Presented at e^+e^- Factories 1999, Sept. 21-24, 1999, KEK Laboratory **CBN 99-29**

Beam Diagnostics and Correction at CESR *

Stuart Henderson † and David Sagan
Laboratory of Nuclear Studies, Cornell University, Ithaca NY, 14853, USA

Abstract

The implementation of a beamloss diagnostic system at CESR is described. This system measures beam motion and RF system parameters at the point of a beamloss. The information is useful in identifying beam current limiting phenomena in CESR. A unified optics measurement and correction package is also described, and examples of its performance are presented.

1 INTRODUCTION

Beam diagnostic instrumentation has been vital in increasing the stored beam current in CESR and improving the luminosity performance. Central to understanding the factors which limit CESR beam currents is the ability to observe beam motion, RF system parameters and vacuum condition on a fast time scale at the point of a beamloss or other transient. Such observations are essential to diagnosing the underlying current limitations. Two systems are in use at CESR to provide beamloss diagnostics – one on a fast turn-by-turn time scale and another on a slower 10-30 Hz time scale.

An important factor responsible for increases in CESR's beam-beam tuneshift parameter in the last several years (see [1, 2] in these proceedings) is the ability to reliably and repeatedly correct the lattice functions and transverse coupling. A new unified optics measurement and correction system has recently been developed at CESR. The system acquires orbit, betatron phase, transverse coupling and dispersion data and includes an analysis package for lattice correction and identification of isolated magnetic field errors.

The implementation of these beam diagnostic and lattice correction tools at CESR is described.

2 BEAM CURRENT LIMITATIONS

Since 1996 CESR beam currents have been limitetd by two factors: a longitudinal coupled bunch instability and RF system trips associated with high-current operation. At any given time it is important to identify the beam current limitations so that corrective action can be taken.

To put the issue in perspective, a review of a recent running period is instructive. A new superconducting RF cavity was installed in CESR in March 1999 [1]. A short 6-week running period in April-May 1999 was devoted to commissioning the new cavity, exploring beam current limits and running beams for the CHESS (Cornell High-Energy Synchrotron Source) experimental users. During this commissioning period, we undertook aggressive beam processing of the RF system and exploration of current limitations. During this 47 day period, we logged 301 beamlosses (6.4/day) and 250 RF system trips (5.3/day). The rates during CHESS running (when beam currents were not being pushed as aggressively) were 2.8 beamlosses/day and 1.9 RF trips/day. Clearly, with this volume of beamlosses, there is much to be learned from the beam and RF system at the point of a beamloss. In order to understand the limitations and causes of beamlosses and RF trips a diagnostic system was developed in 1996.

3 BEAMLOSS DIAGNOSTICS

The beamloss diagnostic system is described in detail in [3]. The sytem is a turn-by-turn data acquisition system based on COMET^{TM} VME multichannel analog digitizer boards manufactured by Omnibyte Inc.[1] Four input channels per board are digitized with individual 12-bit ADCs at rates up to 5 MHz. Digitized data from each channel is sequentially stored in its own 128 kbyte buffer. Each channel has a buffer depth of 64k samples. For beamloss monitoring the boards are operated in a post-trigger mode in which data is continuously written in a circular buffer fashion until the sytem is triggered to stop when a beamloss occurs.

Horizontal and vertical beam positions as well as bunch current obtained from a single dedicated BPM are recorded for a selection of bunches in CESR. In addition, RF system diagnostic signals (forward power, reflected power, cavity field, cavity phase and tuning angle) for each cavity are digitized along with the beam position information. In this way, beamloss events in which the RF sytem trips can be studied to determine if beam motion contributed to the trip, or if the beam was stable at the point of a trip.

A second general diagnostic system records a set of signals (available from the CESR database) at a slower rate (10-30 Hz). This system is quite useful for analyzing problems which are too slow for the turn-by-turn system to observe.

