

Turn-by-Turn Imaging of the Transverse Beam Profile in PEP-II

Alan S. Fisher^{*}, Mark Petree^{*},
Richard Kraus[†], Yat-Shan Au[¶], and Benny Chan[¶]

^{*}*Stanford Linear Accelerator Center, 2575 Sand Hill Road, Menlo Park, CA 94025, USA*

[†]*Department of Physics, University of Nevada, Reno, NV 89557, USA*

[¶]*Physics Department, California Institute of Technology, Pasadena, CA 91125, USA*

Abstract. During injection or instability, the transverse profile of an individual bunch in a storage ring can change significantly in a few turns. However, most synchrotron-light imaging techniques are not designed for this time scale. We have developed a novel diagnostic that enhances the utility of a fast gated camera by adding, inexpensively, some features of a dual-axis streak camera, in order to watch the turn-by-turn evolution of the transverse profile, in both x and y . The beam's elliptical profile is reshaped using cylindrical lenses to form a tall and narrow ellipse—essentially the projection of the full ellipse onto one transverse axis. We do this projection twice, by splitting the beam into two paths at different heights, and rotating the ellipse by 90° on one path. A rapidly rotating mirror scans these vertical “pencils” of light horizontally across the photocathode of the camera, which is gated for 3 ns on every N^{th} ring turn. A single readout of the camera captures 100 images, looking like a stroboscopic photograph of a moving object. We have observed the capture of injected charge into a bunch and the rapid change of beam size at the onset of a fast instability.

Keywords: Imaging, Transverse, Turn-by-turn

PACS: 41.85.Ew, 41.85.Ar

INTRODUCTION

At times, the profile of an electron or positron bunch in a storage ring changes rapidly, on time scales as fast as one turn (microseconds). For example, an injected charge enters at an angle to the stored-beam orbit, oscillates about it, and damps into equilibrium over many milliseconds. Similarly, instabilities can cause rapid motion of the beam center or an increase in its size, on a millisecond time scale.

Several techniques provide partial information about such events. Beam-position monitors and similar pickups used for transverse feedback can measure the centroid for each turn, and, in the case of feedbacks, for each bunch. However, changes in beam size cannot be seen. Synchrotron-light or x-ray imaging shows both the centroid and the profile, but transverse measurements typically use video cameras, which average over all bunches and over a video frame interval (33 ms for the American NTSC standard or 40 ms for the European CCIR).

More elaborate cameras are sometimes used. Fast-gated cameras can pick out one bunch from a full ring, but capture new frames at only a few hertz. Streak cameras, normally used for longitudinal profiles with picosecond resolution, can also be set for

a slow sweep triggered with multiple gates. With the beam's x or y dimension aligned with the entrance slit, the camera can then measure changes in one transverse dimension over time. However, this puts a costly (approximately \$200,000) picosecond tool to work at microsecond or millisecond time scales.

Here we present a novel synchrotron-light diagnostic that enhances the utility of the gated camera by adding, inexpensively, some features of a dual-axis streak camera, in order to watch the turn-by-turn evolution of the transverse profile, simultaneously in both x and y . The beam's elliptical transverse profile is reshaped using cylindrical lenses to form a tall and narrow ellipse—effectively the projection of the full ellipse onto one transverse axis. We do this projection twice, by splitting the beam into two paths at different heights, and then rotating one ellipse by 90° before reshaping it. These projections are directed onto a rapidly rotating mirror placed close to the photocathode of the gated camera. The mirror scans these vertical “pencils” of light horizontally across the camera, which is gated for nanoseconds on every ring turn (or every N^{th} ring turn) to produce typically 100 images of each projection. These appear side by side, forming one stripe across the camera for each projection. All these images are captured on a single camera readout, resembling a stroboscopic photograph of a moving object that highlights changes in centroid position or beam size.

