

Lessons learned from positron-electron project low level rf and longitudinal feedback

J. Fox, T. Mastorides, C. Rivetta, and D. Van Winkle

Stanford Linear Accelerator Center, Stanford, California 94309, USA

D. Teytelman

Dimtel, Inc., San Jose, California, USA

(Received 14 November 2008; revised manuscript received 16 December 2009; published 19 May 2010; corrected 30 June 2010)

The Positron-Electron Project II (PEP-II) B Factory collider ended the final phase of operation at nearly twice the design current and 4X the design luminosity. In the ultimate operation state, eight 1.2 MW radio-frequency (rf) klystrons and 12 accelerating cavities were added beyond the original implementation, and the two storage rings were operating with longitudinal instability growth rates roughly 5X in excess of the original design estimates. From initial commissioning there has been continual adaptation of the low level rf (LLRF) control strategies, configuration tools, and some new hardware in response to unanticipated technical challenges. This paper offers a perspective on the original LLRF and longitudinal instability control design, and highlights via two examples the system evolution from the original design estimates through to the final machine with 1.2×10^{34} luminosity. The impact of unanticipated signals in the coupled-bunch longitudinal feedback and the significance of nonlinear processing elements in the LLRF systems are presented. We present valuable “lessons learned” which are of interest to designers of next generation feedback and impedance controlled LLRF systems.

DOI: [10.1103/PhysRevSTAB.13.052802](https://doi.org/10.1103/PhysRevSTAB.13.052802)

PACS numbers: 29.27.Bd, 29.20.db

I. INTRODUCTION

The 1991/1993 PEP-II Conceptual Design Reports planned an electron-positron collider with luminosity of 3×10^{33} via a 2.14 A positron low energy ring (LER) and 0.75/1.5 A electron high energy ring (HER). The collider was initially commissioned in 1998 and ran for 10 years. In the final weeks of operation (March/April 2008), the storage rings reached currents of 3.2 A (LER) and 2.1 A (HER) with luminosity of 1.2×10^{34} .

As illustrated in Table I, the machine was continuously developed and achieved operating currents and luminosity well above the original design goals. The combination of additional rf cavities (additional impedances) and increased currents meant that the higher order mode (HOM) driven growth rates for longitudinal coupled-bunch instabilities in the HER were roughly 3X the original design estimates, and in the LER 5.6X the original estimates. Successfully controlling these increased instability growth rates were a significant challenge for the coupled-bunch instability control systems. Similarly, the low level rf (LLRF) systems incorporated direct and comb-loop feedback techniques to minimize the fundamental mode impedance presented to the beam which drove low-mode coupled-bunch instabilities [1,2]. Operating the machines at multiples of the design currents and with extra rf cavities greatly increased the required performance of these control loops.

The PEP-II storage rings routinely ran with low-mode (cavity fundamental-driven) instabilities with growth rates of 1.2 ms^{-1} HER and 3.0 ms^{-1} LER, which correspond to

e-folding intervals of 7 and 10 synchrotron cycles. These instability growth rates are in conjunction with the LLRF impedance control loops operating at their limits. These cavity fundamental-driven instabilities had to be controlled via a dedicated low-mode “woofer” control loop. The rapid instability growth rates were unanticipated by the system designers, as the design estimates and initial simulations had these cavity fundamental-driven instabilities fully damped by the LLRF control loops.

The successful operation of the machine, with the original LLRF and longitudinal feedback designs ultimately operating well above the design specifications, presents interesting examples of the methods used to estimate the performance required in the systems, as well as the design techniques and implementation approaches used. We also can learn from the manner in which the systems were adapted and reconfigured over the course and evolution of the PEP-II lifetime. This sort of historical review highlights some important experience gained by the PEP-II team.

The breadth of the PEP-II experience cannot be usefully discussed in a single paper. Instead, we choose one example each from the broadband feedback and from the LLRF system to expand for discussion. These examples illustrate how the original design estimation missed some very significant details, and how in the course of PEP-II operation unexpected difficulties led to significant insights and new approaches which allowed higher machine performance. We highlight some very significant experience of importance to future machines via these examples and a compilation of publication references.

TABLE I. Summary of rf configurations used in PEP-II from initial commissioning through the end of operations in April 2008. The design luminosity and current goals were achieved in run 1. At the cessation of operations the machine was operating with currents at 1.8X design LER, 2X design HER, and with 4X design luminosity. During this interval the HER reconfigured the rf stations such that a 4-cavity station was reconfigured as a 2-cavity station, and subsequent HER station additions were 2-cavity stations. Over this course of operations the rf configurations were in continuous development following varying installed complements of rf stations, rf cavities, gap voltages, machine synchrotron tunes, etc., which required unique LLRF configurations.

Year (run)	LER rf stations	LER rf cavities	HER rf stations	HER rf cavities	HER I (A)	LER I (A)	Luminosity
1998	2	4	4 (+ 1 parked)	16 (+ 4 parked)	0.6 A	1.0 A	1.2×10^{33}
Run 1	2	4	5	20	0.9 A	1.5 A	3.0×10^{33}
Run 2	3	6	5	20	1.0 A	1.7 A	4.4×10^{33}
Run 3	3	6	6	22	1.1 A	1.9 A	6.3×10^{33}
Run 4	3	6	8	26	1.5 A	2.5 A	9.0×10^{33}
Run 5a	4	8	9	26	1.7 A	3.0 A	1.0×10^{34}
Run 5b	4	8	9	26	1.9 A	2.9 A	1.2×10^{34}
Run 6	4	8	11	28	1.9 A	3.0 A	1.2×10^{34}
Run 7	4	8	11	28	2.1 A	3.2 A	1.2×10^{34}

II. BROADBAND COUPLED-BUNCH FEEDBACK—UNDERSTANDING THE IMPACT OF NOISE IN THE PROCESSING CHANNEL, AND DYNAMIC RANGE ISSUES

Control of HOM-driven coupled-bunch instabilities was a central challenge of the PEP-II design. It was recognized during the early design estimation that there could be unstable coupled-bunch motion in both transverse and longitudinal planes. The design approach relied on damping the cavity HOM impedances via external waveguides and loads, so that the remaining HOM impedances would be greatly reduced (though also spread across a much

greater frequency span resulting from the low external Q of the cavity/damping loads). To keep single bunch stability parameters satisfied, the design proposed filling almost all the possible rf buckets; as a result the coupled-bunch instability systems would need to push the bandwidth of the processing channels and control (kicker) elements to 119 MHz.

The rf system development progressed in parallel with the initial design and development of the broadband feedback systems. Data from rf cavity tests was available to estimate the residual impedances driving instabilities that would have to be controlled by the coupled-bunch feedback systems [3]. This impedance data could be scaled, and folded in frequency to compute the estimated growth rates for the HOM modes expected in PEP-II [4]. As presented in Fig. 1, the damped cavity HOMs were anticipated to excite many bands of coupled-bunch modes. These bands, which spanned dozens of revolution harmonics, led the designers to concentrate on all-mode or bunch-by-bunch control techniques. This choice of an all-mode system meant that, even if the installed rf cavities had slightly different HOM frequencies than predicted from the low-power model, the resulting unstable modes would still be controlled.

A. Longitudinal feedback estimation and design

To combat these longitudinal instabilities, the broadband feedback system design developed a reprogrammable digital processing architecture to control coupled-bunch (HOM-driven) instabilities in both light sources and factory colliders [5]. The block diagram of the longitudinal feedback system (LFB) is depicted in Fig. 2. The system designers used several types of modeling [6,7] and ma-




FIG. 1. (Color) Growth rate estimates from impedance measurements for 3 A LER. Impedance data from [3], growth rate estimates from [4].

