

Instrumentation and Diagnostics for PEP-II*

Alan S. Fisher

Stanford Linear Accelerator Center, Stanford University
M.S. 17, P.O. Box 4349, Stanford, California 94309

Abstract

PEP-II is a 2.2-km-circumference collider with a 2.1-A, 3.1-GeV positron ring (the Low-Energy Ring) 1 m above a 1-A, 9-GeV electron ring (the High-Energy Ring); both rings are designed to allow an upgrade to 3 A. Since June 1997, we have had three runs totaling 14 weeks to commission the full HER, reaching a current of 0.75 A. Positrons were transported through the first 90 m of the LER in January 1998, with full-ring tests planned for the summer. This workshop provides a timely opportunity to review the design of the beam diagnostics and their performance, with an emphasis on what works, what doesn't, and what we're doing to improve it. This paper discusses: the synchrotron-light monitor, including both transverse imaging onto a CCD camera and longitudinal measurements with a streak camera; beam-position monitors, with processors capable of 1024-turn records, FFTs, and phase-advance measurements; tune measurements with a spectrum analyzer, including software for peak tracking; measurements of both the total ring current and the charge in each bucket, for real-time control of the fill; and beam-loss monitors using small Cherenkov detectors for measuring losses from both stored and injected beam.

Invited talk presented at 8TH Beam Instrumentation Workshop (BIW 98), Stanford, California, May 4-7, 1998.

* Work supported by Department of Energy contract DE AC03 76SF00515 for SLAC and DE-AC03-76SF00098 for LBNL.

Instrumentation and Diagnostics for PEP-II*

Alan S. Fisher

*Stanford Linear Accelerator Center, Stanford University
M.S. 17, P.O. Box 4349, Stanford, California 94309*

Abstract. PEP-II is a 2.2-km-circumference collider with a 2.1-A, 3.1-GeV positron ring (the Low-Energy Ring) 1 m above a 1-A, 9-GeV electron ring (the High-Energy Ring); both rings are designed to allow an upgrade to 3 A. Since June 1997, we have had three runs totaling 14 weeks to commission the full HER, reaching a current of 0.75 A. Positrons were transported through the first 90 m of the LER in January 1998, with full-ring tests planned for the summer. This workshop provides a timely opportunity to review the design of the beam diagnostics and their performance, with an emphasis on what works, what doesn't, and what we're doing to improve it. This paper discusses: the synchrotron-light monitor, including both transverse imaging onto a CCD camera and longitudinal measurements with a streak camera; beam-position monitors, with processors capable of 1024-turn records, FFTs, and phase-advance measurements; tune measurements with a spectrum analyzer, including software for peak tracking; measurements of both the total ring current and the charge in each bucket, for real-time control of the fill; and beam-loss monitors using small Cherenkov detectors for measuring losses from both stored and injected beam.

INTRODUCTION

The PEP-II *B* Factory (1) is a 2.2-km-circumference, two-ring, e^+e^- collider under construction at the Stanford Linear Accelerator Center (SLAC) in the tunnel of the original PEP single-ring collider. The project is a collaboration with the Lawrence Berkeley and Lawrence Livermore National Laboratories (LBNL and LLNL). Its goal is the study of CP violation by tracking the decay of $B\bar{B}$ meson pairs produced with nonzero momentum in the lab frame, the design involves two rings at different energies; both rings require large currents for high luminosity. The 2.1-A, 3.1-GeV positron ring (the low-energy ring, or LER) runs 1 m above the 9-GeV, 1-A electron ring (the high-energy ring, or HER). At one interaction point (IP), the LER comes down to the height of the HER, and the two beams collide with zero crossing angle in the BaBar detector. Table 1 lists several of the parameters for PEP-II operation.

Because the HER reuses the PEP-I magnets (although with a new, low-impedance, vacuum chamber), it began commissioning first, in May 1997, and has accumulated 14 weeks of full-ring operation through the end of the January 1998 run. At that point, the

* Supported by the U.S. Department of Energy under contracts DE-AC03-76SF00515 for SLAC and DE-AC03-76SF00098 for LBNL.

TABLE 1. PEP-II Parameters.

Parameter	HER	LER	Unit
Circumference	2199.318		m
Revolution frequency	136.312		kHz
Revolution time	7.336		μ s
RF frequency	476		MHz
Harmonic number	3492		
Number of full buckets	1658		
Bunch separation	4.20		ns
Luminosity	3×10^{33}		$\text{cm}^{-2} \cdot \text{s}^{-1}$
Center-of-mass energy	10.58		GeV
Current	0.99 (3 max)	2.16 (3 max)	A
Energy	9.01 (12 max at 1 A)	3.10 (3.5 max)	GeV
RF voltage	14.0	3.4	MV
Synchrotron tune	0.0449	0.0334	
Betatron tunes (x,y)	24.617, 23.635	38.570, 36.642	
Emittances (x,y)	49.18, 1.48	65.58, 1.97	nm-rad
Bend radius in arc dipoles	165	13.75	m
Bend radius in SLM dipole	165	43.45	m
Critical energy in arc dipoles	9.80 (23.23 max)	4.81 (6.92 max)	keV
Critical energy in SLM dipole	9.80 (23.23 max)	1.52 (2.19 max)	keV

maximum current reached 750 mA in 1222 bunches. Both horizontal and vertical feedback were running; longitudinal feedback had been commissioned but was not running since power amplifiers were out for repair. The full RF voltage (15 MV) was available, with low-level feedback to stabilize the output. This paper reviews the design of the beam diagnostics and discusses their performance during commissioning.

