmatrix} \sqrt{2J_I} \cos \phi_I \\ -\sqrt{2J_I} \sin \phi_I \\ \sqrt{2J_{II}} \cos \phi_{II} \\ -\sqrt{2J_{II}} \sin \phi_{II} \\ \sqrt{2J_{III}} \cos \phi_{III} \\ -\sqrt{2J_{III}} \sin \phi_{III} \end{pmatrix} \equiv N \cdot \vec{J}, \quad (1)$$

where N is a matrix that transforms the single-turn transfer matrix M into a block-diagonal rotation R :

$$N^{-1}MN = R. \quad (2)$$

The Cartesian variables are defined as follows. x and y are, respectively, horizontal and vertical transverse coordinates in a plane perpendicular to the reference trajectory at a given point s ; p_x and p_y are the normalized transverse canonical momenta:

$$p_x = \frac{\gamma m \dot{x} + qA_x}{P_0} \quad (3)$$

(and similarly for p_y), where γ is the relativistic factor for a particle with rest mass m and charge q , the dot denotes the time derivative, A_x is the x component of the electromagnetic vector potential, and P_0 is a chosen reference momentum. The longitudinal coordinate is defined by

$$z = \beta_0 c(t_0 - t), \quad (4)$$

where t is the time that the particle arrives in the plane perpendicular to the reference trajectory at s , β_0 is the normalized velocity of a particle with momentum P_0 , and a particle traveling along the reference trajectory with normalized velocity β_0 arrives at the point s at time t_0 . The energy deviation δ is defined by

$$\delta = \frac{E}{P_0 c} - \frac{1}{\beta_0}, \quad (5)$$

where E is the total energy of the particle (for a particle with the reference momentum, $\delta = 0$).

The beam emittances are given by averaging the oscillation amplitudes (betatron and synchrotron actions) over all particles in the beam:

$$\epsilon_k = \langle J_k \rangle, \quad (6)$$

where $k = \text{I, II, or III}$ indicates the degree of freedom. Now consider a particle that undergoes an instantaneous change in energy, for example, resulting from photon emission. The energy deviation δ changes by some amount $\Delta\delta$. In the approximation that the photon is emitted directly along the instantaneous direction of motion of the particle (valid for ultrarelativistic particles), the transverse Cartesian coordinates and momenta, and the longitudinal coordinate, do not change:

$$\vec{x} \mapsto \vec{x} + \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \Delta\delta \end{pmatrix}. \quad (7)$$

The change in the action-angle variables may be expressed:

$$\vec{J} \mapsto \vec{J} + N^{-1} \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \Delta\delta \end{pmatrix}. \quad (8)$$

The condition for there to be no excitation of the mode II emittance (generally identified with the vertical emittance) can then be written:

$$N_{36}^{-1} = N_{46}^{-1} = 0, \quad (9)$$

where N_{ij}^{-1} is the (i, j) component of the inverse of N .

Generally, it is easier to give a direct physical interpretation to the components of the single-turn transfer matrix M than to the components of the normalizing matrix N (which can be constructed from the eigenvalues of M). For example, the components of M outside the 2×2 block diagonals are related to betatron coupling (M_{13}, M_{14} etc.), to horizontal dispersion (M_{15}, M_{16} , etc.), and to vertical dispersion (M_{35}, M_{36} , etc.). Although the condition in Eq. (9) does represent a constraint on the components of M , it is not easy to express this constraint directly in terms of the components of M . In general, it is possible for every component of M to be nonzero, while still satisfying Eq. (9). Physically, this would correspond to a case where the dispersion is parallel to the axis for motion in mode I. Then, under a change in energy of a particle, only mode I motion is excited.

It should be noted that, although we refer to axes corresponding to motion in each of the (transverse) normal modes, a beam excited in one or other normal mode will, in general, describe an ellipse, rather than a straight line, in coordinate space when observed over many turns at a particular location in a storage ring. Properly, one can only speak of a mode ‘‘axis’’ in the special case that the ellipse has infinitesimal width. That is, the ellipse is essentially a line. However, if the coupling in the ring is not strong, then the ellipses described by normal mode motion will have very narrow width, and it is reasonable to associate the major axis of the ellipse with an axis of the normal mode motion. Some justification for this assertion is given in Sec. IV, where we present and discuss the experimental results.

Often, low-emittance tuning procedures in storage rings address the vertical dispersion and betatron coupling separately. However, if one can measure (and minimize) the

component of the dispersion along the mode II axis, then it should be possible to minimize the vertical emittance by correcting a single quantity (the mode II dispersion) rather than two separate quantities. Besides offering a more direct and concise correction procedure, there is the possibility of achieving a better correction. With a limited number of correction elements in the ring, it is conceivable that the optimum conditions for low emittance may be achieved not by trying to minimize either (or both) the vertical dispersion and betatron coupling, but by trying to align the dispersion with the mode I axis.

There are numerous formalisms for describing coupled motion in accelerator beam lines, including storage rings; see, for example, Refs. [8–10]. The significance of the mode II dispersion for the emittance is indicated by the fact that, in the formalism of [9], the mode II emittance in a storage ring may be estimated from

$$\epsilon_{\text{II}} \approx C_q \gamma^2 \frac{I_{5,\text{II}}}{I_2}, \quad (10)$$

where $C_q \approx 3.832 \times 10^{-13}$ m is the quantum constant, γ is the relativistic factor of the beam, and I_2 and $I_{5,\text{II}}$ are synchrotron radiation integrals:

$$I_2 = \oint \frac{ds}{\rho^2}, \quad I_{5,\text{II}} = \oint \frac{\mathcal{H}_{\text{II}}}{|\rho|^3} ds. \quad (11)$$

ρ is the radius of curvature of the orbit, and the lattice function \mathcal{H}_{II} is defined by

$$\mathcal{H}_{\text{II}} = \gamma_{\text{II}} \eta_{\text{II}}^2 + 2\alpha_{\text{II}} \eta_{\text{II}} \eta'_{\text{II}} + \beta_{\text{II}} \eta_{\text{II}}'^2, \quad (12)$$

where α_{II} , β_{II} , γ_{II} , η_{II} , and η'_{II} are generalizations of the uncoupled Twiss parameters and dispersion functions to the coupled case. In the case of weak coupling, we can expect the Twiss parameters to take values close to those in the case of zero (or corrected) coupling; however, the mode II dispersion may be significantly different from the vertical dispersion. This may occur, for example, if coupling rotates the mode II axis so that it has a horizontal component at a location where there is large horizontal dispersion: the mode II dispersion will then acquire a significant contribution from the horizontal dispersion, even if the vertical dispersion is zero.

It is important to note that the emittance computed using Eq. (10) takes into account the effects of both vertical dispersion and betatron coupling. An experimental method to measure and minimize directly the mode II dispersion, rather than the vertical dispersion and betatron coupling separately, therefore offers the possibility of an elegant and effective procedure for low-emittance tuning.

B. Measurement of mode II dispersion

The BPMs in a storage ring are usually designed to return the horizontal (x) and vertical (y) position of the beam with respect to a specified “laboratory” coordinate

system. The coordinate values are obtained by processing the signals on the buttons induced by a beam. Generally, the signals on the BPM buttons have a nonlinear dependence on the position of the beam. The precise form of the dependence depends on the geometry of the buttons and the vacuum chamber. However, if the beam position varies over a small range, then the button signals can be assumed to vary linearly with beam position. This assumption is needed for calibrating BPMs using normal mode beam excitation. The justification for this can be obtained from the simulation and experimental results.

Alignment errors and errors in the signal amplification and processing electronics lead to systematic errors in the coordinate values returned by a BPM. For simplicity, we consider in this section a linear model in which the errors may be represented by a gain matrix, g :

$$\begin{pmatrix} x \\ y \end{pmatrix}_{\text{measured}} = g \cdot \begin{pmatrix} x \\ y \end{pmatrix}_{\text{actual}}, \quad (13)$$

where

$$g = \begin{pmatrix} g_{xx} & g_{xy} \\ g_{yx} & g_{yy} \end{pmatrix}. \quad (14)$$

Ideally, g is the identity matrix. In practice, the components of g can vary independently of each other. The value of g_{yx} is particularly significant for correction of the vertical dispersion: if it is nonzero, then a purely horizontal motion of the beam resulting from a change in beam energy (as during a dispersion measurement) can also appear in the measured y coordinate. An attempt to “correct” the measured vertical dispersion would then introduce vertical dispersion where none was, in fact, present.

There are various techniques that can be applied to determine the gain errors (either in a linear, or in a more general nonlinear model), including ORM analysis [5] and BPM gain mapping [11,12]. Measurements of the BPM gain errors will be valid only over a limited period of time due to such things as alignment changes, drifts in electronics, and changes in the orbit coupled with the nonlinear response of a BPM, etc. Some of the measurement techniques (e.g. ORM analysis) involve significant time and effort. Other techniques, particularly those based on analysis of turn-by-turn data, may be accomplished more quickly. However, it will generally be beneficial to have a range of methods available that may provide results that are either directly comparable or complementary. Given the need to repeat measurements periodically because of variation in machine conditions, techniques in which data may be collected within a few minutes have a strong advantage.

One such technique is to calibrate the BPM from normal mode motion of the beam. By observing turn-by-turn motion with the beam excited in one or the other transverse normal mode, it is possible to calibrate the BPMs so that the values returned are the beam coordinates along the

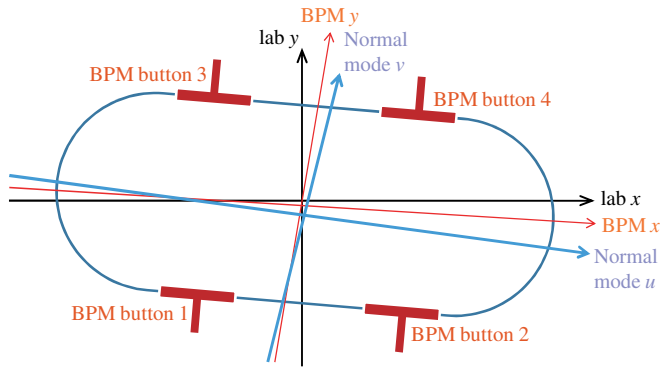


FIG. 1. Coordinate systems in a BPM. The laboratory coordinate system is defined with respect to the survey positions. BPM alignment and electronics errors mean that the coordinates returned by the BPM are with respect to the “BPM axes” which are not perfectly aligned with the laboratory system. Coupling means that resonant excitation of the transverse modes leads to motion along axes different from the laboratory or BPM axes. The four BPM buttons are labeled according to the convention used in CESR.

normal mode axes. As discussed in Sec. II A, these are the coordinates that are most relevant for low-emittance tuning. The relevant coordinate systems that can be defined in a BPM are illustrated schematically in Fig. 1.

Consider a BPM with four buttons. Let the signal on button i when a bunch passes through the BPM be denoted b_i ($i = 1, 2, 3$, or 4). If the coupling is not too large, the beam motion in normal mode I will lie along a line that we refer to as the u axis, and motion in mode II will lie along a line that we refer to as the v axis.