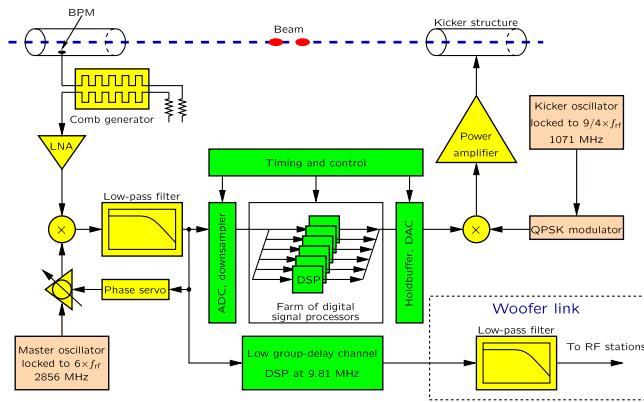


FIG. 2. (Color) Block diagram of the longitudinal broadband feedback system developed for use at PEP-II/DAFNE/ALS. The system uses up to 500 MHz input and output sampling rates for a programmable DSP based array processor. In the PEP-II implementation the beam signals are detected at 3 GHz, while the output signal is translated to a band centered at 1071 MHz.

chine measurements [8] to develop estimates of the required noise floors, gains, and output powers required for the various installations. The key formalism used to design the control filters was a finite impulse response (FIR) bandpass filter (typically with 4–16 coefficients), specified to implement a net overall 90° phase shift at the synchrotron frequency with zero DC gain. The expense of the broadband 1–2 GHz kicker power amplifiers led to every effort to minimize the installed output power. The gain of the digital signal processing (DSP) filter in conjunction with the rf function gains provides the necessary total loop gain with some extra margin [9].

During the earliest phases of system estimation, the designers were concerned that the system noise floor [rf phase detector noise + analog/digital (A/D) converter noise] had to not saturate the kicker amplifiers at the operating gain. Care was taken to develop low phase noise oscillators and receivers so that the controlled damped beam would damp down to the noise floor of the processing channel (roughly 2% of the cm bunch length). This noise floor and damping behavior was validated in the PEP-II system simulations, and in the lab testing of the various components. The necessary operating gains were estimated, and there was still some extra gain margin in the system.

The control filter gain at the synchrotron frequency (shown in Fig. 3) is directly related to the damping rate achieved by the feedback channel. The installed output amplifier power, in conjunction with the number of output kickers and their transfer impedance, determines the maximum kicker voltage which can be developed. The installed power effectively sets the maximum allowed beam disturbance amplitude from which unsaturated (linear) actuator response is possible. The initial simulation efforts studied the behavior of the system in both saturated and linear

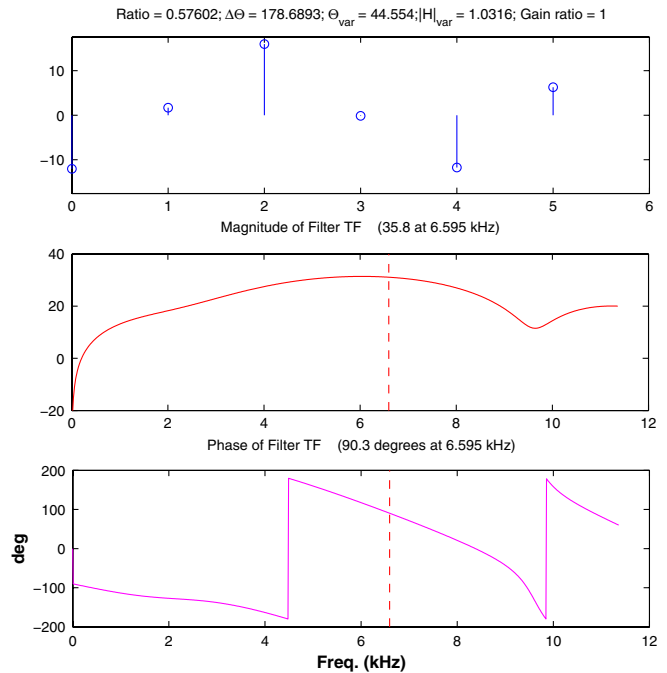


FIG. 3. (Color) Example DSP control filter implemented for each bunch in the PEP-II HER. This is a 6-tap FIR filter, optimized to have 35 dB gain at the 6.5 kHz synchrotron frequency with overall net 90° phase shift including system delays. The gain at 720 Hz is roughly 28 dB below the system operating gain and rapidly falls to zero at DC.

(unsaturated) states to better understand the limit of stability of the controlled system [6].

The Advanced Light Source (ALS) was the first installation where this programmed DSP feedback system was commissioned [10]. Many important system design choices were validated in operation at the ALS, and the operation there provided important experience in measuring system performance, instability growth rates, developing techniques to time the system kickers and fast channels to 20 picosecond accuracy. Based on this operating experience at the ALS the initial PEP-II commissioning was expected to be uneventful.

B. PEP-II broadband feedback commissioning and operational performance

Figure 4 presents measurements of the LER HOM growth rates measured in PEP-II (complete results for both HER and LER are presented in [11]). The most unstable band of modes seen in the machine is the largest impedance seen in the cavity test data (aliased into the baseband between 107 and 110 MHz of Fig. 1), and the expected growth rates are in very good agreement. The damping rates achieved by the DSP system are also shown in Fig. 5. Because of the increased impedances (the extra rf cavities), at the ultimate currents the DSP filter systems were running with system gains roughly 3X to 6X higher




FIG. 4. (Color) LER HOM-driven coupled-bunch modal growth rates for modes 790–810. Data from various runs is scaled by the number of installed rf cavities for consistency (from [11]).




FIG. 5. (Color) LER HOM-driven coupled-bunch modal damping rates (growth plus feedback damping) (from [11]).

than originally estimated for the design 0.9 A (HER) and 1.5 A (LER) currents [11]. This increased gain meant that the noise floor in the processing channels was also amplified 3X to 6X over the original design estimates, and the allowable range for driven motion of the beam is also reduced by this same factor of 3 to 6.

C. Unexpected longitudinal feedback system implications from the LLRF system

The initial PEP-II commissioning measurements were very surprising—while the system performed well damping coupled-bunch instabilities, and the achieved damping rates were adequate to control the beam, the amount of “noise” in the processing channel was much greater than anticipated. Studies showed this was due to signals in the LLRF system, at a broad band of frequencies from well below up through the synchrotron frequency, which drove the beam longitudinally through the rf cavities.

The rf high voltage (HV) power supplies in PEP-II were three phase systems with silicon controlled rectifier (SCR)

commutation to regulate the rf output voltage as demanded by the rf power regulating loops. These power supplies had limited output filtering due to concerns with energy stored in the output filter capacitors and possible damage to the klystron in an arc [12]. As a result, the klystron rf output had significant (2%) ripple components at 720 Hz (the supply fundamental component) as well as multiple harmonics and subharmonics. Because of the limited stored energy in the high voltage power supplies (HVPS) filters, line noise, regulator transients, and single or few AC cycle transients had significant sensitivity in the cavity accelerating voltages and the regulation of the accelerating rf power.

The ALS commissioning experience had not suggested these issues. In retrospect, the ALS rf and LLRF systems operated at such a relatively low power (single 330 kW klystron) compared to the 1.2 MW of each PEP-II klystron, that the ALS klystron HV power supply could be much better filtered and quieter than the energy storage-limited PEP-II rf power supplies. The PEP-II implementation also had multiple rf stations, and each could add noise to the accelerating voltage seen by the beam.