LER commissioning also began that month with the injection and transport of beam through the first 90 m of the ring. The first run of the complete LER is planned for July 1998, followed by high-current commissioning and colliding-beam studies in the fall. BaBar will be installed in early 1999.

SYNCHROTRON-LIGHT MONITOR

Synchrotron radiation (SR) in the visible and near ultraviolet (600–200 nm) has been used to measure HER beam profiles in all three dimensions. The measurements are made in the middle of Region 7, a high-dispersion point in the middle of a HER arc. A second synchrotron-light monitor (SLM) for the HER was planned near the start of the arc, where the dispersion is low, but has not been built due to budget limitations. An SLM is also planned for the LER in Region 2. We first present the HER system, and then discuss the modifications needed for the LER.

HER Synchrotron-Light Monitor

The high current in each ring leads to a high SR power on the first mirror. The ring's design does not permit the solutions common in synchrotron light sources. Space in the narrow tunnel does not allow backing the mirror away to reduce the heat load. Access to SR is restricted, with few ports, all presenting little impedance to the beam. To reduce the power along the SR stripe, the beam is incident on the mirror at 4° to grazing, giving

a maximum power (for nominal energy and 3 A) of 200 W/cm in Arc 7 of the HER (and 19 W/cm for the LER in IR-2). We detail the arrangement for the HER.

HER arcs are almost entirely filled by the 5.4-m dipoles, with a quadrupole, corrector and sextupole taking up much of the rest of the 7.6-m half cell (see Fig. 1). The intense SR fan strikes the water-cooled outer wall of the chamber. The first mirror (Fig. 2), mounted in the vacuum chamber on the arc's outer wall, reflects the light horizontally across the chamber to the downstream inner corner. The mirror is slightly rotated to recess the upstream edge behind the opening in the chamber wall, so that it doesn't receive power at normal incidence. The downstream edge sticks slightly into the chamber, shading the leading edge of the chamber as it resumes downstream of the mirror.

At 200 W/cm, the mirror cannot be cooled sufficiently to obtain adequate flatness for good imaging. Instead, note that the high-power SR fan at the critical energy is 15 times narrower than the visible fan we image. When the electrons travel on axis, a 4-mm-high slot along the mid-plane of the 7-cm-long mirror passes the x-ray fan, while visible light reflects from the surfaces above and below. Because of grazing incidence, the x rays never reach the bottom of the slot, which tapers from 0 to 5 mm in depth, but travel past the mirror to dump their heat into a thermally separate absorber (Fig. 2). The residual




FIGURE 1. HER and LER beamlines in the middle of Arc 7 in (a) an elevation view and (b) a plan view, showing path of the HER synchrotron light going down to the enclosure on the optical table below the HER dipole. Part of the light continues from the table to the penetration leading up to the streak-camera lab.




FIGURE 2. The slotted first mirror (M1) and the x-ray absorber, both mounted in the wall of the HER chamber.

heat load of 1 W/cm², due largely to scattered SR, causes a temperature variation across the surface of less than 1°C.

However, the electron beam will not always be correctly positioned. Then we do not demand that the mirror be suitable for imaging, but only that it not exceed its yield strength. We can then steer the electrons back to their proper orbit and wait for the mirror to cool. Both the mirror and the dump are made of Glidcop[®] (copper strengthened with a dispersion of fine aluminum-oxide particles) with water-cooling channels, following techniques (2) developed for the Advanced Light Source. An ANSYS thermal analysis of a beam hitting the mirror 2.5 mm above the top of the slot shows that the temperature for a 3-A HER beam rises from 35°C to 160°C, and the stress rises to 90% of yield (3).

Two 45° mirrors, with a fused-silica window in between, transport the light from M1 to imaging optics in a nitrogen-filled

enclosure on an optical table located below the HER dipole (Fig. 1(a)). This location was chosen to get good resolution from a short, stable optical path. The dipole itself provides radiation shielding. Both mirrors are motorized for remote adjustment, to correct for changes in the beam's orbit.

Fig. 1(b) shows the imaging scheme, designed to compensate for the effect of the slot. In geometric optics, a slot or aperture placed in the plane of a lens (like a camera iris) serves only to restrict uniformly the amount of light reaching the image plane without otherwise affecting the image. Here, the first focusing mirror F1 images the slot onto the second focusing mirror F2, which then images the beam onto a CCD camera. Another camera images M1, so that we can center the SR on the slot. The third focusing element, a motorized lens F3, adjusts the focus for different beam orbits.

We also considered two effects of diffraction. Table 2 shows the loss in resolution due to the small vertical dimension of the beam. To reduce this effect, the light at all

TABLE 2. Resolution of the SLM for measurements at 300 nm. Here, the diffraction spot size, $0.26(\rho\lambda^2)^{1/3}$, uses a larger experimental coefficient (from LEP) rather than the calculated value of 0.21. The image size is given by the quadrature addition of the source size and the diffraction size.

	HER Mid-Arc 7	LER Mid-Arc 7	LER IR-2
Radius of curvature in dipole [m]	165	13.75	43.45
Diffraction spot size σ_d [μm]	64	28	41
Electron/positron beam size σ_x [μm]	1000	622	2003
Electron/positron beam size σ_y [μm]	176	216	161
σ_y / σ_d	2.8	7.7	3.9
Image size σ_{image} [μm]	187	218	166
$\sigma_{\text{image}} / \sigma_y$	1.06	1.01	1.03



(a) Feedback running in all planes.