* Work supported by the National Science Foundation.
† Email: sdh9@cornell.edu

[1] Omnibyte, 245 West Roosevelt Road, West Chicago, Illinois, 60185

3.1 Longitudinal Instability

A series of beamlosses in November 1998 had the signature shown in Figure 1. The top plot shows the horizontal beam position (arbitrary units) for 1024 turns prior to a beamloss, and the bottom shows the corresponding FFT (plotted as the natural logarithm of the power versus the fractional tune). The FFT shows a large peak at the synchrotron tune

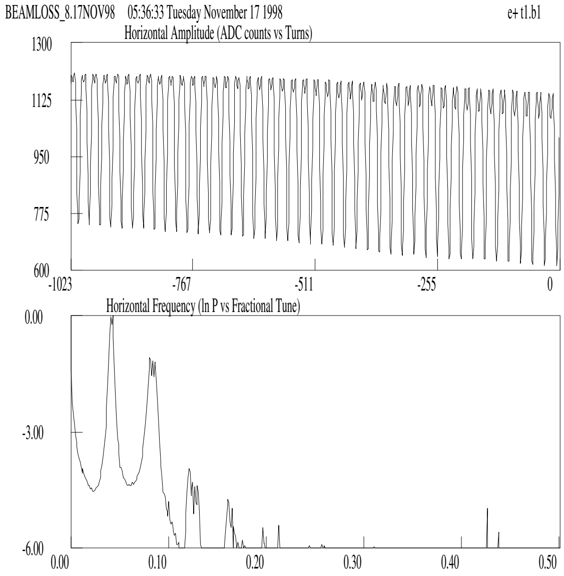

Figure 1: Longitudinal Instability. The top plot shows the horizontal beam position versus turn number prior to a beamloss. The bottom plot shows the FFT of the turn-by-turn position data plotted versus fractional tune. The synchrotron tune is 0.045.

(0.045) and higher harmonics. The synchrotron motion is readily observed in the horizontal beam position measurement due to dispersion at the pickup. The longitudinal motion grows without bound until the beam is lost. The higher harmonics of the synchrotron tune are due to the large horizontal displacement at the beam pickup due to the electrostatic pretzel. Thus, a beamloss due to a longitudinal instability has an obvious, readily identifiable signature. Since we needed longitudinal feedback to maintain beam stability and these beamlosses were occuring at beam currents lower than had been obtained before, the longitudinal feedback system was suspected. Indeed, it was discovered that the longitudinal feedback system's pedestal correction circuit was periodically resetting and subsequently exciting longitudinal motion producing a beamloss.

3.2 RF Waveguide Arc

In the commissiong period mentioned above for the third of CESR's SRF cavities (designated W1), beam currents were limited at one point due to a limit in the power which could be transmitted through the RF window and coupler. At this point, numerous beamlosses occured due to RF system trips at a hard limit of 140 kW in travelling wave power through the W1 cavity. These trips had the signature shown in Figure 2. The plot shows the forward power, cavity field and reflected power for the W1 cavity, and for comparison the reflected power for the E2 cavity (located on the other side of CESR) for $\sim$56 turns bracketing the beamloss. (The x-axis in this plot is sample number. Each sample is $\sim$280 ns and there are 9 samples/turn.) Additional information indicated that the cause for the trip was high reflected power detected at the RF klystron. The earliest indication in Figure 2 is a small rapid increase in the W1 reflected power signal, accompanied by a drop in the cavity field. Subsequently the reflected power rapidly increases.