Assuming a linear response, which will be the case if the amplitude of excitation is not too large, the relationship between a change $(\Delta u, \Delta v)$ of the beam from its closed orbit position and the corresponding change in the button signals can be written as

$$\begin{pmatrix} \Delta b_1 \\ \Delta b_2 \\ \Delta b_3 \\ \Delta b_4 \end{pmatrix} = \begin{pmatrix} \left(\frac{\partial b_1}{\partial u}\right)_v & \left(\frac{\partial b_1}{\partial v}\right)_u \\ \left(\frac{\partial b_2}{\partial u}\right)_v & \left(\frac{\partial b_2}{\partial v}\right)_u \\ \left(\frac{\partial b_3}{\partial u}\right)_v & \left(\frac{\partial b_3}{\partial v}\right)_u \\ \left(\frac{\partial b_4}{\partial u}\right)_v & \left(\frac{\partial b_4}{\partial v}\right)_u \end{pmatrix} \begin{pmatrix} \Delta u \\ \Delta v \end{pmatrix}. \quad (15)$$

The matrix elements $(\partial b_i/\partial u)_v$ and $(\partial b_i/\partial v)_u$ can be obtained by observing the beam when the beam is shaken at the mode I and mode II frequencies via resonant excitation of the beam. When the beam is shaken at a frequency corresponding to mode I, the position $(\Delta u_j, \Delta v_j)$ on the j^{th} turn will have $\Delta v_j = 0$. Similarly, when shaking at the mode II frequency, Δu_j will be zero. When shaking the beam, the absolute changes Δu and Δv will not be known. However, we can determine the correlations between the changes in button signals from the readings taken during

the excitation in each normal mode. For example, $(\frac{\partial b_2}{\partial b_1})_v$ is the change in the signal at button 2 per unit change in the signal at button 1, for motion in mode I (v constant). Since

$$\left(\frac{\partial b_2}{\partial b_1}\right)_v = \frac{\left(\frac{\partial b_2}{\partial u}\right)_v}{\left(\frac{\partial b_1}{\partial u}\right)_v}, \quad (16)$$

we can write Eq. (15) in the form

$$\begin{pmatrix} \Delta b_1 \\ \Delta b_2 \\ \Delta b_3 \\ \Delta b_4 \end{pmatrix} = BC \begin{pmatrix} \Delta u \\ \Delta v \end{pmatrix}, \quad (17)$$

where

$$B = \begin{pmatrix} 1 & 1 \\ \left(\frac{\partial b_2}{\partial b_1}\right)_v & \left(\frac{\partial b_2}{\partial b_1}\right)_u \\ \left(\frac{\partial b_3}{\partial b_1}\right)_v & \left(\frac{\partial b_3}{\partial b_1}\right)_u \\ \left(\frac{\partial b_4}{\partial b_1}\right)_v & \left(\frac{\partial b_4}{\partial b_1}\right)_u \end{pmatrix}, \quad (18)$$

and

$$C = \begin{pmatrix} \left(\frac{\partial b_1}{\partial u}\right)_v & 0 \\ 0 & \left(\frac{\partial b_1}{\partial v}\right)_u \end{pmatrix}. \quad (19)$$

The components of B can be determined by correlating the changes in signals measured on different buttons while resonantly exciting the beam in one or the other of the transverse normal modes. Without knowing the absolute amplitude of the motion, the components of C cannot be determined. However, they may be estimated, e.g., from modeling the BPM response to horizontal and vertical beam motion. Errors in the values for the components of C will lead to systematic errors in the measurement of the normal mode dispersion. In particular, an error in $(\frac{\partial b_1}{\partial v})_u$ will lead to a systematic error, corresponding to some scale factor, in the measurement of the mode II dispersion. However, since the ultimate goal is to achieve zero mode II dispersion, this scale factor is not expected to be significant. Put another way, if the mode II dispersion is zero, then the systematic error will not affect the measurement. The systematic error will mean that the effect of the correction will not be exactly as expected, and some iteration will be needed; however, if the error is not too large, the correction should still converge.

Given values for the components of B (from measurements) and C (from modeling), we can determine C^{-1} and B^{-1} (defined such that $B^{-1} \cdot B = I_2$, with I_2 being the 2×2 identity matrix). Then

$$\begin{pmatrix} \Delta u \\ \Delta v \end{pmatrix} = C^{-1} B^{-1} \begin{pmatrix} \Delta b_1 \\ \Delta b_2 \\ \Delta b_3 \\ \Delta b_4 \end{pmatrix}. \quad (20)$$

Equation (20) represents a BPM calibration. It allows us to determine the changes in position along the normal mode axes in a BPM, from the changes in button signals. Applied to the button signals taken during a dispersion measurement, Eq. (20) allows us to determine the normal mode dispersion. The dispersion measurement may be made in the conventional way, by observing the change in beam position on the (calibrated) BPMs with respect to a change in beam energy induced by variation of the frequency of the rf cavities.

Note that, since the BPM is calibrated directly from normal mode motion of the beam, the normal mode dispersion measurement will be insensitive to BPM tilt and gain errors. Such errors can, in some circumstances, be a limiting factor in emittance tuning using conventional techniques. Nonlinear variation of the BPM buttons response to changes in beam position (i.e. changes that are large compared with the amplitude of the resonant excitation) may mean that the calibration is a function of the closed orbit position. The calibration may also be a function of such things as the beam current, the temperature of the electronics, etc. However, since the calibration is straightforward and (with the appropriate hardware for performing a resonant excitation) can be accomplished very quickly (within a few minutes), a calibration can be performed immediately before each dispersion measurement.

C. Correction of mode II dispersion and minimization of mode II emittance

Once the BPMs have been calibrated to allow measurement of the mode II dispersion, correction of the dispersion may be accomplished using conventional techniques. For example, a set of components (vertical steering magnets, and skew quadrupoles) that affect the mode II dispersion may be identified, and a response matrix between these components and the mode II dispersion at the BPMs constructed using a model, or from measurements in the machine. Inverting the response matrix (by singular value decomposition) and applying the inverted matrix to the measured mode II dispersion gives the changes required to correct the mode II dispersion.

It is important to note that, typically, BPMs are located in field-free regions where synchrotron radiation is not generated. However, the mode II emittance will be minimized by correcting the mode II dispersion in locations (dipoles, undulators, and wigglers) where synchrotron radiation is produced. In principle, this ought to be taken into account in the correction scheme, e.g., by placing a higher weight on those BPMs that are closest to the points where large amounts of synchrotron radiation are produced. In

weighting the BPMs for the correction procedure, there should also be some account taken of the optics. The mode II dispersion at a location with large horizontal dispersion will show a strong response to even a small amount of coupling: the coupling tilts the normal mode axes, and even a small tilt of the mode II axis will lead to a large component of the horizontal dispersion appearing along the mode II axis. On the other hand, where the horizontal dispersion is small, the mode II dispersion will be relatively insensitive to changes in coupling. Low-emittance storage rings generally need low horizontal dispersion in the dipoles (and in any insertion devices); but the horizontal dispersion may reach large values between the dipoles. To minimize the mode II emittance, therefore, a larger weight should be given to the mode II dispersion close to the dipoles and insertion devices, both because this is where the mode II emittance is generated, and because this is where the mode II dispersion will likely be least sensitive to coupling. In practice, we find from simulations of CEsrTA that, although weighting the BPMs has some effect on the results of tuning based on mode II dispersion correction, there does not seem to be a very strong impact. It is possible that with further optimization of the way in which the weighting is applied, or in other lattices, there may be a more significant effect. The results from simulations and experiments presented in the following sections have been obtained without applying any weighting to the BPMs.

It should also be noted that, in general, emittance is generated not just by the dispersion η_{II} , but also by gradient of the dispersion, η'_{II} ; however, it is usually only possible to measure the dispersion and not its gradient. The inability to measure η'_{II} may limit the effectiveness of tuning for low emittance by correcting the mode II dispersion. It may be hoped, however, that if the BPMs are not too far apart, then zero (or very low) values for η_{II} will imply low values also for η'_{II} . The possible effectiveness of low-emittance tuning based on correction of only η_{II} in CEsrTA will be demonstrated in simulations in the following section.

III. SIMULATIONS

A. CEsrTA

To illustrate the use of normal mode calibration of the BPMs for low-emittance tuning of a storage ring, we present the results of some simulations based on a model of CEsrTA. The results of some experiments on CEsrTA are presented in Sec. IV.

CesrTA is a synchrotron with 768 m circumference that can store either electrons or positrons. The beam energy can be varied between 2 and 5 GeV. For all the studies described here, the energy was 2 GeV. Typical lattice functions (Twiss parameters and dispersion) are shown in Fig. 2. With these lattice functions, the natural emittance of CEsrTA at 2 GeV is approximately 2.7 nm.




FIG. 2. Lattice functions in CsrTA. Top: horizontal (black) and vertical (red) beta functions. Bottom: horizontal dispersion, with BPM locations indicated by blue circles, and skew quadrupole locations indicated by vertical dashed lines.

The main purpose of CsrTA is to investigate electron cloud effects in a parameter regime with low vertical emittance (< 20 pm). The storage ring was developed from the CESRc ring by a number of modifications, including removal of the detector (CLEO) and relocation of the superferric damping wigglers to dispersion-free regions of the lattice. Since the original CESRc ring had not been designed to reach the low emittances needed for CsrTA, a significant part of the CsrTA program consisted of development of techniques for achieving vertical emittance below 20 pm on a reliable and routine basis. This includes improvement of the instrumentation including new BPM electronics capable of turn-by-turn measurements with resolution of a few microns, and fast beam-size monitors to allow emittance measurements with resolution of a few picometers.

Simulations for this paper were carried out in an extended version of AT [13], an accelerator modeling code running in Matlab. The model allows a range of magnetic field strength and alignment errors to be simulated, including skew quadrupole fields, tilts (around the beam axis), and vertical displacements. The x and y coordinates (in the laboratory frame) of a particle (representing the centroid of a bunch) tracked through the lattice can be recorded over any number of turns at any specified locations. Since the calibration technique that we apply is based on readings from individual buttons, the model button signals were

generated as functions of the beam x and y coordinates. The “BPM coordinates” can then be calculated as functions of the button signals. By this means, the calibration process can be simulated, and the calibration matrices B and C for each BPM calculated from the results of tracking over multiple turns in the model.

B. BPM model

The calibration technique that we describe here depends on the BPM button signals having a linear dependence on the beam position, over the range of beam position variations that occur during the calibration process. However, over the range of possible beam positions, the BPM button signals have a complicated nonlinear response, that depends on the geometry of the vacuum chamber in the region of the BPM, and on the geometry of the buttons themselves. To include nonlinear effects in the simulations without needing to implement the full complex dependence of the button signals on the beam position, we used a simplified model in the simulations, in which the signal b_i on BPM button i is given by

$$b_i = \frac{g_i}{\sqrt{(x - x_i)^2 + (y - y_i)^2}}, \quad (21)$$

where (x, y) are the beam coordinates (with respect to the reference trajectory) and (x_i, y_i) are the coordinates of the

button location (again with respect to the reference trajectory). g_i is a gain factor, that will include (for example) a dependence on the bunch charge. In principle, if the bunch charge is known accurately, then the signals from only two buttons are needed to determine the beam coordinates. However, in the case that the buttons are symmetrically positioned at

$$x_i = \pm x_0, \quad y_i = \pm y_0 \quad (22)$$

(where for x_i , the plus sign holds for $i = 2$ and $i = 4$, and for y_i , the plus sign holds for $i = 3$ and $i = 4$), then averaging the button signals leads to the following expressions for the beam coordinates:

$$x = \frac{1}{8x_0} \left[\left(\frac{g_1}{b_1} \right)^2 - \left(\frac{g_2}{b_2} \right)^2 + \left(\frac{g_3}{b_3} \right)^2 - \left(\frac{g_4}{b_4} \right)^2 \right], \quad (23)$$

$$y = \frac{1}{8y_0} \left[\left(\frac{g_1}{b_1} \right)^2 + \left(\frac{g_2}{b_2} \right)^2 - \left(\frac{g_3}{b_3} \right)^2 - \left(\frac{g_4}{b_4} \right)^2 \right]. \quad (24)$$

BPM alignment (including tilt) errors may be implemented in the simulation by applying some appropriate transformation to the button coordinates (x_i, y_i) in Eq. (21), while continuing to use the nominal values for the button coordinates in calculating the beam coordinates from the button signals using Eqs. (23) and (24). Similarly, gain errors may be implemented by applying some variation to the values of g_i in Eq. (21), while continuing to use the nominal values for the gain coefficients in calculating the beam coordinates from the button signals using Eqs. (23) and (24).