(b) Longitudinal feedback turned off.



(c) Horizontal feedback off.



(d) Vertical feedback off.

FIGURE 3. Transverse images of the HER beam.

three emission points is taken near horizontally defocusing quadrupoles, where the beams are large vertically. For a point source, diffraction from the slot causes some narrowing of the full width at half maximum (FWHM) of the image and creates tails, but the effect is small for our 4-mm slot. However, when this pattern is convolved with a narrow ($\sigma_y/\sigma_d = 2$) Gaussian electron beam, the tails disappear and the distribution broadens by 6% over the FWHM without a slot. The increase is less with PEP's broader beams.

This system, installed for our September run, soon produced beam images. For example, at high current and without feedback, the SLM displayed strong oscillations on the electron beam. Bunch-by-bunch transverse (4) and longitudinal (5) feedback decreased the motion and the spot size (Fig. 3). We use a video digitizer to determine the transverse beam size from the images. As the beam current varies, we set the electronic shutter of the CCD (Pulnix TM-7EX) to adjust for the light level; color filters are not usually inserted.

Problems also became apparent early on. First, we found that a manufacturing error led to poor alignment between M1 and the axis of the exit tube. Some corrective bending of the chamber was needed in the clean room before installation. Afterward, a beam bump using horizontal corrector magnets let us position the electron beam to center the light reflected from M1 as it enters the exit tube.

More seriously, on the TV monitor we found a second, somewhat distorted, image of the electrons above the main one (the one we show in Fig. 3); the images are separated by about the height of the TV screen. By scanning the electrons vertically, in order to move the region of illumination above or below the slot in M1, we could see that the two images came from the two halves of the mirror. The

shapes and separation of the two images did not vary with beam current, suggesting that the distortion was not thermal in origin. We also were unable to reach best focus within the range of the motorized lens F3, again suggesting a mirror distortion.

We next tried to determine the cause of the distortion. The M1 assembly, with its internal cooling channels, was first brazed at SLAC before being sent out for nickel plating and optical polishing at SESO in Marseilles, France. To keep the mirror balanced during polishing, SESO asked us not to attach the long stainless-steel tubes that bring the water through the vacuum housing to the back side of the mirror. Instead, we attached short stumps of tubing before polishing, and welded the tubes to the stumps after the mirror returned. Additional complexity resulted from PEP's policy of forbidding welds or brazes between water and vacuum: an "air guard" (a volume connected to outside air) must surround the water weld.

In November, after the run, we removed the mirror assembly for inspection, and took it to LLNL to make an interferogram for comparison with one done at SESO. Although M1 had been flat to $\lambda/30$ rms (using HeNe-laser red, 633 nm, and including both halves) after polishing, it was no longer so; the variations on one half alone were $\lambda/5$ rms. The worst deformation was on the front face opposite the insertion for a cooling tube on the back side. It appears that shrinkage while this weld was cooling deformed the mirror's front face. It also appeared that the mirror had folded slightly, so that the upper and lower halves were no longer on the same plane. When M1 is mounted inside the vacuum chamber (where we were unable to test it), this effect may be compounded by additional stress from pressing the mirror's support bracket against the mounting surface on the chamber, and by the weight and stiffness of the cooling tubes (although the tubes had been coiled to add flexibility).

After studying these problems, we remounted the mirror in December. A bellows just upstream of the source point was removed to place a survey telescope along the path of the synchrotron light, to re-survey and adjust the mirror's position. Then, two mirrors diverted the optical path into the aisle so that a reticle (focusing target) could be used as a source point. The CCD camera was repositioned to focus with F3 in mid-range.

The new camera position allowed us to focus the synchrotron light, during January, by scanning the focusing stage and finding a minimum spot size. We also bumped the orbit horizontally, to vary the distance along the tangent from the source point to M1. However, the minimum was at different settings for the vertical and horizontal planes, and for both, the measured beam size was larger than expected. The double image from the two mirror halves was still present.

For the long term, we have discussed two approaches. We might rebuild the mirror, with suitable modifications. The rear plate, where the cooling tubes attach, needs to be stiffer, and we must further relieve the mirror of the weight and stress of the long lines. The risk, of course, is that a new design carries the risk of new problems.

Alternatively, we are investigating the use of adaptive optics, a new technology in which a deformable ("rubber") mirror inverts the wavefront errors measured by a wavefront sensor. The method was originally developed for the military and later for astronomy. Japan's KEK *B* Factory has devised its own version and tested it at the Photon Factory, in order to simplify the primary mirror of its SLM (6). (In addition, a special weak dipole was added to provide a source point with lower SR power.)

Now these components are about to become more commercial and less astronomical in both size and price. One firm (7) has a DARPA (Defense Advanced Projects Research Agency) contract to commercialize a 1-cm-diameter deformable mirror with 325 actuators on a 0.5-mm grid, driven by a multichannel DAC and matched to a wavefront sensor. We are now discussing the possibility of SLAC participation in a "beta" test of this system for PEP's October 1998 run. The deformable mirror would be located at

M1's image plane on the optical table, with the sensor just downstream to form a closed feedback loop. A PC can calculate the correction and control the loop directly, or, for greater bandwidth, the loop can be run by a digital signal processor (DSP) on the wavefront-sensor board. Since the deformation we see is largely independent of time and temperature, a static correction may be sufficient, but subtler, time-dependent effects as the beam current changes may become apparent after the static error is corrected.