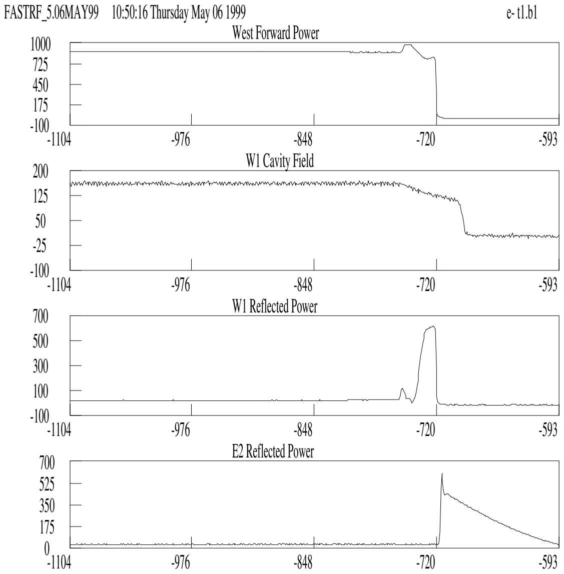

Figure 2: W1 waveguide arc event. Shown are the forward power, cavity field and reflected power for the W1 cavity and the reflected power for the E2 cavity. The x-axis is sample number, where a sample is taken every 280 ns. The plot spans 56 CESR turns ($\sim$150 μs).

The forward power is then removed as the transmitter turns off due to high VSWR. The beam continues to circulate for another 20 turns or so, but there is no reflected power from the beam driving the cavity or from the cavity releasing its stored energy, as is observed in the E2 cavity. These observations are consistent with an arc in the RF waveguide. Forward power no longer propagates to the cavity so the cavity field begins to drop. Meanwhile, RF power is reflected due to the improper impedance match, so the reflected power rises rapidly. Some of this large reflected power reaches the transmitter due to a non-ideal RF circulator and the transmitter trips on high VSWR, turning off forward power to all the CESR cavities. The arc travels along the waveguide, eventually reaches the RF coupler at which point the cavity field drops suddenly (sample 700) as the cavity is effectively "shorted." The beam continues to circulate for about 20 turns (off the right-hand side of this plot). The circulating beam provides power to the cavity which goes into the arc or is reflected from the arc, but in any case doesn't register as reflected power outside the cavity. In contrast, the reflected power from the E2 cavity rises abruptly once the forward power is removed since the beam is driving the cavity, and the cavity is releasing its stored energy.

It was believed that these arcs were occuring in the section of the waveguide which makes the transition to cryogenic temperature, and was thought to be due to adsorbed gases on the waveguide surface. The cavity was warmed to 50K to remove adsorbed hydrogen. Following this warmup and subsequent high-power pulsed processing the beam power delivered through W1 was increased from 140 kW to 260 kW[1].

3.3 Quench from Cavity Regulation

During the W1 SRF cavity commissioning we also encounterd a period with numerous W1 quench trips. A quench is signalled when the cavity field drops below a minimum value. Figure 3 shows three traces over a 206 sec period (left) and an expanded window of 8 sec (right) bracketing a beamloss which occured during injection of beam into CESR. The lower trace shows the beam current, the middle the reflected power from the W1 cavity, and the top the cavity field. The cavity regulation maintains a constant field until the reflected power becomes large enough, at which point the controls switch from field to reflected power regulation. During these tests the SRF cavities were running at gradients lower than design (5-6 MV/m). As a result, at operating currents the beam-cavity system is overcoupled and the reflected power grows quadratically with beam current. As the beam current increases the cavity regulation can be seen switching from field to reflected power on the middle trace just prior to the beamloss. Once this happens, the cavity field drops with increasing beam current and becomes low enough that the quench threshold is crossed. In fact, the last points before the beam is lost in the expanded view of the W1 cavity field are seen to drop. This difficulty has been eliminated by a modification to the RF controls which inhibits power regulation if beam current is present.

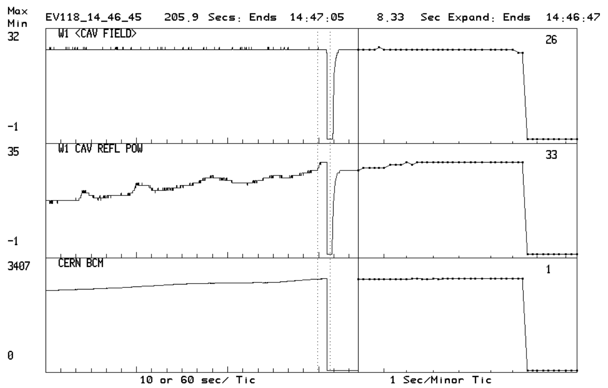

Figure 3: The W1 cavity field and reflected power and the total beam current are shown bracketing a W1 quench trip.