Simulated calibration data, obtained by implementing Eqs. (21), (23), and (24) in a tracking code, are shown in Figs. 3 and 4. The data were generated by determining the 6D phase space coordinates of a particle with, in this case, nonzero mode I oscillation amplitude (using the eigenvectors of the single-turn matrix), and zero amplitude in the other two oscillation modes. The particle was tracked for

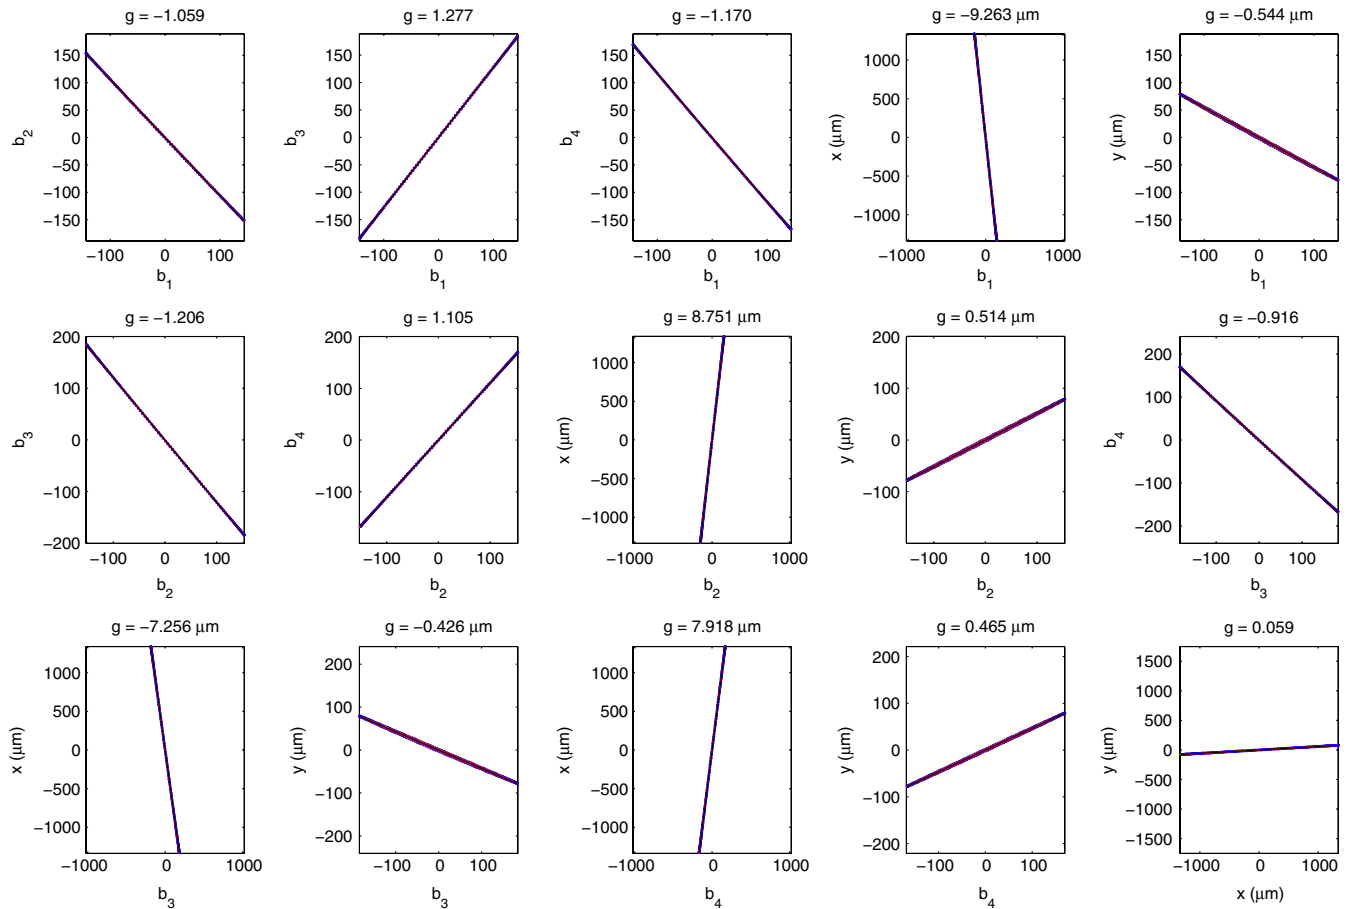


FIG. 3. Simulation of turn-by-turn button signals on a BPM (7W). Each plot shows a correlation between two BPM button signals or beam coordinates (returned by the BPM), obtained by tracking a bunch performing coherent oscillations in normal mode I. An ideal model (with no errors) of the CesrTA lattice is used: tracking data are shown as blue points; fitted ellipses are shown in red; the major axes of the ellipses are shown as black lines. The ellipses are more clearly visible when coupling is present (see Fig. 4). Button signals are calculated using Eq. (21), with relative gain errors applied to the buttons as follows: $g_2/g_1 = 1.063$, $g_3/g_1 = 1.282$, $g_4/g_1 = 1.170$. The button locations are given by $x_0 = 20$ mm, $y_0 = 10$ mm. x and y beam coordinates are calculated from Eqs. (23) and (24), but using nominal gain values $g_1 = g_2 = g_3 = g_4$.

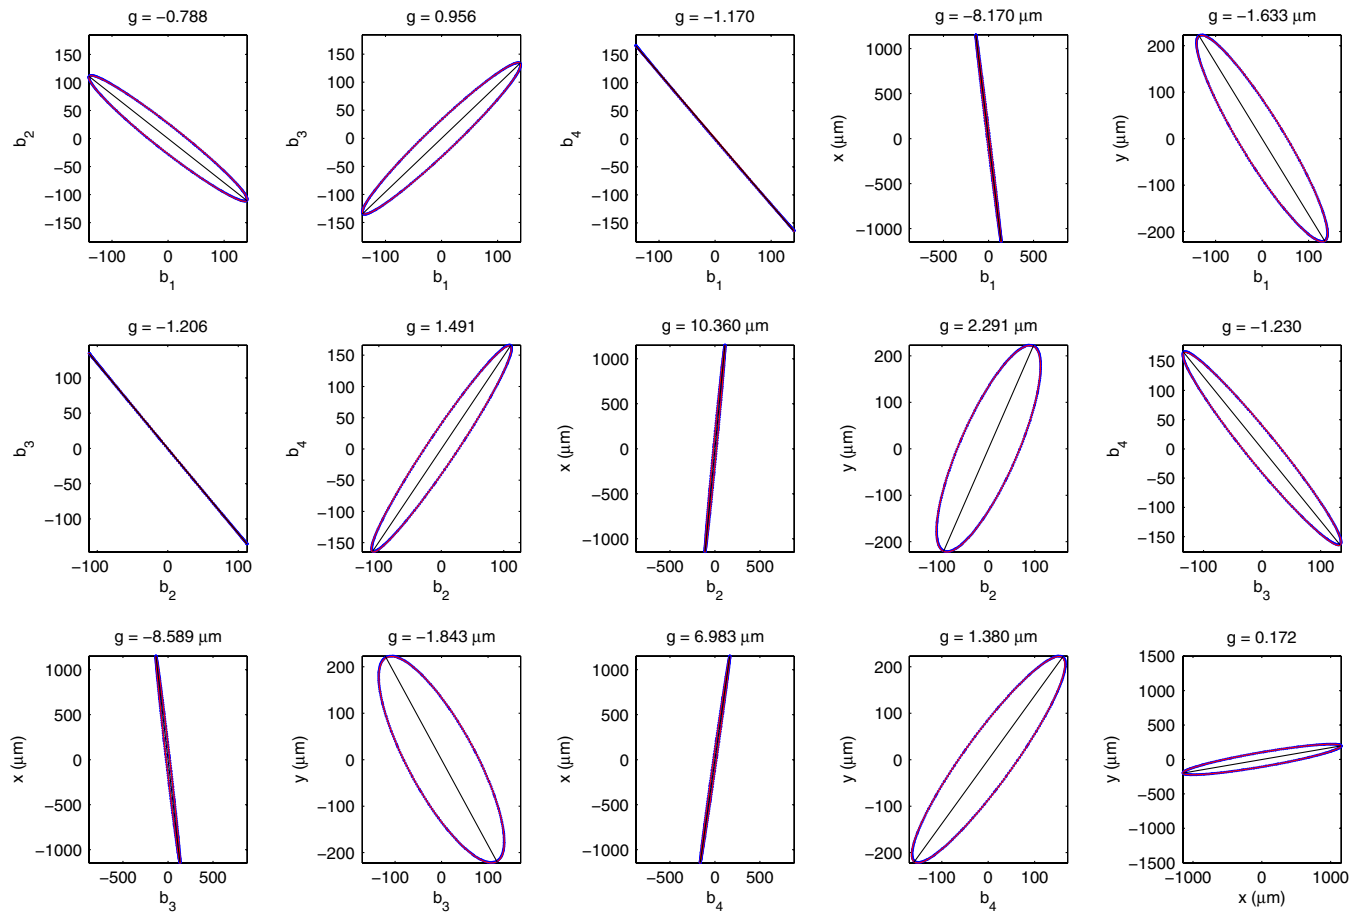


FIG. 4. As Fig. 3, but with skew quadrupole 48W turned on with strength $k = -0.023 \text{ m}^{-2}$.

several hundred turns in an ideal model of the lattice, with the button signals at the BPMs, calculated using Eq. (21), recorded on each turn. Figure 3 shows a set of correlation plots (between BPM button signals, or beam coordinates) for BPM 7W with relative gain errors applied to the buttons. The x and y beam coordinates are calculated from the button signals by applying Eqs. (23) and (24), but with “nominal” gain values $g_1 = g_2 = g_3 = g_4$. The relative button gain errors, used for calculating the button signals from the actual beam position in the simulation, have been chosen roughly to reproduce the experimental data that we show later in Fig. 7. Note that there is reasonably good agreement between the simulation and the experimental data in the gradients of the correlations between the button signals, and between the x coordinate and the button signals. The agreement between the simulation and the experimental data in the gradients involving the y coordinate, however, is poor: this is likely the result of coupling in the real machine, which is not included in the model. Without any coupling or gain errors, there will be no variation in the y coordinate resulting from oscillation in normal mode I. Figure 4 shows the results of the same simulation, except with one skew quadrupole (SKQ48W)

turned on, with strength $k = -0.023 \text{ m}^{-2}$. (Figure 4 is to be compared with Fig. 12, which shows the corresponding experimental data.) With coupling generated by the skew quadrupole, the beam motion in normal mode I is now no longer along a line in coordinate space: instead, the beam describes an ellipse, which results in ellipses clearly visible in the correlation plots between the button signals. Again, there is good agreement in the gradients of the major axes of the ellipses, between the simulation (Fig. 4) and the experimental data (Fig. 12), except for the plots involving the y coordinate.