LER Synchrotron-Light Monitor

LER measurements were originally planned (8) for the middle of Arc 7, with the LER's primary mirror directly above that of the HER. This arrangement would allow us to combine the optics for both rings on one optical table in the tunnel, and to transport the light from both up to the same optics room and streak camera. However, the aluminum vacuum chambers in the LER arcs are quite different from the copper chambers of the HER. In LER arcs, the SR diverging from the beam downstream of each dipole enters an antechamber; 2/3 of the photons strike a water-cooled photon stop 6 m beyond the bend (see Fig. 1(a)), while the remainder, emitted in the downstream part of the bend, continue to the next photon stop. To extract the photons, we would need a vertical slot in a photon stop; in addition, we would have to modify the design for the primary mirror M1 and its mount.

To economize, we are instead now planning to install the LER's SLM in IR-2. Because this is the straight section of Region 2, with the IP at the center, the LER has dipoles that steer the positrons down to the level of the electrons, then horizontally across the IP, and finally back up. Downstream of the IP, one horizontal bend followed by a long drift has been selected for the SLM. The SR from the various bends makes the heat load in this straight larger than the others; consequently, the chambers use the same octagonal copper extrusions as the HER arcs. We already have a copy of the HER's M1, originally built for the second HER SLM; this mirror will fit without modification into a chamber which will be fitted with a mirror-mounting flange and a light-exit port. The disadvantage of choosing IR-2 is that the LER system cannot share the HER's optical table, controls, or streak-camera room, and there is a much longer distance to any possible site for the streak camera.

The dispersion at the SLM location in IR-2 is low, as it also is in the original Arc-7 location. No second SLM had been planned for the LER since there is no suitable high-dispersion bend.

Streak-Camera Measurements

In addition to these transverse-profile measurements, the January run included longitudinal studies using a streak camera. The HER light is split in front of the camera; one half is transported through an 11-m penetration to a ground-level optics room, as shown in Fig. 1(b). For single bunches, we measured bunch length versus current and RF voltage. Multibunch studies looked at longitudinal instabilities as we varied the current and fill pattern. These results are presented in a separate paper (9) at this Workshop.

BEAM-POSITION MONITORS

The beam-position monitor system (10) must provide a wide dynamic range, with the ability to measure both a multibunch beam in a full ring (up to 8×10^{10} e^\pm with 238-MHz spacing) and a single bunch of 5×10^8 during injection. To tune injection for a top-off fill, we plan to inject a single small bunch into the ion-clearing gap, and measure its position without interference from a 3-A stored beam. This requires a gated measurement and careful impedance matching to reduce round-trip reflections on the long cables to the processors. Table 3 compares the requirements to bench-test results.

Each BPM uses four 15-mm-diameter pickup buttons, matched to 50Ω and designed to contribute a low impedance to the ring. They are located at each quadrupole, with ≈ 300 sets per ring. To avoid synchrotron radiation, they are placed $\pm 45^\circ$ off the horizontal (or, in the octagonal HER arc chambers and the elliptical LER arc chambers, at points where 45° field lines terminate on the wall).

Except near the interaction and injection points, only one plane is measured— x only at QFs, y at QDs—in order to economize on cables and processors. In a filter-isolator box (FIB) next to the quad, the four button signals are filtered at $2f_{RF}$ (952 MHz) and combined in pairs (top and bottom, or left and right). Out-of-band power is matched into a load. For two-plane measurements, special FIBs filter but do not combine. An isolator on each FIB output limits a second pass of signal reflected from residual mismatch at the processor input.

The RInQ (ring I&Q) processor is a CAMAC module receiving one x and one y BPM signal from each ring. To limit cable length (for both cost and high-frequency loss), most RInQs are in the tunnel, in crates under the HER dipoles for shielding. Each channel has a 20-MHz bandpass filter, which sets the time resolution to 20 ns; a programmable input attenuator, to adjust for dynamic range; an in-phase and quadrature (I&Q) demodulator, to convert the signal to baseband cosine and sine components, while avoiding dependence on the phase of the RF reference; a track and hold with a gate opened once per turn and centered on any selected bucket; and a 14-bit digitizer. The wide bandwidth argued for direct conversion rather than using an intermediate frequency. A digital signal processor (a TI320C31 DSP) computes the position and records 1024 turns (either consecutive or every N^{th}) for display, averaging or a fast Fourier transform (FFT). Both the digital and RF circuits share a single printed-circuit board, made in two sections with a different dielectrics.

To economize on cables and processors, the HER and LER signals are multiplexed for demodulation and digitizing. To measure a small charge in one ring while the other is full, a SPST switch is followed in series by a high-isolation DPST switch; the paths are separated and shielded upstream of the second switch. In retrospect, this choice was less economical than it appeared: the isolation was difficult to achieve, while cellular tele-

phones have rapidly reduced the cost of components for this frequency range.

The RInQ includes two on-board, tunable, digitally synthesized frequency sources, clocked by the ring RF. One is used for the local oscillator (LO) for the I&Q; the other provides a calibration signal introduced into the signal path by a 10-dB

TABLE 3. BPM resolution for single-bunch fills, with and without averaging over 1024 turns. Measurements were made on the bench and so do not include the beat-frequency effect discussed in the text.