IMPLEMENTATION

Emission Points and Transport Lines

In PEP-II at the Stanford Linear Accelerator Center [1], synchrotron light monitors (SLMs) [2] collect visible light from electrons in an arc dipole of the high-energy ring (HER) and from positrons at a bend 30 m downstream of the interaction point (IP) in the low-energy ring (LER). For each SLM, a video camera in the tunnel records beam images. In addition, a transport line re-images part of this light onto an optical table (Fig. 1) in a hutch outside the shielding wall. Both transport lines include an out-of-plane periscope rotating the beam ellipse through 90° , so that the image has its major axis (the x -axis at the emission point) oriented vertically. This becomes the initial image to be manipulated by cylindrical lenses in the next section.

Cylindrical Lenses

Figure 1 shows the optical path on the LER table. (The HER has a similar arrangement, built but not yet commissioned, that differs only in the image location and in the magnification factors needed.) First, a polarizer retains the component of the incoming light that is horizontally polarized in the plane of the source dipole. To avoid chromatic blurring from the lenses in both the transport line and on the table, an image-quality color filter passes a band centered at 450 nm (blue) with a full width at half maximum of 80 nm.

A beamsplitter directs the light along two paths. The first, for measurement of the major axis of the incoming beam ellipse, retains its vertical orientation. After the beamsplitter, a periscope adjusts the height of the optical axis, without rotating the

image. Then two cylindrical lenses re-image the ellipse onto the photocathode of our gated camera, with different horizontal and vertical magnifications (Table 1).

Horizontal focusing needs a cylindrical lens with a short focal length placed close to the camera, in order to strongly demagnify the minor axis of the ellipse. To resolve 128 images in a 512-pixel-wide image, the image should have an rms width of about 1 pixel (24 μm). A nominal 100-mm focal length (before adjusting for wavelength) is as short as possible while allowing sufficient room for the rotating mirror that scans the thin vertical ellipse across the camera. The camera and the 100-mm lens are then placed far enough from the initial image to get the needed horizontal demagnification.

The major-axis measurements use the bottom half of the 12-mm by 12-mm photocathode, and so the vertical imaging should reduce the major axis to a typical rms height of 0.5 mm. Another standard catalog lens can give this vertical demagnification over the distance already chosen.

The second path measures the beam's minor axis. An out-of-plane periscope in the corner of the table simultaneously changes the optical-axis height, turns the path through 90°, and rotates the ellipse by 90°. The sizes on the camera should be the same as before: 0.5 mm rms vertically and 1 pixel rms horizontally. As the major and minor paths merge just before the rotating mirror, the minor path passes above the rim of the last mirror of the major path and images onto the top half of the camera.

By adding a second path, the initial image is relayed to the camera with four different magnifications that all must focus on the same plane. In addition, the minor path needs more difficult magnifications, since we must stretch the smaller axis of the elongated ellipse (Table 1) and shrink the larger one. The vertical magnification could not be achieved with a single catalog lens, but it could be done with a lens pair, which offers two more degrees of freedom (focal length and separation).