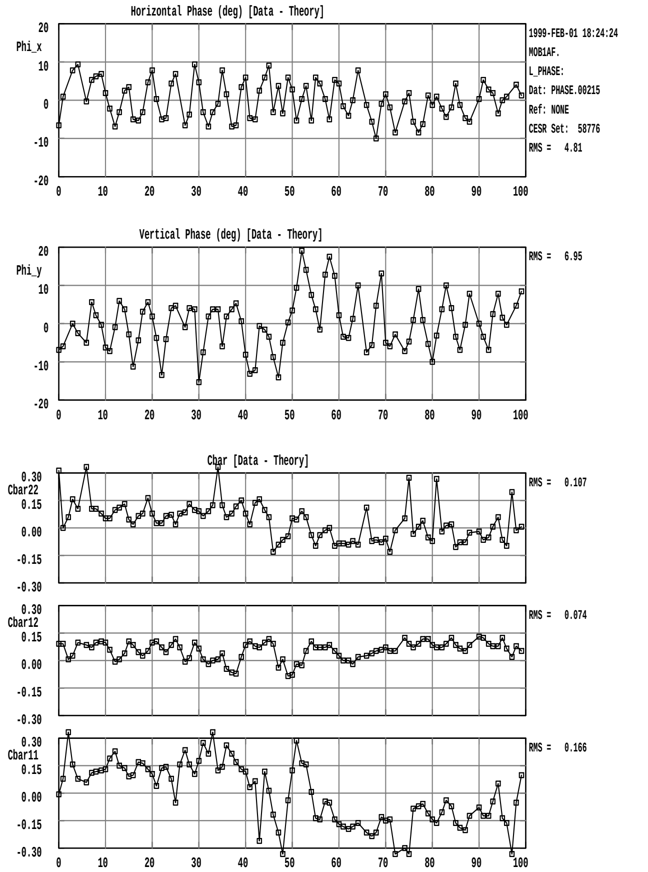

Figure 4: Betatron phase and coupling data taken after loading a new experimental lattice. The top two plots are the difference between the measured and theoretical horizontal and vertical betatron phases. The bottom three plots are the coupling parameters (see text).

4 LATTICE CORRECTION AND DIAGNOSTICS

4.1 System Description and Lattice Correction

A unified optics measurement, analysis and correction package has been developed. The system acquires orbit, betatron phase, transverse coupling and dispersion data. The betatron phase and coupling data is acquired by shaking the beam in both planes at a betatron sideband and digitizing the response at each beam button on every CESR BPM detector [4, 5]. The digitized data is analyzed to determine the relative phase between shaker and detector for the phase data. Analysis of out-of-shaking-plane position information yields the coupling data. Dispersion data is acquired by modulating the CESR RF frequency and measuring the orbit difference.

An example of betatron phase and coupling data acquired recently is shown in figure 4. The Figure shows a measurement taken upon loading a new experimental accelerator lattice. The top two plots display the difference between the measured betatron phase advance and that specified by the theoretical lattice (in degrees) for the horizontal and vertical planes, respectively, as functions of the BPM number. The difference in degrees is plotted versus BPM number. For orientation, BPM 0 is adjacent to the interaction point (IP), BPM 50 is opposite the IP, and BPM 99 is adjacent the IP on the other side. Betatron phase errors of this magnitude translate into β-function errors on the order of 30-40%. The bottom plots shows the three coupling parameters [6, 5] $\overline{C}_{22}$ (top), $\overline{C}_{12}$ (middle) and $\overline{C}_{11}$ (bottom). The parameter $\overline{C}_{22}$ is related to the tilt angle of the real space ellipse and $\overline{C}_{12}$ is related to the ratio of minor to major axes. For orientation, full coupling implies peak $\overline{C}$ parameters of $\sim$0.7.