Averaging [Turns]	Charge [e^\pm]	Required [μm]	Measured [μm]
1	5×10^8	1000	<100
1	10^{10}	100	<20
1024	10^{10}	15	<1

directional coupler at the input. We calibrate the ADC pedestals, the gain ratios of the I and Q channels, their phase offset from 90°, and the top/bottom and left/right gain ratios. Since PEP must operate as a “factory,” these calibrations must be performed with beam stored in the rings, by measuring during the gap. The calibration source is tuned slightly off the LO, many measurements are made, and the points are fitted to a sine at the beat frequency, in order to measure the gains and pedestals of both I and Q. Both the LO and the calibration source are tuned off 952 MHz to further isolate the calibration from the stored beam. The multibunch fill has a widely spaced spectrum, allowing a narrow-band measurement at any 238-MHz harmonic (such as 952), but the single-bunch case has lines spaced at 136 kHz, requiring a separate broadband calibration, in which the source sweeps through the channel’s pass band.

Although this calibration process works well on the bench, many of the BPMs on the HER have been troubled by a false beam oscillation at the beat frequency between the LO and 952 MHz, or at twice this frequency. The oscillation’s amplitude is equivalent to beam motion of 60 to 300 μm peak to peak. We normally program the LO to operate just off 952 MHz, typically from 21 Hz to 1 kHz off, because the beat frequency is a signature of a false oscillation, while a DC error would be indistinguishable from real beam position, and the error can be made to disappear by averaging the position over complete periods. The problem is related to the calibration; for example, as the LO phase slowly beats, the signal moves from cosine to sine, and an error in the I-to-Q gain ratio would show up at the second harmonic of the beat. DSP software changes have improved the situation, and recently there have been indications that the problem is, at least in part, related to the new power supplies for the PEP crates, which are different from those on the test bench.

During the January run, we commissioned a new BPM application to measure the phase advances and beta functions at BPMs around the ring. The beam is driven at the tune frequency in one plane, say x , using the tune system. All BPMs record the same 1024 turns. The DSP for the k^{th} BPM then fits the responses x_{kn} for turn n to

$$x_k = A_{x,k} \cos(v_x \omega_{\text{rev}} t + \phi_{x,k}) , \quad (1)$$

giving the amplitudes $A_{x,k}$ and phase advances $\phi_{x,k}$ around the ring. Network traffic is kept low, since only one phase and amplitude per BPM are reported back to the Control System. Although the beta functions $\beta_{x,k}$ can be found from $\beta_{x,k} = A_{x,k}^2 / \epsilon_x$, a better way (11) that is independent of BPM calibrations uses

$$\phi_{x,k} - \phi_{x,k-1} = \int_{s_{k-1}}^{s_k} \frac{ds}{\beta_x(s)} , \quad (2)$$

along with the known optics and measured phase advances from BPM k to BPMs $k-1$ and $k+1$. This method has allowed us to tune out beta beats with small changes in the IP quadrupoles.

TUNE MONITOR

The HER tune monitor, shown in Fig. 4, takes signals from dedicated BPM-type pickup buttons, passes them through a downconverter and into a digital spectrum analyzer. The LER monitor will be a duplicate. (In addition, each ring has a second dedicated set of buttons reserved for special measurements, such as high-frequency spectra to examine bunch dynamics.)

The button signals are combined with 180° hybrids to form a sum signal and horizontal and vertical difference signals; an RF switch selects two of these three. A computer-controlled step attenuator, followed by an amplifier, allows for a wide dynamic range, from a single bunch of 5×10^8 electrons to 1658 bunches of 8×10^{10} , 4.2 ns apart (a full ring except for a 5% ion-clearing gap).

Up to this point, the components are all broadband, to keep the pulses narrow. A fast GaAs switch (actually a pair of switches in series, to reduce leakage from other bunches) can then gate the signal from one or more bunches, or pass the signal from the entire ring. Several effects, such as the fast ion instability (12), may cause the tune to vary along the bunch train following the ion-clearing gap, and the gate lets us measure these effects. Another important case will occur during collisions: the first and last bunches in the train will not experience one parasitic crossing (near-collision with a bunch from the other beam) 0.63 m (half the bunch spacing) away from the IP, because the other beam has a gap rather than a bunch as they approach or leave the IP. Compared to the other bunches, these will feel a different beam-beam kick and so will have different tunes. Another application of the gate is to measure the tune of a specific bunch while feedback is turned off for that bunch alone.

The two-channel spectrum analyzer (Hewlett Packard 89410A) uses digital signal processors (DSPs) and an FFT to compute high-resolution spectra from 0 to 10 MHz. To bring the signals from the pickups into range, the front end includes mixers at $2f_{RF}$.