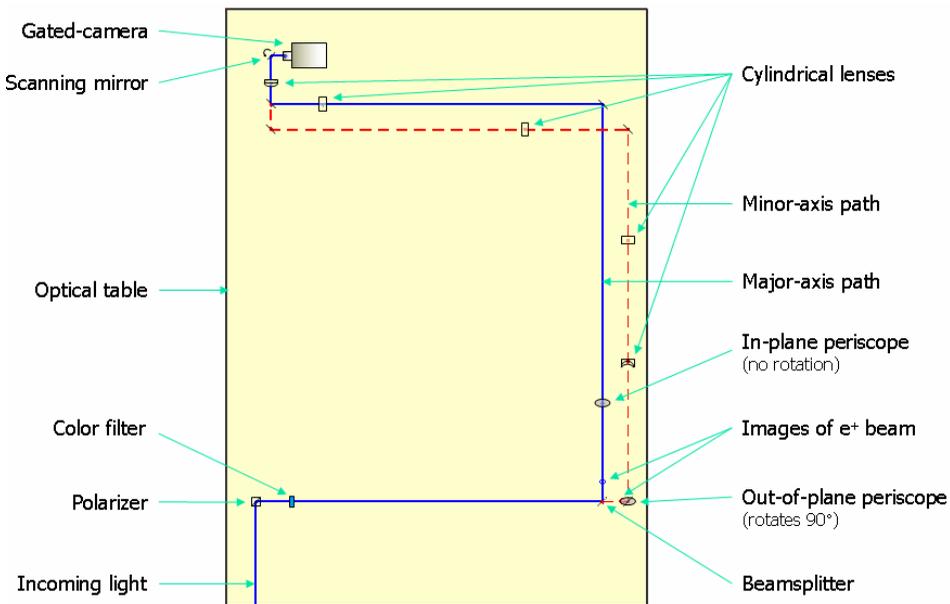


FIGURE 1. The layout of the optical table for the LER.

TABLE 1. Magnifications. A camera pixel corresponds to 24 μm on the 12-mm photocathode.

	Major-Axis Path		Minor-Axis Path	
	Horizontal	Vertical	Horizontal	Vertical
Image-to-camera distance (mm)	2013.1		2138.1	
Initial image size (mm rms, in collision)	0.38	3	3	0.38
Magnification	0.0545	0.167	0.0383	1
Image size on camera (mm rms)	0.021	0.5	0.115	0.38
Nominal focal lengths (mm)	100	250	-1000	1000, 700

Horizontally, although the minor path needs a stronger demagnification, the major path's 100-mm lens is too close to the camera to use different lenses for each. Also, note that the minor path is 125 mm longer than the major path. The vertical imaging can readily accommodate the extra path. Horizontally, we turn this distance into an opportunity for additional demagnification, by adding a cylindrical lens to relay the initial image to the plane 125 mm away, where the 100-mm lens forms an image of the photocathode. Because this extra distance is short, and we want to avoid the distortions of very small F numbers, we use a negative lens (Table 1), which places a virtual image on this conjugate plane and gains a longer distance from the initial image to the negative lens. This arrangement cannot reduce the major axis to the ideal width of 1 pixel, but is sufficient to resolve individual turns.

Rotating Mirror

Because this imaging concept was devised a month before a scheduled PEP shutdown in the summer of 2004, our initial experiments needed a mirror that was rapidly available "off the shelf." We found something suitable under the counter. A bar-code scanner used for supermarket checkout counters had the essential features—a tiny octagonal mirror with 3-mm-high by 13-mm-wide facets, a servo-controlled DC motor with speeds from 6 to 80 Hz, and a trigger output synchronized with the rotation—all packaged in a 50-mm \times 40-mm \times 20-mm box.

This mirror demonstrated turn-by-turn imaging, but with limitations. The 3-mm height did not allow simultaneous imaging of both x and y profiles. The motor was free running, not triggered by PEP timing. Instead, the camera was triggered by an output from the scanner, capturing only steady-state behavior.

We subsequently obtained a scanning galvanometer, a flat mirror mounted on the shaft of a DC motor (GSI Lumonics model VM500S, Fig. 2). These are common in laser scanners for light shows,

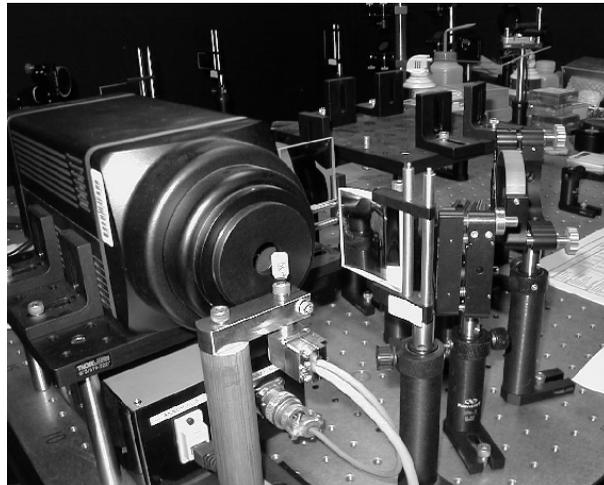


FIGURE 2. The galvanometer mirror, with its motor below, sits in front of the camera. At right, the final cylindrical lens.