Figure 6 presents the receiver power spectrum of a controlled damped 1700 mA HER beam. The spectrum of this beam phase error or baseband front-end signal shows all sorts of features including the synchrotron resonance, rf system klystron HV power supplies, noise within the LLRF system processing channels, and phase reference distribution systems. Part of the perturbing signals and

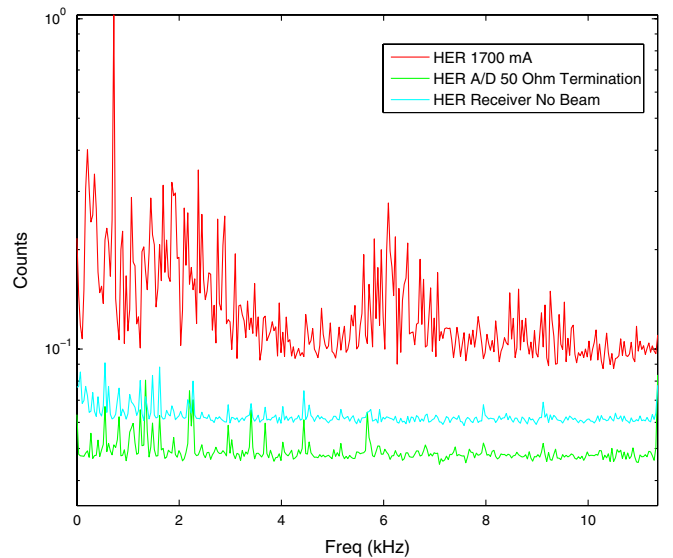


FIG. 6. (Color) Power spectrum of detected closed-loop HER beam motion (rms A/D Counts). A/D quantizing noise is 0.4 counts rms and the combined noise with the phase detector receiver (no beam) is 0.6 counts rms. It is important to realize that the quantization noise from the 8 bit A/D is negligible compared to the signals on the beam, and a 6 bit A/D system would suffice with the same performance.

noise in the phase error signal are impressed on the beam. While these signals are filtered through the bandpass DSP control filters, (reducing the power away from ω_s in the kicker system), the overall impact of the low-frequency signals from the rf system was problematic. In the same plot, the quantizing noise of the analog-to-digital converter (ADC) and the receiver noise spectrum when no beam is present are depicted.

Additional insight of the signals involved in the longitudinal feedback channel can be found by looking in the time domain at the front-end and back-end signals. Figure 7 (upper) shows the HER baseband front-end signal (ADC, from Fig. 2) for a single turn while the machine is operating stably at 1800 mA. This figure shows the gap transient, of roughly 4 degrees at the rf frequency, which fills roughly 1/3 of the dynamic range of the ADC. Each bunch rides on a unique synchronous phase. The bunch longitudinal coordinates are processed by the DSP filters, which are bandpass functions which remove the DC synchronous phase position from each bunch sequence, and provide gain around the synchrotron frequency. Figure 7 (lower) shows the output of the DSP processing [digital-to-analog converter (DAC), from Fig. 2] for this same turn. It is important to observe that the structure of the gap transient is removed, and about 40% of dynamic range of the output channel (+127/−128 DAC counts) is used from the noise and residual motion of the beam at the synchrotron frequency.

The systems operate in this equilibrium without difficulty, running the output power stages with signal components from the noise and driven motion perturbations of the beam. The true HOM instability signal is damped to the noise floor of the ADC as seen in Fig. 6, and the majority of

the power in the processing channel is from driven motion and broadband noise. The dynamic range around the operating point, with 35 dB gain of the processing filter, is such that a few A/D counts of synchrotron frequency motion through the processing channel fully saturate the output stage (this corresponds to a few tenths of an rf degree of phase motion). The high gain is necessary to have adequate damping, as the HER HOM growth rates are roughly 3X the original design estimates from the additional cavities and higher currents.

At the end of PEP-II operations, at 2100 mA in the HER, the operation of the HER broadband longitudinal feedback system began to reach an effective gain limit due to saturation effects in the power stages from 720 Hz and impulsive noise on the beam from the rf HV power supplies and other rf system disturbances. The system exhibited short transient excitations at the synchrotron frequency from transient effects in one or more of the rf stations, which drove barycentric (mode zero) longitudinal motion for short bursts.

D. Understanding the mechanism of transient beam loss from runaway HOM excited motion

Figures 8–16 show an interesting fault file in the HER at 1800 mA in which an impulsive low-frequency transient saturates the feedback with mode zero signals, leading to loss of HOM control and eventual loss of the beam. This sort of beam loss was very hard to diagnose as the measured growth and damping rates always showed excellent margins in operation at the same current, and the behavior was not a steady-state situation. The ring would operate for extended periods (weeks) without any anomalous loss of




FIG. 7. (Color) HER front-end and back-end signals of the longitudinal feedback system for a single turn while the HER system is operating with nominal beam parameters at 1800 mA. The upper plot shows the phase error signal for all the bunches. The lower plot depicts the base band signal driving all the individual bunches at the same turn.




FIG. 8. (Color) Time-domain fault file from the HER showing the data at the output of the DSP filters (the output signals from the DSP baseband processing with dynamic range +127/−128 DAC counts) The transient content is significant enough to pass through the control filter and saturate the power stage near 1000 turns in the data set. The 5000 turns of the recording is 36 ms long and is from an 1800 mA HER fill.




FIG. 9. (Color) Time-domain fault file from the HER showing the data at the input of the DSP filters. The individual bunch coordinates are shown during the same interval as Fig. 8. At roughly 1000 turns in the data set impulsive rf power supply noise (or phase reference noise), with amplitudes of a few A/D counts, becomes problematic as a barycentric motion (mode zero) gets excited at an amplitude which saturates the broadband feedback channel.

HOM control, and then suddenly exhibit spurious beam loss due to HOM motion.

To understand the sequence of events which lead to beam loss in this transient, it helps to start with Fig. 8, which is a time-domain fault file from the HER showing the data at the output of the DSP baseband processing (Hold buffer/DAC, in Fig. 2, before heterodyning up to




FIG. 10. (Color) Front-end and back-end signal of the HER longitudinal feedback system for turn 500. The upper plot shows the phase error signal for all the bunches. The lower plot depicts the base band signal driving all the individual bunches at the same turn. Figures 11–14 show the same signal for turns 502 to 512 corresponding to half a period of the approximately 6 kHz synchrotron oscillation. At this moment the kick signal is almost completely saturated at the full-scale value of +127.




FIG. 11. (Color) Front-end and back-end signal of the longitudinal feedback system for turn 502 of the time sequence shown in Figs. 8 and 9. At this moment the kick signal is mostly positive, and hitting the full-scale saturation value of +127.

the output 1071 MHz quad phase shift keying carrier). This transient shows impulsive low-frequency noise in the rf systems appearing at the output of the DSP control filter. The impulsive excitations are driving the output stage to the full-scale dynamic range of +127/−128 DAC counts (a significantly greater level than seen in Fig. 7). The transient content at the synchrotron frequency in each bunch is significant enough to pass through the control filter and saturate the power stage beginning near turn 400 in the data set. Figure 9 shows the data for 4800 turns in the HER at the input of the DSP filters (low-pass filter/ADC, Fig. 2) for the same transient before the beam abort. The




FIG. 12. (Color) Front-end and back-end signal of the longitudinal feedback system for turn 506 of the time sequence shown in Figs. 8 and 9. At this moment the kick signal around the turn is hitting both positive and negative saturation, though the average value is close to zero and centered.




FIG. 13. (Color) Front-end and back-end signal of the longitudinal feedback system for turn 510 of the time sequence shown in Figs. 8 and 9. At this interval the kick signal is largely saturated at the negative full-scale value of -128 .

individual bunch coordinates are shown during the same interval as Fig. 8. The beginning of the data (turns <300) shows the gap transient with all motion damped to the noise floor (similar to the behavior depicted in Fig. 7).

We can look in greater detail in Figs. 10–14 at the front-end and back-end signals at turns 500, 502, ..., 512, which present sampling in time of roughly one-half of the synchrotron period or $80 \mu\text{sec}$. We see that the front-end signals shift up and down as a common synchronous phase motion over this interval, and the output DSP correction signals saturate from this mode zero (barycentric) motion of the beam.