The heart of the analysis package is an optimizer which minimizes a merit function specifically constructed for the problem at hand. The minimization algorithm is a conjugate gradient method of Fletcher-Reeves/Polak-Ribiere which is contained in Numerical Recipes (subroutine FRPRMN). A typical analysis of the data in figure 4 would use all quadrupole strengths and interaction region quadrupole rotation angles as variables. (The IR quadrupoles are mounted on rotators. Non-zero rotation angles are needed to compensate the coupling introduced by the detector solenoid field). The optimizer adjusts the variables in an attempt to fit the data by minimizing a merit function, M, which is constructed as follows:

$$
\begin{aligned}
M \;=\; & \sum_{i\in\, det} w_x \left[\phi_{x,i}(\mathrm{data}) - \phi_{x,i}(\mathrm{fit})\right]^2 + \\
& w_y \left[\phi_{y,i}(\mathrm{data}) - \phi_{y,i}(\mathrm{fit})\right]^2 + \\
& w_{\overline{C}} \left[\overline{C}_{12,i}(\mathrm{data}) - \overline{C}_{12,i}(\mathrm{fit})\right]^2 + \\
& \sum_{j\in\, quads} w_k \left[k_j(\mathrm{fit}) - k_j(\mathrm{theory})\right]^2 + \\
& \sum_{n\in\, IRquads} w_\theta \left[\theta_n(\mathrm{fit}) - \theta_n(\mathrm{theory})\right]^2 \qquad (1)
\end{aligned}
$$

where $\phi_{x,i}$ is the horizontal betatron phase at detector i, $\phi_{y,i}$ is the vertical betatron phase at detector i, $\overline{C}_{12,i}$ is the coupling parameter at detector i (the $\overline{C}_{11}$ and $\overline{C}_{22}$ data are not used because, for instrumental reasons, they tend to be noisier), k_j is the strength of quadrupole j, and θ_n is the rotation angle of IR quadrupole n, and the w's are the user-defined relative weights. The "data" refers to measured values, and the "fit" is obtained from a lattice model of the ring. The "fit" quadrupole strengths and IR quadrupole rotation angles are used to compute the fit betatron phases and coupling, which in turn are used to calculate the merit function. The inclusion of corrector terms (the last two terms in equation 1) prevents degeneracies from leading to pathological solutions, such as may occur when two correctors have a similar effect on the fit data.

Lattice correction is then a matter of removing the errors predicted by the fitting results. After the correction based on the data in figure 4 another optics measurement was performed and yielded the results shown in figure 5. Already, at this point, the lattice is rather well corrected. The RMS phase errors have been reduced by a factor of 2.5 (horizontal) and 3.2 (vertical). The measurement, analysis and correction process may take as little as five minutes so that multiple lattice corrections are often performed in a reasonable amount of time.