The analyzer includes a tracking generator to excite the beam with a swept sine or broadband noise when needed. No separate excitation structures are needed. Instead, for transverse excitation, this signal is summed with the input to the power amplifiers for the stripline dampers of the transverse feedback system (13). For longitudinal excitation, the signal will be added to a signal sent by the longitudinal feedback system (14) to modulate the ring RF for control of low-frequency modes; this “sub-woofer link” was partly commissioned during the January run. The drive can be switched on or off separately for each plane, and also passes through another 2-ns gate, which can chop the

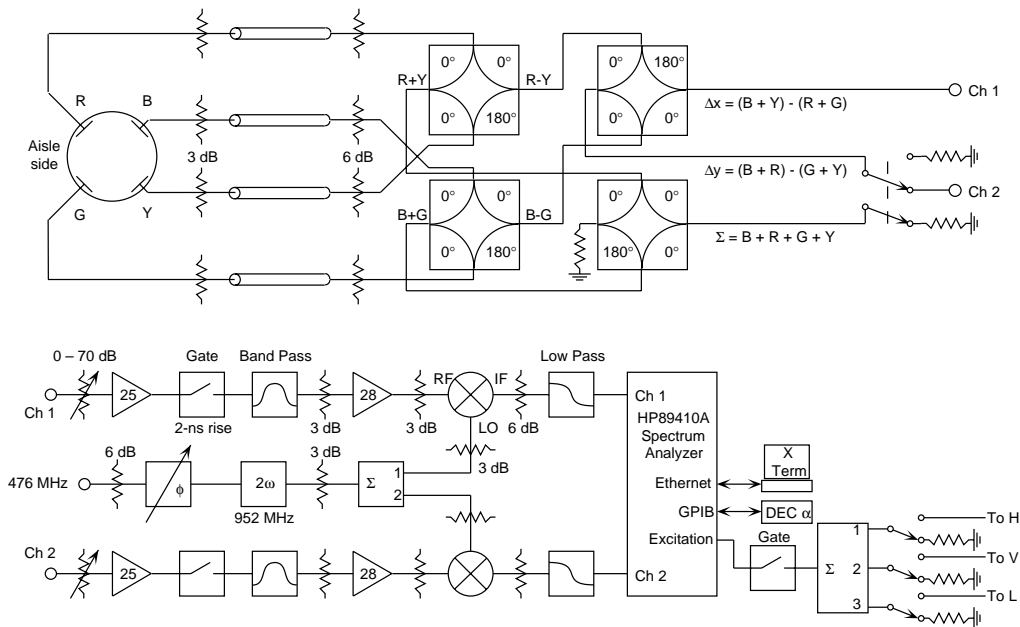


FIGURE 4. The PEP-II tune monitor (one per ring).

excitation in order to drive all or part of the bunch train. To study multibunch instabilities, for example, we can drive one bunch while measuring the response of the following bunches. When the LER is complete, we can excite bunches in one ring and measure the corresponding bunches in the other.

In October, PEP multibunch spectra were filled with peaks as the current was raised. Much of this was due to noise imposed on the beam from the RF system, driving large synchrotron and synchro-beta oscillations that were too big to be controlled by longitudinal feedback. The spectra cleaned up markedly in January after commissioning various feedback loops to control the outputs of the RF stations.

The spectrum analyzer interfaces through both GPIB—for data transfer and remote commands—and ethernet—for FTP data transfer and for an X-window display of an image of the front panel, showing the two traces and allowing control with a mouse. Its internal processor runs both built-in functions like peak finding and user programs in Instrument Basic; both can be executed on command from PEP Control-System software, with results returned to the Control System. During January, we commissioned routines that record the tunes periodically for a history buffer, and that follow tune peaks during scans of chromaticity or xy coupling, to automatically make a correlation plot. To avoid reporting the wrong peak in a complex spectrum, the algorithm tracks the evolution of the two peaks (one x and one y) initially identified by the user. Future routines may control the beam-excitation signal to measure the peak with the minimum drive.

In July 1998, when the LER is commissioned, we plan to have our first run with both electrons and positrons. The tune monitor will be used to help bring the two beams into collision. We will excite one beam transversely with a sine wave, and look with the tune monitor for the coupling of this motion to the other beam as we scan the relative positions of the two beams at the IP. A correlation plot showing the amplitude of the response against position will find the best alignment for collisions. For good sensitivity, we will set the spectrum analyzer to a narrow resolution bandwidth (1 Hz). To get a larger amplitude for the driving or responding beam, we can choose to use the tune frequency of one ring, but we must avoid confusion with self-excited motion.

CURRENT MONITORS

A commercial DC current transformer (DCCT) (15) measures the total current in the HER with a 5- μ A resolution over a 1-s integration time and a full-scale value of 5 A. (For comparison, a 1-A current with a 3-hour lifetime drops by 93 μ A/s, and injecting 5×10^8 e^\pm adds 11 μ A.) A second unit is now being installed in the LER. Our DCCT housing places it outside the vacuum envelope, provides an electrical gap directing DC wall currents around the transformer core, and capacitively bypasses the gap for higher frequencies to present a low impedance to the beam. An integrating voltmeter (Keithley 2002) reads the DCCT output, and will use an input scanner to alternate between HER and LER.

The nominal bunch pattern has 1658 buckets, each separated by 4.2 ns (two RF periods) and filled to as much as 8×10^{10} e^\pm for a 3-A beam. To balance the beam-beam kicks, the variation in charge per bunch within each ring must be held to $\pm 2\%$. Such tight control requires a second current diagnostic, the bunch-current monitor (BCM) (16), to measure the charge in each of the 3492 RF buckets. A third system, the bunch-injection controller (BIC), then plans the filling sequence, while a fourth, the master pattern generator (MPG) implements it.