The exact source of this large disturbance is not known from this data sequence, but it must originate in an rf




FIG. 14. (Color) Front-end and back-end signal of the longitudinal feedback system for turn 512 of the time sequence shown in Figs. 8 and 9. At this moment, roughly $1/2$ a synchrotron cycle later than Fig. 10, the kick signal is almost completely saturated at the negative full-scale value of -128 .

station, perhaps in a HVPS which has some jitter or regulation dropout for a line cycle due to some transient on the AC mains or an internal mechanism. The modulation of the cavity rf voltage is significant enough that the beam follows the vector shift of accelerating voltage, and the transient has significant excitation of the synchrotron frequency. Excited by perturbation and noise from the rf stations, this damped motion of beam mode 0 temporarily saturates the broadband longitudinal beam control channel. This temporary saturation of the HOM control channel reduces the overall gain in the feedback system for all the modes controlled by the system. In particular, the dominant unstable beam modes (around mode 800) start to grow at a growth rate defined by the reduced gain of the saturated longitudinal feedback system. Once some HOM modes grow in amplitude so that the control filter is saturated, there is no option of recapturing the motion and damping the HOM-driven instability even as the transient saturation forced by mode 0 ceases. As seen Figs. 8 and 9, a small group of bunches around bunch 450 become unstable. Figure 15 shows the front-end and back-end signals for turn 1000 where it is possible to verify that bunches around 450 became unstable saturating the bunch-by-bunch controller for those particular bunches. Other transients analyzed have shown similar patterns where small numbers of bunches become unstable along the filled turn, though in patterns where other locations around the turn grow unstably.

It is important to notice that the original transient disturbance of the beam is mode 0 induced by rf station perturbations. Because of the limited dynamic range of the broadband longitudinal feedback (limited due to satu-

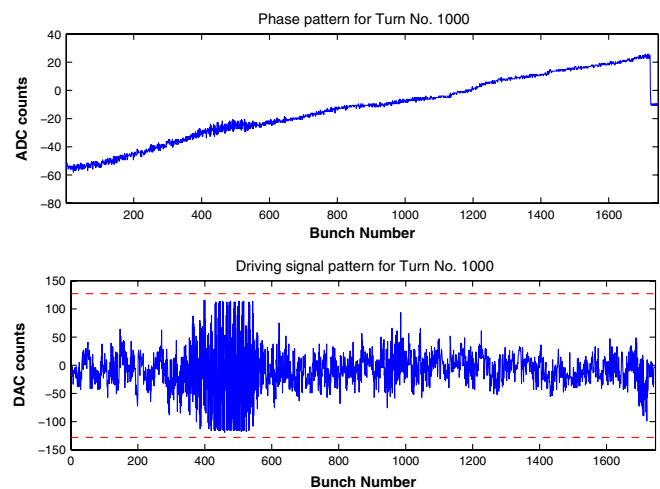


FIG. 15. (Color) Front-end and back-end signal of the longitudinal feedback system for turn 1000 of the time sequence shown in Figs. 8 and 9. The upper plot shows a set of unstable bunches around bunch 450. The lower plot depicts the response of the back-end signal to those unstable bunches. The large phase error drives the back-end signal to positive and negative saturation for the unstable bunches.

rated output power), this transient saturates the control path for the dominant HOM beam modes. The low-order beam modes and mode 0 are stabilized in this system by the low group delay woofer and are stable (damping) during the transient after the initial driven excitation.

The runaway beam motion is also interesting viewed on a modal domain as in Fig. 16. This figure presents the same numeric data as Fig. 9 processed to show the data of the input bunch motion transformed to the modal domain. Once HOM control is lost for a few bunches from power stage saturation, the runaway HOMs in the band centered at mode 800 grow exponentially. Mode 800 is the largest unstable mode driven by the cavity HOMs impedance, as seen in Fig. 1. The patterns seen in Fig. 16 are similar for many abort transients analyzed. As seen in Fig. 16, modes around 800 become unstable though the growth rate seen is lower than the natural (free) growth rate of those modes (the saturated feedback channel is reducing the growth rates but cannot reduce them to zero). At the end of the transient when the amplitude of the phase oscillation is very large, low-order beam modes eventually became unstable because the phase detector in the longitudinal system becomes blinded to the low-mode motion in the presence of large amplitude saturating HOM motion.

In conclusion it is important to remark that this operational limit described is mostly due to the equilibrium operation of the system near to the limits imposed by the system dynamic range or dynamic power.

The behavior where the HOM modes first become unstable and then at the end of the transient the low-order modes grow and become unstable can be qualitatively understood from Video 1. This animation is a movie sequence of the data presented in Fig. 8 and illustrates the dynamics in the front-end and back-end signals for this transient (Fig. 15 shows one instant of this sequence at turn 1000). It is interesting to watch the initial disturbance, and then the growth of a few bunches with mode 800 HOM-driven motion. The system dynamics is very complicated,






FIG. 16. (Color) Time-domain fault file (same data as Figs. 8 and 9) processed to show the data of the input bunch motion transformed to the modal domain. Once HOM control is lost for a few bunches from power stage saturation, the runaway HOMs in the band centered at mode 800 grow exponentially. Mode 800 is the largest impedance in the cavity HOMs.



VIDEO 1. The animation shows the phase detected input signals, and the computed correction signals, for the transient studied in Figs. 8–15. The transient lasts 35 ms and noise saturated behavior leads to loss of control and loss of the beam.

for several thousand turns the motion is confined to just a few bunches, and the remainder stay well controlled. But by turn 3300 the loss of bunch control begins to spread all around the turn, and beam is finally lost on low-mode motion as the woofer control path saturates. The exact origin of the source of the barycentric motion is not captured or recorded in the beam motion fault file—we can only infer what must be happening in the rf systems from the motion of the beam.

It is also interesting to look at the animation in Video 2, another animated sequence of 4800 turns in the machine. This recording was a snapshot taken to study the noise floor and transients in the stored beam, and there is no beam loss



VIDEO 2. The animation shows the system behavior for a different 35 ms of operation of the HER at similar conditions as Video 1. While excited in a manner leading to saturation of the control path, this transient ends with the feedback system regaining control with no loss of the beam.

at the end of the recording. However, we see that in operation the barycentric mode zero motion is bursting on and off, fully saturating the feedback channel at turns 500, 2500, and 4000. In this example, the beam system recovers from this saturation, and the HOM control recovers and the beam is not lost. The insight from recordings such as this helped form the understanding of the role of rf station transients and perturbations in defining the operational limits of the HOM control systems in HER and LER. It is interesting that, while the LER had higher absolute growth rates for HOM impedance driven motion, requiring higher feedback channel gain than the HER, the HER system, with the greater number of rf stations and higher power requirements, ran into the control limit due to rf system noise.

As these transients were related to all sorts of disturbances in the AC mains, and in the dynamics and jitter of individual HVPS regulators, they depended on the particular operating configuration of stations deployed in the HER, and on the health of the HVPS SCR stacks and regulators. In operation each station and HVPS would have unique noise and transient contributions. The relationship to the exact operating stations, which stations had HVPS SCR repairs, etc., was not initially understood by the group trying to understand the limits of the HOM feedback system. To their view these effects which randomly saturated the system occurred infrequently and appeared and disappeared over time for no obvious reason. Understanding the sequence and origins of these transient beam losses was difficult, and required patience to understand the larger environment in which the beam was interacting.

E. Lessons learned: Broadband coupled-bunch longitudinal feedback

The original designers never anticipated that the control limits of the system would be reached from impulsive noise and transients generated in the rf systems, and it was only in the last year of operations that this mechanism was finally seen and understood from fault-file data. As this mechanism was finally comprehended, several new LLRF (direct-loop) and low-mode (woofer) configurations were developed which had better rejection of the 720 Hz and harmonic perturbations and, hence, better regulation of the accelerating voltage. If higher current operation of the HER had been anticipated or required, more dramatic changes in the HVPS regulators and a different gain partitioning in the broadband feedback (through additional high-power output amplifiers) would have been necessary. Another control approach, which separated the low-mode control (in the woofer) from the frequency bands controlled by the HOM system, is possible to implement, though not consistent with the bunch-by-bunch time-domain filter architecture that was the core backbone of the DSP processing farm. This split-band processing ap-

proach was not developed in the original design, instead the low-mode woofer was seen as a parallel adjunct channel that might be useful to allow more robust control at the ultimate currents in the LER.