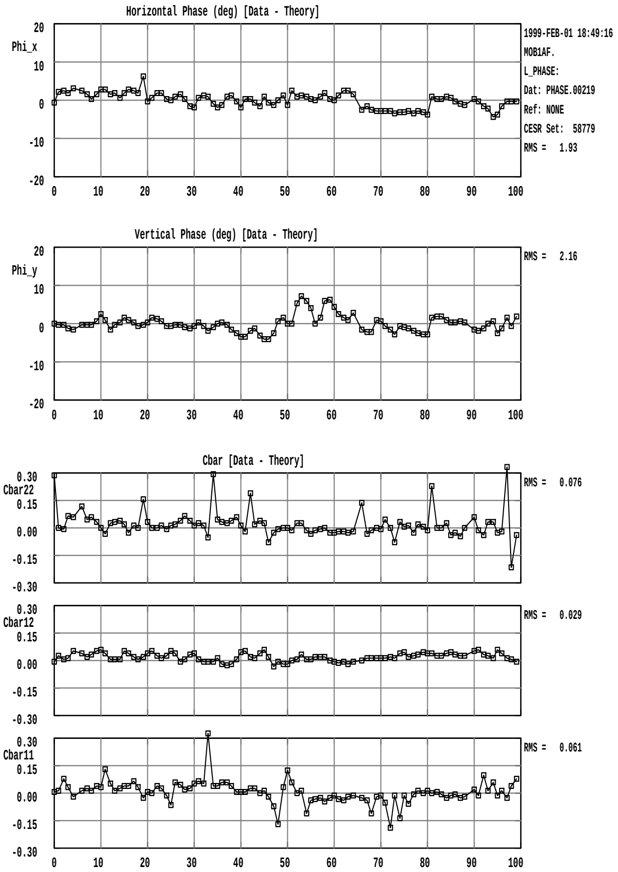

Figure 5: Betatron phase and coupling data taken after correction of the lattice based on the phase/coupling data shown in figure 4.

4.2 *Isolated Magnetic Field Errors*

In a case when one expects one or only a few isolated magnetic field errors, a "wave" analysis [7, 6] can be quite useful. In such an analysis the data (orbit, betatron phase or coupling) is fit over a region of beam detectors to a freely propagating wave, assuming no errors are present. When the fit is propagated around the machine, isolated errors appear as a location at which the fit deviates from the data. In fact, if only a single error is present, then the exact location of the error may be determined.

Such an example is shown in Figure 6. After a scheduled maintenance period in October 1998, we found it difficult to correct the local transverse coupling with only IR quadrupole angle adjustments. A "wave" analysis of the

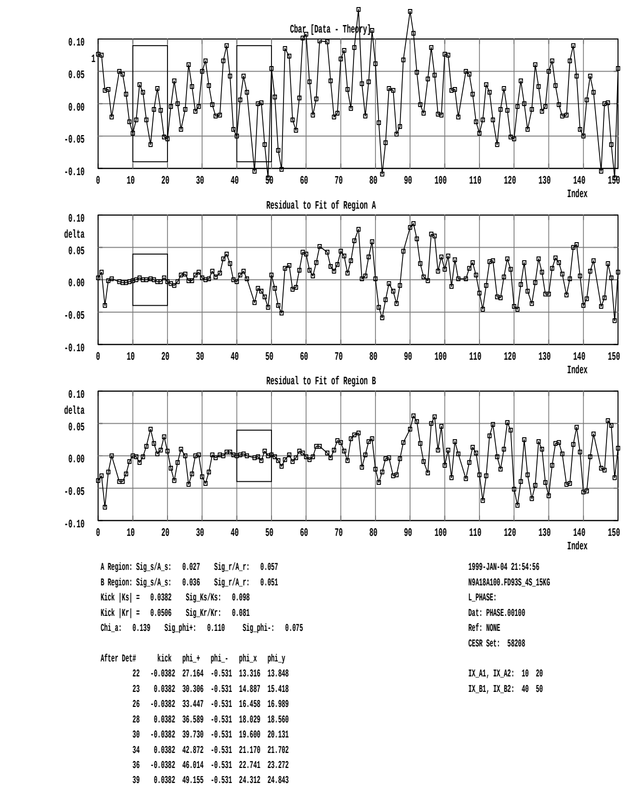

Figure 6: Wave analysis shows a coupling error located at detector 34. The top plot shows $\overline{C}_{12}$ data. The middle shows the residuals obtained by fitting a coupling wave to beam detectors 10 - 20, and the bottom plot shows residuals obtained by fitting to beam detectors 40-50.

coupling data obtained following correction with IR quad angle adjustments is shown in the figure. The $\overline{C}_{12}$ data is shown in the top plot, and the large coupling error is evident. (In these figures, the detectors "wrap-around" at detector 100, that is, trajectories are shown for one and a half times around the circumference.) A freely propagating coupling wave is fit to the data between detectors 10 and 20 and is propagated around the ring. The middle plot shows the residual of the fit wave and the data. As expected the residual is zero in the fit region. The residual deviates from zero at detector 34 in one direction and detector 3 in the other direction. The location at which the residual deviates is a candidate location for a coupling error. The bottom plot shows the residual obtained by fitting the data between detectors 40 and 50, also relevaling a large error at detector 34, as well as a smaller one near detector 60.