displays, and printers. A voltage controls the rotation angle, and so a ramp from a waveform generator can be triggered to capture, for example, beam injection. For a low moment of inertia, the mirror must be small and thin. Our metal mirror is 14.2 mm by 8.7 mm, small enough to move with the necessary speed but taller than the 12-mm photocathode. The assembly—mirror, motor, and driver—costs \$1500.

There are both optical and mechanical constraints on the mirror location. It must be squeezed between the final, horizontally focusing, lens and the camera. The thin horizontal focus means a large convergence angle for rays heading to the cathode, requiring a wide mirror if far from the camera. If the mirror is near the camera, the necessary angular rotation speed increases, and the camera's own housing interferes. In fact, we had to replace this housing, a tube for mounting a camera lens, with a short protective sleeve. The eventual compromise put the mirror 35 mm from the cathode.

Triggers and Timing

Our PI-MAX camera from Princeton Instruments has an image intensifier that can be gated for intervals down to 2 ns, followed by a CCD with a 512 by 512 pixel array. A single picture may involve multiple gates with any separation. The software supplied with the camera then reads out the full CCD at a maximum rate of 2 Hz. The software can acquire single pictures or sequences of pictures with the option of scanning the gate width or delay. In focus mode, the camera takes pictures continually without saving them and with constant settings.

Consequently, a trigger at 2 Hz initiates the mirror rotation, followed by a comb of camera triggers spaced by one or more ring turns ($7.34 \mu\text{s}/\text{turn}$). These triggers must be carefully timed relative either to a given ring bucket or to injection from the linac. Since PEP bunches are normally separated by 4.2 ns (every 2nd bucket), the camera's gate width is set to 3 ns, which accommodates the gate's rise and fall time and the different lengths of the two optical paths.

Several trigger modes are provided by PEP programmable delay units (PPDUs). Triggers are linked to the fill pattern of the stored beam or to the arrival of an injected pulse. They can be single shots or a comb repeating at the ring-turn rate. Ring-turn combs are continuous, while injection combs continue for tens of milliseconds, depending on the injection rate. The timing relative to a bunch is arbitrary and must be measured.

The PPDUs drive two digital delay generators (Fig. 3). The first, DG1, runs at the slow 2-Hz rate, while DG2 provides a fast pulse train for the gates. Both units

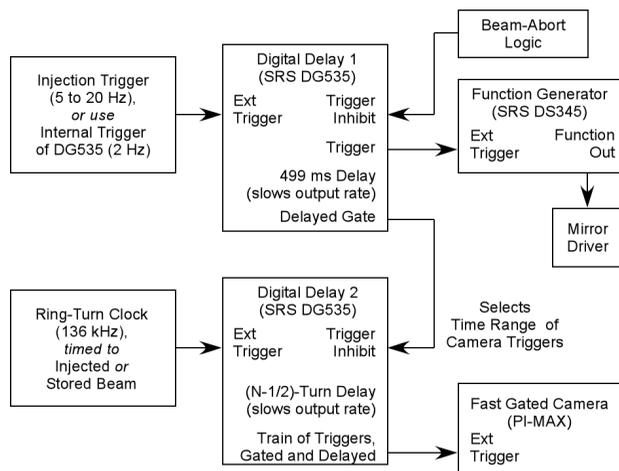


FIGURE 3. Timing for injected- or stored-beam measurements.