C. Emittance calculation

Exact values for the horizontal, vertical, and longitudinal emittances in the model, taking into account field and alignment errors leading to dispersion and coupling, are obtained in AT using the envelope method [14]. As an initial test of the theory, the emittance for 1000 seeds of skew quadrupole errors distributed around the ring was calculated in three different ways: (1) using the envelope method; (2) estimated using Eq. (10); (3) estimated using Eq. (10) but with a modified version of Eq. (12), in which

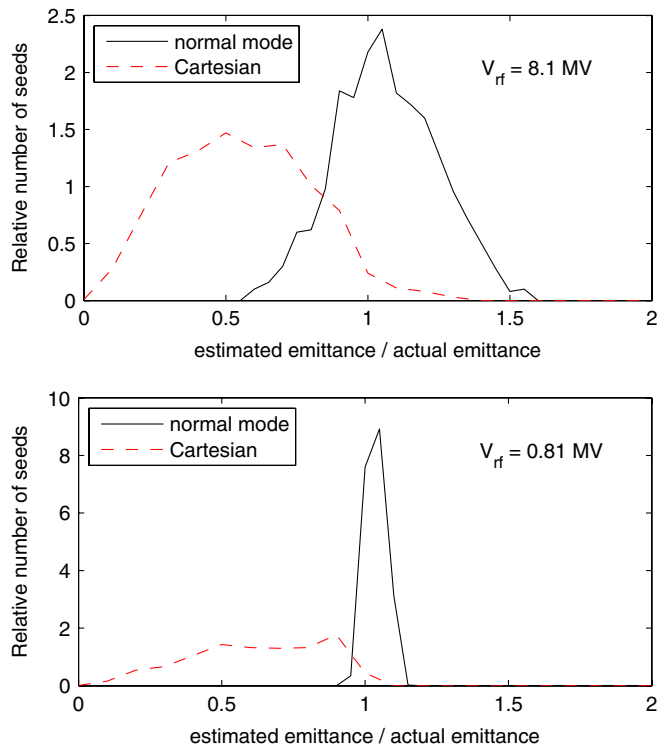


FIG. 5. Emittance estimated from dispersion in CesiumTA with (top) 8.1 MV rf voltage, and (bottom) 0.81 MV rf voltage. The lines show the distribution (over 1000 seeds of machine errors) of the ratio of the estimated mode II emittance to the actual mode II emittance; the vertical axes are scaled so that each curve has unit area. In the normal mode case, the mode II dispersion is used, and the emittance estimated using Eq. (10). In the Cartesian case, a modified version of Eq. (10) is used, in which the mode II dispersion is replaced by the vertical dispersion in Eq. (12). In both cases, the exact emittance is calculated using the envelope method. The approximations involved in Eq. (10) are valid for small values of the synchrotron tune (which is 0.074 at 8.1 MV rf, and 0.023 at 0.81 MV rf).

the mode II dispersion is replaced by the vertical (y) dispersion. Figure 5 shows the distribution of the ratio of the estimated (methods 2 and 3) to the exact (method 1) mode II emittances over 1000 seeds of random skew quadrupole errors. Results are shown for two cases: with the nominal rf voltage of 8.1 MV in CesiumTA, and with a reduced rf voltage of 0.81 MV.

If we consider first the case of the reduced rf voltage, we see that Eq. (10) gives a very good estimate of the true mode II emittance, while the emittance calculated from modified equations (using the vertical rather than the mode II dispersion) shows a very large spread. This illustrates the fact that the mode II dispersion takes account all of the significant contributions to the emittance generation.

For larger rf voltages, the spread in the results from Eq. (10) increases: this is because the approximations required to derive this equation start to break down. In particular, the derivation assumes that any correlations between the longitudinal and transverse coordinates of

particles in the bunch can be neglected. Although this is true for most storage rings, correlations between longitudinal and transverse coordinates can appear when the synchrotron tune becomes large. A large synchrotron tune results from strong longitudinal focusing from the rf cavities, and a large momentum compaction factor. If we consider an ellipse in longitudinal phase space, the rf focusing will give a “tilt” to the ellipse, while the momentum compaction factor as the particles move around the rest of the ring will have the effect both of advancing particles around the ellipse (the longitudinal phase advance), and modifying the shape of the ellipse. The stronger the rf focusing, and the larger the momentum compaction factor, the larger the tilt of the longitudinal phase space ellipse will be. For a real distribution, this indicates that there will be a significant correlation between the energy of a particle and its longitudinal position (coordinate) within the bunch. Where dispersion is present, this will in turn lead to a correlation between the transverse and longitudinal coordinates: this is not taken into account in the derivation of Eq. (10). In CesiumTA, the synchrotron tune is about 0.023 with an rf voltage of 0.81 MV, increasing to 0.074 at an rf voltage of 8.1 MV. The effect of the larger synchrotron tune on the accuracy of Eq. (10) can be clearly seen in Fig. 5.

The fact that the relationship between mode II emittance and mode II dispersion expressed by Eq. (10) starts to break down for large rf voltages, suggests that correction of the mode II dispersion may be a more effective technique for correcting the emittance in storage rings at lower values of the synchrotron tune. However, the simulation results presented in the following section suggest that the technique should still be useful even in CesiumTA with the nominal rf voltage. The purpose of the simulations described in the present section is principally to illustrate the relevance of the mode II dispersion for the mode II emittance, but the results also give some indication of potential limitations of the use of mode II dispersion for low-emittance tuning in particular (somewhat unusual) parameter regimes.

D. Emittance tuning

Emittance tuning simulations were performed to investigate the likely mode II emittance that might be achieved with the magnitude of magnet alignment and BPM gain errors expected in practice in CesiumTA. Random errors were applied, with values having Gaussian distributions, with rms as shown in Table I. Note that the BPM alignment and gain errors are applied as variations in the button coordinates (x_i, y_i) and gain coefficients g_i , as described in Sec. III B.

With errors applied in the model, the correction procedure was implemented as follows.

(1) An initial orbit correction using the vertical steering magnets was applied.

TABLE I. The rms values for error distributions.

Error type	Distribution rms
Dipole tilt	300 μ rad
Quadrupole tilt	300 μ rad
Quadrupole vertical alignment	250 μ m
Sextupole vertical alignment	250 μ m
BPM vertical alignment	100 μ m
BPM tilt	20 mrad
BPM gain	0–10%
Dispersion measurement resolution	12 mm

(2) The normal modes were identified from the single-turn transfer matrix; particles in each of the normal modes (with initial values of the dynamical variables taken from the real parts of the eigenvectors of the single-turn transfer matrix) were tracked for 50 turns, with the button signals at each BPM recorded on each turn.

(3) The calibration matrix B , defined in Eq. (18), for each BPM was determined from correlations between the turn-by-turn button signals. Nominal values, based on electromagnetic modeling of the BPM buttons and the design of the electronics, were assumed for the calibration matrix C .

(4) The mode II dispersion was calculated at each BPM from the change in (calibrated) BPM reading with respect to energy. That is, the change in button readings corresponding to a change in closed orbit with respect to energy were found, and then the change in coordinates calculated using Eq. (20).

(5) A dispersion correction was applied by setting the strengths of 26 skew quadrupoles (with locations shown in Fig. 2), determined from the nominal (i.e. ideal model) response matrix between the skew quadrupole strengths and the mode II dispersion at the BPMs.

(6) A further iteration of the dispersion measurement and correction was applied (by repeating steps 2–5).

For simulations involving a large number of seeds, the total time required for tracking becomes significant. The time required for tracking depends mainly on the number of turns used to collect the BPM calibration data. In step 2, the number of turns was chosen to minimize the tracking time required in the simulation, while still achieving good accuracy for the determination of the normal modes. It was found that increasing the number of turns beyond 50 did not have any significant impact on the results. In the experimental studies on CEsrTA, we used the full 1024 turns of data returned by the BPMs, simply because the data were available, and there would be no significant advantage in reducing the number of turns. We have not investigated how, in practice, the accuracy of the results depends on the number of turns of data collected.

The tuning simulation was repeated using a similar procedure to that given above, but using the vertical dispersion instead of the mode II dispersion. The results from each simulation, as distributions of the final emittance over

1000 seeds of random errors with different magnitude of BPM gain errors, are shown in Fig. 6.

Clearly, a dispersion correction (following an orbit correction) alone is insufficient to achieve the goal of mode II emittance less than 20 pm if the correction is based on the vertical dispersion. This is consistent with experience of CEsrTA operation: low-emittance tuning has generally been accomplished using correction of both (vertical) dispersion and coupling, with each correction calculated and applied separately. However, if the mode II dispersion is used, then dispersion correction (following an orbit correction) seems capable, at least in simulation, of reducing the emittance well below the goal of 20 pm. Note that the technique seems robust, in that for all 1000 seeds of random errors, the procedure converged to a stable solution.

Also shown in Fig. 6 are the distributions of the vertical dispersion and the mode II dispersion (rms of values measured at the BPMs, including simulated noise), after tuning based on correction of the respective dispersion in each case. For the vertical dispersion, the results are again sensitive to the BPM gain errors, and in fact the gain errors limit the quality of the dispersion correction that can be achieved. The mode II dispersion measurement is insensitive to the (systematic) BPM gain errors, but appears to be significantly affected by the (random) measurement errors: a resolution of 12 mm is assumed for the dispersion measurement. Because the measured value of the mode II dispersion following correction is limited by the resolution of the dispersion measurement, the mode II emittance following correction is poorly correlated with the measured rms dispersion. It is therefore difficult to set a goal for the dispersion correction: however, from the distributions in Fig. 6, it appears that if a mode II dispersion of around 15 mm (rms measured at the BPMs) can be achieved, then a mode II emittance below 20 pm has a good probability (in the simulations, 86% of seeds gave a mode II emittance below 20 pm after correction of the mode II dispersion).

It is also significant to note that the emittance achieved after minimizing the vertical dispersion is dependent on the BPM gain errors: larger BPM gain errors tend to lead to larger final mode II emittances. This is to be expected, since gain errors in the BPMs will couple horizontal dispersion into the vertical. With large BPM coupling, the vertical dispersion measurement will be subject to significant systematic errors, limiting the quality of the correction. However, the correction based on normal mode dispersion is completely insensitive to BPM gain errors: this is a consequence of the fact that the BPMs are calibrated using normal mode excitation of the beam. The calibration technique effectively determines the appropriate relationship between changes in button signal and changes in beam position, for the BPM to read the beam motion along the normal mode axes.