The central difficulty in feedback for the mode zero excitation is still that the rf system has megavolts to push the beam around, while the broadband feedback system has only a few kilovolts. As a result, very small modulations in the rf cavity voltages can completely saturate the HOM control. Addressing this situation with very expensive HOM broadband power amplifiers is not a cost effective direction (the final PEP-II installation had \$1 000 000 worth of broadband power amplifiers)—instead it is necessary to improve the regulation of the rf HVPS and the effectiveness of the cavity voltage regulation in the LLRF direct loops and other regulators. The residual excitations can then be controlled via feedback through the woofer path as originally planned. It is also necessary to have diagnostics which can indicate if line transients or other transient disturbances are occurring in an rf system, so that the origin of this sort of transient beam motion can be efficiently identified.

The designers of the broadband longitudinal feedback chose to develop a programmable system based on fixed-instruction DSP microprocessors, and knew that there was value in flexibility in the control filters. Over the course of the PEP-II operations, this central core signal processing system remained unchanged, though the actual control filters and DSP operating codes did evolve in response to various operating conditions (changes in synchrotron tune, concerns about noise rejection of out of band signals, etc.). The original design incorporated some flexible memory structures that were exploited in system diagnostics to excite the beam (very useful for narrowband single bunch excitation as part of timing the back-end kicker structures to the circulating beam [13]) as well as a beam signal recorder that was the central feature in the development of transient-domain beam diagnostics [14]. The flexibility and modularity of the original design was very important over the full lifetime of the project.

One area where the original design did change significantly over the lifetime of the machine involved the high-power beam line components. The flexibility to add additional power amplifiers was part of the original design, and as currents increased the number of installed power amplifiers was increased as anticipated. However, thermal difficulties with beam line components and cables were much worse than anticipated. As shown in Fig. 17, as currents increased over 2 A there were many difficulties with coaxial feedthroughs, power cables, and connectors.

The implementation of the LER beam line kicker vacuum elements was changed due to thermal management concerns. Here the original drift-tube kicker design [15,16] was based on HER and LER operating currents consistent with the design report, and as currents increased a damped



FIG. 17. (Color) Photo of a damaged high-power directional coupler installed on the PEP-II beam line to instrument beam-driven kicker power. While designed to operate at 5 KW power levels, the kicker loads, cables, and feedthroughs were problematic in operation. Three different styles of high-power rf connectors were commissioned in the PEP-II runs, each had high-power difficulties. The beam induced power was spread in a very broad band extending to 20 GHz. Seen in the figure is a damaged and destroyed 9/16 DIN connector and directional coupler port.

cavity-style kicker [17–19] was implemented to help with the beam induced heating. The basic drift-tube kicker design performed well from initial commissioning, and the original HER drift-tube kickers served until the final operation at twice the design current.

III. THE IMPACT OF NONLINEAR ELEMENTS IN THE LLRF FEEDBACK PATHS

The original designers of the PEP-II LLRF systems were very concerned about the large fundamental mode impedances driving unstable coupled-bunch motion at low modes as the cavities were detuned with increasing currents. The designs proposed featured direct and comb impedance controlling feedback loops [1]. In this approach it is necessary to have sufficient loop gain to reduce the effective impedance of the cavity rf fundamental. The design concept was developed using a frequency domain technique, and the essential topology of the two impedance control loops, the limits on group delay, the basic structure of the 2nd order IIR notch filter, the required direct-loop gain, etc., were studied [20,21]. The essential elements of the PEP-II LLRF and rf station design are shown in Fig. 18.

It was understood that a nonlinear element in this loop would have a significant impact on the effectiveness of the impedance control. The power klystron was an obvious candidate for nonlinear behavior. Linear and nonlinear (time-domain) simulation tools were used to estimate system performance, define stable operating points for the rf feedback, and estimate stability limits [20,22]. During this

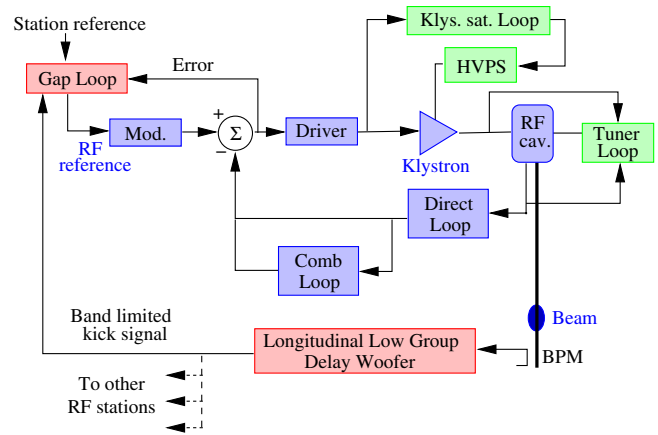


FIG. 18. (Color) Block diagram of an impedance controlled rf system used at PEP-II. In the actual implementation, each rf station had one 1.2 MW klystron and 2 or 4 HOM-damped rf cavities. The LLRF systems used direct and comb loops, klystron HV regulating loops, gap feed-forward loop, and other low-frequency regulators. The broadband longitudinal processing is implemented once per ring, as is the low group delay woofer processing. From the LGDW band-limited kick information is sent to selected stations around the ring via fiber optic digital data links.

design phase the high-power 1.2 MW PEP-II klystron was still in development, and data from a lower power PEP-I klystron was used to model a nonlinear power transfer characteristic in the time-domain simulation.

These initial simulation efforts were essential in developing the overall LLRF structure. However, they did not attempt to quantify stability limits of the systems (beyond producing system trajectories of 10 or 20 ms duration), and did not study the robustness of the system stability to small changes in the klystron or LLRF dynamic responses. They did not attempt to model the trade-offs between rf station stability (stability of the direct and comb loops) vs the stability of the beam dynamics. The essential focus of these initial simulation studies was to validate the topology of the proposed system design, using the broadband direct feedback loop in conjunction with a digital IIR comb filter. Using these tools, it was estimated that the cavity-fundamental driven coupled-bunch modes in the HER would be stable for the design current, while for the LER at design current it was anticipated that the cavity driven low modes would be stable in conjunction with the operation of the planned broadband feedback system.

To allow for some extra control margin for low coupled-bunch modes, and as an insurance policy, the original LLRF design [23,24] included a dedicated control path from the broadband longitudinal damping system, in which a band-limited version of the longitudinal HOM correction signal was driven through the LLRF system. This link allows the rf system and cavities to serve as a very powerful low-frequency beam kicker (this path was named the “woofer channel”).

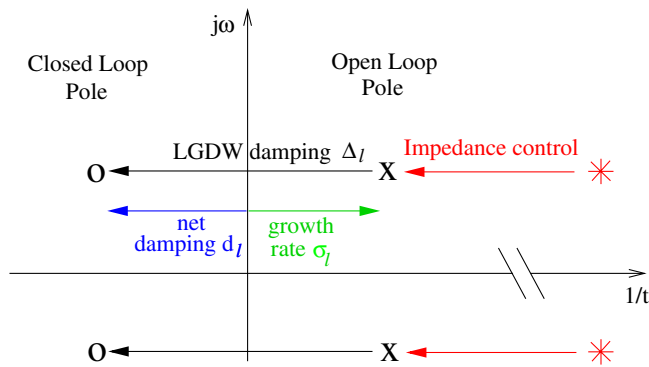


FIG. 19. (Color) The conceptual behavior of the most unstable longitudinal eigenmode driven from the cavity fundamental impedance as modified by the direct and comb loops of the LLRF feedback and the LGDW. Both systems are necessary to achieve stable behavior at high currents. If the impedance reduction of the LLRF control loops is less than as designed (trajectory from * to X), the maximum damping from the LGDW channel may not be sufficient to bring the closed-loop pole O to the left of the imaginary axis as required for a stable system.