With a relative accuracy of 0.5%, the BCM in each ring updates measurements at 60 Hz. This rate compares to 40 Hz for interleaved injection into both rings (and 60 Hz for e^+ only, or 120 Hz for e^- only). In each ring we sum and filter the signals from a set of four BPM-type buttons, using a microstrip combiner with a 2-period comb filter at $3f_{RF}$ (1428 MHz). The filter is designed to avoid crosstalk from adjacent bunches. A mixer at $3f_{RF}$ brings this signal down to baseband, giving a DC to 1-GHz “video” output. The operating frequency is a compromise: a lower frequency would not allow the high video bandwidth and low adjacent-bunch crosstalk; at a higher frequency, the mixer output would be more sensitive to synchrotron oscillations, and some of the button signals would come from propagating modes in the beampipe.

The video goes to a VXI crate, where an 8-bit track-and-hold ADC clocked at f_{RF} digitizes every bucket at the same phase. This data stream is divided among 12 Xilinx field-programmable gate arrays. Even after this “decimation,” the rate remains too high for processing. Each Xilinx downsamples by processing only one bucket out of 8 in each turn, so that it takes 8 turns to sample the entire ring. The data for each bucket is summed over 256 measurements in each 60-Hz interval, to improve the resolution and to average over many synchrotron oscillations, then is written into a table in a reflected (dual-port) memory. In addition, lifetime measurements of individual bunches require an accuracy of 0.05% in 1 s, to allow the detection of a lossy bunch (≤ 10 minute lifetime) at a rate useful for operator adjustments; consequently, the VXI processor maintains a second table with sums over 1-s intervals.

The BIC, in an adjacent VME crate, reads the memories of both rings, and also their DCCT voltages and exact measurement times using the Keithley’s GPIB interface. It normalizes the individual bunch currents to the DCCT and calculates lifetimes. Over an EPICS interface to the control system, it displays the DC currents, lifetimes, and bunch charges, and receives the user’s desired fill pattern. User settings for the BCM go in the opposite direction, from EPICS to the BIC and through the reflected memory to the BCM. Based on the charge measurements, the BIC lists the injection sequence for the MPG, which controls the injector linac to fill the appropriate buckets in the rings with bunches of selected sizes. New charge is usually injected in a pattern of nine zones around the ring, so that the bunch can damp before more charge arrives in a zone. The complete fill-control system was commissioned during the January run.

BEAM-LOSS MONITORS

A network of 100 beam-loss monitors (BLMs) detects losses at selected points (collimators, septa, and selected quadrupoles) around the rings. The system covers a wide dynamic range, from high losses to well stored beams, provides reasonable localization, and allows the measurement of injection loss. The output is used for machine tuning, for loss histories, and for the rapid detection of high losses requiring a beam abort.

We have chosen a Cherenkov detector, using a small (16 mm diameter) photomultiplier with 2-ns-wide pulses (comparable to the bucket spacing). The Cherenkov radiator is an 8-mm-diameter, 10-mm-long, fused-silica cylinder placed against the fused-silica PMT window, with optical grease on the interface. The opposite end and the cylindrical surface are aluminized for internal reflection. The assembly is enclosed in 1 cm of lead to avoid synchrotron-radiation background, but remains small enough to be moved for commissioning and troubleshooting. Using the ring magnets as shielding, the BLMs can be placed for preferential sensitivity to HER or LER, or they

can be exposed to losses from both rings. As the LER is installed, some BLMs are being moved from their initial placement on the HER.

BLM processors (BLMPs), ten-channel CAMAC modules distributed around the rings in the BPM crates, process each BLM signal through two input circuits that together provide a wide dynamic range. To measure low loss rates, the PMT pulses pass through a discriminator and are counted over 1-s or 8-ms intervals. A different procedure is needed for high losses, to avoid high count rates and pulse pile-up, and for injection losses, since a counter would record only a single count even if the injected bunch hits the wall by a BLM. For these situations, the PMT signal is integrated by a 10- μ s (about one ring turn) RC filter, and a peak detector then saves the maximum for 8 ms. A multiplexer scans the channels and digitizes these lossiest-turn readings, which are available to the Control System on request. The peak detectors are then reset.

If the integrated signal exceeds a programmable threshold, it is possible to abort one or both rings (determined by two programmable abort-enable bits) by firing a kicker on each ring to send the beam into a dump. The BLM processor then records the triggering channel and, through a daisy chain linking all the processors, causes all BLMs to freeze their most recent readings. Several other faults, such as a loss of RF or the closure of a valve, can also fire this abort system.

Another daisy-chain signal provides a 100- μ s gate around injection time. During this interval, the BLM network is inhibited from aborting the stored beam, since faulty injection is a more likely source of a large loss (and stored-beam losses will persist after the gate). To measure injection loss, the multiplexer timing is restarted to digitize the outputs of the peak detectors within 1 ms after this inhibit interval. This scheme measures the loss not on the first turn, but on the worst turn, since the injected bunch may not scrape until a later turn, depending on its betatron phase at the obstacle. The Control System then acquires the readings within 3 ms of injection.

ACKNOWLEDGMENTS

I would like to thank the many people from all three laboratories who have worked with me in developing PEP's diagnostics, as well as the many others involved in the building, installing, and commissioning the machine. The list is too long to repeat here, but you know who you are. I appreciate your help and the collegial and cooperative spirit in which it was offered. Of course, we still have to finish the LER, collide the beams, install the detector,...