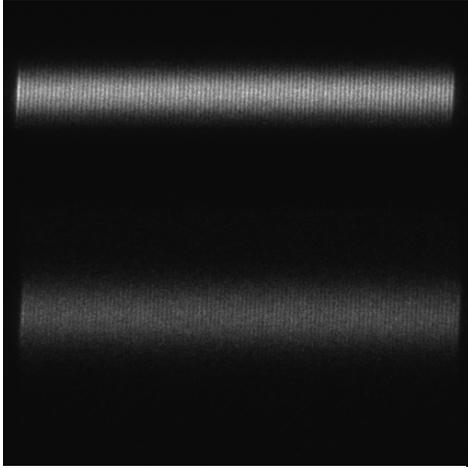


FIGURE 4. Stable 2.9-A positron beam. 100 images of one bunch from the train, taken on consecutive turns ($73 \mu\text{s}$ total). Dark bands between images may not show here due to the reduction of the photo.

repeatedly gates the camera. The rate can be reduced to every N^{th} ring turn by adding delay to one output. Depending on the measurement, the train can be locked either to the stored bunches or to injection. These time delays must be carefully measured to 1 ns. We first stored or injected a single bunch and scanned a wide gate to find the range that lit up the camera. Then this range was scanned with a narrower gate, and the process was repeated. The trigger inhibit is driven by the gate from DG1, to delay the first camera gate until the mirror is in the right position, and to limit the length of the train of gates.

IMAGING POSITRONS IN PEP'S LER

Stored Beam

Figure 4 shows a stable stored beam imaged on 100 consecutive turns. The images, all identical, form two stripes. The upper set shows the major axis of the beam stretched vertically, with the minor axis compressed horizontally. The lower trace shows the minor axis for the same bunch on the same turns. The lower set is a bit dimmer, both because the demagnification does not compress the light quite as tightly (Table 1), and because the beamsplitter in use does not split the light quite equally.

After a HER abort, the LER exhibited the transverse instability of Fig. 5, which displays one bunch 125 times on every 80^{th} turn, for a total span of 73 ms. Centroid motion is clearly visible in both planes. The frequency of the instability may be aliased by the sample rate.

include a trigger-inhibit input, which halts outputs when held at ground.

To measure stored beam, DG1 is triggered internally at a fixed 2 Hz. For injection, a pulse from the PPDU gives a single trigger for each injection. To reduce the trigger rate to the camera's 2 Hz, one of DG1's outputs is delayed. (All outputs must finish before a new sequence begins.) DG1 triggers the mirror's waveform generator, and produces a gate of adjustable width, used to define the length of the pulse train from DG2. The trigger inhibit for DG1 is connected to logic from the beam-abort system. When imaging an instability leading to an abort, the camera runs steadily in focus mode, and then halts at the abort.

The PPDU's pulse train triggers DG2 at a ring-turn rate. Its output then

Beam Aborts

In December 2005, PEP experienced numerous fast beam aborts. Characteristically, the beam-position monitors, which record the final 2800 turns in a circular buffer, showed a sudden small burst of transverse motion that damped in 1 ms, followed 3 ms later by an abort triggered by a rapid loss of LER current. The initial motion was well below threshold for an abort, and had damped before the beam loss began.

To better understand these events, we set the camera to capture every 80th turn for 125 images, spanning a 73-ms interval. The camera ran at its top speed, 2 Hz, and in focus mode so that pictures are taken continually but not saved. A logic level from the abort system inhibited triggers on an abort, to hold the final picture. The time scale was carefully chosen: with 500 ms between consecutive pictures, we had a 1 in 7

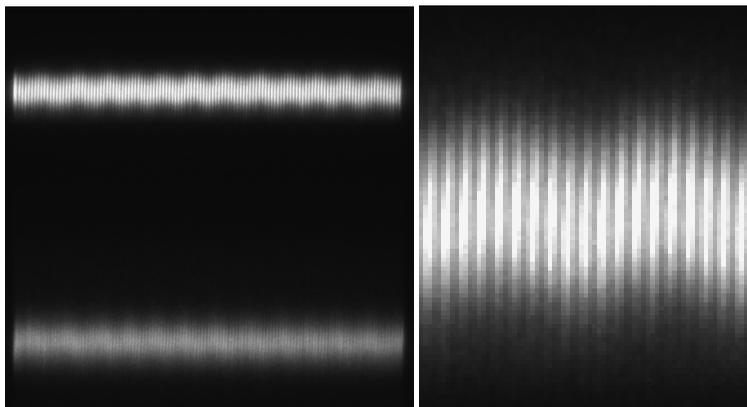


FIGURE 5. Transversely unstable beam, showing periodic centroid motion. (a) 125 images of one bunch, taken on every 80th turn (73 ms total). (b) The middle region of the upper trace of (a), magnified.