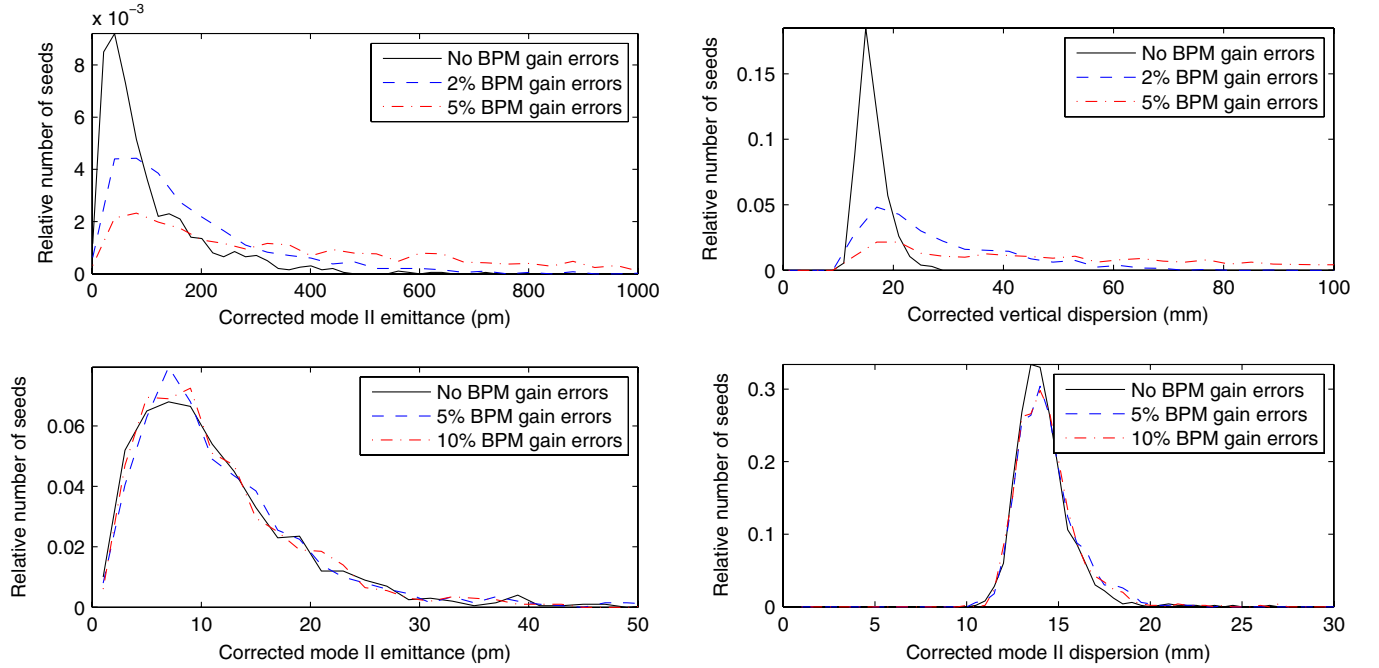


FIG. 6. Distribution (from 1000 seeds of machine errors with rms shown in Table I) of final mode II emittance (top and bottom, left) and dispersion (top and bottom, right) after correction of dispersion measured at the BPMs. The vertical axes are scaled so that each curve has unit area. Top: emittance (left) and vertical dispersion (right) after tuning based on correction of vertical (y) dispersion. Bottom: emittance (left) and mode II dispersion (right) after tuning based on correction of mode II dispersion. In each case, the dispersion is the rms of the dispersion measured at the BPMs (including measurement errors, corresponding to a resolution of 12 mm). Note the different horizontal scales between the top plots and the bottom plots. Note also that the emittance distribution for the correction based on vertical dispersion is sensitive to the magnitude of the BPM gain errors, whereas the emittance distribution for a correction based on mode II dispersion is insensitive to the BPM gain errors. Black lines correspond to no BPM gain errors; blue lines correspond to 2% (vertical dispersion correction) or 5% (normal mode dispersion correction) gain errors; and red lines correspond to 5% (vertical dispersion correction) or 10% (normal mode dispersion correction) gain errors.

Similar results have been obtained for simulations of other storage rings, e.g., the KEK Accelerator Test Facility [15].

IV. EXPERIMENTAL RESULTS

A. BPM calibration

Typical calibration data for a single BPM (BPM 7W) are shown in Fig. 7. The plots are produced by recording the raw button signals for 1024 consecutive turns, while the beam is resonantly excited in normal mode I. A similar set of plots is produced for resonant excitation in mode II. Each plot shows the correlation between a pair of buttons. The control system also records the horizontal (x) and vertical (y) position of the beam on each turn, based on a nominal calibration of the BPM: the correlations between the coordinates and the button signals are also shown. The amplitude of the beam oscillation is more than 2 mm peak to peak, in the horizontal. There is some oscillation visible in the vertical, about 100 μm peak to peak: this could be a consequence of some coupling in the machine, or some error in the nominal BPM calibration, or both. The correlations between the button signals are almost perfectly

linear, with very little scatter: we expect our assumption of a linear relationship between the beam position and button signals that we used in Sec. II B to be valid. The gradient of the correlations between each pair of buttons provides the values for the calibration matrix B , Eq. (18). The gradients are positive or negative according to whether the buttons are on the same horizontal side or opposite horizontal sides of the vacuum chamber (see Fig. 1). Note that calibration data for all BPMs in the ring can be collected simultaneously. Fitting the gradients to the correlation plots is a straightforward and very fast procedure.

To test the calibration of a BPM, we can inspect the Fourier spectrum of the beam coordinates returned by the BPM over a number of consecutive turns during resonant excitation of a normal mode. If the BPM is incorrectly calibrated to read the normal mode coordinates, the spectrum of both coordinates will contain peaks at the resonant frequency corresponding to the excited mode. For a correctly calibrated BPM, the spectrum of the mode I coordinate will contain a peak only at the mode I frequency, and the spectrum of the mode II coordinate will contain a peak only at the mode II frequency. Figure 8 shows the Fourier

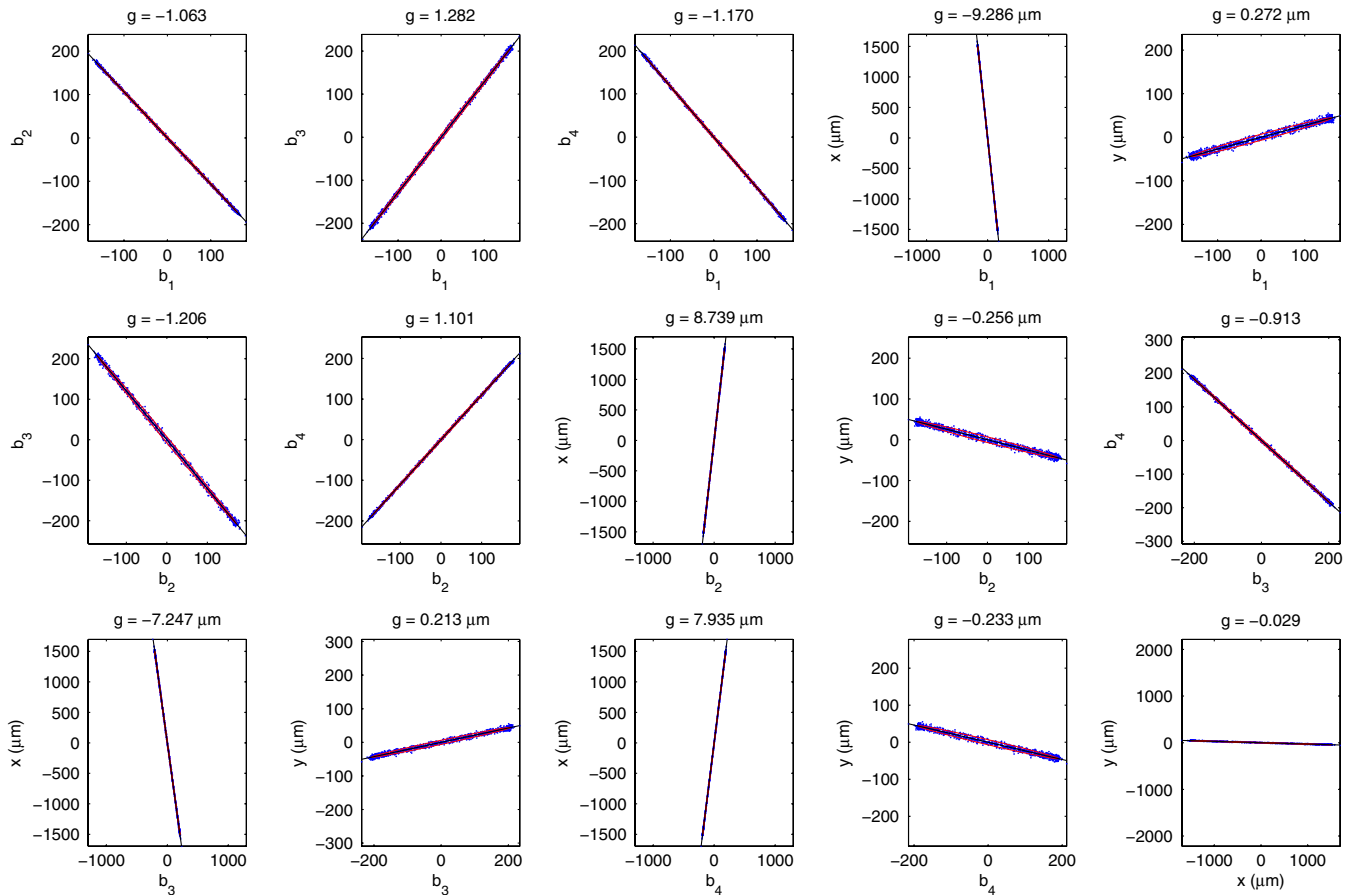


FIG. 7. Correlation plots for the button signals from BPM 7W during resonant excitation of normal mode I. Turn-by-turn signals from the BPM buttons (and the coordinates calculated using the nominal calibration) are shown as blue points. The points in each correlation plot are fitted with a red ellipse; the major axis of the ellipse is shown with a black line. The ellipses described by the data points are more obvious if the coupling is increased: compare with Fig. 12. Compare also with the simulation results shown in Fig. 3.




FIG. 8. Fourier spectra of beam position coordinates returned by BPM 7W over 1024 consecutive turns, with mode I resonant excitation. Note that the (fractional part of the) mode I tune is 0.59, and the (fractional part of the) mode II tune is 0.63.

spectra of the coordinates returned by BPM 7W over 1024 consecutive turns, with resonant beam excitation in mode I. We see that a peak corresponding to the “horizontal” tune (the mode I resonant frequency) appears in the spectra of both the x and y coordinates. However, using the normal mode calibration derived from the data shown in part in Fig. 7, the mode I frequency peak is entirely absent from the mode II spectrum. Similarly, Fig. 9 shows the Fourier spectra of the coordinates returned by BPM 7W over 1024 consecutive turns, with resonant beam excitation in mode II. With the nominal calibration, a peak corresponding to the mode II resonant frequency is clearly visible in the spectra of both the x and y coordinates. However, with the calibration based on observation of the turn-by-turn data, the mode II resonant frequency peak is entirely absent from the spectrum of the mode I coordinate. Note that there is a significant peak at the mode I resonant frequency, visible in the mode I coordinate: this likely results from real beam motion. These observations provide confidence in the BPM calibration obtained from observation of turn-by-turn button signals during resonant excitation in one or the other of the normal modes.




FIG. 9. Fourier spectra of beam position coordinates returned by BPM 7W over 1024 consecutive turns, with mode II resonant excitation. Note that the (fractional part of the) mode I tune is 0.59, and the (fractional part of the) mode II tune is 0.63.

B. Mode II dispersion response to changes in skew quadrupole strength

To test the measurement of the mode II dispersion using calibrated BPMs, and also to validate the model, the change in mode II dispersion was measured, following known changes in selected skew quadrupole strengths. Figure 10 shows the measured mode II dispersion after a standard tuning procedure, involving correction of the closed orbit, vertical dispersion, and betatron coupling. The skew quadrupole strengths in a model were adjusted to fit the measured mode II dispersion. Here, an ideal model, with all magnet strengths set to their design values, and no alignment or BPM errors, was used for the fitting. Using a response matrix (inverted by singular value decomposition), strengths of all skew quadrupoles were determined that gave the best match between the mode II dispersion in the model, and the mode II dispersion measured in the machine. The strengths of the skew quadrupoles and the mode II dispersion found in this fit in this case are also shown in Fig. 10. The rms (over the values at the BPMs) measured mode II dispersion is 21 mm, and the fitted skew quadrupole strengths are generally less than $2 \times 10^{-3}/\text{m}^2$. Note that there is a significant residual between the fitted dispersion and the measured dispersion: this suggests that there may be only a limited ability to correct the mode II dispersion using the skew quadrupoles.