REFERENCES

1. *PEP-II: An Asymmetric B Factory*, Conceptual Design Report, LBL-PUB-5379, SLAC-418, CALT-68-1869, UCRL-ID-114055, UC-IIRPA-93-01, June 1993.
2. DiGennaro, R., and Swain, T., *Nucl. Instrum. Methods* **A291**, 313–318 (1990).
3. Daly, E.F., Fisher, A.S., Kurita, N.R., Langton, J. B., "Mechanical Design of the HER Synchrotron-Light-Monitor Primary Mirror for the PEP-II B Factory," *Proc. IEEE Particle Accelerator Conf.*, Vancouver, BC, May 1997 (IEEE Press, Piscataway, NJ), in press.
4. Barry, W., Byrd, J., Corlett, J., Fahmie, M., Johnson, J., Lambertson, G., Nyman, M., Fox, J., and Teytelman, D., "Design of the PEP-II Transverse Coupled-Bunch Feedback System," *Proc. IEEE Particle Accelerator Conf.*, Dallas, TX, May 1995 (IEEE Press, Piscataway, NJ, 1996).

5. Teytelman, D., Fox, J., Hindi, H., Limborg, C., Linscott, I., Prabhakar, S., Sebek, J., Young, A., Drago, A., Serio, M., Barry, W., and Stover, G., "Beam Diagnostics Based on Time-Domain Bunch-by-Bunch Data,"; Prabhakar, S., Teytelman, D., Fox, J., Young, A., Corredoura, P., and Tighe, R., "Commissioning Experience from HER PEP-II Longitudinal Feedback," in these Proceedings.
6. Mitsuhashi, T., Hiramatsu, S., Takeuchi, N., Itoh, M., and Yatagai, T., "A Design of Synchrotron Radiation Monitor for KEK B-Factory," *Proc. 11th Symp. Accelerator Technology and Science*, SPring-8, Ako, Hyogo, Japan, 21–23 Oct. 1997.
7. MEMS Optical, Inc., Huntsville, Alabama.
8. Fisher, A.S., Alzofon, D., Arnett, D., Bong, E., Daly, E., Gioumousis, A., Kulikov, A., Kurita, N., Langton, J., Reuter, E., Seeman, J.T., Wienands, H.U., Wright, D., Chin, M., Hinkson, J., Hunt, D., and Kennedy, K., "Diagnostics Development for the PEP-II B Factory," *Beam Instrumentation: Proceedings of the Seventh Workshop*, Argonne, IL, May 1996, AIP Conf. Proc. **390** (Amer. Inst. Phys., Woodbury, NY, 1997), pp. 248-256.
9. Fisher, A.S., Assmann, R.W., Lumpkin, A.H., Zotter, B., Byrd, J., and Hinkson, J., "Streak-Camera Measurements of the PEP-II High-Energy Ring," in these Proceedings.
10. Aiello, G.R., Johnson, R.G., Martin, D.J., Mills, M.R., Olsen, J.J., and Smith, S.R., "Beam Position Monitor System for PEP-II," *Beam Instrumentation: Proceedings of the Seventh Workshop*, Argonne, IL, May 1996, AIP Conf. Proc. **390** (Amer. Inst. Phys., Woodbury, NY, 1997), pp. 341-349. Smith, S.R., Aiello, G.R., Hendrickson, L.J., Johnson, R.G., Mills, M.R., and Olsen, J.J., "Beam Position Monitor System for PEP-II"; Johnson, R., Smith, S., Kurita, N., Kishiyama, K., and Hinkson, J., "Calibration of the Beam-Position-Monitor System for the SLAC PEP-II B Factory," *Proc. IEEE Particle Accelerator Conf.*, Vancouver, BC, May 1997 (IEEE Press, Piscataway, NJ), in press. Johnson, R., Smith, S., Aiello, G., "Performance of the Beam-Position Monitor System for the SLAC PEP-II B-Factory," in these Proceedings.
11. Castro, P., et al, "Betatron Function Measurement at LEP Using the BOM 1000-Turns Facility," *Proc. IEEE Particle Accelerator Conf.*, Washington, DC, May 1993 (IEEE Press, Piscataway, NJ, 1993), pp. 2103–2105.
12. Byrd, J., "An Initial Search for the Fast Ion Instability in the ALS," PEP-II AP-Note 95.49, 28 Aug. 1995.
13. Barry, W., Byrd, J., Corlett, J., Fahmie, M., Johnson, J., Lambertson, G., Nyman, M., Fox, J., and Teytelman, D., "Design of the PEP-II Transverse Coupled-Bunch Feedback System," in *Proc. IEEE Particle Accelerator Conf.*, Dallas, TX, May 1995 (IEEE, Piscataway, NJ).
14. Teytelman, D., Fox, J., Hindi, H., Limborg, C., Linscott, I., Prabhakar, S., Sebek, J., Young, A., Drago, A., Serio, M., Barry, W., and Stover, G., "Beam Diagnostics Based on Time-Domain Bunch-by-Bunch Data,"; Prabhakar, S., Teytelman, D., Fox, J., Young, A., Corredoura, P., and Tighe, R., "Commissioning Experience from HER PEP-II Longitudinal Feedback," in these Proceedings.
15. Parametric Current Transformer, Bergoz Precision Beam Instrumentation, Crozet, France.
16. Chin, M.J., Hinkson, J.A., "PEP-II Bunch-by-Bunch Current Monitor," *Proc. IEEE Particle Accelerator Conf.*, Vancouver, BC, May 1997 (IEEE Press, Piscataway, NJ), in press.