A. PEP-II LLRF system commissioning and operational performance

Initial rf system commissioning was successful [25]. As currents increased in both HER and LER, the growth rates of the fundamental-driven longitudinal modes were 4 to 5 times greater than expected from the simulation estimates using either linear models or time-domain nonlinear models. While it was possible to control these instabilities in the HER at 1300 mA, and LER at 1500 mA through the design woofer configuration, there were great concerns that there would be a loss of control margin for modest increases of current.

The very fast low-mode instabilities led to the development of a “Low Group Delay Woofer” feedback channel [26,27], in which there was FIR filter control of the low modes independently from the HOM control filters. The commissioning of the low group delay woofer provided much better low-mode control and allowed higher currents, but there were still concerns about stability for increases in operating currents. Figure 19 conceptually illustrates the necessary combined action of the LLRF feedback and the LGDW feedback to achieve beam stability. The essential problem was that the impedance control from the LLRF system left low-mode growth rates at such a fast rate that the woofer channels were marginally able to control them due to group delay and gain limits.

B. Understanding performance limits in LLRF impedance control

The unusually fast growth rates were interpreted as signs of poor impedance control from the direct and comb feedback loops. Considerable effort was invested in trying to understand what mechanism was causing this reduction in

the achieved impedance control. The primary diagnostic which showed the problem was the measurement of the beam instability low-mode growth rates, but there were not diagnostics to quantify what in the LLRF system was causing these effects. The essential conclusion was that the effective gain of the direct and comb loops was less than predicted by the linear tools and models.

As it was always understood that the klystron behavior was nonlinear, this obvious and easy to identify nonlinear component was hypothesized to be the culprit causing the ineffective impedance control loops. The original station and beam dynamics models had shown an impact from a nonlinear power transfer characteristic. Attempts were made to try running the klystrons at a less-saturated operating point. (This was possible at modest beam currents, where the loss in efficiency was made up in extra heating in the collectors. This operating point was only possible for the SLAC-developed klystrons with full-power collectors, and was not possible in the Marconi and Phillips klystrons.) Machine measurements with only a fraction of the operating klystrons shifted to a different operating point were made, but with only a fraction of the net impedance possibly effected, the change in system dynamics would be small. The net result on the low-mode growth rates, as measured on the beam, was not very dramatic, and the measurements did not clearly show any real difference in growth rates vs klystron operating point.

Without better system models to understand the dynamics and impact of the klystron large-signal saturation curves, it really was not possible to predict what should or should not be expected from the klystron large-signal characteristic and variations in the operation point. Yet the need was very real to address this situation, as the low-mode instabilities were clearly going to become uncontrollable at foreseeable operating currents.

In response to these concerns a “klystron linearizer” was developed which used yet another feedback topology and control technique around the klystron to force the large-signal and small-signal gains to a fixed value [28]. While the linearity performance of this technique could be studied via a high-power test stand and spare klystron, the impact on the beam instability growth rates could only be evaluated on the actual storage ring. Because of the multiple rf cavities and rf stations, the impact of the technique could really only be quantified if all (or almost all) of the rf stations in the storage ring were equipped with the linearized klystrons. This was a significant technology and operational investment. Surprisingly, the experiment in the real machine showed the linearizers did not have the full effect predicted. This discrepancy between the machine and the modeling was not understood at the time.

There were also issues with stability of the direct and comb loops as currents increased. The rf system dynamics are parameter dependent on the cavity detuning and the operating point of the klystron amplifier. In operation

direct-loop and comb-loop parameters were feed-forward adapted to track the change in the cavity detuning with current. The loops became unstable as the loop operating points moved with klystron power. This amount of change in the system dynamics with operating point was not anticipated by the designers, and had not been incorporated into the system modeling and simulation as part of the design phase.

The operational difficulties, with a continual trade-off between station stability and instability growth rates, became a difficult issue as currents increased and margins were lost. Model-based configuration techniques were developed to allow the online configuration and tuning of the rf direct and comb loops using closed-loop transfer functions, taken with beam in the machine [29]. While this approach allowed better stability margins, this was a very intensive effort and time-consuming task. For every change of current, or operating configuration of operating and parked stations, gap voltages, etc., a series of loop optimizations had to be made for each station, as the dynamics of each station was unique.

C. Identifying the limiting nonlinear element

These concerns drove renewed investment in the nonlinear time-domain rf station-beam model. It was restructured to allow close comparisons between machine measurements and the modeling—the identical time-domain tools were used to measure low-mode instability growth in both the machine and the simulation [30]. This more detailed simulation revealed some subtle persistent deviations between the physical and simulated systems. In conjunction with high-power klystron test stand measurements, a consistent deviation in the frequency response of the small-signal gain between model and physical system was understood [31]. A medium power solid-state amplifier in the direct feedback path was eventually revealed as the source.

In the LLRF system, the entire processing chain must faithfully provide linear response for small modulation signals which can be 60 or 90 dB below the high-power fundamental (Fig. 20). As it is the small modulation signals which provide the impedance control feedback, the impact of a nonlinear element can be very significant. The design and development team did not realize the significance of testing all the processing chain for these sorts of nonlinear effects. The medium power amplifier was specified as having spurious harmonics better than -60 dBc and the amplifier was uneventfully tested for gain uniformity and frequency response as part of system design.

For the initial 7 years of operation these driver amplifiers had never been a source of any trouble or curiosity (the focus was often the power stages). As the impact of these nonlinear elements became better understood, new measurement techniques were developed to allow lab testing of system elements under realistic conditions. As an example,

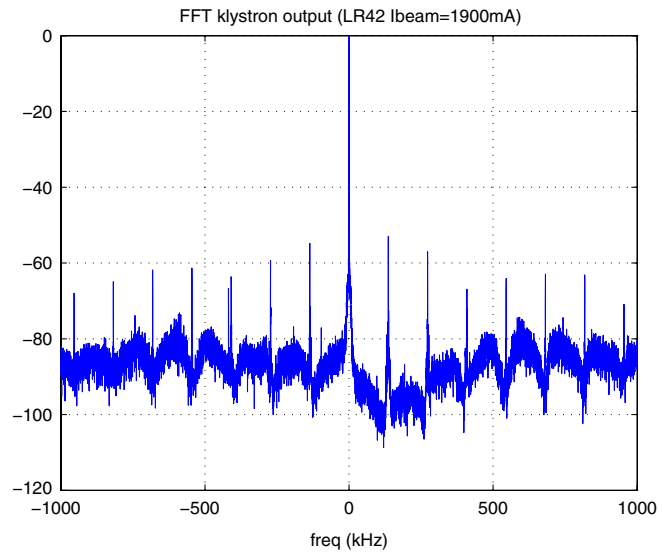


FIG. 20. (Color) Power spectrum of signals in the klystron output during closed-loop operation. ± 7 revolution harmonics are visible around the 476 MHz carrier.

a swept frequency response is a common laboratory measurement for rf components. Similarly, it is common to specify spurious responses in a power spectrum relative to a single carrier signal. To quantify the linearity of the LLRF components, new two-tone and swept small-signal plus large-signal carrier tests were developed. As shown in Fig. 21, two signals are presented to characterize the transfer function of the low-power klystron drive amplifier to modulation signals. The amplifier is tested using a large-signal power carrier in conjunction with a small test signal (-30 dB below the carrier). The small-signal gain compression is obvious and very significant. To compare with the original single sweep frequency test conducted on this amplifier, the large-signal response is included in Fig. 21.




FIG. 21. (Color) Large and small-signal transfer function measurement of the original LLRF driver amplifier. The large-signal response is a single swept test frequency, the small signal is a swept signal -30 dB below a fixed 476 MHz carrier.




FIG. 22. (Color) Swept sideband image responses for 13 amplifiers “A” from [31].