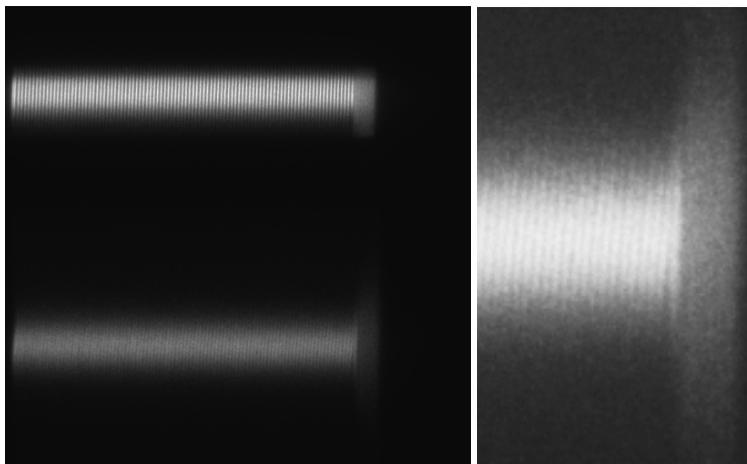


FIGURE 6. Sudden beam-blow-up in the LER, followed by an abort. (a) 125 images of one bunch, taken on every 80th turn (73 ms total). (b) A magnified view of the lower trace of (a), starting just before the instability. Centroid motion is visible in one image before the blow-up begins.

chance of finding the abort in the picture. A longer span would include few images of the abort itself, while a shorter time would have a lower probability of capture. With 125 images, the dark spaces between the images are just discernable, although they may not reproduce well here.

Figure 6(a) shows the beam blowing up in both transverse directions over the 4 ms of the abort, after which the photo is dark. Figure 6(b) expands the final milliseconds of the

lower (minor-axis) trace. The instability begins with a small transverse deflection, corresponding to the time when the BPMs record transverse motion. As the beam blows up, the dark spaces between the images fill in completely, since the narrow dimension of each image also widens. (During the blow-up shown, the major axis appears clipped at the top; this was found later to be due to a partly closed iris that was passing the normal beam size but not the severely increased size during the instability.) Ultimately, the cause of the instability was traced to an improperly seated RF seal at a bellows near one RF cavity. Part of the seal protruded into the beampipe, exciting high-order modes.

A deficiency of this technique for observing beam aborts is the 1 in 7 chance of catching the event, since we are using the camera as a continuous circular buffer despite its long readout time. An alternative is possible, although we have not tested it since it sacrifices single-bucket and turn resolution. An ordinary CCD video camera makes a far better (and much cheaper) circular buffer, since its chip continues to integrate while the image is scanned from storage in adjacent dark pixels. Depending on the arrangement of this buffer, all illuminated pixels integrate light continuously, except for the brief vertical-blanking interval. A non-interlaced CCD with interline transfer could capture aborts with high probability. This method would lose bunch and turn information, merging all the bunch images into a continuous envelope. The extra light would compensate for not having an image intensifier.

Injection

The original motivation for the rotating-mirror technique was the diagnosis of backgrounds in PEP's *BABAR* detector from faulty injection. Also, we have speculated that the bunch-by-bunch transverse feedback, which responds to the centroid of each bunch, may be kicking the stored beam too much as it responds to the oscillations of the injected pulse damping into the same bucket. The first step is to capture an injection event.