After making the initial measurement, the strength of skew quadrupole 48W was changed by $\Delta k = -0.023 \text{ m}^{-2}$. The change in mode II dispersion and the changes in fitted skew quadrupole strengths are shown in Fig. 11. The horizontal dispersion in the lattice at the location of this skew quadrupole is close to zero. Therefore, the vertical

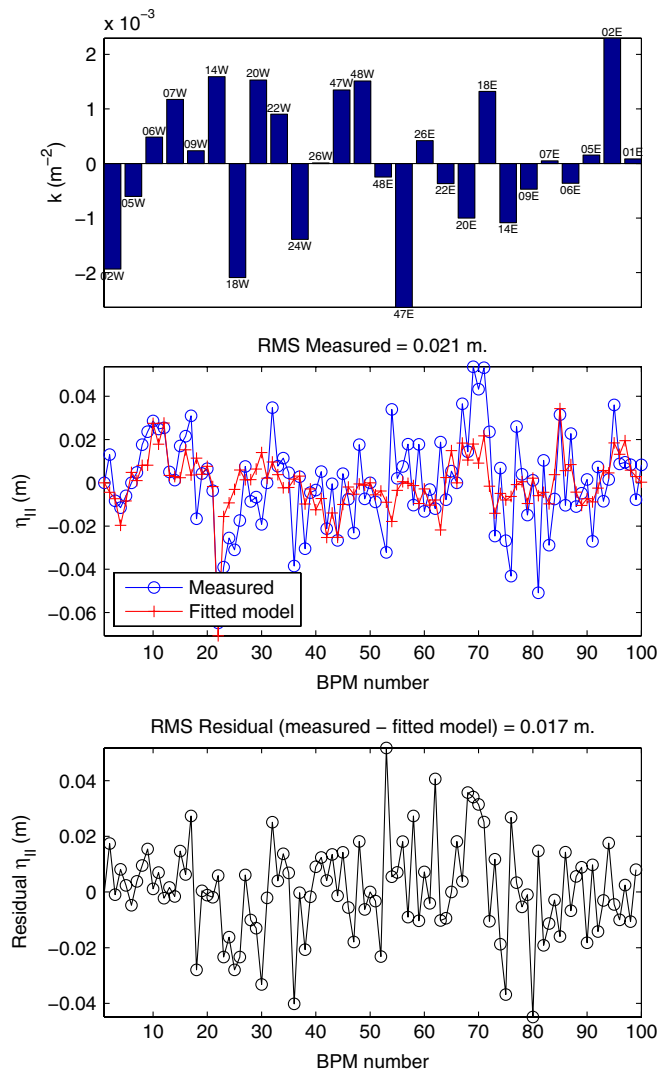


FIG. 10. Measured mode II dispersion after correction of orbit, (vertical) dispersion, and coupling. The middle plot shows a comparison between the measured dispersion, and the dispersion in a model with skew quadrupole strengths adjusted to fit the measurement. The final skew quadrupole strengths found from the fitting are shown in the top plot (the order of the quadrupoles is the same as shown in Fig. 2); the bottom plot shows the residual between the measurement and the fitted model. The rms is the root mean square over the measurements at the BPMs.

dispersion is not expected to change when the strength of the skew quadrupole is varied. This was verified using conventional measurement techniques (using the nominal BPM calibrations). However, there is a large change in the mode II dispersion (peak at 0.6 m). This is primarily a result of the coupling introduced by the skew quadrupole which tilts the normal mode axes. Thus, where the horizontal dispersion is large, a significant component appears along the mode II axis.

The strengths of the skew quadrupoles in a model of the machine were adjusted to fit the mode II dispersion measured after the change in skew quadrupole 48W, and the

strengths compared with the values fitted to the dispersion measured before changing skew quadrupole 48W. The changes in fitted skew quadrupole strengths are shown in the top plot in Fig. 11. The change in fitted strength of skew quadrupole 48W (-0.022 m^{-2}) agrees well with the change made in the machine (-0.023 m^{-2}). A number of other skew quadrupoles are fitted with smaller, but not insignificant, strengths. This may be the result of discrepancies between the model and the actual machine condition, or of some degeneracy between the effects of the skew quadrupoles on the mode II dispersion.

Note that the dispersion was measured by excitation of longitudinal oscillations of the beam. Low-level

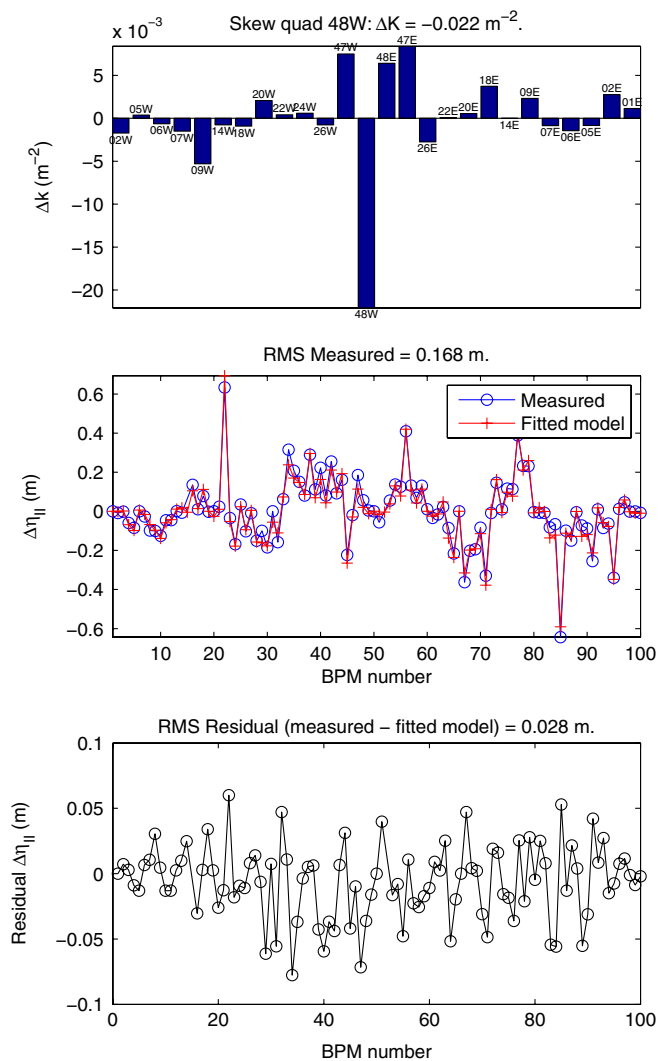


FIG. 11. Effect of change in strength of skew quadrupole 48W on mode II dispersion in CEsrTA. Top: skew quadrupole strengths in a model fitted to the measured mode II dispersion. Center: comparison between measured mode II dispersion and the mode II dispersion in the fitted model. Bottom: residual between measured and fitted mode II dispersion. The rms is the root mean square over the measurements at the BPMs.

processing of the button signals recorded during the excitation returns an amplitude and phase for the signal (at the synchrotron frequency) on each button. The normal mode calibration data (obtained by observation of button signals during excitation of the transverse modes) can then be applied directly to the button amplitude and phase data, to return the mode I and mode II dispersion. This technique, known as an “AC dispersion measurement” [16], has been shown to produce results that match very closely the measurement of dispersion using more conventional techniques. This is to be contrasted to the common procedure for measuring dispersion in a storage ring which involves measuring the change in the closed orbit when the rf frequency is varied. In CEsrTA, changing the rf frequency is a relatively slow process because of the superconducting rf system and there is a very limited range over which the rf frequency can be varied. The AC dispersion measurement, on the other hand, requires no variation in the rf parameters; only a few thousand turns (2^{15} turns are generally used when making the measurement on CEsrTA) of turn-by-turn button signals are required, which makes this measurement technique extremely fast. Results agree closely with dispersion measurement using rf frequency variation, whether using the standard or the normal mode BPM calibration.

Figure 12 shows the correlation plots between the button signals on BPM 7W during resonant excitation of mode I, with the strength of skew quadrupole 48W changed by $\Delta k = -0.023 \text{ m}^{-2}$. Compared with Fig. 7, which shows the same measurement, but with skew quadrupole 48W set to its nominal value after machine tuning, the effects of coupling are now very pronounced. Distinct ellipses are produced when signals from pairs of buttons are plotted against each other over successive turns. The approximation of weak coupling is a poor one, and the normal mode motion cannot be associated with simple axes in coordinate space. Nevertheless, we can still obtain an effective “normal mode calibration” of the BPMs, by taking the gradients of the major axes of the ellipses to define the components of the calibration matrix B . The justification for this lies in the agreement between the measured mode II dispersion response to changes in skew quadrupole strengths, and the response predicted from the model, as shown, for example, in Fig. 11. Furthermore, if the normal mode calibration is used in correcting the mode II dispersion, then on successive iterations the coupling should become weaker, and the weak coupling assumption will become a better approximation.

The validity of the model is further illustrated in Fig. 13. This shows the variation in the skew quadrupole strengths in a model fitted to the measured mode II dispersion, over a range of applied changes in strength of skew quadrupole 48W. The agreement between the applied change and the fitted strength is good over a wide range of values, even where there is significant coupling in the machine.

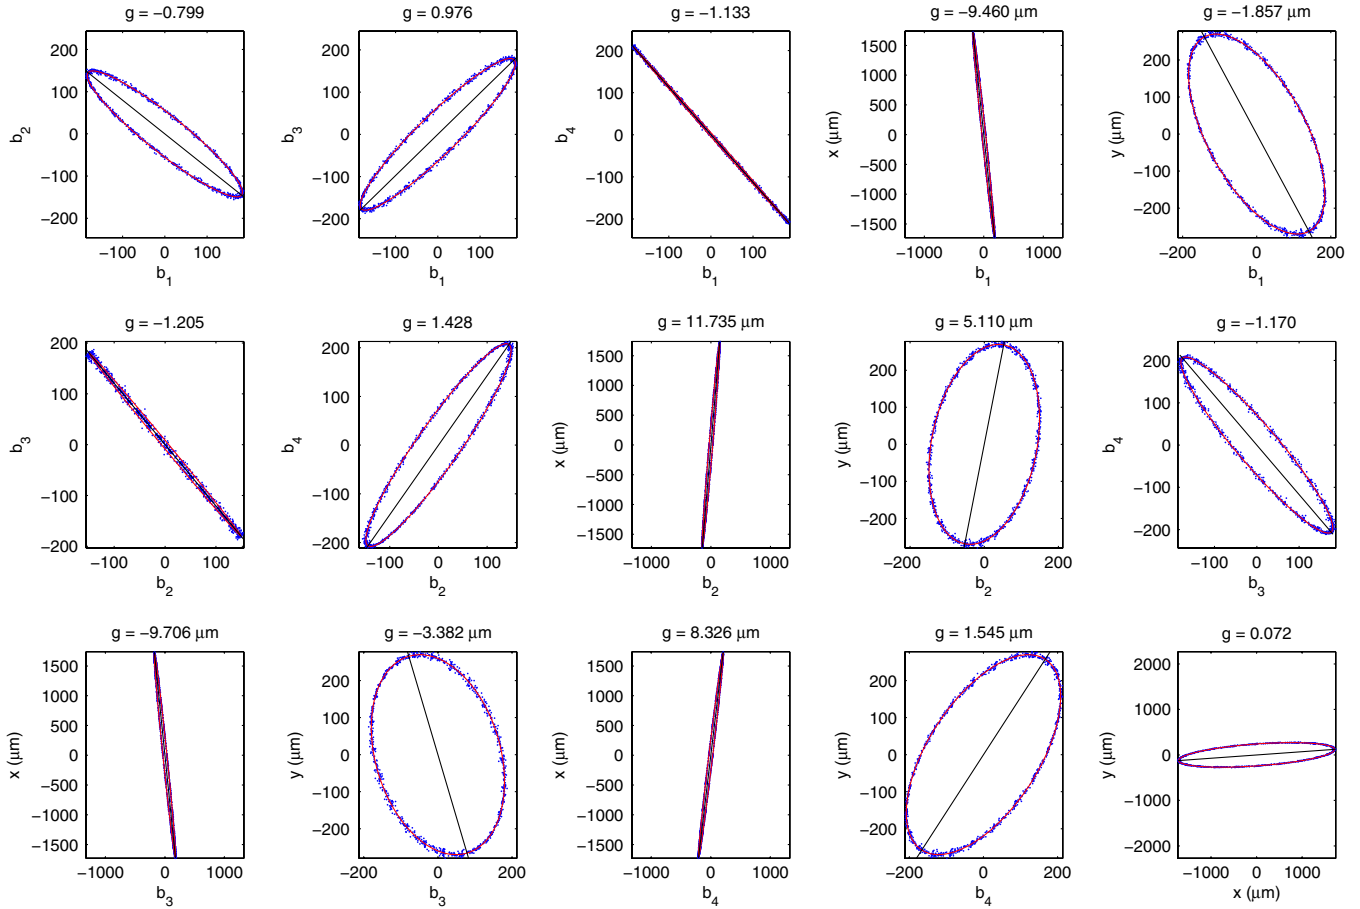


FIG. 12. Correlation plots for the button signals from BPM 7W during resonant excitation of normal mode I, with the strength of skew quadrupole 48W changed by $\Delta k = -0.023 \text{ m}^{-2}$. The blue points show the data from the BPM; in each correlation plot, the points are fitted with a red ellipse; the major axis of each ellipse is shown with a black line. Compare with the simulation results shown in Fig. 4.