The other test developed to quantify the degree of non-linearity in the existing LLRF amplifiers was a carrier plus swept single-sideband image test. Here a large-signal carrier and a swept upper sideband tone at a lower level are impressed at the amplifier input. Instead of measuring the response at the excitation sideband frequency, a spectrum analyzer is used to look at maximum power across the band of interest while the excitation sideband signal is swept. A perfectly linear system would display no power at the image frequencies to the left of the center 476 MHz carrier.

Figure 22 shows the responses for 13 of the installed commissioned (“type A”) amplifiers. Besides the large variations in response between amplifiers, the image signals are roughly -10 to -15 dB below the sideband level (some amplifiers are much worse than this). In the LLRF application, this level of intermodulation would generate interfering signals above and below the 476 MHz fundamental that would transfer modulation from upper to lower revolution harmonics. Figure 23 shows this same sideband test for alternate amplifiers B and C. It is interesting that there is a difference in structure between the amplifiers, but both show better than -25 dB image suppression (for frequencies less than 1 MHz from the carrier one amplifier is clearly better with over -30 dB suppression).

This type of testing provided a wealth of dynamic and nonlinear information about the installed LLRF components. Using the information from the small-signal and carrier frequency response test (Fig. 21) and adding this amplifier’s nonlinear behavior to the rf station-beam simulation allowed the model to predict the actual rapid growth rates which had been such an operational difficulty. From this insight gained with the models, and the lab measurements, these low-power nonlinear amplifiers were replaced in all the rf systems. With this change the LER instability growth rates reduced, and again agreed with the model predictions based on the responses of the new amplifiers [32].




FIG. 23. (Color) Swept sideband image response for amplifier “B” and “C” from [31].

D. System improvements based on more complete models and system dynamics understanding

The simulation model also inspired the development of new rf configuration tools [29], and allowed new control approaches which trade off the stability of the rf station while increasing the stability of the beam. Implemented in the final 2 years of PEP-II, the impact of the “comb rotation” control techniques [30] and the model-machine




FIG. 24. (Color) Modeled (solid) and measured (dashed) LER low-mode instability growth rates for various operating configurations. The 30% reduction in cavity fundamental-driven growth rates from run 5 through to run 6 is due to replacing nonlinear drive amps and implementing comb rotation configurations. Error bars are not included for the April 2008 data because they crowd the image. They are comparable in magnitude with the error bars from the run 6 data. The final LER configuration developed (run 7) was estimated to allow 3600 mA operation in the LER.

agreement is shown in Fig. 24. The growth rates vs current are shown for the various configurations used in the last three years of PEP-II operation. The final configuration, which ran at 3100 mA in the LER, had excellent agreement with the model predictions [32].

E. Lessons learned: LLRF

The PEP-II LLRF architecture incorporated numerous programmable elements, and flexible modular packaging allowed some system functions to be expanded and modified over the operational history. One central feature of the original design was a flexible digital memory with associated analog A/D and D/A functions within the LLRF processing loops. These features were anticipated to be very important as part of a baseband network analyzer used to inject time-domain excitation sequences into the rf systems. Used in operation with beam, these excitations and recorded responses were analyzed and the frequency response and stability margins of the closed-loop systems were estimated. These features were critical in the development of online configuration tools, which allowed the response of an operating system to be quantified as the system dynamics changed with current. This memory was also the heart of a fault-file recorder system, in which transient data could be triggered in nominal operation of a station, or triggered as part of a machine fault sequence.

The original design was prescient in incorporating these features, which turned out to be essential for developing the techniques for high-current station configuration and for understanding some of the very complex system behavior from power supply ripple, line noise injection, and other system perturbations. Understanding the origins of an interfering signal can be very difficult, as once it is injected into a station, the information is impressed on the beam, and appears in the processing loops of all the other stations of that ring. Diagnosing the origin of low-frequency power supply noise, for example, is not obvious from signals recorded from the beam. For this reason several extra diagnostic modules to record HVPS signals were developed during PEP-II operations [33], as well as the addition of several new signal sources in the fault-file system.

The original implementation of the fault files did miss one very significant operational complexity. Each rf station in PEP-II ran as an independent rf station, with station-specific set points and station-specific loop configurations. The action of the beam acts to couple all the stations together through the circulating current and the signals impressed in the rf cavities. The system clocks in each station were phase locked to each other, as required to keep the multiple rf stations in phase alignment. However, there was no master time fiducial or time alignment signal which counted revolutions in the ring, or injected some system-common marker into the fault files. As a result, each station recorded time sequences with unknown time offsets with respect to the other stations. For example, one HER rf

station might fault on a HVPS fault, and trip the station, initiating a fault-file recording in that station. This would initiate a beam transient and beam loss, with all the other stations eventually tripping off on the loss of circulating current. The fault files would be generated in each individual station, and each would have a unique time alignment with respect to the others. Understanding the order of a complex fault sequence could be very difficult—a very skilled person can look for signatures of cavity voltages, etc., which could hint at the time of a common event, such a discreet loss of some part of the beam, and in this way find some time alignment of the many files. This sort of detective work was necessary to understand the source of some of the mysterious transients that would cause rf system trips, and in some cases the files were too complicated to sort out efficiently. Future system designers should implement a common clock or synchronizing marker which can be distributed and inserted into the many independent system fault recorders.

IV. SUMMARY

The PEP-II longitudinal feedback designers did foresee the essential requirements and implement many important features, among them transient-domain diagnostics used to quantify modal growth/damping rates. Major unforeseen surprises included the difficulties with thermal management of the beam induced power in the kicker structures. The limits of high-current instability control in the HER were understood in the last year of operation to be due to noise and HV power supply impulsive transients in the rf systems exciting barycentric beam-driven motion. This driven motion then saturated the output stages via the control filter path. This effect and dynamic range impact of driven motion was never anticipated in the system design phase. The designers were focused on the instabilities driven by the HOM impedances, and the necessary growth rates (gain), and viewed the system noise floor and beam motion detector as the important elements limiting the gain. While this was true for the installations at DAFNE and the ALS, the PEP-II experience was completely different.

There are many “lessons learned” in the LLRF experience. The system designers were correct in their suspicion that a nonlinear element would reduce the impedance control loop gain. However, the initial focus on the klystron power transfer nonlinearity missed the importance of another significant nonlinear element (the low-power driver amplifier). Before this mechanism was understood, the instability growth rates were attacked with a special LGDW feedback channel. A complex klystron linearizer was developed and investigated. But the complete understanding of the source of the fast growth rates was elusive for many years, despite considerable efforts and measurements.

The most available diagnostics from the machine were the physical modal growth rates, which could not be directly compared to the original simulation model results. It was the availability of the high-power test stand frequency response data (taken as part of the linearizer development) and the second generation nonlinear time-domain model that led to close comparison of the physical system and the models. The subtle differences led to the understanding of the effects. More importantly, the simulation model then offered a means to estimate the required performance of alternative technical components. This approach and the confidence in the simulation results led to the rapid specification, test, and commissioning of replacement driver amplifiers.

The LLRF effort also revealed the difficulties in configuration management, the amount of skilled resources needed to cope with the complexity of rf system operations due to individual station dynamics with station by station unique configurations. While the fault-file analysis was essential in understanding the system in operation, it took multiyear investment of very skilled people to understand the complex dynamics. The magnitude of this investment was never anticipated [34].

The PEP-II experience shows the essential value of system modeling and analysis. It was vital for all the project phases including design, commissioning, and operations. We could not identify useful information about what was happening from machine measurements and fault-file data alone. Alternative machine configurations were developed in the simulations and this testing and evaluation saved time during physical “machine development” studies. The project experience also shows the architectural usefulness of reconfigurable signal processing implemented in DSP and field programmable gate array techniques. This flexibility was exploited numerous times in the LFB and LGDW [35].

With the completion of the PEP-II program this experience is being transferred to new projects. A commercial multibunch feedback platform is now available based on the experience and techniques developed in these projects [36]. The LLRF modeling tools are being adapted to study similar systems in commissioning for the LHC and future light sources [37].