For the stored-beam measurements, all bunches in the ring exhibited similar behavior, and so timing the triggers was not difficult. Getting an image of injection demands careful measurements of the delay from the injection trigger in order to find the one bunch receiving the charge.

The difficulty is compounded by the tiny amount of charge injected, less than 10% of a full PEP-II bucket. Because PEP injection is vertical, the injected charge oscillates along the minor axis of the stored bunch. To direct all the light along the minor-axis path, the beamsplitter was removed. Fig. 7(a) shows the minor-axis view of the first injection into an empty bucket. The camera captured 80 consecutive turns (587 μ s), but the injection takes place on the 13th turn, leaving the first 12 dark. Afterward, the bunch oscillates widely about the stored-beam orbit. In Fig. 7(b), after 12 turns with a quiet stored beam, a second injection arrives and starts its oscillation.

Figure 8(a) shows a longer time interval, again with 80 images but now triggered on every 10th turn. Some damping of the oscillation is evident, although a full damping time is an order of magnitude longer. Figure 8(b) shows the second injection on the same time scale.

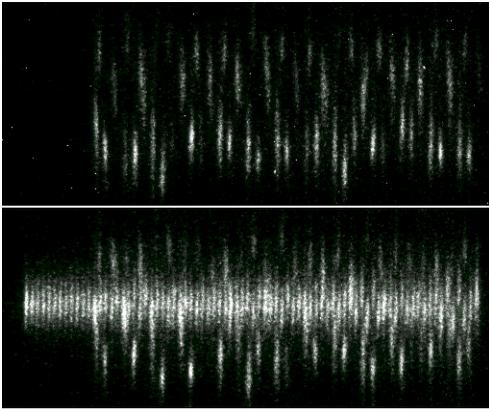


FIGURE 7. Minor-axis image of injected charge arriving on the 13th of 80 consecutive turns. In (a) the bucket is initially empty, while in (b) the second injection merges with the previously stored bunch.

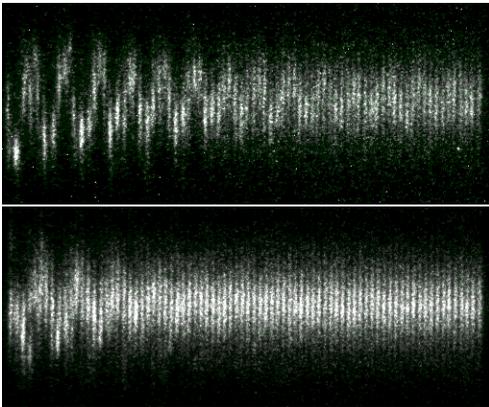


FIGURE 8. Similar to Fig. 7, but imaging every 10th turn. Again, in (a) the bucket is initially empty, while (b) shows the second injection.

CONCLUSION

The new technique presented here obtains multiple images of the transverse projections of an electron or positron bunch in a storage ring in order to capture the beam's evolution over consecutive turns, or for time scales up to 100 ms. A combination of cylindrical optics and a fast rotating mirror imaged the beam onto a fast gated and intensified camera, providing observations of beam instabilities, beam-loss events, and injection.

ACKNOWLEDGMENTS

Students participated in this project with support from summer programs of the U.S. Department of Energy and SLAC (Summer Undergraduate Laboratory Internships, for R.K.), and of Caltech (Summer Undergraduate Research Fellowships,

for Y.-S.A. and B.C.). Prof. Frank Porter of Caltech provided support for the initial experiments. We appreciate the support of PEP's accelerator-physics and operations groups.

REFERENCES

1. U.S. Dept. of Energy contract DE-AC03-76SF00515.
2. Fisher, A.S., "Instrumentation and Diagnostics for PEP-II," *Proc. Beam Instrumentation Workshop*, Stanford, CA, 4-7 May 1998, edited by R.O. Hettel et al., AIP Conference Proceedings 451, Woodbury, NY: American Institute of Physics, 1998, pp. 95-109.