FIG. 13. Variation in skew quadrupole strengths in a model fitted to the measured mode II dispersion, over a range of applied changes in strength of skew quadrupole 48W (SKQ48W). The fitted versus applied strength of skew quadrupole 48W (black line, with data points indicated by circles) shows a linear behavior, with gradient close to 1 (shown by a dashed line). Other skew quadrupoles (colored lines) show some (approximately) linear response, but with much smaller gradients.

The fitted model strengths of other skew quadrupoles also vary (approximately) linearly with the strength of skew quadrupole 48W applied in the machine, but the dependence is relatively weak. There are three skew quadrupoles for which the fitted gradients are between 20% and 30% of the gradient fitted for skew quadrupole 48W: these are the skew quadrupoles 47W, 48E, and 47E, which are exactly the skew quadrupoles in the same region of the machine as skew quadrupole 48W (which is at $s = 377 \text{ m}$ in Fig. 2). As already mentioned, the fitted changes in strength of these skew quadrupoles may be the result of some degeneracy, or the result of some discrepancy between the model and the machine: further studies using a carefully calibrated model would be needed to identify the exact cause. However, it appears to be possible to use the model to predict reasonably well the effect of changes in skew quadrupole strength on the mode II dispersion in the machine. This suggests that a response matrix (between skew quadrupole strengths and mode II dispersion measured at the BPMs) calculated from the ideal model should provide a reasonable basis for correcting the mode II dispersion in the machine.

C. Low-emittance tuning by correction of mode II dispersion

To explore the effect of the mode II dispersion on the vertical emittance in CEsrTA, two attempts were made at correcting the measured mode II dispersion. The beam diagnostics in CEsrTA includes a fast (turn-by-turn) x-ray synchrotron radiation beam-size monitor (xBSM) [17], with resolution for beam-size measurements less than 10 μm . Unfortunately, on the first attempt at mode II dispersion correction, this monitor was unavailable, which meant that an estimate of the emittance could not be made. However, the results of dispersion and coupling measurements (shown in Table II) appeared promising. Initially, the storage ring was tuned using the conventional procedure, involving correction of the orbit, dispersion, and coupling by optimizing the strengths of the vertical steering magnets and skew quadrupoles. Following this initial tuning, all skew quadrupoles were turned off. The normal mode calibration factors were measured, and the mode II dispersion was measured. The coupling, characterized by the rms value of \bar{C}_{12} [9,18,19], was also measured. Using the response matrix between the skew quadrupoles and mode II dispersion, at the BPMs, computed from the ideal model, the skew quadrupole strengths needed to correct the measured mode II dispersion were calculated. These strengths were applied, and the BPMs recalibrated. The measurements of mode II dispersion and coupling (rms \bar{C}_{12}) were then repeated. As can be seen from Table II, there was a significant reduction in both the measured mode II dispersion (in line with the prediction from the model), and the coupling.

Once the xBSM was available, a further attempt at low-emittance tuning using mode II dispersion correction was made. On this occasion, the following procedure was applied.

TABLE II. Results of an initial test of mode II dispersion correction in CEsrTA. The rms values are the root mean squares of the measurements at all BPMs.

Machine conditions	rms η_{II} (mm)	rms \bar{C}_{12}
After initial tuning, all skews off	31	0.014
After mode II dispersion correction	18	0.007

TABLE III. Results of a test of emittance tuning in CEsrTA using correction of mode II dispersion. The rms of the mode II dispersion is calculated over the values measured at the BPMs. The vertical beam size is measured using an x-ray beam-size monitor (xBSM). The vertical dispersion (η_y) is measured at the BPM closest to the source point for the xBSM.

Correction step	Machine conditions	rms η_{II} (mm)	rms vertical beam size (μm)	η_y (mm)	ϵ_y (pm)
1	Following initial tuning (orbit, dispersion and coupling correction)	38	21	6	14
2	All skew quadrupoles turned off	32	27	3	24
3	After first mode II dispersion correction	32	22	8	14
4	After second mode II dispersion correction	31	28	13	21

(1) The storage ring was tuned using the conventional procedure, based on correction of orbit, vertical dispersion, and betatron coupling, using vertical steering magnets and skew quadrupoles. The emittance was estimated from measurements of the vertical beam size using the xBSM.

(2) All skew quadrupoles were turned off, and the emittance again estimated from the xBSM.

(3) Normal mode BPM calibration data were collected using the results of the calibration, the mode II dispersion at the BPMs was measured. Using a response matrix computed from the design model, skew quadrupole strengths to correct the mode II dispersion were calculated and applied. Note that equal weighting was given to the dispersion measured at each BPM. The emittance was again estimated from the xBSM.

(4) A further iteration of normal mode calibration and correction was applied, and a final estimate of the emittance made from the xBSM.

The results are shown in Table III. The emittance is estimated using the formula

$$\sigma_y = \sqrt{\beta_y \epsilon_y + \sigma_\delta^2 \eta_y^2}, \quad (25)$$

using the nominal vertical beta function at the xBSM source point of $\beta_y = 30$ m and the rms energy spread at 2 GeV of $\sigma_\delta = 9 \times 10^{-4}$. The vertical dispersion at the xBSM was measured by observing the change in the centroid beam position with respect to a change in beam energy induced by a change in rf frequency. We assume that the betatron coupling makes a negligible contribution to the vertical beam size.

At first sight, the results appear disappointing. First, the measured mode II dispersion decreased after turning off the skew quadrupoles following the initial (conventional) emittance tuning. However, this is not inconsistent with an increase in the emittance, since the mode II dispersion is measured only at the BPMs, and not in the dipoles and wigglers where it really has significance for the emittance. It should also be noted that the changes in mode II dispersion that were measured were small compared with the accuracy of the measurement (of approximately 12 mm, based on the repeatability of the measurement with no changes in machine condition between measurements). Second, the attempt at correcting the mode II dispersion

appears to have had little effect on the dispersion. But this is in fact consistent with the model, given that the measured mode II dispersion could not be well fitted using the model. Figure 14 shows the strengths of the skew quadrupoles in the model fitted to the measured dispersion (top), a comparison between the mode II dispersion fitted in the model and the mode II dispersion measured in the machine (center) and the difference between the measurement and the fit (bottom). In this case, the model achieves a poor fit to the measurement: the strengths of the skew quadrupoles are weak, and the residual between the measurement and the fitted model is almost as large as the measured mode II dispersion. This indicates that the sources of mode II dispersion in the machine (including orbit distortion, orbit offsets in the sextupoles, and quadrupole tilts) are not well fitted by the skew quadrupoles. In this case, the skew quadrupoles are not expected to be effective in correcting the mode II dispersion.

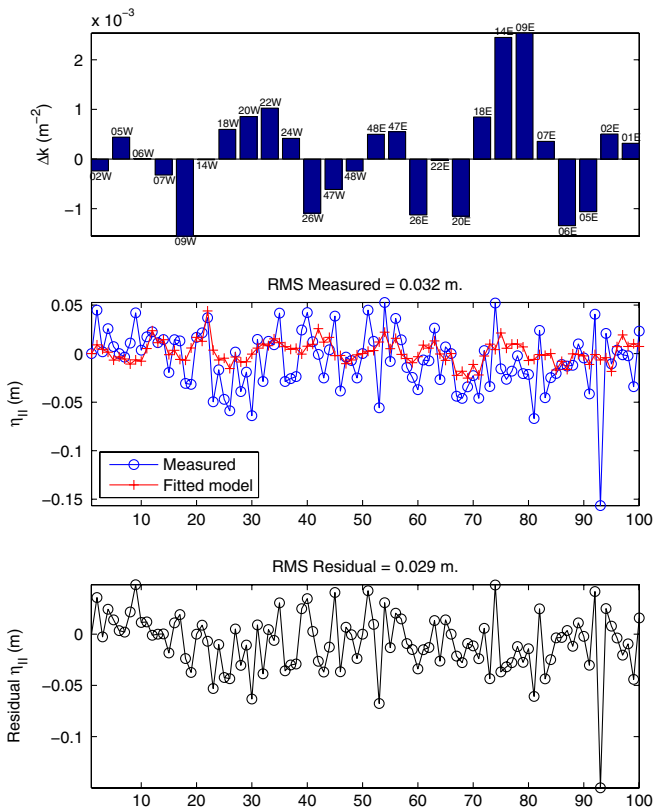


FIG. 14. Results of fitting a model to a measurement of the mode II dispersion. The measurement was made after all skew quadrupoles were turned off, following an initial low-emittance tuning using conventional techniques (vertical orbit and dispersion correction, and coupling correction). Top: fitted skew quadrupole strengths. Middle: measured and fitted mode II dispersion. Bottom: residual between measured and fitted mode II dispersion. Note that in this case, a poor fit is achieved in the model to the mode II dispersion. This suggests that the skew quadrupoles would not be effective in correcting the real sources of mode II dispersion in the machine.

It must also be noted that after both corrections, the measured rms mode II dispersion is almost twice that expected from the simulations (see Fig. 6). This could be because the measurement error was larger on this occasion than previously, or because the model was not sufficiently close to the machine. Although there appeared to be good agreement between the model and the machine for the effect of changing one particular skew quadrupole (SKQ48W), this was for relatively large changes in the skew quadrupole strength. It is possible that for precise correction of the mode II dispersion, at the level achieved in the simulations, a more accurate model of the machine would be required.

Nevertheless, it appears that the first correction, although it had little apparent effect on the mode II dispersion, did lead to some reduction of the vertical emittance. But a further attempt at correction appears to have led to an increase in the vertical emittance. There are several possible explanations for this. One is that the correction, by attempting to reduce the overall rms dispersion, may have resulted in a reduction in the mode II dispersion at some locations, but at the cost of a small increase in the mode II dispersion at others. Since the mode II dispersion tends to be small where the horizontal dispersion is small, i.e., in the dipoles and wigglers, the most likely locations for a small increase in the mode II dispersion are at the locations where synchrotron radiation is generated: this would then lead to an increase in the mode II emittance. One way to address this issue would be to apply a higher weight to the dispersion measurements at the BPMs close to the wigglers and dipoles, and a lower weight to the dispersion measurements at locations where the horizontal dispersion is large (but where no synchrotron radiation is generated).

Another possible explanation could simply be that the mode II dispersion measurement is not sufficiently accurate, or that the response matrix (based on the ideal model) is not sufficiently close to the actual machine. Based on reproducibility tests, the accuracy of the mode II dispersion measurement is believed to be approximately 12 mm: this is large compared to the changes in mode II dispersion seen after applying the corrections. In this case, it is possible that there was in fact a real increase in the mode II dispersion, with a consequent increase in the emittance.