ACKNOWLEDGMENTS

Skilled and motivated people were the most important factor for the success of the PEP-II LLRF and LFB. Space limits our ability to name individually all of the skilled and essential contributors. We especially thank our collaborators at SLAC, LNF-INFN, KEK, and LBL who contributed so much to the feedback program, and our colleagues and collaborators at CERN who were central in the LLRF programs of the past and future. We thank the SLAC accelerator department, Klystron department, and Accelerator operations groups for support of the PEP-II

measurement and experimental efforts, and credit the SLAC ARD department for the consistent encouragement to develop new techniques in these areas. This work was supported by the U.S. Department of Energy under Contract No. DE-AC02-76SF00515.

-
- [1] D. Boussard, CERN Report No. CERN-SL-91-2.
 - [2] D. Boussard, in the 3rd Workshop on the Tau-Charm Factory, Marbella, Spain, 1993.
 - [3] R. A. Rimmer, J. Byrd, M. Irwin, and D. A. Goldberg, in *Proceedings of the 5th European Particle Accelerator Conference (EPAC 96), Sitges, Spain, 1996* (IOP Publishing, London, 1996), pp. 2035–2037; S. Myers, A. Pacheco, R. Pascual, C. Petit-Jean-Genaz, and J. Poole, *Proceedings of the 5th European Particle Accelerator Conference (EPAC 96), Sitges, Spain, 1996* (IOP Publishing, London, 1996), Vol. 3.
 - [4] S. Prabhakar, Doctoral thesis, Stanford University, 2001.
 - [5] J. Fox *et al.*, in the *9th Beam Instrumentation Workshop (BIW 2000)*, Cambridge, Massachusetts, 2000 (Report No. SLAC-PUB-8410, 2000, p. 8).
 - [6] D. Briggs *et al.*, in *Proceedings of IEEE 1991 Particle Accelerator Conference (APS Beam Physics)*, San Francisco, California, 1991, pp. 1407–1409.
 - [7] H. Hindi, S. Prabhakar, J. Fox, and D. Teytelman, Reports No. SLAC-PUB-7716 and No. SLAC-PEP-II-AP-NOTE-97-20, 1997, p. 4. Poster presented at Proceedings of the Particle Accelerator Conference, Vancouver, BC, Canada, 1997 (IEEE, New York, 1997).
 - [8] J. D. Fox *et al.*, in Proceedings of the 4th European Particle Accelerator Conference (EPAC 94), London, England, 1994, pp. 1619–1621.
 - [9] SLAC Report No. SLAC-R-418.
 - [10] G. D. Stover, *et al.*, in *Proceedings of the 18th Particle Accelerator Conference, New York, 1999* (IEEE, New York, 1999), Vol. 2, pp. 1213–1215.
 - [11] T. Mastorides, C. Rivetta, J. D. Fox, and D. V. Winkle, *13th Beam Instrumentation Workshop (BIW08)*, Lake Tahoe, California, 2008 (Report No. SLAC-PUB-13289).
 - [12] R. Cassel and M. N. Nguyen, in Proceedings of the Particle Accelerator Conference, Vancouver, BC, Canada, 1997 (Ref. [7]), p. 3482.
 - [13] D. Teytelman and J. Fox, *AIP Conf. Proc.* **648**, 474 (2002).
 - [14] J. Fox *et al.*, in Proceedings of the 18th Particle Accelerator Conference, New York, 1999 (Ref. [10]), pp. 636–640.
 - [15] J. M. Byrd, J. Johnson, G. Lambertson, and F. Voelker, in Stanford 1992, Proceedings, B Factories, pp. 220–223.
 - [16] F. Voelker and J. R. Johnson, Report No. SLAC-ABC-69.
 - [17] F. Marcellini, M. Tobiyama, P. MacIntosh, J. Fox, H. Schwarz, D. Teytelman, and A. Young, *AIP Conf. Proc.* **648**, 458 (2002).
 - [18] P. McIntosh *et al.*, in *Proceedings of the 20th Particle Accelerator Conference, Portland, OR, 2003* (IEEE, New York, 2003), p. 3141.
 - [19] R. Boni, A. Gallo, A. Ghigo, F. Marcellini, M. Serio, and M. Zobov, Part. Accel. **52**, 95 (1996); A. Gallo, R. Boni, A. Ghigo, F. Marcellini, M. Serio, and M. Zobov, prepared

- for International Workshop on Collective Effects and Impedance for B Factories (CEIBA 95), Tsukuba, Japan, 1995; A. Gallo, M. Migliorati, and L. Palumbo, *Nucl. Instrum. Methods Phys. Res., Sect. A* **404**, 223 (1998).
- [20] F. Pedersen, in Stanford 1992, Proceedings, B Factories, pp. 192–207.
- [21] S. T. Craig, Internal Reports, AECL Research, Accelerator Physics Branch (Chalk River), 1992.
- [22] R. Tighe and P. Corredoura, in *Proceedings of the Particle Accelerator Conference, Dallas, TX, 1995* (IEEE, New York, 1995), pp. 2666–2668.
- [23] P. Corredoura, R. Claus, L. Sapozhnikov, H. Schwarz, R. Tighe, and C. Ziomek, in Proceedings of the Particle Accelerator Conference, Dallas, TX, 1995 (Ref. [22]), pp. 2672–2674.
- [24] P. Corredoura, Proceedings of the 18th Particle Accelerator Conference, New York, 1999 (Ref. [14]), Vol. 1, pp. 435–439.
- [25] P. Corredoura *et al.*, in Proceedings of the Particle Accelerator Conference, Vancouver, BC, Canada, 1997 (Ref. [7]), p. 2359.
- [26] D. Teytelman, D. Van Winkle, and J. Fox, in *Proceedings of the 21st Particle Accelerator Conference, Knoxville, 2005* (IEEE, Piscataway, NJ, 2005), p. 1069.
- [27] L. Beckman *et al.*, in Proceedings of the 20th Particle Accelerator Conference, Portland, OR, 2003 (Ref. [18]).
- [28] J. Fox, T. Mastorides, D. Teytelman, D. Van Winkle, Y. Zhou, and A. Gallo, in Proceedings of the 21st Particle Accelerator Conference, Knoxville, 2005 (Ref. [26]), p. 2660.
- [29] D. Teytelman, in Proceedings of the 21st Particle Accelerator Conference, Knoxville, 2005 (Ref. [26]), p. 2863.
- [30] C. Rivetta, T. Mastorides, J. D. Fox, D. Teytelman, and D. Van Winkle, *Phys. Rev. ST Accel. Beams* **10**, 022801 (2007).
- [31] J. D. Fox, T. Mastorides, C. H. Rivetta, and D. Van Winkle, in *Proceedings of the 2007 Particle Accelerator Conference, Albuquerque, New Mexico* (IEEE, Albuquerque, New Mexico, 2007), p. 2451.
- [32] T. Mastorides, C. Rivetta, J. D. Fox, D. V. Winkle, and D. Teytelman, *Phys. Rev. ST Accel. Beams* **11**, 062802 (2008).
- [33] David N. Brown, SLAC internal engineering note.
- [34] D. Van Winkle, J. Fox, and D. Teytelman, in Proceedings of the 21st Particle Accelerator Conference, Knoxville, 2005 (Ref. [26]), p. 3931.
- [35] A. Drago, A. Gallo, A. Ghigo, M. Zobov, J. D. Fox, and D. Teytelman, *Phys. Rev. ST Accel. Beams* **6**, 052801 (2003).
- [36] Dimtel, Inc., San Jose, CA USA, <http://dimtel.com/index>.
- [37] T. Mastorides, C. Rivetta, J. D. Fox, D. V. Winkle, P. Baudrenghien, and J. Tuckmantel, in *Proceedings of the 11th European Particle Accelerator Conference, Genoa, 2008* (EPS-AG, Genoa, Italy, 2008), THPC125.