In fact, during this maintenance period, the strength of a horizontal corrector which is located adjacent to detector 34 was increased by adding a few turns of cable around the magnet iron. Guided by this wave analysis, we found that when the corrector was turned off, the coupling changed dramatically. In an attempt to correct this isolated coupling error, the IR quadrupole angles had been adjusted, and in fact themselves became coupling errors, as is evidenced by the wave launched in the interaction region (near detectors 0 and 100) in the middle and lower plots. When these turns were removed, this isolated coupling error was removed and the transverse coupling was successfully corrected.

This analysis of isolated field errors has been useful in many other circumstances as well. For example, it was discovered through such a wave analysis that the strength of one of the rare-earth interaction region quadrupole magnets was $\sim$0.3% different from that assumed in the design. In several other circumstances, bad-bit transitions in quadrupole controller cards led to large quadrupole errors which were easily identified.

4.3 Detector Solenoid Compensation Correction

The detector solenoid compensation is accomplished by rotating three symmetrically placed pairs of IR quadrupoles. Gross solenoid compensation errors are diagnosed and corrected by the transverse coupling measurement and analysis just described. While this method guarantees that the compensation is closed, since beam detectors are lacking at the interaction point it does not guarantee that the coupling parameters are zero at the IP. If the coupling is non-zero at the IP then the beams are differentially displaced due to the crossing angle trajectory which extends throughout the compensation region (and the entire machine due to the pretzel operation of CESR).

There is no straightforward way to adjust the vertical overlap of the beams at the interaction point. A pair of electrostatic vertical separators are located symmetrically about the crossing point diametrically opposite the IP and are powered to provide beam separation at this parasitic crossing point. Adjustment of the phase advance between these separators produces a differential displacement of the two beams at the IP. This phase advance is routinely adjusted to maximize the luminosity. We found that this adjustment resulted typically in a vertical orbit ripple of $\pm$0.3 mm (1-2 σ_v at the IP). This observation suggested that the phase advance adjustment was necessary to compensate for another effect which was separating the beams at the IP.

We found no dependence of the vertical orbit on the crossing angle amplitude, as is expected if the solenoid compensation is closed. These facts suggested that the solenoid compensation was in fact closed, but that the coupling at the IP was non-zero. That is, they suggest that a closed coupling bump had been "built-in" by the routine coupling correction. To investigate this possibility, two closed coupling bumps were constructed by rotating the 3 pairs of IR quadrupoles; rotation of all quadrupoles on one side of the IP in the direction opposite that of those on the other side results in a coupling bump which displaces the beams vertically at the IP (with a crossing angle trajectory); rotation of the quadrupoles on both sides in the same direction yields a vertical angle at the IP. It was found that as the vertical displacement coupling bump was introduced, the phase advance adjustment between vertical separators could be eased, and the ripple reduced. In this way the solenoid compensation was corrected, yielding a reduced vertical orbit ripple. The original motivation for this correction scheme was ultimately to improve the luminosity performance of CESR by achieving the correct solenoid compensation. Indeed, the beam-beam tuneshift parameter improved following this correction [2]

5 ACKNOWLEDGMENTS

The authors would like to thank Sergey Belomestnykh, David Rubin and Stuart Peck for results presented in this paper.

6 REFERENCES

[1] S. Henderson, in these proceedings.

[2] D. Rice, in these proceedings.

[3] S. Henderson, Proc. Part. Acc. Conf. 1997, Vancouver BC, p.2029

[4] D. Sagan, D. Rubin and S. Greenwald, Proc. 1996 Euro. Part. Acc. Conf. p. 1335.

[5] P. Bagley and D. Rubin, Proc. Part. Acc. Conf. 1989, p. 874.

[6] D. Sagan and D. Rubin, Phys. Rev. Spec. Topics, Acc. and Beams, 2, 074001 (1999).

[7] D. Sagan, Cornell CBN 97-01.