Following the experimental studies, further simulations were carried out to investigate the possible reasons for the increase in emittance following a second iteration of the mode II dispersion correction. It was found that using the procedure described in Sec. III D (and which was used also in the experimental studies), the mode II emittance does indeed fluctuate by large amounts on successive iterations: a typical example for a particular seed of machine errors is shown in Fig. 15 (top). Applying a higher weight in the correction to the dispersion measurements at BPMs close to the dipoles and wigglers can reduce the fluctuations to some extent, but the effect is not strong, or completely

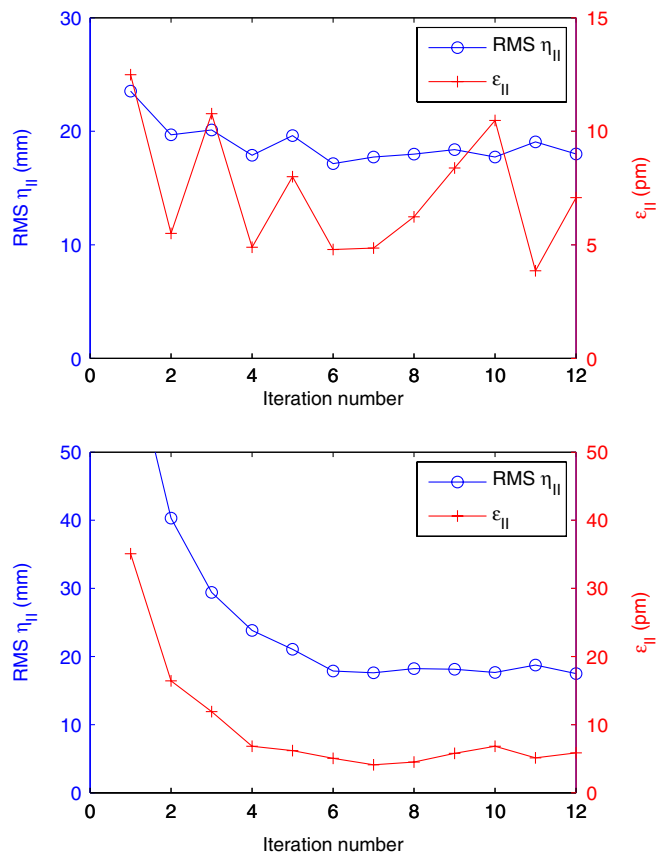


FIG. 15. Results from simulation of iterating a mode II dispersion correction for tuning the mode II emittance. Top: with a scaling factor of 1 applied to the computed changes in skew quadrupole strengths, there are significant fluctuations in the mode II emittance in successive iteration. Bottom: with a scaling factor of 0.4, the fluctuations are smoothed out, and the emittance (and dispersion) converge to some final values. The same magnitude of machine errors was applied as used in the simulations described in Sec. III D. Note that the same seed of random errors was used to generate both top and bottom plots.

reliable. The effects of random errors in the dispersion measurement (which are included in the simulations) can be reduced to some extent simply by applying a scaling factor to the computed correction. That is, the computed changes in skew quadrupole strengths are multiplied by a chosen scaling factor (with value between 0 and 1) before they are applied: this effectively averages the dispersion correction over several iterations. A typical example of the results, for a scaling factor of 0.4 and a particular seed of machine errors, is shown in Fig. 15 (bottom). As expected, a small value for the scaling factor means that a larger number of iterations are required before the emittance converges. However, the fluctuations in the emittance between successive iterations are effectively smoothed out, and the final values achieved are similar to the smallest values achieved with a scaling factor of 1. This modification of the tuning technique (i.e. the implementation of a scaling factor) would be easy to implement in practice; however,

some further investigation would be required to optimize the value of the scaling factor, to minimize the number of iterations required to reach convergence while still achieving a reduction in the emittance fluctuations on successive corrections. In simulations, a value of around 0.4 for the scaling factor seemed to give good results.

Finally, we note that the emittance values estimated from the beam measurements are broadly consistent with the results of simulations presented in Sec. III. With realistic machine errors in CEsrTA, a correction based purely on correction of the orbit and *vertical* dispersion is expected to achieve vertical emittances of order 50 pm or more, depending on the size of the errors and, in particular, the BPM gain errors (see Fig. 6). In the conventional tuning techniques applied at CEsrTA, in particular through orbit and dispersion correction, and correction of the coupling (characterized by \bar{C}_{12}), emittances of around 7 pm have been achieved. Simulations suggest that tuning based on correction of the mode II dispersion could give emittances of that order or below: however, this depends on the seed of random errors. In the distribution shown in Fig. 6, a little over 50% of seeds of machine errors result in a final emittance below 10 pm, while 36% of seeds result in a final emittance between 10 and 20 pm.

V. SUMMARY AND CONCLUSIONS

Calibration of BPMs using normal mode beam excitation offers the possibility of a novel technique for low-emittance tuning in storage rings. This technique has a number of potential advantages over conventional techniques, such as ORM analysis. These advantages include: (i) The calibration data can be collected and processed quickly (within a few minutes) and easily. The only instrumentation requirements are: a means for resonant excitation of transverse beam motion; and BPMs capable of returning turn-by-turn signals from each button. (ii) Since the calibration takes only a short time, it can be performed after any significant change in beam conditions, such as changes in orbit, coupling, or beam current. This ensures that the calibration is always accurate for the present machine conditions. (iii) Using calibrated BPMs, low-emittance tuning consists of minimizing the mode II dispersion, which can be measured directly using conventional techniques (e.g. observing the change in beam position with respect to a change in beam energy). This is again a relatively fast and straightforward procedure. (iv) The technique can be applied as easily to a large ring as to a small ring. (v) Since the BPMs are calibrated directly from observation of normal mode beam motion, the calibrated BPM readings are insensitive to BPM alignment errors, or gain errors arising (for example) from poor characterization of the electronics. Such errors can be difficult to determine with good accuracy, and can vary over time, but can be an important limitation on conventional low-emittance tuning techniques.

Simulations for CEsrTA (and for the KEK Accelerator Test Facility) suggest that the technique is very robust, and could be used to achieve very low beam emittances: the distribution shown in Fig. 6 for the emittance following normal mode dispersion correction has a little over 50% of seeds resulting in a final emittance below 10 pm. The basic principles behind the BPM calibration have been demonstrated in experimental tests at CEsrTA. It has been shown possible to collect good quality data reliably and routinely; and tests of the calibration (principally, comparisons of the frequency spectra of readings from uncalibrated and calibrated BPMs) indicate that following the calibration procedure, the BPMs do indeed read beam motion along the normal mode axes. Predictions of a model (based on an ideal machine) have been shown to be in good agreement with measurements of the response of the mode II dispersion to changes in skew quadrupole strengths. This provides the basis for a low-emittance tuning procedure based simply on the correction of the mode II dispersion using skew quadrupoles.

So far, attempts to apply the normal mode calibration to low-emittance tuning of CEsrTA have met with limited success. Using conventional techniques, emittances of around 7 pm (significantly below the initial project goal) have already been achieved. Simulations suggest that it may be possible, depending on the exact distribution of machine errors, to achieve emittances of that order or less, using correction of the mode II dispersion. However, in a test of the technique, it was found that the measured mode II dispersion could not be well fitted using the available skew quadrupoles. Although the low value of the emittance achieved after conventional tuning (involving correction of the orbit and vertical dispersion, and separate correction of the coupling) was reproduced after tuning based on correction of the mode II emittance, further reduction of the emittance was not demonstrated. It is not possible to make a direct detailed comparison between the mode II dispersion measured in the machine and the mode II dispersion predicted by a model, because the dispersion depends on the exact set of errors present, and these are not known. However, the measured change in mode II dispersion resulting from a known change in skew quadrupole strength is consistent with the change predicted from a model (taking into account that an ideal model has been used, and not a model that has been calibrated against machine measurements). Simulations can be used to predict the distribution of final emittance, after correction of the mode II dispersion, over many seeds of random error: the measured emittance of 14 pm following application of the tuning procedure on the real machine has a reasonable probability within the distribution (shown in Fig. 6) calculated using realistic values for the magnitudes of the errors.

From simulations, it seems possible that the reason the emittance was not further reduced by additional iterations of the correction procedure was because of inaccuracies in

the dispersion measurement, and differences between the response matrix in the model and the machine. However, it appears (from simulation) that an effective way to improve the technique could simply be to apply a scaling factor to the computed skew quadrupole strengths, before applying them to the machine. Other improvements to the correction procedure are also possible, for example, larger weights could be applied to the dispersion measurements at BPMs close to the wigglers and dipoles. Further studies would be needed to optimize the correction procedure: however, the results obtained so far provide a proof of principle for the technique. Our studies have focused on an operating storage ring (CEsrTA) in its present configuration. Although a number of upgrades to the diagnostics and correction systems have been made in recent years to allow the machine to achieve very low vertical emittances, these upgrades have inevitably been limited by the availability of resources, and constraints from the original machine design and construction. It would be interesting to investigate further modifications and upgrades (for example, the possibility of additional skew quadrupole correctors) that may allow achievement of still smaller vertical emittances.

ACKNOWLEDGMENTS

We wish to thank our colleagues at CEsrTA, in particular Mark Palmer, Bob Meller, and Matt Rendina, for helpful discussions and technical support. This work was supported by the Science and Technology Facilities Council, UK; and by the National Science Foundation and the U.S. Department of Energy under Contracts No. PHY-0734867 and No. DE-FC02-08ER41538.

-
- [1] Y. Honda *et al.*, *Phys. Rev. Lett.* **92**, 054802 (2004).
 - [2] Å. Andersson *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **591**, 437 (2008).
 - [3] M. Böge *et al.*, in *Proceedings of the 2009 Particle Accelerator Conference, Vancouver, BC, Canada* (IEEE, Piscataway, NJ, 2009), pp. 2279–2281.
 - [4] R. Dowd, M. Boland, G. LeBlanc, and Y-R. E. Tan, *Phys. Rev. ST Accel. Beams* **14**, 012804 (2011).
 - [5] J. Safranek, *Nucl. Instrum. Methods Phys. Res., Sect. A* **388**, 27 (1997).
 - [6] The ILC Global Design Effort, <http://www.linearcollider.org/about/Publications/Reference-Design-Report>.
 - [7] M. A. Palmer *et al.*, *Proceedings of the 2009 Particle Accelerator Conference, Vancouver, BC, Canada* (Ref. [3]), pp. 4200–4204.
 - [8] D. Edwards and L. Teng, *IEEE Trans. Nucl. Sci.* **20**, 3 (1973).
 - [9] D. Sagan and D. Rubin, *Phys. Rev. ST Accel. Beams* **2**, 074001 (1999).
 - [10] A. Wolski, *Phys. Rev. ST Accel. Beams* **9**, 024001 (2006).
 - [11] K. Satoh and M. Tejima, in *Proceedings of the Particle Accelerator Conference, Dallas, TX* (IEEE, New York, 1995), pp. 2482–2484.

-
- [12] D. L. Rubin *et al.*, *Phys. Rev. ST Accel. Beams* **13**, 092802 (2010).
- [13] A. Terebilo, in *Proceedings of the 2001 Particle Accelerator Conference, Chicago, IL* (IEEE, New York, 2001), pp. 3203–3205.
- [14] K. Ohmi, K. Hirata, and K. Oide, *Phys. Rev. E* **49**, 751 (1994).
- [15] K. Panagiotidis, Ph.D. thesis, University of Liverpool, 2010.
- [16] J. Shanks *et al.*, Proceedings of the 2009 Particle Accelerator Conference, Vancouver, BC, Canada (Ref. [3]), pp. 2754–2756.
- [17] D. Peterson *et al.*, *Proceedings of the 2010 International Particle Accelerator Conference, Kyoto, Japan* (ICR, Kyoto, 2010), pp. 1194–1196.
- [18] D. Sagan, R. Meller, R. Littauer, and D. Rubin, *Phys. Rev. ST Accel. Beams* **3**, 092801 (2000).
- [19] D. Sagan, *Phys. Rev. ST Accel. Beams* **3**, 102801 (2000).